Lagrangian spheres and cyclic quotient T-singularities

Matthew Ryan Buck
School of Mathematical Sciences
Lancaster University

A thesis submitted for the degree of
Doctor of Philosophy

August, 2025

To Darcey, Mum, and Dad.

Abstract

We study the Lagrangian isotopy classification of Lagrangian spheres in the Milnor fibre, $B_{d,p,q}$, of the cyclic quotient surface T-singularity $\frac{1}{dp^2}(1, dpq-1)$. We prove that there is a finitely generated group of symplectomorphisms such that the orbit of a fixed Lagrangian sphere exhausts the set of Lagrangian isotopy classes. Previous classifications of Lagrangian spheres have been established in simpler symplectic 4-manifolds that admit global genus 0 Lefschetz fibrations, which $B_{d,p,q}$ does not. We construct Lefschetz fibrations for which the Lagrangian spheres are isotopic to matching cycles, which reduces the problem to a computation involving the mapping class group of a surface. These fibrations are constructed using the techniques of J-holomorphic curves and Symplectic Field Theory, culminating in the construction of a J-holomorphic foliation by cylinders of T^*S^2. Our calculations provide evidence towards the symplectic mapping class group of $B_{d,p,q}$ being generated by Lagrangian sphere Dehn twists and another type of symplectomorphism arising as the monodromy of the $\frac{1}{p^2}(1, pq-1)$ singularity.

Acknowledgements

There are two people to whom I am most deeply indebted and without this thesis would not exist: my partner Darcey, whose unwavering support through the ups and downs of PhD life made this possible; and my brilliant supervisor Jonny Evans, whose guidance, limitless generosity with his time, and overflowing passion for mathematics inspired me. I thank my mum and dad for their love, support, and encouragement that enabled me to reach this point. Thanks to my friends at Lancaster for facilitating a relaxed environment, and sharing many laughs over the years. Similarly, thanks to Kai Hugtenburg and Ilaria Di Dedda.

I enjoyed many moments spent discussing maths, exploring, and sharing a drink at several conferences with my peers Nikolas Adaloglou and Johannes Hauber. Thanks to them and the wider pinwheels crew for making the discussions we had so much fun. Special thanks also to Richard Siefring for pointing me in the right direction when I was confused about some aspects of their work on intersection theory.

I am grateful to the Engineering and Physical Sciences Research Council whose funding supported my PhD.

Declaration

I declare that the work presented in this thesis is, to the best of my knowledge and belief, original and my own work. The material has not been submitted, either in whole or in part, for a degree at this, or any other university. This thesis does not exceed the maximum permitted word length of 80,000 words including appendices and footnotes, but excluding the bibliography. A rough estimate of the word count is: 32907

Matthew Ryan Buck

Contents

Chapter 1

Introduction

1.1 Lagrangian isotopy problems

Let (M, ω) be a symplectic manifold. That is, a smooth manifold M equipped with a closed, and non-degenerate 2-form ω. In symplectic topology, one is often concerned with the study of certain submanifolds called *Lagrangians*.

Definition 1.1.1. A Lagrangian submanifold L of (M, ω) is a half-dimensional submanifold[1] on which the restriction of the symplectic form to the tangent bundle TL vanishes. That is, $\dim L = \frac{1}{2} \dim M$, and $\omega|_{TL} \equiv 0$.

There are various notions of equivalence of Lagrangians, and in this work, we focus on isotopy equivalence. Before stating what this means, recall that a symplectomorphism of (M, ω) is a diffeomorphism $\phi : M \to M$ such that $\phi^* \omega = \omega$. A stronger notion is that of a Hamiltonian diffeomorphism. Let $H : M \to \mathbb{R}$ be a smooth function — referred to as a *Hamiltonian* in symplectic geometry — and consider the vector field equation

$$\iota_V \omega = \omega(V, \cdot) = -\mathrm{d}H.$$

Non-degeneracy of ω implies that this has a unique solution, which we write as V_H. The Hamiltonian flow ϕ_t^H — also called a *Hamiltonian isotopy* — of H is

[1] Elementary linear algebra implies that the dimension of a symplectic manifold is always even.

defined to be the flow of the vector field V_H. Note that $(\phi_t^H)^*\omega = \omega$, so these are symplectomorphisms. The time-1 flow ϕ_1^H is called a *Hamiltonian diffeomorphism* of (M, ω).

Definition 1.1.2. Let $L, L' \subset (M, \omega)$ be two Lagrangian submanifolds. We say that L and L' are

1. *Lagrangian isotopic* if there exists a smooth isotopy L_t such that $L_0 = L$, $L_1 = L'$, and L_t is Lagrangian for each t.

