The embedded contact homology of toric contact manifolds

by

Keon Choi

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Mathematics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Michael Hutchings, Chair
Professor Denis Auroux
Professor Raphael Bousso

Fall 2013

The embedded contact homology of toric contact manifolds

Abstract

The embedded contact homology of toric contact manifolds

by

Keon Choi

Doctor of Philosophy in Mathematics

University of California, Berkeley

Professor Michael Hutchings, Chair

Embedded contact homology (ECH) is an invariant of a contact three-manifold. In Part I of this thesis, we provide a combinatorial description of the ECH chain complex of certain "toric" contact manifolds. This is an extension of the combinatorial description appearing in [11] and [12]. ECH capacities are invariants of a symplectic four-manifold with boundary, which give obstructions to symplectically embedding one symplectic four-manifold with boundary into another. In Part II of this thesis, we compute the ECH capacities of a large family of symplectic four-manifolds with boundary, called "concave toric domains". Examples include the (nondisjoint) union of two ellipsoids in $\mathbb{R}^4$. We use these calculations to find sharp obstructions to certain symplectic embeddings involving concave toric domains. This is a joint work with D. Cristofaro-Gardiner, D. Frenkel, M. Hutchings and V. G. B. Ramos.

Contents

Acknowledgments

I would like to thank my advisor Professor Hutchings for giving me guidance and support throughout my Ph.D. I thank Professor Auroux and Professor Bousso for reading my thesis and giving me helpful suggestions. I thank Vinicius Ramos for many helpful discussions and suggestions in preparing this thesis. I also thank my collaborators, D. Cristofaro-Gardiner, D. Frenkel, M. Hutchings and V. G. B. Ramos, for helpful discussions and for granting permissions to include our joint work here.

Part I

Combinatorial description of embedded contact homology of toric contact manifolds

Chapter 1

Introduction

Given a three-manifold Y equipped with a nondegenerate contact form λ, the embedded contact homology (ECH) of (Y, λ) is the homology of a chain complex generated by certain unions of Reeb orbits and whose differential counts certain embedded holomorphic curves in $\mathbb{R} \times Y$. This paper aims to provide a combinatorial description of ECH chain complexes for a class of "toric" contact manifolds.

The inspiration for this work comes from the two papers by Hutchings and Sullivan [11,12] where the notions of "polygonal paths" and "rounding corners" were introduced. These were used to describe the ECH generators and differentials of T^3 with certain contact forms as well as closely related instances of the periodic Floer homology. We extend these notions and show similar results for $I \times T^2$ and T^3, where I is an interval and both are equipped with general torus-invariant contact forms.

We remark that the homology of the ECH chain complex can be computed indirectly: ECH, Heegaard Floer homology and Seiberg-Witten Floer homology are isomorphic to each other [3, 16, 25] and there is a combinatorial formulation of Heegaard Floer homology [21]. However, it is of theoretical interest to understand the ECH chain complex itself. More practically, computation of contact geometric invariants such as ECH spectrum requires more information about the chain complex as such information is lost under the above isomorphism.

In this section, we introduce the class of contact manifolds investigated in the main part of this paper and state the main theorem. Section 2 proves the main theorem. Section 3 applies the main result to describe the ECH chain complex of T^3 with torus-invariant contact forms.

1.1 Embedded contact homology

We briefly review the definition of ECH necessary for stating the main theorem. For details, see Section 2.1. Given a closed oriented three-manifold Y, a contact form on Y is a 1-form λ on Y satisfying $\lambda \wedge d\lambda > 0$. Then, λ determines a contact structure $\xi = \mathrm{Ker}(\lambda)$, which

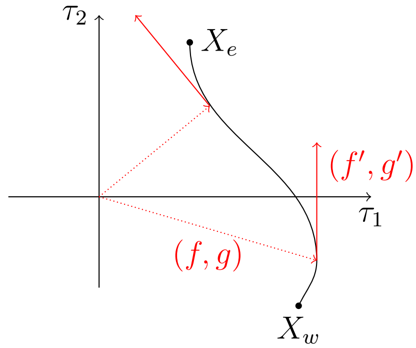

Figure 1.1: An example of a path (f, g) satisfying (1.2.1).

is an oriented two-plane field, and the Reeb vector field R characterized by $d\lambda(R,\cdot) = 0$ and $\lambda(R) = 1$. Assume that λ is nondegenerate, which means that all Reeb orbits of λ are "cut out transversely". We also fix a generic admissible almost complex structure J on the symplectization $\mathbb{R} \times Y$. This means that J is $\mathbb{R}$-invariant, $J(\partial_s) = R$ where s denotes the $\mathbb{R}$ coordinate, and J sends ξ to itself so that $d\lambda(v, Jv) > 0$ for $0 \neq v \in \xi$. See [10, Section 1.3] for details.

An orbit set γ in the homology class $\Gamma \in H_1(Y)$ is a finite set of pairs $\{(\gamma_i, m_i)\}$ where γ_i are distinct embedded Reeb orbits and m_i are positive integers such that $[\gamma] := \sum m_i[\gamma_i] = \Gamma$. We say that γ is *admissible* if $m_i = 1$ whenever γ_i is hyperbolic. Then, the ECH chain complex $ECC_*(Y, \lambda, \Gamma, J)$ is generated (over $\mathbb{Z}/2$ coefficients) by admissible orbit sets in the homology class Γ. Let $H_2(Y, \alpha, \beta)$ denote the set of 2-chains Σ in Y with $\partial\Sigma = \sum_i m_i[\alpha_i] - \sum_j n_j[\beta_j]$, modulo boundaries of 3-chains. If $\alpha = \{(\alpha_i, m_i)\}$ and $\beta = \{(\beta_j, n_j)\}$ are two orbit sets and $Z \in H_2(Y, \alpha, \beta)$, we associate to them an integer $I(\alpha, \beta, Z)$ called the *ECH index*.

Let (Σ, j) be a punctured compact Riemann surface and consider a (J-)holomorphic map $u : (\Sigma, j) \to (\mathbb{R} \times Y, J)$. A positive/negative end of u is an association of a puncture of Σ to a (possibly multiply covered) Reeb orbit ρ so that, near that puncture, u is asymptotic to $\mathbb{R} \times \rho$ with $s \to \pm\infty$, respectively. A holomorphic curve from α to β is a holomorphic map u whose total multiplicity of positive ends at covers of α_i is m_i and whose total multiplicity of negative ends at covers of β_i is n_i, with no other ends. The ECH differential coefficient $\langle\partial\alpha, \beta\rangle$ between two generators α and β counts holomorphic curves u from α to β with $I(\alpha, \beta, [\mathrm{im}(u)]) = 1$. We will sometimes use $C := \mathrm{im}(u)$ to refer to the holomorphic curve u and $I(C)$ to denote $I(\alpha, \beta, [C])$.

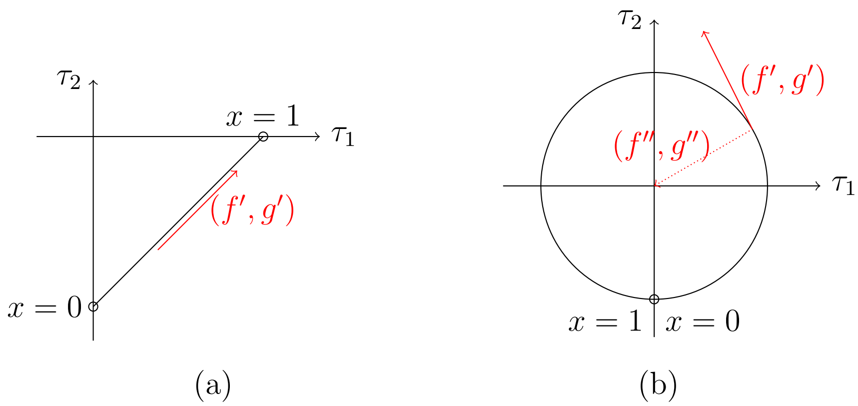

Figure 1.2: Graphs of (f, g) taken from (S^3, λ_{std}) and (T^3, λ_n).

1.2 Toric contact manifold $(I \times T^2, \lambda)$

Let $I = [X_w, X_e] \subset \mathbb{R}$ be an interval with coordinate x. Choose a pair of generic real-valued smooth functions f and g on I and suppose $(f, g) : I \to \mathbb{R}^2$ satisfies the pointwise condition

$$(f, g) \times (f', g') = fg' - f'g > 0. \tag{1.2.1}$$

Here, $f' = df/dx, g' = dg/dx$ and $\times$ denotes the usual cross product in $\mathbb{R}^2$. Figure 1.1 shows an example of (f, g) satisfying the condition (1.2.1).

Consider the oriented three-manifold $I \times T^2$ where $T^2 = (\mathbb{R}/\mathbb{Z})^2$ has coordinates t_1 and t_2 and the orientation is given by the ordered basis $\{\partial_{t_1}, \partial_x, \partial_{t_2}\}$. Consider a 1-form

$$\bar{\lambda} = -gdt_1 + fdt_2 \tag{1.2.2}$$

on $I \times T^2$. Equation (1.2.1) implies that $\bar{\lambda}$ is a contact form. The contact structure $\bar{\xi}$ of $\bar{\lambda}$ is

$$\bar{\xi} = \text{span}\{\partial_x, -f\partial_{t_1} - g\partial_{t_2}\} = \text{span}\{\partial_x\} \oplus \text{span}\{\partial_{t_1}, \partial_{t_2}\}$$

and the Reeb vector field $\bar{R}$ of $\bar{\lambda}$ is

$$\bar{R} = \frac{f'\partial_{t_1} + g'\partial_{t_2}}{(fg' - f'g)} \in \text{span}\{\partial_{t_1}, \partial_{t_2}\}.$$

Hence, the graph of (f, g) illustrates how $\bar{\xi}$ and $\bar{R}$ are rotating. Note that we have an S^1-family of closed Reeb orbits at each $x \in I$ where

$$f'/g' \in \mathbb{Q} \cup \{\infty\}.$$

Example 1.2.1. We present (f, g) taken from some standard contact manifolds.

(i) Consider $\mathbb{C}^2 = \mathbb{R}^4$ with coordinates z_i for $i = 1, 2$ and the standard symplectic form. Under the new coordinates (r_1, r_2, t_1, t_2) given by $r_i = |z_i|^2$ and $t_i = \arg(z_i)/2\pi \in \mathbb{R}/\mathbb{Z}$,

$$\omega_{std} = \pi \sum_i dr_i dt_i.$$

Define a Liouville vector field X_L by $X_L := \frac{1}{\pi}\sum_i r_i \partial_{r_i}$ so that

$$\bar{\lambda} := \iota_{X_L}\omega_{std} = \sum r_i dt_i.$$

Hence, restricting $\bar{\lambda}$ to $S^3 = \{(1-x, x, t_1, t_2) \in \mathbb{C}^2 | x \in [0, 1]\}$ gives the standard contact form

$$\lambda_{std} = (1-x)dt_1 + xdt_2$$

on S^3. The graph of (f, g) corresponding to λ_{std} on $(0, 1) \times T^2 \subset S^3$ is shown in Figure 1.2 (a). Recall that λ_{std} is degenerate and we have an S^2-family of Reeb orbits in (S^3, λ_{std}). This is reflected by the graph of (f, g) having a constant rational slope of one, giving rise to $((0, 1) \times S^1)$-family of orbits in $(0, 1) \times T^2$.

(ii) Consider $T^3 = (\mathbb{R}/\mathbb{Z})^3$ with coordinates x, t_1, t_2 and a contact form

$$\lambda_n = (\cos 2n\pi x)dt_1 + (\sin 2n\pi x)dt_2$$

for some $n \geq 1$. Then, (f, g) corresponding to λ_n on $(0, 1) \times T^2 \subset (\mathbb{R}/\mathbb{Z}) \times T^2$ is shown in Figure 1.2 (b), where n is the number of times the graph of (f, g) rotates around the origin. If $n = 1$, we can embed (T^3, λ_1) into $(\mathbb{C}^\times)^2 \subset \mathbb{C}^2$ similarly to above. This time, let $X_L = \frac{1}{\pi}\sum_i (r_i - 2)\partial_{r_i}$ on $(\mathbb{C}^\times)^2$ and restrict to $T^3 = S \times T^2$ where $S = \{(r_1, r_2) \in (0, \infty)^2 | \sum_i |r_i - 2|^2 = 1\}$.

Definition 1.2.2. (Convexity) Let $I = [X_w, X_e]$ and consider $I \times T^2$.

(a) A contact form $\bar{\lambda} = -gdt_1 + fdt_2$ on $I \times T^2$ is said to be *convex* (respectively *concave*) at $x = x_0$ if

$$(f', g') \times (f'', g'') > 0 \text{ (respectively } < 0)$$

at $x = x_0$. If $(f', g') \times (f'', g'') = 0$ at some $x = x_0$, then we call x_0 a *point of inflection.* We say that $\bar{\lambda}$ is convex (respectively concave) if $\bar{\lambda}$ is convex (respectively concave) at all $x \in I$.

(b) A Reeb orbit of $\bar{\lambda}$ at $x = x_0$ is said to be *convex* (respectively *concave*) if $\bar{\lambda}$ is convex (respectively concave) at $x = x_0$.

In Figure 1.2, all orbits of (T^3, λ_n) are convex while the orbits of (S^3, λ_{std}) are neither convex nor concave since $(f', g') \times (f'', g'') = 0$. In Figure 1.3, the orbits at $x = x_{conv}$ are convex and the orbits at $x = x_{conc}$ are concave.

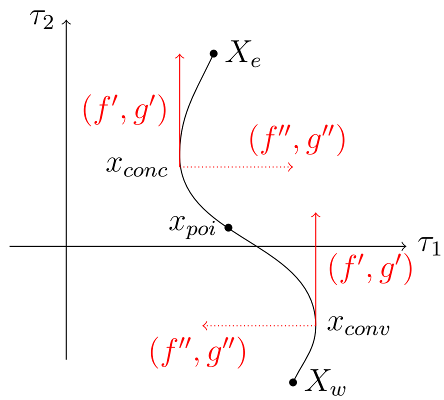

Figure 1.3: (f, g) for some $\bar{\lambda}$ with a point of inflection at $x = x_{poi}$.

We note that, even though $(I \times T^2, \bar{\lambda})$ contains infinitely many S^1-families of Reeb orbits, we may regard only finitely many S^1-families as relevant and disregard the rest by using a filtered version of ECH. We recover the usual ECH by a direction limit argument (see Section 2.1). Then, for a generic choice of f and g, all relevant Reeb orbits are either convex or concave. Moreover, using a filtered version of ECH allows the following perturbation of $\bar{\lambda}$. Recall that defining ECH requires a nondegenerate contact form λ. By a general Morse-Bott argument as in [2], one can perturb $\bar{\lambda}$ to λ so that each of the relevant S^1-families of Reeb orbits gives two Reeb orbits of λ and no other relevant Reeb orbits. These two orbits correspond to the two critical points of the auxiliary Morse function on S^1 and one can show that one of these two orbits is positive hyperbolic while the other is elliptic. See Section 2.1 for more details on the Morse-Bott argument and the definition of preferred perturbations of λ, which we call "good" perturbations (Definition 2.1.2). Throughout the paper, $\bar{\lambda}$ will denote a T^2-invariant contact form $-g dt_1 + f dt_2$ on $I \times T^2$ and λ will denote a good perturbation of $\bar{\lambda}$. We will say that a Reeb orbit of λ is convex/concave if it comes from an S^1-family of convex/concave Reeb orbits of $\bar{\lambda}$.

We point out that the usual ECH differential counts holomorphic curves in a symplectization of a closed contact three-manifold. Here, we count holomorphic curves in $\mathbb{R} \times (I \times T^2)$ that do not intersect $\mathbb{R} \times \{X_w, X_e\} \times T^2$. A version of the maximum principle (Lemma 2.2.1) guarantees that we still have Gromov compactness for such moduli spaces as in [1].

1.3 Combinatorial representation

In this section, we define combinatorial objects that will be used to state the main theorem.

Definition 1.3.1. Let $I = [X_w, X_e]$ be an interval.

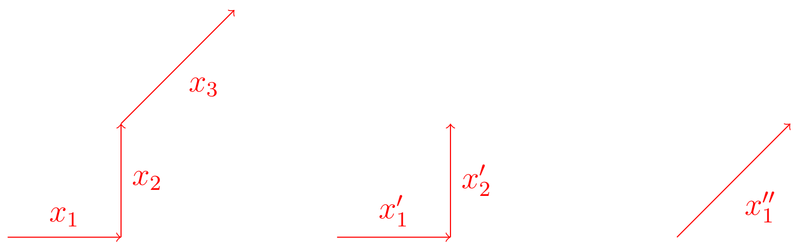

Figure 1.4: Examples of IP paths $\mathcal{P}, \mathcal{P}'$ and $\mathcal{P}''$.

(a) An *(abstract) integral polygonal path* $\mathcal{P}$, or an *IP path*, on I is an n-tuple (v_i) where $n \geq 0$ and each v_i is a pair (w_i, x_i) such that:

  (i) for $1 \leq i \leq n$, $w_i \in \mathbb{Z}^2$ is a primitive vector and $x_i \in I$, and

  (ii) for $1 \leq i < n$, $x_i \leq x_{i+1}$ with equality only if $w_i = w_{i+1}$.

  Each v_i is called an *edge* of $\mathcal{P}$ at $x = x_i$. We will treat v_i also as a vector in $\mathbb{Z}^2$ when convenient and write $x(v_i) := x_i$.

(b) A *realization* of an IP path $\mathcal{P}$ with n edges is (the image of) a continuous map $\phi : [0, n] \to \mathbb{R}^2$ satisfying:

  (i) $\phi(0) \in \mathbb{Z}^2$ and

  (ii) for each $1 \leq i \leq n$, $\phi|_{[i-1,i]}$ is linear and $\phi(i) = \phi(i-1) + v_i$.

(c) A *decoration* of an IP path $\mathcal{P}$ is an association of each edge of $\mathcal{P}$ with one of the labels in $\{\check{e}, \check{h}, \hat{e}, \hat{h}\}$.

Note that a realization of an IP path is unique up to a translation by $\mathbb{Z}^2 \subset \mathbb{R}^2$. One can depict an IP path by its realization ϕ and by marking $\phi([i-1, i])$ with x_i for each $1 \leq i \leq n$. Figure 1.4 depicts an IP path $\mathcal{P}$ consisting of three edges v_i with $x_i = x(v_i)$, an IP path $\mathcal{P}'$ consisting of two edges v'_j with $x'_j = x(v'_j)$ and an IP path $\mathcal{P}''$ consisting of one edge v''_1 with $x''_1 = x(v''_1)$. For examples of decorations, see Figure 1.5 (b).

Lemma 1.3.2. *Let $I = [X_w, X_e]$ be an interval, let $\bar{\lambda} = -g dt_1 + f dt_2$ be a T^2-invariant contact form on $I \times T^2$ and let λ be a good perturbation λ of $\bar{\lambda}$. There is a natural way to assign a unique IP path on I, denoted $\mathcal{P}_\gamma$, to each orbit set γ of λ. In addition, γ induces a decoration of $\mathcal{P}_\gamma$ uniquely up to transposing labels on two edges v and v' with $x(v) = x(v')$.*

Proof. We can write an orbit set γ of λ in the "ordered product notation"

$$\gamma = \gamma_1 \gamma_2 \cdots \gamma_n,$$

where each γ_i is an embedded orbit at $x = x(\gamma_i)$ and $x(\gamma_i)$ is non-decreasing. This representation is unique up to transposing an elliptic γ_i and a hyperbolic γ_j with $x(\gamma_i) = x(\gamma_j)$.

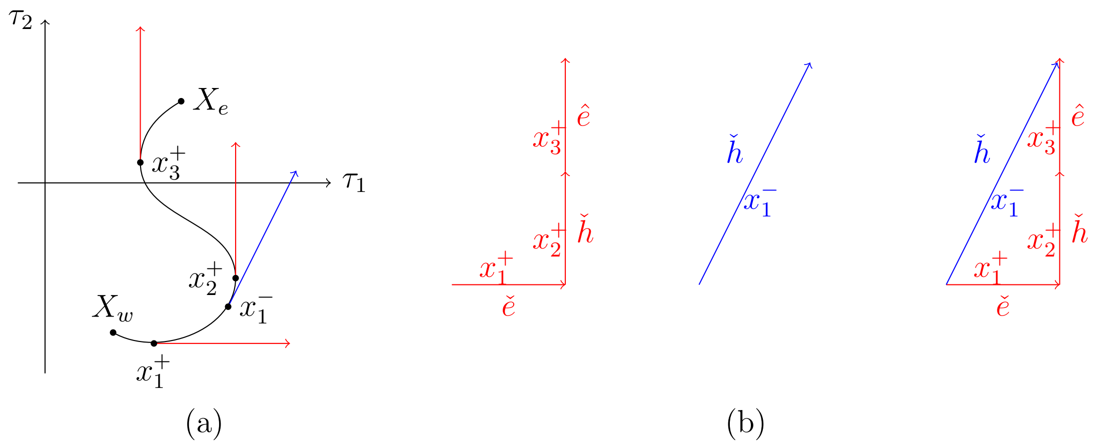

Figure 1.5: Two orbit sets of λ, IP paths associated to each of them, and an IP region between them, with induced decorations.

Then, $\mathcal{P}_\gamma = (v_i)$ where v_i is the pair $([\gamma_i], x(\gamma_i))$ with $[\gamma_i] \in H_1(I \times T^2) = \mathbb{Z}^2$ and $x(\gamma_i) \in I$. We can also label each v_i according to whether γ_i is elliptic convex ($\check{e}$), hyperbolic convex ($\check{h}$), elliptic concave ($\hat{e}$) or hyperbolic concave ($\hat{h}$). See Section 1.2 for the four types of Reeb orbits of λ. □

Definition 1.3.3. We call $\mathcal{P}_\gamma$, as in Lemma 1.3.2, the IP path *associate to* γ. The decoration of $\mathcal{P}_\gamma$ as in Lemma 1.3.2 is called an *induced decoration*.

Figure 1.5 (a) shows the graph of (f, g) for a contact form $\bar{\lambda} = -g dt_1 + f dt_2$ on $I \times T^2$. Each arrow corresponds to an embedded orbit appearing in two orbit sets α and β of a good perturbation of $\bar{\lambda}$. In the ordered product notation, $\alpha = \alpha_1\alpha_2\alpha_3$ where α_i are embedded orbits of λ at $x(\alpha_i) = x_i^+$, and α_1, α_2 and α_3 are elliptic convex $\check{e}$, hyperbolic convex $\check{h}$ and elliptic concave $\hat{e}$, respectively. Similarly, $\beta = \beta_1$ where β_1 is an embedded orbit of λ at $x(\beta_1) = x_1^-$ and it is hyperbolic convex. In Figure 1.5 (b), the IP paths associated to α and β are drawn in red and blue, respectively, along with an induced decoration. The last picture will be explained shortly.

Remark 1.3.4. Let $I = [X_w, X_e]$ and consider $I \times T^2$.

(a) Suppose we fixed a contact form $\bar{\lambda} = -g dt_1 + f dt_2$ on $I \times T^2$ and a good perturbation λ of $\bar{\lambda}$.

  (i) Not all IP paths on I are associated to orbit sets of λ: if $\mathcal{P}$ is associated to an orbit set of λ, $x(v)$ must satisfy $f'(x(v))/g'(x(v)) \in \mathbb{Q} \cup \{\infty\}$ for each $v \in \mathcal{P}$ and $x(v)$ determines $v \in \mathbb{Z}^2$, since it must be a positive multiple of $(f', g') \in \mathbb{R}^2$ at $x = x(v)$.

  (ii) Even if $\mathcal{P}$ satisfies the above conditions, not all decorations of $\mathcal{P}$ can be induced from an orbit set of λ since the convexity of λ at $x = x_i$ determines whether v_i should be labeled by a check ($\vee$) or a hat ($\wedge$).

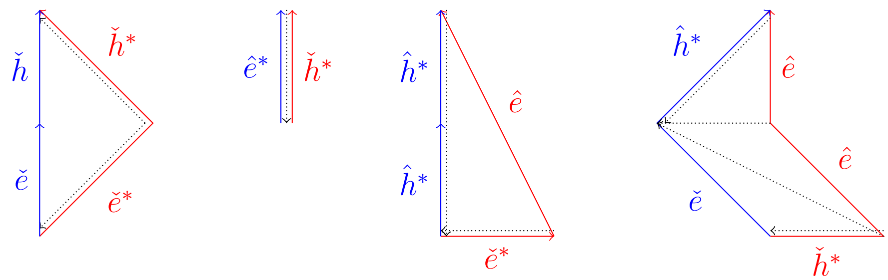

Figure 1.6: Decorated IP regions: slice classes are drawn in dotted lines and extreme edges are marked with asterisks.

(b) On the other hand, given a decorated IP path $\mathcal{P}$ on I, it is easy to find (f, g) satisfying (1.2.1) and a good perturbation λ of $\bar{\lambda} = -g dt_1 + f dt_2$ so that, for some orbit set α of λ, $\mathcal{P} = \mathcal{P}_\alpha$ with an induced decoration.

We prefer to consider all IP paths on I and all decorations on them without reference to a particular λ.

Definition 1.3.5. Let $I = [X_w, X_e]$ be an interval and let $\mathcal{P}^+ = (v_i^+)$ and $\mathcal{P}^- = (v_j^-)$ be two IP paths on I with

$$\sum_i v_i^+ = \sum_j v_j^- \in \mathbb{Z}^2. \tag{1.3.1}$$

(a) The *(abstract) integral polygonal region* $\mathcal{R}$, or the *IP region*, on I between $\mathcal{P}^+$ and $\mathcal{P}^-$ is the pair $(\mathcal{P}^+, \mathcal{P}^-)$. We write $\partial^+\mathcal{R} := \mathcal{P}^+$ and $\partial^-\mathcal{R} := \mathcal{P}^-$. Each edge of $\partial^+\mathcal{R}$ and $\partial^-\mathcal{R}$ is called a *positive edge* and a *negative edge*, respectively. Positive and negative edges of $\mathcal{R}$ are called *edge*s of $\mathcal{R}$ and the set of edges of $\mathcal{R}$ is denoted $\partial^\pm\mathcal{R}$.

(b) Let $\mathcal{R}$ be an IP region with m edges and let (v_k) be an ordering of $\partial^\pm\mathcal{R}$ so that $x(v_k)$ is nondecreasing. A *realization* of $\mathcal{R}$ is (the image of) the continuous map $\Phi : [0,1] \times [0,m] \to \mathbb{R}^2$ satisfying: $\Phi([0,1] \times \{0\}) = p \in \mathbb{Z}^2$ and for each $1 \le k \le m$,

  (i) $\Phi|_{[0,1]\times[k-1,k]}$ is linear.

  (ii) If v_k is a positive edge, then $\Phi(1,k) = \Phi(1,k-1) + v_k$ and $\Phi(0,k) = \Phi(0,k-1)$.

  (iii) If v_k is a negative edge, then $\Phi(0,k) = \Phi(0,k-1) + v_k$ and $\Phi(1,k) = \Phi(1,k-1)$.

(c) A *decoration* of an IP region $\mathcal{R}$ is a decoration of IP paths $\partial^+\mathcal{R}$ and $\partial^-\mathcal{R}$.

Note that a realization of an IP region is unique up to a translation by $\mathbb{Z}^2$ as well as reordering (v_k) by transposing two edges v and v' with $x(v) = x(v')$. If $v_{k_0+1}, \cdots, v_{k_0+m_0}$ are all the edges of $\mathcal{R}$ at $x = x_0$, then $\mathrm{im}(\Phi|_{[0,1]\times[k_0,k_0+m_0]})$ is unchanged under a transposition of (v_k) among these m_0 edges. Also, $\phi^+ := \Phi|_{\{1\}\times[0,m]}$ is a realization of $\partial^+\mathcal{R}$ after collapsing each subinterval $[k-1,k]$ where ϕ^+ is constant. Similarly, $\phi^- := \Phi|_{\{0\}\times[0,m]}$ gives a realization of $\partial^-\mathcal{R}$. We can depict an IP region $\mathcal{R}$ by its realization and by marking ϕ^+ and ϕ^- as before. Figure 1.5 (b) and Figure 1.6 show examples of (decorated) IP regions, with extra information which we discuss shortly.

Definition 1.3.6. Let $I = [X_w, X_e]$ and let $\bar{\lambda}$ and λ be contact forms on $I \times T^2$ as in Lemma 1.3.2. Let α and β be orbit sets of λ with $[\alpha] = [\beta] \in \mathbb{Z}^2$. The IP region *associated to* α and β is the IP region between $\mathcal{P}_\alpha$ and $\mathcal{P}_\beta$ and is denoted $\mathcal{R}_{\alpha,\beta}$. An *induced decoration* of $\mathcal{R}_{\alpha,\beta}$ is a decoration of $\mathcal{P}_\alpha$ and $\mathcal{P}_\beta$ induced by α and β, respectively.

Note that the homology condition $[\alpha] = [\beta]$ ensures (1.3.1). In Figure 1.5 (b), $\mathcal{R}_{\alpha,\beta}$ is the triangle between the red path $\mathcal{P}_\alpha$ and the blue path $\mathcal{P}_\beta$.

Definition 1.3.7. Let $I = [X_w, X_e]$ and let $\mathcal{R}$ be an IP region on I.

(a) The *slice class* $\sigma_{\mathcal{R}}(x_0) \in \mathbb{Z}^2$ of $\mathcal{R}$ at $x = x_0$ is

$$\sigma_{\mathcal{R}}(x_0) := -\sum_{\substack{v\in\partial^+\mathcal{R}\\ x(v)\le x_0}} v + \sum_{\substack{w\in\partial^-\mathcal{R}\\ x(w)\le x_0}} w, \tag{1.3.2}$$

simply written as $\sigma(x_0)$ when $\mathcal{R}$ is clear.

(b) An edge v_0 of $\mathcal{R}$ is said to be *west extreme* if $x(v_0) = \min\{x(v)|v \in \partial^\pm\mathcal{R}\}$. It is said to be *east extreme* if $x(v_0) = \max\{x(v)|v \in \partial^\pm\mathcal{R}\}$. West extreme and east extreme edges of $\mathcal{R}$ are collectively referred to as *extreme* edges of $\mathcal{R}$.

Let $x_0 \in I$. It is easy to check that if Φ is a realization of $\mathcal{R}$ and k_0 is the number of edges v of $\mathcal{R}$ with $x(v) \le x_0$, then

$$\sigma_{\mathcal{R}}(x_0) = -\Phi(1,k_0) + \Phi(0,k_0).$$

Figure 1.6 depicts four decorated IP regions $\mathcal{R}$ between $\mathcal{P}^+$ and $\mathcal{P}^-$ along with each distinct slice class drawn in a dotted arrow. We omitted the markings x_i^+'s and x_j^-'s for simplicity. Despite the omission, the slice classes determine the order of the real numbers x_i^+'s and x_j^-'s and, in particular, the extreme edges of $\mathcal{R}$. Here, extreme edges are marked with asterisks only for illustrative purposes. One can check that each $\mathcal{R}$ in fact arises from a pair of orbit sets of λ described in Figure 1.5. This association is unique here despite the omission of x_i^+'s and x_j^-'s, but this is not true in general. For example, in Figure 1.13, omitting x_i on the top horizontal edges will result in ambiguity.

Definition 1.3.8. (Concatenations of IP paths and IP regions)

(a) We say that two IP paths $\mathcal{P} = (v_i)$ with n edges and $\mathcal{P}' = (v'_j)$ with n' edges are *composable* if $x(v_n) < x(v'_1)$, or if $x(v_n) = x(v'_1)$ with $v_n = v'_1 \in \mathbb{Z}^2$. If they are composable, we obtain an IP path with $(n+n')$ edges by concatenating ordered tuples (v_i) and (v'_j). We call this the *concatenation* of $\mathcal{P}$ and $\mathcal{P}'$ and denote it by $\mathcal{P}\mathcal{P}'$.

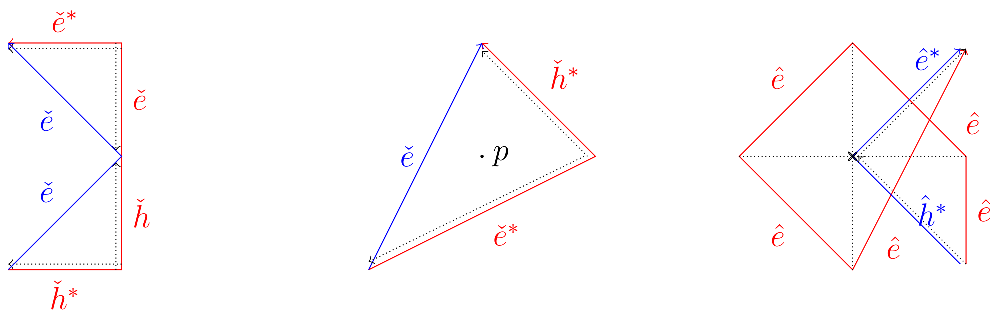

Figure 1.7: A decomposable region, a non-minimal region and a (non-embedded) minimal region.

(b) We say that two IP regions $\mathcal{R}$ and $\mathcal{R}'$ are *composable* if $\partial^+\mathcal{R}$ and $\partial^+\mathcal{R}'$ are composable, $\partial^-\mathcal{R}$ and $\partial^-\mathcal{R}'$ are composable and $\max\{x(v)|v \in \partial^{\pm}\mathcal{R}\} \leq \min\{x(v')|v' \in \partial^{\pm}\mathcal{R}'\}$. In this case, the *concatenation* $\mathcal{R}\mathcal{R}'$ of $\mathcal{R}$ and $\mathcal{R}'$ is the IP region between $\partial^+\mathcal{R}\partial^+\mathcal{R}'$ and $\partial^-\mathcal{R}\partial^-\mathcal{R}'$.

In Figure 1.4, $\mathcal{P}$ is the concatenation of $\mathcal{P}'$ and $\mathcal{P}''$, assuming $x_1 = x_1', x_2 = x_2', x_3 = x_1''$. The first region in Figure 1.7 is a concatenation of two triangular regions. We now describe some special properties of IP regions, which will play a role in the description of the differential.

Definition 1.3.9. Let $I = [X_w, X_e]$. Let $\mathcal{R}$ be an IP region on I with m edges and let $\Phi : [0,1] \times [0,m] \to \mathbb{R}^2$ be a realization of $\mathcal{R}$.

(a) $\mathcal{R}$ is called a *local bigon* if it has two edges v and w and they satisfy $x(v) = x(w)$. $\mathcal{R}$ is said to be *nonlocal* if it is not a local bigon.

(b) $\mathcal{R}$ is said to be *decomposable* if it can be written as a concatenation $\mathcal{R}_1\mathcal{R}_2$ for some IP regions $\mathcal{R}_1$ and $\mathcal{R}_2$, each with a positive number of edges. We say $\mathcal{R}$ is *indecomposable*, otherwise.

