Symplectic Embeddings of Toric Domains

by

Daniel Reid Irvine

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Mathematics)
in the University of Michigan
2020

Doctoral Committee:

Professor Daniel Burns, Jr., Chair
Professor Anthony Bloch
Assistant Professor Jun Li
Professor John Schotland
Professor Alejandro Uribe

Daniel Reid Irvine
dirvine@umich.edu
ORCID ID: 0000-0003-2721-5901

Dedication

This work is gratefully dedicated to my family.

Acknowledgments

I would like to extend my deep gratitude to my advisor, Daniel Burns, for his continued support and encouragement during my graduate career. Beyond his mathematical expertise, he has been a steadfast friend.

Second, I would like to thank Richard Hind for his advisement during both my undergraduate and graduate career. He encouraged me to undertake this ambitious project, and his enthusiasm steered me towards the field of symplectic geometry.

Third, I would like to thank Jun Li for his advice and many helpful conversations.

Finally, I would like to thank all the instructors that guided me through my mathematical career. I would not be where I am today if not for them.

Table of Contents

List of Figures

Abstract

This work examines a special class of symplectic manifolds known as toric domains. The main problem under consideration is embedding one toric domain into another in a way that preserves the symplectic structure. The problem of when a symplectic embedding exists is largely open, except for simple manifolds. One important example manifold, which is also a toric domain, is the symplectic polydisc, which is a product of 2-dimensional discs. We say that a polydisc is *stabilized* when it is crossed with several copies of the complex plane (which can be viewed as a disc of infinite radius). The main result of this work is a characterization of when a large class of stabilized polydiscs embed into other stabilized polydiscs. This result incorporates high-dimensional polydisc embeddings. In examining these stabilized polydisc embeddings we consider J*-holomorphic curves*, which are solutions to a PDE. A crucial step in the proof of the embedding result will be constructing J-holomorphic curves with specific boundary conditions. It turns out that the curves constructed for this theorem persist under the process of stabilizing. This makes such curves very useful for obstructing other stabilized embedding problems, which gives an avenue for future research.

The 4-dimensional symplectic embedding problem can be studied quantitatively using *weight vectors*. Heuristically, a weight vector represents the radii of an optimal packing of symplectic balls into a toric domain, and this optimal ball packing will be described algorithmically in this work. Conversely, if one is given a weight vector, one can use this algorithm to generate a toric domain having the given vector as its weight vector. We also discuss how to pass between two classes of toric domains, called concave and convex, and examine the effect that this transformation has on weight vectors.

Weight vectors are useful because they serve as a shorthand for an infinite sequence of symplectic capacities known as *ECH capacities*, but the computation of ECH capacities from a weight vector is difficult. Another main result of this work gives a method of simplifying weight vectors to make the computation of the corresponding ECH capacity sequence easier.

CHAPTER 1

Introduction

Let (X_1, ω_1) and (X_2, ω_2) be symplectic manifolds of the same dimension. A *symplectic embedding* is a smooth embedding $\varphi\colon X_1 \to X_2$ such that $\varphi^*\omega_2 = \omega_1$. Symplectic embeddings will be the principal object of study in this work. Despite the relatively simple definition, symplectic embedding problems are really only understood for the simplest of examples, and even then only in low dimensions.

Definition 1.0.1. The space $\mathbb{C}^n$ with coordinates $z_j = x_j + iy_j$ is a symplectic manifold when equipped with the *standard symplectic form*. This form is

$$\omega_{std} = \sum_{j=1}^{n} dx_j \wedge dy_j. \tag{1.0.1}$$

Notice that this manifold has (real) dimension $2n$. In this work, n will always denote half the dimension of the symplectic manifold under consideration.

Definition 1.0.2. The *symplectic ellipsoid of areas* r_1 *through* r_n is the set

$$E(r_1, ..., r_n) = \left\{ (z_1, \dots, z_n) \in \mathbb{C}^n \,\middle|\, \sum_{j=1}^{n} \frac{\pi \cdot |z_j|^2}{r_j} \leq 1 \right\},$$

with the symplectic form induced from $(\mathbb{C}^n, \omega_{std})$.

Definition 1.0.3. The *symplectic polydisc of areas* r_1 *through* r_n is the set

$$P(r_1, ..., r_n) = \left\{ (z_1, \dots z_n) \in \mathbb{C}^n \;\middle|\; \pi \cdot |z_j|^2 \leq r_j \,,\, j = 1, 2, ..., n \right\},$$

with the symplectic form induced from $(\mathbb{C}^n, \omega_{std})$.

Let us now briefly motivate the study of symplectic embeddings. To simplify the discussion, we begin by considering *Liouville domains*, which are compact symplectic manifolds

with boundary such that the symplectic form has a positive primitive on the boundary. The boundary of a Liouville domain carries dynamical information. Consider the tangent space at any point of the boundary. For dimension reasons, there is a direction (generated by a vector in the tangent space) for which the symplectic form vanishes. The integral curves of this direction field will give a foliation of the boundary. These integral curves are called *Reeb orbits*. One can define the *action* of a closed Reeb orbit by integrating the positive primitive of the symplectic form over the orbit. More details about this construction will be given in chapter 4 of this work.

Example 1.0.1. The 4-dimensional ellipsoid $E(a, b)$, with a/b irrational will have only two closed Reeb orbits. Considering the ellipsoid as a subset of $\mathbb{C}^2$, the intersection of the boundary of the ellipsoid with the complex plane $z_j = 0$, $j = 1, 2$ will give a closed Reeb orbit. These closed Reeb orbits will have action πa and πb, respectively.

It turns out that a symplectic embedding detects and reflects the subtle Reeb dynamics of the boundaries of the source and target manifolds. Suppose we have a symplectic embedding of one Liouville domain into another, say $\varphi\colon X_1 \to X_2$. The so-called symplectic cobordism $X_2 \setminus X_1 = X_2 \setminus \varphi(X_1)$ has boundary

$$\partial(X_2 \setminus X_1) = \partial X_1 \sqcup \partial X_2,$$

which obviously carries information about the Reeb dynamics of both X_1 and X_2. We look for J-holomorphic curves in the cobordism $X_2 \setminus X_1$. A J-holomorphic curve is a map from a punctured Riemann surface into $X_2 \setminus X_1$ which satisfies the Floer equation (a PDE). The boundary conditions on the Floer equation ensure that the punctures in the domain of the J-holomorphic curve map to Reeb orbits on the boundary of $X_2 \setminus X_1$. The upshot is that a careful analysis of J-holomorphic curves in $X_2 \setminus X_1$ may obstruct an embedding of X_1 into X_2. In dimension 4 there are useful tools for the analysis of J-holomorphic curves, one of which will be described shortly. Indeed, these tools have led to a detailed study of 4-dimensional embedding problems involving balls, ellipsoids, and polydiscs (see [Sch05], [MS12b], [FM15], [CGFS17], and [Ush19]). Chapters 2 and 3 of this work will focus on four-dimensional symplectic embedding problems. Chapter 4 does solve a high-dimensional embedding problem for a large class of polydiscs. At this time the only other high-dimensional polydisc embedding result is [Gut08].

Even though the dynamical information at the boundary of a symplectic manifold is important, we often consider embeddings that are constructed using an exhaustion method.

Notation. The symbol $U \hookrightarrow V$ means that there exists a symplectic embedding $K \to V$ for

every compact set $K \subset Interior(U)$. Take, for instance, $K = (1-\epsilon)U$ for any $0 < \epsilon < 1$, and for U compact.

One analytical tool that is exclusive to dimension 4 is embedded contact homology (abbreviated ECH), which we shall briefly motivate. ECH is, as the name implies, a homology theory. The ECH chain complex is generated by multiples of embedded Reeb orbits which satisfy a non-degeneracy condition. The chain complex differential counts certain J-holomorphic curves connecting (asymptotic to) two chain complex generators. Since this homology theory is constructed to count the very curves that can be used to obstruct embeddings, it should come as no surprise that ECH carries information about the embedding properties of a symplectic 4-manifold. Indeed, by filtering the ECH chain complex generators by action, one can define a sequence of symplectic invariants called ECH capacities. We will not focus here on the full homology theory. Instead, the capacities will be our main computational tool for obstructing embeddings.

ECH capacities assign to a symplectic 4-manifold X an infinite sequence of invariants $c_*(X)$ with the following properties (cf. [Hut14] §1.2)

1. Each $c_i(X) \in \mathbb{R} \cup \{\infty\}$, and $0 = c_0(X) \leq c_1(X) \leq ... \leq \infty$.

2. (Monotonicity) If (X_1, ω_1) symplectically embeds into (X_2, ω_2) then

$$c_k(X_1) \leq c_k(X_2),$$

   for all $k \geq 0$.

3. (Conformality) Let X_2 be defined by rescaling the symplectic manifold X_1 as $(X_2, \omega_2) = (X_1, r\omega_1)$ for $r > 0$. Then

$$c_k(X_2) = |r| c_k(X_1).$$

The monotonicity property is the workhorse of this invariant. We often use the contrapositive of the monotonicity property to obstruct a symplectic embedding of 4-manifolds. Specifically, if just one ECH capacity of X_1 is larger than the corresponding capacity for X_2, then no embedding $X_1 \hookrightarrow X_2$ exists. Sometimes the monotonicity property can be strengthened to an if and only if statement. In that situation, we say that the ECH capacities give *sharp* obstructions to symplectic embeddings. (A large class of examples for which ECH capacities are sharp will be described in proposition $1.0.1$, below.)

There is another property of ECH capacities that we shall examine in detail.

4. (Disjoint union)
$$c_k\left(\bigsqcup_{i=1}^{m} X_i\right) = \max_{k_1+\ldots+k_m=k} \sum_{i=1}^{m} c_{k_i}(X_i).$$

Let's focus on the case of a disjoint union of $m = 2$ symplectic manifolds. In this case

$$c_k(X_1 \sqcup X_2) = \max_{k_1+k_2=k} \{c_{k_1}(X_1) + c_{k_2}(X_2)\}.$$

This is the inspiration for an operation on sequences.

Definition 1.0.4. Given two infinite sequences u and v, define the *sequence sum operation* as follows
$$(u + v)_k = \max_{i+j=k} \{u_i + v_j\}.$$
This operation gives a new infinite sequence.

Notice that the right-hand side of $c_k(X_1 \sqcup X_2)$ is, by construction, the sequence sum of the ECH capacities for X_1 and X_2. The moral of the story is the following.

Note. The operation of disjoint union on symplectic 4-manifolds effects the operation of sequence sum on the level of ECH capacities.

The operation of sequence sum is associative and commutative. So one could extend this observation to more than just a disjoint union of $m = 2$ manifolds.

It is natural to ask what operation "undoes" the above fact. On the level of symplectic manifolds, can we find an inverse to disjoint union? On the level of sequences, can we find an inverse to sequence addition? To that end, let us try.

Definition 1.0.5. Given two infinite sequences u and v, define the *sequence subtraction operation* as follows
$$(u - v)_k := \inf_{t\geq 0} \{u_{k+t} - v_t\}.$$
This operation gives a new infinite sequence. For the purpose of sequence subtraction, we define $\infty - \infty = \infty$.

As is the case with real numbers, the operation of sequence subtraction is not associative and is not commutative. We would like to understand this operation as it applies to ECH capacities. Because of property (1) of the ECH capacity sequence (above), if we want the result of sequence subtraction to be a sequence of ECH capacities, then the resulting sequence had better be nondecreasing and nonnegative. Hence we should only perform $u - v$ if $u \geq v$ as sequences. This is explained below.

Definition 1.0.6. Given two infinite sequences u and v, say that $u \geq v$ if $u_k \geq v_k$ for all $k \geq 0$. Furthermore, we write $u > v$ if $u \geq v$ and

$$\lim_{k\to\infty}(u_k - v_k) = \infty \text{ or } \infty - \infty.$$

When $u \geq v$ we will say that it is permissible to perform $u - v$. When $u > v$, sequence subtraction is possible and the infimum in the definition can be replaced with a minimum.

Note. In this work, we shall apply the operations of sequence sum and sequence subtraction only to sequences of ECH capacities.

Now that we have set out some properties of ECH capacities, let us give some useful formulas for their computation.

Example 1.0.2. (cf. [Hut14]) The k^{th} ECH capacity of the ellipsoid $E(a, b)$ is given by the $(k + 1)^{st}$ smallest number in the infinite array

$$ia + jb, \qquad i, j \in \mathbb{Z}_{\geq 0}.$$

The zeroeth capacity $c_0(E(a, b))$ is always zero. A special case of this computation gives the ECH capacities for a ball, $B^4(a) = E(a, a)$. The above formula reduces to

$$c_k(B^4(a)) = d \cdot a,$$

where d is the unique non-negative integer satisfying

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

Example 1.0.3. (cf. [Hut14]) The k^{th} ECH capacity of the symplectic polydisc $P(a, b)$ is given by

$$c_k(P(a, b)) = \min\left\{ia + jb \,|i, j \in \mathbb{Z}_{\geq 0}, (i + 1)(j + 1) \geq k + 1\right\}.$$

Example 1.0.2 and example 1.0.3 give useful formulas for computing ECH capacities, but these formulas are computationally expensive. Chapters 2 and 3 of this work find ways to make the computation of these ECH capacities easier.

1.0.1 Summary of Results

We now give a brief summary of the main results of this work. In chapter 2, a special class of symplectic manifolds, known as toric domains, will be defined. Toric domains are special, because the computation of their symplectic invariants (i.e. ECH capacities) often reduce to combinatorial problems for regions in Euclidean space. Specifically, we associate to a toric 4-manifold its *moment image*, which is a region in the first quadrant of the plane. Then we associate to the moment image a *weight vector*. This weight vector will serve as a convenient shorthand for the full sequence of ECH capacities: if two toric domains have the same weight vector, then they have the same sequence of ECH capacities.

Two classes of toric domains will be considered, called concave and convex. In chapter 2, algorithms will be given for computing the weight vector of both a concave toric domain and a convex toric domain. We use these techniques to compute the weight vector of an important example, the Lagrangian bidisc, to more accuracy than was done in [Ram17]. The main result of chapter 2 will be a method for converting a concave toric domain into a convex toric domain. Theorem 2.2.1 will describe how to compute the weight vector of the resulting convex toric domain from the weight vector of the starting concave toric domain.

The following result of Cristofaro-Gardiner shows that concave and convex toric domains are a good class of examples for computing ECH capacities.

Proposition 1.0.1. (Cristofaro-Gardiner, [CG19] §5) ECH capacities are *sharp* for embeddings of concave toric domains into convex toric domains. This means that an embedding of a concave toric domain X_1 into a convex toric domain X_2 is obstructed if and only if

$$c_k(X_1) > c_k(X_2) \qquad \text{for some } k \geq 0.$$

In chapter 3 of this work we consider the *Cremona transform* as an operation on the weight vectors corresponding to convex toric domains. The main result of chapter 3 is theorem 3.1.1, which shows that the Cremona transform does not change the ECH capacity sequence that results from a given weight vector. A repeated application of the Cremona transform (and re-ordering the vector entries) will simplify the weight vector. The purpose of this simplification is to make the computation of ECH capacities easier.

In chapter 4 we turn to a concrete embedding problem. A symplectic manifold is said to be *stablized* when it is crossed with several copies of $\mathbb{C}$. For example, stabilizing a 4-dimensional polydisc, $P(a,b)$ would result in

$$P(a,b) \times \mathbb{C}^N = P(a,b,\underbrace{\infty,...,\infty}_{N}).$$

The main result of chapter 4 is the following.

Theorem. Suppose that $x \geq 2$ and $N \geq 1$. There exists a symplectic embedding

$$P(1,x) \times \mathbb{C}^N \hookrightarrow P(a,b) \times \mathbb{C}^N$$

if and only if either

- $a \geq 2$, or
- $1 \leq a < 2$ and $b \geq x$.

As mentioned above, symplectic embedding problems are largely open in dimensions greater than four. The stabilization procedure provides a nice middle ground between the known techniques of dimension four and the open problems in higher dimensions. Specifically, we can start by considering the un-stabilized symplectic cobordism $P(a,b) \setminus P(1,x)$, look for J-holomorphic curves there, and stabilize everything under consideration (crossing the curves with a point). The benefit of this approach is that we can use the techniques of embedded contact homology to enumerate curves in the 4-dimensional cobordism, before arguing that these curves persist to the stabilized case. The positivity of area of the stabilized curves will then give conditions on a, b in terms of x, the parameters in the above theorem.

1.0.2 Important Past Results

Here we present four landmark past results that shaped the field. Dates will be given for historical context. The first result shows the connection between symplectic geometry and volume-preserving geometry. It was originally formulated in the language of classical mechanics.

Theorem. (Liouville, 1838, [Lio38]) The Hamiltonian flow on the phase space of a mechanical system preserves volume. Consequently, symplectic embeddings preserve volume.

Indeed, on a 2-dimensional symplectic manifold, a symplectic form is an area form, and Liouville's theorem is immediate. For a century and a half it was unknown to what extent symplectic geometry differed from volume-preserving geometry in higher dimensions. The problem was further complicated by a remarkable discovery of Darboux.

Theorem. (Darboux, 1882, [Dar82]) Let (M^{2n}, ω) be a symplectic manifold. About any point $p \in M$ there is a neighborhood U (called a "Darboux chart") that is diffeomorphic to an open set V of $(\mathbb{C}^n, \omega_{std})$, say $\varphi \colon V \to U$, with the property that $\varphi^* \omega = \omega_{std}$.

An important consequence of this theorem is the absence of local invariants in symplectic geometry. Since any two symplectic structures on (M, ω) are locally symplectically equivalent using Darboux charts, then the only possible symplectic invariants of M must be global. Volume is a global invariant of a sympletic manifold, but (in dimensions greater than two) this invariant obviously does not distinguish the symplectic geometry of M from its volume-preserving geometry. At the time of Liouville and Darboux, no other global invariant of symplectic manifolds was known. It was not until Gromov's breakthrough non-squeezing theorem (described below) that mathematics learned of the difference between these geometries and of the importance of embedding problems to studying symplectic invariants.

The above examples of ellipsoids and polydiscs were the first examples studied to try to understand the 4-dimensional symplectic embedding problem. A simple starting point would be to ask the following.

Question. For what values of r_1, r_2, R, r_1', r_2' does there exist a symplectic embedding

$$\text{from} \quad \begin{cases} E(r_1, r_2) \\ \text{or} \\ P(r_1, r_2) \end{cases} \quad \text{into} \quad \begin{cases} B(R) \\ \text{or} \\ E(r_1', r_2') \text{ or } P(r_1', r_2') \end{cases} \quad ?$$

In order to parameterize this problem, we introduce an *embedding function*, which is best explained by example. In this language, Gromov studied the function

$$f_{Gromov}(x) = \inf\{R \,|\, E(1, x) \text{ symplectically embeds into } P(R, \infty)\}, \quad \text{for } x \geq 1,$$

and his celebrated non-squeezing theorem can be stated as follows.

Theorem. (Gromov, 1985, [Gro85]) The function $f_{Gromov}(x)$ is identically one.

The target manifold of this embedding problem has infinite volume, whereas the source has finite volume. One would be able to stretch an ellipsoid to fit into a cylinder in a volume-preserving manner, but not in a way that preserves the symplectic structure. Hence the name "non-squeezing". This result is the first indication that symplectic geometry differs from volume-preserving geometry.

For the four-dimensional problem of embedding $E(1, x)$ into $B^4(R)$ we set

$$f_{EB}(x) = \inf\{R \,|\, E(1, x) \text{ symplectically embeds into } B^4(R)\}, \quad \text{for } x \geq 1.$$

McDuff and Schlenk computed the function f_{EB} for all $x \geq 1$. In particular, they discovered the following.

Theorem. (McDuff-Schlenk, 2012, [MS12b]) For $1 \leq x \leq (\frac{1+\sqrt{5}}{2})^4$, the graph of f_{EB} is shown in figure 1.1.

This portion of the graph is informally called the "Fibonacci staircase," because it the graph resembles a staircase having infinitely many steps, with the corner of each stair step occurring at a ratio of Fibonacci numbers.

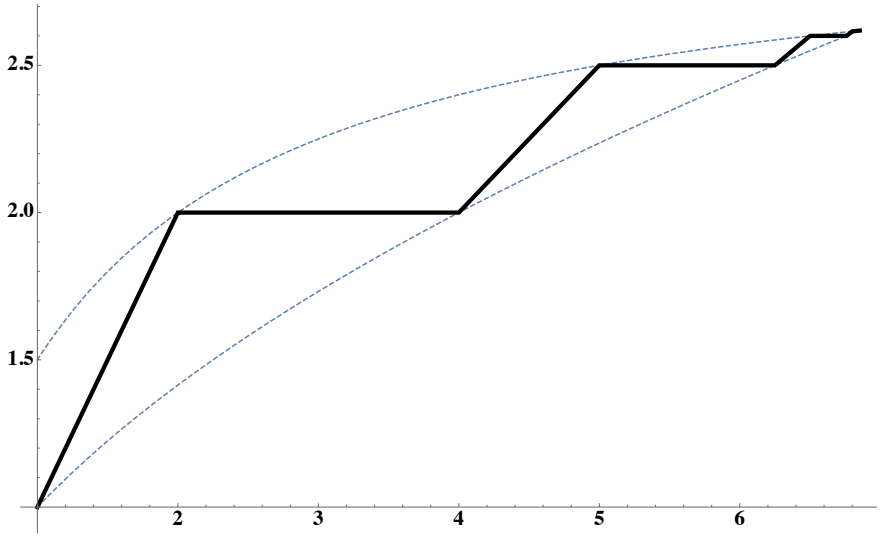

Figure 1.1: A portion of the graph of f_{EB}.

Notice also that the graph in figure 1.1 is bounded by two other curves (which are dashed in the figure). The lower bound is the graph of $y = \sqrt{x}$, the minimal volume of a ball into which $E(1, x)$ could embed. Every time the graph of f_{EB} rises above this volume curve, there is some obstruction to the embedding problem that is more subtle than the volume obstruction. The upper bound is the graph of $y = 3x/(x + 1)$, which is the embedding that one obtains using the technique of "symplectic folding" (cf. [Hin15] for the folding construction). Every time the graph of f_{EB} meets the folding curve, we have that folding gives the optimal embedding of an ellipsoid into a ball. The beautiful graph of f_{EB} showed that the symplectic embedding problem is subtle and rich.

Similarly, for the four-dimensional problem of embedding the ellipsoid $E(1, x)$ into the cube $C^4(R) = P(R, R)$ we define

$$f_{EC}(x) = \inf\{R \,|\, E(1, x) \text{ symplectically embeds into } C^4(R)\}, \quad \text{for } x \geq 1.$$

In [FM15], Frenkel and Müller computed the function $f_{EC}(x)$ for all $x \geq 1$. A portion of the graph is reproduced in figure 1.2. This graph also exhibits interesting staircase-like behavior. The corners of the stair steps in the graph occur at ratios of Pell numbers. This

graph also has a lower bound, given by the graph of $y = \sqrt{x/2}$, which represents the volume of the smallest cube into which $E(1, x)$ can embed. The upper bound is the graph of $y = 2x/(x + 1)$, which represents the embedding that one obtains using the symplectic folding construction.

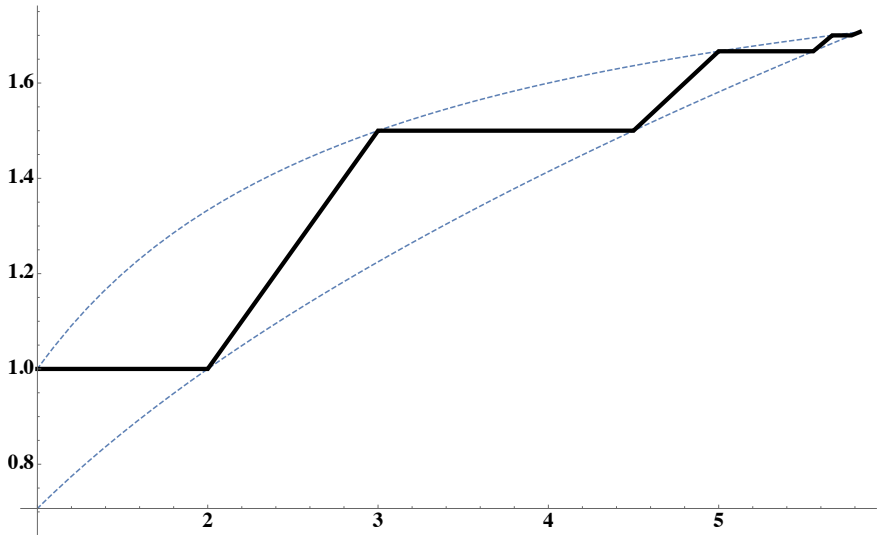

Figure 1.2: A portion of the graph of f_{EC}.

The functions f_{EB} and f_{EC} illustrate the richness of the symplectic embedding problem. These graphs have interesting number theoretic properties (via the Fibonacci and Pell numbers), which we will not examine in this work. The connections to planar geometry and blowups of the complex projective plane will be examined in chapters 2 and 3, respectively.

CHAPTER 2

Toric Domains

In this chapter, we discuss toric domains, which will be our main source of examples of symplectic manifolds. We shall see that for certain 4-dimensional toric domains, the infinite sequence of ECH capacities can be abbreviated by a vector of finite length, called the weight vector. Conversely, given a weight vector, one can produce a toric domain having the given vector as its weight vector. Note, however, that the translation between toric domains and weight vectors is not bijective, because there can be two non-symplectomorphic toric domains with the same weight vector. This is implied by the following exercise.

Exercise. Use example $1.0.2$ and example $1.0.3$ to show that the ellipsoid $E(1,2)$ and the polydisc $P(1,1)$ have the same sequence of ECH capacities. After reading proposition $2.1.3$ or proposition $2.1.4$, one can show that the toric domains $E(1,2)$ and $P(1,1)$ have the same weight vector, which will be $(2;1,1)$.

Nevertheless, we use the weight vector as a convenient shorthand for the full sequence of ECH capacities, which encode the embedding properties of a toric domain, as described in chapter 1.

Definition 2.0.1. A *toric domain* is the preimage of some region in the first orthant of $\mathbb{R}^n$ under the map $\mu\colon\ \mathbb{C}^n \to \mathbb{R}^n$ given by

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

A toric domain inherits the standard symplectic form $(1.0.1)$ from $\mathbb{C}^n$. The map μ is the moment map for the standard torus action on $\mathbb{C}^n$. For this reason, the region in $\mathbb{R}^n$ that is associated to a toric domain X is usually called the *moment image* of X.

In chapters 2 through 3, we focus on four-dimensional toric domains embedded in $\mathbb{C}^2$. In this case, we fix $n = 2$ and consider the preimage of regions in the plane. Specifically, we shall consider the toric domains that are associated to regions in the (closed) first quadrant.

Such a planar region can be either concave or convex, as we define below. We consider the toric domain that results to be invariant under translation of the planar region and the action of $SL(2,\mathbb{Z})$ on $\mathbb{R}^2$. For this reason, we may as well assume that all planar regions under consideration have some edges coinciding with the x-axis or the y-axis.

Definition 2.0.2. A 4-dimensional toric domain is called *concave* or *convex* if its corresponding planar region is bounded by three curves: the line segment from $(0,0)$ to $(0,\eta)$; the line segment from $(0,0)$ to $(\zeta,0)$; and the graph of a continuous function connecting $(0,\eta)$ with $(\zeta,0)$. Specifically, the two possibilities are as follows.

1. (See [CCGF+14].) A *concave* region lies beneath the graph of a convex function $f\colon [0,\zeta] \to [0,\eta]$ such that $f(0) = \eta$ and $f(\zeta) = 0$. Here we interpret the convexity of f non-strictly, meaning $f'' \geq 0$, whenever the second derivative exists. Furthermore $f' < 0$, whenever the first derivative exists.

2. (See [GH17].) A *convex* region lies beneath the graph of a non-increasing concave function $g\colon [0,\zeta] \to [0,\eta]$ such that $g(0) = \eta$. We do not require that $g(\zeta) = 0$, but if $g(\zeta) \neq 0$, then the planar region includes the line segment from $(\zeta,0)$ to $(\zeta, g(\zeta))$.

In both cases, we have $\zeta > 0$ and $\eta > 0$. A more general notion of convexity is used in [CG19], but we will not need that generality here. A special case of the above definition occurs when the planar region is a polygon, as defined below.

Definition 2.0.3. A *polygon* is a closed region in the closed first quadrant of the plane that is either concave or convex, according as to the two cases in definition 2.0.2. Furthermore, the bounding function f or g in that definition must be piecewise-linear.

When speaking of a polygon Δ, we denote the number ζ by $\zeta(\Delta)$ and we denote the number η by $\eta(\Delta)$. A polygon which meets the criteria to be concave or convex will have a vertex at the origin, and the interior angle at that vertex will be right. Some polygons in the plane are obviously concave or obviously convex, but will need to be translated or acted upon by $SL(2,\mathbb{Z})$ in order to meet the criteria of definitions 2.0.2 and 2.0.3. Once the criteria in this definition have been met, we say that the concave or the convex polygon is in *standard position*.

Notation. For brevity, we will often write "the region below the graph of f" as a shorthand for "the region in the (closed) first quadrant bounded by the graph of f, and the two line segments connecting the origin to the intercepts of f".

For the purpose of this chapter, we will set X_Δ to be the toric domain that is obtained by pulling back either a concave polygon, in which case we call X_Δ *concave*; or a convex

polygon Δ, in which case we call X_Δ *convex*. We sometimes also assume that Δ is *integral*, meaning all of its vertices are at points of $\mathbb{Z}^2 \subset \mathbb{R}^2$. More generally, we could require Δ to be *rational*, meaning all of its vertices are at points of $\mathbb{Q}^2 \subset \mathbb{R}^2$. If the moment image of a toric domain is not a polygon, but is still concave or convex, we apply those adjectives to the toric domain, too.

Example 2.0.1. Consider the triangle in the first quadrant with vertices at points $(0,0)$, $(a,0)$, and $(0,a)$. The preimage of this triangle under the moment map will be the 4-ball with symplectic area a, namely $B^4(a)$. Such a ball is extremely important for symplectic ball-packing examples, and more generally for symplectic embedding problems. For this reason, we should give a name to the polygon that corresponds to $B^4(a)$. Notice that the polygon is an isosceles right triangle with legs on the axes. We call this polygon a *standard triangle of size* a.

Example 2.0.2. Given $0 < a \leq b$, the symplectic ellipsoid of radii $\sqrt{a}, \sqrt{b}$ (sometimes called the symplectic ellipsoid of areas a and b) is given in definition 1.0.2. The ellipsoid $E(a,b)$ is a toric domain with moment image being a triangle with vertices at $(0,0)$, $(a,0)$, and $(0,b)$. This triangle is both concave and convex. Hence $E(a,b)$ is both concave toric and convex toric. Notice also that $E(a,a) = B^4(a)$.

Example 2.0.3. Given $0 < a \leq b$, the symplectic polydisc of radii $\sqrt{a}, \sqrt{b}$ (also called the symplectic polydisc of areas a and b) is given in definition 1.0.3. The polydisc $P(a,b)$ is a convex toric domain with moment image being a rectangle of width a and height b. This moment image is a convex polygon bounded by the graph of $y = b$ together with the vertical line $x = a$. When $a = b$, the resulting manifold $C^4(a) := P(a,a)$ is called the symplectic cube of radius $\sqrt{a}$.

2.1 ECH Capacities of Toric Domains

Recall that ECH capacities are a strictly 4-dimensional symplectic capacity. When applied to a toric domain, we can compute ECH capacities directly from the moment image. In fact, we can abbreviate the infinite sequence of ECH capacities by a tuple of numbers. In the case when X_Δ is a toric domain with moment image a standard triangle or a rational convex or concave polygon, the tuple of numbers that determines the ECH capacities will be finite in length. This finite vector is called the *weight vector* associated to X_Δ. If the moment image is not a standard triangle or a rational polygon, the *weights* of the toric domain can still be defined, but will be an infinite sequence. We will pursue this issue further following proposition 2.1.2.

Notation. For sections 2.1.1 and 2.1.3, Δ will always denote a rational polygon, and X_Δ will denote the associated toric domain. In the case when Δ is concave, a recursive, geometrical construction gives a vector that we shall call the *positive weight vector*. In the case when Δ is convex, we get a *negative weight vector*. These two constructions are detailed below. When the notation Ω or X_Ω is used, it denotes a more general toric domain, not necessarily rational. Again, we view these constructions as associating a vector (or an infinite sequence of real numbers) to a polygon.

We will also use the following continuity property of ECH capacities of a toric domain.

Lemma 2.1.1. (cf. Lemma 2.3 of [CCGF+14]) For a fixed value of k, $c_k(X_\Omega)$ is a continuous function of Ω with respect to the Hausforff metric.

2.1.1 ECH Capacities of Concave Toric Domains

We begin with the concave case, explaining how to obtain a positive weight vector, denoted $(a_1, ..., a_N)$. The algorithm will terminate if we assume that the concave polygon Δ is rational or is a standard triangle. The following procedure has been adapted from [CCGF+14], and the procedure is originally inspired by [MS12b].

Proposition 2.1.1. (The positive weight vector of a concave polygon) Let Δ be a rational concave polygon in standard position with $\zeta, \eta > 0$ or let Δ be a standard triangle.

1. If Δ is a *standard triangle* of size a, then the positive weight vector is (a), and this algorithm terminates.

2. Otherwise, find the largest *standard triangle* inscribed within Δ. It is acceptable for some edges of Δ to coincide with some edges of the standard triangle. The size of this triangle is the first entry of the positive weight vector, a_1.

3. The line $x + y = a_1$ meets the graph of f (the function that bounds Δ) in a line segment from a point we call $(x_2, a_1 - x_2)$ to a point $(x_3, a_1 - x_3)$ with $x_2 \leq x_3$. Let Δ'_2 denote the (closed) portion of Δ that is above the line $x + y = a_1$ and left of the line $x = x_2$. Let Δ'_3 denote the (closed) portion of Δ that is above the line $x+y = a_1$ and right of the line $x = x_3$.

4. Apply two transformations to the subset Δ'_2. First apply the translation $(x, y) \mapsto (x, y - a_1)$, which places one vertex of Δ'_2 at the origin. Then act on Δ'_2 by $\begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix} \in SL(2, \mathbb{Z})$. The result of these two transformations will be a new rational concave

polygon in standard position, which we call Δ_2. If $x_2 = 0$, then we interpret Δ_2 as the empty set.

5. On a separate set of coordinate axes, apply two transformations to the subset Δ_3'. First apply the translation $(x, y) \mapsto (x - a_1, y)$, which places one vertex of Δ_3' at the origin. Then act on Δ_3' by $\begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \in SL(2, \mathbb{Z})$. The result of these two transformations will be a new rational concave polygon in standard position, which we call Δ_3. If $x_3 = a_1$, then we interpret Δ_3 as the empty set.