2. *symplectic isotopic* if there exists an isotopy $\phi_t : M \to M$ of symplectomorphisms such that $\phi_0 = \mathrm{id}_M$ and $\phi_1(L) = L'$.

3. *Hamiltonian isotopic* if there exists a Hamiltonian isotopy ϕ_t^H such that $\phi_1^H(L) = L'$.

These properties are listed in increasing order of strength, that is

$$\text{Hamiltonian isotopic} \Longrightarrow \text{symplectic isotopic} \Longrightarrow \text{Lagrangian isotopic.}$$

In general, these are strict implications; two Lagrangian isotopic submanifolds are not necessarily symplectic isotopic, and so on.

We focus on a particular instance of this isotopy classification: that of Lagrangian 2-spheres in a certain family of symplectic 4-manifolds. Elementary symplectic geometry shows that the above notions of isotopy equivalence are the same for Lagrangian spheres of dimension $n > 1$. This is due to basic isotopy extension results [38, §3.4] combined with the fact that the first de Rham cohomology group vanishes: $H^1(S^n) = 0$, for $n > 1$.

Classifications of Lagrangian spheres have been achieved in various cases: starting with Hind [23] who proved that every Lagrangian sphere in $S^2 \times S^2$ is Hamiltonian isotopic to the anti-diagonal[2] $\bar{\Delta} = \{(x, -x) \in S^2 \times S^2\}$. Evans [15]

[2] Here the symplectic form on $S^2 \times S^2$ is obtained by taking the direct product of a volume form on S^2 with itself, so the factors $S^2 \times \{\text{point}\}$ and $\{\text{point}\} \times S^2$ have equal symplectic area. We'll call this the *monotone* form.

then proved a similar theorem for Lagrangian spheres in certain del Pezzo surfaces.[3] These theorems give examples of symplectic manifolds where Lagrangian *knotting* does not occur, that is, two Lagrangian spheres are Lagrangian isotopic if, and only if, they are smoothly isotopic.

On the other hand, Seidel proved [44] that Lagrangian knotting occurs in general. The proof revolves around a special symplectomorphism associated to a Lagrangian sphere L called a generalised Dehn twist[4] τ_L. Many of Seidel's papers give the explicit construction of the Dehn twist, and we refer to [48].[5] The headline is that the squared twist τ_L^2 is smoothly isotopic to the identity map, but it may, or may not, be symplectically isotopic depending on the ambient symplectic manifold (M, ω). For example, in T^*S^2 equipped with the canonical symplectic form ω_{can} — defined in Equation (1.3.5) — any iterate τ^k of the Dehn twist about the zero-section is *not* isotopic to the identity map through symplectomorphisms. In fact, more is true: τ generates the *symplectic mapping class group* of $(T^*S^2, \omega_{\text{can}})$, which is defined to be the group of connected components of the group of compactly-supported symplectomorphisms, see [43]. However, the squared Dehn twist about the antidiagonal in $S^2 \times S^2$ (again with the monotone form) is symplectically isotopic to the identity map.

Despite this, classifications of Lagrangian spheres in manifolds where knotting occurs have been achieved. Consider the Milnor fibre of the A_n surface singularity. That is, the complex manifold W_n given by the equation

$$W_n = \{(z_1, z_2, z_3) \mid z_1^2 + z_2^2 + z_3^{n+1} = 1\} \subset \mathbb{C}^3.$$

Equipped with the restriction of the symplectic form $\omega_{\mathbb{C}^3} = \frac{i}{2}\sum_{i=1}^3 \mathrm{d}z_i \wedge \mathrm{d}\bar{z}_i$, this is a symplectic 4-manifold. In [24], Hind proves that any Lagrangian sphere in

[3]A symplectic del Pezzo surface is either $S^2 \times S^2$ or $\mathbb{CP}^2$ blown-up in $n < 9$ generic points equipped with an anticanonical Kähler form [15, Definition 1.3]. Evans proves that no Lagrangian knotting occurs in the $n = 2, 3, 4$ cases.

[4]In the case $n = 1$ this is the classical Dehn twist.

[5]Although, the authoritative reference for generalised Dehn twists should probably Seidel's thesis [42].

W_1 or W_2 is Lagrangian isotopic to one obtained from a finite set of "standard" Lagrangian spheres by applying Dehn twists about these standard spheres. Wu [57] then extended this result to all W_n. One can rephrase these results as saying that the only symplectic knotting of Lagrangian spheres in the A_n Milnor fibres comes from Dehn twists. The main result of this thesis proves the corresponding result for a similar class of symplectic manifolds, which are also Milnor fibres of complex surface singularities, see Theorem 1.2.2 for the statement.