(c) A lattice point $p \in \mathbb{Z}^2$ is *internal* to Φ if there is an open ball $U \subset \text{int}([0,1] \times [0,m])$ so that $p \in \Phi(U)$. We say $\mathcal{R}$ is *minimal* if Φ contains no internal lattice point.

Note that the definition of minimality does not depend on the choice of a realization Φ. Each of the regions in Figure 1.6, including the bigon, is nonlocal. See also Figure 1.10 and Figure 1.11 for the distinction between local and nonlocal bigons. Figure 1.7 demonstrates decomposability and minimality.

Remark 1.3.10. Let $I = [X_w, X_e]$ and let $\bar{\lambda}$ and λ be contact forms on $I \times T^2$ as in Lemma 1.3.2. Let α and β be orbit sets of λ and let J be a small perturbation of the admissible almost complex structure $\bar{J}$ in (2.1.11). There is a natural parallel between the following features of a holomorphic curve C from α to β and the aforementioned features of an IP region $\mathcal{R}_{\alpha,\beta}$ between $\mathcal{P}_\alpha$ and $\mathcal{P}_\beta$:

(a) The *slice* $\mathcal{S}_C(x_0)$ of C at $x = x_0 \notin \{x(\alpha_i), x(\beta_j)\}$ is

$$\mathcal{S}_C(x_0) := C \cap (\mathbb{R} \times \{x_0\} \times T^2),$$

oriented outward normal first as a boundary of $C \cap (\mathbb{R} \times [X_w, x_0] \times T^2)$. Then, $\mathcal{S}_C(x_0)$ defines a homology class $[\mathcal{S}_C(x_0)] \in \mathbb{Z}^2 = H_1(\mathbb{R} \times I \times T^2)$ and $\sigma_{\mathcal{R}_{\alpha,\beta}}(x_0) = [\mathcal{S}_C(x_0)]$.

(b) An end of C at a (possibly multiply covered) orbit ρ_0 is *west/east extreme* if $x(\rho_0) = \min / \max\{x(\rho) | \rho \text{ is an end of } C\}$.

(c) C is a *local cylinder* if it has one positive end and one negative end and they have the same x-coordinate. C is *nonlocal* if it is not a local cylinder.

(d) We say that C is *irreducible* if its domain is connected. An irreducible C is analogous to an indecomposable $\mathcal{R}_{\alpha,\beta}$. See Proposition 2.2.7 and Corollary 2.2.8 for a precise statement.

(e) A genus zero C is analogous to a minimal $\mathcal{R}_{\alpha,\beta}$. See Proposition 2.2.10 for a precise statement.

One can draw a similar parallel between these features of IP regions and features of tropical curves in a view by Taubes [24] and Parker [20]. See also Remark 2.2.16.

An important notion regarding an IP region is positivity. It is related to the intersection positivity of holomorphic curves.

Definition 1.3.11. (Positivity) Let $I = [X_w, X_e]$ and let $\mathcal{R}$ be an IP region on I. Consider a realization $\Phi : [0,1] \times [0,m] \to \mathbb{R}^2$ of $\mathcal{R}$ with usual orientations on $[0,1] \times [0,m] \subset \mathbb{R}^2$ and $\mathbb{R}^2$.

(a) $\mathcal{R}$ is said to be *positive* if $\Phi|_{[0,1]\times[k-1,k]}$ is either degenerate or orientation-preserving for each $1 \leq k \leq m$.

(b) Let $\bar{\lambda}$ be a T^2-invariant contact form on $I \times T^2$ and let $\bar{R}$ denote the Reeb vector field of $\bar{\lambda}$. $\mathcal{R}$ is said to be *positive with respect to* $\bar{\lambda}$ if

$$\bar{R}(x) \times \sigma_{\mathcal{R}}(x) \geq 0 \tag{1.3.3}$$

for all $x \in I$.

We note that the definition of positivity does not depend on the choice of a realization Φ. All IP regions depicted previously are in fact positive. In Figure 1.8, the first IP region is positive and the next two are not. All IP regions in Figure 1.6 are positive with respect to $\bar{\lambda}$ given in Figure 1.5.

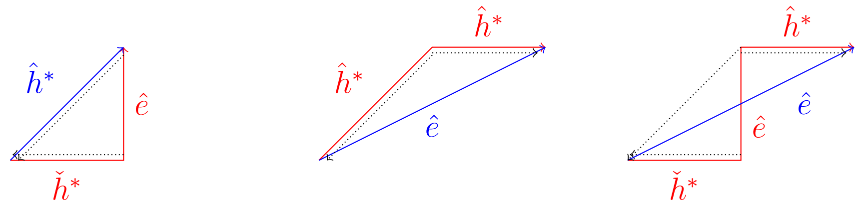

Figure 1.8: Three IP regions: the first is positive, the next two are not.

Remark 1.3.12. Let $I = [X_w, X_e]$ and consider $I \times T^2$.

(a) If α and β are orbit sets of a T^2-invariant contact form $\bar{\lambda}$ on $I \times T^2$ and $\mathcal{R}_{\alpha,\beta}$ is positive with respect to $\bar{\lambda}$, then the definition of Φ and (1.3.3) imply that $\mathcal{R}_{\alpha,\beta}$ is positive.

(b) Conversely, if $\mathcal{R}$ is a positive IP region on I and $v = v' \in \mathbb{Z}^2$ whenever v and v' are two edges of $\mathcal{R}$ with $x(v) = x(v')$, it is easy to find a T^2-invariant contact form $\bar{\lambda}$ on $I \times T^2$ so that:

  (i) $\mathcal{R} = \mathcal{R}_{\alpha,\beta}$ for some orbit sets α and β of a good perturbation of $\bar{\lambda}$, and

  (ii) $\mathcal{R}$ is positive with respect to $\bar{\lambda}$.

  See also Remark 1.3.4.

We now define a combinatorial analogue of the ECH index for a decorated IP region.

Definition 1.3.13. (Index of an IP region) Let $\mathcal{R}$ be a decorated IP region. We define

$$I(\mathcal{R}) := 2Area(\mathcal{R}) + \sum_{v \in \partial^+\mathcal{R}} CZ(v) - \sum_{v \in \partial^-\mathcal{R}} CZ(v) \tag{1.3.4}$$

where $Area(\mathcal{R})$ is the (signed) area of a realization of $\mathcal{R}$ with respect to the standard volume form on $\mathbb{R}^2$ and $CZ(v)$ is $1, 0, 0$, and -1 if v is labeled $\check{e}, \check{h}, \hat{h}$ and $\hat{e}$, respectively.

The three regions in Figure 1.8 have $2Area(\mathcal{R}) = 1, -1$ and 0, respectively, and $I(\mathcal{R}) = 0, 0$ and 0, respectively.

Definition 1.3.14. Let $\mathcal{R}$ be a decorated nonlocal IP region.

(a) We say that an edge v of $\mathcal{R}$ is:

  (i) S^1-loose and $\mathbb{R}$-loose if v is a positive edge labeled $\check{e}$, or a negative edge labeled $\hat{e}$.

  (ii) S^1-tight and $\mathbb{R}$-loose if v is a positive edge labeled $\check{h}$, or a negative edge labeled $\hat{h}$.

  (iii) S^1-loose and $\mathbb{R}$-tight if v is a positive edge labeled $\hat{h}$, or a negative edge labeled $\check{h}$.

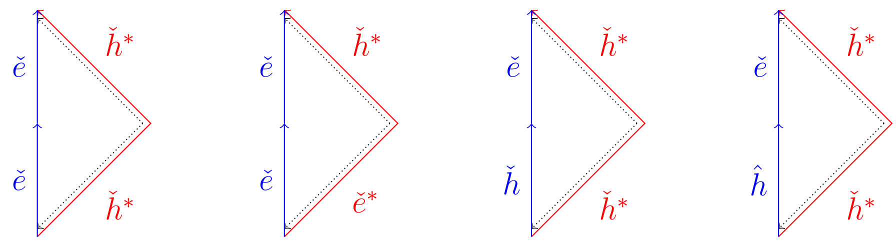

Figure 1.9: Four decorations of $\mathcal{R}$. The first is minimal and the next three are obtained from it by S^1-relaxing an edge, S^1-relaxing a different edge and $\mathbb{R}$-relaxing an edge, respectively.

(iv) S^1-tight and $\mathbb{R}$-tight if v is a positive edge labeled $\hat{e}$, or a negative edge labeled $\check{e}$.

(b) S^1*-relaxing* an edge v refers to replacing the label of v from S^1-tight to S^1-loose while keeping $\mathbb{R}$-tightness.

(c) $\mathbb{R}$*-relaxing* an edge v refers to replacing the label of v from $\mathbb{R}$-tight to $\mathbb{R}$-loose while keeping S^1-tightness.

(d) The *minimal* decoration of $\mathcal{R}$ is the decoration of $\mathcal{R}$ where all edges are S^1-tight, all extreme edges are $\mathbb{R}$-loose and all non-extreme edges are $\mathbb{R}$-tight.

From Definition 1.3.13, it is easy to check that relaxing an edge corresponds to increasing the index $I(\mathcal{R})$ by one. Also, the minimal decoration gives the smallest index $I(\mathcal{R})$ among all decorations of $\mathcal{R}$ where extreme edges are labeled $\mathbb{R}$-loose: we will see why we impose such a condition in Corollary 2.2.3. In Figure 1.9, the indices are $I(\mathcal{R}) = 0, 1, 1$ and 1, respectively.

Remark 1.3.15. From the S^1 Morse-Bott theory perspective, S^1-relaxing an edge corresponds to removing the restriction of a holomorphic curve having an end at a particular fixed orbit of an S^1-family. However, in general, the above definitions regarding decorations are made without any association to a specific λ on $I \times T^2$ or its orbit sets (see Remark 1.3.4). In particular, the minimal decoration of an IP region *does not* refer to a decoration induced from a particular pair of orbit sets. Similarly, $\mathbb{R}$-relaxing an edge of a decorated IP region *does not* refer to replacing an orbit appearing in a pair of orbit sets.

1.4 The main theorem

Theorem 1.4.1. *Let $I = [X_w, X_e]$, let $\bar{\lambda} = -g dt_1 + f dt_2$ be a T^2-invariant contact form on $I \times T^2$ and let $\bar{J}$ be a generic admissible almost complex structure on $\mathbb{R} \times I \times T^2$. Suppose (λ, J) is a good perturbation of $(\bar{\lambda}, \bar{J})$. (See Definition 2.1.3). If α and β are admissible orbit sets of λ, then $\langle \partial\alpha, \beta \rangle \neq 0 \in \mathbb{Z}/2$ if and only if there exist orbit sets γ_1 and γ_2 such that:*

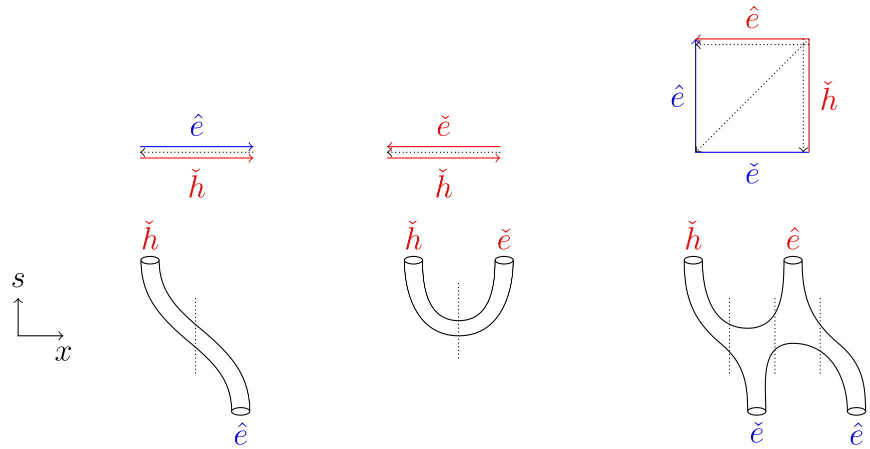

Figure 1.10: IP regions corresponding to nonzero differential.

(a) $\alpha = \gamma_1 \alpha' \gamma_2$ *and* $\beta = \gamma_1 \beta' \gamma_2$ *in the ordered product notation,*

(b) $\mathcal{R}_{\alpha',\beta'}$ *is positive with respect to* $\bar{\lambda}$*,*

(c) $\mathcal{R}_{\alpha',\beta'}$ *is nonlocal, indecomposable and minimal with two extreme edges, and*

(d) an induced decoration of $\mathcal{R}_{\alpha',\beta'}$ *can be obtained from the minimal decoration of* $\mathcal{R}_{\alpha',\beta'}$ *by* S^1*-relaxing one edge.*

We remark that if one induced decoration of $\mathcal{R}_{\alpha',\beta'}$ can be obtained from the minimal decoration by S^1-relaxing one edge, then any induced decoration of $\mathcal{R}_{\alpha',\beta'}$ can be obtained this way.

Example 1.4.2. Here are some (decorated) IP regions $\mathcal{R}_{\alpha,\beta}$ associated to a pair of admissible orbit sets α and β with $\langle \partial\alpha, \beta \rangle \neq 0$. These are illustrated in Figure 1.10 along with members of $\mathcal{M}(\alpha, \beta)$:

(i) A nonlocal bigon $\mathcal{R}_{\alpha,\beta}$ with one positive edge and one negative edge. This corresponds to a holomorphic cylinder with one positive end and one negative end.

(ii) A nonlocal bigon $\mathcal{R}_{\alpha,\beta}$ with two positive edges. This corresponds to a holomorphic cylinder with two positive ends.

(iii) An IP region $\mathcal{R}_{\alpha,\beta}$ with (possibly multiple) positive edges and (possibly multiple) negative edges (in the picture, two positive edges and two negative edges). This corresponds to general holomorphic curves (in fact, spheres). The number of positive ends can be any positive number and the number of negative ends can be any number, as long as the total number is at least two. See Figure 1.13.

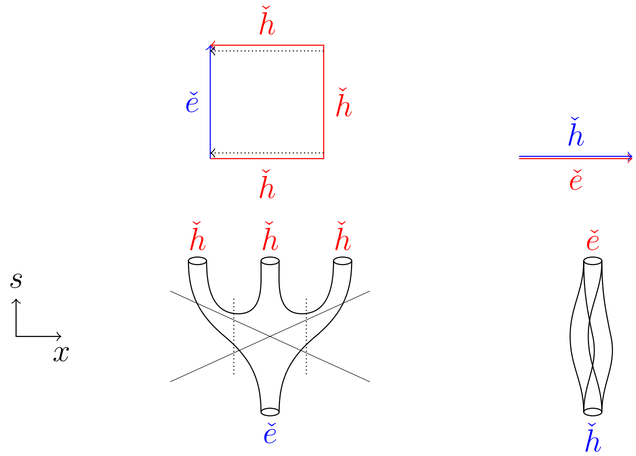

Figure 1.11: $I(\mathcal{R}) = 1$ IP regions corresponding to zero differential coefficient.

Here are some (decorated) IP regions $\mathcal{R}_{\alpha,\beta}$ associated to a pair of admissible orbit sets α and β with $\langle \partial\alpha, \beta \rangle = 0$. These are illustrated in Figure 1.11 along with members/non-members of $\mathcal{M}(\alpha, \beta)$:

(iv) An IP region $\mathcal{R}_{\alpha,\beta}$ whose induced decoration can be obtained from the minimal decoration by $\mathbb{R}$-relaxing an edge. $\mathcal{M}(\alpha, \beta)$ is empty for a good perturbation (λ, J), although it may (and will in some cases) be nonempty in general.

(v) A local bigon $\mathcal{R}_{\alpha,\beta}$. This corresponds to a Morse flow within an S^1-family of Reeb orbits. Such holomorphic curves exist in pairs.

Remark 1.4.3. We make a few remarks in comparison with [12] and [11]. Let $I = [X_w, X_e]$ and consider $\bar{\lambda}$ and λ on $I \times T^2$ as in Theorem 1.4.1.

(a) If $\bar{\lambda}$ is convex, e.g. (T^3, λ_n) in Example 1.2.1 (ii), then $\mathcal{R}_{\alpha',\beta'}$ as in Theorem 1.4.1 cannot have any non-extreme positive edges and the two extreme edges must be positive. In [12], β' is said be obtained from α' by *rounding a corner*. Similarly, if $\bar{\lambda}$ is concave, then $\mathcal{R}_{\alpha',\beta'}$ cannot have any non-extreme negative edges and the two extreme edges must be negative. This is “dual” to rounding a corner as in [11]. See Figure 1.12 and Remark 1.4.4.

(b) If $\bar{\lambda}$ is convex, a holomorphic curve C with $I(C) = 1$ and no negative ends must have exactly two positive ends. For a general T^2-invariant contact form $\bar{\lambda}$ can support an $I(C) = 1$ holomorphic curves with no negative ends but with an arbitrary number of positive ends. See Figure 1.13 for an example. One can similarly construct an $I(C) = 1$ holomorphic curve with arbitrary number of positive and negative ends as described in Example 1.4.2 (iii).

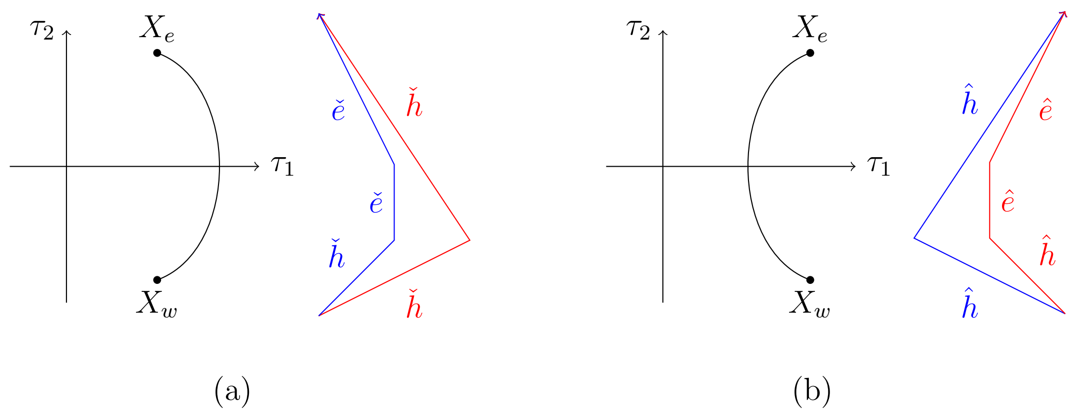

Figure 1.12: (a) Rounding a corner for a convex $\bar{\lambda}$ and (b) its dual operation for $\bar{\lambda}^\vee$.

(c) In contrast to rounding corners, for a general $\bar{\lambda}$, $\mathcal{R}_{\alpha,\beta}$ with $\langle \partial\alpha, \beta \rangle \neq 0$ may not be embedded. For example, the last IP region in Figure 1.7 is positive, irreducible and minimal and the decoration can be obtained from the minimal decoration by S^1-relaxing the east extreme edge. As in Remark 1.3.12, one can give a T^2-invariant contact form $\bar{\lambda}$ and a perturbation λ of $\bar{\lambda}$ so that this decorated IP region is associated to a pair of admissible orbit sets α and β of λ.

Remark 1.4.4. (Duality) We observe that the criteria of Theorem 1.4.1 is symmetric in the following sense. Let $I = [X_w, X_e]$ and let $f, g, f^\vee$ and $g^\vee : I \to \mathbb{R}$ be such that the graphs of (f, g) and $(f^\vee, g^\vee)$ are reflections of each other about some straight line. Furthermore, suppose that both

$$\begin{aligned} \bar{\lambda} &= -g dt_1 + f dt_2, \qquad \text{and} \\ \bar{\lambda}^\vee &= -g^\vee dt_1 + f^\vee dt_2 \end{aligned}$$

define contact forms on $I \times T^2$. See Figure 1.12 for a simple example where (f, g) and $(f^\vee, g^\vee)$ are reflections about a vertical line. Also, suppose two IP regions $\mathcal{R}$ and $\mathcal{R}^\vee$ are reflections of each other about the same straight line so that:

- $\partial^+\mathcal{R} = \partial^-\mathcal{R}^\vee$, $\partial^-\mathcal{R} = \partial^+\mathcal{R}^\vee$, and
- the convexity of the labels are reversed.

as illustrated in Figure 1.12. Then, $\mathcal{R}$ satisfies the conditions of Theorem 1.4.1 if and only if $\mathcal{R}^\vee$ does. This gives the duality between the differential for the ECH chain complex of $(I \times T^2, \lambda)$ and the differential for the ECH chain complex of $(I \times T^2, \lambda^\vee)$.

Here, it is important that both $\bar{\lambda}$ and $\bar{\lambda}^\vee$ are contact. To illustrate this, if $\bar{\lambda}$ is convex, there can be a nonlocal bigon with two positive edges satisfying the conditions of Theorem 1.4.1, as in Example 1.4.2 (ii). However, it is easy to see that any $\bar{\lambda}^\vee$ which is dual to $\bar{\lambda}$

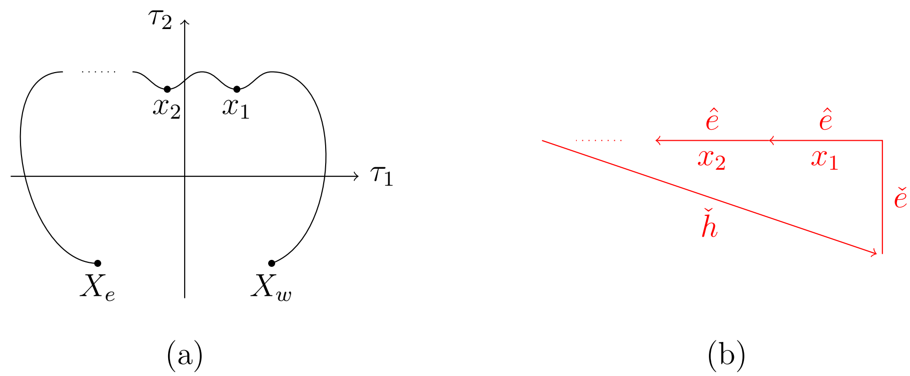

Figure 1.13: $\mathcal{R}$ can have arbitrarily many top horizontal $\hat{e}$ edges.

in the above sense is necessarily *not contact*. In fact, if such $\bar{\lambda}^{\vee}$ is contact, then Theorem 1.4.1 would imply that there is a holomorphic cylinder with two negative ends contributing to a nonzero differential in the ECH chain complex of $(I \times T^2, \lambda^{\vee})$. This certainly does not happen.

In the next section, we show that the criteria in Theorem 1.4.1 are necessary, primarily using the ECH index computation and a version of intersection positivity. We show that the criteria are sufficient by using induction and reducing the case of general holomorphic curves to the case of holomorphic spheres with two or three punctures, which were analyzed by Taubes in [23].

Chapter 2

Proof of the main theorem

2.1 Preliminaries

Embedded contact homology

We continue with the review of ECH from Section 1.1, following [7,10]. Let (Y, λ) be a three-manifold with a nondegenerate contact form λ and fix $\Gamma \in H_1(Y)$ and a generic admissible almost complex structure J on $\mathbb{R} \times Y$. Recall that the ECH chain complex $ECC_*(Y, \lambda, \Gamma, J)$ is generated by admissible orbit sets in the homology class Γ. To describe the moduli spaces of interest, it is convenient to use the notion of holomorphic currents. We say that two holomorphic curves C and C' are equivalent if C is obtained from C' by a pre-composition with a bi-holomorphic map on its domain. Then, a *holomorphic current* $\mathcal{C}$ is a finite set of pairs $\{(C_k, d_k)\}$ where C_k are equivalent classes of distinct irreducible somewhere injective holomorphic curves in $(\mathbb{R} \times Y, J)$ with positive and negative ends at Reeb orbits and d_k are positive integers. We say that a holomorphic current $\mathcal{C}$ is "somewhere injective" if $d_k = 1$ for each k and say that $\mathcal{C}$ is "embedded" if it is somewhere injective, each C_k is embedded and C_k are pairwise disjoint.

Let $\alpha = \{(\alpha_i, m_i)\}$ and $\beta = \{(\beta_j, n_j)\}$ be two orbit sets in the homology class Γ and let $Z \in H_2(Y, \alpha, \beta)$. A holomorphic current from α to β in the homology class Z is a holomorphic current whose total multiplicity of positive ends at α_i is m_i, the total multiplicity of negative ends at β_j is n_j, with no other ends and $[\mathcal{C}] = Z$. Let $\mathcal{M}(\alpha, \beta, Z)$ denote the moduli space of such holomorphic currents.

We now define the ECH index $I(\alpha, \beta, Z) \in \mathbb{Z}$, also denoted $I(\mathcal{C})$ if $\mathcal{C} \in \mathcal{M}(\alpha, \beta, Z)$. Let τ be a symplectic trivialization of ξ over each of the Reeb orbits α_i and β_j. Then,

$$I(\alpha, \beta, Z) := c_\tau(Z) + Q_\tau(Z) + CZ_\tau^I(\alpha, \beta), \tag{2.1.1}$$

for $c_\tau(Z), Q_\tau(Z)$ and $CZ_\tau^I(\alpha, \beta)$ which we describe next. For more details, see [7,10]. $c_\tau(Z) = \langle c_1(\xi, \tau), Z \rangle$ is the relative Chern class: if S is a representative of Z, this is the count of zeroes of a section of $\xi|_S$ which is constant with respect to τ near ends. The second term $Q_\tau(Z)$ is the relative intersection pairing, which is the algebraic intersection number between S

and a push-off of S, satisfying certain conditions near ends. The third term CZ_τ^I is the Conley-Zehnder term

$$CZ_\tau^I(\alpha,\beta) = \sum_i \sum_{k=1}^{m_i} CZ_\tau(\alpha_i^k) - \sum_j \sum_{k=1}^{n_j} CZ_\tau(\beta_j^k), \tag{2.1.2}$$

where $CZ_\tau(\rho) \in \mathbb{Z}$ is the Conley-Zehnder index of the Reeb orbit ρ with respect to τ. Compare this with the Fredholm index of a holomorphic curve C with k positive ends at $\rho_1^+, \cdots, \rho_k^+$ and l negative ends at $\rho_1^-, \cdots, \rho_l^-$:

$$\operatorname{ind}(C) = -\chi(\Sigma) + 2c_\tau([C]) + CZ_\tau^{ind}(C). \tag{2.1.3}$$

where

$$CZ_\tau^{ind}(C) = \sum_{i=1}^{k} CZ_\tau(\rho_i^+) - \sum_{j=1}^{l} CZ_\tau(\rho_j^-). \tag{2.1.4}$$

Here are some important properties of the ECH index [10, Section 3.4]. Let α and β be orbit sets of (Y,λ) and let $Z \in H_2(Y,\alpha,\beta)$. Then,

(a) (Well-defined) $I(\alpha,\beta,Z)$ does not depend on the choice of the trivialization τ.

(b) (Index ambiguity formula) If $Z, Z' \in H_2(Y,\alpha,\beta)$, then

$$I(\alpha,\beta,Z) - I(\alpha,\beta,Z') = \langle c_1(\xi) + 2PD(\Gamma), Z - Z' \rangle \tag{2.1.5}$$

(c) (Additivity) If γ is another orbit set in the homology class Γ and if $W \in H_2(Y,\beta,\gamma)$, then

$$I(\alpha,\gamma,Z+W) = I(\alpha,\beta,Z) + I(\beta,\gamma,W) \tag{2.1.6}$$

(d) (Index inequality) If a holomorphic curve C from α to β is somewhere injective, then

$$\operatorname{ind}(C) \leq I(C) \tag{2.1.7}$$

with equality only if C is embedded and the multiplicity of the ends at any orbit satisfies a certain "partition condition" (see Definition 2.1.1.)

(e) (Trivial cylinders) If $\mathcal{C}$ is contains no trivial cylinders and $\mathcal{T}$ is a union of trivial cylinders, then

$$I(\mathcal{C} \cup \mathcal{T}) \geq I(\mathcal{C}) + 2\#(\mathcal{C} \cap \mathcal{T}). \tag{2.1.8}$$

Let C be a holomorphic curve from $\alpha = \{(\alpha_i, m_i)\}$ to $\beta = \{(\beta_j, n_j)\}$. For each i, C has ends at covers of α_i with total multiplicity m_i. This gives a partition of m_i denoted by $p_i^+(C)$. We similarly define the partition $p_j^-(C)$ of n_j for each j.

Definition 2.1.1. For each embedded Reeb orbit ρ and $m \geq 1$, we define special partitions $p^+_\rho(m)$ and $p^-_\rho(m)$. (See Lemma 2.2.5 for partitions relevant to us, or [10, Section 3.9] for the general definition.) We say that C satisfies the *partition condition* if $p^+_i(C) = p^+_{\alpha_i}(m_i)$ and $p^-_j(C) = p^-_{\beta_j}(n_j)$.

Let α and β be admissible orbit sets and let $\mathcal{M}_1(\alpha, \beta)$ be the moduli space of holomorphic currents $\mathcal{C}$ with $I(\alpha, \beta, [\mathcal{C}]) = 1$. The key consequence of property (d) and property (e) above is that, for a generic J, any $\mathcal{C} \in \mathcal{M}_1(\alpha, \beta)$ can be written as the disjoint union $C' \sqcup \mathcal{T}$ where $\mathcal{T}$ is trivial and C' is an irreducible embedded holomorphic curve with $\mathrm{ind}(C') = 1$. We also have that $\mathcal{M}_1(\alpha, \beta)/\mathbb{R}$ is compact by a modified version of Gromov compactness (See [10, Section 5.3]). Hence, we can define the differential coefficient between α and β by

$$\langle \partial\alpha, \beta \rangle := \#\mathcal{M}_1(\alpha, \beta)/\mathbb{R}.$$

Lastly, we define an action $\mathcal{A}(\alpha)$ of an orbit set $\alpha = \{(\alpha_i, m_i)\}$ by

$$\mathcal{A}(\alpha) := \sum_i m_i \int_{\alpha_i} \lambda.$$

If $u : (\Sigma, j) \to (\mathbb{R} \times Y, J)$ is a holomorphic curve from α to β, we have

$$\mathcal{A}(\alpha) - \mathcal{A}(\beta) = \int_\Sigma u^*(d\lambda) \geq 0$$

by Stokes' theorem and so the ECH chain complex is filtered by the action of its generators. For each $L > 0$, the filtered ECH chain complex ECC^L_* is the subcomplex of ECC_* which consists only of generators with action less than L. We can recover ECH_* as the direct limit of ECH^L_* as $L \to \infty$. For many subsequent arguments, we rely on being able to disregard any orbit with action greater than L. Hence, throughout the paper, we will always assume a filtered version of ECH for some fixed $L > 0$.

Morse-Bott theory

We now return to the contact manifold $(I \times T^2, \bar{\lambda})$ where $I = [X_w, X_e]$ and $\bar{\lambda} = -g dt_1 + f dt_2$ is a T^2-invariant contact form. In order to define the ECH of this contact manifold, we need to perturb the degenerate contact form $\bar{\lambda}$ to a nondegenerate contact form λ. Recall also that the definition of ECH requires the choice of a generic admissible almost complex structure J on $\mathbb{R} \times (I \times T^2)$. The goal of this section is to describe the perturbations and almost complex structures that will result in a nice combinatorial description of the ECH chain complex.

Before describing the perturbation, we parametrize all S^1-families of Reeb orbits simultaneously by the following function Θ, extending [11, Appendix A]. For an S^1-family $\bar{\rho}$ at $x = x_0$, let $T_0 := \{x_0\} \times T^2$ and suppose $\rho \in \bar{\rho}$ has the homology class $(p, q) \in \mathbb{Z}^2 = H_1(T_0)$. Let $wedge(p, q)$ be a wedge of p-fold covered circle at $(\mathbb{R}/\mathbb{Z}) \times \{0\}$ and q-fold covered circle at

$\{0\} \times (\mathbb{R}/\mathbb{Z})$ in $T_0 = (\mathbb{R}/\mathbb{Z})^2$. Let S be any surface in T_0 such that $\partial S = \rho \cup (-wedge(p,q))$. For example, one can choose the trapezoid with vertices at $(0,0), (r,0), (p+r,q)$ and $(0,q)$ in the universal cover $\mathbb{R}^2 \to (\mathbb{R}/\mathbb{Z})^2$ for an appropriate r. Then, we define $\Theta : \bar{\rho} \to \mathbb{R}/\mathbb{Z}$ by

$$\Theta(\rho) := \int_S dt_1 dt_2 \in \mathbb{R}/\mathbb{Z} \tag{2.1.9}$$

where the integral is independent of S modulo $\mathbb{Z}$. Explicitly,

$$\Theta(\rho) = (t_1, t_2) \times (p, q) + pq/2$$

for any $(x_0, t_1, t_2) \in \rho$. From here on, we always identify $\bar{\rho}$ as $\mathbb{R}/\mathbb{Z}$ using Θ. We note that, if we change the identification of the fiber $T^2 = (\mathbb{R}/\mathbb{Z})^2$ by $SL(2,\mathbb{Z})$, then Θ changes (simultaneously for all $\bar{\rho}$) by $\Theta \mapsto \Theta + c$ for $c = 0$ or $1/2$. In particular, we may later choose a convenient identification without affecting the analysis.

We now discuss the S^1 Morse-Bott theory following [2]. Assume generic f and g. Since there are only finitely many S^1-families of Reeb orbits of $\bar{\lambda}$ with action less than L, we describe the perturbation on a small neighborhood of each such family. Let $\bar{\rho}$ be an S^1-family of Reeb orbits at $x = x_0$ with action less than L. Regard a Morse function $H_{\bar{\rho}} : \mathbb{R}/\mathbb{Z} \to \mathbb{R}$ with two critical points as a function on $\{x_0\} \times T^2 = \cup_{\rho \in \bar{\rho}} \rho$. Extend $H_{\bar{\rho}}$ to a function $\tilde{H}_{\bar{\rho}}$ on $(x_0 - \epsilon, x_0 + \epsilon) \times T^2$ with a compact support and $\partial_x \tilde{H}_{\bar{\rho}} = 0$ near $\{x_0\} \times T^2$. Then, for $\epsilon, \eta > 0$ sufficiently small,

$$\lambda := (1 + \eta \tilde{H}_{\bar{\rho}})\bar{\lambda} \tag{2.1.10}$$

is a contact form on $(x_0 - \epsilon, x_0 + \epsilon) \times T^2$ with two nondegenerate Reeb orbits at $\operatorname{crit} H_{\bar{\rho}}$ and no other Reeb orbits of action less than L. Recall that the contact structure $\bar{\xi}$ of $\bar{\lambda}$ is a trivial symplectic 2-plane bundle with a fiber $\operatorname{span}\{\partial_x, -f\partial_{t_1} - g\partial_{t_2}\}$. We will always use this trivialization in this paper and call this τ. We compute that

$$\bar{R} = \frac{f'\partial_{t_1} + g'\partial_{t_2}}{(fg' - f'g)},$$
$$\mathcal{L}_{\partial_x} \bar{R} = \frac{f'g'' - f''g'}{(fg' - f'g)^2}(-f\partial_{t_1} - g\partial_{t_2}).$$

Hence, with respect to τ, the linearized flow of $\bar{R}$ along ρ, parametrized by ν, is $\nu \mapsto \begin{pmatrix} 1 & 0 \\ r\nu & 1 \end{pmatrix}$ with $r > 0$ if $f'g'' - f''g' > 0$ (convex) and with $r < 0$ if $f'g'' - f''g' < 0$ (concave). We conclude that:

- If $\bar{\rho}$ is convex, then λ has an elliptic orbit $\check{e}$ at $\max H_{\bar{\rho}}$ whose linearized return map is a small positive rotation with respect to τ, and a positive hyperbolic orbit $\check{h}$ at $\min H_{\bar{\rho}}$ whose linearized return map does not rotate with respect to τ.
- If $\bar{\rho}$ is concave, then λ has a positive hyperbolic orbit $\hat{h}$ at $\max H_{\bar{\rho}}$ whose linearized return map does not rotate with respect to τ, and an elliptic orbit $\hat{e}$ at $\min H_{\bar{\rho}}$ whose linearized return map is a small negative rotation with respect to τ.