6. The positive weight vector already has first entry a_1. Adjoin to this vector the weight vectors for Δ_2 and Δ_3, which are defined recursively by re-starting this algorithm at step (1). The union of vector entries in this definition is an unordered union with repetitions.

See figure 1 in [CCGF+14] for an illustrative example of this procedure.

Example 2.1.1. Let $q \geq 1$ be rational. As explained in example 2.0.2 above, the ellipsoid $E(1, q)$ is a rational concave toric domain whose moment image is a triangle. It was shown (in [MS12b] appendix A) that the positive weight vector of $E(1, q)$ is given by the Farey expansion of the rational number q. Consequently, the elements of the positive weight vector for $E(1, q)$ are at most 1.

As a concrete example, the reader can compute using the algorithm above that the ellipsoid $E(1, 25/9)$ has positive weight vector

$$\left(1, 1, \frac{7}{9}, \frac{2}{9}, \frac{2}{9}, \frac{2}{9}, \frac{1}{9}, \frac{1}{9}\right).$$

Remark. There is another sequence associated to $E(1, q)$ in [MS12b], which is the continued fraction expansion of the rational number q. Let us write this continued fraction expansion as $q = [\ell_0; \ell_1, ..., \ell_m]$. This sequence is closely related to the weight expansion described above; the ℓ_j are given by the multiplicities of the weights. More precisely, the integer ℓ_0 is the number of times that 1 is repeated in the weight expansion of $E(1, q)$, and so on. We will not use the continued fraction expansion, but we include it here to distinguish it from the weight sequence. In example 2.1.1 above, we easily see that

$$\frac{25}{9} = [2; 1, 3, 2].$$

Once the positive weight vector has been computed for a concave toric domain X_Δ, say that it is $(a_1, ..., a_N)$. One can then compute the k^{th} ECH capacity for X_Δ using the

formula

$$c_k(X_\Delta) = c_k\left(\coprod_{j=1}^{N} B(a_j)\right). \tag{2.1.1}$$

Based on the observation in section 1, we can view the ECH capacities of a disjoint union as a sequence sum, as in definition 1.0.4. Let c_* denote the sequence of ECH capacities. Then

$$c_*(X_\Delta) = c_*(B(a_1)) + ... + c_*(B(a_N)). \tag{2.1.2}$$

Here we are using the associativity of the sequence sum operation. More precisely, this formula is equivalent to

$$c_k(X_\Delta) = \max\left\{\sum_{j=1}^{N} a_j d_j \,\middle|\, \sum_{j=1}^{N} d_j^2 + d_j \leq 2k\right\}, \tag{2.1.3}$$

where the maximum is taken over all N-tuples of non-negative integers $d_1, ..., d_N$.

The above procedure deconstructs a concave toric domain into smaller triangles of sizes given by the resulting positive weight vector. Let us, instead, re-formulate this algorithm to be more constructive. In so doing, we will exhibit a ball-packing of the given concave toric domain. First we set up an algebraic framework.

2.1.2 Realizing the positive weight vector as a ball packing

Consider the two matrices

$$A = \begin{bmatrix} 1 & 0 \\ -1 & 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix}^{-1}$$

and

$$B = \begin{bmatrix} 1 & -1 \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}^{-1}.$$

The monoid of words in the alphabet $\{A, B\}$ includes the identity element I (which we consider to be a word of length zero) together words of positive length. Some elements of this monoid are

$$\{I, A, B, AA, AB, BA, BB, AAA, AAB, ABA, ABB, BAA, BAB, BBA, BBB, ...\}.$$

This monoid is totally ordered by the lexicographic ordering, which is illustrated in the list above. Each of these matrices acts on the standard triangle of size a to get a list of so-called *deformed triangles*. The total ordering gives a convenient way to organize these deformed

triangles. Denote a deformed triangle by $\Delta_m^{(j)}$. The subscript m denotes the length of the word in the above monoid that acts on the standard triangle. The superscript j denotes the position in the total ordering restricted to words of length m. Since the size a of the standard triangle is arbitrary, we also decorate a with the same subscript and superscript when referring to the size of the deformed triangle. Recapitulating, $\Delta_m^{(j)}$ is some product of m matrices applied to the standard triangle of size $a_m^{(j)}$. We abuse notation slightly and say that this $a_m^{(j)}$ is the size of the deformed triangle. The list of the first fifteen deformed triangles follows in figures 2.1 through 2.4.

We won't need to draw triangles with $m \geq 4$ in the examples that follow, but it should be evident from the description how to generate such triangles. Each comes from a product of matrices A and B applied to a standard triangle.

Now we can re-formulate the positive weight algorithm of Proposition 2.1.1 as a triangle-packing construction, as opposed to a deconstructive algorithm.

Proposition 2.1.2. (The positive weight vector as a ball-packing) Let Δ be a concave polygon in standard position with $\zeta, \eta > 0$ (from the definition of concave polygon), or let Δ be a standard triangle. This will be an iterative algorithm, broken into several parts.

m=0 iteration:

1. If Δ is a *standard triangle* of size a, then the positive weight vector is (a), and this algorithm terminates.

2. Otherwise, record the largest value of $a_0^{(1)}$ such that $\Delta_0^{(1)}$ is inscribed within Δ.

m=1 iteration:

3. Translate all points of the deformed triangle $\Delta_1^{(1)}$ by the vector $(0, a_0^{(1)})$. Record the largest value of $a_1^{(1)}$ such that this translated triangle lies within the polygon Δ.

4. Translate the deformed triangle $\Delta_1^{(2)}$ by the vector $(a_0^{(1)}, 0)$. Record the largest value of $a_1^{(2)}$ such that this translated triangle lies within the polygon Δ.

m=2 iteration:

5. Translate the deformed triangle $\Delta_2^{(1)}$ by the vector $(0, a_0^{(1)} + a_1^{(1)})$. Record the largest value of $a_2^{(1)}$ such that this translated triangle lies within the polygon Δ.

6. Translate the deformed triangle $\Delta_2^{(2)}$ by the vector $(a_1^{(1)}, a_0^{(1)} - a_1^{(1)})$. Record the largest value of $a_2^{(2)}$ such that this translated triangle lies within the polygon Δ.

7. Translate the deformed triangle $\Delta_2^{(3)}$ by the vector $(a_0^{(1)} - a_1^{(2)}, a_1^{(2)})$. Record the largest value of $a_2^{(3)}$ such that this translated triangle lies within the polygon Δ.

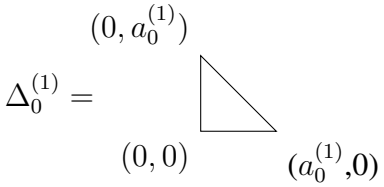

Figure 2.1: The $m = 0$ triangle with vertices labeled.

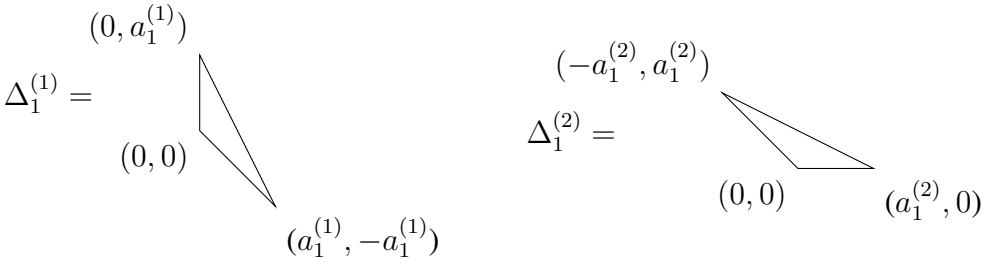

Figure 2.2: $m = 1$ triangles.

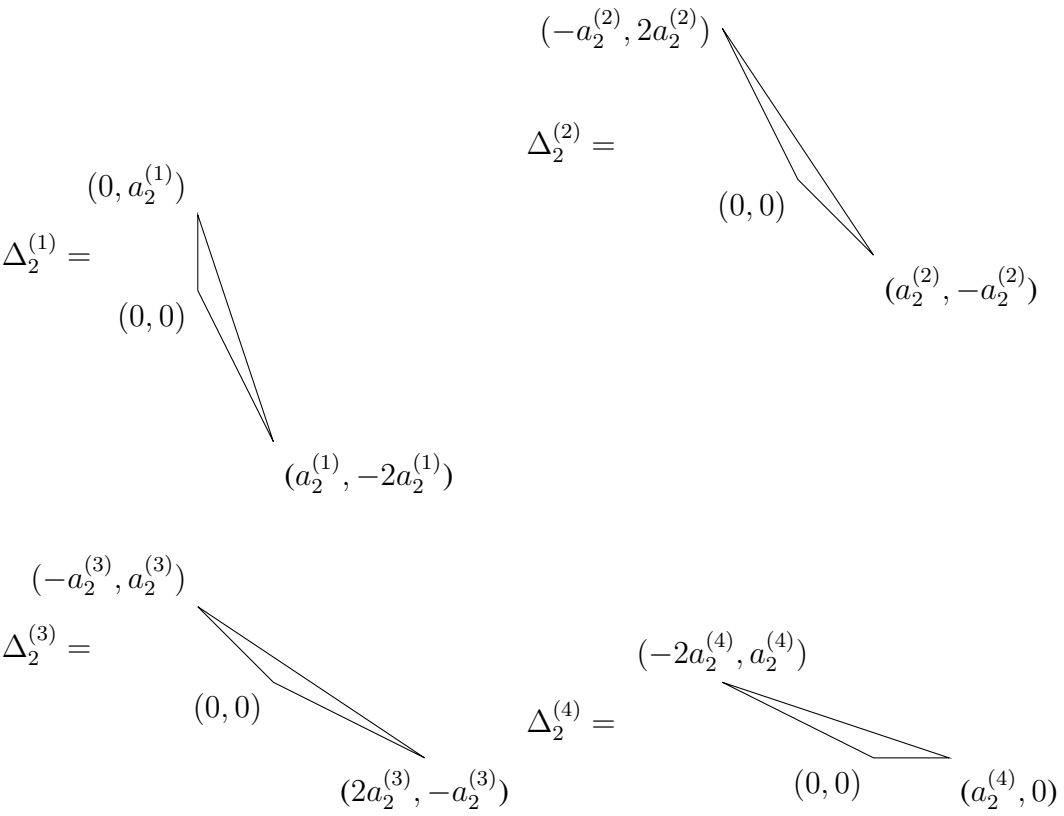

Figure 2.3: $m = 2$ triangles.

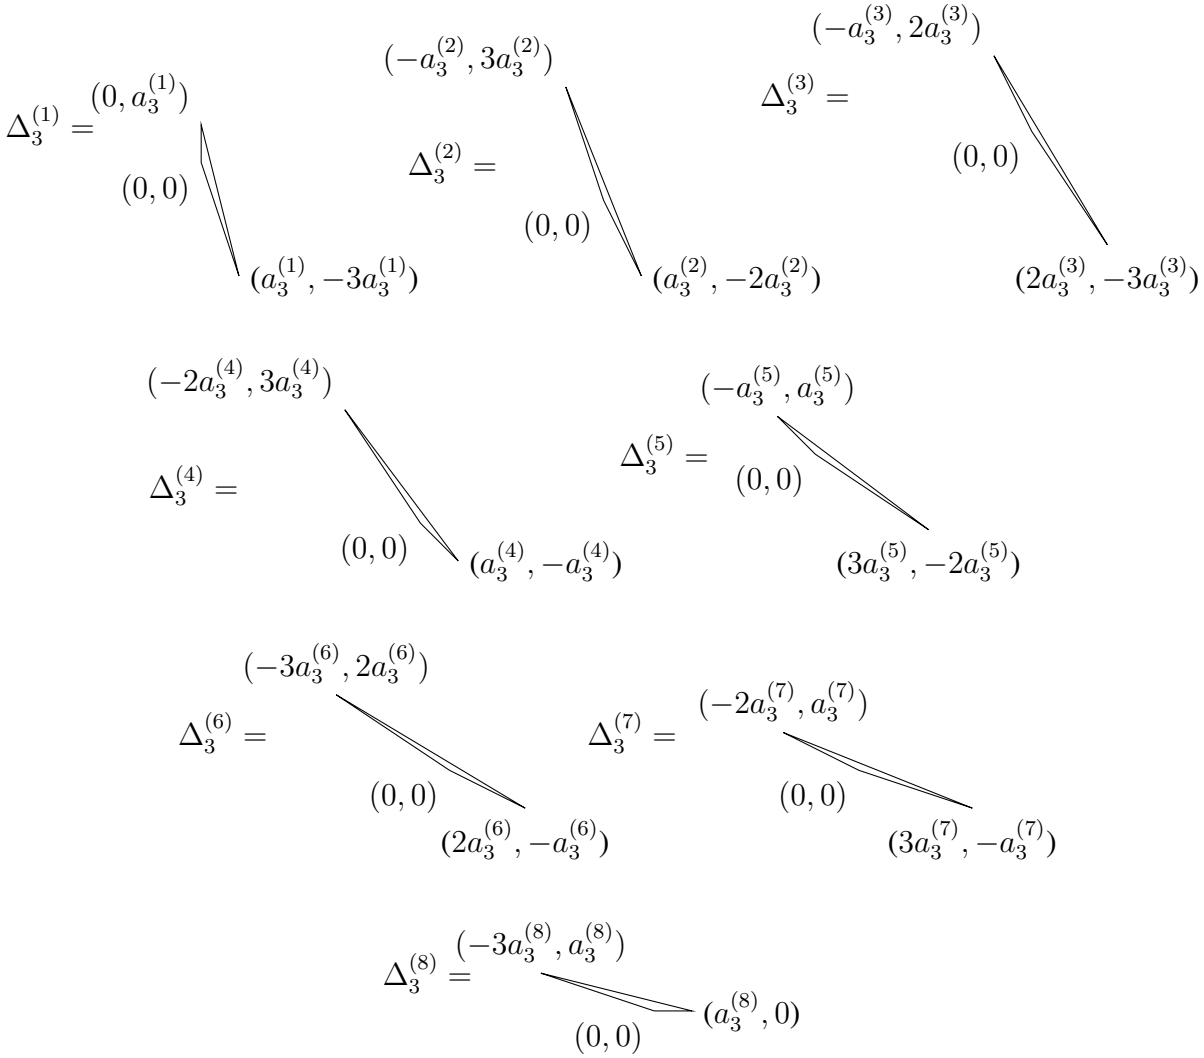

Figure 2.4: $m = 3$ triangles.

8. Translate the deformed triangle $\Delta_2^{(4)}$ by the vector $(a_0^{(1)} + a_1^{(2)}, 0)$. Record the largest value of $a_2^{(4)}$ such that this translated triangle lies within the polygon Δ.

m=3 iteration:

9. Translate the deformed triangle $\Delta_3^{(1)}$ by the vector

$$\begin{bmatrix} 0 \\ a_0^{(1)} + a_1^{(1)} + a_2^{(1)} \end{bmatrix}.$$

Record the largest value of $a_3^{(1)}$ such that this translated triangle lies within the polygon Δ.

10. Translate the deformed triangle $\Delta_3^{(2)}$ by the vector

$$\begin{bmatrix} a_2^{(1)} \\ a_0^{(1)} + a_1^{(1)} - 2a_2^{(1)} \end{bmatrix}.$$

Record the largest value of $a_3^{(2)}$ such that this translated triangle lies within the polygon Δ.

11. Translate the deformed triangle $\Delta_3^{(3)}$ by the vector

$$\begin{bmatrix} a_1^{(1)} - a_2^{(2)} \\ a_0^{(1)} - a_1^{(1)} + 2a_2^{(2)} \end{bmatrix}.$$

Record the largest value of $a_3^{(3)}$ such that this translated triangle lies within the polygon Δ.

12. Translate the deformed triangle $\Delta_3^{(4)}$ by the vector

$$\begin{bmatrix} a_1^{(1)} + a_2^{(2)} \\ a_0^{(1)} - a_1^{(1)} - a_2^{(2)} \end{bmatrix}.$$

Record the largest value of $a_3^{(4)}$ such that this translated triangle lies within the polygon Δ.

13. Translate the deformed triangle $\Delta_3^{(5)}$ by the vector

$$\begin{bmatrix} a_0^{(1)} - a_1^{(2)} - a_2^{(3)} \\ a_1^{(2)} + a_2^{(3)} \end{bmatrix}.$$

Record the largest value of $a_3^{(5)}$ such that this translated triangle lies within the polygon Δ.

14. Translate the deformed triangle $\Delta_3^{(6)}$ by the vector

$$\begin{bmatrix} a_0^{(1)} - a_1^{(2)} + 2a_2^{(3)} \\ a_1^{(2)} - a_2^{(3)} \end{bmatrix}.$$

Record the largest value of $a_3^{(6)}$ such that this translated triangle lies within the polygon Δ.

15. Translate the deformed triangle $\Delta_3^{(7)}$ by the vector

$$\begin{bmatrix} a_0^{(1)} + a_1^{(2)} - 2a_2^{(4)} \\ a_2^{(4)} \end{bmatrix}.$$

Record the largest value of $a_3^{(7)}$ such that this translated triangle lies within the polygon Δ.

16. Translate the deformed triangle $\Delta_3^{(8)}$ by the vector

$$\begin{bmatrix} a_0^{(1)} + a_1^{(2)} + a_2^{(4)} \\ 0 \end{bmatrix}.$$

Record the largest value of $a_3^{(8)}$ such that this translated triangle lies within the polygon Δ.

mth iteration

17. Continue this procedure for as long as necessary to compute $a_m^{(j)}$ for $m \geq 4$. The m^{th} iteration involves packing at most 2^m triangles in Δ. The translation vector for the triangle $\Delta_m^{(j)}$ is a vertex of a triangle that was constructed in the $m-1$ iteration.

The result

18. The (possibly finite) list given by $(a_0^{(1)}, a_1^{(1)}, a_1^{(2)}, ..., a_m^{(1)}, ...)$ is the positive weight expansion for Δ.

This constructive algorithm has its advantages. First, using the definition of X_Δ, it is immediate that

$$\text{vol}(X_\Delta) = \text{area}(\Delta) = \sum_{m,j} \text{area}(\Delta_m^{(j)}) = \frac{1}{2} \sum (a_m^{(j)})^2, \tag{2.1.4}$$

where the rightmost sum is taken over the terms $a_m^{(j)}$ in the positive weight expansion of Δ. Again, this sum will have finitely many terms when Δ is a rational concave polygon or a standard triangle. When the sum has infinitely many terms, it must still converge if $\mathrm{vol}(X_\Delta)$ is finite.

In example 2.1.1, we explain that the positive weight vector of the ellipsoid $E(1,q)$ is given by the Farey expansion of q. The above volume formula reproduces formula (1.2.2) from [MS12b]. Here we have given another proof to the formula, using triangle packings of a triangle instead of square packings of a rectangle.

The second advantage of the constructive algorithm is that it can be repeated *ad infinitum*. As mentioned above, the positive weight vector for X_Ω can be defined even when Ω is not a rational polygon. For more general concave Ω, the algorithms in Proposition 2.1.1 and Proposition 2.1.2 should not be expected to terminate. In this case, the positive weights give an infinite sequence of numbers, say (a_i). One can then compute the ECH capacities to any desired degree of accuracy by truncating this infinite sequence to length N, and then using the formula (2.1.3) and lemma 2.1.1. The constructive algorithm makes it clear that such a general Ω can be packed with infinitely many triangles, as long as the sequence of areas of these triangles vanishes to second order.

Third, one could run this constructive algorithm in reverse: given a (valid) positive weight vector, proposition 2.1.2 explains how to position and size triangles in the first quadrant, the outer boundary of which will define a concave region of the plane. Hence we can construct a concave toric domain with a given vector as its positive weight vector. Note that there can be multiple toric domains with the same weight vector, however.

Example 2.1.2. We present an example computation of the ball-packing algorithm in Proposition 2.1.2. Consider the concave polygon that is bounded by the piecewise-linear function

$$f\colon [0,27] \to [0,27]$$

given by

$$f(x) = \begin{cases} 27 - 4x & \text{if } 0 \leq x \leq 1 \\ \frac{51}{2} - \frac{5x}{2} & \text{if } 1 \leq x \leq 3 \\ 23 - \frac{5x}{3} & \text{if } 3 \leq x \leq 6 \\ 21 - \frac{4x}{3} & \text{if } 6 \leq x \leq 9 \\ \frac{63}{4} - \frac{3x}{4} & \text{if } 9 \leq x \leq 13 \\ \frac{69}{5} - \frac{3x}{5} & \text{if } 13 \leq x \leq 18 \\ \frac{51}{5} - \frac{2x}{5} & \text{if } 18 \leq x \leq 23 \\ \frac{27}{4} - \frac{x}{4} & \text{if } 23 \leq x \leq 27 \end{cases}$$

The polygon bounded by the graph of f and the x- and y-axes is shown in figure 2.5. That figure also shows the deformed triangles that fill the polygon to give positive weight vector $(18, 6, 6, 2, 2, 2, 2, 1, 1, 1, 1, 1, 1, 1, 1)$. This example polygon is designed to be symmetrical about the line $y = x$ and to include one of each of the deformed triangles shown in figures 2.1 through 2.4.

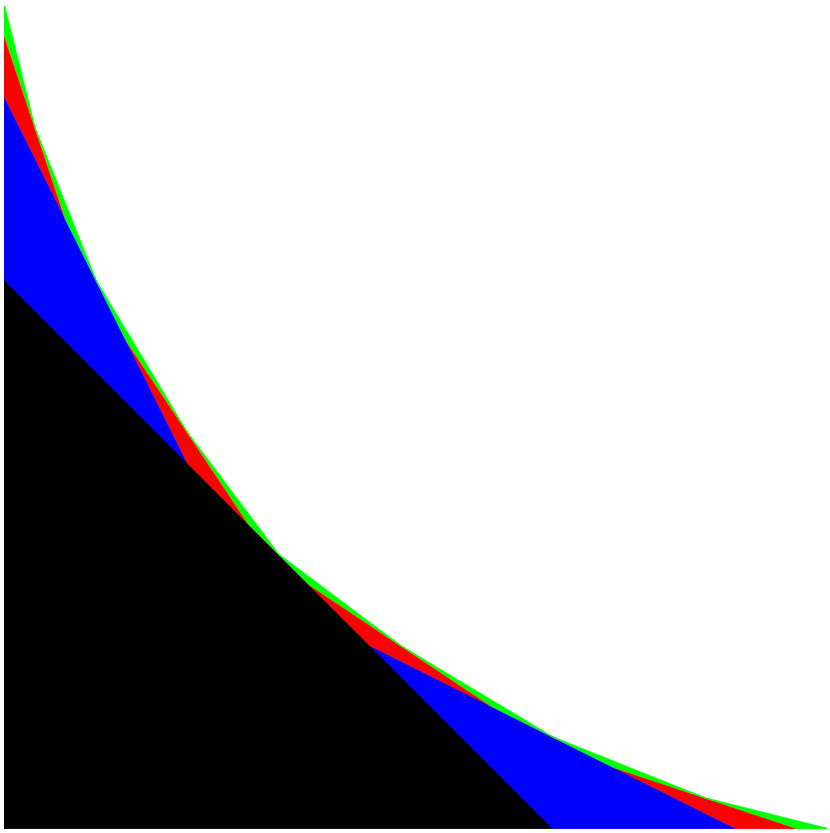

Figure 2.5: An example of a concave toric domain.

2.1.3 ECH Capacities of Convex Toric Domains

In this section we explain how to obtain the *negative weight vector* from a convex toric domain, and vice versa. The negative weight vector is a vector with $k+1$ entries, denoted $(b; a_1, ..., a_k)$. The semicolon in this notation should distinguish a negative weight vector from a positive weight vector. The algorithm below will be recursive and will rely on Proposition 2.1.1. For this reason we continue to assume that Δ is rational or is a standard triangle. The following procedure has been adapted from [CG19].

Proposition 2.1.3. (The negative weight vector of a convex polygon) Let Δ be a rational convex polygon in standard position, or let Δ be a standard triangle.

1. If Δ is a *standard triangle* of size b, then the negative weight vector is $(b; 0)$, and this algorithm terminates.

2. Otherwise, find the smallest *standard triangle* that circumscribes Δ. It is acceptable for some edges of Δ to coincide with some edges of the standard triangle. The size of this triangle is the first entry of the negative weight vector, b. Let Δ_1 denote the standard triangle of size b superimposed over the polygon Δ in standard position.

3. The line $x + y = b$ meets the curve that bounds Δ in a line segment from a point we call $(x_2, b - x_2)$ to a point $(x_3, b - x_3)$ with $x_2 \leq x_3$. Let Δ_2' denote the (closed) portion of $\Delta_1 \setminus \Delta$ that is below the line $x + y = b$ and left of the line $x = x_2$. Let Δ_3' denote the (closed) portion of $\Delta_1 \setminus \Delta$ that is below the line $x + y = b$ and right of the line $x = x_3$.

4. Apply two transformations to the subset Δ_2'. First apply the translation $(x, y) \mapsto (x, y-b)$, which places one vertex of Δ_2' at the origin. Then act on Δ_2' by $\begin{bmatrix} -1 & -1 \\ 1 & 0 \end{bmatrix} \in SL(2, \mathbb{Z})$. The result of these two transformations will be a new rational concave polygon in standard position, which we call Δ_2. If $x_2 = 0$, then we interpret Δ_2 as the empty set.

5. On a separate set of coordinate axes, apply two transformations to the subset Δ_3'. First apply the translation $(x, y) \mapsto (x - b, y)$, which places one vertex of Δ_3' at the origin. Then act on Δ_3' by $\begin{bmatrix} 0 & 1 \\ -1 & -1 \end{bmatrix} \in SL(2, \mathbb{Z})$. The result of these two transformations will be a new rational concave polygon in standard position, which we call Δ_3. If $x_3 = b$, then we interpret Δ_3 as the empty set.

6. The negative weight vector already has first entry $(b;\ )$. Append after the semicolon the positive weight vectors for Δ_2 and Δ_3, which are defined using the algorithm in Proposition 2.1.1. The union of vector entries in this definition is an unordered union with repetitions.

Once the negative weight vector has been computed for a convex toric domain X_Δ, say that it is $(b; a_1, ..., a_k)$. Then we can algorithmically compute the ECH capacities using the sequence subtraction operation given in definition 1.0.5.

$$c_*(X_\Delta) = c_*(B(b)) - c_*\left(\bigsqcup_{j=1}^{k} B(a_j)\right).$$

Futhermore, recall that we can view the ECH capacities of a disjoint union as a sequence sum, as in definition 1.0.4. Combining this with the above formula gives

$$c_*(X_\Delta) = c_*(B(b)) - (c_*(B(a_1)) + ... + c_*(B(a_k))).$$

2.1.4 Realizing the negative weight vector as a ball packing

We wish to reformulate the algorithm of proposition 2.1.3 as a ball packing construction. As in section 2.1.2, we start with a list of deformed triangles, which will be packed into the moment image of a convex toric domain. Consider the two matrices

$$C = \begin{bmatrix} 0 & 1 \\ -1 & -1 \end{bmatrix}$$

and

$$D = \begin{bmatrix} -1 & -1 \\ 1 & 0 \end{bmatrix}$$

that were mentioned in proposition 2.1.3. We prepend the letters C and D to each of the words in the alphabet of words from section 2.1.2. This gives two more alphabets

$$\{CI, CA, CB, CAA, CAB, CBA, CBB, CAAA, ...\}$$

and

$$\{DI, DA, DB, DAA, DAB, DBA, DBB, DAAA, ...\}.$$

Each of the words in the above two alphabets represents a product of matrices that can act on the standard triangle $\Delta(a)$ to give a new deformed triangle, which is then translated.

The convex region Ω can be packed with these deformed triangles to give a constructive (as opposed to a deconstructive) algorithm for computing the negative weight vector.

Since we have already parameterized several of the deformed triangles in the concave case, we shall save some time by performing both of the following two procedures separately to each of the deformed triangles of proposition 2.1.2:

$$\text{act by the matrix } C, \text{ then translate by the vector } \begin{bmatrix} 0 \\ b \end{bmatrix},$$

and

$$\text{act by the matrix } D, \text{ then translate by the vector } \begin{bmatrix} b \\ 0 \end{bmatrix}.$$

This produces a new deformed triangle, which we denote by Λ^C or Λ^D, in order to distinguish from the deformed triangles in section 2.1.2. Specifically, if we act on the translated, deformed triangle $\Delta_m^{(j)}$ by the matrix C, and then translate up by b units, we call the resulting triangle $\Lambda_m^{C,(j)}$, we denote the size of this deformed triangle by $a_m^{C,(j)}$, and so on.

Proposition 2.1.4. (The negative weight vector as a ball-packing) Let Δ be a convex polygon in standard position with $\zeta, \eta > 0$ (from the definition of convex polygon), or let Δ be a standard triangle. This will be an iterative algorithm, broken into several parts.

m=0 iteration:

1. If Δ is a *standard triangle* of size b, then the negative weight vector is $(b; 0)$, and this algorithm terminates.

2. Otherwise, record the largest value of b such that $\Delta(b)$ circumscribes Δ.

3. Record the largest value of $a_0^{C,(1)}$ such that the translated, deformed triangle

$$\Lambda_0^{C,(1)} = \left\{ \begin{bmatrix} 0 \\ b \end{bmatrix}, \begin{bmatrix} 0 \\ -a_0^{C,(1)} + b \end{bmatrix}, \begin{bmatrix} a_0^{C,(1)} \\ -a_0^{C,(1)} + b \end{bmatrix} \right\}$$

   (given here as a list of its vertices) lies within the region $\Delta(b) \setminus Int(\Delta)$.

4. Record the largest value of $a_0^{D,(1)}$ such that the translated, deformed triangle

$$\Lambda_0^{D,(1)} = \left\{ \begin{bmatrix} b \\ 0 \end{bmatrix}, \begin{bmatrix} -a_0^{D,(1)} + b \\ 0 \end{bmatrix}, \begin{bmatrix} -a_0^{D,(1)} + b \\ a_0^{D,(1)} \end{bmatrix} \right\}$$

   lies within the region $\Delta(b) \setminus Int(\Delta)$.

m=1 iteration:

5. Record the largest value of $a_1^{C,(1)}$ such that the translated, deformed triangle

$$\Lambda_1^{C,(1)} = \left\{ \begin{bmatrix} a_0^{C,(1)} \\ -a_0^{C,(1)} + b \end{bmatrix}, \begin{bmatrix} a_0^{C,(1)} + a_1^{C,(1)} \\ -a_0^{C,(1)} - a_1^{C,(1)} + b \end{bmatrix}, \begin{bmatrix} a_0^{C,(1)} - a_1^{C,(1)} \\ -a_0^{C,(1)} + b \end{bmatrix} \right\}$$

lies within the region $\Delta(b) \setminus Int(\Delta)$.

6. Record the largest value of $a_1^{D,(2)}$ such that the translated, deformed triangle

$$\Lambda_1^{D,(1)} = \left\{ \begin{bmatrix} -a_0^{D,(1)} + b \\ a_0^{D,(1)} \end{bmatrix}, \begin{bmatrix} -a_0^{D,(1)} - a_1^{D,(2)} + b \\ a_0^{D,(1)} + a_1^{D,(2)} \end{bmatrix}, \begin{bmatrix} -a_0^{D,(1)} + b \\ a_0^{D,(1)} - a_1^{D,(2)} \end{bmatrix} \right\}$$

lies within the region $\Delta(b) \setminus Int(\Delta)$.

m=2 iteration:

7. Record the largest value of $a_2^{C,(1)}$ such that the translated, deformed triangle

$$\Lambda_2^{C,(1)} = \left\{ \begin{bmatrix} a_0^{C,(1)} + a_1^{C,(1)} \\ -a_0^{C,(1)} - a_1^{C,(1)} + b \end{bmatrix}, \begin{bmatrix} a_0^{C,(1)} + a_1^{C,(1)} + a_2^{C,(1)} \\ -a_0^{C,(1)} - a_1^{C,(1)} - a_2^{C,(1)} + b \end{bmatrix}, \right.$$

$$\left. \begin{bmatrix} a_0^{C,(1)} + a_1^{C,(1)} - 2a_2^{C,(1)} \\ -a_0^{C,(1)} - a_1^{C,(1)} + a_2^{C,(1)} + b \end{bmatrix} \right\}$$

lies within the region $\Delta(b) \setminus Int(\Delta)$.

8. Record the largest value of $a_1^{D,(2)}$ such that the translated, deformed triangle lies within the region $\Delta(b) \setminus Int(\Delta)$.

mth iteration:

9. Continue this procedure as long as necessary to compute $a_m^{C,(j)}$ and $a_m^{D,(j)}$ for $m \geq 3$. As mentioned above, one generates the triangle $\Lambda_m^{C,(j)}$ by acting on the deformed triangle $\Delta_m^{(j)}$ with the matrix C, then translating by the vector $(0, b)$. One similarly generates the triangle $\Lambda_m^{D,(j)}$ by acting on the deformed triangle $\Delta_m^{(j)}$ with the matrix D, then translating by the vector $(b, 0)$. The reader may have noticed that some triangles were omitted from this recursive construction. This is because the convex polygon Δ is assumed to be bounded by the graph of a non-increasing function (whose graph passes the vertical line test). Some of the triangles Λ^C and Λ^D would violate this property of the graph, hence cannot be packed into the region $\Delta(b) \setminus \Delta$. If

the reader uses a more general notion of convex toric domain (such as the definition given in [CG19]) then all triangles Λ^C, Λ^D may be used.

Notice that this algorithm packs with triangles the region within the standard triangle of size b, but outside the graph of the function g (which bounds Δ). Since we assumed that g bounds some area (i.e. $\zeta > 0$ in Definition 2.0.2), we must have that the area of the standard triangle of size b is at least the area of Δ. Provided Δ is not itself a standard triangle, this inequality of areas must be strict. Recapitulating, when Δ is not a standard triangle, we have

$$\frac{b^2}{2} - area(\Delta) > 0.$$

Furthermore, this region $\Delta(b) \setminus \Delta$ is packed by triangles of positive area, which gives the tail of the negative weight vector. Let us say that the entire negative weight vector is $(b; a_1, ..., a_k)$. Area considerations further imply that

$$b^2 > a_1^2 + ... + a_k^2, \tag{2.1.5}$$

because the triangle of size b is packed with triangles of sizes $a_1, ..., a_k$. The inequality is strict because Δ has positive area. This result (2.1.5) is called the volume condition on negative weight vectors.