Definition 2.1.2. (A good perturbation) Let $I = [X_w, X_e]$ and $\bar{\lambda} = -gdt_1 + fdt_2$ be a T^2-invariant contact form on $I \times T^2$. Fix $0 < \delta < 1/3$ and $L > 0$. Let $\epsilon, \eta > 0$. Let Ξ denote the (finite) set of S^1-families of Reeb orbits with action less than L and for each $\bar{\rho} \in \Xi$, let $H_{\bar{\rho}}$ and $\tilde{H}_{\bar{\rho}} = \tilde{H}_{\bar{\rho}}(\epsilon)$ be as above. We say that λ defined by

$$\lambda = \left(1 + \eta \sum_{\bar{\rho} \in \Xi} \tilde{H}_{\bar{\rho}}\right) \bar{\lambda}$$

is a *good perturbation* of $\bar{\lambda}$ if, for each $\bar{\rho} \in \Xi$:

(i) $H_{\bar{\rho}} : (\mathbb{R}/\mathbb{Z}) \to \mathbb{R}$ has exactly two critical points and

- if $\bar{\rho}$ is convex, $H_{\bar{\rho}}$ attains the minimum at 0 and the maximum at δ, and
- if $\bar{\rho}$ is concave, $H_{\bar{\rho}}$ attains the maximum at 0 and the minimum at $-\delta$.

(ii) ϵ is sufficiently small that

- $\bar{\lambda}$ does not have any point of inflection on $(x(\bar{\rho}) - \epsilon, x(\bar{\rho}) + \epsilon)$.
- If α and β are orbit sets with action less than L and if $\bar{R}(x)$ is a multiple of $[\alpha]-[\beta]$ for some $x \in (x(\bar{\rho}) - \epsilon, x(\bar{\rho}) + \epsilon)$, then $x = x(\bar{\rho})$.

(iii) ϵ and $\eta = \eta(\epsilon, \tilde{H}_\rho)$ are sufficiently small for (2.1.10).

(iv) η is sufficiently small that the linearized return angle ϕ of the elliptic orbit from $\bar{\rho} \in \Xi$ satisfies $|\phi| < 2\pi/\lceil L/\mathcal{A}(\rho)\rceil$.

Condition (ii) is used for the positivity lemma 2.2.1. Condition (iv) is the simplifying assumption for the Conley-Zehnder indices. Condition (i) is used in the last step of Section 2.2 to rule out $\mathbb{R}$-relaxing as mentioned in Example 1.4.2 (iv).

To describe the almost complex structure, consider the admissible almost complex structure $\bar{J}$ on the symplectization $\mathbb{R} \times (I \times T^2, \bar{\lambda})$ defined by $\bar{J}(\partial_s) = \bar{R}$ and

$$\bar{J}(\partial_x) = -f\partial_{t_1} - g\partial_{t_2}. \tag{2.1.11}$$

This has the property

$$\bar{R} \times \bar{J}(\partial_x) = \frac{1}{fg' - f'g}(f', g') \times (-f, -g) = 1 > 0. \tag{2.1.12}$$

We pick a generic admissible almost complex structure J on the symplectization $\mathbb{R} \times (I \times T^2, \lambda)$, which is a small perturbation of $\bar{J}$.

Definition 2.1.3. Let $\bar{\lambda}$ be a T^2-invariant contact structure on $I \times T^2$ and $\bar{J}$ be the distinguished almost complex structure for $\bar{\lambda}$ satisfying (2.1.12). We say that the pair (λ, J) is a *good perturbation* of $(\bar{\lambda}, \bar{J})$ if λ is a good perturbation of $\bar{\lambda}$ and, additionally, λ and J are sufficiently close to $\bar{\lambda}$ and $\bar{J}$ in the sense of Lemma 2.2.1, Proposition 2.2.21 and Proposition 2.3.5.

This requirement is necessary to prove positivity of relevant IP regions and to relate the holomorphic curves in the perturbed setup to those in the unperturbed setup for the Morse-Bott complex.

2.2 Proof of necessity

In this section, we fix $I = [X_w, X_e]$, a T^2-invariant contact form $\bar{\lambda}$ on $I \times T^2$ with the Reeb vector field $\bar{R}$, the distinguished almost complex structure $\bar{J}$ on $\mathbb{R} \times (I \times T^2)$ defined by (2.1.11) and a good perturbation λ of $\bar{\lambda}$. After Lemma 2.2.1, we will also assume J is sufficiently close to $\bar{J}$ that the assertions of Lemma 2.2.1 holds.

We start by proving the following important property satisfied by any IP region $\mathcal{R}_{\alpha,\beta}$ associated to orbit sets α and β with nonempty $\mathcal{M}(\alpha, \beta)$. This is an adaptation of [11, Lemma 3.11].

Lemma 2.2.1. *(Positivity) Let α and β be orbit sets of λ in the homology class Γ and suppose J is sufficiently close to $\bar{J}$. Suppose $C \in \mathcal{M}(\alpha, \beta)$ is a holomorphic curve from α to β. Then, $\mathcal{R}_{\alpha,\beta}$ is positive with respect to $\bar{\lambda}$, i.e. for all $x \in I$,*

$$\bar{R}(x) \times \sigma(x) \geq 0.$$

Moreover, if λ is unperturbed near $\{x_0\} \times T^2$, then the equality holds at $x = x_0$ if and only if $\mathcal{S}_C(x_0) = \emptyset$.

In particular, the equality condition applied to $x = X_w$ and $x = X_e$ implies that we have Gromov compactness for the moduli space of holomorphic curves that do not intersect $\mathbb{R} \times \{X_w, X_e\} \times T^2$. This result can also be interpreted as intersection positivity of C with the leaves of symplectic foliation given by $\mathbb{R}$ cross the Reeb flow [11].

Remark 2.2.2. The second assertion fails if λ is perturbed near $\{x_0\} \times T^2$. For example, a holomorphic curve corresponding to an auxiliary Morse flow of $H_{\bar{\rho}}$ satisfies $\bar{R}(x) \times \sigma(x) = 0$ for all $x \in I$ but it does not even stay within $\mathbb{R} \times \{x(\bar{\rho})\} \times T^2$. All holomorphic curves are affected similarly under the perturbation.

Proof. Suppose λ is unperturbed near $\{x_0\} \times T^2$ and x_0 is regular for the projection $\pi_x|_C : \Sigma \to I$. Let ν parametrize a component $\mathcal{S}'$ of $\mathcal{S}_C(x_0)$ and consider the ordered basis

$$\{\partial_s, \bar{R}, \partial_x, \bar{J}\partial_x\}$$

of $T_p(\mathbb{R} \times I \times T^2)$ at any point $p \in \mathcal{S}'$. By the orientation convention on slices, $J^{-1}\partial_\nu \in \text{span}\{\partial_s, \partial_x\}$ has a positive ∂_x component. Hence, $\partial_\nu \in \text{span}\{\bar{R}, \bar{J}\partial_x\}$ has a positive $\bar{J}\partial_x$ component for J sufficiently close to $\bar{J}$. By (2.1.12),

$$\bar{R}(x) \times \partial_\nu > 0.$$

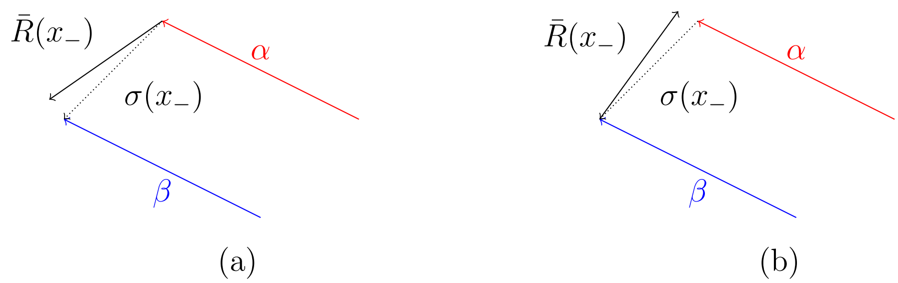

Figure 2.1: Two possible scenarios just before $R(x) \times \sigma(x) = 0$.

We obtain both results by integrating this along ν and summing over all components. For a non-regular x_0, take the limit of this result for regular x_i's with $\lim_{i\to\infty} x_i \to x_0$.

We now argue for the first assertion when $x_0 \in I_\epsilon := [x(\rho) - \epsilon, x(\rho) + \epsilon]$ for some Reeb orbit ρ. Since $\sigma|_{I_\epsilon}$ can jump only at $x = x(\rho)$ and only by a multiple of $\bar{R}(x(\rho))$, the function

$$A(x) := \bar{R}(x) \times \sigma(x)$$

is defined continuously on I_ϵ and is non-negative at the two endpoints of I_ϵ. Suppose $A(x_0) = 0$ for some $x_0 \in I_\epsilon$ with $\sigma(x_0) \neq 0$. By condition (ii) of Definition 2.1.2, $x_0 = x(\rho)$. Since $\bar{\lambda}$ does not have any point of inflection on I_ϵ, $A(x)$ cannot take a minimum on $I_\epsilon \setminus \{x(\rho)\}$, so $A(x)$ must stay non-negative throughout I_ϵ. □

Corollary 2.2.3. *Let $\mathcal{R}_{\alpha,\beta}$ be as in Lemma 2.2.1 with an induced decoration and suppose that $\sigma(x(\rho) \pm \epsilon)$ are not both zero for some embedded orbit $\rho \in \alpha \cup \beta$. If $\bar{R}(x(\rho)) \times \sigma(x(\rho)) = 0$, then $\mathcal{R}_{\alpha,\beta}$ must have at least one $\mathbb{R}$-loose edge at $x = x(\rho)$. Furthermore, if $\mathcal{R}_{\alpha,\beta}$ is indecomposable, it has extreme edges at $x = x(\rho)$.*

Proof. Write $x_\pm := x(\rho) \pm \epsilon$. By the assumption, $\sigma(x_+) = k[\rho]$ and $\sigma(x_-) = l[\rho]$ for some integers k and l. By symmetry, assume $l \neq 0$ and define

$$B(x) := \bar{R}(x) \times \sigma(x_-).$$

By Lemma 2.2.1, $B(x_-) > 0$. If $B(x_+) \geq 0$, then by condition (ii) of Definition 2.1.2, $B(x)$ is positive on $[x_-, x(\rho)) \cup (x(\rho), x_+]$ and zero at $x = x(\rho)$. This contradicts genericity of $\bar{\lambda}$. Thus, $B(x_+) < 0$.

Now assume ρ is convex, i.e. $\bar{R}(x_-) \times [\rho] > 0$. We have $l > 0$ since

$$0 < \bar{R}(x_-) \times [\rho] = (1/l) \cdot B(x_-)$$

and $k \leq 0$ since

$$0 \leq \bar{R}(x_+) \times \sigma(x_+) = (k/l) \cdot B(x_+).$$

At $x = x(\rho)$, $\mathbb{R}$-loose edges are positive edges and by (1.3.2), the number of positive edges at $x = x(\rho)$ is at least $l - k > 0$. For the second assertion, if $\mathcal{R}_{\alpha,\beta}$ is indecomposable, kl cannot be strictly negative so $k = 0$. This completes the proof for $l \neq 0$ and a convex ρ. The other cases can be argued similarly. Figure 2.1 illustrates two possible slice classes at $x = x_-$ for a convex ρ and a concave ρ, respectively. □

We introduce some notations which are convenient when dealing with indices.

Definition 2.2.4. (a) The "signed" combinatorial Conley-Zehnder index for an edge v of an IP region $\mathcal{R}$ is defined as

$$cz_{\mathcal{R}}(v) := \pm CZ(v) \text{ if } v \in \partial^{\pm}\mathcal{R}.$$

We simply write $cz(v)$ when $\mathcal{R}$ is clear. Similarly, if C is a holomorphic curve, then we define the "signed" Conley-Zehnder index for a positive/negative end of C at $\rho^{\pm}$ by

$$cz_C^{ind}(\rho^{\pm}) := \pm CZ_{\tau}(\rho^{\pm}).$$

We simply write $cz^{ind}(\rho^{\pm})$ when C is clear.

(b) We rewrite the combinatorial ECH index (1.3.4) of an IP region $\mathcal{R}$ as

$$I(\mathcal{R}) = I^a(\mathcal{R}) + I^c(\mathcal{R}) \tag{2.2.1}$$

where

$$I^a(\mathcal{R}) := 2Area(\mathcal{R}) - \#\{\text{edges of } \mathcal{R}\} \tag{2.2.2}$$

and

$$I^c(\mathcal{R}) := \sum_{v \in \partial^{\pm}\mathcal{R}} (cz(v) + 1). \tag{2.2.3}$$

Similarly, rewrite the Fredholm index formula (2.1.3) as

$$\begin{aligned} \text{ind}(C) &= [2g(C) - 2 + \#\{\text{ends of } C\}] + 0 + \sum_{\rho} cz^{ind}(\rho) \\ &= 2g(C) - 2 + \sum_{\rho} (cz^{ind}(\rho) + 1), \end{aligned} \tag{2.2.4}$$

where the sum is over positive and negative ends ρ of C and $c_\tau([C]) = 0$ since τ is the restriction of a global trivialization of ξ.

The upshot is that each summand in (2.2.3) and in the summation of (2.2.4) is nonnegative, as we show next.

Lemma 2.2.5. *Let ρ be an embedded orbit of λ. Write $\rho = \check{e}, \check{h}, \hat{h}$ or $\hat{e}$ depending on whether it is elliptic convex, hyperbolic convex, hyperbolic concave or elliptic concave.*

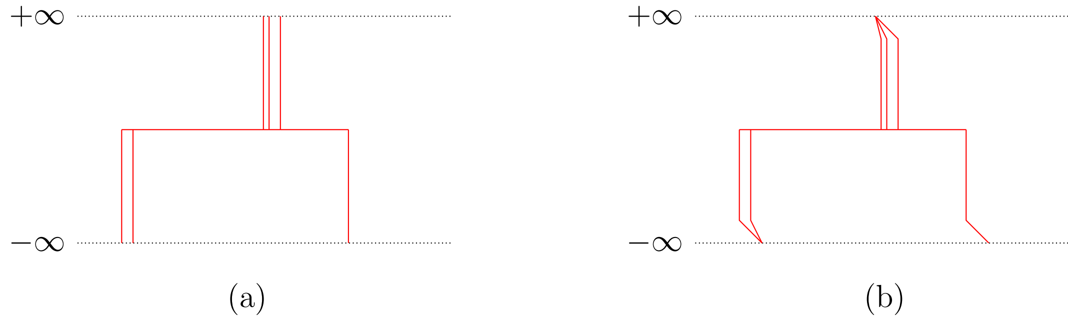

Figure 2.2: Image of S under the projection to $\mathbb{R} \times I$ before and after the final step.

(a) If $\mathcal{A}(\rho^k) < L$ for some $k \geq 1$, then

$$CZ_\tau(\rho^k) = \begin{cases} 1 & \text{if } \rho = \check{e}, \\ 0 & \text{if } \rho = \check{h}, \\ 0 & \text{if } \rho = \hat{h}, \\ -1 & \text{if } \rho = \hat{e}. \end{cases}$$

(b) The special partitions relevant to us are:

$$\begin{aligned} p^+_{\hat{e}}(m) = p^-_{\check{e}}(m) &= (m), \\ p^+_{\check{e}}(m) = p^-_{\hat{e}}(m) &= (1, \cdots, 1), \\ p^\pm_{\check{h}}(m) = p^\pm_{\hat{h}}(m) &= (1, \cdots, 1). \end{aligned} \tag{2.2.5}$$

Proof. (a) follows from condition (iv) of Definition 2.1.2. We refer to [10] or [7] for (b). ☐

Proposition 2.2.6. *(ECH index computation) Let α and β be orbit sets and consider $\mathcal{R}_{\alpha,\beta}$ with an induced decoration. If $Z \in H_2(I \times T^2, \alpha, \beta)$, the ECH index $I(\alpha, \beta, Z)$ is independent of Z and*

$$I(\alpha, \beta, Z) = I(\mathcal{R}_{\alpha,\beta}).$$

We write $I(\alpha, \beta) = I(\alpha, \beta, Z)$.

Proof. Since ξ is trivial and the generator $[T^2]$ of $H_2(I \times T^2) \cong \mathbb{Z}$ has algebraic intersection number zero with every orbit, by the index ambiguity formula (2.1.5),

$$I(\alpha, \beta, Z) - I(\alpha, \beta, Z') = \langle c_1(\xi) + 2PD(\Gamma), Z - Z' \rangle = 0,$$

so $I(\alpha, \beta, Z)$ is independent of Z.

Let $\alpha = \{(\alpha_i, m_i)\}$ and $\beta = \{(\beta_j, n_j)\}$. We construct a surface S in $[-\infty, \infty] \times (I \times T^2)$ (embedded except at $\pm\infty$) to represent Z. We start with half-cylinders

$$[0, \infty] \times \{x(\alpha_i) + \epsilon/k\} \times \pi_{T^2}(\alpha_i)$$

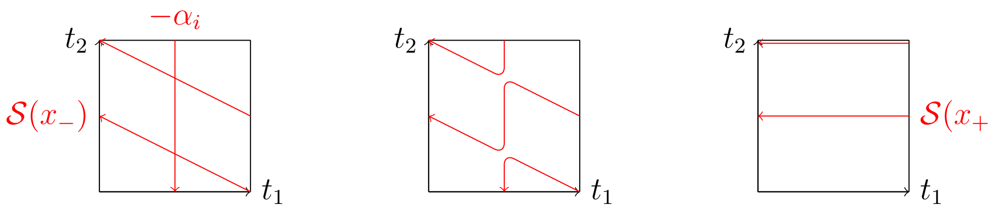

Figure 2.3: A surgery contributing $[\alpha_i] \times [\mathcal{S}(x_-)] = 2$ to $c_1(NS, \tau)$.

for each α_i and $1 \leq k \leq m_i$, and half-cylinders

$$[-\infty, 0] \times \{x(\beta_j) - \epsilon/k\} \times \pi_{T^2}(\beta_j)$$

for each β_j and $1 \leq k \leq n_j$. The $\pm\epsilon/k$ terms are arbitrarily chosen perturbations to ensure the half-cylinders are pairwise disjoint. We construct S as a union of the above half-cylinders and a movie of curves $\mathcal{S}(x)$ in $\{0\} \times \{x\} \times T^2$. Away from the half-cylinders, $\mathcal{S}(x)$ is a (possibly empty) disjoint union of straight embedded curves in T^2. See Figure 2.2 (a) for the projection of S to $\mathbb{R} \times I$. In this example, $\alpha = \{(\alpha_1, 3)\}$ with $x(\alpha_1) = 1$ and $\beta = \{(\beta_1, 2), (\beta_2, 1)\}$ with $x(\beta_1) = 0$ and $x(\beta_2) = 2$. The fiber over each point away from the trivalent vertices is a disjoint union of straight embedded curves.

Suppose there is exactly one half-cylinder between x_- and x_+, say $[0, \infty] \times \{x_0\} \times \pi_{T^2}(\alpha_i)$. We obtain $\mathcal{S}(x_+)$ from $\mathcal{S}(x_-)$ as follows:

(i) If $\mathcal{S}(x_-)$ and α_i are parallel, then simply add or remove a component to/from $\mathcal{S}(x_-)$.

(ii) Otherwise, we perform a "surgery": the boundary of the half-cylinder at $\{0\} \times \{x_0\} \times T^2$ is $\{x_0\} \times \pi_{T^2}(-\alpha_i)$. We resolve each intersection of $\mathcal{S}(x_-)$ and $-\alpha_i$ and linearly interpolate $\mathcal{S}(x)$ between x_0 and x_+. See Figure 2.3.

The case of a half-cylinder for β_j is similar. We have constructed a surface with boundaries $\{+\infty\} \times \{x(\alpha_i) + \epsilon/k\} \times \pi_{T^2}(\alpha_i)$ and $\{-\infty\} \times \{x(\beta_j) - \epsilon/k\} \times \pi_{T^2}(\beta_j)$. As a final step, we deform this surface so that it has boundaries $\{\infty\} \times \alpha_i$ and $\{-\infty\} \times \beta_j$. We can keep the projection of this surface to T^2 unchanged during the deformation while the projection to $\mathbb{R} \times I$ is interpolated linearly between (a) and (b) in Figure 2.2.

We now compute each of the three terms in the ECH index formula (2.1.1) using S. Since τ is the restriction of a global trivialization of ξ,

$$c_\tau(Z) = 0.$$

For the Q_τ term, the ends of S have writhe zero by construction, so

$$Q_\tau(\alpha, \beta) = c_1(NS, \tau),$$

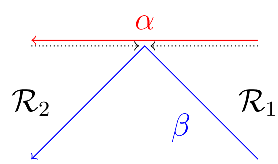

Figure 2.4: A decomposable IP region $\mathcal{R}_{\alpha,\beta} = \mathcal{R}_1\mathcal{R}_2$. ($\mathcal{R}_i$ may be bigons.)

where NS is the normal bundle to S. See [7, Section 2.7] for details. It has a section $\pi_{NS}(\partial_s + \partial_x)$ which is non-vanishing except at the points of resolution. At each half-cylinder $[0,\infty] \times \{x_0\} \times [\rho]$ or $[-\infty, 0] \times \{x_0'\} \times [\rho]$, we check that the sign of the zeroes agrees with the sign of $[\rho] \times [\mathcal{S}(x_-)]$ so the signed count is simply $[\rho] \times [\mathcal{S}(x_-)] = [\rho] \times \sigma(x_-)$. Hence, the signed count for the kth half-cylinder is equal to the area of $\Phi([0,1] \times [k-1,k])$ for a realization Φ of $\mathcal{R}_{\alpha,\beta}$. By summing over all half-cylinders, we get

$$c_1(NS, \tau) = 2Area(\mathcal{R}).$$

Finally, by Lemma 2.2.5, the Conley-Zehnder term is

$$\sum_i \sum_{k=1}^{m_i} CZ_\tau(\alpha_i^k) - \sum_j \sum_{k=1}^{n_j} CZ_\tau(\beta_j^k) = \sum_{v \in \partial^\pm \mathcal{R}} cz(v).$$

□

Proposition 2.2.7. *Let α and β be admissible orbit sets with $I(\alpha,\beta) = 1$. If $C \in \mathcal{M}(\alpha,\beta)$ is irreducible, then $\mathcal{R}_{\alpha,\beta}$ is indecomposable.*

Proof. Since C is somewhere injective (see Section 2.1), the index inequality (2.1.7) implies that $I(C) = \text{ind}(C) = 1$ and C satisfies the partition condition. Hence, by the Fredholm index formula (2.2.4),

$$\sum_\rho (cz^{ind}(\rho) + 1) \leq 3. \tag{2.2.6}$$

Write $\mathcal{R}_{\alpha,\beta} = \mathcal{R}_1 \cdots \mathcal{R}_n$ with indecomposable $\mathcal{R}_j$ and suppose $n \geq 2$. Since C is irreducible, the equality condition of Lemma 2.2.1 implies that the east extreme edges of $\mathcal{R}_1$ and the west extreme edges of $\mathcal{R}_2$ occur at the same $x = x(\rho)$ for some ρ. Note that each of $\mathcal{R}_1$ and $\mathcal{R}_2$ is either local or has an $\mathbb{R}$-loose edge at $x = x(\rho)$ by Corollary 2.2.3. Either way, each has an $\mathbb{R}$-loose edge $x = x(\rho)$. By symmetry, assume that ρ is convex, which means that $\mathbb{R}$-loose edges are positive edges. See Figure 2.4.

We consider an induced decoration of $\mathcal{R}_{\alpha,\beta}$. Since α is admissible, at most one of the positive edges at $x = x(\rho)$ can be hyperbolic. Also, if two of them are elliptic, then by the partition condition with $p_{\tilde{e}}^+ = (1, \cdots, 1)$ in Lemma 2.2.5, the two edges belong to distinct ends of C and give at least two summands in (2.2.6) with $(cz^{ind}(\rho)+1) = 2$. Since $\text{ind}(C) = 1$, this cannot happen and $\mathcal{R}_{\alpha,\beta}$ must have exactly two positive edges at $x = x(\rho)$, with one being hyperbolic and the other elliptic. Moreover, all other edges of $\mathcal{R}_{\alpha,\beta}$ have $(cz(v)+1) = 0$. In particular, no other edges can be extreme edges of a nonlocal IP region, which are necessarily $\mathbb{R}$-loose by Corollary 2.2.3. This implies that all $\mathcal{R}_j$ must be local.

In order to have $I(\mathcal{R}_{\alpha,\beta}) = 1$, $\mathcal{R}_{\alpha,\beta}$ must have an odd number of hyperbolic edges and since β is also admissible, all negative edges of $\mathcal{R}_{\alpha,\beta}$ must be elliptic while exactly one of the positive edges is hyperbolic. This makes $I(\mathcal{R}_{\alpha,\beta}) = -1$, which is a contradiction. □

Corollary 2.2.8. *Let α and β be admissible orbit sets with $I(\alpha, \beta) = 1$.*

(a) If we can write $\alpha = \alpha'\rho$ and $\beta = \beta'\rho$ for some embedded orbit ρ, there is a bijection

$$\mathcal{M}(\alpha, \beta) \cong \mathcal{M}(\alpha', \beta').$$

The same conclusion holds for $\alpha = \rho\alpha'$ and $\beta = \rho\beta'$.

(b) If $\mathcal{C} \in \mathcal{M}(\alpha, \beta)$ is reducible, then α and β can be written as in (a).

Proof. (a) Each distinct holomorphic current $\mathcal{C}'$ from α' to β' gives a distinct $\mathcal{C}' \cup (\mathbb{R} \times \rho) \in \mathcal{M}(\alpha, \beta)$. It remains to show that every $\mathcal{C} \in \mathcal{M}(\alpha, \beta)$ arises this way. Recall from Section 2.1 that $\mathcal{C} \in \mathcal{M}(\alpha, \beta)$ contains a single component C' with $I(C') = 1$ and all other components are trivial. Hence, if there are no trivial cylinders at $x = x(\rho)$, then all ends at $x = x(\rho)$ must be ends of C' and the IP region associated to C' can be written as $\mathcal{R}''\mathcal{R}_{\rho,\rho}$ for some IP region $\mathcal{R}''$. This contradicts the conclusion of Proposition 2.2.7.

(b) Write $\mathcal{C} \in \mathcal{M}(\alpha, \beta)$ as $C' \sqcup \mathcal{T}$ where $\mathcal{T}$ is trivial. Let $\mathcal{R}'$ be the IP region corresponding to C' and suppose $T \in \mathcal{T}$ is a trivial cylinder with ends at some embedded orbit ρ. Since $T \cap C' = \emptyset$, $[\rho] \times [\sigma_{\mathcal{R}'}(x(\rho))] = 0$. By Corollary 2.2.3, C' has extreme ends at ρ.

□

In view of Proposition 2.2.7 and Corollary 2.2.8, we assume for the rest of the section that α and β are admissible orbit sets with $I(\alpha, \beta) = 1$ and that $\mathcal{R}_{\alpha,\beta}$ is indecomposable.

Lemma 2.2.9. *Suppose $\mathcal{R}$ is a positive indecomposable IP region. Then $I^a(\mathcal{R})$ is even and $I^a(\mathcal{R}) \geq -2$ with equality if and only if $\mathcal{R}$ is minimal.*

Proof. Let Φ be a realization of $\mathcal{R}$. Both assertions follow from

$$I^a(\mathcal{R}) = 2\#\{\text{internal lattice points of } \Phi\} - 2,$$

which is the consequence of Pick's theorem. □

Proposition 2.2.10. *Let α and β be admissible orbit sets with $I(\alpha, \beta) = 1$ and suppose $\mathcal{R}_{\alpha,\beta}$ is nonlocal and indecomposable. If $\mathcal{M}(\alpha, \beta) \neq \emptyset$, then:*

(a) $\mathcal{R}_{\alpha,\beta}$ has a single west extreme edge and a single east extreme edge.

(b) $\mathcal{R}_{\alpha,\beta}$ *is minimal and an induced decoration of* $\mathcal{R}_{\alpha,\beta}$ *can be obtained from the minimal decoration of* $\mathcal{R}_{\alpha,\beta}$ *by* S^1*-relaxing one edge or* $\mathbb{R}$*-relaxing one non-extreme edge.*

(c) Any $C \in \mathcal{M}(\alpha, \beta)$ *has genus zero.*

Proof. (a) Consider $\mathcal{R}_{\alpha,\beta}$ with an induced decoration. Since $\mathcal{R}_{\alpha,\beta}$ is positive, $I^c(\mathcal{R}_{\alpha,\beta}) \leq 3$ by Lemma 2.2.9 and (2.2.1). Since $\mathcal{R}_{\alpha,\beta}$ is indecomposable, each extreme edge v is $\mathbb{R}$-loose by Corollary 2.2.3, and so $(cz(v)+1) \geq 1$. If v and v' are two distinct west extreme edges of $\mathcal{R}_{\alpha,\beta}$ and v'' is an east extreme edge of $\mathcal{R}_{\alpha,\beta}$, then $3 \leq (cz(v)+1) + (cz(v')+1) + (cz(v'')+1) \leq 3$, so $cz(v) = cz(v') = 0$ and both v and v' must be hyperbolic. This contradicts the admissibility of α or β. A similar argument holds for the multiplicity of east extreme edges.

(b) As in (a), we have $I^c(\mathcal{R}_{\alpha,\beta}) \geq (cz(v)+1) + (cz(v')+1) \geq 2$ for the two extreme edges v and v'. Hence, from Lemma 2.2.9,

$$-2 \leq I^a(\mathcal{R}_{\alpha,\beta}) = I(\mathcal{R}_{\alpha,\beta}) - I^c(\mathcal{R}_{\alpha,\beta}) \leq -1.$$

Since $I^a(\mathcal{R}_{\alpha,\beta})$ is even, we conclude that $I^a(\mathcal{R}_{\alpha,\beta}) = -2$ and $I^c(\mathcal{R}_{\alpha,\beta}) = 3$. The result follows from the fact that $I^c(\mathcal{R}_{\alpha,\beta}) \geq 2$ with equality only if $\mathcal{R}_{\alpha,\beta}$ is minimally decorated.

(c) By the hypothesis and Corollary 2.2.8, C is irreducible, somewhere injective and nonlocal. Since each extreme end of C contributes $(cz^{ind}(\rho)+1) \geq 1$ in (2.2.4) and $\mathrm{ind}(C) = 1$, we must have $g(C) = 0$.

□

Proposition 2.2.10 almost proves the necessity part of Theorem 1.4.1. See the non-examples in Example 1.4.2 for the cases we still need to consider. We deal with these cases in the next section using an argument from Morse-Bott theory which exploits condition (i) of a good perturbation λ and a good perturbation pair (λ, J). (Definition 2.1.2 and Definition 2.1.3.)

A Morse-Bott argument

Before proceeding with the argument, we first establish some definitions and notations. In this section, consider $I = [X_w, X_e]$, a T^2-invariant contact form $\bar{\lambda}$ on $I \times T^2$ and the distinguished almost complex structure $\bar{J}$ defined by (2.1.11). For each S^1-family of orbits $\bar{\rho}$ of $\bar{\lambda}$, let $\bar{\rho}(\theta_0) \in \bar{\rho}$ denote the orbit corresponding to $\theta_0 \in \mathbb{R}/\mathbb{Z}$ via Θ. An *orbit set* of $\bar{\lambda}$ in the homology class $\Gamma \in H_1(I \times T^2)$ is a finite set of pairs $\{(\bar{\gamma}_i(\theta_i), m_i)\}$, where $\bar{\gamma}_i$ is an S^1-family of orbits and $\theta_i \in \mathbb{R}/\mathbb{Z}$, so that $\sum_i m_i[\bar{\gamma}_i(\theta_i)] = \Gamma$. We can write it in the ordered product notation

$$\bar{\gamma}_1(\theta_1) \cdots \bar{\gamma}_n(\theta_n)$$

where $x(\bar{\gamma}_i)$ is nondecreasing. We denote this orbit set by $\bar{\gamma}(\theta)$ where $\bar{\gamma} = \bar{\gamma}_1 \cdots \bar{\gamma}_n$ is called a *family orbit set* and $\theta = (\theta_i) \in (\mathbb{R}/\mathbb{Z})^n$. Note that there is a unique way to write a family orbit set in the ordered product notation, while θ is unique only up to transposing θ_i and θ_j with $\bar{\gamma}_i = \bar{\gamma}_j$.

Definition 2.2.11. Let $I = [X_w, X_e]$ be an interval.

(a) A *partial decoration* of an IP path $\mathcal{P}$ on I is an association of each edge of $\mathcal{P}$ with one of the labels $\{\vee, \wedge\}$.

(b) Let $\mathcal{R}$ be an IP region on I. A *partial decoration* of $\mathcal{R}$ is a partial decoration of $\partial^+\mathcal{R}$ and $\partial^-\mathcal{R}$.

(c) An edge v of a partially decorated IP region $\mathcal{R}$ on I is said to be $\mathbb{R}$-*loose* if v is a positive edge labeled $\vee$ or a negative edge labeled $\wedge$. v is $\mathbb{R}$-*tight* otherwise.

A decoration of an IP path $\mathcal{P}$ gives a partial decoration of $\mathcal{P}$ by forgetting e/h labels and just keeping check ($\vee$) or hat ($\wedge$) labels.

Lemma 2.2.12. *Let $I = [X_w, X_e]$ be an interval and let $\bar{\lambda}$ be a T^2-invariant contact form on $I \times T^2$. There is a natural way to assign a unique IP path $\mathcal{P}_{\bar{\gamma}}$ on I to each family orbit set $\bar{\gamma}$ (or each orbit set $\bar{\gamma}(\theta)$) of $\bar{\lambda}$. Moreover, $\bar{\gamma}$ (or $\bar{\gamma}(\theta)$) induces a unique partial decoration on $\mathcal{P}_{\bar{\gamma}}$.*

Proof. Let $\mathcal{P} = (v_i)$ with $v_i = [\bar{\gamma}_i(0)] \in \mathbb{Z}^2$ and $x(v_i) = x(\bar{\gamma}_i)$. Label each v_i as $\vee$ if $\bar{\gamma}_i$ is convex and as $\wedge$ otherwise. □

Definition 2.2.13. We call $\mathcal{P}_{\bar{\gamma}}$ as in Lemma 2.2.12 the IP path *associated to* $\bar{\gamma}$ (or $\bar{\gamma}(\theta)$). We say that the partial decoration of $\mathcal{P}_{\bar{\gamma}}$ in Lemma 2.2.12 is *induced* by $\bar{\gamma}$ (or $\bar{\gamma}(\theta)$).