Example 2.1.3. We present an example computation of the ball-packing algorithm in proposition 2.1.4. Consider the concave polygon that is bounded by the piecewise-linear function

$$g\colon [0, 27] \to [0, 27]$$

given by

$$g(x) = \begin{cases} 27 - \frac{x}{4} & \text{if } 0 \leq x \leq 4 \\ \frac{138}{5} - \frac{2x}{5} & \text{if } 4 \leq x \leq 9 \\ \frac{147}{5} - \frac{3x}{5} & \text{if } 9 \leq x \leq 14 \\ \frac{63}{2} - \frac{3x}{4} & \text{if } 14 \leq x \leq 18 \\ 42 - \frac{4x}{3} & \text{if } 18 \leq x \leq 21 \\ 49 - \frac{5x}{3} & \text{if } 21 \leq x \leq 24 \\ 69 - \frac{5x}{2} & \text{if } 24 \leq x \leq 26 \\ 108 - 4x & \text{if } 26 \leq x \leq 27 \end{cases}$$

The polygon bounded by the graph of g and the x- and y-axes is shown in figure 2.6. That figure also shows the deformed triangles that fill the region exterior to the polygon

and inside the standard triangle $\Delta(36)$. The negative weight vector is

$$(36; 9, 9, 6, 6, 2, 2, 2, 2, 1, 1, 1, 1, 1, 1, 1, 1).$$

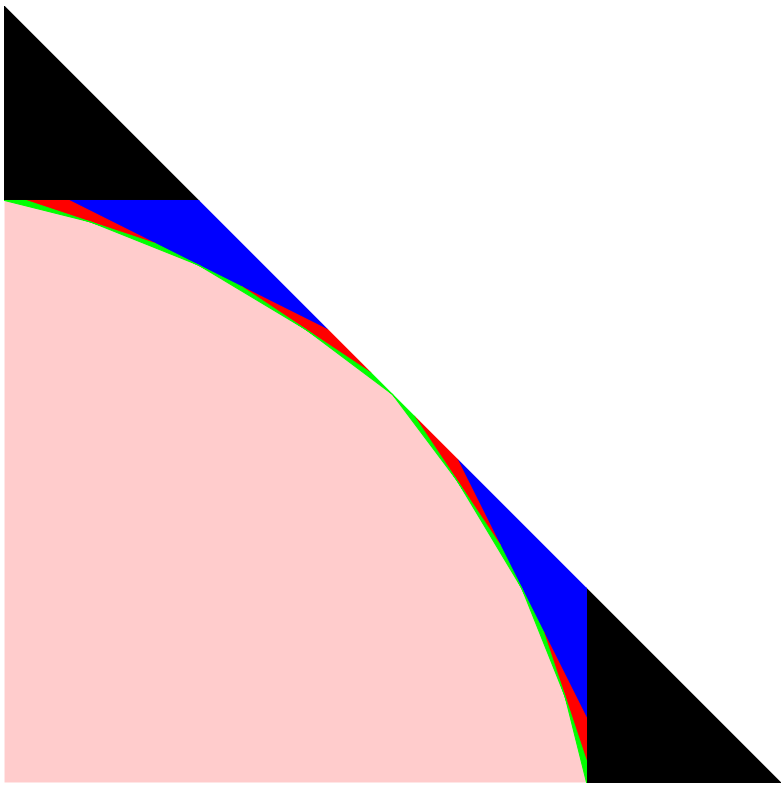

Figure 2.6: An example of a convex toric domain.

2.2 From Concave to Convex

In this section we give a procedure for converting a concave toric domain into a convex toric domain. If one already knows the positive weight vector of the given concave domain, one can easily find the negative weight vector of the convex domain that results from this procedure. We have already noted that an ellipsoid is both concave and convex. The procedure that follows will convert an ellipsoid into the same ellipsoid. We examine this situation further in example 2.2.3, below. Let Ω_c be a concave region in the first quadrant, as described in definition 2.0.2.

We are assuming that Ω_c is bounded by a function $f\colon [0, \zeta] \to [0, \eta]$ such that $f(0) = \eta$ and $f(\zeta) = 0$. We also assume that $f' < 0$ and $f'' \geq 0$, whenever these derivatives exist.

Since we need to operate on this graph by linear transformations, it will be convenient to parameterize the graph as a pair $(t, f(t))$ with domain $0 \leq t \leq \zeta$.

It is immediate from the definition of f that

$$\min_{[0,\zeta']} t = 0, \qquad \text{for any } 0 \leq \zeta' \leq \zeta, \tag{2.2.1}$$

and this minimum obviously occurs at $t = 0$. Furthermore, since f is decreasing,

$$\min_{[\zeta',\zeta]} f(t) = 0, \qquad \text{for any } 0 \leq \zeta' \leq \zeta, \tag{2.2.2}$$

and this minimum obviously occurs at $t = \zeta$.

Proposition 2.2.1. Let Ω_c be a concave region in the first quadrant that is bounded by the graph of $f\colon [0, \zeta] \to [0, \eta]$, as described in definition 2.0.2. Let Ω_v be the region in the first quadrant that is bounded by the axes and by the parametric curve $(\zeta - t, \eta - f(t))$ for $0 \leq t \leq \zeta$. Then Ω_v will be a convex region, as defined in definition 2.0.2. If Ω_c was a polygon, then Ω_v will be a polygon.

Proof. As mentioned above, we could view the concave region Ω_c as bounded by a parametric curve $(t, f(t))$ for $0 \leq t \leq \zeta$. This parametric curve begins at the point $(0, \eta)$ when $t = 0$, and it ends at the point $(\zeta, 0)$ when $t = \zeta$. We compare this to the parametric curve $(\zeta - t, \eta - f(t))$ for $0 \leq t \leq \zeta$. This parametric curve starts at $(\zeta, 0)$ when $t = 0$, and it ends at $(0, \eta)$ when $t = \zeta$. This computation shows that the curve $(\zeta - t, \eta - f(t))$ moves from right to left in the first quadrant.

In addition, the graph that defines Ω_v is just the graph of f rotated $180°$ (or π radians) about the point $(\zeta/2, \eta/2)$, which is the midpoint of the line through $(0, \eta)$ and $(\zeta, 0)$. This rotation reverses the concavity of the graph, making Ω_v a convex region.

The final claim in the proposition is clear, because if f is piecewise-linear, then the parametric curve $(\zeta - t, \eta - f(t))$ will be piecewise-linear. □

We defined our concave region using a function of a single variable, whose graph we then parameterized. A similar procedure as above will work on a concave region that is bounded by a more general parametric curve. This is made explicit in the following.

Proposition 2.2.2. Suppose that Ω_c is a concave region in the first quadrant that is bounded by the axes and a parametric curve segment $(p(t), q(t))$ with domain $0 \leq t \leq T$. We assume that parametric curve starts at the point $(0, \eta)$ when $t = 0$ and ends at the point $(\zeta, 0)$ when $t = T$. There is a corresponding convex region Ω_v in the first quadrant that is

bounded by the axes and parametric curve $(\zeta - p(t), \eta - q(t))$. If Ω_c was a polygon, then Ω_v will be a polygon.

The proof is entirely the same as for proposition 2.2.1.

Returning to proposition 2.2.1, we can re-parameterize the curve $(\zeta - t, \eta - f(t))$ by making the substitution $t \mapsto \zeta - s$ for $0 \leq s \leq \zeta$. This gives a new parameterization

$$(s, \eta - f(\zeta - s)), \quad 0 \leq s \leq \zeta,$$

which travels from left to right and has the benefit of resembling the graph of a function of a single variable. We summarize this in the following.

Corollary 2.2.1. Let Ω_c be a concave region in the first quadrant that is bounded by the graph of $f\colon [0, \zeta] \to [0, \eta]$, as described in definition 2.0.2. Let Ω_v be the region in the first quadrant that is bounded by the axes and by the graph of $g\colon [0, \zeta] \to [0, \eta]$ given by

$$g(x) = \eta - f(\zeta - x).$$

Then Ω_v will be a convex region, as defined in definition 2.0.2. Furthermore $g(\zeta) = 0$. If Ω_c was a polygon, then Ω_v will be a polygon.

The reason for converting a concave region Ω_c into a convex region Ω_v is that it gives us a useful computational tool.

Theorem 2.2.1. Let Ω_c and Ω_v be as described in proposition 2.2.1 or proposition 2.2.2. Say that the positive weight expansion of Ω_c is given by the (possibly finite) list

$$(a_1, a_2, a_3, ..., a_N, ...),$$

where, in this expansion, we have included terms which are zero. Then the negative weight expansion of Ω_v is given by the (possibly finite) list

$$(b; b_2, b_3, ...) := (\zeta + \eta - a_1; \eta - a_1, \zeta - a_1, a_2, a_3, ..., a_N, ...).$$

Note that some entries of this negative weight expansion may be omitted for reasons of symmetry, and this situation will be discussed in the proof.

Proof. We use the parameterization $(t, f(t))$ for the boundary of Ω_c and the parameterization $(\zeta - t, \eta - f(t))$ for the boundary of Ω_v, both of which are described in proposition 2.2.1. We describe the computation of both relevant weight expansions in several steps.

Step 1: The first step in computing the positive weight expansion for Ω_c is to find the size of the largest standard triangle inscribed within Ω_c. This amounts to computing

$$a_1 := \min_{[0,\zeta]}\{t + f(t)\}. \tag{2.2.3}$$

By our description of the algorithm in proposition 2.1.1, this minimum occurs at the input $t = x_2$.

On the other hand, the first step in computing the negative weight expansion for Ω_v is to find the size of the smallest standard triangle that circumscribes Ω_v. This amounts to computing

$$b := \max_{[0,\zeta]}\{\zeta - t + \eta - f(t)\} = \zeta + \eta + \max_{[0,\zeta]}\{-t - f(t)\} = \zeta + \eta - a_1. \tag{2.2.4}$$

(In the last step of (2.2.4), we converted the maximum to a minimum using the fact that $x \mapsto -x$ is a homeomorphism $\mathbb{R} \to \mathbb{R}$.) By construction, this minimum occurs at $t = x_2$.

Step 2: Next, we need to find the second entry, a_2, of the positive weight expansion of Ω_c. Look back to the algorithm presented in proposition 2.1.1. We shall assume that $x_2 \neq 0$ in that algorithm, so that $a_2 \neq 0$. (If it is the case that $x_2 = 0$, then steps 2 and 4 in this proof can be omitted. The negative weight vector will not have the term $\zeta - a_1$ and the fourth term will be $a_2 = 0$.) The term a_2 is computed by first shifting the graph of $(t, f(t))$, $0 \leq t \leq x_2$, downward by a_1 units, then acting on this shifted graph by a linear transformation that results in a new parameterized curve, $(t, t + f(t) - a_1)$, $0 \leq t \leq x_2$. Finding the size of the largest standard triangle inscribed in this curve amounts to computing

$$a_2 := \min_{[0,x_2]}\{2t + f(t) - a_1\} = \min_{[0,x_2]}\{2t + f(t)\} - a_1. \tag{2.2.5}$$

We shall use this result in a future step, but for now, let us say that the minimum in equation (2.2.5) occurs at time $t = t_2$.

On the other hand, the second step in computing the negative weight expansion for Ω_v is to shift the graph of $(\zeta - t, \eta - f(t))$ left by b units and act by a linear transformation, which results in the graph of the parametric curve

$$(\eta - f(t), -\zeta + t + b - \eta + f(t)) = (\eta - f(t), t + f(t) - a_1).$$

We have already noted that the graph that defines Ω_v is traced out from right to left. The linear transformation that we apply in this step again reverses the orientation of the graph. Luckily, this will cause a minimum to occur on the domain $[0, x_2]$, which is exactly what

we need for the computation of

$$b_2 := \min_{[0,x_2]}\{\eta - f(t) + t + f(t) - a_1\} = \min_{[0,x_2]}\{t + \eta - a_1\} = \eta - a_1, \tag{2.2.6}$$

using equations (2.2.1) and (2.2.4). By the results of (2.2.1), this minimum occurs at $t = 0$. The reverse in graph orientation is also why we perform these computations in a slightly different order than is described in proposition 2.1.3.

Step 3: Next, we need to find the third entry, a_3, of the positive weight expansion of Ω_c. Look back to the algorithm presented in proposition 2.1.1. We shall assume that $x_3 \neq a_1$ in that algorithm, so that $a_3 \neq 0$. (If it is the case that $x_3 = a_1$, then steps 3 and 5 in this proof can be omitted.) The term a_3 is computed by first shifting the graph of $(t, f(t))$ left by a_1 units, and acting on this shifted graph by a linear transformation that results in a new parameterized curve $(t - a_1 + f(t), f(t))$. Finding the size of the largest standard triangle inscribed in this curve amounts to computing

$$a_3 := \min_{[x_3,\zeta]}\{2f(t) + t - a_1\} = \min_{[x_3,\zeta]}\{2f(t) + t\} - a_1. \tag{2.2.7}$$

We will use this computation in a later step, but for now let us say that the minimum in equation (2.2.7) occurs at time $t = t_3$.

On the other hand, the third step in computing the negative weight expansion for Ω_v involves shifting the graph of $(\zeta - t, \eta - f(t))$ downward by b units and acting by a linear transformation which results in the graph of the parametric curve

$$(-\zeta + t - \eta + f(t) + b, \zeta - t), \qquad x_3 \leq t \leq \zeta.$$

Again by the reverse in orientation, this graph is traced out from left to right on the domain $[x_3, \zeta]$. Finding the size of the largest standard triangle inscribed in this curve amounts to computing

$$b_3 := \min_{[x_3,\zeta]}\{-\zeta + t - \eta + f(t) + b + \zeta - t\} = \min_{[x_3,\zeta]}\{-\eta + f(t) + b\} = \zeta - a_1, \tag{2.2.8}$$

using equations (2.2.2) and (2.2.4). Also by (2.2.2), this minimum occurs at $t = \zeta$.

Step 4: For the remaining steps in this proof, we focus only on computing the negative weight expansion of Ω_v, since that is the purpose of this theorem. At this point in the negative weight algorithm, all domains under consideration are actually concave. So the remainder of this proof will explain why certain terms in the negative weight expansion of Ω_v reduce to terms in the positive weight expansion of Ω_c. It is important to note that the

minimum of (2.2.6) occurs at $t = 0$. This will imply that the region considered in step 2 will not be split into two pieces by the intersection of the graph of $(\eta - f(t), -\zeta + t + b - \eta + f(t))$ with the standard triangle $\Delta(b_2)$. Instead, there is only one region (under this graph) that must be transformed in this step of the iterative procedure. We must shift the graph leftward by $b_2 = \eta - a_1$ units and act by a linear transformation that results in the graph of the curve

$$(a_1 + t + b - \zeta - \eta, t + b - \zeta - \eta + f(t)), \qquad 0 \leq t \leq x_2.$$

Finding the size of the largest standard triangle inscribed within the region beneath this graph amounts to computing

$$b_4 := \min_{[0,x_2]}\{t + t - a_1 + f(t)\} = \min_{[0,x_2]}\{2t + f(t)\} - a_1 =: a_2.$$

As in equation (2.2.5), this minimum occurs at $t = t_2$.

Step 5: Here we note that the minimum of (2.2.8) occurs at $t = \zeta$. This will imply that the region considered in step 3 will not be split into two pieces by the intersection of the graph of $(-a_1 + t + f(t), \zeta - t)$ with the standard triangle $\Delta(b_3)$. Instead, there is only one region (under this graph) that must be transformed in this step of the iterative procedure. We must shift the graph downward by $b_3 = \zeta - a_1$ units and act by a linear transformation that results in the graph of the curve

$$(-a_1 + t + f(t), f(t)), \qquad x_3 \leq t \leq \zeta.$$

Finding the size of the largest standard triangle inscribed within the region beneath this graph amounts to computing

$$b_5 := \min_{[x_3,\zeta]}\{-a_1 + t + 2f(t)\} = \min_{[x_3,\zeta]}\{t + 2f(t)\} - a_1 =: a_3.$$

As in equation (2.2.7), this minimum occurs at $t = t_3$.

Conclusion: In steps 4 and 5 of this proof, the computation of b_4 and b_5 for the convex toric domain Ω_v reduced exactly to the computation of a_2 and a_3 for the concave toric domain Ω_c. Furthermore, the minima occurred at exactly the same arguments. For this reason, the computation of further b_i in the negative weight vector must reduce exactly to the computation of some a_j in the positive weight vector. This is what we intended to show. □

Remark. The reader should pay close attention to when certain entries in the positive weight vector of Ω_c vanish, because this will cause entries in the negative weight vector of Ω_v to

change or be omitted from the computation. This situation can be avoided by slightly perturbing the function that defines the outer boundary of Ω_c, if necessary.

Example 2.2.1. Let Ω_c be the concave region that is described in example 2.1.2. The moment image is a polygon that is pictured in figure 2.5. In that example, we computed the positive weight vector of Ω_c to be

$$(18, 6, 6, 2, 2, 2, 2, 2, 1, 1, 1, 1, 1, 1, 1, 1),$$

and this polygon was designed so that none of the first fifteen entries of the weight vector vanish. (In other words, none of the triangles in the triangle packing of figure 2.5 have area zero.) Furthermore, the function $f\colon [0, 27] \to [0, 27]$ that defines the outer boundary of Ω_c has intercepts $\zeta = 27 = \eta$.

By applying proposition 2.2.1 (or, equivalently, corollary 2.2.1) we obtain a convex toric domain Ω_v, which has moment image given by the polygon in figure 2.6. We can now apply theorem 2.2.1 to quickly and easily compute the negative weight vector of Ω_v to be

$$(27 + 27 - 18; 27 - 18, 27 - 18, 6, 6, 2, 2, 2, 2, 2, 1, 1, 1, 1, 1, 1, 1, 1).$$

This exactly coincides with our earlier computation in example 2.1.3.

Example 2.2.2. The Lagrangian bidisc is the subset of $\mathbb{R}^4$ given by

$$L_B = \left\{ (x_1, y_1, x_2, y_2) \in \mathbb{R}^4 \,\middle|\, x_1^2 + x_2^2 \leq 1,\ y_1^2 + y_2^2 \leq 1 \right\}.$$

Notice that the Lagrangian bidisc is not the same as the symplectic polydisc of definition 1.0.3. It was proved in [Ram17] that the Lagrangian bidisc L_B is symplectomorphic to the toric domain with moment image bounded by the axes and the graph of the parametric curve

$$(2\sin(t/2) - t\cos(t/2),\ 2\sin(t/2) + (2\pi - t)\cos(t/2)), \qquad 0 \leq t \leq 2\pi.$$

This curve starts at the point $(0, 2\pi)$ when $t = 0$ and ends at the point $(2\pi, 0)$ when $t = 2\pi$. For this reason we set $\zeta = \eta = 2\pi$. Let us call the concave region bounded by this curve Ω_c. Here we compute the first fifteen entries of the weight sequence of Ω_c, using the notation

of proposition 2.1.2. (Alternatively, these entries could be named a_1 through a_{15}.)

$$a_0^{(1)} = 4$$

$$a_1^{(1)} = 3\sqrt{3} - 4 \qquad a_1^{(2)} = 3\sqrt{3} - 4$$

$$a_2^{(1)} = 4\sqrt{2} - 3\sqrt{3} \qquad a_2^{(2)} = -4 - 3\sqrt{3} + 10\sqrt{\frac{5}{8} + \frac{\sqrt{5}}{8}}$$

$$a_2^{(3)} = -4 - 3\sqrt{3} + 10\sqrt{\frac{5}{8} + \frac{\sqrt{5}}{8}} \qquad a_2^{(4)} = 4\sqrt{2} - 3\sqrt{3}$$

$$a_3^{(1)} = -4\sqrt{2} + 10\sqrt{\frac{5}{8} - \frac{\sqrt{5}}{8}} \qquad a_3^{(2)} = -4\sqrt{2} - 3\sqrt{3} + 14\cos\left(\frac{3\pi}{14}\right)$$

$$a_3^{(3)} = -3\sqrt{3} - 5\sqrt{\frac{1}{2}\left(5 + \sqrt{5}\right)} + 16\cos\left(\frac{\pi}{8}\right)$$

$$a_3^{(4)} = -4 - 5\sqrt{\frac{1}{2}\left(5 + \sqrt{5}\right)} + 14\cos\left(\frac{\pi}{14}\right)$$

$$a_3^{(5)} = -4 - 5\sqrt{\frac{1}{2}\left(5 + \sqrt{5}\right)} + 14\cos\left(\frac{\pi}{14}\right)$$

$$a_3^{(6)} = -3\sqrt{3} - 5\sqrt{\frac{1}{2}\left(5 + \sqrt{5}\right)} + 16\cos\left(\frac{\pi}{8}\right)$$

$$a_3^{(7)} = -4\sqrt{2} - 3\sqrt{3} + 14\cos\left(\frac{3\pi}{14}\right) \qquad a_3^{(8)} = -4\sqrt{2} + 10\sqrt{\frac{5}{8} - \frac{\sqrt{5}}{8}}$$

With these computations, we can use proposition 2.1.2 to visualize the positive weight vector as a triangle packing of the moment image. This triangle packing is pictured in figure 2.7. Notice that the triangles do not fully fill the area of the moment image. That is because the moment image is not a rational concave polygon, which implies that the sequence of weights should be infinite. The figure shows the triangle packing given by only the first fifteen weights, computed above. Next, we can apply proposition 2.2.2 to the concave region Ω_c to generate a convex region Ω_v, which is bounded by the coordinate axes and the graph of the parametric curve

$$\left(2\pi - 2\sin(t/2) + t\cos(t/2),\ 2\pi - 2\sin(t/2) - (2\pi - t)\cos(t/2)\right), \qquad 0 \le t \le 2\pi.$$

Some entries of the negative weight vector of the convex region Ω_v can be easily computed using theorem 2.2.1.

$$(4\pi - 4; 2\pi - 4, 2\pi - 4, ...),$$

and after the entry $2\pi - 4$, the negative weight vector coincides with entries $a_1^{(1)}$ onwards

of the positive weight vector above.

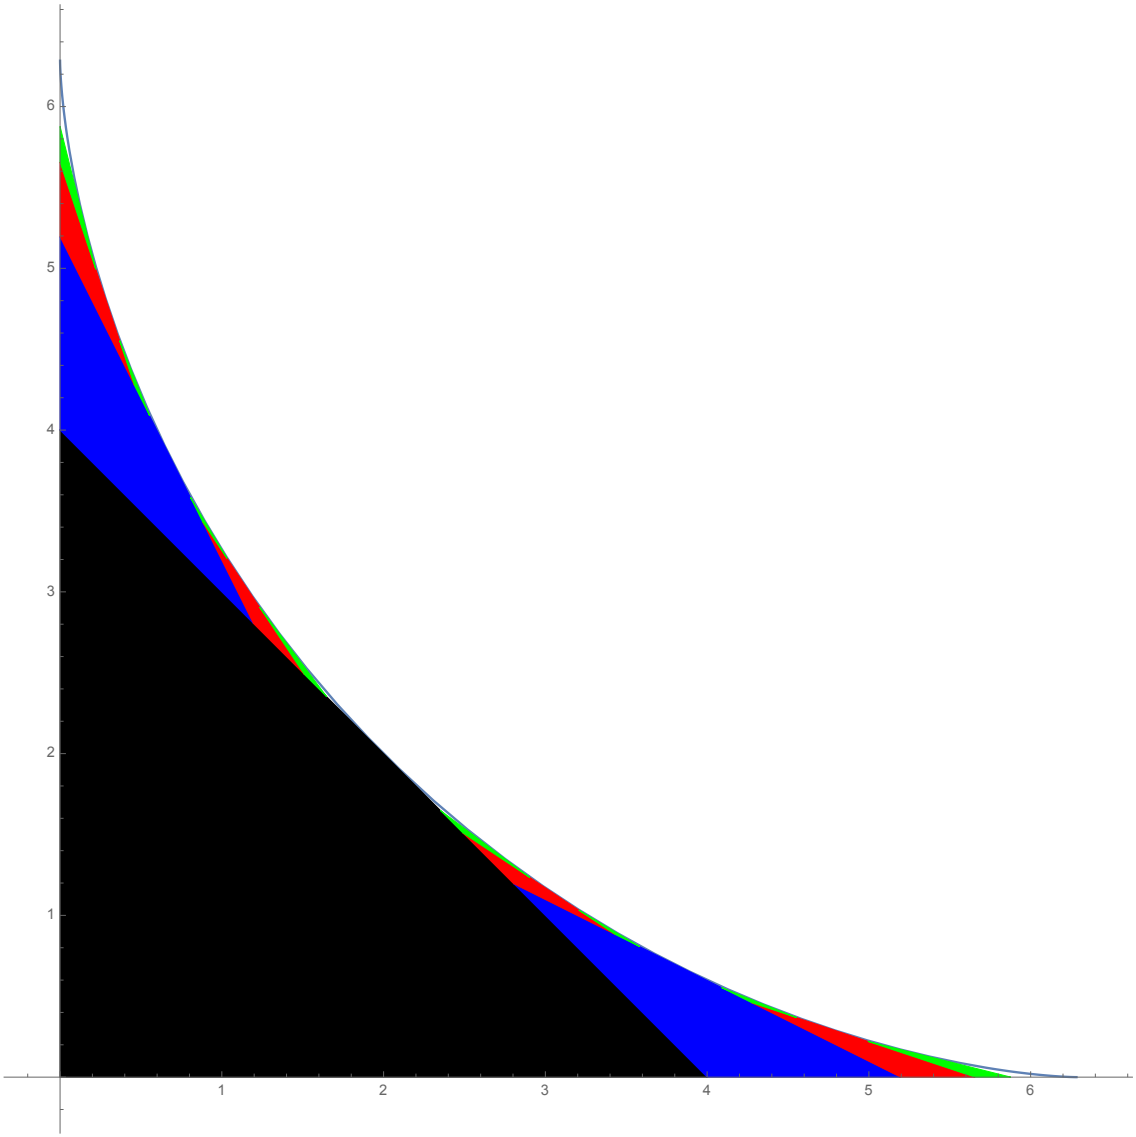

Figure 2.7: The triangle packing that gives the positive weight vector for the Lagrangian bidisc, as described in example 2.2.2. The outermost curve is the graph of $(2\sin(t/2) - t\cos(t/2),\ 2\sin(t/2) + (2\pi - t)\cos(t/2))$.

Remark. Let $q \geq 1$ be rational. The ellipsoid $E(1, q)$ can be treated as either a concave toric domain or a convex toric domain, and we examine this further in the following example. First we consider what this concave-convex duality says about ball packing problems. The moment image of $E(1, q)$ is depicted in figure 2.8 as the region below the linear function $f\colon [0, 1] \to [0, q]$, given by $f(x) = q - qx$. This line segment connects the point $(0, q)$ with the point $(1, 0)$. The other line in figure 2.8 is the graph of $g\colon [0, q] \to [0, q]$, $g(x) =$

$q - x$. This line segment connects $(0, q)$ to $(q, 0)$. The region below g represents the moment image of $E(q, q) = B(q)$, which is obviously the smallest standard triangle that circumscribes the graph of f. The fact that $E(1, q)$ is both a concave and a convex toric domain implies the following about packing triangles: one can pack the region below the graph of f with triangles if and only if one can pack the region below the graph of g with triangles that avoid the (pink) shaded region. This is geometrically very obvious, but it gives a non-obvious result about symplectic ball packings of the ellipsoid. The (pink) shaded region in figure 2.8 is affine equivalent to $E(q - 1, q)$, which obviously includes into $E(q, q) = B(q)$. In terms of ball packings, we have the following.

Proposition 2.2.3. Let $\mathcal{B}$ be any disjoint union of symplectic balls.

$$\mathcal{B} \hookrightarrow E(1, q) \quad \Longleftrightarrow \quad \mathcal{B} \hookrightarrow (B(q) \setminus E(q - 1, q)).$$

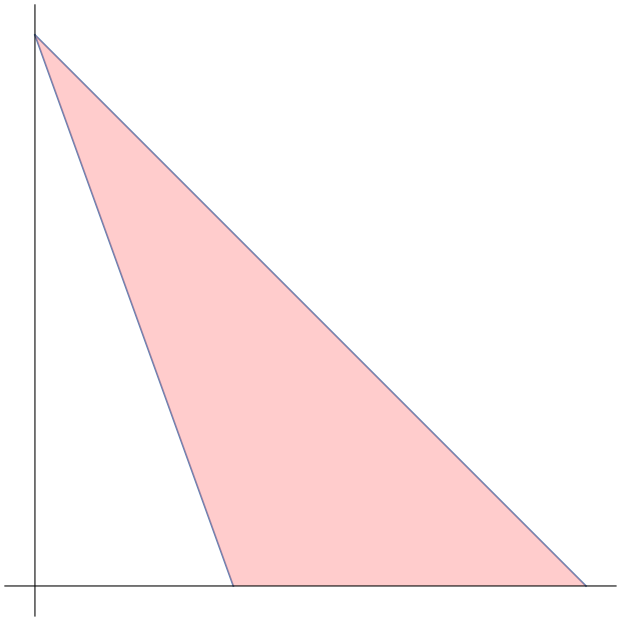

Figure 2.8: The ellipsoid $E(1, q)$ can be viewed as either a concave toric domain or a convex toric domain.

Example 2.2.3. Let Ω_c be the ellipsoid $E(1, q)$, with $q \geq 1$ rational and in lowest terms. The moment image of $E(1, q)$ is a triangle with intercepts $\zeta = 1$ and $\eta = q$. The positive weight vector of $E(1, q)$ is computed in many sources, including [MS12b]. It begins with entries

$$(\underbrace{1, ..., 1}_{\lfloor q \rfloor \text{ copies}}, \underbrace{q - \lfloor q \rfloor, ..., q - \lfloor q \rfloor}_{\lfloor 1/(q - \lfloor q \rfloor) \rfloor \text{ copies}}, ...).$$

As described in example 2.1.1, the entries are given by the Farey expansion of q, and this vector has finite length. If $q - \langle q \rangle = 0$, these entries are omitted.

Applying proposition 2.2.1 to this toric domain gives the exact same ellipsoid, $E(1, q)$, which we now treat as a convex toric domain.

$$(q; q-1, \underbrace{1, ..., 1}_{\lfloor q \rfloor - 1 \text{ copies}}, \underbrace{q - \lfloor q \rfloor, ..., q - \lfloor q \rfloor}_{\lfloor 1/(q - \lfloor q \rfloor) \rfloor \text{ copies}}, ...).$$

Since these two weight vectors correspond to the same toric domain, they must give the exact same sequence of ECH capacities. This proves the following.

Corollary 2.2.2. Let $q \geq 1$ be rational and written in lowest terms. The (finite length) positive weight vector

$$(\underbrace{1, ..., 1}_{\lfloor q \rfloor \text{ copies}}, \underbrace{q - \lfloor q \rfloor, ..., q - \lfloor q \rfloor}_{\lfloor 1/(q - \lfloor q \rfloor) \rfloor \text{ copies}}, ...) \tag{2.2.9}$$

generates the same sequence of ECH capacities as the (finite length) negative weight vector

$$(q; q-1, \underbrace{1, ..., 1}_{\lfloor q \rfloor - 1 \text{ copies}}, \underbrace{q - \lfloor q \rfloor, ..., q - \lfloor q \rfloor}_{\lfloor 1/(q - \lfloor q \rfloor) \rfloor \text{ copies}}, ...). \tag{2.2.10}$$

The vector (2.2.10) comes from the vector (2.2.9) by applying theorem 2.2.1.

CHAPTER 3

Blowup Forms on $\mathbb{C}\mathrm{P}^2$

In this chapter we continue to consider symplectic 4-manifolds. As suggested by proposition $1.0.1$, convex toric domains are a convenient class of examples to consider as the target of a symplectic embedding. In this chapter, we will probe the embedding properties encoded by negative weight vectors by considering embeddings of balls into convex toric domains. To do this, we define the notion of a blowup form on $\mathbb{C}\mathrm{P}^2$, and we compare and contrast with the constructions of the previous chapter.

Definition 3.0.1. Let $\mathbb{C}\mathrm{P}^2(\mu)$ denote the complex projective plane, equipped with the Fubini-Study symplectic form, normalized so that the area of a line is $2\pi\mu$. We can blow up $\mathbb{C}\mathrm{P}^2(\mu)$ to obtain a new symplectic manifold, and we say that this blowup is of size ν if the area of the exceptional divisor is $2\pi\nu$

The main dictionary between the ball-packing problems of chapter 2 and blowups of $\mathbb{C}\mathrm{P}^2$ is given by the following result of McDuff and Polterovich.

Proposition 3.0.1. (McDuff and Polterovich, [MP94]) There exists a symplectic embedding of balls $B^4(w) \hookrightarrow B^4(\mu)$ if and only if there exists a symplectic form on $\mathbb{C}\mathrm{P}^2\#\overline{\mathbb{C}\mathrm{P}^2}$ in the class dual to $\mu L - wE$ (where L denotes the line class and E denotes the class of the exceptional divisor). More generally, let $\mathcal{B}$ denote a finite disjoint union of k symplectic 4-balls. There exists a symplectic embedding $\mathcal{B} \hookrightarrow B^4(\mu)$ if and only if there exists a symplectic form on the k-fold blowup of $\mathbb{C}\mathrm{P}^2$ in the class dual to $\mu L - \sum a_i E_i$ (where E_i denote the exceptional divisors and the a_i are given by the radii of the balls in the collection $\mathcal{B}$).

Because of this result, we should examine symplectic forms on blowups of $\mathbb{C}\mathrm{P}^2$. Let us write $\omega_{(b;a_1,\ldots,a_k)}$ for the symplectic form that is obtained from $\mathbb{C}\mathrm{P}^2(b)$ by blowups of size a_1 through a_k. Now, we shall momentarily suppress mention of the symplectic form, and consider the underlying manifold.

Let $M_k = \mathbb{C}\mathrm{P}^2 \# k\overline{\mathbb{C}\mathrm{P}^2}$ denote an k-fold blowup of $\mathbb{C}\mathrm{P}^2$. The homology $H_2(M_k, \mathbb{Z})$ is easily computed. Let us name the canonical basis $\{L, E_1, ..., E_k\}$, where $L = [\mathbb{C}\mathrm{P}^1]$ is the image of the homology class of a line in $\mathbb{C}\mathrm{P}^2$ under the inclusion $H_2(\mathbb{C}\mathrm{P}^2) \to H_2(M_k)$ and where $E_1, ..., E_k$ are the homology classes of the exceptional divisors. We assume that each E_j has self-intersection number -1, so that these blowups are generic (no iterated blowups). We also note that L has self-intersection number 1.

Definition 3.0.2. A vector of the form $(b; a_1, ..., a_k)$ is said to *encode* a cohomology class $c \in H^2(M_k)$ if the following conditions hold.

$$\frac{1}{2\pi}\langle c, L\rangle = b,$$

and

$$\frac{1}{2\pi}\langle c, E_j\rangle = a_j \qquad \text{for} \qquad 1 \leq j \leq k.$$

We emphasize that this interpretation of a vector only specifies how a certain cohomology class pairs with a basis for $H_2(M_k)$. It would be nice to have a specific symplectic form in this class. Specifically, we seek a so-called blowup form within the cohomology class encoded by a vector.

Definition 3.0.3. A *blowup form* on M_k is a symplectic form for which there exist pairwise disjoint embedded symplectic spheres in classes $L, E_1, ..., E_k$.