Consider $I \times T^2$ with a T^2-invariant contact form $\bar{\lambda}$ and an almost complex structure $\bar{J}$ on $\mathbb{R} \times I \times T^2$. A $\bar{J}$-holomorphic curve $\bar{C}$ from $\{(\bar{\alpha}_i(\theta_i^+), m_i)\}$ to $\{(\bar{\beta}_j(\theta_j^+), n_j)\}$ is a $\bar{J}$-holomorphic curve whose positive ends at covers of $\bar{\alpha}_i(\theta_i^+)$ have total multiplicity m_i and whose negative ends at covers of $\bar{\beta}_j(\theta_j^-)$ have total multiplicity n_j, with no other ends.

Definition 2.2.14. The IP region *associated* to a pair of family orbit sets $\bar{\alpha}$ and $\bar{\beta}$ is the IP region between $\mathcal{P}_{\bar{\alpha}}$ and $\mathcal{P}_{\bar{\beta}}$ and is denoted $\mathcal{R}_{\bar{\alpha},\bar{\beta}}$. An induced partial decoration of $\mathcal{R}_{\bar{\alpha},\bar{\beta}}$ is an induced partial decoration of $\mathcal{P}_{\bar{\alpha}}$ and $\mathcal{P}_{\bar{\beta}}$.

For each S^1-family of embedded orbits $\bar{\rho}$, let $H_{\bar{\rho}}$ be a generic Morse function on $\bar{\rho} \cong S^1$. For $m \geq 1$, let $\bar{\rho}^m := \{m\text{-fold cover of } \rho | \rho \in \bar{\rho}\}$ and let $H_{\bar{\rho}^m} = H_{\bar{\rho}}$ be a Morse function on $\bar{\rho}^m$ under the identification (m-fold cover of ρ) $\leftrightarrow \rho$. A *$\bar{J}$-holomorphic building* $\bar{C}$ with $H_{\bar{\rho}^m}$ is a sequence of $\bar{J}$-holomorphic curves $\bar{C}^1, \cdots, \bar{C}^k$ such that:

(i) Each end of $\bar{C}^i$ converges to ϱ for some $\varrho \in \bar{\rho}^m$.

(ii) For $1 < i < k$, there is a bijection between the negative ends of $\bar{C}^i$ and the positive ends of $\bar{C}^{i+1}$. For each such pair, both ends converge to orbits in the same $\bar{\rho}^m$ and there is a downward flow of $H_{\bar{\rho}^m}$ from the negative end of $\bar{C}^i$ to the positive end of $\bar{C}^{i+1}$.

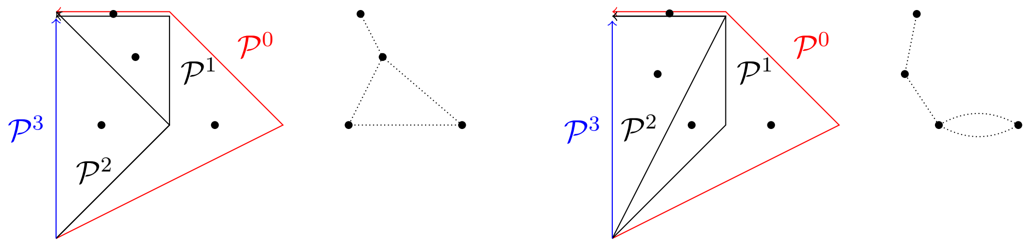

(a) (b)

Figure 2.5: Two partitions of $\mathcal{R} = (\mathcal{P}^0, \mathcal{P}^3)$ and the respective collapsed dual graphs.

(iii) For each positive end of $\bar{C}^1$ at some $\varrho \in \bar{\rho}^m$, there is a (possibly constant) downward flow of $H_{\bar{\rho}^m}$ from a critical point of $H_{\bar{\rho}^m}$ to ϱ. For each negative end of $\bar{C}^k$ at some $\varrho \in \bar{\rho}^m$, there is a (possibly constant) downward flow of $H_{\bar{\rho}^m}$ from ϱ to a critical point of $H_{\bar{\rho}^m}$.

Definition 2.2.15. Suppose $\mathcal{P}^i$ is an IP path on I for each $0 \leq i \leq k$ and $\mathcal{R}^i$ is a positive IP region on I between $\mathcal{P}^{i-1}$ and $\mathcal{P}^i$ for each $1 \leq i \leq k$. We write each $\mathcal{R}^i$ as

$$\mathcal{R}^i = \mathcal{R}^i_1 \cdots \mathcal{R}^i_{n_i}$$

where each $\mathcal{R}^i_j$ is an indecomposable IP region.

(a) We call the list $(\mathcal{R}^i_j) = ((\mathcal{R}^i_{j_i}))$ a *partition* of the IP region $\mathcal{R}$ between $\mathcal{P}^0$ and $\mathcal{P}^k$.

(b) The *dual graph* of a partition $(\mathcal{R}^i_j)$ is the graph characterized by:

- There is a vertex for each $\mathcal{R}^i_j$ and
- For a pair of vertices corresponding to $\mathcal{R}^i_j$ and $\mathcal{R}^{i+1}_{j'}$, there is an edge between them for each *shared* edge v between $\mathcal{R}^i_j$ and $\mathcal{R}^{i+1}_{j'}$, i.e. $v \in \partial^- \mathcal{R}^i_j$ and $v \in \partial^+ \mathcal{R}^{i+1}_{j'}$.

(c) The *collapsed dual graph* of $(\mathcal{R}^i_j)$ is obtained from the dual graph of $(\mathcal{R}^i_j)$ by the following procedure: for each vertex p corresponding to a local bigon of $(\mathcal{R}^i_j)$:

- If there is an edge e between p and another vertex p', then identify p and p' and remove the loop corresponding to e.
- Otherwise, remove p.

Figure 2.5 shows two partitions of an IP region $\mathcal{R}$ between $\mathcal{P}^0$ and $\mathcal{P}^3$. Each partition contains five IP regions, including one local bigon and one nonlocal bigon. The bigon (depicted with a gap) between $\mathcal{P}^0$ and $\mathcal{P}^1$ is nonlocal while the bigon (depicted with no gap) between $\mathcal{P}^1$ and $\mathcal{P}^2$ is local. Since the above $\mathcal{R}$ has one internal lattice and each nonlocal $\mathcal{R}^i_j$ in its partition is minimal, the (collapsed) dual graph contains one cycle. A decomposable $\mathcal{R}$ has a disconnected (collapsed) dual graph.

Remark 2.2.16. Compare this with the view by Taubes [24] and Parker [20], where a tropical curve is related to a dual graph of a certain triangulation of (a realization of) an IP region $\mathcal{R}$.

Lemma 2.2.17. *Consider* $(I \times T^2, \bar{\lambda})$*, the admissible almost complex structure* $\bar{J}$ *on* $\mathbb{R} \times I \times T^2$ *as in* (2.1.11) *and a Morse function* $H_{\bar{\rho}^m}$ *on each family of Reeb orbits* $\bar{\rho}^m$ *as in the definition of* $\bar{J}$*-holomorphic building. Let* $\bar{C}$ *be a* $\bar{J}$*-holomorphic building with* $H_{\bar{\rho}^m}$*. There is a natural way to assign to* $\bar{C}$ *a unique IP region* $\mathcal{R}$ *and a unique partition of* $\mathcal{R}$*. Moreover,* $\bar{C}$ *induces a unique partial decoration of each* $\mathcal{R}^i_j$*.*

Proof. Let $\bar{C} = (\bar{C}^1, \cdots, \bar{C}^k)$ where $\bar{C}^i$ is a $\bar{J}$-holomorphic curve from $\bar{\alpha}^i(\theta^i_+)$ to $\bar{\beta}^i(\theta^i_-)$. By the definition of a $\bar{J}$-holomorphic building with $H_{\bar{\rho}^m}$, $\bar{\alpha}^{i+1} = \bar{\beta}^i$ for each $1 \leq i < k$. Hence, we can set $\mathcal{P}^i := \mathcal{P}_{\bar{\alpha}^i}$ and $\mathcal{P}^k = \mathcal{P}_{\bar{\beta}^k}$ and each region between $\mathcal{P}^i$ and $\mathcal{P}^{i+1}$ is associated to $\bar{C}^i$, hence, is positive by Lemma 2.2.1. $\bar{C}^i$'s induce partial decorations of $\mathcal{P}^i$'s by Lemma 2.2.12 and hence, of each $\mathcal{R}^i_j$. $\mathcal{R}$ is the region between $\mathcal{P}^0$ and $\mathcal{P}^k$. □

Definition 2.2.18. We call $\mathcal{R}$, as in Lemma 2.2.17, the IP region *associated* to $\bar{C}$ and $(\mathcal{R}^i_j)$ the partition of $\mathcal{R}$ *associated* to $\bar{C}$.

We present two key lemmas. The first lemma restricts the complexity of a partition associated to a $\bar{J}$-holomorphic building.

Lemma 2.2.19. *Consider* $(I \times T^2, \bar{\lambda})$*,* $\bar{J}$ *and* $H_{\bar{\rho}^m}$ *as in Lemma 2.2.17. Let* $\bar{C}$ *be a* $\bar{J}$*-holomorphic building with* $H_{\bar{\rho}^m}$ *and let* $(\mathcal{R}^i_j)$ *be a partially decorated partition of* $\mathcal{R}$ *associated to* $\bar{C}$*. Suppose that* $\mathcal{R}$ *is minimal with* l $\mathbb{R}$*-loose edges and* $(\mathcal{R}^i_j)$ *contains* m *nonlocal regions. Then,*

$$m \leq l - 1.$$

The equality holds only if $\mathcal{R}$ *is indecomposable and each nonlocal* $\mathcal{R}^i_j$ *has exactly two* $\mathbb{R}$*-loose edges.*

Proof. We count the number of $\mathbb{R}$-loose edges in the partition. First, each nonlocal $\mathcal{R}^i_j$ contains at least 2 $\mathbb{R}$-loose edges. On the other hand, since $\mathcal{R}$ has no internal lattice point, the collapsed dual graph does not contain any cycles and, thus, has at most $m - 1$ edges. Each edge in the collapsed dual graph gives a shared edge v between two nonlocal IP regions $\mathcal{R}^i_j$ and $\mathcal{R}^{i'}_{j'}$ and by the definition of $\mathbb{R}$-tightness, v is $\mathbb{R}$-loose for exactly one of $\mathcal{R}^i_j$ or $\mathcal{R}^{i'}_{j'}$. Comparing these two counts, we get $2m \leq l + (m - 1)$. □

The second lemma is used to exploit the particular choice of auxiliary Morse functions in Definition 2.1.2. It is an adaptation of [11, Lemma A.2]:

Lemma 2.2.20. *(*Θ*-constraint) Consider* $(I \times T^2, \bar{\lambda})$ *and the admissible almost complex structure* $\bar{J}$ *on* $\mathbb{R} \times I \times T^2$ *defined by* (2.1.11)*. Let* $\bar{\alpha}(\theta^+) = \bar{\alpha}_1(\theta^+_1) \cdots \bar{\alpha}_m(\theta^+_m)$ *and* $\bar{\beta}(\theta^-) =$

$\bar{\beta}_1(\theta_1^-)\cdots\bar{\beta}_n(\theta_n^-)$ *be orbit sets of* $\bar{\lambda}$ *and let* $\bar{C}$ *be a* $\bar{J}$*-holomorphic curve from* $\bar{\alpha}(\theta^+)$ *to* $\bar{\beta}(\theta^-)$. *Then,*

$$\Theta(\bar{C}) := \sum_{i=1}^{m} \theta_i^+ - \sum_{j=1}^{n} \theta_j^- = 0 \in \mathbb{R}/\mathbb{Z}. \tag{2.2.7}$$

Proof. Recall $\bar{J}(\partial_x) = -f\partial_{t_1} - g\partial_{t_2}$ and consider any $p \in \mathbb{R}\times I\times T^2$. We check that $dsdx - dt_1dt_2$ annihilates $(v, \bar{J}v)$ for any $v \in T_p(\mathbb{R}\times I\times T^2)$: if $v = a\partial_s + b\bar{R} + c\partial_x + d\bar{J}(\partial_x)$,

$$dsdx(v, \bar{J}v) = -ad + bc$$

and

$$dt_1dt_2(v, \bar{J}v) = (bc - ad)(dt_1dt_2(\bar{R}, \bar{J}(\partial_x)) = bc - ad.$$

Hence,

$$\int_C dsdx = \int_C dt_1dt_2 = \int_{(\pi_{T^2})_*C} dt_1dt_2 \equiv \sum_{i=1}^{m}\theta_i^+ - \sum_{j=1}^{n}\theta_j^-$$

by the definition of Θ. On the other hand, let $\tilde{\varepsilon} > 0$ be small and define

$$I_{\tilde{\varepsilon}} := I \setminus \bigcup_{\rho\in\bar{\alpha}\cup\bar{\beta}} (x(\rho) - \tilde{\varepsilon}, x(\rho) + \tilde{\varepsilon})$$

and $C_{\tilde{\varepsilon}} := C \cap (\mathbb{R}\times I_{\tilde{\varepsilon}}\times T^2)$. Since $\int_C dsdx < \infty$ and $\partial C_{\tilde{\varepsilon}}$ does not have any ∂_x component,

$$\int_C dsdx = \lim_{\tilde{\varepsilon}\to 0}\int_{C_{\tilde{\varepsilon}}} dsdx = \lim_{\tilde{\varepsilon}\to 0}\int_{\partial C_{\tilde{\varepsilon}}} sdx = 0.$$

□

We now return to the proof of the necessity part of Theorem 1.4.1. Consider $(I\times T^2, \bar{\lambda})$ and the admissible almost complex structure $\bar{J}$ on $\mathbb{R}\times I\times T^2$ by (2.1.11). For each S^1-family of embedded orbits $\bar{\rho}$, let $H_{\bar{\rho}}$ be a Morse function as in Definition 2.1.2. For each S^1-family of m-fold covered orbits $\bar{\rho}^m$, $m > 1$, let $H_{\bar{\rho}^m} = H_{\bar{\rho}}$ be a Morse function on $\bar{\rho}^m$ with the identification (m-fold cover of ρ) $\leftrightarrow \rho$.

Proposition 2.2.21. *Let* (λ_n, J_n) *be a sequence of generic perturbations of* $(\bar{\lambda}, \bar{J})$ *converging to* $(\bar{\lambda}, \bar{J})$ *such that each* λ_n *is a good perturbation of* $\bar{\lambda}$ *and* J_n *is an admissible almost complex structure for* λ_n*. Let* α *and* β *be admissible orbit sets of* λ_1 *(and hence any* λ_n*) with* $I(\alpha,\beta) = 1$*. If* $\mathcal{R}_{\alpha,\beta}$ *is indecomposable and positive with respect to* $\bar{\lambda}$ *but an induced decoration of* $\mathcal{R}_{\alpha,\beta}$ *can be obtained from the minimal decoration by* $\mathbb{R}$*-relaxing a non-extreme edge, then* $\mathcal{M}^{J_n}(\alpha,\beta) = \emptyset$*, for* n *sufficiently large.*

Proof. Suppose there exist J_n-holomorphic curves C_n from α to β for all n. By Proposition 2.2.10, all the C_n's have genus zero and we can pass to a subsequence so that all the C_n's have the same partitions at the ends. The compactness argument as in [2] shows that (C_n) converges to a $\bar{J}$-holomorphic building $\bar{C}$ with $H_{\bar{\rho}^m}$. The glued surface of $\bar{C}$ also has genus zero and there is a bijection between:

(i) a positive end of C_n at a cover of an orbit of type $\check{e}$ or $\hat{h}$, and a flow of $H_{\bar{\rho}^m}$ from $\max H_{\bar{\rho}^m}$ to a positive end of $\bar{C}^1$,

(ii) a positive end of C_n at a cover of an orbit of type $\check{h}$ or $\hat{e}$, and a positive end of $\bar{C}^1$ at $\min H_{\bar{\rho}^m}$,

(iii) a negative end of C_n at a cover of an orbit of type $\hat{e}$ or $\check{h}$, and a flow of $H_{\bar{\rho}^m}$ from a negative end of $\bar{C}^k$ to $\min H_{\bar{\rho}^m}$,

(iv) a negative end of C_n at a cover of an orbit of type $\hat{h}$ or $\check{e}$, and a negative end of $\bar{C}^k$ at $\max H_{\bar{\rho}^m}$.

Consider the partition $(\mathcal{R}^i_j)$ associated to $\bar{C}$. Since an IP region associated to a trivial current only contributes local bigons to the partition, we ignore all the levels of $\bar{C}$ that are trivial as currents, i.e. multiply covered trivial cylinders, and rename the remaining levels as $\bar{C}^i$ for $1 \leq i \leq k$ for some $k \geq 1$. By Lemma 2.2.19, $(\mathcal{R}^i_j)$ contains at most two nonlocal IP regions, so $k \leq 2$.

Suppose $k = 1$. Since every edge of $\mathcal{R}_{\alpha,\beta}$ is labeled S^1-tight, there can be no Morse flows at any end of $\bar{C}^1$. Hence, there is a contribution of $+\delta$ to $\Theta(\bar{C}^1)$ for each $\mathbb{R}$-loose edge (a positive edge labeled $\check{h}$ or a negative edge labeled $\hat{h}$) and a contribution of zero for each $\mathbb{R}$-tight edge (a positive edge labeled $\hat{e}$ or a negative edge labeled $\check{e}$). Thus, $\Theta(\bar{C}^1) = +3\delta \neq 0$. This is a contradiction and we conclude that $k = 2$.

By the equality condition of Lemma 2.2.19, there are exactly two nonlocal IP regions, say $\mathcal{R}^1_{j_1}$ and $\mathcal{R}^2_{j_2}$, in $(\mathcal{R}^i_j)$, each with two extreme edges and sharing exactly one edge v_0 between them. Let $\bar{C}^i_{j_i}$ be the corresponding nontrivial holomorphic curves and $\bar{\rho}$ be the S^1-family of orbits of $\bar{\lambda}$ at $x = x(v_0)$. Then, $\bar{C}^1_{j_1}$ has a negative end at $\bar{\rho}(\theta^-)$ and $\bar{C}^2_{j_2}$ has a positive end at $\bar{\rho}(\theta^+)$ for some $\theta^\pm \in \mathbb{R}/\mathbb{Z}$.

First, assume $\bar{\rho}$ is convex, so that v_0 is non-extreme for $\mathcal{R}^1_{j_1}$ and extreme for $\mathcal{R}^2_{j_2}$. Hence, after matching each summand of $\Theta(\bar{C}^1_{j_1})$ in (2.2.7) with the edges of $\mathcal{R}^1_{j_1}$, v_0 contributes $-\theta^-$ to $\Theta(\bar{C}^1_{j_1})$, each of the two extreme edges of $\mathcal{R}^1_{j_1}$ contribute $+\delta$, and all the other edges contribute zero. Hence, $\theta^- = 2\delta$. Similarly for $\Theta(\bar{C}^2_{j_2})$, v_0 contributes θ^+ to $\Theta(\bar{C}^2_{j_2})$, its other extreme edge contributes $+\delta$ and all other edges contribute zero. Hence, $\theta^+ = -\delta$. But $H_{\bar{\rho}}$ has the maximum at $\theta = 0$ and minimum at $\theta = \delta$, so there cannot be a Morse flow from 2δ to $-\delta$. Hence, the $\bar{J}$-holomorphic building $\bar{C}$ as described does not exist and we conclude that for large enough n, C_n does not exist. If $\bar{\rho}$ is concave, the argument is similar: there cannot be a Morse flow from $\theta^- = \delta$ to $\theta^+ = -2\delta$ since $H_{\bar{\rho}}$ has the maximum at $\theta = -\delta$ and minimum at $\theta = 0$. □

Lastly, we deal with local bigons.

Lemma 2.2.22. *Consider $(I \times T^2, \bar{\lambda})$ and the admissible almost complex structure $\bar{J}$ on $\mathbb{R} \times I \times T^2$ by (2.1.11). Let λ be a good perturbation of $\bar{\lambda}$ and J be a generic admissible almost complex structure which is a small perturbation of $\bar{J}$. If α and β are admissible orbit sets with $I(\alpha, \beta) = 1$ and $\mathcal{R}_{\alpha,\beta}$ is a local bigon, then $\langle \partial\alpha, \beta \rangle = 0$.*

Proof. First, consider the sequence (λ_n, J_n) of generic perturbations of $(\bar{\lambda}, \bar{J})$ converging to $(\bar{\lambda}, \bar{J})$ and suppose each λ_n is a good perturbation. Suppose the edges of $\mathcal{R}_{\alpha,\beta}$ occur at $x = x(\bar{\rho})$ for an S^1-family of orbits $\bar{\rho}$ and by symmetry, assume $\bar{\rho}$ is convex. By a Morse-Bott argument as in [2], for large enough n, there are two J_n-holomorphic cylinders (modulo $\mathbb{R}$) from $\check{e}$ to $\check{h}$, corresponding to the two Morse flows of $H_{\bar{\rho}}$ from $\max H_{\bar{\rho}}$ to $\min H_{\bar{\rho}}$.

Now for the given (λ, J), consider deforming it to $(\bar{\lambda}, \bar{J})$ via (λ_r, J_r) for $r \in [0, 1]$. Consider the moduli space of J_r-holomorphic curves $\mathcal{M}^{J_r}(\alpha, \beta)$. It is possible that at discrete values of r, the moduli space contains a broken holomorphic curve. However, by the equality condition of Lemma 2.2.1, all components must stay within $\mathbb{R} \times (x(\bar{\rho}) - \epsilon, x(\bar{\rho}) + \epsilon) \times T^2$. Since each component C' of the broken holomorphic curve has $I(C') = I(\alpha', \beta') \geq 0$, any component C' with $I(C') = 0$ must have the same positive and negative end, i.e. it is trivial cylinder. Hence, this moduli count stays the same and $\#\mathcal{M}(\alpha, \beta)/\mathbb{R} = 0$. □

Combining the results of Proposition 2.2.10, Proposition 2.2.21 and Lemma 2.2.22 proves the necessity part of the theorem.

2.3 Proof of sufficiency

In this section, we show that if admissible orbit sets α and β satisfy the conditions of Theorem 1.4.1, then the mod 2 count of $\mathcal{M}(\alpha, \beta)/\mathbb{R}$ is indeed nonzero. By Corollary 2.2.8, we may assume that $\mathcal{R}_{\alpha,\beta}$ is indecomposable and positive. We use induction on the number of edges of $\mathcal{R}_{\alpha,\beta}$ where each step involves partitioning an IP region into two smaller IP regions. Before we proceed with the induction, we need to establish the invariance of the moduli count under certain deformations of (λ, J).

Invariance of the moduli count

For this section, we fix $I = [X_w, X_e]$, a decorated positive IP region $\mathcal{R}$ and $L > 0$ and consider various contact forms $\bar{\lambda}$ and λ on $I \times T^2$.

Definition 2.3.1. Let $\bar{\lambda} = -g dt_1 + f dt_2$ be a contact form on $I \times T^2$ and $\mathcal{R}$ is a fixed decorated IP region on I. We say that $\bar{\lambda}$ *supports* $\mathcal{R}$ *(via* ϕ*)* if there is a reparametrization ϕ of I such that:

(i) $\mathcal{R}$ is positive with respect to $\phi^*\bar{\lambda}$ and

(ii) $\mathcal{R}$ is associated to a pair of orbit sets α and β of a good perturbation of $\phi^*\bar{\lambda}$ with an induced decoration.

Lemma 2.3.2. *Let $\bar{\lambda}_0$ and $\bar{\lambda}_1$ be two contact forms on $I \times T^2$ which support $\mathcal{R}$ via ϕ_0 and ϕ_1, respectively. Then, there is a path $r \mapsto \bar{\lambda}_r$, $r \in [0, 1]$ from $\bar{\lambda}_0$ to $\bar{\lambda}_1$ of contact forms supporting $\mathcal{R}$.*

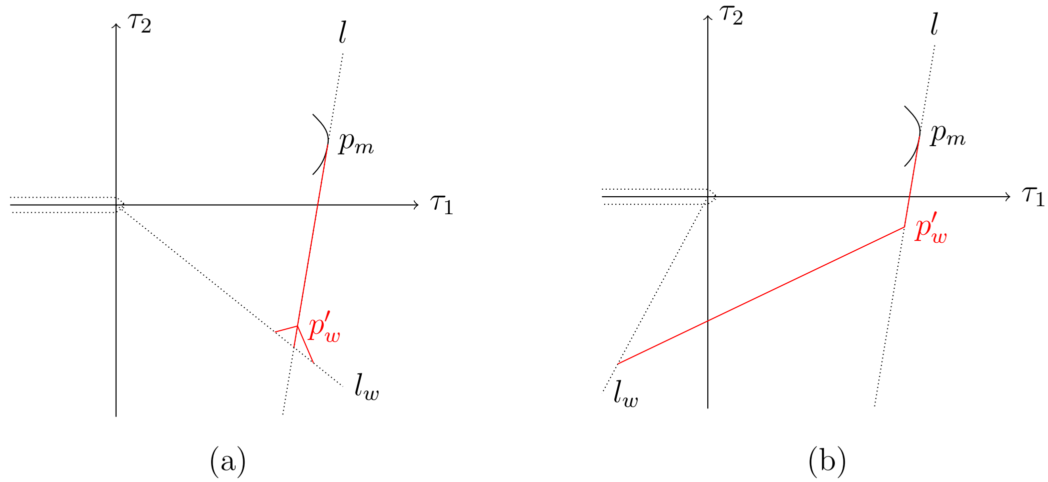

Figure 2.6: Part of the graph of (F, G).

Before proving this lemma, we describe an auxiliary function $\psi : [x_w, x_e] \to \mathbb{R}^+$ associated to a contact form $\bar{\lambda} = -g dt_1 + f dt_2$ on $[x_w, x_e] \times T^2$ whose Reeb vector field $\bar{R}$ satisfies that: there is a vector $0 \neq \sigma \in \mathbb{Z}^2$ such that $\bar{R}(x) \times \sigma \geq 0$ for all $x \in [x_w, x_e]$. For simplicity, we assume $\sigma = (-1, 0)$. Hence, $\bar{\lambda}$ being a contact form with $\bar{R}(x) \times \sigma \geq 0$ translates to $g(x)$ being an increasing function and the graph of (f, g) rotating clockwise in $\mathbb{R}^2 \setminus \{\text{negative } \tau_1\text{-axis}\}$. Let $x_m \in (x_w, x_e)$ be any point where $\bar{R}(x_m)$ is a multiple of $\bar{R}(x_w) + \bar{R}(x_e) \in \mathbb{R}^2$. Consider the tangent line l to (f, g) at $p_m := (f(x_m), g(x_m))$. We claim that there is a continuous path $(F, G) : [0, T] \to \mathbb{R}^2$ for some $T > 0$, consisting of three linear paths connecting the four distinct points $p_w, p'_w, p'_e, p_e \in \mathbb{R}^2$ defined as follows:

(i) Start at $p_w := (c_w f(x_w), c_w g(x_w))$ for some $c_w \in \mathbb{R}^+$ and travel in $\bar{R}(x_w)$-direction to a point $p'_w \in l$.

(ii) Then, travel in $\bar{R}(x_m)$ direction to $p'_e \in l$ so that p_m lies between p'_w and p'_e.

(iii) Then, travel in $\bar{R}(x_e)$ direction to $p_e := (c_e f(x_e), c_e g(x_e))$ for some $c_e \in \mathbb{R}^+$.

To prove the claim, suppose l intersects $l_w := \{(cf(x_w), cg(x_w) | c \in \mathbb{R}^+\}$ at $(c_0 f(x_w), c_0 g(x_w))$ for some $c_0 \in \mathbb{R}^+$. Then, we set c_w to be slightly smaller, equal to, or slightly larger than c_0 depending on whether $\bar{R}(x_w) \times \bar{R}(x_m)$ is positive, zero, or negative. See Figure 2.6 (a), where the three short red segments illustrate these three possibilities. If l does not intersect l_w, any sufficiently large $c_w \in \mathbb{R}^+$ will work: (1.2.1) and $\bar{R}(x_w) \times \sigma \geq 0$ ensures that the path starting at c_w in the direction $\bar{R}(x_w)$ intersects l. See Figure 2.6 (b). We have described the path (F, G) from p_w to p_m and we similarly construct the path (F, G) from p_m to p_e. Obtain a smooth path $(\tilde{F}, \tilde{G}) : [0, T] \to \mathbb{R}^2$ from (F, G) by smoothing the corners of (F, G) at p'_w and p'_e and define $\tilde{\psi} : [x_w, x_e] \to \mathbb{R}^+$ so that the image of $(\tilde{\psi} f, \tilde{\psi} g)$ agrees with the image of $(\tilde{F}, \tilde{G})$. Finally, let ψ be a smooth C^1-small perturbation of $\tilde{\psi}$ so that:

(i) $\psi' \equiv 0$ in a small neighborhood of x_w and x_e and

(ii) (ψg) is a strictly increasing function on $[x_w, x_e]$.

These conditions will be used in part (c) of the proof of Lemma 2.3.2 to ensure that the new contact form $\tilde{\lambda}$ obtained using these auxiliary functions satisfies that: (i) there are orbit sets whose associated decorated IP paths are $\partial^+\mathcal{R}$ and $\partial^-\mathcal{R}$, and (ii) $\mathcal{R}$ is positive with respect to $\tilde{\lambda}$. Note that $-\psi g dt_1 + \psi f dt_2$ is "minimally fluctuating" in the sense that it does not have any extra Reeb orbits except those that are absolutely necessary to be able to support $\mathcal{R}$. We proceed with the proof now.

Proof. (of Lemma 2.3.2) We consider the following four type of paths $r \mapsto \bar{\lambda}_r$ of contact forms supporting $\mathcal{R}$. We will show that any two contact forms supporting $\mathcal{R}$ can be constructed as a composition of these.

(a) If ϕ is a re-parametrization of I, then we can deform $\bar{\lambda}$ to $\phi^*\bar{\lambda}$ via

$$\bar{\lambda}_r := [r\phi + (1-r)\,\mathrm{id}]^*\bar{\lambda}.$$

We note that each $(I \times T^2, \bar{\lambda}_r)$ is contactomorphic to another, but the distinguished $\bar{J}_r$ defined by (2.1.11) depends on the parametrization of I, so it is nontrivial that the moduli count of $\bar{J}_r$-holomorphic curves is the same.

(b) If v_w and v_e are west and east extreme edges of $\mathcal{R}$, then using a path of diffeomorphisms $\phi_r : [X_w, X_e] \to [x_w(r), x_e(r)]$ where

$$x_w(r) = (1-r)X_w + r(x(v_w) - \epsilon), \qquad x_e(r) = (1-r)X_e + r(x(v_e) + \epsilon)$$

and where each $\phi_r|_{[x(v_w),x(v_e)]} = \mathrm{id}$, we can deform $\bar{\lambda}$ to $\phi_1^*\bar{\lambda}$, which is "uninteresting" outside of $[x(v_w), x(v_e)]$.

(c) Let $\bar{\lambda} = -g dt_1 + f dt_2$ be a contact form supporting $\mathcal{R}$ via id and let $\bar{R}(x)$ be the Reeb vector field of $\bar{\lambda}$. We deform $\bar{\lambda}$ to a "minimally fluctuating" contact form supporting $\mathcal{R}$ using auxiliary functions ψ. Write

$$I \setminus \{x(v) | v \in \partial^{\pm}\mathcal{R}\} = [X_w, x_1) \cup (x_1, x_2) \cup \cdots \cup (x_{k-1}, x_k) \cup (x_k, X_e].$$

For each $1 \le i \le k-1$, obtain an auxiliary function $\psi_i : [x_i, x_{i+1}] \to \mathbb{R}^+$ for $\bar{\lambda}|_{[x_i,x_{i+1}]}$ with $\sigma = \sigma(x_i + \epsilon)$ as described above. Also, let $\psi_0 : [X_w, x_1] \to \mathbb{R}^+$ and $\psi_k : [x_k, X_e] \to \mathbb{R}^+$ be constant. Re-scale ψ_i for each $1 \le i \le k$ so that $\psi_i(x_i) = \psi_{i-1}(x_i)$. Patching these ψ_i gives a smooth function $\psi : [X_w, X_e] \to \mathbb{R}^+$. The properties of each auxiliary function ψ_i implies that $-\psi g dt_1 + \psi f dt_2$ defines a contact form supporting $\mathcal{R}$. For $r \in [0,1]$, let

$$\bar{\lambda}_r := [(1-r) + r\psi](-g dt_1 + f dt_2).$$

Since $\bar{\lambda}_0$ and $\bar{\lambda}_1$ are both contact forms supporting $\mathcal{R}$, so is $\bar{\lambda}_r$ for every $r \in [0,1]$.

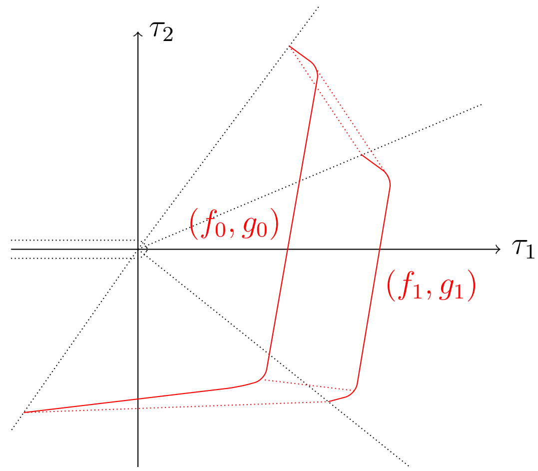

Figure 2.7: Interpolating two "minimally fluctuating" contact forms.

(d) Let $\bar{\lambda}_0 = -g_0 dt_1 + f_0 dt_1$ and $\bar{\lambda}_1 = -g_1 dt_1 + f_1 dt_1$ be two "minimally fluctuating" contact forms supporting $\mathcal{R}$ as constructed in (c). Suppose that, for all $x \in I$, the angle between $\bar{R}_0(x)$ and $\bar{R}_1(x)$ is small, i.e. $\left|\bar{R}_0(x) \times \bar{R}_1(x)\right| / \left|\bar{R}_0(x)\right| \left|\bar{R}_1(x)\right| < \varepsilon_0$ for some small $\varepsilon_0 > 0$. Then, for sufficiently small ε_0,

$$\bar{\lambda}_r := (1-r)\bar{\lambda}_0 + r\bar{\lambda}_1$$

is a contact from supporting $\mathcal{R}$. Figure 2.7 illustrates this interpolation on $[x_i, x_{i+1}]$ with $\sigma(x_i + \epsilon) = (-1, 0)$.

Given any contact form $\bar{\lambda}$ supporting $\mathcal{R}$, we may assume that it supports $\mathcal{R}$ via $\phi = \text{id}$ using part (a). We assume that $\bar{\lambda}$ does not have any Reeb orbits (of action less than L) outside of the interval $[x(v_w), x(v_e)]$ where v_w and v_e are east and west extreme edges of $\mathcal{R}$, using part (b). We can also assume that $\bar{\lambda}$ satisfies the conditions of (d): use part (c) with each auxiliary ψ_i sufficiently close to id and re-parametrize I if necessary, using (a) again. Hence, we can connect any two $\bar{\lambda}_0$ and $\bar{\lambda}_1$ using (d) after these simplifying assumptions. □

We now define $\bar{J}_r$ for each $\bar{\lambda}_r$ by (2.1.11), and choose a path of good perturbations (λ_r, J_r) of $(\bar{\lambda}_r, \bar{J}_r)$. Let $r_0 \in [0,1]$ be such that λ_{r_0} supports $\mathcal{R}$ via $\phi = \text{id}$. If ρ is an embedded orbit of λ_{r_0}, then $(\phi_r)_*\rho$ is an embedded orbit of λ_r for all $r \in [0,1]$. Let α and β be a pair of orbit sets of λ_{r_0} whose associated IP region with an induced decoration is $\mathcal{R}$. We define the moduli space

$$\mathcal{M}^r := \mathcal{M}^{J_r}((\phi_r)_*\alpha, (\phi_r)_*\beta)$$

of J_r-holomorphic currents from $(\phi_r)_*\alpha$ to $(\phi_r)_*\beta$. The following is an adaptation of [11, Lemma 3.15]:

Lemma 2.3.3. *Consider $\mathcal{R}, \alpha, \beta, \lambda_r$ and $\mathcal{M}^r$ as above. Suppose $\mathcal{R}$ is nonlocal, indecomposable and minimal and has exactly two $\mathbb{R}$-loose edges. Then*

$$\#(\mathcal{M}^0/\mathbb{R}) = \#(\mathcal{M}^1/\mathbb{R}).$$

Remark 2.3.4. A key assumption in Lemma 2.3.3 is that $\bar{\lambda}_r$ never violates positivity with respect to $\mathcal{R}$. If $\bar{\lambda}_{r+\varepsilon}$ violates positivity, it is possible to have $\#(\mathcal{M}^{r-\varepsilon}/\mathbb{R}) \neq 0$ and $\mathcal{M}^{r+\varepsilon} = \emptyset$. See equation (2.3.2) for where this assumption is used.

Proof. Away from a discrete set $\{r_i\} \subset (0,1)$, the moduli space

$$\widetilde{\mathcal{M}} = \bigcup_{r\in[0,1]} \mathcal{M}^r$$

forms a two-dimensional manifold with an $\mathbb{R}$-action. There are two types of (possible) bifurcation points: the first type is where $\bar{\lambda}_{r_i}$ has a Reeb orbit of action less than L and whose linearized return map is id. This happens if $(f', g') \in \mathbb{Q} \cup \{\infty\}$ at a point of inflection, which was avoided for a generic (f, g) but which cannot be avoided for a generic path (f_r, g_r). The other type occurs where $J_{r'_j}$ is not generic for $(\{r'_j\} \times Y, \lambda_{r'_j})$ so that a $J_{r'_j}$-holomorphic curve C with $I(C) = 0$ can exist. We can arrange that at each bifurcation point r_i of the first type, the almost complex structure J_{r_i} is generic and there is exactly one S^1-family $\bar{\rho}_{r_i}$ of Reeb orbits of $\bar{\lambda}_{r_i}$ whose linearized return map is id.

Case 1. Let $r_0 \in (0,1)$ be a bifurcation point of the first type and suppose there is a broken J_{r_0}-holomorphic curve C from α to β. Similarly to $\bar{J}$-holomorphic buildings, C partitions $\mathcal{R}$ into $(\mathcal{R}^i_j)$, but here, Lemma 2.2.19 needs to be modified since an intermediate edge v may occur at $x = x(\bar{\rho}_{r_0})$ in which case v may be extreme for both IP regions of $(\mathcal{R}^i_j)$ sharing v. We claim that $(\mathcal{R}^i_j)$ cannot contain such an edge so that the conclusion of Lemma 2.2.19 still holds.

Suppose that there are two distinct nonlocal IP regions $\mathcal{R}_1$ and $\mathcal{R}_2$ of $(\mathcal{R}^i_j)$ which share an edge v and that v is extreme for both $\mathcal{R}_1$ and $\mathcal{R}_2$. Let $x_0 := x(\bar{\rho}_{r_0})$ and without loss of generality, suppose v is a positive edge for $\mathcal{R}_1$ and a negative edge for $\mathcal{R}_2$. Since both $\mathcal{R}_1$ and $\mathcal{R}_2$ are positive IP regions, v must be east extreme for one and west extreme for the other. By symmetry, assume v is east extreme for $\mathcal{R}_1$ and west extreme for $\mathcal{R}_2$ and let $\bar{R}_r$ denote the Reeb vector field of $\bar{\lambda}_r$. Positivity of $\mathcal{R}_1$ and $\mathcal{R}_2$ with respect to $\bar{\lambda}_{r_0}$ implies that

$$\begin{aligned} \bar{R}_{r_0}(x) \times v > 0 \quad & \text{for } x_0 - \epsilon < x < x_0, \\ \bar{R}_{r_0}(x) \times v = 0 \quad & \text{for } x = x_0, \\ \bar{R}_{r_0}(x) \times v > 0 \quad & \text{for } x_0 < x < x_0 + \epsilon. \end{aligned} \tag{2.3.1}$$

Also, since v is the slice class of $\mathcal{R}$ at $x = x_0$ for all r near r_0, positivity of $\mathcal{R}$ with respect to $\bar{\lambda}_r$ implies that, for all $r \neq r_0$ near r_0,

$$\bar{R}_r(x) \times v > 0 \quad \text{for } x_0 - \epsilon < x < x_0 + \epsilon. \tag{2.3.2}$$

For a generic path $r \mapsto (\bar{\lambda}_r, \bar{J}_r)$, equations (2.3.1) and (2.3.2) cannot all be satisfied and this proves the claim.

Hence, by Lemma 2.2.19, C contains exactly one nonlocal component C' which must be somewhere injective since each of its extreme ends has multiplicity one. We have that $I(C') \geq 1$ by genericity of J_{r_0} and any other non-trivial local components C'' must also have $I(C'') \geq 1$. Therefore, C does not have any other non-trivial components and we do not have a bifurcation at $r = r_0$.

Case 2. Let $r_0' \in (0, 1)$ be a bifurcation point of the second type and let $C \in \mathcal{M}^{r_0}$ be a broken J_{r_0}-holomorphic curve. This time, a standard application of Lemma 2.2.19 implies that C contains one nonlocal (somewhere injective) component C'. By genericity of the path $r \mapsto (\bar{\lambda}_r, \bar{J}_r)$, $I(C') \geq 0$. Suppose C contains a nontrivial local component C'' with $I(C'') = 1$. Since $I(C') = 0$, the IP region $\mathcal{R}_{C'}$ associated to the positive and negative orbit sets of C' are admissible and so are α and β. Hence, by Proposition 2.2.7, C'' must be a cylinder with positive and negative ends of multiplicity one. As in Lemma 2.2.22, C'' corresponds to an auxiliary Morse flow of $H_{\bar{\rho}}$ for some $\bar{\rho}$. These flows occur in pairs and by the standard gluing arguments as in [17], we have that the mod 2 count of $\mathcal{M}^r/\mathbb{R}$ does not change during this bifurcation. □

Base cases for induction

Here, we list some IP regions $\mathcal{R}_{\alpha,\beta}$ for the cases where we already know the differential coefficient:

Proposition 2.3.5. *Let α and β be admissible orbit sets and suppose $\mathcal{R}_{\alpha,\beta}$ is indecomposable and minimal. Further suppose that $\mathcal{R}_{\alpha,\beta}$ has exactly one S^1-loose edge and that it is of one of the following types:*

(i) A nonlocal bigon with one positive and one negative edge.

(ii) A nonlocal bigon with two positive edges.

(iii) A triangle formed by the two positive extreme edges with all the edges labeled convex. If there are multiple non-extreme edges, they are all elliptic.

(iv) A triangle formed by the two negative extreme edges with all the edges labeled concave. If there are multiple non-extreme edges, they are all elliptic.

Then, $\mathcal{M}(\alpha, \beta)/\mathbb{R} \cong \{pt\}$ and, in particular, $\langle \partial\alpha, \beta \rangle \neq 0$.

Proof. Each of the above cases except for case (i) is a special case of the main results in [11, 12]. We include the proof from there for completeness.

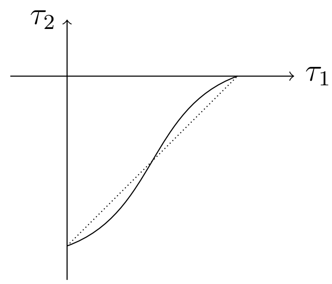

Figure 2.8: A graph of (f, g) for $\bar{\lambda}$ for case (i).

(i) Using Lemma 2.3.3, assume $\bar{\lambda} = -g dt_1 + f dt_2$ where

$$f(x) = x(1 + H(x))$$
$$g(x) = (x - 1)(1 + H(x))$$

for $x \in (0, 1)$ and

$$H(x) = \pm\eta x(x - 1/2)(x - 1)$$

with a plus sign if $x(\alpha_1) < x(\beta_1)$ and minus otherwise. See Figure 2.8 for the graph of (f, g) for $H(x)$ with the plus sign and compare this with $(0, 1) \times T^2 \subset S^3$ with λ_{std} as in Example 1.2.1. Let $\tilde{H}(x, \theta)$ be a perturbation of the Morse function $H(x, \theta) := H(x)$ on $(0, 1) \times \mathbb{R}/\mathbb{Z}$ such that:

- $\tilde{H}(x, \theta)$ has four critical points at $(x_M, \delta), (x_M, 0), (x_m, 0)$ and $(x_m, -\delta)$ where x_M and x_m are the local maximum and minimum of $H(x)$ and
- $\tilde{H}(x_M, \cdot)$ and $\tilde{H}(x_m, \cdot)$ satisfy the same conditions of Definition 2.1.2 as $H_{\bar{\rho}}$ for a convex $\bar{\rho}$ and $H_{\bar{\rho}}$ for a concave $\bar{\rho}$, respectively.

Then, using Lemma 2.3.3 and by changing coordinates (t_1, t_2) if necessary, we may regard λ as a perturbation of λ_{std} by the auxiliary Morse function $\tilde{H}$. By a standard Morse-Bott argument as in [2], the unique $(\mathbb{R} \times I \times S^1)$-family of J_{std}-holomorphic cylinders gives a unique J-holomorphic curve from the orbit at (x_M, δ) to the orbit at $(x_m, 0)$ as well as one from the orbit at $(x_M, 0)$ to the orbit at $(x_m, -\delta)$.

(ii) We compare $\mathcal{M}(\alpha, \beta)$ with the moduli space of holomorphic cylinders in $(\mathbb{R} \times S^2 \times S^1, \bar{J}_0)$ considered by Taubes in [23, Theorem A.1(c)], where $\bar{J}_0$ is an $\mathbb{R} \times S^1 \times S^1$-invariant almost complex structure. More precisely, we identify

$$\mathbb{R} \times [X_w, X_e] \times T^2 = \mathbb{R} \times (x_1, x_2) \times S^1 \times S^1 \subset \mathbb{R} \times S^2 \times S^1$$

and deform our J to a perturbation of $\bar{J}_0$ using Lemma 2.3.3. Then, the unique member of $\mathcal{M}(\alpha, \beta)/\mathbb{R}$ is obtained from the unique $(\mathbb{R} \times S^1)$-family of $\bar{J}_0$-holomorphic curves by the usual Morse-Bott argument [2].

(iii) Let $C \in \mathcal{M}(\alpha,\beta)$. The partition condition (2.2.5) on elliptic convex negative ends implies that C has only one negative puncture so C has three punctures regardless of the number of non-extreme edges of $\mathcal{R}_{\alpha,\beta}$. Therefore, we can compare $\mathcal{M}(\alpha,\beta)$ with the moduli space of three-punctured spheres in $(\mathbb{R}\times S^2\times S^1, \bar{J}_0)$ in [23, Theorem A.2]. By a Morse-Bott argument, the unique member of $\mathcal{M}(\alpha,\beta)/\mathbb{R}$ comes from the unique $(\mathbb{R}\times S^1\times S^1)$-family of $\bar{J}_0$-holomorphic spheres. The auxiliary Morse flow occurs at $\bar{\rho}$ where, either $\beta=\beta_1$ is the hyperbolic orbit from $\bar{\rho}$, or one of α_i's is the elliptic orbit from $\bar{\rho}$.

(iv) This case is similar to (c) but with the identification

$$(-\mathbb{R})\times[X_e,X_w]\times T^2=\mathbb{R}\times[x_1,x_2]\times S^1\times S^1\subset\mathbb{R}\times S^2\times S^1.$$

□

Induction step

In this section, we complete the proof of Theorem 1.4.1 using induction.

Proof. (of sufficiency part of Theorem 1.4.1) Let α and β be as in the hypothesis of Theorem 1.4.1 and assume $\mathcal{R}_{\alpha,\beta}$ is indecomposable using Corollary 2.2.8. If $\mathcal{R}_{\alpha,\beta}$ is a bigon, the theorem holds by Proposition 2.3.5. Let $n>2$ and suppose we have shown the theorem holds whenever $\mathcal{R}_{\alpha,\beta}$ has less than n edges. Let w_1 and w_2 be edges of $\mathcal{R}_{\alpha,\beta}$ so that $x(w_1)$ and $x(w_2)$ are the two smallest entries from $(x_i^+)\cup(x_j^-)$. Since $\mathcal{R}_{\alpha,\beta}$ has one west extreme edge and one east extreme edge, w_2 is not extreme. In particular, $x(w_1)<x(w_2)$. We also have that $w_1\neq\pm w_2\in\mathbb{Z}^2$: otherwise, $\mathcal{R}_{\alpha,\beta}$ is forced to be a bigon by Corollary 2.2.3. By symmetry, we assume that w_1 is labeled convex and consider the two cases depending on the convexity of w_2.

Case 1. (w_2 is labeled convex.) Since w_1 is $\mathbb{R}$-loose by the hypothesis, $w_1\in\mathcal{P}_\alpha$. Since w_2 is not extreme, it must be $\mathbb{R}$-tight by condition (d) of the hypothesis, and hence $w_2\in\mathcal{P}_\beta$. After a change of basis, we may assume $w_1=(1,1)$ and $w_2=(0,1)$. For convenience, let us use the slope of the underlying vector as the subscript, i.e. write v_1, x_1, v_∞ and x_∞ for $w_1, x(w_1), w_2$ and $x(w_2)$, respectively. See Figure 2.9 (a) for the graph of (f,g) for $\bar{\lambda}=-g dt_1+f dt_2$.

Using Lemma 2.3.3, we assume that there are no points of inflection between x_1 and x_∞. We may also assume that $\lambda|_{[X_w,x_1]}$ is convex. Let $X_w<x_{-1}<x_0<x_1$ be x-coordinates so that $\bar{R}(x_{-1})\in\mathbb{R}^2$ is a positive multiple of $(1,-1)$ and $\bar{R}(x_0)$ is a positive multiple of $(1,0)$ (See Figure 2.9 (a).) With an abuse of notation, let v_{-1} denote the IP path with one edge at $x=x_{-1}$, and let it also denote the edge itself. Similarly, let v_0 denote the IP path with one edge at $x=x_0$ as well as the edge itself.

Write $\mathcal{P}_\alpha=v_1\mathcal{P}^+$ and $\mathcal{P}_\beta=v_\infty\mathcal{P}^-$ and let $\mathcal{P}^0:=v_{-1}v_1\mathcal{P}^+$ and $\mathcal{P}^2:=v_0\mathcal{P}^-$. Let $\tilde{\alpha}$ and $\tilde{\beta}$ be two orbit sets of λ such that $\mathcal{P}_{\tilde{\alpha}}=\mathcal{P}^0$ and $\mathcal{P}_{\tilde{\beta}}=\mathcal{P}^2$. We examine each of the nonzero

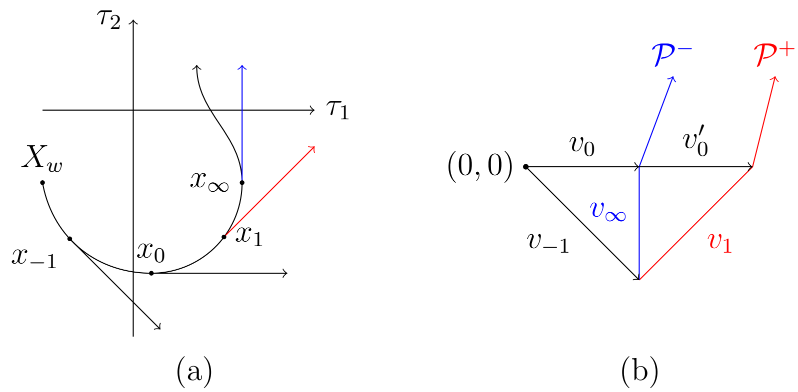

Figure 2.9: Case 1 of the induction step.

summands of

$$\langle \partial^2 \tilde{\alpha}, \tilde{\beta} \rangle = \sum_\gamma \langle \partial \tilde{\alpha}, \gamma \rangle \langle \partial \gamma, \tilde{\beta} \rangle = 0. \tag{2.3.3}$$

Consider an IP path $\mathcal{P}^1 = \mathcal{P}_\gamma$ for an orbit set γ such that there is a holomorphic curve from $\tilde{\alpha}$ to γ as well as one from γ to $\tilde{\beta}$. Since the IP region $\mathcal{R}$ between $\mathcal{P}^0$ and $\mathcal{P}^2$ is minimal with three $\mathbb{R}$-loose edges, by Lemma 2.2.19, $\mathcal{P}^1$ partitions $\mathcal{R}$ into two nonlocal IP regions $\mathcal{R}^1_{j_1}$ and $\mathcal{R}^2_{j_2}$, each with two extreme edges, and they share a single edge v. Since each $\mathcal{R}^i_{j_i}$ is nonlocal, $x(v) > x_{-1}$. We claim that, in fact, $x(v) \geq x_0$. Suppose otherwise, i.e. $x_{-1} < x(v) < x_0$. Then, v has slope between -1 and 0 and it must be the first edge of $\mathcal{P}^1$. Hence, as is clear from Figure 2.9, the realization of $\mathcal{P}^1$ starting at $(0,0) \in \mathbb{Z}^2$ necessarily intersects the realization of $\mathcal{P}^0$ starting at $(0,0)$. This contradicts positivity of $\mathcal{R}^1_{j_1}$. Hence, $x(v) \geq x_0$ and, in particular, both v_{-1} and v_0 must belong to the same IP region $\mathcal{R}' \in (\mathcal{R}^1_{j_1}, \mathcal{R}^2_{j_2})$.

Next, suppose that the east extreme edge w'_e of $\mathcal{R}'$ has $x(w'_e) > x_\infty$. Then, any edge w' of $\mathcal{R}'$ with $x_{-1} < x(w') \leq x_\infty$ is a negative edge and hence,

$$\sigma_{\mathcal{R}'}(x_\infty + \epsilon) = -v_{-1} + \sum_{x_{-1} < x(w') \leq x_\infty} w' = (p,q) \neq 0$$

with $p \geq 0$. This violates positivity of $\mathcal{R}'$ at $x = x_\infty + \epsilon$, and we conclude $x(w'_e) \leq x_\infty$. Recall also that $(\mathcal{R}^1_{j_1}, \mathcal{R}^2_{j_2})$ contains exactly four $\mathbb{R}$-loose edges: v_{-1}, v_1, the east extreme edge of $\mathcal{R}$ and one instance of v. Hence, w'_e must be either v_1 or the $\mathbb{R}$-loose instance of v. To summarize, if γ is from a nonzero summand in (2.3.3), one of the IP regions $\mathcal{R}'$ in the partition of $\mathcal{R}$ by $\mathcal{P}_\gamma$ must look like the following:

(i) (Case $w'_e = v_1$.) $\partial^- \mathcal{R}'$ cannot be just v_0, so $\partial^- \mathcal{R}' = v_0 v'_0$ with $v'_0 = (1,0) \in \mathbb{Z}^2$ and $x(v'_0) = x_0$.

(ii) (Case $w'_e = v$.) $\partial^- \mathcal{R}'$ must be just v_0, so $v = v_\infty$.

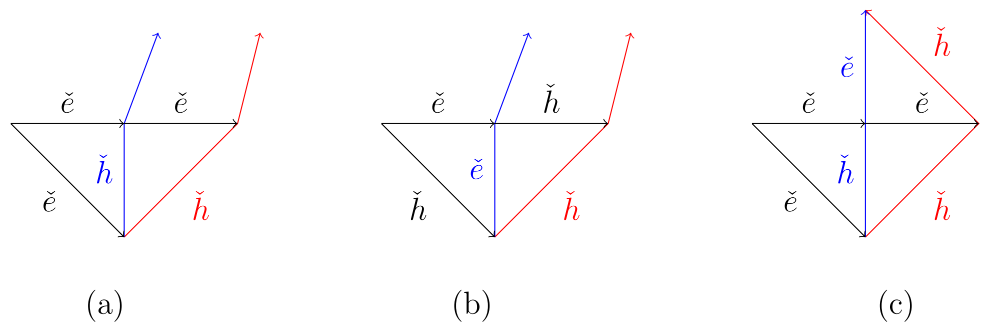

Figure 2.10: (a) and (b) are two possible decorations of $(\mathcal{R}^i_{j_i})$. (c) illustrates an extra induction step required for the decoration as in (b).

See Figure 2.9 where both v_1 and v_∞ are indicated. This completely describes all possible IP paths $\mathcal{P}_\gamma$ for nonzero summands of (2.3.3).

We now consider decorations of $\mathcal{R}$ and $(\mathcal{R}^i_{j_i})$. We can make $I(\mathcal{R}) = 2$ by extending the induced decorations of $\mathcal{P}_\alpha$ and $\mathcal{P}_\beta$ as follows:

- Label v_0 as $\check{e}$.
- Label v_{-1} as $\check{e}$ if any edge of $\mathcal{P}_\beta$ at $x = x_\infty$ is labeled $\check{h}$. Label it $\check{h}$ otherwise.

Imposing $I(\mathcal{R}') = 1$ determines the label of v in the above two cases as follows:

(i) (Case $w'_e = v_1$.) $v = v'_0$ is labeled $\check{h}$ if both v_1 and v_{-1} are labeled $\check{h}$; it is labeled $\check{e}$ otherwise.

(ii) (Case $w'_e = v$.) $v = v_\infty$ is labeled $\check{e}$ if v_{-1} is $\check{h}$; it is labeled $\check{h}$ otherwise.

See Figure 2.10 (a) and (b) for two possible decorations.

Let η and η' be orbit sets whose associated decorated IP paths are $\mathcal{P}_\eta = v_0 v'_0 \mathcal{P}^+$ and $\mathcal{P}_{\eta'} = v_{-1} v_\infty \mathcal{P}^-$ with prescribed decorations as above. We have shown that if $\langle \partial\tilde{\alpha}, \gamma\rangle\langle\partial\gamma, \tilde{\beta}\rangle \neq 0$ in (2.3.3), then $\gamma = \eta$ or $\gamma = \eta'$. If v'_0 is labeled $\check{e}$ as in Figure 2.10 (a), then by the induction hypothesis and Proposition 2.3.5,

$$\langle \partial\tilde{\alpha}, \eta\rangle \neq 0, \quad \langle \partial\eta, \tilde{\beta}\rangle \neq 0, \quad \langle \partial\eta', \tilde{\beta}\rangle \neq 0$$

so we have

$$\langle \partial\tilde{\alpha}, \eta'\rangle = \langle \partial\alpha, \beta\rangle \neq 0.$$

If v'_0 is labeled $\check{h}$ as in Figure 2.10 (b), then we apply another induction step with $\alpha_{new} = v_{-1}v_1$ and $\beta_{new} = v_0 v'_0$. After another change of basis, $\mathcal{R}_{\alpha_{new},\beta_{new}}$ looks like Figure 2.10 (c). Now $\langle \partial\tilde{\alpha}, \eta\rangle \neq 0$ follows from the above case.

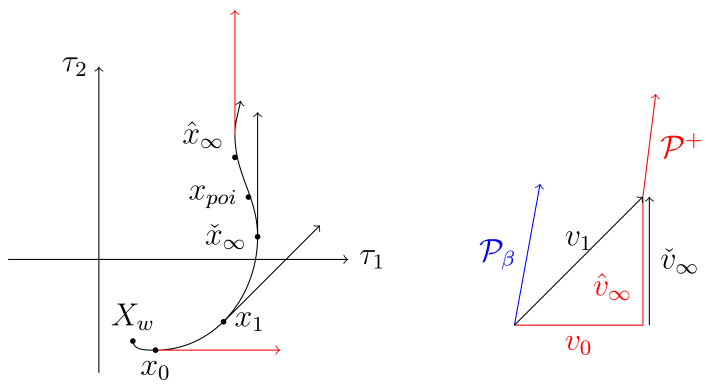

Figure 2.11: Case 2 of the induction step.

Case 2. (w_2 is labeled concave.) In this case, both w_1 and w_2 are from $\mathcal{P}_\alpha$. After a change of basis, we assume $w_1 = (1,0)$ and $w_2 = (0,1)$ and write $v_0, x_0, \hat{v}_\infty$ and $\hat{x}_\infty$ for $w_1, x(w_1), w_2$ and $x(w_2)$, respectively. Also let $x_1 \in (x_0, x_\infty)$ be such that $\bar{R}(x_1)$ is a multiple of $(1,1) \in \mathbb{Z}^2$.

Using Lemma 2.3.3, we may assume that there is exactly one point of inflection x_{poi} between x_0 and x_∞ and that $f'(x_{poi})/g'(x_{poi}) \approx 0$. Let $\check{v}_\infty = (0,1)$ and $\check{x}_\infty$ be the unique point between x_0 and $\hat{x}_\infty$ with $\bar{R}(\check{x}_\infty) \in \mathbb{R}^2$ proportional to $(0,1)$. See Figure 2.11 (a). Write $\mathcal{P}_\alpha = v_0\hat{v}_\infty\mathcal{P}^+$ and let $\mathcal{P}^0 := v_0\check{v}_\infty\mathcal{P}^+$ with $x(\check{v}_\infty) = \check{x}_\infty$ and $\mathcal{P}^2 := \mathcal{P}_\beta$. We examine each of the nonzero summands of

$$\langle \partial^2\tilde{\alpha}, \beta\rangle = \sum_\gamma \langle \partial\tilde{\alpha}, \gamma\rangle\langle \partial\gamma, \beta\rangle = 0 \tag{2.3.4}$$

where $\tilde{\alpha}$ is an orbit set such that $\mathcal{P}_{\tilde{\alpha}} = \mathcal{P}^0$, i.e. it is obtained from α by replacing one orbit at $\hat{x}_\infty$ with an orbit at $\check{x}_\infty$.

Consider an IP path $\mathcal{P}^1 = \mathcal{P}_\gamma$ for an orbit set γ such that there is a holomorphic curve from $\tilde{\alpha}$ to γ, as well as one from γ to β. By Lemma 2.2.19, there are two nonlocal $\mathcal{R}^i_{j_i}$, each with two extreme edges, and they share a single edge v. As before, the IP region $\mathcal{R}$ between $\mathcal{P}^0$ and $\mathcal{P}^2$ has exactly three $\mathbb{R}$-loose edges and so $\check{v}_\infty$ must be an extreme edge of $\mathcal{R}' \in (\mathcal{R}^1_{j_1}, \mathcal{R}^2_{j_2})$. There are two cases:

(i) ($\check{v}_\infty$ is the west extreme edge of $\mathcal{R}'$.) Positivity of $\mathcal{R}'$ at $x = \hat{x}_\infty + \epsilon$ forces that $\mathcal{R}'$ is a bigon with the extreme east edge $v = \hat{v}_\infty$.

(ii) ($\check{v}_\infty$ is the east extreme edge of $\mathcal{R}'$.) Note v_0 is the only other edge of $\mathcal{R}$ with $x(\cdot) < \check{x}_\infty$. Since v_0 and $\check{v}_\infty$ do not form a bigon, we need $x(v) < \check{x}_\infty$ and $\mathcal{R}'$ must be a triangle formed by v_0, v and $\check{v}_\infty$. In particular, $v = v_1 := (1,1) \in \mathbb{Z}^2$ with $x(v_1) = x_1$.

This completely describes all possible IP paths $\mathcal{P}_\gamma$ for nonzero summands of (2.3.4).

We now consider decorations of $\mathcal{R}$ and $(\mathcal{R}^i_{j_i})$. Note that $\mathcal{R}$ differs from $\mathcal{R}_{\alpha,\beta}$ by replacing the edge at $\hat{x}_\infty$ to the one at $\check{x}_\infty$. Hence, we can make $I(\mathcal{R}) = 2$ by:

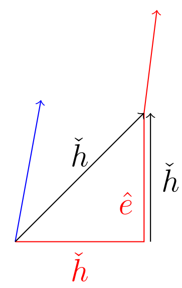

Figure 2.12: A decoration for Case 2.

- Label $\check{v}_\infty$ as $\check{e}$ if $\hat{v}_\infty$ is labeled $\hat{h}$; label it $\check{h}$ otherwise.
- Use an induced decoration of $\mathcal{R}_{\alpha,\beta}$ for all other edges of $\mathcal{R}$.

See Figure 2.12 for one possible decoration. Imposing $I(\mathcal{R}') = 1$ determines the label of v in each of the above two cases as follows:

(i) ($\check{v}_\infty$ is the west extreme edge of $\mathcal{R}'$) $v = \hat{v}_\infty$ is labeled $\hat{h}$ if $\check{v}_\infty$ is labeled $\check{e}$; v is labeled $\hat{e}$ otherwise.

(ii) ($\check{v}_\infty$ is the east extreme edge of $\mathcal{R}'$) $v = v_1$ is labeled $\check{h}$ if both v_0 and $\check{v}_\infty$ are labeled $\check{e}$; v is labeled $\check{e}$ otherwise.

Let η be the orbit set whose associated decorated IP path $\mathcal{P}_\eta$ is $v_1\mathcal{P}^+$ decorated as above. We have shown that if γ is an orbit set with $\langle \partial\tilde{\alpha}, \gamma\rangle\langle\partial\gamma, \beta\rangle \neq 0$, then $\gamma = \eta$ or $\gamma = \alpha$. By the induction hypothesis and Proposition 2.3.5,

$$\langle \partial\tilde{\alpha}, \eta\rangle \neq 0, \quad \langle \partial\eta, \beta\rangle \neq 0, \quad \langle \partial\tilde{\alpha}, \alpha\rangle \neq 0$$

so we have

$$\langle \partial\alpha, \beta\rangle \neq 0.$$

□

Chapter 3

ECC of T^3

In the previous section, we have considered a contact manifold $(I \times T^2, \lambda)$. In this section, we show that, after small modifications, the same combinatorial description applies to the closed manifold $T^3 = \mathbb{R}/\mathbb{Z} \times T^2$ with a contact form λ, which is a small perturbation of a T^2-invariant contact form.

3.1 Preliminaries

Consider $(f, g) : \mathbb{R} \to \mathbb{R}^2$ satisfying $(f, g) \times (f', g') > 0$ as before so that $\tilde{\lambda} = -g dt_1 + f dt_2$ is a contact form on $\mathbb{R} \times T^2$. Suppose further that (f, g) is $\mathbb{Z}$-periodic and regard it as a function on $\mathbb{R}/\mathbb{Z}$ with coordinate x. Consider $\bar{\lambda} = -g dt_1 + f dt_2$ on $T^3 = \mathbb{R}/\mathbb{Z} \times T^2$ with coordinates (x, t_1, t_2). Then, $(T^3, \bar{\lambda})$ is a contact manifold similar to $(I \times T^2, \bar{\lambda})$ previously discussed. In this section, we discuss some differences from the previous treatment.

Definition 3.1.1. For each S^1-family $\bar{\rho}$ of Reeb orbits of $(T^3, \bar{\lambda})$ with action less than L, consider slightly modified Morse functions $H_{\bar{\rho}}$ on $\bar{\rho}$ as follows: Let $\tilde{\delta} = \delta/N$ for some $N \gg 0$. Then,

- $H_{\bar{\rho}}$ attains the maximum at $-\tilde{\delta}$ and the minimum at $+\delta$ if $\bar{\rho}$ is convex.
- $H_{\bar{\rho}}$ attains the maximum at $-\delta$ and the minimum at $+\tilde{\delta}$ if $\bar{\rho}$ is concave.

We say that a contact form λ on T^3 is a good perturbation of $\bar{\lambda}$ if it satisfies conditions (ii) - (iv) of Definition 2.1.2 and condition (i) with the above Morse functions $H_{\bar{\rho}}$ instead. Let $\bar{J}$ be defined using (2.1.11), as before. We say that a pair (λ, J) of a contact form on $\mathbb{R}/\mathbb{Z} \times T^2$ and a generic admissible almost complex structure J on $\mathbb{R} \times (\mathbb{R}/\mathbb{Z} \times T^2)$ is a good perturbation of $(\bar{\lambda}, \bar{J})$ if λ is a good perturbation of $\bar{\lambda}$ and (λ, J) is sufficiently close to $(\bar{\lambda}, \bar{J})$ in the sense of Lemma 3.2.2 in addition to Lemma 2.2.1, Proposition 2.2.21 and Proposition 2.3.5.

Throughout this section, assume λ is a good perturbation of $\bar{\lambda}$ and (λ, J) is a good perturbation of $(\bar{\lambda}, \bar{J})$. We will also use the same notation $\bar{\lambda}$ and λ for a $\mathbb{Z}$-periodic contact form on $\mathbb{R} \times T^2$ and a contact form on $\mathbb{R}/\mathbb{Z} \times T^2$.

We induce order on $\mathbb{R}/\mathbb{Z}$ from $([0,1), <)$. Then, similarly to the ordered product notation of an orbit set of $(I \times T^2, \lambda)$ we can write an orbit set α of $(\mathbb{R}/\mathbb{Z} \times T^2, \lambda)$ in the ordered product notation. Also, consider the covering map $\pi : \mathbb{R} \times T^2 \to \mathbb{R}/\mathbb{Z} \times T^2$. Any orbit $\tilde{\rho}$ of $(\mathbb{R} \times T^2, \lambda)$ projects to an orbit ρ of $(\mathbb{R}/\mathbb{Z} \times T^2, \lambda)$.