There are three necessary conditions for a vector $(b; a_1, ..., a_k)$ to encode the cohomology class of a blowup form. First, the entries $b, ..., a_k$ of this vector should be positive, since they represent symplectic areas. Note that any instance of $a_j = 0$ could be ignored, but we sometimes want to pad out the tail of such a vector with zeroes. This non-negativity condition will make the symplectic form ample, in the description below. Second, one obviously cannot blow up $\mathbb{C}\mathrm{P}^2$ by balls whose collective volume exceeds that of the manifold being blown up. This observation imposes the following *volume condition* on vectors vectors representing a blowup form:

$$a_1^2 + ... + a_k^2 < b^2. \tag{3.0.1}$$

Third is the so-called "Gromov inequality," which is that $a_i + a_j < b$ for all $1 \leq i \neq j \leq k$.

The reason we seek such a blowup form is because of proposition 3.0.1, which calls for a form within a certain class. Recall that in proposition 2.2.3 we argued why a symplectic embedding of balls into an ellipsoid is equivalent to a symplectic embedding of balls into a ball, which avoids a certain ellipsoid. An analogous result holds for embedding any concave toric domain into a convex toric domain.

Proposition 3.0.2. (Cristofaro-Gardiner, [CG19]) Let X be a concave toric domain with positive weight vector $(a_1, ..., a_N)$. Let Y be a convex toric domain with negative weight vector $(b; b_1, ..., b_k)$. There exists a symplectic embedding

$$X \hookrightarrow Y$$

if and only if there exists a symplectic embedding

$$\left(\bigsqcup_{i=1}^{N} B^4(a_i)\right) \sqcup \left(\bigsqcup_{j=1}^{k} B^4(b_j)\right) \hookrightarrow B^4(b).$$

Now let us combine the results of proposition 3.0.1 with proposition 3.0.2.

Proposition 3.0.3. Let X be a concave toric domain with positive weight vector $(a_1, ..., a_N)$. Let Y be a convex toric domain with negative weight vector $(b; b_1, ..., b_k)$. There exists a symplectic embedding

$$X \hookrightarrow Y$$

if and only if there exists a blowup form on the $N + k$-fold blowup of $\mathbb{C}\text{P}^2$ in the class represented by the vector $(b; b_1, ..., b_k, a_1, ..., a_N)$.

We shall use this result to translate between symplectic embeddings and cohomology classes on blowups of $\mathbb{C}\text{P}^2$. The purpose of this translation is to prove a result about ECH capacities of convex toric domains. First we must explain a procedure known as the Cremona transform.

3.1 The Cremona Transform

We begin by defining the Cremona transform to be an operation on vectors. Since the vectors appearing in this chapter have a concrete geometrical meaning, we shall shortly explain how to interpret the Cremona transform geometrically.

Definition 3.1.1. Given a vector of the form $(b; a_1, ..., a_k)$, first pad out the tail of the vector with zeroes, if necessary, so that $k \geq 3$. We define the *defect* to be $r = b - a_1 - a_2 - a_3$. The *Cremona transform* of the vector $(b; a_1, ..., a_k)$ is the new vector

$$(b + r; a_1 + r, a_2 + r, a_3 + r, a_4, a_5, ..., a_k).$$

Example 3.1.1. The vector $(4; 2, 2, 1)$ has defect $r = 4 - 2 - 2 - 1 = -1$. The Cremona

transform of this vector is

$$(4-1; 2-1, 2-1, 1-1) = (3; 1, 1, 0).$$

On the other hand, suppose we begin with the vector $(3; 1, 1, 0)$. Then we have defect $r = 3 - 1 - 1 - 0 = 1$. The Cremona transform of this vector is

$$(3+1; 1+1, 1+1, 0+1) = (4; 2, 2, 1).$$

We have computed that $(4; 2, 2, 1) \mapsto (3; 1, 1, 0) \mapsto (4; 2, 2, 1)$ under the operation of Cremona transform. This computation shows that the vectors $(4; 2, 2, 1)$ and $(3; 1, 1, 0)$ are in involution, with respect to the procedure of Cremona transform.

Remark. We now describe a geometrical interpretation of the Cremona transform, in accordance with the interpretation of a vector as encoding a cohomology class, as in definition 3.0.2, above. Note the basis for $H_2(M_k, \mathbb{Z})$ given in that definition. This discussion is adapted from [CGFS17]. Recall that, given a vector u in an inner-produce space, the map

$$r_u(x) = x - 2\frac{\langle x, u\rangle}{\langle u, u\rangle}u,$$

is a geometrical reflection, hence an involution. In a similar way, we choose a class $A \in H_2(M_k, \mathbb{R})$ having $A \cdot A \neq 0$ and we define

$$r_A(B) = B - 2\frac{A \cdot B}{A \cdot A}A.$$

This map will be an involution of $H_2(M_k, \mathbb{R})$. In particular, when $A \cdot A \in \{\pm 1, \pm 2\}$, the map r_A is an automorphism of $H_2(M_k, \mathbb{Z})$. Towards this end, we consider the classes

$$A_{0,0} = L - E_1 - E_2 - E_3$$

and

$$A_{i,j} = E_i - E_j \qquad \text{for} \qquad 1 \leq i < j \leq n.$$

These classes were chosen to have self intersection number -2, meaning that the map $r_{A_{*,*}}$ will be an automorphism. In fact, the formula for $r_{A_{*,*}}$ simplifies to

$$r_{A_{i,j}}(B) = B + (A_{i,j} \cdot B)A_{i,j} \qquad \text{for} \qquad (i, j) = (0, 0) \text{ or } 1 \leq i < j \leq n.$$

In the basis $\{L, E_1, ..., E_n\}$ for $H_2(M_k, \mathbb{Z})$ we can now give a geometrical interpretation to

the maps $r_{A_{*,*}}$.

1. $r_{A_{i,j}}$ for $i, j \neq 0$ has the effect of transposing the classes of two exceptional divisors. Geometrically the class $A_{i,j}$ can be represented by a smoothly embedded J-holomorphic sphere. The smooth Dehn-Seidel twist along this sphere is a diffeomorphism inducing $r_{A_{i,j}} = B + (A_{i,j} \cdot B)A_{i,j}$ with regards to the Picard-Lefschetz formula.

2. $r_{A_{0,0}}$ is the Cremona transform on homology. Geometrically $A_{0,0}$ can also be represented by a smoothly embedded sphere, and $r_{A_{0,0}}$ is again a smooth Dehn-Siedel twist along this sphere.

Note that the k-fold blowup of $\mathbb{C}\mathrm{P}^2(b)$ with symplectic form $\omega_{(b;a_1,...,a_k)}$ is a constructible example of a blowup form in the cohomology class represented by the vector $(b; a_1, ..., a_k)$, provided the above necessary conditions are met. Recall that the volume of $\mathbb{C}\mathrm{P}^2(b)$ by blowups of size $a_1, ..., a_k$ is $b^2 - a_1^2 - ... - a_k^2$. We have the following.

Exercise. Given a vector $(b; a_1, ..., a_k)$ that represents a blowup form, it is an easy calculation to see that the Cremona transform does not change the volume of the underlying symplectic manifold, because

$$2br - 2(a_1 + a_2 + a_3)r - 2r^2 = 0.$$

These are the extra terms in the volume computation, after applying the Cremona transform.

Definition 3.1.2. A vector $(b; a_1, ..., a_k)$ is said to be *reduced* if the tail is ordered so that

$$a_1 \geq a_2 \geq ... \geq a_k, \tag{3.1.1}$$

and if we also have

$$a_1 + a_2 + a_3 \leq b. \tag{3.1.2}$$

In other words, the defect of a reduced vector is non-negative.

Definition 3.1.3. One applies the *standard Cremona move* to a vector v by composing the following two maps

1. The map that permutes the tail of v so that it is ordered as in (3.1.1).

2. The map $v \mapsto \begin{cases} Cremona(v) & \text{if defect}(v) < 0 \\ v & \text{if defect}(v) \geq 0 \end{cases}$

Specifically, the order of operations is the map (2) followed by the map (1). Let us say that two weight vectors are *Cremona equivalent* if you can get from the first vector to the second vector using a finite number of standard Cremona moves.

Note that the first map in this definition corresponds to the geometric Dehn-Seidel twist $r_{A_{i,j}}$ described in part (1) of the preceding remark. Note that the second map in this definition either does nothing to the vector v or it corresponds to the Dehn-Seidel twist $r_{A_{0,0}}$ described in part (2) of the preceding remark. This explains why the Cremona transform is an involution. The standard Cremona move is not an involution, because it eventually terminates in a reduced vector. That this process terminates is explained in the following lemma.

Lemma 3.1.1. The following facts apply to the standard Cremona move.

1. The standard Cremona move is a piecewise-linear and continuous map on vectors, $\mathbb{R}^{n+1} \to \mathbb{R}^{n+1}$.

2. The vectors which are fixed by the standard Cremona move are exactly the reduced vectors.

3. Suppose that $v = (b; a_1, ..., a_k)$ and $v' = (b'; a'_1, ..., a'_k)$ are such that v is ordered as in $(3.1.1)$ and $v \neq v'$. If $v \mapsto v'$ under the standard Cremona move, then

   - $b' < b$, and
   - $a'_j \leq a_j$ for $1 \leq j \leq k$, with at least one inequality being strict.

Proof. The first fact is obvious. For the second fact, it is also obvious that the re-ordering map in the standard Cremona move will fix a vector if and only if the vector is already ordered as in (3.1.1). We examine which vectors are fixed by the other map in the standard Cremona move. Note that a reduced vector is one whose defect is non-negative. This is equivalent to saying that a vector is reduced if and only if it is fixed by the map

$$v \mapsto \begin{cases} Cremona(v) & \text{if defect}(v) < 0 \\ v & \text{if defect}(v) \geq 0 \end{cases}.$$

Finally, we prove the third fact in this lemma. Assume that $v \mapsto v'$ under the standard Cremona move, but $v \neq v'$. Then the fact (2) that we just proved implies that v was not reduced, and since v is ordered, we must have that defect$(v) < 0$. So the standard Cremona

move will involve a Cremona transform applied to v. By definition, the Cremona transform takes $v = (b; a_1, ..., a_k)$ to

$$Cremona(v) = (b + r; a_1 + r, a_2 + r, a_3 + r, a_4, a_5, ..., a_k).$$

The result, after re-ordering the tail, is assumed to be v'. Since the first entry b' of v' equals b plus a negative number, we get that $b' < b$. This proves the first bullet point of the fact (3). Since we started with

$$a_1 \geq a_2 \geq a_3,$$

adding the same number r to each of these three terms gives

$$a_1 + r \geq a_2 + r \geq a_3 + r.$$

We need to decide whether these elements of v' are re-ordered by the standard Cremona move. There are four possibilities.

i. If $a_3 + r \geq a_4$, then no re-ordering is necessary. We set

$$a'_1 = a_1 + r, \quad a'_2 = a_2 + r, \quad a'_3 = a_3 + r, \quad a'_j = a_j \text{ for } 4 \leq j \leq k.$$

Then v' satisfies the second bullet point of fact (3).

ii. Otherwise, we must have $a_3 + r < a_4$. Since $a_2 + r \geq a_3 + r$ it may be the case that

$$a_1 + r \geq a_2 + r \geq a_4 > a_3 + r,$$

whence we set

$$a'_1 = a_1 + r, \quad a'_2 = a_2 + r, \quad a'_3 = a_4, \quad a'_j = a_3 + r \text{ for some } 4 \leq j \leq k.$$

as a result of the re-ordering in the standard Cremona move. Then $a'_3 = a_4 \leq a_3$ and $a'_1 < a_1$. The term $a_3 + r$ gets moved to the right in the re-ordered v', but sinve v was ordered, we still satisfy fact (3).

iii. Another possibility is

$$a_1 + r \geq a_4 > a_2 + r \geq a_3 + r.$$

The analysis is the same as (ii), except that two terms get moved to the right by

re-ordering.

iv. The only remaining possibility is

$$a_4 > a_1 + r \geq a_2 + r \geq a_3 + r.$$

The analysis is the same as (ii), except that three terms get moved to the right by re-ordering.

□

Example 3.1.2. The vector $(4; 2, 2, 1)$ is not reduced. The vector $(3; 1, 1, 0)$ is reduced. As shown above, the Cremona transform maps $(4; 2, 2, 1) \mapsto (3; 1, 1, 0)$, and we stop at this step, because we have arrived at a reduced vector. This operation was one standard Cremona move.

We now have two uses for the vector notation $(b; a_1, ..., a_k)$.

1. $(b; a_1, ..., a_k)$ could represent the negative weight vector of a convex toric domain; or
2. $(b; a_1, ..., a_k)$ could represent a cohomology class for the k-fold blowup of $\mathbb{C}\mathrm{P}^2$.

We described above how the Cremona transform can be applied to vectors which represent cohomology classes. We would like, instead, to apply the Cremona transform to the negative weight vector of a convex toric domain. We can identify these two vectors because of Proposition 3.0.3. This identification will be explained further in the proof of the main theorem of this chapter, which follows.

Theorem 3.1.1. Cremona equivalent negative weight vectors give rise to the same ECH capacity sequence. For this reason, we may as well compute the ECH capacities of a convex toric domain X_Ω using the reduced negative weight vector that corresponds to Ω.

Proof. We first explain why it suffices to consider negative weight vectors with $k = 3$ nonzero terms in the tail. If there are more than $k = 3$ terms in the tail, the Cremona transform only affects the first three terms of the tail of a vector. If a vector has fewer than $k = 3$ nonzero terms in its tail, one can perturb the vector by an arbitrarily small amount, ϵ, to ensure that it has precisely 3 nonzero terms in the tail. For instance, a vector that originally has one term in its tail, say $(b; a)$ can be perturbed to $(b + 2\epsilon, a, \epsilon, \epsilon)$, and this perturbation does not change the defect of the vector. After completing the following proof on a vector with $k = 3$ entries, we let $\epsilon \to 0$. This perturbation will not change the ECH capacities of Ω, because of the continuity property of ECH capacities (lemma 2.1.1).

With this simplifying assumption, we let $v = (b; a_1, a_2, a_3)$ denote the negative weight vector that corresponds to the convex polygon Ω. Let $v' = (b'; a_1', a_2', a_3')$ be a vector that is Cremona equivalent to v. (We may as well assume that v' is the *reduced* vector that is Cremona equivalent to v, which is guaranteed to exist because of lemma 3.1.1.) One can use proposition 2.1.4 to generate a convex polygon that has v' as its negative weight vector. Let us call that polygon Ω'. We now must show that X_Ω and $X_{\Omega'}$ have the same sequence of ECH capacities.

To say that X_Ω and $X_{\Omega'}$ have the same sequence of ECH capacities is equivalent to the following ball packing problem:

$$\mathcal{B} \hookrightarrow X_\Omega \quad \Longleftrightarrow \quad \mathcal{B} \hookrightarrow X_{\Omega'},$$

where $\mathcal{B}$ is any finite disjoint union of symplectic 4-balls. Recall that, in Chapter 2, we explained how one can view the symplectic manifolds involved in these ball packings as moment images. In this case, packing 4-balls into a convex toric domain reduces to packing triangles into a convex region of the plane. From this point of view, it is easy to see that we can sub-divide the balls in the collection $\mathcal{B}$ into smaller balls, if necessary, so that the radii of balls in $\mathcal{B}$, call them $w_1, ..., w_N$, satisfy

$$\max\{w_1, ..., w_N\} \leq \min\{a_1, a_2, a_3, a_1', a_2', a_3'\}.$$

By proposition 3.0.2, there exists a symplectic embedding $\mathcal{B} \hookrightarrow X_\Omega$ if and only if there exists a symplectic embedding

$$\left(\bigsqcup_{i=1}^{N} B^4(w_i)\right) \sqcup \left(\bigsqcup_{j=1}^{3} B^4(a_j)\right) \hookrightarrow B^4(b). \tag{3.1.3}$$

Similarly, there exists $\mathcal{B} \hookrightarrow X_{\Omega'}$ if and only if there exists

$$\left(\bigsqcup_{i=1}^{N} B^4(w_i)\right) \sqcup \left(\bigsqcup_{j=1}^{3} B^4(a_j')\right) \hookrightarrow B^4(b'). \tag{3.1.4}$$

Next, we apply proposition 3.0.1 to these putative embeddings. The embedding (3.1.3) exists if and only if there exists a blowup form on the $(N+3)$-fold blowup of $\mathbb{C}\mathrm{P}^2$ in the class represented by the vector $(b; a_1, a_2, a_3, w_1, ..., w_N)$. Similarly, the embedding (3.1.4) exists if and only if there exists a blowup form on the $(N+3)$-fold blowup of $\mathbb{C}\mathrm{P}^2$ in the class represented by the vector $(b'; a_1', a_2', a_3', w_1, ..., w_N)$.

Recapitulating, we have reduced the theorem to proving that there exists a blowup form

in the class represented by $\alpha = (b; a_1, a_2, a_3, w_1, ..., w_N)$ if and only if there exists a blowup form in the class represented by $\alpha' = (b'; a_1', a_2', a_3', w_1, ..., w_N)$. Karshon and Kessler [KK17] developed the following algorithm to decide when a cohomology class contains a blowup form. They showed that a blowup form exists in class α if and only if α reduces under standard Cremona moves to a vector with positive entries in its tail. Similarly, a blowup form exists in class α' if and only if α' reduces under standard Cremona moves to a vector with positive tail. But obviously α reduces under standard Cremona moves to a vector with positive tail if and only if α' does, because $\alpha \mapsto \alpha'$ under a single Cremona transformation. (Recall that we re-order vector entries after doing the Cremona transformation in the "standard" move.) This proves the equivalence of (3.1.3) and (3.1.4) and proves the theorem.

□

Example 3.1.3. We can now push corollary 2.2.2 even further. Let $q \geq 1$ be rational and written in lowest terms. Recall that (2.2.10) gives the negative weight vector for $E(1, q)$ treated as a convex toric domain. Let $\rho = \lfloor \frac{q-1}{2} \rfloor$. Let $\upsilon = \lfloor q \rfloor - 1 - 2\rho$, which will be 0 or 1. Applying the cremona transform ρ times to the vector (2.2.10) gives the negative weight vector

$$(q-\rho; q-\rho-1, \underbrace{1}_{\upsilon \text{ copies}}, \underbrace{q-\lfloor q \rfloor, ..., q-\lfloor q \rfloor}_{\lfloor 1/(q-\lfloor q \rfloor) \rfloor \text{ copies}}, ...).$$

We have proved the following

Corollary 3.1.1. The (finite length) positive weight vector

$$(\underbrace{1, ..., 1}_{\lfloor q \rfloor \text{ copies}}, \underbrace{q-\lfloor q \rfloor, ..., q-\lfloor q \rfloor}_{\lfloor 1/(q-\lfloor q \rfloor) \rfloor \text{ copies}}, ...)$$

generates the same sequence of ECH capacities as the (finite length) negative weight vector

$$(q-\rho; q-\rho-1, \underbrace{1}_{\upsilon \text{ copies}}, \underbrace{q-\lfloor q \rfloor, ..., q-\lfloor q \rfloor}_{\lfloor 1/(q-\lfloor q \rfloor) \rfloor \text{ copies}}, ...).$$

This corollary has the potential to greatly reduce the length of the resulting negative weight vector, when compared to the starting positive weight vector.

Example 3.1.4. The ellipsoid $E(1, 23/2)$ has positive weight vector

$$(\underbrace{1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1}_{\lfloor 23/2 \rfloor = 11}, 1/2, 1/2).$$

In the context of example 3.1.3, we compute $\rho = 5$ and $\upsilon = 0$. Hence, by corollary 3.1.1, this positive weight vector generates the same sequence of ECH capacities as the following, much shorter, negative weight vector:

$$\left(\frac{13}{2}; \frac{11}{2}, \frac{1}{2}, \frac{1}{2}\right).$$

Note that one may be given a positive weight vector that satisfies the hypothesis of corollary 3.1.1 without knowing that it comes from an ellipsoid.

3.2 Operations on Polygons

In this section we examine some operations that can be done to planar polygons to produce new planar polygons. For the remainder of this chapter, we consider only convex integral polygons, meaning that each vertex of the polygon lies at an integer lattice point, $\mathbb{Z}^2 \subset \mathbb{R}^2$. We then explain the effect that these operations have on the negative weight vector.

3.2.1 Blowing Up

Consider a convex integer polygon Δ. To *blow up* the polygon Δ means to locate a $45 - 45 - 90$ triangle having side lengths 1,1, and $\sqrt{2}$ inside of the polygon Δ such that two of the sides of this triangle are on the boundary of Δ. (In other words, the triangle is not completely in the interior.) Excise this triangle, making the remaining side of the triangle into a new edge of the resulting polytope Δ'.

It should be clear from this definition that a compact polygon can only be blown up a finite number of times.

Example 3.2.1. The square in the first quadrant having vertices (0,0),(2,0),(0,2), and (2,2) contains a triangle in its upper-right corner. This triangle is shaded in the left side of figure 3.1. Blowing up will give a new polygon, which is the convex hull of the points (0,0),(2,0),(2,1),(1,2), and (0,2). This polygon is also shown in figure 3.1.

Proposition 3.2.1. Suppose that Δ is a convex polygon with negative weight vector $(b; a_1, ..., a_n)$. A blow-up operation performed on Δ results in a new polygon Δ' with negative weight vector $(b; a_1, ..., a_n, 1)$ up to Cremona equivalence.

Proof. Blowing up the convex integer polygon is obviously equivalent to packing a triangle of size one into the polygon. This has the effect of adding a 1 to the tail of the negative weight vector. □

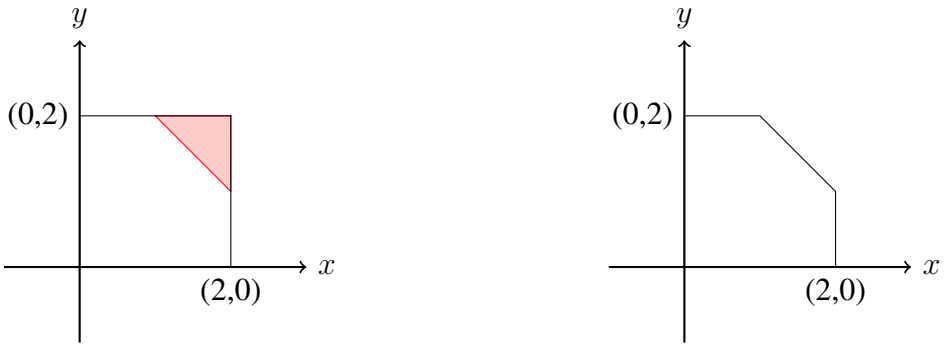

Figure 3.1: An example of blowing up a convex polygon.

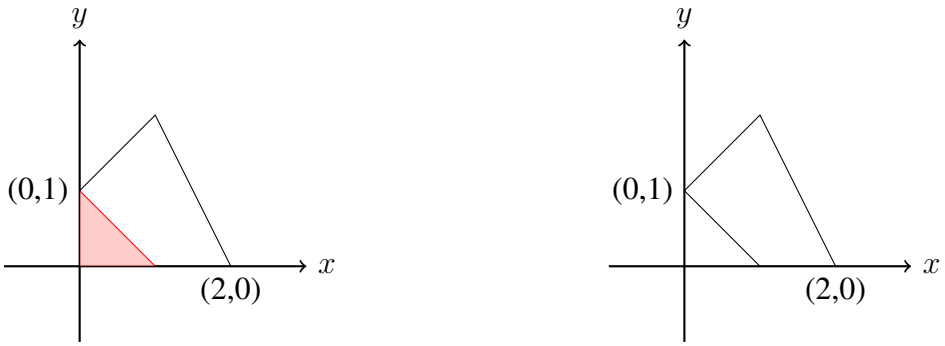

Figure 3.2: Another example of blowing up a convex polygon.

Example 3.2.2. The square in figure 3.1 has negative weight expansion $(4; 2, 2)$. Blowing up the square results in the polygon on the right side of figure 3.1. This polygon must have negative weight expansion $(4; 2, 2, 1)$, which we showed in a previous example to be Cremona equivalent to $(3; 1, 1, 0)$. This is indeed the negative weight expansion of the polygon on the right side of the figure, as the reader can check.

Example 3.2.3. Consider the polygon on the left side of figure 3.2. This polygon contains an isosceles right triangle of leg length 1 in its lower left corner. This triangle is shaded in the left side of figure 3.2. Blowing up by excising this shaded triangle will give a new polygon, which is shown on the right side of figure 3.2. Before blowing up, the negative weight vector of the left polygon was $(3; 1, 1, 1, 1)$. After blowing up, the negative weight vector of the right polygon is $(3; 1, 1, 1, 1, 1)$, as the reader can check.

3.2.2 Dualizing

Again, we will focus on convex integer polygons containing only a single interior lattice point. for the purpose of the following definition, we may as well assume that the only interior lattice point is the origin (0,0).

Definition 3.2.1. Given an integer polygon Δ with the origin as the only interior lattice point, say that Δ has precisely t lattice points on its boundary. We form the *dual polygon* $\Delta^\vee$ (using a process known as *dualizing*) as follows.

- List the integer lattice points that make up the boundary of Δ in clockwise order. These points should be written as vectors $\vec{p}_1, ..., \vec{p}_t$ in $\mathbb{Z}^2$.
- Let $\vec{q}_i = \vec{p}_{i+1} - \vec{p}_i$. All subscripts are taken modulo t, so that, e.g., $\vec{p}_{t+1} = \vec{p}_1$.
- Let $\Delta^\vee$ be the convex hull of $\vec{q}_1, ..., \vec{q}_t$.

By construction, $\Delta^\vee$ will be a convex integer polygon with the origin as its only interior lattice point.

The following classical result shows that the polygons we are considering are a good class for dualizing.

Theorem. (Scott, [Sco76]) Up to the action of $SL(2, \mathbb{Z})$ and translation by lattice points, there are precisely 16 integer polygons in the plane having precisely one interior lattice point.

For the sake of brevity, we refer to the sixteen convex integer polygons having only a single interior lattice point as *Fano polygons*, because of their connection to Fano varieties (which we will not examine here). The sixteen Fano polygons are drawn in figure 3.3. In that figure, note that the two polygons with the same label are dual to each other, and the polygons G through J are self-dual. For a computation of the negative weight vector of polygons E and C, see example 3.2.2 and example 3.2.3. The next theorem is not necessary to the analysis, but it helps to organize things.

Theorem. (Poonen, Rodriguez-Villegas, [PRV00]) Let $\ell(\Delta)$ denote the number of lattice points on the boundary of the convex integer polygon Δ. If Δ contains only a single interior lattice point, then so does $\Delta^\vee$, and

$$\ell(\Delta) + \ell(\Delta^\vee) = 12.$$

This means that, of the sixteen Fano polygons, we may as well group them according to the number of lattice points on the boundary. Such a polygon must have at least three lattice points on its boundary (in which case it is a triangle); hence at most 9, because of the symmetry in the formula $\ell(\Delta) + \ell(\Delta^\vee) = 12$. This shows that

$$\ell(\Delta) \in \{3, ..., 9\}.$$

Note. We can group the sixteen Fano polygons into three groups:

1. If $\ell(\Delta)$ is odd, then $\ell(\Delta^\vee)$ is odd.
2. If $\ell(\Delta) = 6$, then $\ell(\Delta^\vee) = 6$. In this case Δ will be self-dual.
3. If $\ell(\Delta) \in \{4, 8\}$, then $\ell(\Delta^\vee) \in \{8, 4\}$.

These facts can also be verified by inspecting all of the 16 possible polygons, drawn below.

Finally, to state the main result of this section, we need to define an operation of addition of weight sequences.

Notation. Let us denote the negative weight vector of a convex polygon Δ as $w(\Delta)$. If the tail of a weight vector has repetitions, abbreviate those repetitions with a superscript, which is not to be read as an exponent. For example $(3; 1, 1, 1)$ is abbreviated $(3; 1^3)$.

Definition 3.2.2. Given the negative weight vector of two convex polygons, say

$$w(\Delta) = (b; a_1, ..., a_n) \qquad \text{and} \qquad w(\Delta') = (b'; a'_1, ..., a'_n),$$

define the operation of *weight vector addition* by

$$w(\Delta) + w(\Delta') = (\max\{b, b'\}; \underbrace{a_1, ..., a_n, a'_1, ..., a'_n}_{\text{unordered list w. repetitions}}).$$

The addition of weight sequences results in a list of numbers that resembles a weight sequence, but we are not claiming that the resulting weight sequence has geometric significance.

Proposition 3.2.2. In cases (1) and (2) of the above note, we have $w(\Delta)+w(\Delta^\vee) = (3; 1^6)$. In case (3) of the above note, we may have $w(\Delta) + w(\Delta^\vee) = (4; 2^2, 1^5)$, which is Cremona equivalent to $(3; 1^6)$.

Proof. Check all of the sixteen Fano polygons, which are drawn below. □

Note. Recall the operation of sequence summation, from definition 1.0.4. Unfortunately, it is not the case that $c(X_\Delta) + c(X_{\Delta^\vee})$ always corresponds to the ECH capacities for $(3; 1^6)$. An explicit counterexample follows. The second ECH capacity for the toric domain corresponding to $(3; 1^6)$ is $c_2 = 3$. The first three ECH capacities for the toric domain corresponding to $(3; 1^3)$ are $c_0 = 0$, $c_1 = 2$, and $c_2 = 3$. As was noted in the fact above, the polygon (and the toric domain) corresponding to $(3; 1^3)$ is self-dual. So we sum $c(X_\Delta)+c(X_{\Delta^\vee})$

by adding the sequence $\{0, 2, 3\}$ to itself. This summation gives $\{0, 2, 4\}$, and the $c_2 = 4$ term differs from the calculation for $(3; 1^6)$.

This counterexample shows that adding weight sequences does not have the geometric interpretation of taking the disjoint union of the corresponding toric domains. So we need to distinguish the operations of sequence addition from the operation of weight sequence addition.

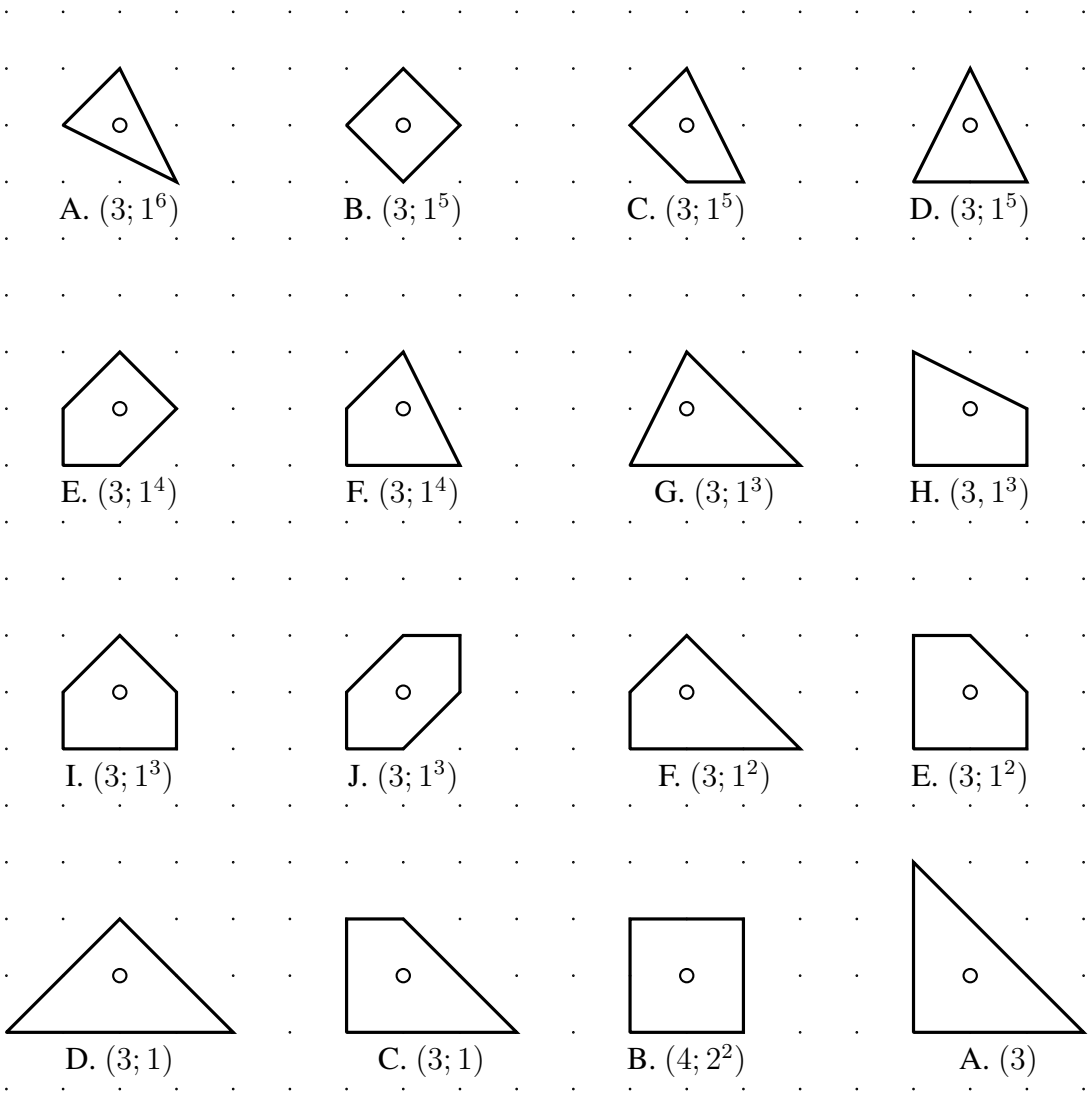

Figure 3.3: The sixteen Fano polygons with their negative weight vectors.

CHAPTER 4

Embeddings of Stabilized Polydiscs

Recall that we defined the symplectic ellipsoid in definition 1.0.2 and the symplectic polydisc in definition 1.0.3. We repeat these definitions here using real coordinates. We work in $\mathbb{C}^n \cong \mathbb{R}^{2n}$, with coordinates $x_1, y_1, ..., x_n, y_n$ and with the standard symplectic form

$$\omega_{std} = \sum_{i=1}^{n} dx_i \wedge dy_i.$$

Definition 4.0.1. The *symplectic polydisc of areas* r_1 *through* r_n is the set

$$P(r_1, ..., r_n) = \left\{ (x_1, y_1, \dots, x_n, y_n) \in \mathbb{R}^{2n} \;\middle|\; \pi \cdot \left(x_j^2 + y_j^2\right) \leq r_j \,,\; j = 1, 2, ..., n \right\},$$

with the induced symplectic form

Definition 4.0.2. The *symplectic ellipsoid of areas* r_1 *through* r_n is the set

$$E(r_1, ..., r_n) = \left\{ (x_1, y_1, \dots, x_n, y_n) \in \mathbb{R}^{2n} \;\middle|\; \sum_{j=1}^{n} \frac{\pi \cdot \left(x_j^2 + y_j^2\right)}{r_j} \leq 1 \right\},$$

with the induced symplectic form.

In each of the above definitions, we make it a convention to order the areas as $r_1 \leq r_2 \leq ... \leq r_n$, because a permutation of the coordinates (x_i, y_i) would be a linear symplectomorphism.

These are the two most fruitful examples of toric domains where symplectic embedding problems have been studied in (real) dimension four. When examining higher-dimensional embedding problems, we use the following construction.

Definition 4.0.3. A *stabilized polydisc* is the symplectic manifold $P(r_1, r_2) \times \mathbb{R}^{2n-4}$ with the symplectic form inherited from $\mathbb{R}^{2n}$ viewed as a product symplectic manifold.

Notice that in this definition we always use a 4-dimensional polydisc in order to make the resulting stabilized polydisc a manifold of dimension $2n$, $n \geq 3$.

In this chapter shall prove the following main theorem about stabilized polydiscs.

Theorem 4.0.1. (Main theorem) Suppose that $x \geq 2$ and $n \geq 3$. There exists a symplectic embedding

$$P(1, x) \times \mathbb{R}^{2n-4} \hookrightarrow P(a, b) \times \mathbb{R}^{2n-4}$$

if and only if either

- $a \geq 2$, or
- $1 \leq a < 2$ and $b \geq x$.

Proof. (outline) At first we will focus on the case $n = 3$. At the conclusion of this chapter (section 4.7), we will discuss how all of the proofs generalize to higher dimensions. So for now, we fix $n = 3$.

Let us immediately dispense with the easy cases of this theorem. The "if" ($\Longleftarrow$) portion of this theorem follows from the work of Hind and Kerman ([HK14] §4). It was later shown in [Hin15] that these embeddings are possible using symplectic folding in the case $a > 2$. Next, for the "only if" claim ($\Longrightarrow$), we assume that a symplectic embedding, of the sort mentioned in the theorem, exists. If the capacity a happens to satisfy $a \geq 2$, then the first bullet point of the conclusion is proved. Otherwise we must show if $1 \leq a < 2$ then $b \geq x$ to prove the second bullet point of the conclusion. We prove this claim by contradiction, and this will be the content of the rest of the chapter.

We can, however, relax the assumption $x \geq 2$. We show that it suffices to consider, instead, the case $x > 2$. When $x = 2$ and $a \geq 2$, then there is no contradiction to the claim, as above. If $x = 2$ and $1 \leq a < 2$, then we are assuming

$$P(1, 2) \times \mathbb{R}^2 \hookrightarrow P(a, b) \times \mathbb{R}^2 .$$

For $0 < \lambda < 1$, inclusion clearly gives

$$P(1 - \lambda, 2) \times \mathbb{R}^2 \hookrightarrow P(a, b) \times \mathbb{R}^2 .$$

Re-scaling the polydiscs by $\frac{1}{1-\lambda}$ gives

$$P\left(1, 2/(1 - \lambda)\right) \times \mathbb{R}^2 \hookrightarrow P(a/(1 - \lambda), b/(1 - \lambda)) \times \mathbb{R}^2 .$$

Because $a < 2$, we can choose λ so that $a/(1 - \lambda) < 2$. Then we have reduced the case of

$x = 2$ to the case of $x > 2$. This is just a specific example of why the above definition of $\hookrightarrow$ on page 2 is useful. $\square$

Recaptitulating, for the remainder of this chapter, we shall assume that $x > 2$, that

$$P(1, x) \times \mathbb{R}^2 \hookrightarrow P(a, b) \times \mathbb{R}^2$$

and that $1 \leq a < 2$, but $b < x$.

We can also use the main theorem to obstruct embeddings of four-dimensional polydiscs, as the following result illustrates.

Corollary 4.0.1. Suppose that $P(1, x)$ symplectically embeds into $P(a, b)$. Then, by taking a product with the identity map, we find that $P(1, x) \times \mathbb{R}^2 \hookrightarrow P(a, b) \times \mathbb{R}^2$. Assume further that $x \geq 2$. Then the main theorem implies that either $a \geq 2$ or $1 \leq a < 2$ and $b \geq x$.

We compare this corollary to a result of Hutchings, which has been modified to match the notation of this paper.

Theorem 4.0.2. (Hutchings, [Hut16]) Suppose that $P(1, x)$ symplectically embeds into $P(a, b)$ with $1 \leq a \leq b$. Assume further that

$$1 \leq x \leq \frac{2(b/a)}{1 + \frac{(b/a)-1}{4\left\lceil \frac{b}{a} \right\rceil - 1}} \quad (\geq 2). \tag{4.0.1}$$

Then $b \geq x$.

The upper bound $(4.0.1)$ is always at least 2. So we may as well assume $1 \leq x \leq 2$, and conclude $b \geq x$. Together, Theorem 4.0.2 and Corollary 4.0.1 give a full picture of four-dimensional embeddings $P(1, x)$ into $P(a, b)$ for all values of x. The two results reach similar conclusions.

Specifically, let us define a four-dimensional embedding function f, in the spirit of [MS12b], also known as an embedding capacity function. For $x \geq 1, a \geq 1$, we set

$$f(x, a) := \inf_b \{\text{there exists a symplectic embedding } P(1, x) \text{ into } P(a, b), \text{ with } b \geq a\}.$$

The two preceding results show that for $1 \leq a < 2$, $f(x, a) = x$. The case of $a \geq 2$ is only partially known because this is the regime of symplectic folding (see proposition 4.4.4 and Figure 7.2 in [Sch05]).

An outline of the sections of this chapter follows. In section 4.1 we examine the Reeb dynamics of a symplectic ellipsoid, and write down a formula for the generalized Conley-Zehnder index. In section 4.2 we compute the generalized Conley-Zehnder indices for

Reeb orbits on a polydisc, and describe a special neighborhood of a Lagrangian torus, along which we shall *stretch the neck*. In section 4.3, we assert the existence of a holomorphic curve in a compactification of $P(a,b) \times \mathbb{R}^2$ with negative end asymptotic to a skinny ellipsoid. We then stretch the neck, and in sections 4.4 and 4.5 we examine the geometry of the aforementioned curve. The moral of this story is that the energy of a holomorphic curve must be distributed in a precise way. In section 4.6 we prove the existence theorem of the holomorphic curve from section 4.3. The proof in section 4.6 can be read before reading all of section 4.3, if so desired. Finally, in section 4.7 we explain how all of the results of this chapter extend to dimension $2n \geq 6$.

In proving the main theorem, we will introduce many parameters. For easy reference, let us define these parameters upfront in a list that has some interdependence:

- n will always denote half the (real) dimension of the ambient symplectic manifold.
- x, a, and b are capacities of the polydiscs in the statement of the main theorem.
- d is an integer parameter satisfying $2d+1 >> b$.
- S is a very large real number satisfying $da + b << S$. Because of the definition of $\hookrightarrow$, we can and shall replace $P(1,x) \times \mathbb{R}^2$ with $P(1,x,S)$.
- $\epsilon > 0$ is small enough that $\frac{(2d+1)\epsilon}{2} < x$ (This ϵ will be fixed throughout the proof, and the ϵ in the definition of $\hookrightarrow$ on page 2 will not factor into the proof).
- $\delta_1, \delta_2, \delta_3, ..., \delta_n$ are very small real numbers satisfying $\delta_2, \delta_3, ..., \delta_n < \delta_1 << \epsilon$ and such that
$$\frac{\epsilon - \delta_i}{\epsilon - \delta_j} \quad \text{for } 1 \leq i \neq j \leq n, \tag{4.0.2}$$
are all irrational.

4.1 The Symplectic Ellipsoid

In this section we consider the Reeb dynamics on the boundary of a six-dimensional symplectic ellipsoid, which is an example of a toric domain. When the ratio of the ellipsoid areas (e.g. semi-major axis to semi-minor axis) is irrational, then the Reeb dynamics on the boundary are easy to understand. Moreover, when the ellipsoid is sufficiently eccentric (a so-called "skinny" ellipsoid) then there will only be one Reeb orbit of any consequence. This is precisely the setup we shall leverage when producing a holomorphic curve in the sequel.

Notice that the six-dimensional ellipsoid $E(r_1, r_2, r_3)$, as in definition 4.0.2, can also be expressed as the toric domain for the tetrahedral region below the face $\mathcal{R}$ that is shown in figure 4.1a. This ellipsoid inherits the standard symplectic form from $\mathbb{C}^3 \cong \mathbb{R}^6$, setting $z_j = x_j + iy_j$, and this form restricts to a contact form α on $\partial E(r_1, r_2, r_3)$. Specifically

$$\alpha = \frac{1}{2}\left(x_1 dy_1 - y_1 dx_1 + x_2 dy_2 - y_2 dx_2 + x_3 dy_3 - y_3 dx_3\right).$$

This contact form generates a Reeb vector field

$$v = \frac{2\pi}{r_1}\left(x_1\frac{\partial}{\partial y_1} - y_1\frac{\partial}{\partial x_1}\right) + \frac{2\pi}{r_2}\left(x_2\frac{\partial}{\partial y_2} - y_2\frac{\partial}{\partial x_2}\right) + \frac{2\pi}{r_3}\left(x_3\frac{\partial}{\partial y_3} - y_3\frac{\partial}{\partial x_3}\right),$$

along with a contact structure ξ defined by $\xi = ker(\alpha)$.

For the remainder of section 4.1, assume that the triple (r_1, r_2, r_3) satisfies the following condition

$$\frac{r_i}{r_j} \in \mathbb{R}\setminus\mathbb{Q} \quad \text{for all } 1 \leq i \neq j \leq 3. \tag{4.1.1}$$

Then the boundary $\partial E(r_1, r_2, r_3)$ corresponds to the constant energy level set of a Hamiltonian mechanical system (i.e. a sum of three independent oscillators) that is completely integrable. The components of the moment map μ, above, obviously Poisson commute, and fibers of this map are either T^3 or T^2 or T^1, as will be described shortly. It follows that there exist linear (action angle) coordinates on each torus fiber such that the Hamiltonian flow is straight-line motion. Choose a Hamiltonian (involving all coordinates of μ) that generates the Reeb flow. Then we examine the three options for the fibers of the moment map μ.

- A point (z_1, z_2, z_3) within the open face of $\mathcal{R}$ will have no coordinates vanishing, and a fiber of the moment map over such a point will be T^3. A Reeb orbit in such a fiber cannot be closed, because the assumption (4.1.1) implies that $T\cdot(r_1, r_2, r_3) \in \mathbb{R}^3\setminus\mathbb{Q}^3$ for every real $T > 0$.

- A fiber over the three edges of $\mathcal{R}$ (but not over a vertex) will be T^2. A Reeb orbit in such a fiber cannot be closed, because the assumption (4.1.1) implies that the slope is irrational.

- Finally, a Reeb orbit in a fiber over any vertex of $\mathcal{R}$ must be closed, because such a fiber is just $T^1 = S^1$.

The conclusion is that the irrationality condition on the capacities precludes closed Reeb orbits on ∂E, save for the three orbits γ_1, γ_2, γ_3 that occur in the planes $\{|z_i| = |z_j| =$

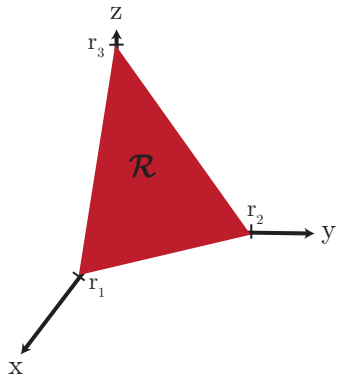

(a) The region that generates the ellipsoid $E(r_1, r_2, r_3)$.

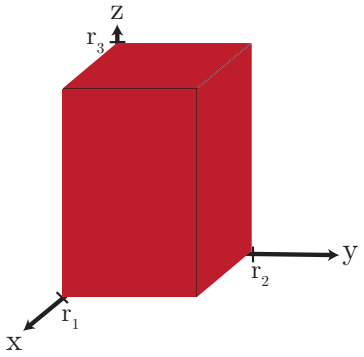

(b) The region that generates the polydisc $P(r_1, r_2, r_3)$. The edges of the cuboid have been exaggerated here.

Figure 4.1: The (solid) regions in the first octant that generate two important toric domains.

$0,\ i \neq j\}$. Furthermore, the condition (4.1.1) applied to the capacities r_1, r_2 would be sufficient to preclude closed Reeb orbits if the ellipsoid were four dimensional, $E(r_1, r_2)$, a case we shall revisit later. Henceforth, the phrase "Reeb orbit" will always mean *closed* Reeb orbit.

Example 4.1.1. In the six-dimensional case, the convention $r_1 \leq r_2 \leq r_3$ means that we may name $\gamma = \gamma_1$ the *short orbit*. In coordinates $(x_1, y_1, x_2, y_2, x_3, y_3)$ as above, the Reeb orbits are parameterized as

$$\gamma_1(t) = \left(\frac{r_1}{\pi}\cos\left(\frac{2\pi t}{r_1}\right), \frac{r_1}{\pi}\sin\left(\frac{2\pi t}{r_1}\right), 0, 0, 0, 0\right) \qquad \text{with period } r_1,$$

$$\gamma_2(t) = \left(0, 0, \frac{r_2}{\pi}\cos\left(\frac{2\pi t}{r_2}\right), \frac{r_2}{\pi}\sin\left(\frac{2\pi t}{r_2}\right), 0, 0\right) \qquad \text{with period } r_2,$$

$$\gamma_3(t) = \left(0, 0, 0, 0, \frac{r_3}{\pi}\cos\left(\frac{2\pi t}{r_3}\right), \frac{r_3}{\pi}\sin\left(\frac{2\pi t}{r_3}\right)\right) \qquad \text{with period } r_3.$$

The time t Reeb flow $\text{Flow}_t\colon \partial E \to \partial E$ is easily seen to be the linear map

$$\text{Flow}_t(x_1, ..., y_3) = \begin{bmatrix} Rot(2\pi t/r_1) & 0 & 0 \\ 0 & Rot(2\pi t/r_2) & 0 \\ 0 & 0 & Rot(2\pi t/r_3) \end{bmatrix} \cdot \begin{bmatrix} x_1 \\ \vdots \\ y_3 \end{bmatrix},$$

where $Rot(\varphi)$ denotes the real 2×2 matrix of rotation through angle φ,

$$Rot(\varphi) = \begin{bmatrix} cos(\varphi) & -sin(\varphi) \\ sin(\varphi) & cos(\varphi) \end{bmatrix}.$$

Along any Reeb orbit (the short orbit $\gamma = \gamma_1$, in particular) the Reeb flow gives a *linearized return map* $P_\gamma \colon T_{\gamma(0)}\,\mathbb{C}^3 \to T_{\gamma(T)}\,\mathbb{C}^3$, which is defined to be the derivative of the time T (the period) flow. The trivialization of these tangent spaces come from the inclusion into $\mathbb{C}^3$ (with the linearization acting trivially on the normal vector to $\partial E|_\gamma$).

Definition 4.1.1. The Reeb orbit γ is said to be *nondegenerate* whenever P_γ does not have 1 as an eigenvalue. Furthermore, γ is called *elliptic* if the eigenvalues of P_γ lie within S^1, or γ is called *hyperbolic* if the eigenvalues of P_γ are real.

4.1.1 Generalized Conley-Zehnder Indices

We define an index formula for a continuous path of symplectic matrices, which may have 1 as an eigenvalue. This is known as the generalized Conley-Zehnder Index, or the Robbin-Salamon index of the path of matrices. This definition is taken from [Gut14], suitably modified.

Definition 4.1.2. If γ is a closed (possibly degenerate) Reeb orbit with period T, we might view γ as a map having domain $\mathbb{R}\,/\,T$, a circle. Let τ denote a trivialization of the contact structure of the image space along γ. With respect to the trivialization τ, the linearized Reeb flow is a symplectic matrix, denoted $P_\gamma(s)$ for $s \in \mathbb{R}\,/\,T$. Specifically, we regard $P_\gamma(s) \colon (\xi_{\gamma(0)}, \omega_{std}) \to (\xi_{\gamma(s)}, \omega_{std})$ as a $(2n-2) \times (2n-2)$ matrix. To the path of such matrices $\{P_\gamma(s) \,|\, 0 \leq s \leq T\}$, we associate a *generalized Conley-Zehnder index*, denoted $CZ(\gamma)$. The Conley-Zehnder index is computed by regarding the graph of $P_\gamma(s)$ as a path of Lagrangian subspaces of $(\mathbb{R}^{2n-2} \times \mathbb{R}^{2n-2}, (-\omega_{std}) \times \omega_{std})$ and then taking the Robbin-Salamon index of this graph relative to the diagonal $\Delta = \{(x,x) \in \mathbb{R}^{2n-2} \times \mathbb{R}^{2n-2}\}$.

Note that the generalized Conley-Zehnder index is half-integer valued, but it coincides with the traditional Conley-Zehnder index in the case when the Reeb orbit γ is nondegenerate, a fact we use now. The following two claims are immediate consequences of this definition. Recall that we are using the trivialization coming from the inclusion of an ellipsoid into $\mathbb{C}^n$. (See, for example, [Ust99] p.14.)

Claim. Suppose $A \colon [0,T] \to \mathbb{R}^{2n_1}$ and $B \colon [0,T] \to \mathbb{R}^{2n_2}$ are two continuous paths of symplectic transformations. Identify $Sp(2n_1) \times Sp(2n_2)$ as the subgroup of block-diagonal

matrices within $Sp(2n_1+2n_2)$. Then the Conley-Zehnder index enjoys the following direct sum identity.

$$CZ\left(t \mapsto A \oplus B\right) = CZ\left(A\right) + CZ\left(B\right).$$

Claim. A rotation in the plane by angle $2\pi t$ is the transformation $Rot(2\pi t)$ defined above. The path $t \mapsto Rot(2\pi t)$ for $0 \leq t \leq T$ has Conley-Zehnder index

$$CZ(Rot(2\pi t)) = \lfloor T \rfloor + \lceil T \rceil = 2\lfloor T \rfloor + 1.$$

Definition 4.1.3. The total rotation angle, T, of such an elliptic Reeb orbit is called the *monodromy angle*.

Putting these above two claims together, we get the following important result about the short orbit $\gamma = \gamma_1$.

Lemma 4.1.1. For the ellipsoid $E(r_1, r_2, r_3)$, the one parameter group of diffeomorphisms given by the time t flow along γ induces a symplectic linear map $\phi_t \colon \xi_{\gamma(0)} \to \xi_{\gamma(t)}$. Note that r_1 is the period of this short obit. For $0 \leq t \leq r_1$, the path $t \mapsto \phi_t$ has Conley-Zehnder index

$$CZ(\gamma) = 4 + 2\left\lfloor \frac{r_1}{r_2} \right\rfloor + 2\left\lfloor \frac{r_1}{r_3} \right\rfloor.$$

More generally, consider $E(r_1, ..., r_n)$ with all $r_i/r_j \in \mathbb{R} \setminus \mathbb{Q}$. Let $\gamma^{(r)}$ denote the r-fold cover of the short orbit, which is the path $\gamma(t)$ for $0 \leq t \leq r \cdot r_1$. The corresponding path of symplectic transformations has Conley-Zehnder index

$$CZ(\gamma^{(r)}) = 2r + 2\left\lfloor \frac{r \cdot r_1}{r_2} \right\rfloor + ... + 2\left\lfloor \frac{r \cdot r_1}{r_n} \right\rfloor + n - 1.$$

Proof. The second claim from Ustilovsky [Ust99] above implies that re-scaling the rotation $Rot(2\pi t)$ as $Rot(2\pi rt/r_2)$ should re-scale the Conley-Zehnder index as $2\lfloor rT/r_2 \rfloor + 1$, and similarly for rotations in the other coordinate planes. Specifically, in the first coordinate plane we get $2\lfloor rr_1/r_1 \rfloor = 2r$, because the period of the orbit we consider is $T = r_1$. In the first coordinate plane there is no term $+1$, because the time one map is the identity. Using the first claim of Ustilovsky and the above expression for Flow_t, for $0 \leq t \leq r_1$, we have

$$CZ\left(\gamma^{(r)}\right) = 2r + \left[\left(2\left\lfloor \frac{r \cdot r_1}{r_2} \right\rfloor + 1\right) + ... + \left(2\left\lfloor \frac{r \cdot r_1}{r_n} \right\rfloor + 1\right)\right].$$

The terms in square brackets comprise a sum of $(n-1)$ summands, which are identical except for the denominator inside the floor function. The desired formula follows by gathering terms. $\square$

Finally, this short orbit will have action πr_1. The monodromy angle scales as r_1/r_2 for $E(r_1, r_2)$.

4.2 A Lagrangian Torus

Recall that we are assuming that $P(1,x) \times \mathbb{R}^2 \hookrightarrow P(a,b) \times \mathbb{R}^2$. In this section, we regard the symplectic polydiscs under consideration as toric domains, and we compactify the target polydisc. A neighborhood of a Lagrangian torus in the source polydisc will exhibit interesting Reeb dynamics (as did the ellipsoid in the previous section). We shall write down a Conley-Zehnder formula for these Reeb orbits in this section. Finally, we will consider pseudoholomorphic curves in (a compactification of) $P(a,b) \times \mathbb{R}^2$.

Now, let $S >> da+b$. As mentioned above, because of the definition of $\hookrightarrow$, we replace $P(1,x) \times \mathbb{R}^2$ with $P(1,x,S)$. The symplectic polydisc $P(1,x,S)$ is a toric domain for the cuboid in the first octant having side lengths $1, x$, and S. (See figure 4.1b.) We define U to be the subset of $P(1,x,S)$ defined by

$$U = \mu^{-1}(\,(1-\epsilon,1) \times (x-(2d+1)\epsilon, x) \times (S/2, S)\,).$$

In symplectic polar coordinates, one can view a six-dimensional toric domain as toric fibers over a base given by a region in the first octant, as above. Let us view the symplectic polydisc $P(1,x,S)$ in symplectic polar coordinates:

$$R_i = \pi|z_i|^2 \qquad\qquad \theta_i = arctan(y_i/x_i)$$

for $z_j = x_j + iy_j$ in the identification $\mathbb{C}^n \cong \mathbb{R}^{2n}$. Then there is a Lagrangian torus

$$L = \left\{ \pi|z_1|^2 = 1 - \frac{\epsilon}{2},\ \pi|z_2|^2 = x - \frac{(2d+1)\epsilon}{2},\ \pi|z_3|^2 = \frac{3S}{4} \right\}$$

inside the set U.

Using the moment description above, we can give a complete description of the Reeb orbits on ∂U, similar to the characterization given at the beginning of section 4.1. Under the moment map, U maps to a parallelepiped (cuboid) P away from the coordinate axes of the first octant of $\mathbb{R}^3$. The Reeb flow on ∂U is always tangent to the Lagrangian torus fibers of the moment map, and the direction is given by the normal vector to P. Hence, we find periodic Reeb orbits (appearing in a 2-dimensional family which will foliate the fiber torus) precisely when the normal vector to P is a multiple of a nonzero integer vector.

Remark. Unfortunately, in this moment description, the unit tangent bundle will not be smooth. This can be remedied by defining Σ to be a smoothing of ∂U obtained in a special way, described in the next paragraph. Note that U admits a symplectic embedding into T^*T^3 which takes L to the zero section. We use this identification to regard U as a tubular neighborhood of the zero section in the cotangent bundle of the Lagrangian torus L, defined by a Finsler metric. Then the Reeb orbits on ∂U correspond to geodesics on the base. More precisely, if $c(t)$ is a geodesic, then (c, dc^*) will be a Reeb orbit in ∂U, and vice-versa. Such closed geodesics are parameterized by $H_1(T^3, \mathbb{Z}) \setminus (0,0,0)$. In this way we see that closed Reeb orbits on ∂U are in bijective correspondence with some triples of integers. Our distinguished torus has dimension 3, which means that geodesics on the torus appear in 2-dimensional families (given by translating the geodesics). Using the described correspondence in general, Reeb orbits in ∂U must appear in $(n-1)$-dimensional families.

Define a Hamiltonian on the cotangent bundle of the Lagrangian torus L, with fiber coordinates y_1, y_2, y_3 by the formula

$$H = \left(\frac{|y_1|^p}{\epsilon^p} + \frac{|y_2|^p}{((2d+1)\epsilon)^p} + \frac{2|y_3|^p}{S^p} \right)^{1/p},$$

where p is some large integer. Then the set $H^{-1}(1) \subset T^*T^3$, centered at the Lagrangian torus L, will be a smoothing of ∂U. We name this smoothing Σ and henceforth consider only Reeb orbits on Σ. Note that for p large, Σ is close to ∂U in the Hausdorff topology, and the Reeb orbits on Σ approximate the orbits on ∂U.

The Reeb orbits along Σ are pseudo-hyperbolic, and the formulas in the previous section will not suffice to compute the Conley-Zehnder index.

Lemma 4.2.1. Let $c\colon [0,1] \to \Sigma$ be a closed Reeb orbit on Σ in homology class $(k, \ell, m) \in H_1(T^3) \setminus (0,0,0)$. With respect to the trivialization coming from the inclusion into $\mathbb{C}^3$, the generalized Conley-Zehnder index of such an orbit is

$$CZ(c) = 2k + 2\ell + 2m + 1.$$

Proof. The so-called symplectic shear matrix $A(t) = \begin{bmatrix} 1 & -t/2 \\ 0 & 1 \end{bmatrix}$, for $t \in [0,1]$, has generalized Conley-Zehnder index $CZ(A(t)) = \frac{1}{2}$. With respect to the given trivialization, the

orbit $c(t)$ has linearized Reeb flow given as a block-diagonal

$$c(t) = \begin{bmatrix} e^{it} & & \\ & A(t) & \\ & & A(t) \end{bmatrix}, \qquad t \in [0,1].$$

Now, using the product and shear properties of Lemma 26 of [Gut14], we find that

$$CZ(c) = 2(k+m+\ell) + \frac{1}{2} + \frac{1}{2},$$

as desired. Note that because the Reeb orbits along Σ come in Morse-Bott families, we needed the added generality of the definition given in section 4.1.1.

□

Now, choose $\delta_2, \delta_3 < \delta_1 << \epsilon$ very small, and so that the condition (4.0.2) in the list at the beginning of this chapter holds. Define

$$E = E(\epsilon - \delta_1, (2d+1)(\epsilon - \delta_2), (2d+1)(\epsilon - \delta_3)).$$

Then E symplectically embeds into U, and this embedding clearly extends to the closure of the ellipsoid. So we will not denote it by a hooked arrow. Furthermore

$$X = \mathbb{C}\mathrm{P}^1(a) \times \mathbb{C}\mathrm{P}^1(b) \times \mathbb{C}$$

is a compactification of the first factor of $P(a,b) \times \mathbb{R}^2$ endowed with symplectic form

$$\omega_X = a^2\omega_{FS} + b^2\omega_{FS} + \omega_{std}. \tag{4.2.1}$$

Here ω_{FS} is the Fubini-Study form on $\mathbb{C}\mathrm{P}^1$ and ω_{std} is the standard symplectic form on $\mathbb{C} \cong \mathbb{R}^2$. We now have

$$E \to U \hookrightarrow P(a,b) \times \mathbb{R}^2 \hookrightarrow X.$$

The reader may wonder as to why the $\mathbb{C}$ factor remains in X. In future sections, we will consider holomorphic curves in X, whose projection onto the $\mathbb{C}$ factor will always be bounded. More precisely, we shall fix the homology class of such holomorphic curves, which will fix the area and constrain the diameter. In short, there is no reason to worry about compactifying the $\mathbb{C}$ factor.

4.2.1 J-holomorphic curves

We conclude this section by describing the almost-complex geometry of the embedded $E \to U \hookrightarrow X$.

Definition 4.2.1. A symplectic manifold (M, ω) is said to have (one or more) *cylindrical end(s)* if there is a compact codimension 0 submanifold K together with a symplectomorphism $M \setminus K \cong \partial K \times (0, \infty)$. Further, we require that M be equipped with an almost-complex structure J which is translation-invariant on $M \setminus K \cong \partial K \times (0, \infty)$.

Definition 4.2.2. Let (M, ω, J) be any symplectic manifold with compatible almost-complex structure. A *J-holomorphic curve* is a map $u \colon (S^2 \setminus \Gamma, j) \to (M, J)$ satisfying the following PDE (called the Floer equation):

$$du \circ j = J \circ du.$$

Here Γ is a finite collection of punctures and j is a complex structure that makes $S^2 \setminus \Gamma$ into a Riemann surface. Finally, in this paper any J-holomorphic curve u will have image in a manifold M having a cylindrical end, $M \setminus K \cong \partial K \times (0, \infty)$, with submanifold K being of contact type (described below). We then put boundary conditions on the above Floer equation to ensure that u is asymptotic to closed Reeb orbits on ∂K at each of the punctures. Concretely, we give each puncture in the domain of u a collar neighborhood with coordinates $(\theta, r) \in S^1 \times (0, \infty)$ such that

$$\lim_{r \to \infty} u(\theta, r) = \gamma(\theta)$$

is a Reeb orbit in $\partial K \times \{\infty\}$.

Definition 4.2.3. A *J-holomorphic plane* is a J-holomorphic curve whose domain is a once-punctured sphere. To fix notation, we will henceforth consider the domain of a J-holomorphic plane to be $\mathbb{C}$ with the standard complex structure.

In $\mathbb{C}\mathrm{P}^1$, let p_∞ denote the point at infinity. Abbreviate $\mathbb{C}\mathrm{P}^1 \times \mathbb{C}\mathrm{P}^1 \times \mathbb{C}\mathrm{P}^1$ as $(\mathbb{P}^1)^3$. Then our ambient manifold X can be realized as $(\mathbb{P}^1)^3 \setminus (\mathbb{C}\mathrm{P}^1 \times \mathbb{C}\mathrm{P}^1 \times p_\infty)$. Now, inside of $(\mathbb{P}^1)^3$ there are two other important divisors, which are

$$D_1 = p_\infty \times \mathbb{C}\mathrm{P}^1 \times \mathbb{C}\mathrm{P}^1 \qquad \text{and} \qquad D_2 = \mathbb{C}\mathrm{P}^1 \times p_\infty \times \mathbb{C}\mathrm{P}^1.$$

Let $[D_1], [D_2] \in H^2((\mathbb{P}^1)^3)$ denote the classes given by $c_1(L(D_1))$ and $c_1(L(D_2))$ (where $L(D_j)$ denotes the complex line bundle over the divisor D_j). Restrict each of $[D_1]$ and $[D_2]$

to $(\mathbb{P}^1)^3 \setminus (\mathbb{C}\mathrm{P}^1 \times \mathbb{C}\mathrm{P}^1 \times p_\infty)$ to obtain two non-degenerate, independent functionals on $H_2(X)$.

Note. $[D_1]$ and $[D_2]$ are dual to the homology classes of the curves

$$C_1 = \mathbb{C}\mathrm{P}^1 \times p_\infty \times \{1\} \qquad \text{and} \qquad C_2 = p_\infty \times \mathbb{C}\mathrm{P}^1 \times \{1\},$$

and the classes $[C_1], [C_2]$ are an integral basis for $H_2(X)$. Notice that $[C_2]$ is homologically equivalent to $[\{1\} \times \mathbb{C}\mathrm{P}^1 \times \{1\}]$ which is disjoint from the divisor $p_\infty \times \mathbb{C}\mathrm{P}^1 \times \mathbb{C}\mathrm{P}^1$. This shows that $[D_1]$ as so defined will pair trivially with $[C_2]$. In fact, after putting the basis curves C_1, C_2 into general position, we find that the intersection matrix $(a_{ij}) = ([D_i] \cdot [C_j])$ will be the 2×2 identity matrix.

Heuristically, we can regard the pairing of an element of $H_2(X)$ as counting the intersection number of the homology class with the Poincaré dual of

$$L_\infty = p_\infty \times \mathbb{C}\mathrm{P}^1(b) \times \mathbb{C} \cup \mathbb{C}\mathrm{P}^1(a) \times p_\infty \times \mathbb{C}$$

within X. It follows from [BEH+03] that we can find an almost-complex structure on $X \setminus E$ with cylindrical end on ∂E and equal to the standard product structure near L_∞. Notice that L_∞ is a union of 4-dimensional submanifolds of the 6-manifold X. As described, there will be a well-defined intersection number obtained by pairing L_∞ with a closed, oriented 2-dimensional submanifold. We will occasionally abuse notation and identify a J-holomorphic curve u with its image in X. Notice that since we are in the J-holomorphic category, provided u does not lie completely in L_∞, then the intersection of u with L_∞ will consist of a finite number of points. We may assume L_∞ remains a J-complex submanifold, under suitable deformations of J, so that each intersection with a (deformed) J-holomorphic curve u will contribute a *positive* intersection number.

Definition 4.2.4. Let $u\colon\ \mathbb{C} \to X \setminus E$ be a J-holomorphic plane. Let C denote the image of u with the boundary orbit on ∂E. Express $[C]$ as an element of $H_2(X, E)$ with respect to the basis given by the restrictions of $[C_1], [C_2]$ in the note above. We say that the ordered pair of intersection numbers of $[C]$ with $[D_1]$ and $[D_2]$ is the *bidegree*, (d_1, d_2), of u. These intersection numbers are computed away from E. Alternatively, we can contract E to a point and compute the intersection numbers in $H_2(X)$ of C with contracted boundary orbits.

As above, let $\gamma = \gamma_1$ denote the short Reeb orbit on ∂E. Recall that a finite energy plane can have its end asymptotic to a *cylindrical end* of a symplectic manifold.

Example 4.2.1. The unit disc in $\mathbb{C}$ is a compact symplectic manifold. It admits a compatible almost-complex structure with a cylindrical end which makes it biholomorphic to all of $\mathbb{C}$, a non-compact manifold. This is the reason why we can regard finite-energy J-holomorphic planes as having domain $\mathbb{C}$ in the definition above. Recapitulating, we can either view a cylindrical end as the attachment of a half cylinder to a symplectic manifold, or a manifold with cylindrical end can be obtained by fixing the almost-complex structure near the boundary of the symplectic manifold.

Another example of a manifold with cylindrical end is a *completion* of a compact symplectic manifold with contact boundary. We view the manifold $X \setminus E$ as a symplectic cobordism, and we associate a tame almost-complex structure to endow the boundary with a cylindrical end. (Some authors define a symplectic cobordism to automatically include completions.) The cylindrical end should be compatible with the Liouville contact structure on ∂E. This notion of completion will be explained further in sections 4.3 and 4.6.

Note. At several points during the course of this chapter, we will restrict the class of almost-complex structures on $X \setminus E$, while keeping the symplectic form (4.2.1) fixed. The restriction of the almost-complex structures will be very involved, because we shall perform two *neck-stretching* operations: one in section 4.3 and one in section 4.6. Nonetheless, the results of symplectic field theory, [BEH$^+$03], ensure that we have an ample supply of almost-complex structures to draw upon. Here is the current list of restrictions on such J.

Definition 4.2.5. Let $\mathcal{J}^\star$ denote the list of almost-complex structures on $(X \setminus E, \omega)$ which are

- compatible (and tame) with respect to the symplectic form ω on X defined in (4.2.1),
- cylindrical at ∂E and compatible with the Liouville contact structure on ∂E (c.f. def. 4.2 of [Hut09]),
- equal to a standard product structure outside of a compact set, making L_∞ into a complex submanifold.