Definition 3.1.2. Let $\alpha = \alpha_1 \cdots \alpha_n$ be an orbit set of $(\mathbb{R}/\mathbb{Z} \times T^2, \lambda)$. We say that an orbit set $\tilde{\alpha}$ of $(\mathbb{R} \times T^2, \bar{\lambda})$ is a *lift* of α if there is a bijection between embedded orbits $\tilde{\alpha}_{i'}$ appearing in the ordered product notation of $\tilde{\alpha}$ and embedded orbits α_i appearing in the ordered product notation of α so that, for each such pair, $\tilde{\alpha}_{i'}$ projects to α_i under π. If β is another orbit set of λ, we say that the pair $(\tilde{\alpha}, \tilde{\beta})$ is an *admissible lift* of (α, β) under π if $\tilde{\alpha}$ and $\tilde{\beta}$ are orbit sets of $(I \times T^2, \lambda) \subset (\mathbb{R} \times T^2, \lambda)$ where I has length less than $1 + 2\epsilon$.

Let $\tilde{C} \in \mathcal{M}(\tilde{\alpha}, \tilde{\beta})$ be a holomorphic curve in $(\mathbb{R} \times (\mathbb{R} \times T^2), J)$. Then, it projects to a holomorphic curve $C \in \mathcal{M}(\alpha, \beta)$ so that $(\tilde{\alpha}, \tilde{\beta})$ is a lift of (α, β). We point out that if a holomorphic curve C in $\mathbb{R} \times (\mathbb{R}/\mathbb{Z} \times T^2)$ has genus zero, then it necessarily lifts to a holomorphic curve $\tilde{C}$ in $\mathbb{R} \times (\mathbb{R} \times T^2)$.

We modify the definitions of combinatorial objects in the following way. Refer to Section 1.3 for more details. By genericity of (f, g), assume that $f'(0)/g'(0) \notin \mathbb{Q} \cup \{\infty\}$ and recall the order on $\mathbb{R}/\mathbb{Z}$ induced from $([0,1), <)$.

Definition 3.1.3. (a) An *IP path* $\mathcal{P}$ on $\mathbb{R}/\mathbb{Z}$ is an n-tuple of edges (v_i), satisfying the conditions of Definition 1.3.1 except that $x(v) \in \mathbb{R}/\mathbb{Z}$.

(b) If $\mathcal{P}^+ = (v_i^+)$ and $\mathcal{P}^- = (v_j^-)$ are two IP paths with $\sum_i v_i^+ = \sum_j v_j^-$ and $\sigma_0 \in \mathbb{Z}^2$ is a vector, then an *IP region* on $\mathbb{R}/\mathbb{Z}$ with a *reference slice class* σ_0 is the triple $(\mathcal{P}^+, \mathcal{P}^-, \sigma_0)$.

(c) If (v_k) is an ordering of $\partial^{\pm}\mathcal{R}$ with non-decreasing $x(v_k)$, a *realization* of an IP region $\mathcal{R}$ is a continuous map Φ from $[0,1] \times \cup_{k \in \mathbb{Z}} [k, k+1]$ to $\mathbb{R}^2$ such that:

- $\Phi(0,0) - \Phi(1,0) = \sigma_0 \in \mathbb{Z}^2$.
- If v_k is a positive edge, then $\Phi(1,k) = \Phi(1,k-1) + v_k$ and $\Phi(0,k) = \Phi(0,k-1)$.
- If v_k is a negative edge, then $\Phi(0,k) = \Phi(0,k-1) + v_k$ and $\Phi(1,k) = \Phi(1,k-1)$.

Here v_k for $k \in \mathbb{Z}$ is interpreted as modulo m where m is the number of edges of $\mathcal{R}$.

(d) The *slice class* of $\mathcal{R} = (\partial^+\mathcal{R}, \partial^-\mathcal{R}, \sigma_0)$ is

$$\sigma(x) := \sigma_0 - \sum_{\substack{v \in \partial^+\mathcal{R} \\ 0 < x(v) \leq x}} v + \sum_{\substack{w \in \partial^-\mathcal{R} \\ 0 < x(w) \leq x}} w.$$

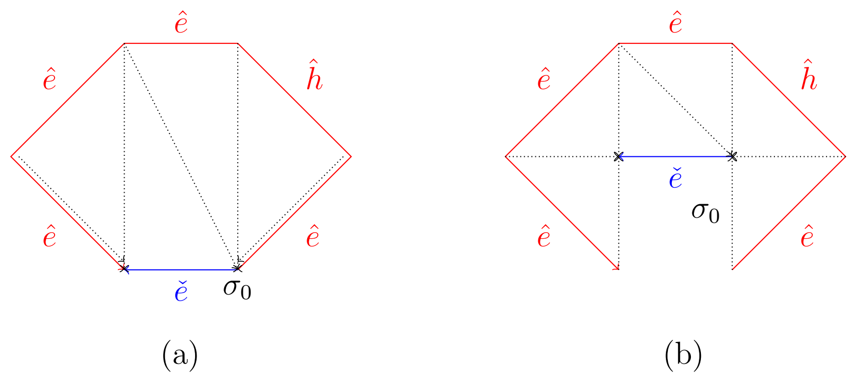

Figure 3.1: Two IP regions $\mathcal{R}$ and $\mathcal{R}'$ with $\partial^{\pm}\mathcal{R} = \partial^{\pm}\mathcal{R}'$ but $\sigma_0 = 0$ and $\sigma_0 = (0, 1)$

(e) If $\Phi([0, 1] \times \{k\})$ is a single point for some realization Φ and some $0 \leq k \leq m$, we say that $\mathcal{R}$ is *simply connected.*

Note that a realization Φ of $\mathcal{R}$ with $\partial^+\mathcal{R} = (v_i^+)$ is "periodic with monodromy" $\sum_i v_i^+ \in \mathbb{Z}^2$: if $\mathcal{R}$ has m edges, then $\Phi(\cdot, \cdot + m) = \Phi(\cdot, \cdot) + \sum_i v_i^+ \in \mathbb{R}^2$.

Definition 3.1.4. Given an orbit set $\alpha = \alpha_1 \cdots \alpha_n$ in $(\mathbb{R}/\mathbb{Z} \times T^2, \lambda)$, the IP path $\mathbb{R}/\mathbb{Z}$ *associated to* α is the n-tuple (v_i) where $v_i = ([\alpha_i], x(\alpha_i))$. We denote it by $\mathcal{P}_\alpha$. Given a holomorphic curve $C \in \mathcal{M}(\alpha, \beta)$, the IP region on $\mathbb{R}/\mathbb{Z}$ *associated to* C is $(\mathcal{P}_\alpha, \mathcal{P}_\beta, [\sigma_C(x_0)])$ denoted $\mathcal{R}_C$.

Note that, as opposed to $\mathcal{R}_C$, $\mathcal{R}_{\alpha,\beta}$ lacks the information about the reference slice class and hence, ambiguous. This is related to the index ambiguity, which we discuss below. Figure 3.1 shows realizations of two IP regions $\mathcal{R}$ and $\mathcal{R}'$ with the same positive and negative edges, but $\sigma_0 = 0$ in (a), whereas $\sigma_0 = (0, 1)$ in (b).

Definition 3.1.5. (ECH index) Let $I(\mathcal{R})$ be the combinatorial ECH index where $Area(\mathcal{R}) := Area(\mathrm{im}(\Phi|_{[0,1]\times[0,m]}))$ in (1.3.4). Here, Φ is a realization of $\mathcal{R}$. Note this depends on σ_0.

We recall that for $Z, Z' \in H_2(\mathbb{R}/\mathbb{Z} \times T^2, \alpha, \beta)$, the index ambiguity

$$I(\alpha, \beta, Z) - I(\alpha, \beta, Z') = \langle 2PD(\Gamma), Z - Z' \rangle$$

may be nonzero if $Z - Z' \in H_2(\mathbb{R}/\mathbb{Z} \times T^2)$ projects to a nonzero class in $H_1(\mathbb{R}/\mathbb{Z}) \otimes H_1(T^2)$.

Lemma 3.1.6. *Let α and β be orbit sets and let C and C' be holomorphic curves from α to β. Let $\sigma_0 = [\mathcal{S}_C(0)]$ and $\sigma_0' = [\mathcal{S}_{C'}(0)]$. Then,*

$$I(\mathcal{R}_C) = I(C), \qquad I(\mathcal{R}_{C'}) = I(C')$$

and they reflect the index ambiguity by

$$I(\mathcal{R}_C) - I(\mathcal{R}_{C'}) = (\sigma_0 - \sigma_0') \times (2[\alpha]).$$

Proof. Since τ is still the restriction of a global trivialization of $T(\mathbb{R}/\mathbb{Z} \times T^2)$, we have $c_\tau(Z) = 0$ for any $Z \in H_2(\mathbb{R}/\mathbb{Z} \times T^2, \alpha, \beta)$ in (2.1.1) and (2.1.3). Then, the rest is a straightforward modification of Proposition 2.2.6 using a surface S whose slice $\mathcal{S}(0)$ is a disjoint union of straight curves in $\{0\} \times T^2$ in homology class σ_0. The second part is straightforward. □

In Figure 3.1, $\sigma_0 = 0, \sigma_0' = (0, 1)$ and $[\alpha] = [\beta] = (-1, 0)$, resulting in $I(\mathcal{R}_C) = 3$ and $I(\mathcal{R}_{C'}) = 1$.

3.2 The theorem

Theorem 3.2.1. *Let $\bar{\lambda} = -g dt_1 + f dt_2$ be a T^2-invariant contact form on $\mathbb{R} \times T^2$ and $\bar{J}$ be the almost complex structure on $\mathbb{R} \times (\mathbb{R} \times T^2)$ defined by (2.1.11). Let (λ, J) be a good perturbation of $(\bar{\lambda}, \bar{J})$. Given a pair of admissible orbit sets α and β of $(\mathbb{R}/\mathbb{Z} \times T^2, \lambda)$ with $I(\alpha, \beta) = 1$, any holomorphic curve $C \in \mathcal{M}(\alpha, \beta)$ lifts to a holomorphic curve $\tilde{C} \in \mathcal{M}(\tilde{\alpha}, \tilde{\beta})$ in $\mathbb{R} \times (I \times T^2) \subset \mathbb{R} \times (\mathbb{R} \times T^2)$ for an admissible lift $(\tilde{\alpha}, \tilde{\beta})$. In particular,*

$$\langle \partial \alpha, \beta \rangle = \sum_{(\tilde{\alpha}, \tilde{\beta})} \langle \partial \tilde{\alpha}, \tilde{\beta} \rangle$$

where the summation is over distinct (modulo $\mathbb{Z}$) admissible lifts $(\tilde{\alpha}, \tilde{\beta})$ of (α, β).

The $+2\epsilon$ term allows that the east extreme end and the west extreme end of $\tilde{C}$ may occur at the same orbit of λ.

Let $C \in \mathcal{M}(\alpha, \beta)$. We define

$$I^a(\mathcal{R}_C) = Area(\Phi) - 2\#\{\text{edges of } \mathcal{R}\} - 2n$$

and

$$I^c(\mathcal{R}_C) = \sum_{v \in \partial^{\pm} \mathcal{R}_C} (cz(v) + 1)$$

so

$$I(\mathcal{R}_C) = I^a(\mathcal{R}_C) + I^c(\mathcal{R}_C) + 2n.$$

To aid computation, let $\mathcal{R}'$ be a decorated simply connected IP region obtained by "slicing" $\mathcal{R}_C$ along σ_0 and labeling each new edge as S^1-tight and $\mathbb{R}$-loose: if $\sigma_0 = nw$ for a primitive vector w and $n \geq 0$ and $\partial^+ \mathcal{R}_C = (v_i)$, then let

$$\partial^+ \mathcal{R}' := (-w, \cdots, -w, (v_i), w \cdots, w)$$
$$\partial^- \mathcal{R}' := \partial^- \mathcal{R}_C$$

where $-w$ is repeated n times at the beginning and w is repeated n times at the end, each labeled $\check{h}$. It is easy to verify

$$I(\mathcal{R}_C) = I(\mathcal{R}'), \qquad I^a(\mathcal{R}_C) = I^a(\mathcal{R}'), \qquad I^c(\mathcal{R}_C) + 2n = I^c(\mathcal{R}').$$

Since $\mathcal{R}_C$ and $\mathcal{R}'$ have the same slice classes and $\mathcal{R}_C$ satisfies the positivity condition, we have $I^a(\mathcal{R}_C) \geq -2$ and $I^c(\mathcal{R}_C) + 2n \leq 3$.

The following two lemmas show that $\mathcal{R}_C$ must be simply connected. Suppose $\mathcal{R}_C$ is not simply connected. Since $\mathcal{R}_C$ is minimal with $I(\mathcal{R}_C) = 1$, we are forced to have $n = 1$ and $\mathcal{R}_C$ has exactly one edge v with $cz(v) = 0$. Similarly to the Morse-Bott argument in Section 2.2, we claim the following:

Lemma 3.2.2. *Consider $\mathbb{R}/\mathbb{Z} \times T^2$ with a T^2-invariant contact form $\bar{\lambda}$ and the admissible almost complex structure $\bar{J}$ on $\mathbb{R} \times (\mathbb{R}/\mathbb{Z} \times T^2)$ by (2.1.11). Let λ_n be a sequence of good perturbations of $\bar{\lambda}$ and let J_n be a generic admissible almost complex structures for λ_n and suppose (λ_n, J_n) converges to $(\bar{\lambda}, \bar{J})$. Fix a non-simply connected IP region $\mathcal{R}$ with $I(\mathcal{R}) = 1$ and one $\mathbb{R}$-loose edge. If (λ_n, J_n) is sufficiently close to $(\bar{\lambda}, \bar{J})$, then there is no J_n-holomorphic curve whose associated region is $\mathcal{R}$.*

Proof. If there is a sequence of J_n-holomorphic curves C_n whose associated IP region is $\mathcal{R}$, then after passing to a subsequence, C_n converges to a $\bar{J}$-holomorphic building $\bar{C}$ as before and we can consider the partition $(\mathcal{R}^i_j)$ of $\mathcal{R}$ associated to $\bar{C}$. Suppose that each $\mathcal{R}^i_j$ is simply connected. Consider the collapsed dual graph Γ that contains a vertex for each realization of nonlocal IP regions appearing in

$$\Phi([0,1] \times \cup_{k\in\mathbb{Z}}[k, k+1]).$$

Note that Γ is connected and it does contain a cycle since $\mathcal{R}$ is minimal. Since $\mathbb{Z}$ acts on the vertices and edges of Γ by a deck transform, $\Gamma/\mathbb{Z}$ contains one cycle. Then, a modification of Lemma 2.2.19 for the collapsed dual graph with one cycle implies that

$$2m \leq m + l \tag{3.2.1}$$

where m is the number of nonlocal IP regions in the partition of $\mathcal{R}$ and $l = 1$ is the number of $\mathbb{R}$-loose edges of $\mathcal{R}$. Hence, $m = l = 1$, contradicting the assumption that $\mathcal{R}^1_1$ is simply connected. Hence, there is a $\mathcal{R}' \in (\mathcal{R}^i_j)$ which is not simply connected. We claim that all other $\mathcal{R}^i_j$ are simply connected. Suppose $\mathcal{R}' = \mathcal{R}^{i'}_{j'}$ and $\mathcal{R}'' = \mathcal{R}^{i''}_{j''}$ are two non-simply connected IP regions in $(\mathcal{R}^i_j)$. Since $i \neq i'$ necessarily, we assume $i' < i''$ without loss of generality. Since $\partial^+\mathcal{R}''$ is a nontrivial IP path, $\mathcal{P}^{i''-1} = \partial^+\mathcal{R}^{i''}$ is nontrivial as well, and all the lattice points on the realization of $\mathcal{P}^{i''-1}$ must be internal to the corresponding realization of $\mathcal{R}$. This contradicts the minimality of $\mathcal{R}$. Hence, the partition $(\mathcal{R}^i_j)$ contains $(m-1)$ nonlocal IP regions, each with at least two $\mathbb{R}$-loose edges. Moreover, the collapsed dual graph of $(\mathcal{R}^i_j)$ contains no cycle since $\mathcal{R}$ is minimal. Another modification of Lemma 2.2.19 for such a partition implies

$$2(m-1) \leq m - 1 + l, \tag{3.2.2}$$

so we have that $m \leq l + 1 = 2$.

By the Θ-constraint, there must be a nonconstant Morse flow in $\bar{C}$ and since all positive and negative edges of $\mathcal{R}$ are S^1 tight, m must be 2. The argument proceeds similarly to the

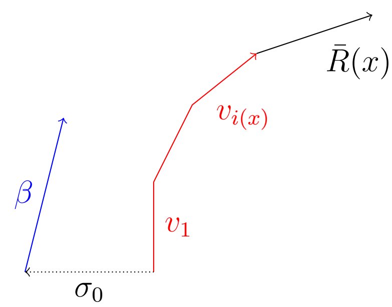

Figure 3.2: $\partial^+\mathcal{R}$ on the interval where $\bar{R}(x) \times v_{i(x)} > 0$.

last step of the proof of Theorem 1.4.1. Let w be the shared edge between the two nonlocal IP regions $\mathcal{R}^1_{i_1}$ and $\mathcal{R}^2_{i_2}$ and by symmetry, assume the S^1-family of orbits $\bar{\rho}$ at $x = x(w)$ is convex. By (3.2.2), the non-simply connected IP region has no $\mathbb{R}$-loose edges. Hence, $\mathcal{R}^1_{i_1}$ is not simply connected, and $\mathcal{R}^2_{i_2}$ is simply connected. Let $\bar{C}^1$ and $\bar{C}^2$ be the nontrivial components of $\bar{C}$ with a negative end and a positive end at $\bar{\rho}(\theta^-)$ and $\bar{\rho}(\theta^+)$, respectively. Then,

$$\Theta(\bar{C}^1) = a_1\tilde{\delta} - \theta^- = 0,$$
$$\Theta(\bar{C}^2) = a_2\tilde{\delta} + \delta + \theta^+ = 0$$

where $a_i \ll N$ are the total number of elliptic edges of $\mathcal{R}_i$. Since $H_{\bar{\rho}}$ has the maximum at $\theta = -\tilde{\delta}$ and the minimum at $\theta = \delta$, there can be no Morse flow from θ^- to θ^+ on $\bar{\rho}$. □

Lemma 3.2.3. *There does not exist any non-simply connected IP region $\mathcal{R}$ such that:*

- *$\mathcal{R}$ contains no $\mathbb{R}$-loose edges, and*
- *$\partial^+\mathcal{R} = \mathcal{P}_\alpha$, $\partial^-\mathcal{R} = \mathcal{P}_\beta$ for some admissible orbit sets α and β with $\langle\partial\alpha, \beta\rangle \neq 0$.*

Proof. Suppose $0 < M < \infty$ is the smallest number such that there is a IP region with M positive edges, satisfying the above conditions. Let $\mathcal{R}_M$ be any such IP region and let $\partial^+\mathcal{R}_M = v_1 \cdots v_M$ be an IP path with an induced decoration. For $x(v_1) < x < 1$, let $i(x)$ denote the largest i with $x(v_i) < x$ and let x_0 be the smallest $x(v_1) < x < 1$ satisfying

$$A(x) := \bar{R}(x) \times v_{i(x)} = 0.$$

Such x_0 exists for the following reason: $A(x)$ is continuous except at each $x = x(v_i)$, but since every v_i is concave, $A(x(v_i) + \epsilon) > 0$. Hence, if x_0 does not exist, $A(x) > 0$ for all $x(v_1) < x < 1$, and $v_i \times v_{i+1} < 0$ for each $1 \leq i < M$ and contradicts that (f, g) rotates around the origin $n \geq 1$ times. See Figure 3.2 for $v_1 \cdots v_{i(x)}$ for any $x(v_1) < x < x_0$. For simplicity, we denote $i(x_0)$ by i_0.

By the choice of x_0, λ is convex at x_0 so we can replace v_{i_0} of $\mathcal{R}_M$ with an edge $\check{v}_0$ at $x = x_0$. We replace the label on this edge from concave to convex, while keeping S^1-tightness. The new decorated IP region $\tilde{\mathcal{R}}$ has $I(\tilde{\mathcal{R}}) = 2$ and is associated to admissible orbit sets $\tilde{\alpha}$ and β. We examine each nonzero summand of

$$\langle \partial^2 \tilde{\alpha}, \beta \rangle = \sum_{\gamma} \langle \partial \tilde{\alpha}, \gamma \rangle \langle \partial \gamma, \beta \rangle = 0. \tag{3.2.3}$$

Let γ be an orbit set which corresponds to a nonzero summand of (3.2.3). Since $\tilde{\mathcal{R}}$ is minimal with exactly one $\mathbb{R}$-loose edge, we may apply similar analysis as in Lemma 3.2.2 to the partition $(\mathcal{R}^i_j)$ of $\tilde{\mathcal{R}}$ by $\mathcal{P}_\gamma$. Here, we have $m \geq 2$ and since (3.2.1) cannot be satisfied, $(\mathcal{R}^i_j)$ must contain a non-simply connected IP region. By (3.2.2), there are exactly two nonlocal regions $\mathcal{R}'$ and $\mathcal{R}''$ in $(\mathcal{R}^i_j)$ and since $\tilde{\mathcal{R}}$ is minimal, without loss of generality $\mathcal{R}'$ is simply connected and $\mathcal{R}''$ is not. By the equality condition of (3.2.2), $\mathcal{R}'$ and $\mathcal{R}''$ share one edge w and since $\mathcal{R}'$ has two $\mathbb{R}$-loose edges, $\check{v}_0$ must be an edge of $\mathcal{R}'$, while all edges of $\mathcal{R}''$ are $\mathbb{R}$-tight.

By the minimality of M, $\mathcal{R}''$ must have at least M positive edges and hence, $\mathcal{R}'$ must be a bigon. One such bigon is between $\check{v}_0$ and v_{i_0}. In order to have another nonzero summand in (3.2.3), there must be a bigon between $\check{v}_0$ and an edge v'_{i_0} with $x(v'_{i_0}) > x(\check{v}_0)$. Replace v_{i_0} of $\mathcal{R}_M$ with v'_{i_0} while keeping S^1-tightness and call the new decorated IP region $\mathcal{R}^{new}_M$. Note that $\mathcal{R}^{new}_M$ is identical to $\mathcal{R}_M$ except $x(v'_{i_0}) > x(v_{i_0})$. We can repeat the above analysis with $\mathcal{R}_M$ replaced by $\mathcal{R}^{new}_M$: the conclusion is that there must be yet another simply connected IP region with M positive edges. Since $\bar{R}(x)$ takes on a multiple of v_{i_0} finitely many times, this cannot continue indefinitely and, at some point, there is only one nonlocal bigon with an edge $\check{v}_0$. This is a contradiction and completes the proof of the claim. □

Proof. (of Theorem 3.2.1) Lemma 3.2.2 and Lemma 3.2.3 show that $\mathcal{R}_C$ is simply connected, i.e. there is $x_0 \in \mathbb{R}/\mathbb{Z}$ such that $\bar{R}(x_0) \times \sigma(x_0) = 0$. If $\sigma(x_0 + \epsilon)$ or $\sigma(x_0 - \epsilon)$ is zero, then we are done by the equality condition of Lemma 2.2.1. Otherwise, by analyzing the ends of C at $x = x_0$ similarly to Proposition 2.2.7, we find that ind(C) has a contribution of at least three to $\sum_\rho (cz^{ind}(\rho) + 1)$ and conclude that $g(C) = 0$.

The genus zero condition implies that C lifts to a holomorphic curve $\tilde{C}$ in $\mathbb{R} \times (\mathbb{R} \times T^2)$. Since all ends of C at ρ with $x(\rho) \neq x_0$ have $cz^{ind}(\rho) = -1$, $\tilde{C}$ must have the west extreme end at some lift $\tilde{x}_0$ of x_0 and the east extreme end at $\tilde{x}_0 + n$ for some integer $n \geq 1$. Note that $\bar{R}(x_0) \times [\mathcal{S}_{\tilde{C}}(\tilde{x}_0 + i + \epsilon)] \geq 0$, for all i by positivity, and

$$0 = \bar{R}(x_0) \times \sigma(x_0) = \bar{R}(x_0) \times \left(\sum_i [\mathcal{S}_{\tilde{C}}(\tilde{x}_0 + i + \epsilon)] \right) \geq 0.$$

Hence, $\bar{R}(x_0) \times [\mathcal{S}_{\tilde{C}}(\tilde{x}_0 + 1 + \epsilon)] = 0$ and $\tilde{C}$ must have an $\mathbb{R}$-loose edge at $\tilde{x}_0 + 1$ by Corollary 2.2.3. Since C has only two $\mathbb{R}$-loose ends, the only $\mathbb{R}$-loose ends of $\tilde{C}$ occur at extreme ends, which means $\tilde{C}$ has the east extreme end at $\tilde{x}_0 + 1$. This complete the proof. □

Part II

Symplectic embeddings into four-dimensional concave toric domains

Chapter 4

Introduction

4.1 ECH capacities

Let (X, ω) be a symplectic four-manifold, possibly with boundary or corners, noncompact, and/or disconnected. Its ECH capacities are a sequence of real numbers

$$0 = c_0(X, \omega) \leq c_1(X, \omega) \leq c_2(X, \omega) \leq \cdots \leq \infty. \tag{4.1.1}$$

The ECH capacities were introduced in [8], see also the exposition in [10]; we will review the definition in the cases relevant to this paper in §6.1.

The following are some key properties of ECH capacities:

(Monotonicity) If there exists a symplectic embedding $(X, \omega) \to (X', \omega')$, then $c_k(X, \omega) \leq c_k(X', \omega')$ for all k.

(Conformality) If $r > 0$ then

$$c_k(X, r\omega) = rc_k(X, \omega).$$

(Disjoint union)

$$c_k\left(\coprod_{i=1}^{n}(X_i, \omega_i)\right) = \max_{k_1+\cdots+k_n=k} \sum_{i=1}^{n} c_{k_i}(X_i, \omega_i).$$

(Ellipsoid) If $a, b > 0$, define the ellipsoid

$$E(a, b) = \left\{(z_1, z_2) \in \mathbb{C}^2 \;\middle|\; \frac{\pi|z_1|^2}{a} + \frac{\pi|z_2|^2}{b} \leq 1\right\}.$$

Then $c_k(E(a, b)) = N(a, b)_k$, where $N(a, b)$ denotes the sequence of all nonnegative integer linear combinations of a and b, arranged in nondecreasing order, indexed starting at $k = 0$.

Here we are using the standard symplectic form on $\mathbb{C}^2 = \mathbb{R}^4$. In particular, define the ball

$$B(a) = E(a, a).$$

It then follows from the Ellipsoid property that

$$c_k(B(a)) = ad \tag{4.1.2}$$

where d is the unique nonnegative integer such that

$$\frac{d^2 + d}{2} \leq k \leq \frac{d^2 + 3d}{2}. \tag{4.1.3}$$

It was shown by McDuff [18], see also the survey [9], that there exists a symplectic embedding $\operatorname{int}(E(a, b)) \to E(c, d)$ if and only if $N(a, b)_k \leq N(c, d)_k$ for all k. Thus ECH capacities give a sharp obstruction to symplectically embedding one (open) ellipsoid into another. It follows from work of Frenkel-Müller [5], see [9, Cor. 11], that ECH capacities also give a sharp obstruction to symplectically embedding an open ellipsoid into a polydisk

$$P(a, b) = \left\{(z_1, z_2) \in \mathbb{C}^2 \;\middle|\; \pi|z_1|^2 \leq a,\ \pi|z_2|^2 \leq b\right\}.$$

On the other hand, ECH capacities do not give sharp obstructions to embedding a polydisk into an ellipsoid. For example, if there is a symplectic embedding $P(1, 1) \to E(a, 2a)$, then ECH capacities only imply that $a \geq 1$, but the Ekeland-Hofer capacities imply that $a \geq 3/2$, see [8, Rmk. 1.8]. Another example is that if there is a symplectic embedding from $P(1, 2)$ into the ball $B(c)$, then both ECH capacities and Ekeland-Hofer capacities only imply that $c \geq 2$; but in fact it was recently shown by Hind-Lisi [6] that $c \geq 3$. In particular, the inclusions $P(1, 1) \to E(3/2, 3)$ and $P(1, 2) \to B(3)$ are "optimal" in the following sense:

Definition 4.1.1. A symplectic embedding $\phi : (X, \omega) \to (X', \omega')$ is *optimal* if there does not exist a symplectic embedding $(X, r\omega) \to (X', \omega')$ for any $r > 1$.

Remark 4.1.2. It follows from the Monotonicity and Conformality properties that if $0 < c_k(X, \omega) = c_k(X', \omega')$ for some k, and if a symplectic embedding $(X, \omega) \to (X', \omega')$ exists, then it is optimal.

4.2 Concave toric domains

We would like to compute more examples of ECH capacities and find more examples of sharp embedding obstructions and optimal symplectic embeddings. An interesting family of symplectic four-manifolds is obtained as follows. If Ω is a domain in the first quadrant of the plane, define the "toric domain"

$$X_\Omega = \left\{z \in \mathbb{C}^2 \;\middle|\; \pi(|z_1|^2, |z_2|^2) \in \Omega\right\}.$$

For example, if Ω is the triangle with vertices $(0,0)$, $(a,0)$, and $(0,b)$, then X_Ω is the ellipsoid $E(a,b)$.

The ECH capacities of toric domains X_Ω when Ω is convex and does not touch the axes were computed in [8, Thm. 1.11], see [10, Thm. 4.14]. Also, the assumption that Ω does not touch the axes can be removed in some and conjecturally all cases. In this paper we consider the following new family of toric domains:

Definition 4.2.1. A *concave toric domain* is a domain X_Ω where Ω is the closed region bounded by the horizontal segment from $(0,0)$ to $(a,0)$, the vertical segment from $(0,0)$ to $(0,b)$, and the graph of a convex function $f:[0,a]\to[0,b]$ with $f(0)=b$ and $f(a)=0$. The concave toric domain X_Ω is *rational* if f is piecewise linear and f' is rational wherever it is defined.

McDuff showed in [18, Cor. 2.5] that the ECH capacities of an ellipsoid $E(a,b)$ with a/b rational are equal to the ECH capacities of a certain "ball packing" of the ellipsoid, namely a certain finite disjoint union of balls whose interior symplectically embeds into the ellipsoid filling up all of its volume. These balls are determined by a "weight expansion" of the pair (a,b). In the present work, we generalize this to give a similar formula for the ECH capacities of any rational concave toric domain. In §4.6 we will give a different formula for the ECH capacities of concave toric domains which are not necessarily rational.

4.3 Weight expansions

Let X_Ω be a rational concave toric domain. The *weight expansion* of Ω is a finite unordered list of (possibly repeated) positive real numbers $w(\Omega)=(a_1,\ldots,a_n)$ defined inductively as follows.

If Ω is the triangle with vertices $(0,0)$, $(a,0)$, and $(0,a)$, then $w(\Omega)=(a)$.

Otherwise, let $a>0$ be the largest real number such that the triangle with vertices $(0,0)$, $(a,0)$, and $(0,a)$ is contained in Ω. Call this triangle Ω_1. The line $x+y=a$ intersects the graph of f in a line segment from $(x_2,a-x_2)$ to $(x_3,a-x_3)$ with $x_2\le x_3$. Let Ω_2' denote the portion of Ω above the line $x+y=a$ and to the left of the line $x=x_2$. By first applying the translation $(x,y)\mapsto(x,y-a)$ to Ω_2' and then multiplying by $\begin{pmatrix}1&0\\1&1\end{pmatrix}\in SL_2(\mathbb{Z})$, we obtain a new domain Ω_2 (which we interpret as the empty set if $x_2=0$). Let Ω_3' denote the portion of Ω above the line $x+y=a$ and to the right of the line $x=x_3$. By first applying the translation $(x,y)\mapsto(x-a,y)$ and then multiplying by $\begin{pmatrix}1&1\\0&1\end{pmatrix}\in SL_2(\mathbb{Z})$, we obtain a new domain Ω_3 (which we interpret as the empty set if $x_3=a$). See Figure 4.1 for an example of this decomposition. Observe that each X_{Ω_i} is a rational concave toric domain. We now define

$$w(\Omega)=w(\Omega_1)\cup w(\Omega_2)\cup w(\Omega_3). \tag{4.3.1}$$

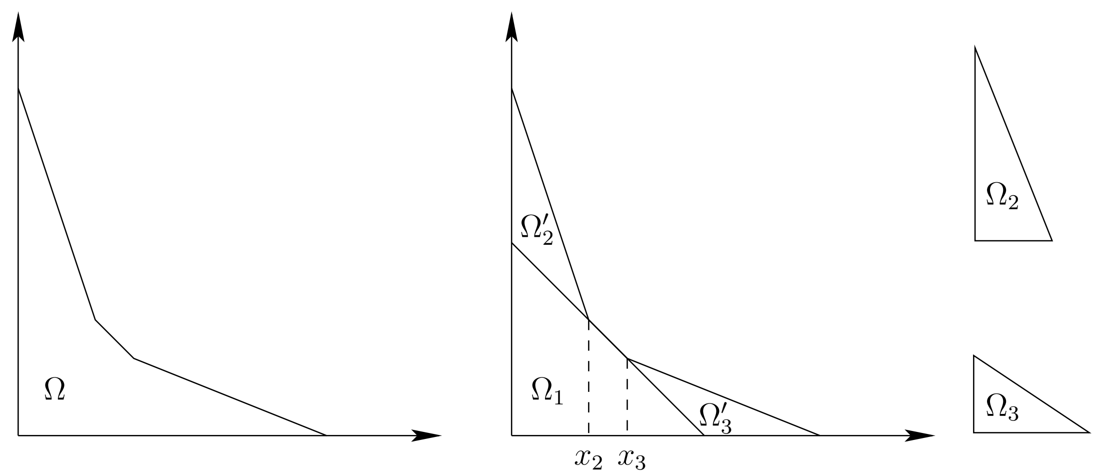

Figure 4.1: The inductive step in the decomposition of a concave toric domain

Here the symbol '$\cup$' indicates "union with repetitions", and we interpret $w(\Omega_i) = \emptyset$ if $\Omega_i = \emptyset$. See §4.4 below for examples of weight expansions.

When Ω is a rational triangle, the weight expansion is determined by the continued fraction expansion of the slope of the diagonal, and in particular $w(\Omega)$ is finite, see [18, §2]. If the upper boundary of Ω has more than one edge, then the upper boundary of each Ω_i will have fewer edges than that of Ω, so by induction $w(\Omega)$ is still finite.

Theorem 4.3.1. The ECH capacities of a rational concave toric domain X_Ω with weight expansion $(a_1, \ldots, a_n)$ are given by

$$c_k(X_\Omega) = c_k\left(\coprod_{i=1}^n B(a_i)\right).$$

Remark 4.3.1. It follows from the Disjoint Union property of ECH capacities, together with the formulas (4.1.2) and (4.1.3) for the ECH capacities of a ball, that

$$c_k\left(\coprod_{i=1}^n B(a_i)\right) = \max\left\{\sum_{i=1}^n a_i d_i \;\middle|\; \sum_{i=1}^n \frac{d_i^2 + d_i}{2} \le k\right\}, \tag{4.3.2}$$

where $d_1, \ldots, d_n$ are nonnegative integers. To compute the maximum on the right hand side of (4.3.2), if we order the weight expansion so that $a_1 \ge \cdots \ge a_n$, then we can assume without loss of generality that $d_i = 0$ whenever $i > k$.

Remark 4.3.2. One can extend Theorem 4.3.1 to concave toric domains which are not rational; in this case the weight expansion is defined inductively as before, but is now an infinite sequence. To prove this extension of Theorem 4.3.1, one can approximate an arbitrary concave toric domain X_Ω by rational concave toric domains whose weight expansion is the portion of the weight expansion of X_Ω obtained from the first n steps, and then use the continuity of the ECH capacities in Lemma 5.2.1 below.

One inequality in Theorem 4.3.1 has a quick proof:

Lemma 4.3.2. If X_Ω is a rational concave toric domain with weight expansion $(a_1, \ldots, a_n)$, then

$$c_k(X_\Omega) \geq c_k\left(\coprod_{i=1}^{n} B(a_i)\right). \tag{4.3.3}$$

To prove Lemma 4.3.2, we will use the following version of the "Traynor trick". Call two domains Ω_1 and Ω_2 in the first quadrant *affine equivalent* if one can be obtained from the other by the action of $SL_2(\mathbb{Z})$ and translation. Let $\triangle(a)$ denote the open triangle with vertices $(0,0)$, $(a,0)$, and $(0,a)$.

Lemma 4.3.3. If T is an open triangle in the first quadrant which is affine equivalent to $\triangle(a)$, then there exists a symplectic embedding $\operatorname{int}(B(a)) \to X_T$.

Proof. It follows from [26, Prop. 5.2] that there exists a symplectic embedding

$$\operatorname{int}(B(a)) \to X_{\triangle(a)}.$$

On the other hand, if Ω_1 and Ω_2 are affine equivalent and do not contain any points on the axes, then X_{Ω_1} is symplectomorphic to X_{Ω_2}. Thus $X_{\triangle(a)}$ is symplectomorphic to X_T and we are done. □

Proof of Lemma 4.3.2. It follows from the definition of the weight expansion that Ω has a decomposition into open triangles $T_1, \ldots, T_n$ such that T_i is affine equivalent to $\triangle(a_i)$ for each i. By Lemma 4.3.3, for each i there is a symplectic embedding $\operatorname{int}(B(a_i)) \to X_{T_i}$. Hence there is a symplectic embedding

$$\coprod_{i=1}^{n} \operatorname{int}(B(a_i)) \to X_\Omega.$$

It then follows from the Monotonicity property of ECH capacities that (4.3.3) holds. □

4.4 Examples and first applications

We now give some examples of how Theorem 4.3.1 can be used to prove that certain symplectic embeddings are optimal.

The following lemma will be helpful. If ℓ is a nonnegative integer, define $w_\ell(\Omega) \subset w(\Omega)$ to be the list of positive real numbers obtained from the first ℓ steps in the inductive construction of the weight expansion. That is, $w_0(\Omega) = \emptyset$ and

$$w_\ell(\Omega) = w(\Omega_1) \cup w_{\ell-1}(\Omega_2) \cup w_{\ell-1}(\Omega_3)$$

for $\ell > 0$.

Lemma 4.4.1. If $w_\ell(\Omega) = (a_1, \ldots, a_m)$, then for any $k \le \ell$,

$$c_k(X_\Omega) = c_k\left(\coprod_{i=1}^{m} B(a_i)\right).$$

Proof. Let $(a_1, \ldots, a_n)$ be the weight expansion for Ω. By Theorem 4.3.1, it is enough to prove that

$$c_k\left(\coprod_{i=1}^{n} B(a_i)\right) = c_k\left(\coprod_{i=1}^{m} B(a_i)\right). \tag{4.4.1}$$

By Remark 4.3.1, the left hand side of (4.4.1) is determined by the k largest numbers in $w(\Omega)$, and the right hand side of (4.4.1) is determined by the k largest numbers in $w_\ell(\Omega)$. It follows from the definition of the weight expansion and induction that the k largest numbers in $w(\Omega)$ are a subset of $w_k(\Omega)$; and the latter is a subset of $w_\ell(\Omega)$ since $k \le \ell$. Thus the two sides of (4.4.1) are equal. □

We now have the following corollary of Theorem 4.3.1.

Corollary 4.4.2. If X_Ω is a rational concave toric domain, let a be the largest real number such that $B(a) \subset X_\Omega$. Then the inclusion $B(a) \subset X_\Omega$ is optimal, so the Gromov width of X_Ω equals a.

Proof. Note that a is just the largest real number such that $\triangle(a) \subset \Omega$. It follows from Lemma 4.4.1 with $\ell = 1$ that $c_1(X_\Omega) = a$. Since $c_1(B(a)) = a$, we are done by Remark 4.1.2. □

Here is a simple example of obstructions to symplectic embeddings in which X_Ω is the domain rather than the target:

Example 4.4.1. Let $a \in (0, 1)$, and let Ω be the quadrilateral with vertices $(0, 0)$, $(1, 0)$, $(a, 1 - a)$ and $(0, 1 + a)$. Then the inclusion $X_\Omega \subset B(1 + a)$ is optimal.

Proof. The weight expansion is $w(\Omega) = (1, a)$. It then follows from equation (4.3.2) that $c_2(X_\Omega) = 1 + a$. Since $c_2(B(1 + a)) = 1 + a$, the claim follows from Remark 4.1.2. □

Another interesting example is the (nondisjoint) union of a ball and a cylinder. Given $0 < a < b$, define $Z(a, b)$ to be the union of the ball $B(b)$ with the cylinder

$$Z(a) = P(\infty, a).$$

That is, $Z(a, b) = X_\Omega$ where Ω is bounded by the axes, the line segment from $(0, b)$ to $(b - a, a)$, and the horizontal ray extending to the right from $(b - a, a)$.

Proposition 4.4.3. The ECH capacities of the union of a ball and a cylinder are given by

$$c_k(Z(a,b)) = \max\left\{db + a\left(k - \frac{d(d+1)}{2}\right) \;\middle|\; d(d+1) \leq 2k\right\} \tag{4.4.2}$$

where d is a nonnegative integer.

Proof. Recall from [8, §4.2] that for any symplectic four-manifold (X,ω), we have

$$c_k(X,\omega) = \sup\left\{c_k\left(X_-, \omega|_{X_-}\right)\right\} \tag{4.4.3}$$

where the supremum is over certain compact subsets $X_- \subset \mathrm{int}(X)$ (namely those for which $(X_-, \omega|_{X_-})$ is a four-dimensional "Liouville domain" in the sense of [8, §1]). It follows immediately that ECH capacities have the following "exhaustion property": if $\{X_i\}_{i\geq 1}$ is a sequence of open subsets of X with $X_i \subset X_{i+1}$ and $\bigcup_{i=1}^{\infty} X_i = \mathrm{int}(X)$, then

$$c_k(X,\omega) = \lim_{i\to\infty} c_k\left(X_i, \omega|_{X_i}\right). \tag{4.4.4}$$

To apply this in the present situation, given a positive integer i, let Ω_i be the quadrilateral with vertices $(0,0)$, $(0,b)$, $(b-a,a)$, and $(b+ia,0)$. Then the interiors of the domains X_{Ω_i} exhaust the interior of $Z(a,b)$. Also, X_{Ω_i} has the same ECH capacities as its interior; this follows for example from (4.4.3). It then follows from the exhaustion property (4.4.4) that

$$c_k(Z(a,b)) = \lim_{i\to\infty} c_k(X_{\Omega_i}). \tag{4.4.5}$$

Assume that $i \geq k$. We now compute $c_k(X_{\Omega_i})$ using Theorem 4.3.1. The weight expansion of Ω_i is

$$w(\Omega_i) = (b, \underbrace{a, \ldots, a}_{i \text{ times}}). \tag{4.4.6}$$

Since $i \geq k$, to compute the maximum in (4.3.2), we can assume that each a weight in (4.4.6) is multiplied by 0 or 1, and the b weight in (4.4.6) is multiplied by $(d^2+d)/2$ for some nonnegative integer d. It then follows that $c_k(X_{\Omega_i})$ equals the right hand side of (4.4.2). It now follows from (4.4.5) that (4.4.2) holds. □

It is interesting to ask when the ellipsoid $E(a,b)$ symplectically embeds into $Z(c,d)$. By scaling, it is equivalent to ask, given $a,b \geq 1$, for which $\lambda > 0$ there exists a symplectic embedding $E(a,1) \to Z(\lambda, \lambda b)$. Of course this trivially holds if λ is sufficiently large that $E(a,1)$ is a subset of $Z(\lambda,\lambda b)$. In some cases this sufficient condition is also necessary:

Corollary 4.4.4. Suppose that (i) $a \in \{1,2\}$ and $b \geq 1$, or (ii) a is a positive integer and $1 \leq b \leq 2$. Then there exists a symplectic embedding $E(a,1) \to Z(\lambda,\lambda b)$ if and only if $E(a,1) \subset Z(\lambda,\lambda b)$.

Proof. We first compute that $E(a,1) \subset Z(\lambda, \lambda b)$ if and only if

$$\lambda \geq \frac{a}{a+b-1}. \tag{4.4.7}$$

Assuming (i) or (ii), we need to show that if there exists a symplectic embedding $E(a,1) \to Z(\lambda, \lambda b)$, then the inequality (4.4.7) holds. By the Monotonicity and Conformality properties of ECH capacities, it will suffice to show that

$$c_a(E(a,1)) = a, \tag{4.4.8}$$
$$c_a(Z(1,b)) = a+b-1. \tag{4.4.9}$$

Now (4.4.8) holds for any positive integer a by the Ellipsoid property. And in both cases (i) and (ii), equation (4.4.9) follows from Proposition 4.4.3, because the maximum in (4.4.2) is realized by $d = 1$. □

Remark 4.4.2. There are many cases in which an ellipsoid $E(a,1)$ symplectically embeds into $Z(\lambda, \lambda b)$ although $E(a,1)$ is not a subset of $Z(\lambda, \lambda b)$. For example, an ellipsoid $E(a,1)$ may embed into a ball $B(c)$ of slightly greater volume, and this is always possible when $a \geq (17/6)^2$, see [19]; if we set $c = \lambda b$, then the ellipsoid is not a subset of $Z(\lambda, \lambda b)$ if we choose b sufficiently large. Moreover, the "symplectic folding" method from [22] can be used to construct examples of symplectic embeddings $E(a,1) \to Z(\lambda, \lambda b)$ where $E(a,1) \not\subset Z(\lambda, \lambda b)$ and also $\mathrm{vol}(E(a,1)) > \mathrm{vol}(B(\lambda b))$, so that $E(a,1)$ does not symplectically embed into the ball $B(\lambda b)$ alone.

Corollary 4.4.4 also has a generalization to symplectic embeddings of an ellipsoid into the union of an ellipsoid and a cylinder, see §7.1.

4.5 Application to ball packings

As a more involved application, we obtain a sharp obstruction to ball packings of the union of certain unions of a cylinder and an ellipsoid. Given positive real numbers a, b and c with $c > a$, define

$$Z(a,b,c) = Z(a) \cup E(b,c).$$

Theorem 4.5.1. Let b, c and $w_1 \geq w_2 \geq \cdots \geq w_n > 0$ be positive real numbers. Assume that $c > 1$ and $b \leq \frac{c}{c-1}$. Then there exists a symplectic embedding

$$\coprod_{i=1}^{n} \mathrm{int}(B(w_i)) \to Z(\lambda, \lambda b, \lambda c)$$

if and only if

$$\lambda \geq \max\{w_1/c, \lambda_1, \ldots, \lambda_n\},$$

where we define

$$\lambda_k = \frac{\sum_{i=1}^k w_i}{k + \frac{b(c-1)}{c}}. \tag{4.5.1}$$

For example, Theorem 4.5.1 gives a sharp obstruction to embedding a disjoint union of balls into the union of a ball and a cylinder, $Z(a,b) = Z(a,b,b)$, as long as $b \leq 2a$.

The outline of the proof of Theorem 4.5.1 is as follows. In §7.2, we will give a symplectic embedding construction to prove:

Proposition 4.5.2. Let b, c and $w_1 \geq w_2 \geq \cdots \geq w_n > 0$ be positive real numbers. Assume that $c > 1$. Define λ_k by (4.5.1). If

$$\lambda \geq \max\{w_1/c, \lambda_1, \ldots, \lambda_n\},$$

then there exists a symplectic embedding

$$\coprod_{i=1}^n \operatorname{int}(B(w_i)) \to Z(\lambda, \lambda b, \lambda c). \tag{4.5.2}$$

This implies the sufficient condition for ball packings in Theorem 4.5.1. We will then use ECH capacities to prove the necessary condition for ball packings in Theorem 4.5.1.

Remark 4.5.1. Unlike Theorem 4.5.1, Proposition 4.5.2 still holds when $b > \frac{c}{c-1}$, but in this case we generally do not know whether better symplectic embeddings are possible. For example, Proposition 4.5.2 implies that one can symplectically embed three equal balls $\operatorname{int}(B(a))$ into $Z(1,3)$ whenever $a \leq 5/3$. However ECH capacities only tell us that if such an embedding exists then $a \leq 2$.

4.6 ECH capacities and lattice points

We now give a different formula for the ECH capacities of a concave toric domain, which is not assumed to be rational. This formula requires the following definitions.

Definition 4.6.1. A *concave integral path* is a polygonal path in the plane, whose vertices are at lattice points, and which is the graph of a convex piecewise linear function $F : [0,B] \to [0,A]$ for some nonnegative integers A, B.

Definition 4.6.2. If Λ is a concave integral path, define $\mathcal{L}(\Lambda)$ to be the number of lattice points in the region bounded by Λ, the line segment from $(0,0)$ to $(0,B)$, and the line segment from $(0,0)$ to $(A,0)$, not including lattice points on Λ itself.

Definition 4.6.3. If X_Ω is the concave toric domain determined by $f : [0,b] \to [a,0]$, and if Λ is a concave integral path, define the Ω*-length* of Λ, denoted by $\ell_\Omega(\Lambda)$, as follows. For each edge e of Λ, let v_e denote the vector determined by e, namely the difference between the right and left endpoints. Let p_e be a point on the graph of f such that the graph of f is contained in the closed half-plane above the line through p_e parallel to e. Then

$$\ell_\Omega(\Lambda) = \sum_{e\in \mathrm{Edges}(\Lambda)} v_e \times p_e. \tag{4.6.1}$$

Here $\times$ denotes the cross product. Note that p_e fails to be unique only when the graph of f contains an edge parallel to e, in which case $v_e \times p_e$ does not depend on the choice of p_e.

Theorem 4.6.1. If X_Ω is any concave toric domain, then its ECH capacities are given by

$$c_k(X_\Omega) = \max\{\ell_\Omega(\Lambda) \mid \mathcal{L}(\Lambda) = k\}. \tag{4.6.2}$$

Here the maximum is over concave integral paths Λ.

Remark 4.6.4. It is interesting to compare Theorem 4.6.1 with the formula for the ECH capacities of *convex* toric domains in [10, Thm. 4.14], in which one *minimizes* a length function over convex paths enclosing a certain number of lattice points.

Example 4.6.5. Let us check that Theorem 4.6.1 correctly recovers $c_k(X_\Omega)$ when Ω is the triangle with vertices $(0,0)$, $(a,0)$ and $(0,b)$, so that $X_\Omega = E(a,b)$.

An equivalent statement of the Ellipsoid property is that $c_k(E(a,b)) = L_k$ where L_k is the smallest nonnegative real number such that triangle bounded by the axes and the line $bx + ay = L_k$ encloses at least $k+1$ lattice points. Call this triangle T_k, and call its upper edge D_k.

To see that L_k agrees with the right hand side of (4.6.2), suppose first that a/b is irrational. There is then a unique lattice point (x_k, y_k) on D_k. We need to show that

$$\max\{\ell_\Omega(\Lambda) \mid \mathcal{L}(\Lambda) = k\} = bx_k + ay_k. \tag{4.6.3}$$

If Λ is a concave integral path, there is a unique vertex $(x,y) \in \Lambda$ such that Λ is contained in the closed half-plane above the line through (x,y) with slope $-b/a$. Then $p_e = (0,b)$ for all edges to the left of (x,y), and $p_e = (a,0)$ for all edges to the right of (x,y). Therefore

$$\ell_\Omega(\Lambda) = bx + ay.$$

If $\mathcal{L}(\Lambda) \le k$, then we must have $bx + ay \le bx_k + ay_k$, since otherwise every lattice point in T_k would be counted by $\mathcal{L}(\Lambda)$. Thus the left hand side of (4.6.3) is less than or equal to the right hand side. To prove the reverse inequality, observe that if Λ is the minimal concave integral path which contains the point (x_k, y_k) and is contained in the closed half-plane above the line D_k, then $(x,y) = (x_k, y_k)$ and $\mathcal{L}(\Lambda) = k$.

Suppose now that a/b is rational. Then D_k may contain more than one lattice point. If Λ is a concave integral path, then there is a unique pair of (possibly equal) vertices $(x,y), (x',y') \in \Lambda$ with $x \le x'$ such that line segment from (x,y) to (x',y') is contained in Λ, and the rest of Λ is strictly above the line through (x, y) with slope $-b/a$. Now if p is any point on the upper edge of Ω, then we have

$$\ell_\Omega(\Lambda) = bx + (x' - x, y' - y) \times p + ay'.$$

We can choose $p = (a, 0)$ for convenience, and this gives

$$\ell_\Omega(\Lambda) = bx + ay.$$

The rest of the argument in this case is similar to the previous case.

One can also deduce the case when a/b is rational from the case when a/b is irrational by a continuity argument using Lemma 5.2.2 below.

4.7 The rest of the paper

Theorems 4.3.1 and 4.6.1, which compute the ECH capacities of concave toric domains, are proved in §5 and §6. The generalization of Corollary 4.4.4 to symplectic embeddings of an ellipsoid into the union of an ellipsoid and a cylinder is given in §7.1. Proposition 4.5.2 and Theorem 4.5.1 on ball packings of the union of an ellipsoid and a cylinder are proved in §7.2 and §7.3.

Acknowledgments. It is a pleasure to thank Daniel Irvine and Felix Schlenk for many helpful discussions.

Chapter 5

The lower bound on the capacities

In this section we use combinatorial arguments to prove half of Theorem 4.6.1, namely:

Lemma 5.0.1. If X_Ω is any concave toric domain, then

$$c_k(X_\Omega) \geq \max\{\ell_\Omega(\Lambda) \mid \mathcal{L}(\Lambda) = k\}. \tag{5.0.1}$$

Here the maximum is over concave integral paths Λ.

5.1 The lower bound in the rational case

The following lemma, together with Lemma 4.3.2, implies Lemma 5.0.1 in the rational case.

Lemma 5.1.1. Let X_Ω be a rational concave toric domain with weight expansion $(a_1, \ldots, a_n)$. Then

$$c_k\left(\coprod_{i=1}^n B(a_i)\right) \geq \max\{\ell_\Omega(\Lambda) \mid \mathcal{L}(\Lambda) = k\}. \tag{5.1.1}$$

Proof. The proof has four steps.

Step 1: Setup. We use induction on n. If $n = 1$, then X_Ω is a ball and we know from Example 4.6.5 that both sides of (5.1.1) are equal. If $n > 1$, let Ω_1, Ω_2, and Ω_3 be as in the definition of the weight expansion in §4.3. By induction, we can assume that the lemma is true for Ω_1, Ω_2, and Ω_3.

Let Λ be a concave integral path with $\mathcal{L}(\Lambda) = k$. We need to show that

$$c_k\left(\coprod_{i=1}^n B(a_i)\right) \geq \ell_\Omega(\Lambda). \tag{5.1.2}$$

To prove this, let W_i denote the disjoint union of the balls given by the weight expansion of Ω_i for $i = 1, 2, 3$. By the definition of the weight expansion we have

$$\coprod_{i=1}^n B(a_i) = \coprod_{i=1}^3 W_i. \tag{5.1.3}$$

In Step 2 we will define concave integral paths Λ_i for $i = 1, 2, 3$, and we write $k_i = \mathcal{L}(\Lambda_i)$. By (5.1.3) and the Disjoint Union property of ECH capacities, we know that

$$c_{k_1+k_2+k_3}\left(\coprod_{i=1}^{n} B(a_i)\right) \geq \sum_{i=1}^{3} c_{k_i}(W_i).$$

By the inductive hypothesis we know that

$$c_{k_i}(W_i) \geq \ell_{\Omega_i}(\Lambda_i).$$

In Steps 3 and 4 we will further show that

$$k_1 + k_2 + k_3 = k \tag{5.1.4}$$

and

$$\sum_{i=1}^{3} \ell_{\Omega_i}(\Lambda_i) = \ell_{\Omega}(\Lambda). \tag{5.1.5}$$

The above four equations and inequalities then imply (5.1.2).

Step 2: Definition of Λ_i*.* The paths Λ_i are obtained from Λ in the same way that the domains Ω_i are obtained from Ω. We now make this explicit in order to fix notation. Let Λ_1 be the maximal line segment with slope -1 from the y axis to the x axis such that Λ is contained in the closed half-space above the line extending Λ_1. Let $(0, A)$ and $(A, 0)$ denote the endpoints of Λ_1. Let Λ_2' denote the portion of Λ to the left of $\Lambda_1 \cap \Lambda$, and let Λ_3' denote the portion of Λ to the right of $\Lambda_1 \cap \Lambda$. Let $T_2 : \mathbb{R}^2 \to \mathbb{R}^2$ be the map obtained by first translating down by A and then multiplying by $\begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix} \in SL_2\mathbb{Z}$. Then $\Lambda_2 = T_2(\Lambda_2')$. Similarly, $\Lambda_3 = T_3(\Lambda_3')$, where $T_3 : \mathbb{R}^2 \to \mathbb{R}^2$ is the map obtained by first translating to the left by A and then multiplying by $\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \in SL_2\mathbb{Z}$.

Step 3: Proof of equation (5.1.4). Since T_2 preserves the lattice, $\mathcal{L}(\Lambda_2)$ is the number of lattice points counted by $\mathcal{L}(\Lambda)$ that are on or above Λ_1 and below Λ_2'. Likewise, $\mathcal{L}(\Lambda_3)$ is the number of lattice points counted by $\mathcal{L}(\Lambda)$ that are on or above Λ_1 and below Λ_3'. The remaining lattice points counted by $\mathcal{L}(\Lambda)$ are those that are below Λ_1, which are exactly the lattice points counted by $\mathcal{L}(\Lambda_1)$.

Step 4: Proof of equation (5.1.5). By construction, there is an injection

$$\phi : \text{Edges}(\Lambda) \to \coprod_{i=1}^{3} \text{Edges}(\Lambda_i).$$

The complement of the image of this injection consists of those edges of Λ_1 that are to the left or to the right of $\Lambda_1 \cap \Lambda$. Denote these two sets of edges by $\text{Left}(\Lambda_1)$ and $\text{Right}(\Lambda_1)$ respectively. We tautologically have

$$\sum_{e \in \phi^{-1}(\text{Edges}(\Lambda_1))} v_e \times p_e = \left(\sum_{\hat{e} \in \text{Edges}(\Lambda_1)} - \sum_{\hat{e} \in \text{Left}(\Lambda_1)} - \sum_{\hat{e} \in \text{Right}(\Lambda_1)} \right) v_{\hat{e}} \times p_{\hat{e}}. \tag{5.1.6}$$

Here if $\hat{e}$ is an edge of Λ_i, then $p_{\hat{e}}$ denotes the (not necessarily unique) point on the upper edge of Ω_i that appears in the formula (4.6.1) for $\ell_{\Omega_i}(\Lambda_i)$. To prove equation (5.1.5), it is enough to show in addition to (5.1.6) that

$$\sum_{e\in\phi^{-1}(\mathrm{Edges}(\Lambda_2))} v_e \times p_e = \sum_{\hat{e}\in\mathrm{Edges}(\Lambda_2)} v_{\hat{e}} \times p_{\hat{e}} + \sum_{\hat{e}\in\mathrm{Left}(\Lambda_1)} v_{\hat{e}} \times p_{\hat{e}} \tag{5.1.7}$$

and

$$\sum_{e\in\phi^{-1}(\mathrm{Edges}(\Lambda_3))} v_e \times p_e = \sum_{\hat{e}\in\mathrm{Edges}(\Lambda_3)} v_{\hat{e}} \times p_{\hat{e}} + \sum_{\hat{e}\in\mathrm{Right}(\Lambda_1)} v_{\hat{e}} \times p_{\hat{e}}. \tag{5.1.8}$$

We will just prove equation (5.1.7), as the proof of (5.1.8) is analogous. Let $e \in \phi^{-1}(\mathrm{Edges}(\Lambda_2))$ and let $\hat{e} = \phi(e)$. We then have

$$v_{\hat{e}} = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix} v_e$$

and

$$p_{\hat{e}} = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix} (p_e - (0,a))$$

where a is as in the definition of the weight expansion of Ω in §4.3. Consequently,

$$v_e \times p_e = v_{\hat{e}} \times p_{\hat{e}} + v_e \times (0,a).$$

Summing over all $e \in \phi^{-1}(\mathrm{Edges}(\Lambda_2))$ gives

$$\sum_{e\in\phi^{-1}(\mathrm{Edges}(\Lambda_2))} v_e \times p_e = \sum_{\hat{e}\in\mathrm{Edges}(\Lambda_2)} v_{\hat{e}} \times p_{\hat{e}} + \sum_{e\in\phi^{-1}(\mathrm{Edges}(\Lambda_2))} v_e \times (0,a). \tag{5.1.9}$$

But the rightmost sum in (5.1.9) agrees with the rightmost sum in (5.1.7), because for $\hat{e} \in \Lambda_1$ one can take $p_{\hat{e}} = (0,a)$, and the total horizontal displacement of the edges in $\phi^{-1}(\mathrm{Edges}(\Lambda_2))$ is the same as the total horizontal displacement of the edges in $\mathrm{Left}(\Lambda_1)$. □

5.2 Continuity

Having proved the lower bound (5.0.1) for rational concave toric domains, we now use a continuity argument to extend this bound to arbitrary concave toric domains.

Recall that the Hausdorff metric on compact subsets of $\mathbb{R}^2$ is defined by

$$d(\Omega_1,\Omega_2) = \max_{p_1\in\Omega_1}\min_{p_2\in\Omega_2} d(p_1,p_2) + \max_{p_2\in\Omega_2}\min_{p_1\in\Omega_1} d(p_2,p_1).$$

Lemma 5.2.1. If k is fixed, then $c_k(X_\Omega)$ is a continuous function of Ω with respect to the Hausdorff metric.

Proof. Fix Ω, and given $r > 0$, consider the scaling $r\Omega = \{(rx, ry) \mid (x, y) \in \Omega\}$. Observe that $X_{r\Omega}$ is symplectomorphic to X_Ω with the symplectic form multiplied by r. It then follows from the Conformality property of ECH capacities that $c_k(X_{r\Omega}) = rc_k(X_\Omega)$. If $\{\Omega_i\}_{i\geq 1}$ is a sequence converging to Ω in the Hausdorff metric, then there is a sequence of positive real numbers $\{r_i\}_{i\geq 1}$ converging to 1 such that

$$r_i^{-1}\Omega \subset \Omega_i \subset r_i\Omega.$$

By the Monotonicity property of ECH capacities, we have

$$r_i^{-1}c_k(X_\Omega) \leq c_k(X_{\Omega_i}) \leq r_i c_k(X_\Omega).$$

It follows that $\lim_{i\to\infty} c_k(X_{\Omega_i}) = c_k(X_\Omega)$. □

Lemma 5.2.2. If k is fixed, then $\max\{\ell_\Omega(\Lambda) \mid \mathcal{L}(\Lambda) = k\}$ is a continuous function of Ω with respect to the Hausdorff metric.

Proof. For k fixed, there are only finitely many concave integral paths Λ with $\mathcal{L}(\Lambda) = k$. Consequently, it is enough to show that if Λ is a fixed concave integral path, then $\ell_\Omega(\Lambda)$ is a continuous function of Ω. By (4.6.1), it is now enough to show that if e is an edge of Λ, then $v_e \times p_e(\Omega)$ is a continuous function of Ω. In fact there is a constant $c > 0$ depending only on v_e such that

$$|v_e \times p_e(\Omega) - v_e \times p_e(\Omega')| \leq cd(\Omega, \Omega').$$

To see this, suppose that $v_e \times p_e(\Omega) < v_e \times p_e(\Omega')$. Write $p_e(\Omega) = (x_0, y_0)$. Every point $(x, y) \in \Omega$ must have $x \leq x_0$ or $y \leq y_0$. The portion of the upper boundary of Ω' with $x \geq x_0$ and $y \geq y_0$ is a path from the line $x = x_0$ to the line $y = y_0$. Let $p' \in \Omega'$ denote the intersection of this path with the line of slope 1 through the point (x_0, y_0). The above path must stay above the triangle bounded by the line $x = x_0$, the line $y = y_0$, and the line through $p_e(\Omega')$ parallel to v_e. It follows that there is a constant c' depending only on v_e such that

$$\min_{p\in\Omega} d(p', p) \geq c' v_e \times (p_e(\Omega') - p_e(\Omega)).$$

□

Proof of Lemma 5.0.1. By Lemmas 4.3.2 and 5.1.1, this holds for rational concave toric domains. The general case now follows from Lemmas 5.2.1 and 5.2.2, since if X_Ω is an arbitrary concave toric domain, then Ω can be approximated in the Hausdorff metric by Ω' such that $X_{\Omega'}$ is a rational concave toric domain. □

Chapter 6

The upper bound on the capacities

To complete the proofs of Theorems 4.3.1 and 4.6.1, we now prove:

Lemma 6.0.3. If X_Ω is any concave toric domain, then

$$c_k(X_\Omega) \leq \max\{\ell_\Omega(\Lambda) \mid \mathcal{L}(\Lambda) = k\}. \tag{6.0.1}$$

Here the maximum is over concave integral paths Λ.

Note that Theorem 4.3.1 follows by combining Lemmas 4.3.2, 5.1.1, and 6.0.3, while Theorem 4.6.1 follows by combining Lemmas 5.0.1 and 6.0.3.

6.1 ECH capacities of star-shaped domains

The proof of Lemma 6.0.3 requires some knowledge of the definition of ECH capacities, which we now briefly review; for full details see [8] or [10]. We will only explain the definition for the special case of smooth star-shaped domains in $\mathbb{R}^4$, since that is what we need here.

Let Y be a three-manifold diffeomorphic to S^3, and let λ be a nondegenerate contact form on Y such that $\operatorname{Ker}(\lambda)$ is the tight contact structure. The *embedded contact homology* $ECH_*(Y,\lambda)$ is the homology of a chain complex $ECC_*(Y,\lambda,J)$ over $\mathbb{Z}/2$ defined as follows. (ECH can also be defined with integer coefficients, see [13, §9], but that is not needed for the definition of ECH capacities.) A generator of the chain complex is a finite set of pairs $\alpha = \{(\alpha_i, m_i)\}$ where the α_i are distinct embedded Reeb orbits, the m_i are positive integers, and $m_i = 1$ whenever α_i is hyperbolic. The chain complex in this case has an absolute $\mathbb{Z}$-grading which is reviewed in §6.3 below; the grading of a generator α is denoted by $I(\alpha) \in \mathbb{Z}$. The chain complex differential counts certain J-holomorphic curves in $\mathbb{R} \times Y$ for an appropriate almost complex structure J; the precise definition of the differential is not needed here. Taubes [25] proved that the embedded contact homology of a contact three-manifold is isomorphic to a version of its Seiberg-Witten Floer cohomology as defined by Kronheimer-Mrowka [15]. For the present case of S^3 with its tight contact structure, this implies that

$$ECH_*(Y,\lambda) = \begin{cases} \mathbb{Z}/2, & * = 0,2,4,\ldots, \\ 0, & \text{otherwise.} \end{cases}$$

We denote the nonzero element of $ECH_{2k}(Y,\lambda)$ by ζ_k.

The *symplectic action* of a chain complex generator $\alpha = \{(\alpha_i, m_i)\}$ is defined by

$$\mathcal{A}(\alpha) = \sum_i m_i \int_{\alpha_i} \lambda.$$

We define $c_k(Y,\lambda)$ to be the smallest $L \in \mathbb{R}$ such that ζ_k has a representative in $ECC_*(Y,\lambda,J)$ which is a sum of chain complex generators each of which has symplectic action less than or equal to L. It follows from [14, Thm. 1.3] that $c_k(Y,\lambda)$ does not depend on J. The numbers $c_k(Y,\lambda)$ are called the *ECH spectrum* of (Y,λ).

If λ is a degenerate contact form on $Y \approx S^3$ giving the tight contact structure, we define

$$c_k(Y,\lambda) = \lim_{n\to\infty} c_k(Y, f_n\lambda) \tag{6.1.1}$$

where $\{f_n\}_{n\geq 1}$ is a sequence of positive functions on Y which converges to 1 in the C^0 topology such that each contact form $f_n\lambda$ is nondegenerate. Lemmas from [8, §3.1] imply that this is well-defined, as explained in [4, §2.5].

Now let $X \subset \mathbb{R}^4$ be a compact star-shaped domain with smooth boundary Y. Then

$$\lambda_{std} = \frac{1}{2}\sum_{i=1}^{2}(x_i dy_i - y_i dx_i)$$

restricts to a contact form on Y, and we define the *ECH capacities* of X by

$$c_k(X) = c_k(Y, \lambda_{std}|_Y). \tag{6.1.2}$$

6.2 The combinatorial chain complex

Let X_Ω be a concave toric domain determined by a convex function $f : [0,a] \to [0,b]$. We assume below that the function f is smooth, $f'(0)$ and $f'(a)$ are irrational, f' is constant near 0 and a, and $f''(x) > 0$ whenever $f'(x)$ is rational. Then ∂X_Ω is smooth. As we will see in §6.3 below, λ_{std} restricts to a degenerate contact form on ∂X_Ω. Similarly to [12], there is a combinatorial model for the ECH chain complex of appropriate nondegenerate perturbations of this contact form, which we denote by $ECC_*^{comb}(\Omega)$ and define as follows.

Definition 6.2.1. A generator of $ECC_*^{comb}(\Omega)$ is a quadruple $\widetilde{\Lambda} = (\Lambda, \rho, m, n)$, where:

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

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

(c) m and n are nonnegative integers.

Here an "edge" of Λ means a segment of Λ of which each endpoint is either an initial or a final endpoint of Λ, or a point at which Λ changes slope.

We define the grading $I^{comb}(\widetilde{\Lambda}) \in \mathbb{Z}$ of the generator $\widetilde{\Lambda} = (\Lambda, \rho, m, n)$ as follows. Let $\Lambda_{m,n}$ denote the path in the plane obtained by concatenating the following three paths:

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

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

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

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

$$I^{comb}(\widetilde{\Lambda}) = 2\mathcal{L}(\Lambda_{m,n}) + h(\widetilde{\Lambda}) \tag{6.2.1}$$

where $h(\widetilde{\Lambda})$ denotes the number of edges of Λ that are labeled 'h'.