Again, this list of conditions will be updated in future sections.

4.3 Stretching the Neck

In this section we describe an important existence theorem of a J-holomorphic curve. This is the curve we shall consider when we stretch the neck along the smoothing Σ of ∂U.

After stretching the neck, the curve limits to a holomorphic building, and we examine the limiting curves that can occur in the building.

Recall that, by assumption, we have, a string of embeddings

$$E(\epsilon - \delta_1, (2d+1)(\epsilon - \delta_2), (2d+1)(\epsilon - \delta_3)) = E \to U \hookrightarrow P(a,b) \times \mathbb{R}^2 \hookrightarrow X.$$

We can now state the main existence theorem of this chapter.

Theorem 4.3.1. In (a completion of) $X \setminus E$ there exist J-holomoprhic planes of bidegree $(d, 1)$ asymptotic to $\gamma_1^{(2d+1)}$. Such curves persist under scaling of the ellipsoid and under a special class of deformations of the almost-complex structure, to be explained in the proof.

The proof of this theorem is long, and will require a technical procedure known as *stretching the neck* along ∂E. In order to not interrupt the flow of the main argument, we have placed the proof of theorem 4.3.1 in section 4.6. The J-holomorphic curve that is produced in theorem 4.3.1 will have area $da + b - (2d+1)(\epsilon - \delta_1)$.

Actually, the main argument will require another neck-stretching procedure, this time along a smoothing Σ of ∂U, which will be described in this section. Note that the special class of deformations of the almost-complex structure that is mentioned in theorem 4.3.1 will allow for this second stretching along Σ. The author hopes that the exposition in this section will make the stretching in section 4.6 easier to understand. It is possible, however, to read section 4.6 before this section. The current state-of-the-art on neck stretching is described in [CM05], but here we are in the simpler case of stretching along a submanifold of contact type.

Definition 4.3.1. We say W is a *contact-type hypersurface* in some symplectic manifold (M^{2n}, ω) if there is a neighborhood of W on which the symplectic form ω is exact, say $\omega = d\lambda$, and the corresponding Liouville vector field v defined by $\omega(v, \cdot) = \lambda(\cdot)$ is transverse (not tangent) to W.

Lemma 4.3.1. In this context, the primitive $\lambda|_W$ is a contact form on W. It induces a contact structure $\xi = \{\lambda = 0\}$, which is a $(2n-2)$-hyperplane distribution on W.

Proof. We must show that $\lambda \wedge (d\lambda)^{n-1} = \lambda \wedge \omega^{n-1}$ is a volume form on W. Now, $ker(\omega|_W)$ is non-trivial, because W is odd-dimensional. Further $ker(\omega|_W) \subset TW^{\perp\omega}$, which is 1-dimensional. We therefore have that $dim(ker(\omega|_W)) = 1$. Choose a non-zero vector field $Y_1 \in ker(\omega|_W)$, and complete this to a frame $\{Y_1, Y_2, ..., Y_{2n-1}\}$ for TW. We evaluate $\lambda \wedge \omega^{n-1}$ on this frame to get a single non-zero term

$$\lambda|_W(Y_1)\omega|_W(Y_2, ..., Y_{2n-1}) = \omega(v, Y_1)\omega(Y_2, ..., Y_{2n-1}) \neq 0.$$

□

Now that we understand the setup, the contact-type hypersurface we consider is $W = \Sigma$, a smoothing of ∂U. Σ separates X into two components, which by abuse of notation we call U and $X \setminus U$, (i.e. when abusing notation we ignore the smoothing that defines Σ). Let v denote the Liouville vector field for Σ. An almost complex structure J on X is said to be *compatible* with the stretching if the contact structure $\xi = \{\lambda = 0\}$ on Σ coincides with

$$T(\Sigma) \cap JT(\Sigma)$$

and if $J(v)$ equals the Reeb vector field of λ. We update the definition of $\mathcal{J}^*$ by requiring the almost-complex structures to be compatible with the stretching along Σ.

Now, for every $N \in \mathbb{N}$, the three stretched manifolds

$$\overline{A_N} := (U \setminus E,\, e^{-N}\omega), \qquad \overline{B_N} := (\Sigma \times [-N, N],\, d(e^t\lambda)), \qquad \overline{C_N} := (X \setminus U,\, e^N\omega) \tag{4.3.1}$$

can be glued along their boundary components to obtain a symplectic manifold, which we denote by $(\overline{X^N}, \omega^N)$. Here t denotes the real coordinate on $[-N, N]$. Choose $J \in \mathcal{J}^*$, and let J^N be the continuous almost-complex structure on $\overline{X^N}$ that equals J on the pieces $\overline{A_N}$ and $\overline{C_N}$ and is translation invariant on piece $\overline{B_N}$. Some perturbation may be necessary to smooth this choice of almost-complex structure, but this will not affect the results below. (See [BEH+03] for more details on this smoothing.) So we may as well assume that all J^N are smooth.

After all this setup, we have that the Reeb orbits on Σ are in 2-dimensional families indexed by triples $(k, \ell, m) \in \mathbb{Z}^3 \setminus \vec{0}$, as described in the remark on page 64. Specifically, k counts the winding in the (x_1, y_1) factor, and so on. The Reeb orbits along Σ may appear in a homology class with $k, \ell, m < 0$ (as opposed to ∂E where winding along the short orbit is described by a non-negative integer). After stretching the neck, the limiting curve (as $N \to \infty$) is a holomorphic building, i.e. a collection of finite energy holomorphic curves in completions of $\overline{C_\infty} := X \setminus U$ and $\overline{A_\infty} := U \setminus E$ with matching asymptotic limits along Σ. We shall name these completions in the compactness theorem below, with less cumbersome notation.

Theorem 4.3.2. (See 10.6 of [BEH+03]) Fix a J in $\mathcal{J}^*$. For each $N \in \mathbb{N}$, let u_N be a J^N-holomorphic curve negatively asymptotic to $\gamma_1^{(2d+1)}$. Fix a representative f_N for each u_N. Then there exists a subsequence of the f_N which converges to a holomorphic building $\mathbf{F}$. The domain of $\mathbf{F}$ is a nodal Riemann sphere (S, j) with punctures, and the building can be described as a collection of finite energy holomorphic maps from the collection of

punctured spheres $S \setminus \{\text{nodes}\}$ into one of the following three completions:

$$A_\infty := (U \setminus E) \cup_\Sigma (\Sigma \times [0, \infty)) \qquad \text{with form } d(e^t\lambda),\ t \in [0, \infty),$$

$$B_\infty := B_\infty = \Sigma \times \mathbb{R}, \qquad \text{with form } d(e^t\lambda),\ t \in \mathbb{R},$$

or

$$C_\infty := (X \setminus U) \cup_\Sigma (\Sigma \times (-\infty, 0]) \qquad \text{with form } d(e^t\lambda),\ t \in (-\infty, 0].$$

We now gather some facts about the completed manifold C_∞. Notice that C_∞ is diffeomorphic to $X \setminus U$, but C_∞ has a specific almost-complex structure near the boundary. This is an example of a manifold with cylindrical end. Consequently J-holomorphic curves into C_∞ can have only negative asymptotic ends. Each curve of $\mathbf{F}$ having target in C_∞ will be asymptotically cylindrical near the punctures in its domain, hence will wind around the Reeb orbits of Σ as $t \to -\infty$. We see that this neck-stretching construction sends all removable singularities of a curve to the Reeb orbit "at $-\infty$". As mentioned above, we can parameterize the Reeb orbits on Σ by three integers. The area of a curve u is defined using the area form coming from the completion of $X \setminus U$, not the re-scaled form on C_N above.

The holomorphic building $\mathbf{F}$ will have a *level structure*, which we now explain. To encode a building $\mathbf{F}$ of height l, we label all of the punctured Riemann surfaces (S, j) by integers $0, 1, ..., l+1$. These labels are called levels. The difference between the levels of any two components of the building that share a node must be one. In this way we think of the levels of the building as heights, and we view curves in the building as glued along nodes to curves one level above or below. An illustration of such a building is given in figure 4.2. Let S_κ denote the union of components of level κ and denote by u_κ the holomorphic curve of $\mathbf{F}$ with domain S_κ. The domain S_κ could be disconnected if bubbling occurs during the neck stretching, but we will not examine bubbling yet. In any case, we have that $u_0 \colon S_0 \to A_\infty$ and $u_\kappa \colon S_\kappa \to B_\infty$ for $1 \leq \kappa \leq l$ and $u_{l+1} \colon S_{l+1} \to C_\infty$. Each node that is shared by S_κ and $S_{\kappa+1}$ is a positive puncture for u_κ and a negative puncture for $u_{\kappa+1}$, and each are asymptotic to the same Reeb orbit. A very important property of the definition from [BEH+03] is that we assume none of the curves u_κ for $1 \leq \kappa \leq l$ consist entirely of trivial cylinders over Reeb orbits, except in the following situation. If a collection of J-holomorphic planes converges to a limiting plane v, then after stretching the neck, this sequence of planes will converge (in the sense of SFT) to a building consisting of v in the highest level together with a trivial cylinder in the symplectization over its limiting orbit. Since we are stretching the J-holomorphic plane that was produced in theorem 4.3.1, we are in this exceptional situation. Notationally, we will not emphasize the number l in the discussion that follows.

Note that part B_∞ is a symplectization and parts A_∞ and C_∞ are symplectic cobordisms (with completions). To wit, a symplectic cobordism is a symplectic manifold M with boundary $\partial M = M^- \sqcup M^+$ such that M^+ has an outward-pointing Liouville vector field and M^- has an inward-pointing Liouville vector field. These Liouville fields define contact forms on $M^- \sqcup M^+$, which give rise to Reeb orbits on ∂M. It is important to note that the Reeb orbits on ∂M come in Morse-Bott families. For this reason, we use the generalized Conley-Zehnder index, as defined in section 4.1.1. Furthermore, we will need to keep track of the dimensions of the families of the Reeb orbits that appear on ∂M. We need to investigate what limiting curves arise in the holomorphic building that we get after stretching the neck along Σ. This investigation will require a virtual index formula for a J-holomorphic curve u in a symplectic cobordism or a symplectization. There is an important quantity that we associate to a curve u, known as the *relative first Chern class*. In parts A_∞ and B_∞ the relative Chern class will turn out to be zero. We examine C_∞ here. Specifically, the symplectic cobordism under consideration is a null-cobordism ($\partial(X \setminus U) \cong \Sigma \sqcup \emptyset$) with completed end $\Sigma \times (-\infty, 0]$. For the full definition of the relative first Chern class in a symplectic cobordism, see §4.2 of [Hut09] (noting that $n = 2$ there). The definition depends on a choice of trivialization τ of the contact structure along the Reeb orbits that are the ends of the J-holomorphic curve u. In our case, curves in C_∞ can have only negative ends that wind along Reeb orbits of Σ. Here we choose a symplectic trivialization of u^*TC_∞, which we deform near the punctures Γ of u so that it agrees with the trivialization on ∂E or Σ that we used to define the Conley-Zehnder index in section 4.1.1 and lemma 4.2.1 (the trivialization coming from inclusion into $\mathbb{C}^3$). Then we define $c_1(u)$ to be the algebraic count of of zeros of a generic section of $\Lambda^n(u^*TC_\infty)$ which is constant with respect to the trivialization on the boundary. The relative first Chern class of a curve u in C_∞ will be denoted $c_\tau(u)$. If u has image C in C_∞ and has bidegree (d_1, d_2), then with this choice of trivialization,

$$c_\tau(u) := c_1(T(C_\infty)|_C, \tau) = 2d_1 + 2d_2.$$

As with the bidegree, the computation of c_τ is designed to ignore the winding of the curve u along its negative ends, where $\Lambda^n(u^*TC_\infty)$ is trivial. Note that the Conley-Zehnder index and the relative first Chern class each depend on the choice of trivialization τ, but the index formula that follows does not depend on the trivialization.

The general index formula for finite-energy genus zero curves asymptotic to Reeb orbits

(taken from [Bou02] Corollary 5.4) is

$$Index(u) = (n-3)(2-s^+-s^-) + 2c_\tau(u) + \sum_{i=1}^{s^+} CZ(\gamma_i^+) + \frac{1}{2}\sum_{i=1}^{s^+}\dim(\gamma_i^+) - \sum_{j=1}^{s^-} CZ(\gamma_j^-) + \frac{1}{2}\sum_{j=1}^{s^-}\dim(\gamma_j^-),$$

where

- n is half the dimension of the ambient space X (usually $n = 3$ in this chapter),
- s^- is the number of negative ends,
- s^+ is the number of positive ends (which will usually be zero in this chapter),
- $CZ(\gamma_j^-)$ is the generalized Conley-Zehnder index of the j^{th} negative end, and
- $\dim(\gamma_i^-)$ counts the dimension of the family of Reeb orbits at each negative end of u.

This index formula predicts the dimension of the moduli space of J-holomorphic curves asymptotic to prescribed Reeb orbits. Henceforth we shall call this Fredholm index the *virtual index*. With the above assumptions (finite energy, genus zero, only negative ends), this formula simplifies to

$$Index(u) = (n-3)(2-s^-) + 4d_1 + 4d_2 - \sum_{j=1}^{s^-}\left[CZ(\gamma_j^-) - \frac{1}{2}\dim(\gamma_j^-)\right]. \tag{4.3.2}$$

More specific instances of this index formula will appear in the arguments that follow.

The virtual index of a *closed* curve of genus 0 is given by the formula

$$Index(u) = (n-3)(2) + 2c_\tau(u). \tag{4.3.3}$$

The closed curve formula is a special case of the above, which can be found in §2.2 of [Hut14], among other sources.

Genericity of J is what allows the dimension of a moduli space of J-holomorphic curves to be predicted by the Fredholm index of a representative. To that end, we need the following lemma.

Lemma 4.3.2. For generic J, any J-holomorphic curve with image in piece C_∞ will have non-negative virtual Fredholm index.

Proof. A J-holomorphic curve $u\colon S^2\setminus\Gamma \to C_\infty$ mapping into C_∞ must have only negative ends winding about some Reeb orbit in Σ. If u is somewhere injective, then it will have non-negative virtual Fredholm index, by genericity of J. Otherwise, u multiply covers a somewhere injective underlying curve, which we call $\tilde{u}\colon S^2\setminus\tilde{\Gamma}\to C_\infty$. Let us say that u is a p-fold covering of $\tilde{u}$. Then there is a ramified covering map $\psi\colon S^2\setminus\Gamma\to S^2\setminus\tilde{\Gamma}$ such that $u=\tilde{u}\circ\psi$. Since these singularities are removable, we can extend ψ to a holomorphic map $\Psi\colon S^2\to S^2$ sending $\Gamma\to\tilde{\Gamma}$. We write s^- for the number of negative ends of u and $\tilde{s}^-$ for the number of negative ends of $\tilde{u}$. Then $p\tilde{s}^- - s^-$ will be the total ramification of ψ over Γ. By the Riemann-Hurwitz formula,

$$p\tilde{s}^- - s^- = \sum_{x\in\Gamma}(m_x-1) \leq \sum_{x\in S^2}(m_x-1) = 2p-2. \tag{4.3.4}$$

We compute the virtual Fredholm index of both u and $\tilde{u}$, using the general index formula. As noted above, a negative end of either curve can be characterized by its homology class, so we write (k_j,ℓ_j,m_j) for the homology class of the j^{th} negative end of u and we write $(\tilde{k}_j,\tilde{\ell}_j,\tilde{m}_j)$ for the j^{th} negative end of $\tilde{u}$. The family of such orbits along Σ has dimension 2. Both curves have a bidegree, and we write (d_1,d_2) for the bidegree of u, and we write $(\tilde{d}_1,\tilde{d}_2)$ for the bidegree of $\tilde{u}$. Now the index formula gives

$$Index(\tilde{u}) = (n-3)(2-\tilde{s}^-) + 4(\tilde{d}_1+\tilde{d}_2) - 2\sum_{j=1}^{\tilde{s}^-}\tilde{k}_j+\tilde{\ell}_j+\tilde{m}_j, \tag{4.3.5}$$

which is non-negative by assumption. Similarly, we find

$$\begin{aligned}
Index(u) &= (n-3)(2-s^-) + 4(d_1+d_2) - 2\sum_{j=1}^{s^-} k_j+\ell_j+m_j \\
&= (3-n)(s^--2) + 4p(\tilde{d}_1+\tilde{d}_2) - 2p\sum_{j=1}^{\tilde{s}^-}\tilde{k}_j+\tilde{\ell}_j+\tilde{m}_j \\
&\geq (3-n)(p\tilde{s}^--2p) + 4p(\tilde{d}_1+\tilde{d}_2) - 2p\sum_{j=1}^{\tilde{s}^-}\tilde{k}_j+\tilde{\ell}_j+\tilde{m}_j \\
&= p\cdot index(\tilde{u}) \geq 0
\end{aligned}$$

using the inequality (4.3.4). □

Lemma 4.3.3. For generic J, any trivial cylinder with image in piece B_∞ will have index $(n-1)$.

Proof. Recall that, by our choice of trivialization, the relative first Chern class term of the index formula vanishes in B_∞. A trivial cylinder is a simple curve in the symplectization layer B_∞ which is topologically $\gamma \times \mathbb{R}$ for some Reeb orbit γ on Σ. Hence γ is both the positive asymptotic end and the negative asymptotic end of such a trivial cylinder. Obviously, the Conley-Zehnder terms in the index formula will cancel out for these two ends, being equal and opposite in sign. This implies that the index of such a trivial cylinder reduces to

$$Index = (n-3)(2-1-1) + (1/2+1/2)(n-1) = n-1$$

□

The following result tells us what limiting curves can appear in the holomorphic building **F** that results from stretching the neck. We make use of the notion of matching curves into components in the pseudoholomorphic building. A matched component is given by formally gluing together some collection of curve components lying in various levels of the building, with identifications being made along paired ends. We treat the domain of such a component as a smooth, connected, punctured Riemann surface. We define the positive ends of a matched component to be any positive ends of the constituent curves which have not already been matched to other curves in the sub-building. Similarly, we define the negative ends of a matched component to be any negative ends of constituent curves that have not already been matched. We are going to need to compute the virtual Fredholm index of these matched components, so we explain here how the Fredholm index behaves under the matching procedure. If a sub-building B is obtained by matching curves $u_1, ..., u_p$ along asymptotic orbits $\gamma_1, ..., \gamma_q$ belonging to families of Reeb orbits in spaces S_j, then

$$Index(B) = \sum_{i=1}^{p} Index(u_i) - \sum_{j=1}^{q} \dim(S_j). \tag{4.3.6}$$

In a similar way, if we were to form a sub-building B by matching other sub-buildings B_i along matching orbits γ_i that belong to families of Reeb orbits in spaces S_i, then we have

$$Index(B) = \sum Index(B_i) - \sum \dim(S_i).$$

Finally, after all these identifications are made, we can compute the index of a single sub-building B by only considering its un-matched positive and negative ends and using the general index formula for a single curve.

Theorem 4.3.3. After stretching the neck along Σ, limiting curves can be matched into components consisting of

- some sub-buildings whose matching ends can be identified to form a plane with a negative end asymptotic to a Reeb orbits on Σ; and
- one finite energy component with positive ends on Σ and with a single negative end asymptotic to γ^{2d+1} on ∂E. Let us call this one limiting component the *special curve*.

Furthermore, the sub-buildings mentioned in this list all have non-negative building index.

Proof. As mentioned above, the curves described in the two bullet points of the theorem are (potentially) several curves identified into connected components, according to the following identifications.

There must be a unique curve in the lowest level of the holomorphic building, because there is only one negative end of the pre-stretched curve, asymptotic to γ_1 by construction. Let us say that the limiting curve in this lowest level is named u_0. Next, we shall identify with u_0 some curves in the symplectization layer, based on matching conditions, and we shall consider these identified curves to be a single component.

Beginning with u_0, locate *all* curves that can be connected to u_0 through a chain of curves with matching ends in B_∞. Identify these curves along their matchings ends, and call the resulting component v_0. Because the negative end of v_0 is uniquely specified by u_0, we know that v_0 can have only positive ends that remain un-matched. Moreover, since the identification glues all matching ends in B_∞, we see that the positive end(s) of v_0 must be asymptotic to Reeb orbits on Σ, the positive end of piece B_∞. Consider the complement of v_0 in the holomorphic building. Let us say that, after identifying matching ends, this complement has exactly M connected components. Name these components $v_1, ..., v_M$. Furthermore, we assume that v_0 is matched as far as possible with curves in B_∞, so that the components v_j for $j \in \{1, ..., M\}$ do not contain any curves that could be matched with v_0.

The curves $v_1, ..., v_M$, described here, exist in the highest level(s) of the holomorphic building **F** from theorem 4.3.2, and the negative ends of these curves must be asymptotic to a Reeb orbit on Σ. Notice that in the gluing construction of the previous paragraph, we are identifying curves in the manifold B_∞ with curves in the manifold A_∞ or C_∞ even though these manifolds are technically disjoint. We make this identification, nonetheless, because the positive cylindrical end of A_∞ (as $t \to \infty$) is equal to the negative cylindrical end of B_∞ (as $t \to -\infty$). Moreover, if we glue the positive ends of v_0 with the negative ends of all the curves $v_1, ..., v_M$, then we recover the original curve that was constructed in theorem 4.3.1. This original curve had genus zero, which forces the identified curve $v_0 \cup ... \cup v_M$

to have genus zero. Consequently, each of $v_0, ..., v_M$ must be genus zero as a matched component.

In proposition 4.6.5 of section 4.6, we explain why a curve in the symplectization layer B_∞ must be a trivial cylinder. Moreover, there can exist no J-holomorpic planes in part A_∞, because the boundary of any such plane would have positive end asymptotic to Σ, and there are no contractible Reeb orbits on Σ. These two facts imply that, after the identifications above, the curves $v_1, ..., v_M$ are not allowed to "turn around", re-enter the symplectization layer, and terminate in A_∞. This fact helps the reader to visualize the curves in C_∞. The key point is that no matter how many times a curve $v_j \in \{v_1, ..., v_M\}$ enters or exits the symplectization layer, it won't affect the Fredholm index. Specifically, we proved in lemma 4.3.2 that the ends of the components $v_1, ..., v_M$ that lie in C_∞ all have non-negative index. These ends may be identified with some trivial cylinders in part B_∞ to form the sub-buildings $v_1, ..., v_M$. We showed in lemma 4.3.3 that these trivial cylinders have index $(n-1)$, which is exactly the dimension of the family of Reeb orbits on Σ, where the matching takes place. Hence, this added index is exactly canceled by the dimension of the matching term in formula (4.3.6). In effect, we need only consider only the negative end and the bidegree of each of $v_1, ..., v_M$ when computing its Fredholm (building) index, and we can ignore the contribution from trivial cylinders. From this analysis we conclude that each of $v_1, ..., v_M$ have non-negative building index.

The curves $v_1, ..., v_M$ have disjoint domains, by construction. In addition, the v_j have finite energy. This discussion leaves, topologically, only one possibility for the curves $v_1, ..., v_M$. Each must be a J-holomorphic plane whose negative end matches with one of the positive ends of v_0. This also leaves only one possibility for v_0. Topologically, v_0 is a pair of pants with M pant legs. Figure 4.2 gives an illustration of these curves.

The curves $v_1, ..., v_M$ within C_∞ must collectively have negative ends that bound a cycle in U (since they are the boundary of v_0 plus a disc in E). If we denote the homology class of the negative end of the curve $v_j \in \{v_1, ..., v_M\}$ by the ordered triple (k_j, ℓ_j, m_j), as described in lemma 4.2.1, then we must have

$$\sum_{j=1}^{M} k_j = 0, \tag{4.3.7}$$

and

$$\sum_{j=1}^{M} \ell_j = 0, \tag{4.3.8}$$

and similarly for the m_j. The negative ends of $v_1, ..., v_M$ must match up with all of the

positive ends of v_0. The preceding calculation gives

$$\sum_{j=1}^{M} k_j + \ell_j + m_j = 0, \tag{4.3.9}$$

where here the sum is taken over all positive ends of v_0. The extra terms in the Conley-Zehnder index formula for any positive end of v_0, coming from the symplectic shears in lemma 4.2.1 will be $\frac{n-1}{2}$. The dimension of the family of Reeb orbits at any positive end of v_0 is $(n-1)$. Once again, the special curve has one negative end asymptotic to $\gamma_1^{(2d+1)}$ and M positive ends on Σ. The first Chern class vanishes for all constituent curves of v_0, hence also for v_0. For the same reason as above, we can ignore the index contribution of any trivial cylinders that were identified to u_0 to form v_0. In all, the general index formula gives

$$\begin{aligned}
Index(v_0) = Index(u_0) &= (n-3)(2-M-1) + 0 \\
&\quad + \sum_{j=1}^{M} CZ(\gamma_j^+) + \frac{1}{2}\sum_{j=1}^{M}(n-1) - [2(2d+1)+n-1] + 0 \\
&= (n-3)(1-M) - 2(2d+1) - (n-1) \\
&\quad + \sum_{j=1}^{M}\left(k_j + \ell_j + m_j + \frac{n-1}{2}\right) + \frac{M}{2}(n-1) \\
&= (n-3) - (n-1) - M(n-3) + M(n-1) - 2(2d+1) \\
&= 2(M-1) - 2(2d+1)
\end{aligned}$$

after incorporating formula (4.3.9). We see that v_0 will have non-negative Fredholm index precisely when $M \geq 2d+2$, which is true by positivity of area of v_0. In a future proof (of theorem 4.4.1) we shall explicitly compute that v_0 is of index zero. □

Lemma 4.3.4. The identified curves mentioned in the bulleted list of theorem 4.3.3 must be asymptotic to a Reeb orbit on Σ with $m = 0$.

Proof. Because Σ in X is of contact type, the symplectic form inherited from X will be exact on a tubular neighborhood of ∂U. In fact, let us identify U with a subset of T^*T^3, as above, so that the exact form on ∂U has a primitive whose integral over a Reeb orbit of Σ in homology class (k, ℓ, m) is given by the following formula (in an arbitrarily large range,

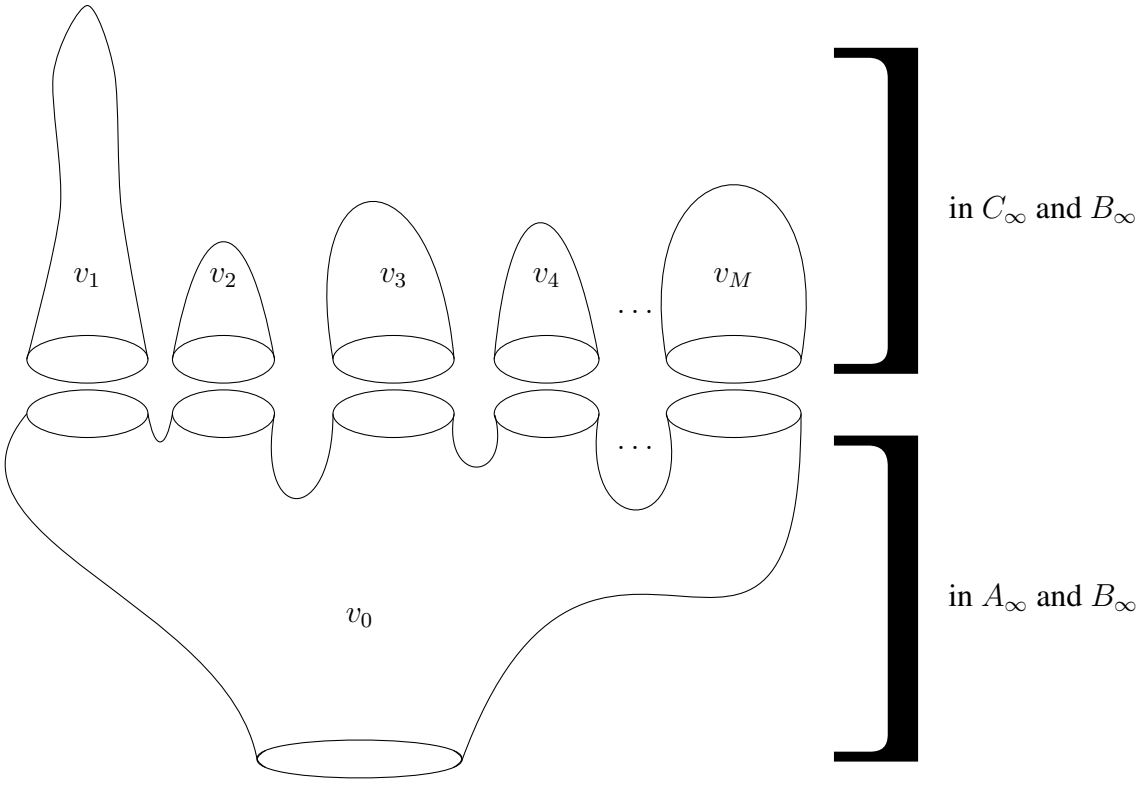

Figure 4.2: The curves $v_0, ..., v_M$.

and up to a small correction due to the smoothing Σ).

$$\mathcal{A} = \frac{\epsilon}{2}|k| + \frac{(2d+1)\epsilon}{2}|\ell| + \frac{S}{4}|m|.$$

This action formula will factor into the formulas for the area of curves in C_∞ and A_∞. Because our original holomorphic plane had area at most $da + b << S$, we see that the special curve in A_∞ cannot have any positive ends asymptotic to Reeb orbits with $m \neq 0$. By the matching conditions, the negative ends of limiting curves in C_∞ also must be asymptotic to Reeb orbits in class $(k, \ell, 0)$.

□

4.4 Analysis of Limiting Curves in C_∞

Now that we understand the procedure of stretching the neck, along with the limiting curves that can result, let us focus on a limiting curve in the part C_∞. In this section, let u denote any of the curves $v_1, ..., v_M$ from the proof of theorem 4.3.3. We showed that such a curve u will in fact be a J-holomorphic plane asymptotic to a single hyperbolic Reeb orbit on its

negative end. Using lemma 4.3.4, the index formula (4.3.2) becomes

$$Index(u) = 0 + 2c_\tau(u) - CZ(k,\ell) + \frac{1}{2}(2) = 4d_1 + 4d_2 - 2(k+\ell), \tag{4.4.1}$$

where c_τ denotes the relative first Chern class of u and CZ denotes the Conley-Zehnder index of the Reeb orbit on the negative end of u. In particular

$$Index(u) = 4d + 4 - 2k - 2\ell \tag{4.4.2}$$

in the case when the curve u has bidegree $(d, 1)$.

Lemma 4.4.1. All limiting planes $u \in \{v_1, ..., v_M\}$ must have virtual index 0 or 2.

Proof. First, the formula (4.4.1) shows that the virtual index must be an even integer. Second, theorem 4.3.3 gives that the curves $v_0, ..., v_M$ must have non-negative index. We noted in the previous section that the Reeb orbits on Σ come in 2-dimensional families, and the limiting curves $v_1, ..., v_M$ are asymptotic to these orbits at the negative end. If the virtual index of u were to exceed the dimension of the family of asymptotic orbits, then other curves in the holomorphic building would necessarily have negative index, which is also precluded by theorem 4.3.3. This leaves only the possibilities of 0 or 2 for the virtual index. □

The following two corollaries (4.4.1 and 4.4.2) and lemma 4.4.2 apply to the identified components $v_1, ..., v_M$ in C_∞.

Corollary 4.4.1. Any limiting component in homology class $(k, 0, 0)$ for some negative integer k, must actually have $k = -1$ and the virtual index must be 2 and the bidegree must satisfy $d_1 = d_2 = 0$

Proof. Assume $\ell = 0$ and $k < 0$. Then the index formula (4.4.1) simplifies to

$$4(d_1 + d_2) + 2|k| = 0 \text{ or } 2,$$

using lemma 4.4.1. Now any positive bidegree $(d_1, d_2) \neq (0, 0)$ would contribute a multiple of 4 to the index and contradict the above equation. So we must have $d_1 = d_2 = 0$. Consequently k must be -1 and the index must be 2.

□

Corollary 4.4.2. Any limiting curve in homology class $(k, 0, 0)$ for some positive integer k must have nonzero bidegree, meaning $(d_1, d_2) \neq (0, 0)$

Proof. The argument is the same as for corollary 4.4.1

□

Lemma 4.4.2. Assume that $x > 2$, $n = 3$, and $a < 2$. If $b < x$, then $\ell \leq 0$ in $\mathbb{Z}$.

Proof. It will only be necessary to prove this assertion in the cases when the bidegrees of u satisfy $(d_1, d_2) = (0, 0)$ or $(d_1, d_2) = (d, 1)$ or $(d_1, d_2) = (d, 0)$ for some $d > 0$ in $\mathbb{Z}$. The following analysis will apply to all cases.

We note that up to a correction of order ϵ, the area of a curve u of bidegree (d_1, d_2) is given by the formula

$$Area(u) = d_1 a + d_2 b - (1k + x\ell) \geq 0, \tag{4.4.3}$$

as all J-holomorphic curves have non-negative area. We subtract from this area inequality one-half of the virtual index above to find

$$d_1(a-2) + d_2(b-2) - \ell(x-1) \geq -1 \tag{4.4.4}$$

which will be the main inequality for this proof. Here we have used Lemma 4.4.1 to restrict the virtual index.

Case 1: If $(d_1, d_2) = (0, 0)$, then the main inequality reduces to

$$-\ell(x-1) \geq -1 \qquad \Longrightarrow \qquad \ell \leq 1/(x-1).$$

Since we have have that $x - 1 > 1$, the integrality of ℓ forces $\ell \leq 0$.

Case 2: If $(d_1, d_2) = (d, 1)$, then the main inequality (4.4.4) reduces to

$$d(a-2) + (b-2) - \ell(x-1) \geq -1 \qquad \Longrightarrow \qquad \underbrace{d(a-2)}_{\text{negative quantity}} + \underbrace{(b-1)}_{\text{small quantity}} \geq \ell \cdot \underbrace{(x-1)}_{\text{positive}}$$

Since $b < x$ we have $\frac{b-1}{x-1} < 1$. Hence

$$\ell \leq \left\lfloor \frac{d(a-2)}{(x-1)} + \frac{(b-1)}{(x-1)} \right\rfloor \leq 0.$$

Repeat this proof, replacing $d_2 = 1$ with $d_2 = 0$, to see that the only remaining case is easier than case 2. □

Finally, we exhibit precisely which finite energy planes occur as limits in C_∞.

Theorem 4.4.1. Assume $a < 2$ and $b < x$. The limiting curves in C_∞ consist of

- A single plane of bidegree $(d, 1)$ which is negatively asymptotic to a Reeb orbit in Σ that is in homology class $(2d + 1, 0, 0)$ and which has virtual index 2;
- A collection of $(2d + 1)$-many finite energy planes of bidegree $(0, 0)$, each of which is negatively asymptotic to a Reeb orbit in Σ that is in homology class $(-1, 0, 0)$. Moreover, these planes have virtual index 2.

Proof. The proof of this theorem will use the notation of theorem 4.3.3, and will rely on the lemmas and corollaries above. Consider the J-holomorphic planes $v_1, ..., v_M$ of theorem 4.3.3. We label the homology class of the negative end of the plane v_j by the ordered triple (k_j, ℓ_j, m_j), for $1 \leq j \leq M$. Similarly, we label the bidegree of the plane v_j by the ordered pair (d_1^j, d_2^j), although some of these degrees may be zero. By definition, the bidegree is a bilinear pairing, and we must have

$$\sum_{j=1}^{M} (d_1^j, d_2^j) = (d, 1), \tag{4.4.5}$$

which is the bidegree of the original curve, before stretching the neck. Since each $d_i^j \geq 0$, we see that the only possibilities for these bidegrees are $(0, 0)$ or $(\cdot, 0)$ or $(*, 1)$ for some non-negative integers $\cdot$ and $*$. This shows that we are in the restricted cases that were considered in the proof of lemma 4.4.2, and the conclusions of that lemma apply here. Combining this result with the equation (4.3.8) implies that $\ell_j = 0$ for $1 \leq j \leq M$. We also know that all $m_j = 0$ by lemma 4.3.4. Hence we may label the homology classes of the negative ends of the planes $\{v_j\}$ by the ordered triple $(k_j, 0, 0)$, for $1 \leq j \leq M$. Notice that, by construction, no $k_j = 0$, because we do not consider the homology class $(0, 0, 0)$ to describe a valid Reeb orbit. We do have a condition (4.3.7) on the k_j, which will allow us now re-index the curves $\{v_1, ..., v_M\}$ so that

- the subset $\{v_1, ..., v_t\}$ consists of curves v_j with $k_j > 0$, and
- the subset $\{v_{t+1}, ..., v_M\}$ consists of curves v_j with $k_j < 0$.

The remainder of this proof will show that these two bullet points correspond to the bullet points in the statement of the theorem. We immediately notice that corollary 4.4.1 applies to the set $\{v_{t+1}, ..., v_M\}$, implying that $k_j = -1$ for these curves, and the bidegrees vanish. Since we have organized the curves according to the sign of the integer k_j, we may decompose the summation (4.3.7) into a more subtle result. We have

$$0 = \sum_{j=1}^{t} k_j - \sum_{j=t+1}^{M} |k_j| \Longrightarrow \sum_{j=1}^{t} k_j = \sum_{j=t+1}^{M} |k_j| = \sum_{j=t+1}^{M} |-1| = M - t. \tag{4.4.6}$$

This calculation also implies that

$$\sum_{j=1}^{M} |k_j| = 2\sum_{j=1}^{t} k_j = 2(M-t). \tag{4.4.7}$$

By definition, the bidegree is a bilinear pairing. The vanishing of certain bidegrees further implies that

$$\sum_{j=1}^{t} (d_1^j, d_2^j) = (d, 1), \tag{4.4.8}$$

which is slightly sharper than (4.4.5). We re-index yet again so that v_1 is the curve with $d_2^1 = 1$, and all other $v_j \in \{v_2, ..., v_t, v_{t+1}, ..., v_M\}$ have $d_2^j = 0$. Corollary 4.4.2 shows that $v_j \in \{v_2, .., v_t\}$ cannot also have $d_1^j = 0$. In other words, all the curves in the set $\{v_1, ..., v_t\}$ have some nonzero (strictly positive) component of bidegree.

On the other hand, we have the curve v_0 with M positive ends and a single negative end asymptotic to the $(2d+1)$-times cover of the Reeb orbit γ_1. The area of v_0 is given by Stokes' Theorem.

$$Area(v_0) = \frac{\epsilon}{2}\sum_{j=1}^{M} |k_j| - (2d+1)(\epsilon - \delta_1).$$

Using formula (4.4.7) and the positivity of area, we get

$$M - t = \frac{1}{2}\sum_{j=1}^{M} |k_j| > (2d+1)\left(\frac{\epsilon - \delta_1}{\epsilon}\right).$$

Since δ_1 may be chosen to be much smaller than ϵ, and since $M - t$ is integral, we find that $M - t \geq 2d + 1$. The curves $v_1, ..., v_t$ have virtual Fredholm index given by formula (4.4.2). Suppose, towards a contradiction, that the sum of all the Fredholm indices of these curves were zero. Then, by non-negativity, each curve must have Fredholm index zero. Since curve v_j has bidegree (d_1^j, d_2^j) and has negative end in class $(k_j, 0, 0)$, we apply the index formula (4.4.1) to find

$$0 = \frac{1}{2}Index(v_j) = 2d_1^j + 2d_2^j - k_j \Longrightarrow k_j = 2d_1^j + 2d_2^j.$$

In a similar way, we can sum the areas of these curves to find

$$\sum_{j=1}^{t} Area(v_j) = da + b - \sum_{j=1}^{t} k_j = da + b - 2d - 2 = d(a-2) + (b-2).$$

The assumption $a < 2$ makes the term $d(a-2)$ negative, and we are free to choose d large enough to ensure that $d(a-2)+(b-2) < 0$, contradicting positivity of area. We must conclude that not all the curves $v_1, ..., v_t$ have Fredholm index zero. In other words,

$$0 < \sum_{j=1}^{t} Index(v_j) = 4d + 4 - 2\sum_{j=1}^{t} k_j = 4d + 4 - 2(M-t).$$

Consequently, $M - t < 2d + 2$. We have now shown

$$2d + 1 \leq M - t < 2d + 2.$$

By integrality, we must have $M - t = 2d + 1$. Substituting this result into the Index computation above gives

$$\sum_{j=1}^{t} Index(v_j) = 4d + 4 - 2(2d+1) = 2. \tag{4.4.9}$$

By lemma 4.4.1, the indices of $v_1, ..., v_t$ can be either 0 or 2. This leaves only one possibility: a single curve among the collection has index 2, and the rest have index zero.

We claim that v_1 is the curve with virtual Fredholm index 2. The proof will be very similar to the argument in the preceding paragraph. Assume, towards a contradiction, that $Index(v_1) = 0$. Then

$$0 = \frac{1}{2} Index(v_1) = 2d_1^1 + 2(1) - k_1 \Longrightarrow k_1 = 2d_1^1 + 2.$$

Substituting this into the area equation of this curve gives

$$Area(v_1) = d_1^1 a + 1b - k_1 = d_1^1(a-2) + (b-2) \tag{4.4.10}$$

The area of the remaining curves must therefore be

$$\sum_{j=2}^{t} Area(v_j) = (d - d_1^1)a - \sum_{j=2}^{t} k_j = (d-d_1^1)a - 2d - 1 + 2d_1^1 + 2 = (d-d_1^1)(a-2) + 1 \tag{4.4.11}$$

Now, we compare equations (4.4.10) and (4.4.11). We showed previously that we may choose d large enough so that $d(a-2)+(b-2) < 0$. But since d_1^1 is a summand of d, increasing d might increase d_1^1, which would make equation (4.4.10) negative; or increasing d might not increase d_1^1 in which case equation (4.4.11) becomes negative. Either case will

contradict the positivity of area. We must conclude that $Index(v_1) = 2$.

Next we show that, in fact, $t = 1$. Assume, towards a contradiction, that $t \geq 2$. Then all of the curves $\{v_2, ..., v_t\}$ must have index zero, because of equation (4.4.9). Consequently

$$0 = \frac{1}{2} Index(v_j) = 2d_1^j - k_j \Longrightarrow k_j = 2d_1^j \qquad \text{for } 2 \leq j \leq t.$$

But any such curve v_j, $2 \leq j \leq t$ will have

$$Area(v_j) = d_1^j a - k_j < 2d_1^j - 2d_1^j = 0,$$

since $a < 2$. Again, this contradicts the positivity of area, and we must conclude that $t = 1$.

This plane v_1 is the only J-curve to touch the divisor L_∞, and it must have bidegree $(d, 1)$, by construction. It must also have $k_1 = M - t = 2d + 1$, by above work. This proves the first bullet point of the theorem statement. Since there are a total of $M = 2d + 1 + t = 2d + 2$ planes in C_∞, we know that the remaining $2d + 1$ curves $\{v_2, ..., v_M\}$ do not meet the divisor at infinity, and we already showed that these planes are in homology class $(-1, 0, 0)$. Equation (4.4.1) shows that each of $\{v_2, ..., v_M\}$ has index 2. This proves the second bullet point of the theorem.

□

4.5 Analysis of the Special Curve

Looking back at Theorem 4.3.3, we see that the negative ends of all curves in C_∞ must match with the positive ends of the so-called special curve, v_0. These matching conditions allow us to use Theorem 4.4.1 to completely characterize the Reeb orbit asymptotics of the special curve. In this section, we will complete the proof of the main theorem.

Corollary 4.5.1. The special curve v_0 must be positively asymptotic to a Reeb orbit in Σ in homology class $(2d + 1, 0, 0)$ and positively asymptotic to an additional $(2d + 1)$-many Reeb orbits, counting multiplicity. The special curve must be negatively asymptotic to a $(2d + 1)$-times cover of the short orbit γ_1 of E.
The area of this special curve is

$$\frac{\epsilon}{2}(2d + 1) - (2d + 1)(\epsilon - \delta_1),$$

Finally, the virtual Fredholm index is zero.

Proof. The claim about Reeb orbits is just a restatement of theorem 4.4.1 and the construc-

tion in theorem 4.3.1. The area calculation is Stokes' theorem, and was explained in the proof of theorem 4.4.1.

The curve that was constructed in theorem 4.3.1 has virtual Fredholm index zero. This curve was stretched to give curves $v_0, ..., v_{2d+2}$. Our special curve is v_0, and it can be matched along its positive end to an additional $2d + 2$ curves of index 2 whose ends are negatively asymptotic to Reeb orbits in 2-dimensional families. This matching is explained in theorem 4.3.3. The matching procedure imposes constraints on the Fredholm indices. The sum of the indices of all the constituent curves $v_0, ..., v_{2d+2}$ must equal the sum of the parent curve plus the sum of the dimensions of orbits at the sites of matching. This implies

$$\underbrace{2 \cdot (2d+2)}_{\text{dim. of matching}} + \underbrace{0}_{Index(\text{parent})} = Index(v_0) + \underbrace{2 \cdot (2d+2)}_{Index(v_1)+...+Index(v_M)}.$$

Consequently $Index(v_0) = 0$. □

Finally, we are in a position to prove the main theorem of this chapter.

Theorem 4.5.1. (Main theorem, remaining case) Suppose that $x \geq 2$, $n \geq 3$, and $a < 2$. If there exists a symplectic embedding

$$P(1, x) \times \mathbb{R}^{2n-4} \hookrightarrow P(a, b) \times \mathbb{R}^{2n-4}$$

then $b \geq x$.

Proof. If such an embedding exists, we may assume, without a loss of generality, that $x > 2$. Assume, towards a contradiction, that $b < x$. Then the results of lemma 4.4.2 and theorem 4.4.1 and corollaries 4.4.1 and 4.5.1 all hold.

We showed in equation (4.4.6) that

$$\sum_{k_i<0} |k_i| = 2d + 1. \tag{4.5.1}$$

As a result of matching imposed by theorem 4.4.1, our limiting building contains $2d + 1$ planes asymptotic to Reeb orbits on Σ in homology class $(-1, 0, 0)$ and having area $1 + O(\epsilon)$. The total area of this limit is at most the area of the plane before stretching, which was computed in section 4.3. This, together with (4.5.1) gives

$$da + b - (2d + 1)(\epsilon - \delta_1) \geq (2d + 1)(1 + O(\epsilon)).$$

Dividing by $2d+1$ and using $a < 2$ gives

$$\frac{b}{2d+1} + \frac{2d}{2d+1} - (\epsilon - \delta_1) > \frac{da}{2d+1} + \frac{b}{2d+1} - (\epsilon - \delta_1) \geq 1 + O(\epsilon).$$

This will give a contradiction when ϵ is sufficiently small, because we initially assumed $b << 2d+1$.

□

4.6 The Proof of Theorem 4.3.1

Now that we have the machinery of neck-stretching, we can prove theorem 4.3.1. This proof will not rely on any of the proofs of prior sections, just the definitions. We restate the theorem here with slightly updated terminology.

Definition 4.6.1. Let us say that $2d+1$ distinct points of a symplectic manifold are *generic relative to* J if there's no closed J-holomorphic curve of index less than $4k+2$ passing through $2k+1$ of the points, for all $0 \leq k \leq d$. Here the notion of curve includes finite energy curves in the completion of the ellipsoid, described below.

Definition 4.6.2. Fix a set of points $p_1, ..., p_{2d+1}$ which are generic with respect to the almost complex structure J_0. A deformation $t \mapsto J_t$, $t \in [0,1]$, of the almost-complex structure is a *generic deformation* if for all $t \in [0,1]$ a curve of index $2k$ passes through $k+1$ of the constraint points, $0 \leq k \leq d$. If a deformation of J_0 is generic, then we call any J_t in this deformation a generic almost-complex structure.

Theorem. In a completion of $X \setminus E$ there exist J-holomoprhic planes of bidegree $(d,1)$, of virtual Fredholm index zero, and which are asymptotic to $\gamma^{(2d+1)}$. Such curves persist under scaling of the ellipsoid and under generic deformations of the almost-complex structure.

We begin with the completion of $X \setminus E$, and we observe that setting $z_3 = 0$ gives an almost-complex submanifold

$$\overline{Y} := \mathbb{C}\mathrm{P}^1(a) \times \mathbb{C}\mathrm{P}^1(b) \setminus E(\epsilon - \delta_1, (2d+1)(\epsilon - \delta_2)).$$

Actually, we will assume that the almost-complex structure is chosen so that $\overline{Y}$ is a complex submanifold of $X \setminus E$. We abbreviate $\tilde{E} := E(\epsilon - \delta_1, (2d+1)(\epsilon - \delta_2)) \subset \mathbb{C}\mathrm{P}^1(a) \times \mathbb{C}\mathrm{P}^1(b)$. Here is a brief outline of the argument that follows.

- We produce a curve C in $\mathbb{C}\mathrm{P}^1(a) \times \mathbb{C}\mathrm{P}^1(b)$ using classical techniques.

- We stretch the neck along the boundary of $\tilde{E}$ to analyze the asymptotics of the curve C.

- We argue that the curve C persists under inclusion into the higher-dimensional space $X \setminus E$, and under scaling of the ellipsoid.

Because of the persistence argument, it suffices to first produce curves in a completion of $\overline{Y}$ (which will have a cylindrical end). This will be the focus of the fourth proposition below. Recall the definition of bidegree in section 4.2. Notice that the ratio of capacities in $\tilde{E}$, namely

$$\theta := \frac{\epsilon - \delta_1}{(2d+1)(\epsilon - \delta_2)}$$

is irrational, because of the assumptions made in (4.0.2). The irrationality of θ ensures that the Reeb orbits of $\partial\tilde{E}$ are isolated, and there will be a short orbit. This number θ will also be the *monodromy angle* of the short orbit γ_1 on $\partial\tilde{E}$, as defined in section 4.1 (see also Lemma 4.1.1). Further, note that the Reeb orbits on $\partial\tilde{E}$ are elliptic and they exist in 0-dimensional families. Theorem 4.3.1 will be proved in the following five propositions.

Proposition 4.6.1. In $\mathbb{C}\mathrm{P}^1(a)\times\mathbb{C}\mathrm{P}^1(b)$ there exists a closed curve of bidegree $(d,1)$ passing through $2d+1$ generic points.

Proof. We first produce a closed embedded curve for the standard J, which is denoted J_0. Such a J_0-curve can be viewed as a graph of a meromorphic function $\mathbb{C}\mathrm{P}^1(b) \to \mathbb{C}\mathrm{P}^1(a)$ which projects to a curve of degree d in $\mathbb{C}\mathrm{P}^1(a)$ and which is biholomorphic to the sphere when projected onto $\mathbb{C}\mathrm{P}^1(b)$. Let us say that the meromorphic function being graphed is of the form f/g, where both f and g are polynomials of degree d. Such polynomials are each specified by $d+1$ parameters, and we must factor out a common scale when dividing f by g. Hence the quotient f/g has $(d+1)+(d+1)-1 = 2d+1$ parameters. This shows that we can produce a J_0-curve through $2d+1$ given points of $\mathbb{C}\mathrm{P}^1(a) \times \mathbb{C}\mathrm{P}^1(b)$. This curve is non-singular, because it is of degree 1 in the second factor.

We present an alternative point of view, by showing that curves in $\mathbb{C}\mathrm{P}^1(a) \times \mathbb{C}\mathrm{P}^1(b)$ passing through $2d+1$ constraint points are in bijective correspondence with curves in the $(2d+1)$-fold blow-up of $\mathbb{C}\mathrm{P}^1(a) \times \mathbb{C}\mathrm{P}^1(b)$. This bijection will be useful for future computations. Fix a generic, ordered set of points $\{p_1, ..., p_{2d+1}\} \subset \mathbb{C}\mathrm{P}^1(a) \times \mathbb{C}\mathrm{P}^1(b)$, and let $\hat{X}$ denote the $(2d+1)$-fold blow-up of $\mathbb{C}\mathrm{P}^1(a) \times \mathbb{C}\mathrm{P}^1(b)$ at those points. Let $\pi\colon \hat{X} \to \mathbb{C}\mathrm{P}^1(a) \times \mathbb{C}\mathrm{P}^1(b)$ denote the projection map. Denote the exceptional divisors by $E_1, ..., E_{2d+1}$. Let A denote a divisor in $\mathbb{C}\mathrm{P}^1(a) \times \mathbb{C}\mathrm{P}^1(b)$ that passes through all $2d+1$ constraint points. The proper transform gives a unique lift to $\hat{A}$ in $\hat{X}$, and $\hat{A}$ determines

an element of $|\pi^{-1}(A) - E_1 - ... - E_{2d+1}|$, which is the projectivization of $\mathscr{O}_{\hat{X}}(\pi^{-1}(A) - E_1 - ... - E_{2d+1})$. Conversely, given a section on $\hat{X}$ of $\mathscr{O}_{\hat{X}}(\pi^{-1}(A) - E_1 - ... - E_{2d+1})$, push it down via π to obtain a section of the complex line bundle $L(A)$ that is defined away from the $2d+1$ constraint points. We know this complex line bundle exists everywhere by assumption. Use the removable singularity theorem to define this section across the $2d+1$ constraint points. This bijection shows that $H^0(\hat{X}, \mathscr{O}(\pi^*(L(A)))) \cong H^0(X, \mathscr{O}(L(A))) = H^0(\mathscr{O}(d_1)) \otimes H^0(\mathscr{O}(d_2))$, for curves that have bidegree (d_1, d_2). The latter space has dimension $(d_1+1)(d_2+1)$ by the Künneth formula. Based on the algebro-geometric definition of genericity of points, we expect a collection of points to be generic if no curve of bidegree $(0,1)$ or $(1,0)$ passes through more than one of the points, along with higher-degree restrictions. The more general requirement is that no curve of bidegree (d_1, d_2) may pass through more than $(d_1+1)(d_2+1) - 1$ of the points. When we consider the special case of bidegree $(d_1, d_2) = (k, 1)$, we have that $(k+1)(1+1) - 1 = 2k+1$. Furthermore, the index formula (4.3.3) implies that a closed genus zero curve of bidegree $(k,1)$ will have virtual Fredholm index $-2+4(k+1) = 4k+2$. This shows the parallel between the above definition of *generic relative to* J and the usual definition of genericity of points. The constraints in the definition are the best possible, because one would expect curves passing through exactly $2k+1$ points to generate a complex of nullity zero. The bijection described in this paragraph gives a formula of $dim[H^0(X, \mathscr{O}(\mathscr{L}_A)] = (d+1)(1+1) = 2d+2$ for curves of bidegree $(d,1)$. Given $2d+1$ generic points, there is a unique section, up to scale, that vanishes at the generic points, because $(2d+2) - (2d+1) = 1$. This shows that there is a unique curve, up to scale, passing through the $2d+1$ given points. This curve will be generically embedded if it avoids a proper discriminant locus, but we will prove it's embedded using the adjunction formula.

Now we view this newly-constructed curve for the standard J_0 as a map

$$u\colon S^2 \to \mathbb{C}\mathrm{P}^1(a) \times \mathbb{C}\mathrm{P}^1(b)$$

(not as a graph), and we apply the four-dimensional adjunction inequality (which is formula 2.7 in [Hut14]). Let $[C]$ denote the class of the image of u. Then

$$\chi(\text{domain of } u) + [C] \cdot [C] - \langle c_1(T(\mathbb{C}\mathrm{P}^1(a) \times \mathbb{C}\mathrm{P}^1(b))), [C] \rangle \geq 0,$$

with equality if and only if the curve u is embedded. In our case, this reduces to

$$2 + \langle (d,1),(d,1)\rangle - \langle (2,2),(d,1)\rangle =$$
$$2 + [d^2(f_1^2) + 2d(f_1 \cdot f_2) + 1(f_2^2)] - [2d(f_1^2) + (2d+2)(f_1 \cdot f_2) + 2(f_2^2)] = 0,$$

where in this calculation f_i^2 denotes the square of the fiber in the i^{th} variable. Being an embedded sphere, such a curve is automatically regular (using Lemma 3.3.3 of [MS12a]). □

The last formula in the above proof is purely topological, and we shall explain in Proposition 4.6.3 how the result will persist under generic deformations of the almost-complex structure. This will show that curves for any J are embedded.

Specifically, for any J, let us define a moduli space

$$\mathcal{M}(J, p_1, ..., p_{2d+1}) := \left\{ (f, (y_1, ..., y_{2d+1})) \,\middle|\, \overline{\partial}_J f = 0, f(y_i) = p_i, \text{ bidegree } (d,1) \right\} / \sim,$$

of closed, constrained J-holomorphic curves identified under $\sim$ by reparameterization of the domain. In the standard case, Proposition 4.6.1 shows that $\mathcal{M}(J_0, p_1, ..., p_{2d+1})$ is nonempty, and representative curves are embedded and automatically regular. To extend this claim to non-standard J, we need to argue that the adjunction formula applies.

We vary the almost complex structure in time, to get a one parameter family J_t, $t \in [0,1]$, with J_0 denoting the standard complex structure. We need to show that the set of points which were chosen to be generic for J_0 remain generic for J_t, $t \in [0,1]$. Since the constraint points remain fixed for all time, we may and shall choose these $2d+1$ points to be in the interior of $\tilde{E}$.

Proposition 4.6.2. There is a second category set of maps $t \mapsto J_t$ such that the points $\{p_1, ..., p_{2d+1}\}$ are generic relative to J_t for all t.

For a proof of this proposition, see lemmas 2.4 through 2.6 and corollary 2.7 of [HK14], replacing $\mathbb{C}\mathrm{P}^2$ with $\mathbb{C}\mathrm{P}^1 \times \mathbb{C}\mathrm{P}^1$.

In other words, it is possible to choose a deformation J_t of the standard structure J_0 that is *generic*, as defined above. We assume such a deformation is given for the remainder of this section. Now we may update the definition of $\mathcal{J}^\star$. In addition to the list of properties at the end of section 4.2, we need $J_t \in \mathcal{J}^\star$ to be such that $\overline{Y}$ is a complex submanifold of $X \setminus E$, and compatible with the stretching along $\partial\tilde{E}$, and generic.

Recall that a holomorphic curve which is not multiply-covered is called *simple.*

Proposition 4.6.3. For any $t \in [0,1]$, the moduli space

$$\mathcal{M}_t := \mathcal{M}(J_t, p_1, ..., p_{2d+1}) = \{(f, (y_1, ..., y_{2d+1})) \\ \left| \overline{\partial}_{J_t} f = 0, f(y_i) = p_i, f \text{ has bidegree } (d,1) \right\} / \sim,$$

of closed J_t-holomorphic curves passing through the generic points $\{p_1, ..., p_{2d+1}\}$ is compact. Curves in this moduli space are identified under reparameterizations of the domain. $\mathcal{M}_t$ has virtual dimension zero. Furthermore, for any $t \in [0,1]$, a representative of $\mathcal{M}_t$ is not nodal, is embedded, is simple, and has ECH index zero.

Proof. The proof of the first claim in this proposition is a standard Gromov compactness argument. Here we are deforming a closed, embedded curve, and a limit of such curves will either be embedded or will bubble. Any bubbling would be of codimension 2, which is precluded for generic deformations of the almost-complex structure, because some component of the bubble tree would have index -2.

This compactness implies that we can apply the adjunction inequality (above) to J_t-holomorphic curves. The conclusion is that for any t, the J_t-curve that is constrained to pass though a set of $(2d+1)$ J_t-generic points will be embedded, and not nodal.

We prove the claim of simple-ness by contradiction. Suppose that a curve

$$C \in \mathcal{M}(J_t, p_1, ..., p_{2d+1})$$

of bidegree $(d,1)$ were multiply-covered. Then C must cover some underlying curve K which has bidegree (i,j) for $i < d$ and $j < 1$, since multiple-coverings will multiply the intersections with L_∞. This underlying curve will have Fredholm index $2c_1 - 2 = 4(i+j) - 2$. This Fredholm index is strictly less than $2(2d+1)$. Hence, for generic J the curve K cannot pass through the $2d+1$ constraint points (with each intersection having codimension 2). This contradicts our construction of C and the genericity of points. We must conclude that C is simple. (Moreover, it is a general fact that a limiting curve is either multiply-covered or is embedded with possibly finitely many, denoted δ-many, point singularities. We shall make use of this fact later.)

We showed in proposition 4.6.2 that the constraint points $p_1, ..., p_{2d+1}$ are generic relative to J_t. The virtual dimension of the moduli space $\mathcal{M}_t$ is computed using the virtual Fredholm index of a representative curve. We are claiming that both the virtual Fredholm index and the ECH index of a representative curve equal zero. To compute the virtual index, we use the formula (4.3.3) for a closed curve of genus zero. Without any constraint

points, a curve C is closed, embedded, and has virtual Fredholm index

$$Index(C) = (2-3)(2-0) + 2(2d+2) = 4d+2.$$

The moduli space of curves corresponding to this Fredholm index can only be counted if $2d+1$ generic constraint points are specified. As noted above, each constraint point has codimension 2. Hence the curve C, when constrained to pass through a generic set of $2d+1$ points, will have virtual Fredholm index zero. Finally, we note that a closed, embedded curve of virtual Fredholm index zero must also have ECH index zero. This fact follows from equation (2.10) of [Hut14]. □

The reason for this computation is as follows. We are now going to stretch the neck along ∂E, which will create, as an intermediate step, almost-complex structures J^N on the manifolds stretched to length N (see theorem 4.6.1 below). Initially (in Proposition 4.6.1) we produced curves for the standard, integrable, almost-complex structure J_0. The point of Proposition 4.6.3 is to show that, after deforming to $J_t := J^N$, the curves persist, and similarly when we take a limit of J_t. For this neck stretching we can re-scale the compact interval as $-\infty \cup (-\infty, 0] \cong [0,1]$. Recall that we assumed that the constraint points are in the interior of $\tilde{E}$. Now, we can stretch the neck of $\mathbb{C}\mathrm{P}^1(a) \times \mathbb{C}\mathrm{P}^1(b)$ along $\partial\tilde{E}$, and use the following SFT theorem to analyze the limiting building.

Theorem 4.6.1. (See 10.6 of [BEH+03]) Fix $2d+1$ constraint points in $\tilde{E}$ and fix a J in $\mathcal{J}^*$. For each $N \in \mathbb{N}$, let u_N be a J^N-holomorphic curve passing through the $2d+1$ constraint points. Fix a representative f_N for each u_N. Then there exists a subsequence of the f_N which converges to a holomorphic building $\mathbf{F}$. The domain of $\mathbf{F}$ is a nodal Riemann sphere (S, j) with punctures, and the building can be described as a collection of finite energy holomorphic maps from the collection of punctured spheres $S \setminus \{\text{nodes}\}$ into one of the following three completions:

$$A^E_\infty := \tilde{E} \cup_{\partial\tilde{E}} (\partial\tilde{E} \times [0,\infty)) \qquad \text{with form } d(e^t\lambda),\ t \in [0,\infty),$$

$$B^E_\infty := \partial\tilde{E} \times \mathbb{R}, \qquad \text{with form } d(e^t\lambda),\ t \in \mathbb{R},$$

or

$$C^E_\infty := \overline{Y} \cup_{\partial\tilde{E}} (\partial\tilde{E} \times (-\infty, 0]) \qquad \text{with form } d(e^t\lambda),\ t \in (-\infty, 0].$$

It will be convenient to view the completion of $\overline{Y}$ from theorem 4.6.1 as a symplectic

null-cobordism. For notational brevity, we set

$$Y := C_{\infty}^{E} = \Big((-\infty, 0] \times \partial\tilde{E} \Big) \cup_{\partial\tilde{E}} \overline{Y}.$$

Then Y is a completion of $\overline{Y}$ with $\partial Y = \partial\tilde{E} \sqcup \emptyset$. In addition, this notation matches the notation of section 4.1 of [Hut09].

As in the discussion following Theorem 4.3.2, curves in Y can have only negative asymptotic ends. For this stretching, we will not focus much on curves in the ellipsoid or in the symplectization, B_{∞}^{E}, other than to note that we have some limiting curves inside and outside the neck. All curves, inside or outside, must have matching asymptotic limits along $\partial\tilde{E}$. A limit of embedded curves must either consist of curves that are embedded or curves that are multiply covered. The problem here is to show that stretching the parent curve C of bidegree $(d, 1)$ does *not* create limiting planes of strictly smaller bidegree. Let $\tilde{C}$ denote any connected component in Y. We now argue the following.

Proposition 4.6.4. The holomorphic building that results from stretching the neck along $\partial\tilde{E}$ consists of simple, embedded components with ECH index and virtual Fredholm index zero. Each component in the highest level, Y, of the building is a holomorphic plane, $\tilde{C}$, that must be negatively asymptotic to the short Reeb orbit on $\tilde{E}$, not the long Reeb orbit.

Proof. We showed in Proposition 4.6.3 that any limiting curves must be simple and not nodal. As mentioned above, a limiting curve that is not multiply covered must be embedded except at finitely many (here $\delta = 0$ many) nodes. We conclude that such a simple curve is embedded.

Next we compute the ECH index and the virtual Fredholm index of $\tilde{C}$. In Proposition 4.6.3, we showed that the parent curve, before stretching, had ECH index zero. This implies that the sum of the ECH indices of the limiting curves, after stretching, must also be zero. Since the ECH index is a non-negative quantity, the ECH index of each limiting curve must be zero. Hence we can write $I_{ECH}(\tilde{C}) = 0$. Now we apply the ECH index inequality from [Hut09]. In the case when $\delta = 0$, this inequality says

$$Index(\tilde{C}) \leq I_{ECH}(\tilde{C}). \tag{4.6.1}$$

Both of these indices are non-negative, because $\tilde{C}$ is simple. Since we have proved that the ECH index is zero, we can conclude that the virtual Fredholm index is zero. Furthermore, this ECH index inequality is actually an equality, which will be helpful for the last part of this proof.

Next, we prove that $\tilde{C}$ must have bidegree $(d, 1)$ and must be negatively asymptotic

to the short Reeb orbit on $\tilde{E}$, not the long Reeb orbit. Because of the summation over negative ends in the virtual index formula (4.3.2), it suffices to consider the case where $\tilde{C}$ has two negative ends. The first is asymptotic to the short Reeb orbit γ_1 and winds r_1 times around this orbit. The second is asymptotic to the long Reeb orbit, which we shall call γ_2, and it winds r_2 times around this long orbit. The Conley-Zehnder indices of these two (multiply-covered) Reeb orbits are equal to

$$CZ(\gamma_1^{(r_1)}) = 2r_1 + 2\lfloor r_1 \cdot \theta \rfloor + 1 \approx 2r_1 + 0 + 1,$$

and

$$CZ(\gamma_2^{(r_2)}) = 2r_2 + 2\left\lfloor \frac{r_2(2d+1)(\epsilon - \delta_2)}{\epsilon - \delta_1} \right\rfloor + 1 \approx 2r_2 + (4d+2)r_2 + 1,$$

as d becomes large compared to ϵ. The index formula (4.3.2) in dimension $n = 4$ with $s^- = 2$ negative ends and with bidegree (d_1, d_2) simplifies to

$$\begin{aligned} Index(\tilde{C}) &= (-1)\cdot(0) + 4d_1 + 4d_2 - CZ(\gamma_1^{(r_1)}) - CZ(\gamma_2^{(r_2)}) + \frac{1}{2}(0+0) \\ &= 4d_1 + 4d_2 - 2 - 2r_1 - 2r_2 - 4dr_2 - 2r_2 \\ &= 4(d_1+d_2) - 2 - 2(r_1+r_2) - 2(2d+1)r_2 \\ &= 4(d_1+d_2) - 2 - 2(2d+1)(r_2+1) \end{aligned}$$

where in the penultimate step we used the fact that $r_1 + r_2 = 2d+1$, by construction. Since $r_1, r_2 \geq 0$, in order to have $Index(\tilde{C}) = 0$ we must have $d_1 + d_2 = d+1$ and $r_2 = 0$. The maximum of each of these degree terms is given by the bidegree of the parent curve, $(d, 1)$. Hence for $\tilde{C}$ we must have $d_1 = d$ and $d_2 = 1$ and its negative end(s) must wind along only the short orbit γ_1.

Finally, it remains to show that $\tilde{C}$ is a holomorphic plane (topologically a disc). We show this by proving that $\tilde{C}$ has only a single negative end asymptotic to a $(2d+1)$-cover of γ_1. It turns out that this is determined indirectly in the machinery of ECH, using "partition conditions," which we shall briefly motivate here. Let $\tilde{C}$ denote any component inside of Y. Again, Y is a symplectic null-cobordism ($\partial Y = \partial\tilde{E} \sqcup \emptyset$), which implies that $\tilde{C}$ has no positive end, and we assume *a priori* that $\tilde{C}$ has a number of negative ends which are all asymptotic to γ_1 on $\partial\tilde{E}$ with *total* multiplicity $2d+1$. It is conceivable for $\tilde{C}$ to be topologically a U-tube with one end wrapping once around γ_1 and the other end wrapping $2d$ times around γ_1. A simple calculation shows that this this configuration is impossible; $\tilde{C}$ must be a 2-handle. See figure 4.3 for an illustration of this non-example.

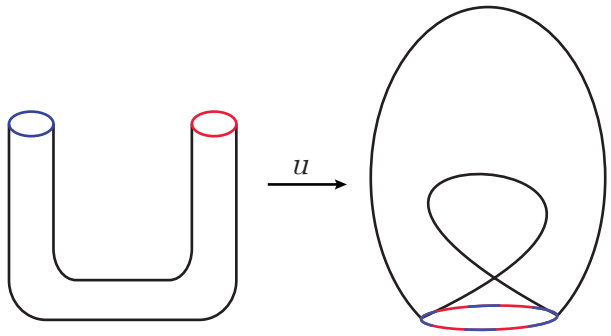

Figure 4.3: A J-holomorphic curve u that is prohibited by the partition conditions.