We define the action $\mathcal{A}^{comb}(\widetilde{\Lambda}) \in \mathbb{R}$ of the generator $\widetilde{\Lambda} = (\Lambda, \rho, m, n)$ by

$$\mathcal{A}^{comb}(\widetilde{\Lambda}) = \ell_\Omega(\Lambda) + an + bm. \tag{6.2.2}$$

One can also define a combinatorial differential on the chain complex $ECC_*^{comb}(\Omega)$ similarly to [12], which agrees with the ECH differential for appropriate perturbations of the contact form and almost complex structures, but we do not need this here. What we do need is the following:

Lemma 6.2.1. For each $\epsilon > 0$, there exists a contact form λ on ∂X_Ω with the following properties:

(a) λ is nondegenerate.

(b) $\lambda = f\lambda_{std}|_{\partial X_\Omega}$ where $\|f - 1\|_{C^0} < \epsilon$.

(c) There is a bijection between the generators of $ECC(\partial X_\Omega, \lambda)$ with $\mathcal{A} < 1/\epsilon$ and the generators of $ECC^{comb}(\Omega)$ with $\mathcal{A}^{comb} < 1/\epsilon$, such that if α and $\widetilde{\Lambda}$ correspond under this bijection, then

$$I(\alpha) = I^{comb}(\widetilde{\Lambda})$$

and

$$|\mathcal{A}(\alpha) - \mathcal{A}^{comb}(\widetilde{\Lambda})| < \epsilon.$$

Lemma 6.2.1 will be proved in §6.3. We can now deduce:

Lemma 6.2.2. For each nonnegative integer k, there exists a generator $\widetilde{\Lambda}$ of $ECC^{comb}(\Omega)$ such that $I^{comb}(\widetilde{\Lambda}) = 2k$ and $\mathcal{A}^{comb}(\widetilde{\Lambda}) = c_k(X_\Omega)$.

Proof. Fix k. For each positive integer n, let λ_n be a contact form provided by Lemma 6.2.1 for $\epsilon = 1/n$. It follows from (6.1.1) and (6.1.2) that we can choose λ_n so that

$$|c_k(X_\Omega) - c_k(\partial X_\Omega, \lambda_n)| < 1/n.$$

By definition, there exists a generator α_n of $ECC_{2k}(\partial X_\Omega, \lambda_n)$ with $\mathcal{A}(\alpha_n) = c_k(\partial X_\Omega, \lambda_n)$. Assume n is sufficiently large that $c_k(X_\Omega) + 1/n < n$. Then $\mathcal{A}(\alpha_n) < n$, so α_n corresponds to a generator $\widetilde{\Lambda}_n$ of $ECC^{comb}(\Omega)$ under the bijection in Lemma 6.2.1, with

$$I^{comb}(\widetilde{\Lambda}_n) = 2k \tag{6.2.3}$$

and

$$|\mathcal{A}^{comb}(\widetilde{\Lambda}_n) - c_k(X_\Omega)| < 2/n. \tag{6.2.4}$$

It follows from (6.2.1) that there are only finitely many generators $\widetilde{\Lambda}$ of $ECC^{comb}(\Omega)$ with $I^{comb}(\widetilde{\Lambda}) = 2k$. Consequently, there exists such a generator $\widetilde{\Lambda}$ which agrees with infinitely many $\widetilde{\Lambda}_n$. It now follows from (6.2.3) and (6.2.4) that $I^{comb}(\widetilde{\Lambda}) = 2k$ and $\mathcal{A}^{comb}(\widetilde{\Lambda}_n) = c_k(X_\Omega)$ as desired. □

Proof of Lemma 6.0.3. Fix k. By the continuity in Lemmas 5.2.1 and 5.2.2, we can assume that Ω is determined by a function $f : [0, a] \to [0, b]$ satisfying the conditions at the beginning of §6.2, such that in addition

$$|f'(0)|, |1/f'(a)| > k. \tag{6.2.5}$$

By Lemma 6.2.2, we can choose a generator $\widetilde{\Lambda} = (\Lambda, \rho, m, n)$ of $ECC^{comb}(\Omega)$ with $I^{comb}(\widetilde{\Lambda}) = 2k$ and $\mathcal{A}^{comb}(\widetilde{\Lambda}) = c_k(X_\Omega)$. It follows from (6.2.5) that $m = n = 0$; otherwise the region bounded by $\Lambda_{m,n}$ and the axes would include at least $k + 1$ lattice points on the axes not in the translate of Λ, so by (6.2.1) we would have $I^{comb}(\widetilde{\Lambda}) > 2k$, which is a contradiction.

Let $k' = \mathcal{L}(\Lambda)$. Then by (6.2.1) we have $k' \le k$, and by (6.2.2) we have $\ell_\Omega(\Lambda) = c_k(X_\Omega)$. Thus

$$c_k(X_\Omega) \le \max\{\ell_\Omega(\Lambda) \mid \mathcal{L}(\Lambda) = k'\}.$$

To complete the proof of Lemma 6.0.3, one could give a combinatorial proof that the right hand side of (6.0.1) is a nondecreasing function of k. Instead we will take a shortcut: by Lemma 5.0.1 we have

$$\max\{\ell_\Omega(\Lambda) \mid \mathcal{L}(\Lambda) = k'\} \le c_{k'}(X_\Omega),$$

and by (4.1.1) we have

$$c_{k'}(X_\Omega) \le c_k(\Omega).$$

Thus the above three inequalities are equalities. □

6.3 The generators of the ECH chain complex

To complete the computations of ECH capacities, our remaining task is to give the:

Proof of Lemma 6.2.1. The proof has five steps.

Step 1. We first determine the embedded Reeb orbits of the contact form $\lambda_{std}|_{\partial X_\Omega}$ and their symplectic actions. Similarly to [10, §4.3], these are given as follows:

- The circle $\gamma_1 = \{z \in \partial X_\Omega \mid z_2 = 0\}$ is an embedded elliptic Reeb orbit with action $\mathcal{A}(\gamma_1) = a$.
- The circle $\gamma_2 = \{z \in \partial X_\Omega \mid z_1 = 0\}$ is an embedded elliptic Reeb orbit with action $\mathcal{A}(\gamma_2) = b$.
- For each $x \in (0, a)$ such that $f'(x)$ is rational, the torus

$$\{z \in \partial X_\Omega \mid \pi(|z_1|^2, |z_2|^2) = (x, f(x))\}$$

is foliated by a Morse-Bott circle of Reeb orbits. Let v_1 be the smallest positive integer such that $v_2 = f'(x)v_1 \in \mathbb{Z}$, write $v = (v_1, v_2)$, and denote this circle of Reeb orbits by $\mathcal{O}_v$. Then each Reeb orbit in $\mathcal{O}_v$ has symplectic action

$$\mathcal{A} = v \times (x, f(x)).$$

In particular, if $\alpha = \{(\alpha_i, m_i)\}$ is a finite set of embedded Reeb orbits with positive integer multiplicities, then α determines a triple (Λ, m, n) satisying conditions (a) and (c) in Definition 6.2.1. The path Λ is obtained by taking the vector v for each Reeb orbit α_i that is in the Morse-Bott circle $\mathcal{O}_v$, multiplied by the covering multiplicity m_i, and concatenating these vectors in order of increasing slope. The integer m is the multiplicity of γ_2 if it appears in α, and otherwise $m = 0$; likewise n is the multiplicity of γ_1 if it appears in α and otherwise $n = 0$. It follows from the above calculations that

$$\mathcal{A}(\alpha) = \ell_\Omega(\Lambda) + an + bm.$$

Step 2. Given $\epsilon > 0$, we can now perturb $\lambda_{std}|_{\partial X_\Omega}$ to $\lambda = f\lambda_{std}|_{\partial X_\Omega}$ where f is C^0-close to 1, so that each Morse-Bott circle $\mathcal{O}_v$ of embedded Reeb orbits with action less than $1/\epsilon$ becomes two embedded Reeb orbits of approximately the same action, namely an elliptic orbit e_v and a hyperbolic orbit h_v; no other Reeb orbits of action less than $1/\epsilon$ are created; and the Reeb orbits γ_1 and γ_2 are unaffected.

Now the generators of $ECC(\partial X_\Omega, \lambda)$ with $\mathcal{A} < 1/\epsilon$ correspond to generators of $ECC^{comb}(\Omega)$ with $\mathcal{A}^{comb} < 1/\epsilon$. Given a generator $\alpha = \{(\alpha_i, m_i)\}$ of $ECC(\partial X_\Omega, \lambda)$ with $\mathcal{A}(\alpha) < 1/\epsilon$, the corresponding combinatorial generator $\widetilde{\Lambda} = (\Lambda, \rho, m, n)$ is determined as follows. The triple (Λ, m, n) is determined as in Step 1. The labeling ρ is defined as follows. Suppose an edge of Λ corresponds to the vector kv where $v = (v_1, v_2)$ is an irreducible integer vector and k is a positive integer. Then either α contains the elliptic orbit e_v with multiplicity k, or α contains the elliptic orbit e_v with multiplicity $k-1$ and the hyperbolic orbit h_v with multiplicity 1. The labeling of the edge is 'e' in the former case and 'h' in the latter case.

To complete the proof of Lemma 6.2.1, we need to show that $I(\alpha) = I^{comb}(\widetilde{\Lambda})$.

Step 3. Let $\alpha = \{(\alpha_i, m_i)\}$ be a generator of $ECC(\partial X_\Omega, \lambda)$. We now review the definition of the grading $I(\alpha)$ in the present context; for details of the grading in general see [10, §3] or [7, §2]. The formula is

$$I(\alpha) = c_\tau(\alpha) + Q_\tau(\alpha) + CZ_\tau^I(\alpha) \tag{6.3.1}$$

where the individual terms are defined as follows. First, τ is a homotopy class of symplectic trivialization of $\xi = \mathrm{Ker}(\lambda)$ over each of the Reeb orbits α_i. Next, $c_\tau(\alpha)$ is the relative first Chern class, with respect to τ, of ξ restricted to a surface bounded by α. That is, if Σ is a compact oriented surface with boundary and $g : \Sigma \to \partial X_\Omega$ is a smooth map such that $g(\partial\Sigma) = \sum_i m_i \alpha_i$, then $c_\tau(\alpha)$ is the algebraic count of zeroes of a section of $g^*\xi$ which on each boundary circle is nonvanishing and has winding number zero with respect to τ. The relative first Chern class is additive in the sense that

$$c_\tau(\alpha) = \sum_i m_i c_\tau(\alpha_i).$$

Next, $Q_\tau(\alpha)$ is the relative self-intersection number; in the present situation this is given by

$$Q_\tau(\alpha) = \sum_i m_i^2 Q_\tau(\alpha_i) + \sum_{i \neq j} m_i m_j \,\mathrm{link}(\alpha_i, \alpha_j). \tag{6.3.2}$$

Here $Q_\tau(\alpha_i)$ is the linking number of α_i with a pushoff of itself via the trivialization τ, and $\mathrm{link}(\alpha_i, \alpha_j)$ denotes the linking number of α_i and α_j. Finally,

$$CZ_\tau^I(\alpha) = \sum_i \sum_{k=1}^{m_i} CZ_\tau(\alpha_i^k)$$

where $CZ_\tau(\alpha_i^k)$ denotes the Conley-Zehnder index of the k-fold iterate of α_i with respect to the trivialization τ. In particular, if γ is an elliptic orbit such that the linearized Reeb flow around γ with respect to the trivialization τ is conjugate to a rotation by $2\pi\theta$ for $\theta \in \mathbb{R}/\mathbb{Q}$, then

$$CZ_\tau(\gamma^k) = 2\lfloor k\theta \rfloor + 1.$$

Step 4. We now calculate the terms that enter into the grading formula (6.3.1) when α is a generator of $ECC(\partial X_\Omega, \lambda)$ with $\mathcal{A}(\alpha) < 1/\epsilon$.

We first choose a trivialization τ of ξ over each embedded Reeb orbit of action less than $1/\epsilon$. There is a distinguished trivialization τ of ξ over γ_1 determined by the disk in the plane $z_2 = 0$ bounded by γ_1. With respect to this trivialization, the linearized Reeb flow around γ is rotation by $-2\pi/f'(a)$, so that

$$CZ_\tau(\gamma_1^k) = 2\lfloor -k/f'(a) \rfloor + 1. \tag{6.3.3}$$

Likewise, there is a distinguished trivialization τ of ξ over γ_2 determined by the disk in the plane $z_1 = 0$ bounded by γ_2. With respect to this trivialization, we have

$$CZ_\tau(\gamma_2^k) = 2\lfloor -kf'(0)\rfloor + 1. \tag{6.3.4}$$

We also have

$$c_\tau(\gamma_i) = 1, \qquad Q_\tau(\gamma_i) = 0$$

for $i = 1, 2$.

We can choose the trivialization τ over the orbits e_v and h_v coming from the Morse-Bott circles so that the linearized Reeb flow around e_v is a slight negative rotation, and the linearized Reeb flow around h_v does not rotate the eigenspaces of the linearized return map. This implies that

$$CZ_\tau(e_v^k) = -1, \qquad CZ_\tau(h_v) = 0 \tag{6.3.5}$$

whenever k is sufficiently small that e_v^k has action less than $1/\epsilon$. We also have

$$\begin{aligned} c_\tau(e_v) = c_\tau(h_v) &= v_1 - v_2, \\ Q_\tau(e_v) = Q_\tau(h_v) &= \operatorname{link}(e_v, h_v) = -v_1 v_2. \end{aligned}$$

Finally, the linking numbers of pairs of distinct embedded Reeb orbits are given as follows. Below, o_v denotes either e_v or h_v.

$$\begin{aligned} \operatorname{link}(\gamma_1, \gamma_2) &= 1, \\ \operatorname{link}(\gamma_1, o_v) &= -v_2, \\ \operatorname{link}(\gamma_2, o_v) &= v_1, \\ \operatorname{link}(o_v, o_w) &= \min(-v_1 w_2, -v_2 w_1). \end{aligned}$$

Step 5. Let α and $\widetilde{\Lambda}$ be as in Step 2; we compute the grading $I(\alpha)$ in terms of $\widetilde{\Lambda} = (\Lambda, \rho, m, n)$.

As in §6.2, let $(0, B)$ and $(A, 0)$ denote the endpoints of Λ. The Chern class calculations in Step 4 then imply that

$$c_\tau(\alpha) = A + B + m + n. \tag{6.3.6}$$

Next, let $\Lambda'_{m,n}$ be defined like $\Lambda_{m,n}$ in §6.2, but with the first path replaced by a horizontal segement from $(0, B+n)$ to $(m, B+n)$, and with the third path replaced by a vertical segment from $(A+m, n)$ to $(A+m, 0)$. Let $R'_{m,n}$ denote the region bounded by $\Lambda'_{m,n}$ and the axes. We then have

$$Q_\tau(\alpha) = 2\operatorname{Area}(R'_{m,n}).$$

This follows by expanding $Q_\tau(\alpha)$ using (6.3.2) and the formulas in Step 4, and then interpreting the result as the area of $R_{m,n}$ computed by appropriately dissecting it into right triangles and rectangles. Let $\mathcal{L}(\Lambda'_{m,n})$ denote the number of lattice points in $R'_{m,n}$, not including lattice points on the translate of Λ. Let E denote the number of lattice points on Λ. By Pick's formula for the area of a lattice polygon, we have

$$2\operatorname{Area}(R'_{m,n}) = 2\mathcal{L}(\Lambda'_{m,n}) + E - 2m - 2n - A - B - 1.$$

Let $e(\widetilde{\Lambda})$ denote the total multiplicity of all elliptic orbits in α. Observe that

$$E = e(\widetilde{\Lambda}) + h(\widetilde{\Lambda}) + 1.$$

Combining the above three equations, we obtain

$$Q_\tau(\alpha) = 2\mathcal{L}(\Lambda'_{m,n}) + e(\widetilde{\Lambda}) + h(\widetilde{\Lambda}) - 2m - 2n - A - B. \tag{6.3.7}$$

Finally, it follows from (6.3.3) and (6.3.4) that

$$\sum_{k=1}^{n} CZ_\tau(\gamma_1^k) + \sum_{k=1}^{m} CZ_\tau(\gamma_2^k) = 2(\mathcal{L}(\Lambda_{m,n}) - \mathcal{L}(\Lambda'_{m,n})) + m + n.$$

By (6.3.5), the sum of the remaining Conley-Zehnder terms in $CZ_\tau^I(\alpha)$ is $-e(\widetilde{\Lambda})$. Thus

$$CZ_\tau^I(\alpha) = 2(\mathcal{L}(\Lambda_{m,n}) - \mathcal{L}(\Lambda'_{m,n})) + m + n - e(\widetilde{\Lambda}). \tag{6.3.8}$$

Adding equations (6.3.6), (6.3.7), and (6.3.8) gives

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

as desired. □

Chapter 7

The union of an ellipsoid and a cylinder

In this section we study symplectic embeddings into $Z(a,b,c)$, which is the union of the cylinder $Z(a)$ with the ellipsoid $E(b,c)$. In §7.1 we give a generalization of Corollary 4.4.4, and in §7.2 and §7.3 we prove Proposition 4.5.2 and Theorem 4.5.1.

7.1 Optimal ellipsoid embeddings

We now prove the following proposition which asserts that certain inclusions of an ellipsoid into the union of an ellipsoid and a cylinder are optimal. This is a generalization of Corollary 4.4.4, which is the case $b = c$.

Proposition 7.1.1. Let a be a positive integer and let b, c and λ be positive real numbers. Assume $c > 1$, $a \geq b/c$, and at least one of the following two conditions:

(i) $a = \lfloor b/c \rfloor + 1$.

(ii) $b \leq \frac{c}{c-1}$.

Then there exists a symplectic embedding $E(a,1) \to Z(\lambda, \lambda b, \lambda c)$ if and only if $E(a,1) \subset Z(\lambda, \lambda b, \lambda c)$.

Proof. Using $c > 1$ and $a \geq b/c$, we calculate that $E(a,1) \subset Z(\lambda, \lambda b, \lambda c))$ if and only if

$$\lambda \geq \frac{ac}{ac + b(c-1)}. \tag{7.1.1}$$

Consequently, as in the proof of Corollary 4.4.4, Proposition 7.1.1 follows from the Ellipsoid axiom and Lemma 7.1.2 below. □

Lemma 7.1.2. Let a be a positive integer and let b and c be positive real numbers with $c > 1$ and $a \geq b/c$. Assume that (i) or (ii) in Proposition 7.1.1 holds. Then

$$c_a(Z(1,b,c)) = a + \frac{b(c-1)}{c}. \tag{7.1.2}$$

Proof. The proof has three steps.

Step 1. We first prove equation (7.1.2) in case (ii) when $c \geq b$.

Referring back to the definition of the weight expansion in §4.3, we have

$$\begin{aligned} X_{\Omega_1} &= B\left(\frac{b(c-1)}{c} + 1\right), \\ X_{\Omega_2} &= E\left(\frac{b(c-1)}{c}, \frac{(c-b)(c-1)}{c}\right), \\ X_{\Omega_3} &= Z(1). \end{aligned}$$

By Theorem 4.3.1 and the Disjoint Union property of ECH capacities, we have

$$c_a(Z(1,b,c)) = \max_{k_1+k_2+k_3=a} \sum_{i=1}^{3} c_{k_i}(X_{\Omega_i}).$$

Now $c_{k_3}(X_{\Omega_3}) = k_3$. Also, it follows from the Ellipsoid property that $c_k(E(\alpha,\beta)) \leq k\alpha$. Since we are assuming that $\frac{b(c-1)}{c} \leq 1$, we deduce that $c_{k_2}(X_{\Omega_2}) \leq k_2$. Thus the maximum is achieved with $k_2 = 0$. Since $1 < \frac{b(c-1)}{c} + 1 \leq 2$, it follows as in (4.4.9) that the maximum is achieved with $k_1 = 1$. Equation (7.1.2) follows.

Step 2. We now prove equation (7.1.2) in case (ii) when $b \geq c$.

Here, in the inductive definition of the weight expansion, the first $\lfloor b/c \rfloor$ steps yield $\lfloor b/c \rfloor$ copies of the ball $B(c)$. The remaining region is $Z(1, b - c\lfloor b/c \rfloor, c)$. Here, if c divides b, then we regard $Z(1, b - c\lfloor b/c \rfloor, c)$ as $Z(1)$. Thus by Theorem 4.3.1 and the Disjoint Union property,

$$c_a(Z(1,b,c)) = \max_{k_1+k_2=a} \left(c_{k_1}\left(E\left(c, c\lfloor b/c \rfloor\right)\right) + c_{k_2}\left(Z(1, b - c\lfloor b/c \rfloor, c)\right)\right). \tag{7.1.3}$$

Step 1 applies to show that

$$c_{k_2}\left(Z(1, b - c\lfloor b/c \rfloor, c)\right) = k_2 + \frac{(b - c\lfloor b/c \rfloor)(c-1)}{c} \tag{7.1.4}$$

whenever $k_2 \geq 1$. Also, it follows from the Ellipsoid property that

$$c_{k_1}(E(c, c\lfloor b/c \rfloor)) = ck_1 \tag{7.1.5}$$

for $k_1 \leq \lfloor b/c \rfloor$. For larger values of k_1, one has to increase k_1 by at least $\lfloor b/c \rfloor \geq 2$ to obtain any increase in $c_{k_1}(E(c, c\lfloor b/c \rfloor))$, and this increase will always equal c. On the other hand, it follows from $b \geq c$ and (ii) that $c \leq 2$. Hence the maximum is attained for $k_1 \leq \lfloor b/c \rfloor$. Since $a \geq \lfloor b/c \rfloor$ and $c > 1$, the maximum is attained with $k_1 = \lfloor b/c \rfloor$. Adding (7.1.4) and (7.1.5) then proves (7.1.2).

Step 3. We now prove equation (7.1.2) in case (i). As in Step 2, the first a steps of the weight expansion yield $a-1$ copies of the ball $B(c)$, together with the ball $B\left(\frac{(b-c(a-1))(c-1)}{c}+1\right)$. It follows from Lemma 4.4.1 that

$$c_a(Z(1,b,c)) = (a-1)c + \frac{(b-c(a-1))(c-1)}{c} + 1.$$

Simplifying this expression gives equation (7.1.2) again. □

Remark 7.1.1. If $c > 1$ and a is a positive integer with $a \leq b/c$, then

$$c_a(Z(1,b,c)) = ac. \tag{7.1.6}$$

This is because the first a steps in the weight expansion yield a copies of the ball $B(c)$, and we can then apply Lemma 4.4.1.

7.2 Construction of ball packings

Proof of Proposition 4.5.2. The proof has three steps.

Step 1. Choose $k \in \{1,\ldots,n\}$ maximizing λ_k. We claim that $\lambda_k \geq w_i$ for all $i > k$.

To see this, use (4.5.1) to compute that

$$\lambda_k - w_{k+1} = \left(k + \frac{b(c-1)}{c} + 1\right)(\lambda_k - \lambda_{k+1}).$$

Since λ_k is maximal, we deduce that $\lambda_k \geq w_{k+1}$. The rest follows from the fact that $w_1 \geq \cdots \geq w_n$.

Step 2. Let Ω be the region for which $X_\Omega = Z(\lambda, \lambda b, \lambda c)$. That is, Ω is bounded by the axes, the line segment from $(0, \lambda c)$ to $\left(\frac{b}{c}(c-1)\lambda, \lambda\right)$, and the horizontal ray extending to the right from the latter point. By Lemma 4.3.3, it suffices to embed disjoint open triangles $T_1, \ldots, T_n$ into Ω, such that T_ℓ is affine equivalent to $\triangle(w_\ell)$ for each ℓ. If $\ell > k$, then by Step 1 we have $\lambda \geq w_\ell$, so we can simply take T_ℓ to be a translate of $\triangle(w_\ell)$ sufficiently far to the right.

Step 3. For $1 \leq \ell \leq k$, we now define the triangle T_ℓ by starting with the triangle $\triangle(w_\ell)$, multiplying by $\begin{pmatrix} 1 & -(\ell-1) \\ 0 & 1 \end{pmatrix} \in SL_2\mathbb{Z}$, and then translating to the right by $\sum_{i=1}^{\ell-1} w_i$. The vertices of T_ℓ are

$$\left(\sum_{i=1}^{\ell-1} w_i, 0\right), \left(\sum_{i=1}^{\ell} w_i, 0\right), \text{ and } \left(\sum_{i=1}^{\ell-1} w_i - (\ell-1)w_\ell, w_\ell\right).$$

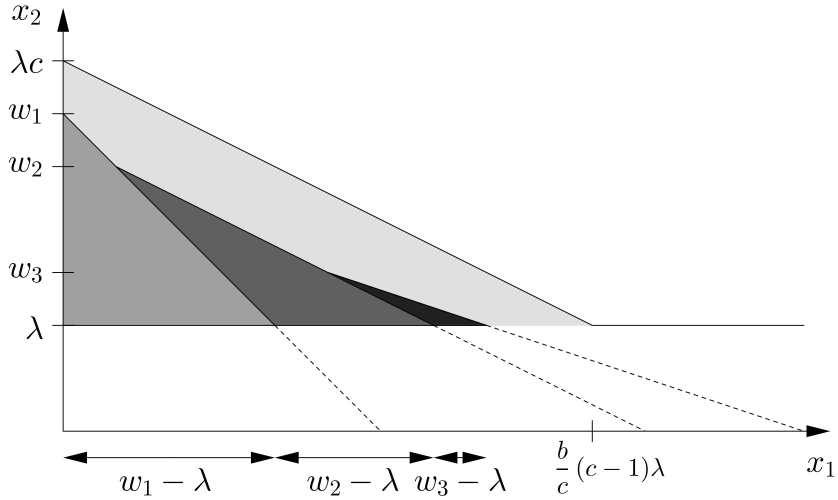

Figure 7.1: The union of the triangles T_ℓ is a subset of Ω

Observe that the right edge of T_ℓ has slope $-1/\ell$; and if $\ell > 1$ then the left edge of T_ℓ is a subset of the right edge of $T_{\ell-1}$. In particular, the triangles $T_1, \ldots, T_k$ are disjoint; and the upper boundary of the union of their closures, call this path Λ, is the graph of a convex function.

To verify that the triangles $T_1, \ldots, T_k$ are contained in Ω, we need to check that the path Λ does not go above the upper boundary of Ω, see Figure 7.1. The initial endpoint of Λ is $(0, w_1)$, which is not above the upper boundary of Ω by our assumption that $\lambda \geq w_1/c$. Next, Λ crosses the horizontal line of height λ at the point $\left(\sum_{\ell=1}^k (w_\ell - \lambda), \lambda\right)$. By convexity, it is enough to check that this point is not to the right of the corner $\left(\frac{b}{c}(c-1)\lambda, \lambda\right)$ of $\partial\Omega$. This holds because

$$\lambda \geq \lambda_k = \frac{\sum_{\ell=1}^k w_\ell}{k + \frac{b}{c}(c-1)}$$

implies that

$$\sum_{\ell=1}^k (w_\ell - \lambda) \leq \frac{b}{c}(c-1)\lambda.$$

□

7.3 The ECH obstruction to ball packings

We now complete the proof of Theorem 4.5.1. By Proposition 4.5.2, it is enough to prove:

Lemma 7.3.1. Under the assumptions of Theorem 4.5.1, if there exists a symplectic embedding

$$\coprod_{i=1}^{n} \operatorname{int}(B(w_i)) \to Z(\lambda, \lambda b, \lambda c),$$

then $\lambda \geq \max\{w_1/c, \lambda_1, \ldots, \lambda_n\}$.

Proof. By the Monotonicity and Conformality properties of ECH capacities, it is enough to show that there is a positive integer k such that

$$c_k\left(\coprod_{i=1}^{n} \operatorname{int}(B(w_i))\right) \geq \max\{w_1/c, \lambda_1, \ldots, \lambda_n\} \cdot c_k(Z(1, b, c)). \tag{7.3.1}$$

By the Disjoint Union axiom, if $1 \leq k \leq n$ then

$$c_k\left(\coprod_{i=1}^{n} \operatorname{int}(B(w_i))\right) \geq \sum_{i=1}^{k} w_i.$$

So to prove (7.3.1), it is enough to show that there exists $k \in \{1, \ldots, n\}$ with

$$\sum_{i=1}^{k} w_i \geq \max\{w_1/c, \lambda_1, \ldots, \lambda_n\} \cdot c_k(Z(1, b, c)). \tag{7.3.2}$$

We will prove this by considering two cases.

Case 1. Assume that $b \leq c$. Then $w_1/c \leq \lambda_1$. Hence

$$\max\{w_1/c, \lambda_1, \ldots, \lambda_n\} = \max\{\lambda_1, \ldots, \lambda_n\}. \tag{7.3.3}$$

We claim now that (7.3.2) holds for $k \in \{1, \ldots, n\}$ maximizing λ_k. To prove this, we need to show that

$$\sum_{i=1}^{k} w_i \geq \lambda_k c_k(Z(1, b, c)).$$

By equation (4.5.1), the above inequality is equivalent to

$$c_k(Z(1, b, c)) \leq k + \frac{b(c-1)}{c}. \tag{7.3.4}$$

Since $b/c \leq 1$, it follows from Lemma 7.1.2 that equality holds in (7.3.4).

Case 2. Assume that $b \geq c$. By Corollary 4.4.2, we have

$$c_1(Z(1, b, c)) = c.$$

Consequently, we can assume without loss of generality that (7.3.3) holds, since otherwise the inequality (7.3.2) holds for $k = 1$. As in Step 1, it is now enough to prove the inequality (7.3.4), where $k \in \{1, \ldots, n\}$ maximizes λ_k.

If $k \geq b/c$, then equality holds in (7.3.4) by Lemma 7.1.2. If $k < b/c$, then the inequality (7.3.4) follows from Remark 7.1.1, since in this case

$$kc < k + \frac{b(c-1)}{c}.$$

□

Bibliography

[1] F. Bourgeois, Y. Eliashberg, H. Hofer, K. Wysocki, and E. Zehnder, *Compactness results in symplectic field theory*, Geom. Topol. **7** (2003), 799–888. MR 2026549 (2004m:53152)

[2] Frédéric Bourgeois, *A Morse-Bott approach to contact homology*, Symplectic and contact topology: interactions and perspectives (Toronto, ON/Montreal, QC, 2001), Fields Inst. Commun., vol. 35, Amer. Math. Soc., Providence, RI, 2003, pp. 55–77. MR 1969267 (2004a:53109)

[3] Vincent Colin, Paolo Ghiggini, and Ko Honda, *Equivalence of Heegaard Floer homology and embedded contact homology via open book decompositions*, Proc. Natl. Acad. Sci. USA **108** (2011), no. 20, 8100–8105. MR 2806645 (2012f:53190)

[4] Daniel Cristofaro-Gardiner and Michael Hutchings, *From one Reeb orbit to two*, 2012, arXiv:1202.4839.

[5] David Frenkel and Dorothee Müller, *Symplectic embeddings of 4-dimensional ellipsoids into cubes*, 2012, arXiv:1210.2266, to appear in J. Symplectic Geom.

[6] Richard Hind and Samuel Lisi, *Symplectic embeddings of polydisks*, 2013, arXiv:1304.3065.

[7] Michael Hutchings, *The embedded contact homology index revisited*, New perspectives and challenges in symplectic field theory, CRM Proc. Lecture Notes, vol. 49, Amer. Math. Soc., Providence, RI, 2009, pp. 263–297. MR 2555941 (2010k:53152)

[8] ———, *Quantitative embedded contact homology*, J. Differential Geom. **88** (2011), no. 2, 231–266. MR 2838266

[9] ———, *Recent progress on symplectic embedding problems in four dimensions*, Proc. Natl. Acad. Sci. USA **108** (2011), no. 20, 8093–8099. MR 2806644 (2012m:53190)

[10] ———, *Lecture notes on embedded contact homology*, 2013, arXiv:1303.5789.

[11] Michael Hutchings and Michael Sullivan, *The periodic Floer homology of a Dehn twist*, Algebr. Geom. Topol. **5** (2005), 301–354. MR 2135555 (2005m:53169)

[12] ———, *Rounding corners of polygons and the embedded contact homology of T^3*, Geom. Topol. **10** (2006), 169–266. MR 2207793 (2006k:53150)

[13] Michael Hutchings and Clifford Henry Taubes, *Gluing pseudoholomorphic curves along branched covered cylinders. II*, J. Symplectic Geom. **7** (2009), no. 1, 29–133. MR 2491716 (2010c:53130)

[14] ———, *Proof of the Arnold chord conjecture in three dimensions 1*, Math. Res. Lett. **18** (2011), no. 2, 295–313. MR 2784673 (2012e:53183)

[15] Peter Kronheimer and Tomasz Mrowka, *Monopoles and three-manifolds*, New Mathematical Monographs, vol. 10, Cambridge University Press, Cambridge, 2007. MR 2388043 (2009f:57049)

[16] Cagatay Kutluhan, Yi-Jen Lee, and Clifford Henry Taubes, *HF=HM I: Heegaard Floer homology and Seiberg-Witten Floer homology*, 2010, arxiv:1007.1979.

[17] Yi-Jen Lee, *Reidemeister torsion in Floer-Novikov theory and counting pseudo-holomorphic tori. I*, J. Symplectic Geom. **3** (2005), no. 2, 221–311. MR 2199540 (2006k:53152)

[18] Dusa McDuff, *The Hofer conjecture on embedding symplectic ellipsoids*, J. Differential Geom. **88** (2011), no. 3, 519–532. MR 2844441 (2012j:53113)

[19] Dusa McDuff and Felix Schlenk, *The embedding capacity of 4-dimensional symplectic ellipsoids*, Ann. of Math. (2) **175** (2012), no. 3, 1191–1282. MR 2912705

[20] Brett Parker, *Holomorphic curves in Lagrangian torus fibrations*, ProQuest LLC, Ann Arbor, MI, 2005, Thesis (Ph.D.)–Stanford University. MR 2707719

[21] Sucharit Sarkar and Jiajun Wang, *An algorithm for computing some Heegaard Floer homologies*, Ann. of Math. (2) **171** (2010), no. 2, 1213–1236. MR 2630063 (2012f:57032)

[22] Felix Schlenk, *Embedding problems in symplectic geometry*, de Gruyter Expositions in Mathematics, vol. 40, Walter de Gruyter GmbH & Co. KG, Berlin, 2005. MR 2147307 (2007c:53125)

[23] Clifford Henry Taubes, *A compendium of pseudoholomorphic beasts in $\mathbb{R} \times (S^1 \times S^2)$*, Geom. Topol. **6** (2002), 657–814 (electronic). MR 1943381 (2005a:32024)

[24] ———, *Pseudoholomorphic punctured spheres in $\mathbb{R}\times(S^1\times S^2)$: properties and existence*, Geom. Topol. **10** (2006), 785–928 (electronic). MR 2240906 (2007f:53120)

[25] ———, *Embedded contact homology and Seiberg-Witten Floer cohomology V*, Geom. Topol. **14** (2010), no. 5, 2961–3000. MR 2746727 (2011m:53182)

[26] Lisa Traynor, *Symplectic packing constructions*, J. Differential Geom. **42** (1995), no. 2, 411–429. MR 1366550 (96k:53046)