The point of the ECH partition conditions is that for an embedded curve in a symplectic cobordism, the multiplicity of its negative end(s) can be computed as a partition $p_j^-(\tilde{C}) = p_{\gamma_j}^-(m_j)$, where the subscript j denotes the j^{th} negative end of $\tilde{C}$, which is asymptotic to a cover of the Reeb orbit γ_j with total covering multiplicity m_j. The multiplicities of these covers will give a partition of the positive integer m_j, hence the name. There is a similar partition condition for the positive end, but in our setup $\tilde{C}$ has no positive end. So we focus on the first negative end of $\tilde{C}$, setting $j = 1$. The total covering multiplicity of γ_1 is $2d + 1$ by construction. So we need to compute the partition

$$p_1^-(\tilde{C}) = p_{\gamma_1}^-(2d + 1).$$

The full definition of these "incoming partition condtions" for a symplectic cobordism is given in the paper [Hut09] with the slightly older notation $p^{in} = p^-$. Here we use the updated notation from [Hut14]. In our case, we need only the fact that $p_{\gamma_1}^-(2d + 1)$ is entirely defined by the monodromy angle θ. Since $\lfloor (2d + 1)\theta \rfloor = 0$, we have that $p_{\gamma_1}^-(2d+1) = (2d+1)$ is the trivial partition. The relevant theorem from [Hut09] says that when the ECH index inequality (4.6.1) is an equality, the partition $p_{\gamma_1}^-$ exactly determines the multiplicity of covering(s) of the negative end(s) of the J-curve. This equality was proved at the beginning of this proof. The computation that $p_{\gamma_1}^-$ is the trivial partition shows that the curve $\tilde{C}$ can have only a single negative end asymptotic to a $(2d + 1)$-cover of γ_1. Hence, in particular, $\tilde{C}$ is a holomorphic plane, and there is only a single limiting curve in the top levels of the holomorphic building that results from stretching the neck. Since we started with a single curve C (before stretching) and we ended with a single curve $\tilde{C}$ (after stretching), we may as well identify $C = \tilde{C}$ within Y, to simplify the notation.

□

This completes the first two bullet points in the proof outline above. We now need to explain why such a curve C in Y can be included into a completion of $X \setminus E$. First, for comparison purposes, we recall the definition of regularity. In dimension four, an almost-complex structure is regular if the linearized normal Cauchy-Riemann deformation operator is surjective at all somewhere injective curves. In particular, there are no somewhere injective curves of negative Fredholm index, and in dimension four spheres of non-negative Fredholm index automatically have a surjective deformation operator. Above we used the fact that there is a second category generic subset of regular almost complex structures. This discussion leads us to, once again, update the conditions on $\mathcal{J}^\star$. Here we need almost-complex structures on the six-dimensional manifold $X \setminus E$ such that the inclusion $\overline{Y} \to X \setminus E$ (equivalently, their respective completions) is a holomorphic embedding with respect to a compatible almost-complex structure on Y for which a curve exists, satisfying the above propositions. Let $\mathcal{J}^\star$ denote the set of almost-complex structures on $X \setminus E$ with this property, in addition to the properties on Y that follow Proposition 4.6.2. We shall have more to say about genericity of the almost-complex structures in a moment, but for now we emphasize that (after the neck stretching along $\partial\tilde{E}$) a generic almost complex structure $J \in \mathcal{J}^*$ admits no J-holomorphic planes of index -2 in all possible bidegrees and with all possible windings along the negative end.

Proposition 4.6.5. If such J-planes exist in a completion Y of $\overline{Y}$, then they persist under inclusion in the higher-dimensional space $X \setminus E$ (suitably completed). Such a stabilized curve will persist under deformations of the complex structure and scaling of the ellipsoid E.

Proof. First, consider a completion of $X \setminus E$ using the inclusion embedding of the ellipsoid. We have already constructed in Proposition 4.6.4 a genus zero J-holomorphic curve C in a completion of $\overline{Y}$, for some J. We need to argue that the stabilized curve $C \times \{pt\}$ includes into (the completion of) $X \setminus E$ without changing the Fredholm index. We refer back to the general index formula for a genus zero curve, found on page 73. The following is a standard argument. Stabilizing the curve C obviously does not change the genus, but it does increase n by 1 and increase each CZ term by 1 for each factor of $\mathbb{R}^2$ in the stabilization. A useful observation of [HK14] is that we can take genus zero curves having $s^+ = 0$ positive ends, so that the index of the stabilized curve reduces to

$$Index(C \times \{pt\}) = 2 - 2(s^-) + Index(C).$$

In particular, the index is unchanged under stabilizations when $s^- = 1$, which is the case of interest here. This implies that the curve from Proposition 4.6.4, before and after stabi-

lization, will have virtual Fredholm index zero. We temporarily call such a stabilized curve a J_0-holomorphic curve, where $J_0 \in \mathcal{J}^*$ is a generic almost-complex structure that results from the neck stretching in proposition 4.6.4. We now have a moduli space

$$\mathcal{M}_0 = \{J_0\text{-holomorphic curves in } X \backslash E \text{ having one negative end asymptotic to } \gamma_1^{2d+1}\}/\sim$$

which has virtual dimension zero, and is nonempty. In this moduli space, curves are identified up to reparameterization of the domain.

After varying J in time, we now have a moduli space

$$\mathcal{M}_t = \{J_t\text{-holo. curves in } X \setminus E \text{ having one negative end asymptotic to } \gamma_1^{2d+1}\}/\sim .$$

The curves in this moduli space are identified up to reparameterization of the domain, and when $t = 0$, we recover the above moduli space $\mathcal{M}_0$. For the remainder of this proof, a generic deformation of the almost-complex structure will be a deformation $t \mapsto J_t \in \mathcal{J}^\star$, $t \in [0,1]$ such that J_t is generic for all times $t \in [0,1]$. We shall explain why a J_0-holomorphic curve curve persists under a generic deformation of the almost-complex structure. We note that for a second-category subset of 1-parameter families of almost-complex structures, the deformation operator for all moduli spaces $\mathcal{M}_t$ is also surjective. This is the analogue of proposition 4.6.2. Moreover, the Fredholm index of a representative curve in $\mathcal{M}_t$ must be even, which precludes curves of negative index even in 1-parameter families.

The above computation shows that

$$\mathcal{M}_{[0,1]} = \{\,(u,t)|\, u \in \mathcal{M}_t,\, t \in [0,1]\}.$$

is a one-dimensional manifold. Hence $\partial\mathcal{M}_{[0,1]}$ consists of a discrete set of points, which we call

$$\partial\mathcal{M}_{[0,1]} \cong \mathcal{M}_0 \sqcup \mathcal{M}_1. \tag{4.6.2}$$

We noted above that for a second-category subset of J_t, the deformation operator for $\mathcal{M}_{[0,1]}$ will be surjective, and representative curves will have non-negative Fredholm index. Not every representative curve of Fredholm index zero will be cut out transversally, however, for all time $t \in [0,1]$. In this higher-dimensional context (dimension ≥ 6), the projection map $\mathcal{M}_{[0,1]} \to [0,1]$ may not be a covering map, but it will give a cobordism from $\mathcal{M}_0$ to $\mathcal{M}_1$. The moduli space $\mathcal{M}_0$ was examined in the preceding propositions, and it was shown to be nonempty. Moreover, at time $t = 0$ we required Y to be holomorphically embedded with respect to J_0. Using (4.6.2), we regard $\mathcal{M}_{[0,1]}$ as a cobordism between $\mathcal{M}_0$

and $\mathcal{M}_1$.

We explain why $M_{[0,1]}$ is sequentially compact, which will imply that each $\mathcal{M}_t$ is compact. To prove sequential compactness, consider a sequence $(u_k, t_k) \in \mathcal{M}_{[0,1]}$ and, after passing to a subsequence, assume that the t_k converge to $t_\infty \in [0,1]$. Then, after possibly taking another subsequence of (u_k, t_k), SFT compactness implies that the u_k converge to a J_{t_∞}-holomorphic building. Note that, by construction, this building contains levels in the completion of $X \setminus E$ and levels in the symplectization layer $\partial E \times \mathbb{R}$. We must show that this building is in fact a J_{t_∞}-holomorphic plane. We do this by showing that any symplectization components must be trivial cylinders. Start with a curve G in the lowest symplectization level. The results of proposition 4.6.4 imply that G must have positive and negative ends asymptotic to γ_1 only. By construction, we also have that the negative end of G must be $\gamma_1^{(2d+1)}$. Let us say that G has s^+ positive ends, with the p^{th} positive end winding a_p times around γ_1, and let us say that G has $s^- = 1$ negative end winding $2d+1$ times around γ_1. The number of positive ends of G is at least the number of negative ends, both counted with multiplicity, which implies that

$$\sum_{p=1}^{s^+} a_p \geq 2d+1.$$

Moreover

$$\left\lfloor \frac{a_p(\epsilon - \delta_1)}{(2d+1)(\epsilon - \delta_j)} \right\rfloor = 0,$$

whenever $2 \leq j \leq n$ and whenever $a_p \leq 2d+1$ because of our choice of $\epsilon, \delta_1, \delta_2, \delta_3$, etc. We will omit these terms from the upcoming Fredholm index formula, knowing that they contribute positively or not at all to the index. The dimension of the elliptical families of Reeb orbits is zero, and in the symplectization layer the first Chern class is zero. Using the trivialization coming from the inclusion of the ellipsoid, we may use the Conley-Zehnder formula from lemma 4.1.1. The virtual Fredholm index formula gives

$$\begin{aligned} Index(G) &= (n-3)(2-s^+-1) + 2c_1(u) + \left(\sum_{p=1}^{s^+} CZ(\gamma_1^{(a_p)})\right) - CZ(\gamma_1^{(2d+1)}) \\ &= (n-3)(1-s^+) + 0 + \left(\sum_{p=1}^{s^+} (2a_p + (n-1))\right) - (2(2d+1) + (n-1)) \\ &= 2(s^+ - 1) + 2\left(\sum_{p=1}^{s^+} a_p - (2d+1)\right) \geq 0 \end{aligned}$$

with equality if and only if $s^+ = 1$, and $\sum_{p=1}^{s^+} a_p = a_1 = 2d + 1$. This happens if and only if G is a trivial cylinder. By induction on levels, all curves in the symplectization layer have nonnegative Fredholm index, and if a symplectization component were not a trivial cylinder, then it must have strictly positive Fredholm index. This would force a component in $X \setminus E$ to be (possibly a multiple cover of) a curve with negative Fredholm index. This configuration is precluded by genericity of $J \in \mathcal{J}^*$. Hence $M_{[0,1]}$ is compact

We claim that the cobordism $\mathcal{M}_{[0,1]}$ is not a null-cobordism. Then $\mathcal{M}_0$ being nonempty implies $\mathcal{M}_1$ is nonempty. We prove this claim by showing that the cardinality of $\mathcal{M}_0$, counted with orientation, is non-zero. Hence a null-cobordism cannot occur. This computation is done in propositions 10 and 11 of [CGH18], but is simpler here because the curves we consider have only a single negative end. A summary of the argument follows. Note that $X \setminus E$ admits an S^1 action that rotates the stabilized factor ($\mathbb{R}^2 \cong \mathbb{C}$ with coordinate z_3). Let us write $\mathcal{J}^*_{S^1}$ for the set of S^1 invariant and admissible almost-complex structures on $X \setminus E$ for which the moduli space $\mathcal{M}_{[0,1]}$ is nonempty. It is proved (in [HK14] section 3.3.2) that such S^1-invariant almost-complex structures exist, making $\mathcal{J}^*_{S^1}$ nonempty. For generic $J \in \mathcal{J}^*_{S^1}$ all J-curves which meet an S^1 orbit exactly once will be regular. One then has automatic transversality for these curves. Once $\mathcal{J}^*_{S^1}$ is known to be nonempty, the assertion is that for some $J \in \mathcal{J}^*_{S^1}$, the curves in the corresponding moduli space $\mathcal{M}_0$ have image in the 4-dimensional, holomorphic slice Y (i.e. $z_3 = 0$). If not, some curve with image not contained entirely in the slice would meet the S^1 orbit more than once, hence in an S^1-family. Consider the projection of such a curve u onto the slice $z_3 = 0$. The projection of this curve is non-injective and the curve u multiply covers its projection. But such a curve u can be excluded by index calculations.

Finally, representatives of $\mathcal{M}_1$ persist under scaling the ellipsoid E. Note that the set of ellipsoid embeddings into X is connected, and one gets from one embedding of E to any other embedding of E through scaling and change of coordinate bases. Re-scaling the ellipsoid does not change the manifold $X \setminus E$ up to diffeomorphism. Furthermore, we can pull all structures (symplectic, almost-complex, etc.) back from the re-scaled $X \setminus E$ to the un-scaled $X \setminus E$. Re-scaling the ellipsoid has no effect on the Fredholm index of curves asymptotic to the ellipsoid, but it does affect the area of curves asymptotic to the ellipsoid. In the proof of this proposition, we have used only Fredholm arguments, making the proof agnostic of scale. This completes the proof of the proposition and the proof of Theorem 4.3.1.

□

4.7 Higher Dimensions

We conclude this chapter by explaining how the above results extend to higher stabilizations, i.e. $n > 3$. In section 4.1, we should consider the $2n$-dimensional ellipsoid

$$E = (\epsilon - \delta_1, (2d+1)(\epsilon - \delta_2), ..., (2d+1)(\epsilon - \delta_n)),$$

where $\delta_2, ..., \delta_n < \delta_1 < \epsilon$ are small, and formula (4.0.2) holds for $1 \leq i \neq j \leq n$. This irrationality of the capacities of the ellipsoid implies that there are only n closed Reeb orbits of E along the coordinate planes. The eccentricity of the ellipsoid E ensures that curves with positive area can only be asymptotic to the short orbit γ_1. (This is the first computaion of Proposition 4.6.4.) The formula for the Conley-Zehnder index of a Reeb orbit on a $2n$-dimensional ellipsoid is given in lemma 4.1.1.

In section 4.2, the source polydisc $P(1, x) \times \mathbb{R}^{2n-4}$ should be replaced by

$$P(1, x, \underbrace{S, ..., S}_{n-2}).$$

This polydisc contains a special subset U defined as a toric domain by

$$U = \mu^{-1}(\,(1-\epsilon, 1) \times (x - (2d+1)\epsilon, x) \times (S/2, S)^{n-2}\,),$$

and the discussion of how to use a Hamiltonian to smooth U to obtain Σ generalizes in the obvious way. The formula for the Conley-Zehnder index of a Reeb orbit on Σ in homology class $(k, \ell, m_1, ..., m_{n-2}) \in \mathbb{Z}^n \setminus \vec{0}$ generalizes to

$$CZ(c) = 2k + 2\ell + 2m_1 + ... + 2m_{n-2} + \frac{n-1}{2}. \tag{4.7.1}$$

We compactify the target polydisc $P(a, b) \times \mathbb{R}^{2n-4}$ as $X = \mathbb{C}\mathbf{P}^1(a) \times \mathbb{C}\mathbf{P}^1(b) \times \mathbb{C}^{n-2}$. The divisor at infinity becomes

$$L_\infty = p_\infty \times \mathbb{C}\mathbf{P}^1(b) \times \mathbb{C}^{n-2} \cup \mathbb{C}\mathbf{P}^1(a) \times p_\infty \times \mathbb{C}^{n-2},$$

and we assume L_∞ is a J-holomorphic submanifold of X. The bidegree of a curve in X is defined by its intersection number with the Poincaré dual of L_∞. We then assume that

$$E \to U \hookrightarrow P(a, b) \times \mathbb{R}^{2n-4} \hookrightarrow X.$$

The existence theorem for curves, Theorem 4.3.1, does not depend on n. The proof of

this theorem is the whole of section 4.6, and the Propositions 4.6.1 through 4.6.4 take place in four dimensions. It is not until Proposition 4.6.5 that stabilization enters. The curve is stabilized by simply including it in the stabilized target manifold. The index computation in Proposition 4.6.5 explicitly mentions n, and still gives a non-negative index when $n > 3$. The conditions for $Index(G) = 0$ are the same for $n > 3$ in Proposition 4.6.5. The conclusion of the computation would be the same. In the penultimate paragraph of Proposition 5, the S^1 action is replaced with a $\mathbb{T}^n$ action, which rotates each stabilized factor. The proof in the reference [CGH18] goes through just the same.

Once the existence of a stabilized curve is ensured, the proof by contradiction of the main theorem can be done in any dimension. The formula (4.7.1) will simplify, because we can show that all $m_j = 0$ as we did in lemma 4.3.4. Notice, in particular, that the Conley-Zehnder index (4.7.1) increases by $1/2$ as n increases by 1. This contribution from the dimension is exactly canceled by the term that involves the dimension of the Reeb orbits in the virtual index formula, but only for a negatively asymptotic end. Specifically, we have shown for a negative end winding about a hyperbolic orbit c on Σ that

$$CZ(c) - \frac{1}{2}\dim(c) = \left(2k + 2\ell + 2m_1 + ... + 2m_{n-2} + \frac{n-1}{2}\right) - \frac{1}{2}(n-1) = 2(k+\ell),$$

in all dimensions $n \geq 3$. Consequently, formula (4.4.2) and the formulas in lemma 4.3.2 still describe the virtual Fredholm index of a plane in $X \setminus U$ with negative end(s) asymptotic to a Reeb orbit on Σ. For this reason, the results of lemma 4.3.2 still hold for $n \geq 3$. (A similar cancellation can be used to show that the so-called "special curve" has virtual Fredholm index zero, but this fact is never used.) This computation essentially reduces the remainder of the proof to the 4-dimensional case.

CHAPTER 5

Conclusion

In this final chapter, the state of the field of symplectic embeddings is considered, and one avenue for future work will be discussed.

Recall the properties of ECH capacities that were described in chapter 1. ECH capacities are one example of the more general notion of a *symplectic capacity*, which is a function

$$\mathfrak{g}\colon \{\text{symp. manifolds with contact boundary}\}/\text{symplectomorphism} \to \mathbb{R}_{\geq 0}$$

satisfying the following axioms

1. (Normalization, optional) The symplectic capacity, when applied to the unit ball, gives
$$\mathfrak{g}(B^{2n}(1)) = 1.$$

2. (Monotonicity) If (X_1, ω_1) symplectically embeds into (X_2, ω_2) then
$$\mathfrak{g}(X_1) \leq \mathfrak{g}(X_2).$$
If this property holds only for exact symplectic embeddings, the capacity $\mathfrak{g}$ is called *exact*.

3. (Conformality) Let X_2 be defined by rescaling the symplectic manifold X_1 as $(X_2, \omega_2) = (X_1, r\omega_1)$ for $r > 0$. Then
$$\mathfrak{g}(X_2) = |r|\mathfrak{g}(X_1).$$

4. (Stabilization, optional) $\mathfrak{g}(X \times \mathbb{C}^N) = \mathfrak{g}(X)$ for $N \geq 1$.

Given a symplectic capacity, one often uses the contrapositive of the monotonicity property to obstruct symplectic embeddings. The stabilization axiom is labeled as optional because not all symplectic capacities will have this property. The stabilization axiom would

obviously be quite useful for exploring stabilized embedding problems, as defined in chapter 4. One example of a symplectic capacity that enjoys the stabilization property is the sequence of Ekeland-Hofer capacities, which we denote by $\mathfrak{g}_k^{EH}$. The Ekeland-Hofer capacity is defined in [EH90]. For Liouville domains, the k^{th} Ekeland-Hofer capacity of X is represented by the action of some closed Reeb orbit on ∂X. We give one example computation here.

Example 5.0.1. The k^{th} Ekeland-Hofer capacity of the ellipsoid $E(a_1, ..., a_n)$ is the k^{th} smallest element of the infinite array $\{ia_j \,|i, j \geq 1\}$, with repetitions allowed. For a stabilized ellipsoid, we allow some $a_j = \infty$.

One can compare this to the computation of the k^{th} ECH capacity of an ellipsoid $E(a, b)$ in example 1.0.2. The formula for the k^{th} ECH capacity involves both areas (or radii) of the ellipsoid, whereas the k^{th} Ekeland-Hofer capacity sees one area at a time.

It was previously conjectured that the Ekeland-Hofer capacities are *sharp* (also called *complete invariants*) for the 4-dimensional ellipsoid embedding problem. This would mean that an embedding of $E(a, b)$ into $E(c, d)$ is obstructed if and only if it is precluded by Ekeland-Hofer capacities or volume. (In other words, the Ekeland-Hofer capacities and the volume capacity would generate the collection of generalized symplectic capacities on the space of open 4-dimensional symplectic ellipsoids.) It was shown in [McD09] that Ekeland-Hofer capacities are not sharp. This implies that in 4 dimensions, Ekeland-Hofer capacities are less informative than ECH capacities, which are known to be sharp (for ellipsoids embedding into ellipsoids) but are unique to dimension 4. When the Ekeland-Hofer capacities are applied to a stabilized 4-manifold, they, by construction, give the same information as for the 4-dimensional non-stabilized manifold. For this reason, we can't expect the Ekeland-Hofer capacities to be a supple invariant on stabilized ellipsoids.

The standard method for creating a symplectic capacity was described in chapter 1. One supposes that a symplectic embedding of $\varphi\colon X_1 \to X_2$ exists, and one looks for J-holomorphic curves in the symplectic completion of the cobordism $X_2 \setminus \varphi(X_1)$. The existence of such a curve can be secured from neck stretching (as was done in proposition 4.6.4), or from some homology theory. ECH capacities, for example, use embedded contact homology to count curves. Once a curve is produced, positivity of area can obstruct the supposed embedding. In order to enumerate such curves, it is useful to construct a (nonempty) moduli space of dimension zero. If one wishes to create a symplectic capacity that enjoys the stabilization property, one could study J-holomorphic curves in symplectic cobordisms whose Fredholm indices are unchanged under the process of stabilization (wherein the J-holomorphic curve is crossed with a point). The class of curves that were

studied in proposition 4.6.5 have Fredholm index unchanged under stabilization. Namely, these curves have genus zero, no positive ends, and a single negative end. Hence, this would need to be the class of curves constructed in the non-stabilized 4-dimensional symplectic cobordism if one wishes to stabilize the curves and the cobordism. Currently, this is the only known method of constructing curves which behave well under stabilization. To build a capacity from this observation, one needs to construct a homology theory from such curves.

Recently, a symplectic capacity was defined in [Sie19b] that enjoys the stabilization property and that looks to be a very informative invariant when applied to stabilized ellipsoids. As mentioned above, the embedding obstructions come from curves which will stabilize, and these curves are counted using linearized contact homology, CH_{lin}. The resulting capacity, denoted $\mathfrak{g}_{\mathfrak{b}}$ relies on the techniques of rational symplectic field theory (RSFT). The reader should note that there is not currently a consensus on which construction of RSFT gives an adequate framework for virtual perturbation of J-holomorphic curves. The paper [Sie19b] describes how the capacities $\mathfrak{g}_{\mathfrak{b}}$ could be reformulated to rely on Floer theory, instead of RSFT. The details of this reformulation have not been fully established. Nonetheless, the initial computations of $\mathfrak{g}_{\mathfrak{b}}$ have lent credence to the following useful property.

Conjecture. The capacities $\mathfrak{g}_{\mathfrak{b}}$ are sharp (i.e. a complete invariant) for the problem of embedding a stabilized ellipsoid into a stabilized ellipsoid

$$E(a,b) \times \mathbb{C}^N \hookrightarrow E(c,d) \times \mathbb{C}^N .$$

Currently, the capacities $\mathfrak{g}_{\mathfrak{b}}$ are very difficult to compute, with the current procedure being recursive. The field should consider computing $\mathfrak{g}_{\mathfrak{b}}$ for other toric domains, and simplifying the methods that currently exist for computation. This would be extremely helpful for examining the stabilized embedding problem for other toric domains. The capacities $\mathfrak{g}_{\mathfrak{b}}$ are not sharp for the problem of embedding stabilized polydiscs. In [Sie19b], a computation of the capacity $\mathfrak{g}_{\mathfrak{b}}$ predicts that if

$$P(1,2) \times \mathbb{C} \hookrightarrow P(c,c) \times \mathbb{C},$$

then $c \geq 3/2$. An application of theorem 4.0.1 from chapter 4 gives the improved bound of $c \geq 2$, and this bound is the best possible.

One such stabilized embedding problem, which we shall call the stabilized Fibonacci staircase, has been studied by Cristofaro-Gardiner, Hind, and McDuff. Recall that a portion

of the graph of the function

$$f_{EB}(x) = \inf\{R \,|\, E(1,x) \text{ symplectically embeds into } B^4(R)\}, \quad \text{for } x \geq 1$$

is shown in figure 1.1. Let us define a stabilized version of this function by

$$f_{EB}^N(x) := \inf\{R \,|\, E(1,x) \times \mathbb{C}^N \text{ symplectically embeds into } B^4(R) \times \mathbb{C}^N\}, \quad \text{for } x \geq 1.$$

In three papers, [CGH18], [McD18], [CGHM18], Cristofaro-Gardiner, Hind, and McDuff computed the value of f_{EB}^N for certain $x \geq 1$. In particular, two of those authors proved the folliowing.

Theorem. ([CGH18]) The functions f_{EB} and f_{EB}^N coincide for $1 \leq x \leq \left(\frac{1+\sqrt{5}}{2}\right)^4$.

Because the graphs of f_{EB} and f_{EB}^N coincide on the given interval, the stabilized embedding function f_{EB}^N is still bounded below by a volume curve, $y = \sqrt{x}$. This is interesting, because stabilized ellipsoids have infinite volume, and therefore an embedding of stabilized ellipsoids should not be obstructed by volume. It is currently not understood why the volume obstruction from the 4-dimensional ellipsoid embedding problem persists to the stabilized embedding problem.

Some values of the function f_{EB}^N can be computed using Siegel's capacities $\mathfrak{g}_\mathfrak{b}$, and this is done in [Sie19a].

As a future avenue of research, the field might consider the stabilized Pell staircase. Recall that the Pell staircase is the graph of the function

$$f_{EC}(x) = \inf\{R \,|\, E(1,x) \text{ symplectically embeds into } C^4(R)\}, \quad \text{for } x \geq 1,$$

and the relevant portion of the graph, for $1 \leq x \leq (1+\sqrt{2})^2$, is shown in figure 1.2. Analogously, one can stabilize the manifolds involved in this embedding problem to make the function

$$f_{EC}^N(x) = \inf\{R \,|\, E(1,x) \times \mathbb{C}^N \text{ symplectically embeds into } C(R) \times \mathbb{C}^N\}, \quad \text{for } x \geq 1.$$

Based on the results cited above, we make the following.

Conjecture. The graph of $f_{EC}^N(x)$ coincides with the graph of $f_{EC}(x)$ for $1 \leq x \leq (1+\sqrt{2})^2$.

The procedure for proving this conjecture would be similar to the technique described for proving theorem 4.0.1 above. For certain ratios of Pell numbers, $1 \leq p \leq (1+\sqrt{2})^2$,

we study J-holomorphic curves in the cobordism $C(f_{EC}(p)) \setminus E(1,p)$. If the curves have genus zero and only a single end asymptotic to the ellipsoid, then their Fredholm indices are unchanged under the process of stabilization. This would imply that such curves of Fredholm index zero could be enumerated in the four-dimensional symplectic cobordism and meaningfully included (crossed with a point) in the corresponding stabilized symplectic cobordism. The symplectic capacities $\mathfrak{g}_\mathfrak{b}$, if computable, could assist in counting such curves.

In the context of the conjecture, notice also that the graph of $y = f_{EC}(x)$ on the interval $x \in [1, (1+\sqrt{2})^2]$ must stay below the "folding curve" $y = 2x/(x+1)$. Analogously, f_{EC} must stay above the volume curve, $y = \sqrt{x/2}$. Consequently, volume plays a non-obvious role in the stabilized problem. For $x >> (1+\sqrt{2})^2$, the "folding curve" of $y = 2x/(x+1)$ may coincide with f^N_{EC}. If so, folding would be the optimal embedding technique for very skinny stabilized ellipsoids into stabilized cubes. Finally, it would be interesting to know which obstructions present in the four-dimensional embedding function f_{EC} do not persist under stabilization to f^N_{EC}.

Bibliography

[BEH+03] Frédéric Bourgeois, Yakov Eliashberg, Helmut Hofer, Krzysztof Wysocki, and Eduard Zehnder. Compactness results in symplectic field theory. *Geom. Topol.*, 7:799–888, 2003.

[Bou02] Frédéric Bourgeois. *A Morse-Bott approach to contact homology*. PhD thesis, Stanford University, https://www.math.u-psud.fr/~bourgeois/papers/thesis.pdf, 2002. Accessed on June 01, 2019.

[CCGF+14] Keon Choi, Daniel Cristofaro-Gardiner, David Frenkel, Michael Hutchings, and Vinicius Gripp Barros Ramos. Symplectic embeddings into four-dimensional concave toric domains. *Journal of Topology*, 7(4):1054–1076, 05 2014.

[CG19] Daniel Cristofaro-Gardiner. Symplectic embeddings from concave toric domains into convex ones. *J. Differential Geom.*, 112(2):199–232, 06 2019.

[CGFS17] Daniel Cristofaro-Gardiner, David Frenkel, and Felix Schlenk. Symplectic embeddings of four-dimensional ellipsoids into integral polydiscs. *Algebr. Geom. Topol.*, 17(2):1189–1260, 2017.

[CGH18] Daniel Cristofaro-Gardiner and Richard Hind. Symplectic embeddings of products. *Comment. Math. Helv.*, 93(1):1–32, 2018.

[CGHM18] Daniel Cristofaro-Gardiner, Richard Hind, and Dusa McDuff. The ghost stairs stabilize to sharp symplectic embedding obstructions. *J. Topol.*, 11(2):309–378, 2018.

[CM05] Kai Cieliebak and Klaus Mohnke. Compactness for punctured holomorphic curves. *J. Symplectic Geom.*, 3(4):589–654, 2005. Conference on Symplectic Topology.

[Dar82] Gaston Darboux. Sur le problème de Pfaff. *Bull. Sci. Math*, 6:14–36,49–68, 1882.

[EH90] Ivar Ekeland and Helmut Hofer. Symplectic topology and Hamiltonian dynamics. II. *Math. Z.*, 203(4):553–567, 1990.

[FM15] David Frenkel and Dorothee Müller. Symplectic embeddings of 4-dimensional ellipsoids into cubes. *J. Symplectic Geom.*, 13(4):765–847, 2015.

[GH17] Jean Gutt and Michael Hutchings. Symplectic capacities from positive S^1-equivariant symplectic homology. *Algebraic & Geometric Topology*, 18, 07 2017.

[Gro85] Mikhail Gromov. Pseudo holomorphic curves in symplectic manifolds. *Invent. Math.*, 82(2):307–347, 1985.

[Gut08] Larry Guth. Symplectic embeddings of polydisks. *Invent. Math.*, 172(3):477–489, 2008.

[Gut14] Jean Gutt. Generalized Conley-Zehnder index. *Ann. Fac. Sci. Toulouse Math. (6)*, 23(4):907–932, 2014.

[Hin15] Richard Hind. Some optimal embeddings of symplectic ellipsoids. *J. Topol.*, 8(3):871–883, 2015.

[HK14] Richard Hind and Ely Kerman. New obstructions to symplectic embeddings. *Invent. Math.*, 196(2):383–452, 2014.

[HK18] Richard Hind and Ely Kerman. Correction to: New obstructions to symplectic embeddings [MR3193752]. *Invent. Math.*, 214(2):1023–1029, 2018.

[Hut09] Michael Hutchings. The embedded contact homology index revisited. In *New perspectives and challenges in symplectic field theory*, volume 49 of *CRM Proc. Lecture Notes*, pages 263–297. Amer. Math. Soc., Providence, RI, 2009.

[Hut14] Michael Hutchings. Lecture notes on embedded contact homology. In *Contact and symplectic topology*, volume 26 of *Bolyai Soc. Math. Stud.*, pages 389–484. János Bolyai Math. Soc., Budapest, 2014.

[Hut16] Michael Hutchings. Beyond ECH capacities. *Geom. Topol.*, 20(2):1085–1126, 2016.

[KK17] Yael Karshon and Liat Kessler. Distinguishing symplectic blowups of the complex projective plane. *J. Symplectic Geom.*, 15(4):1089–1128, 2017.

[Lio38] Joseph Liouville. Note sur la théorie de la variation des constantes arbitraires. *Journ. de Math.*, 3:342–349, July 1838.

[McD09] Dusa McDuff. Symplectic embeddings of 4-dimensional ellipsoids. *J. Topol.*, 2(1):1–22, 2009.

[McD18] Dusa McDuff. A remark on the stabilized symplectic embedding problem for ellipsoids. *Eur. J. Math.*, 4(1):356–371, 2018.

[MP94] Dusa McDuff and Leonid Polterovich. Symplectic packings and algebraic geometry. *Invent. Math.*, 115(3):405–434, 1994. With an appendix by Yael Karshon.

[MS12a] Dusa McDuff and Dietmar Salamon. *J-holomorphic curves and symplectic topology*, volume 52 of *American Mathematical Society Colloquium Publications*. American Mathematical Society, Providence, RI, second edition, 2012.

[MS12b] Dusa McDuff and Felix Schlenk. The embedding capacity of 4-dimensional symplectic ellipsoids. *Ann. of Math. (2)*, 175(3):1191–1282, 2012.

[MS19] Dusa McDuff and Kyler Siegel. Counting curves with local tangency constraints. https://arxiv.org/abs/1906.02394, Jun 2019. Preprint.

[PRV00] Bjorn Poonen and Fernando Rodriguez-Villegas. Lattice polygons and the number 12. *Amer. Math. Monthly*, 107(3):238–250, 2000.

[Ram17] Vinicius Gripp Barros Ramos. Symplectic embeddings and the lagrangian bidisk. *Duke Math. J.*, 166(9):1703–1738, 06 2017.

[Sch05] Felix Schlenk. *Embedding problems in symplectic geometry*, volume 40 of *De Gruyter Expositions in Mathematics*. Walter de Gruyter GmbH & Co. KG, Berlin, 2005.

[Sco76] Paul R. Scott. On convex lattice polygons. *Bull. Austral. Math. Soc.*, 15(3):395–399, 1976.

[Sie19a] Kyler Siegel. Computing higher symplectic capacities 1. https://arxiv.org/abs/1911.06466, Nov 2019. Preprint.

[Sie19b] Kyler Siegel. Higher symplectic capacities. https://arxiv.org/abs/1902.01490, Feb 2019. Preprint.

[Ush19] Michael Usher. Infinite staircases in the symplectic embedding problem for four-dimensional ellipsoids into polydisks. *Algebr. Geom. Topol.*, 19(4):1935–2022, 2019.

[Ust99] Ilya Ustilovsky. *Contact homology and contact structures on S(4m+1)*. PhD thesis, Stanford University, 1999.