1.2 Cyclic quotient surface singularities and $B_{d,p,q}$

Let $n > 1$ be an integer and let $a \geq 1$ be coprime to n. Consider the action of the group of n-th roots of unity, Γ_n, on $\mathbb{C}^2$ with weights $(1, a)$:

$$\mu \cdot (x, y) = (\mu x, \mu^a y).$$

The quotient space $\mathbb{C}^2/\Gamma_n$ is a singular manifold called the cyclic quotient surface singularity of type $\frac{1}{n}(1, a)$. Now let $d, p, q > 0$ be integers with $p > q$ coprime. The A_{dp-1} singularity is the variety

$$\{(z_1, z_2, z_3) \in \mathbb{C}^3 \mid z_1 z_2 = z_3^{dp}\},$$

and the group Γ_p acts on it with weights $(1, -1, q)$. The cyclic quotient singularity $\frac{1}{dp^2}(1, dpq - 1)$ is analytically isomorphic to the quotient A_{dp-1}/Γ_p via the map $\mathbb{C}^2/\Gamma_{dp^2} \to A_{dp-1}/\Gamma_p : (x, y) \mapsto (x^{dp}, y^{dp}, xy)$. The symplectic manifold $B_{d,p,q}$ is defined to be the Milnor fibre of this singularity. Its symplectic structure is inherited from $\mathbb{C}^3$ since the Γ_p action is by symplectomorphisms. As a result, we can explicitly write

$$B_{d,p,q} := \left(z_1 z_2 - \prod_{i=1}^{d} (z_3^p - i) \right) \Big/ \Gamma_p.$$

Remark 1.2.1. Observe that we can recover the case of the Milnor fibres of the A_{d-1} singularities, W_{d-1}, by setting $p = 1$ and $q = 1$.[6]

[6]After a holomorphic change of coordinates $(\zeta_1, \zeta_2) = (\frac{1}{2}(z_1 + z_2), \frac{1}{2i}(z_1 - z_2))$ the equation

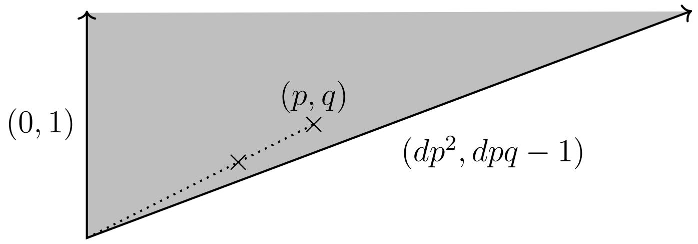

Figure 1.1: A fundamental action domain for $B_{d,p,q}$. The labels indicate the primitive integer direction of each arrow and ray. Drawn here is the case $(d,p,q) = (2,2,1)$.

Following §7.4 of [17], $B_{d,p,q}$ can be equipped with a Hamiltonian system that gives it the structure of an almost toric manifold. A fundamental action domain for this system is shown in Figure 1.1. This picture determines $B_{d,p,q}$ up to symplectomorphism, and is useful for visualising its topology. Over the dotted lines connecting the crosses live $d-1$ Lagrangian 2-spheres, and over the line connecting the unique vertex of the wedge to the leftmost cross lives a Lagrangian CW-complex called a (p,q)-pinwheel. This is a Lagrangian immersion of a 2-disc, which is an embedding on the interior, and maps p-to-1 on the boundary [30, 19]. We shall refer to these $d-1$ spheres as the *standard* Lagrangian spheres in $B_{d,p,q}$, and denote their associated Dehn twists by $\tau_1, \ldots, \tau_{d-1}$. The CW-complex formed by the topological wedge sum of these $d-1$ spheres and the (p,q)-pinwheel is called the *Lagrangian skeleton* of $B_{d,p,q}$.

Similarly to the case of Lagrangian spheres, one can define a symplectomorphism associated to a Lagrangian (p,q)-pinwheel, which we write as $\tau_{p,q}$. The construction is carried out in Chapter 5. The only thing to note for now is that $\tau_{p,q}$ is a compactly-supported symplectomorphism of $B_{d,p,q}$ with support in a neighbourhood of the pinwheel itself.

We are now ready to state the main result.

for $A_{d-1} = B_{d,1,1}$ becomes $\zeta_1^2 + \zeta_2^2 + P(z_3)$ for some polynomial P. An isotopy from $P(z) = -\prod_{i=1}^{d}(z-i)$ to $P(z) = z^d - 1$ through polynomials with pairwise distinct non-zero roots then yields the claimed identification.

Theorem 1.2.2. *Let $L \subset B_{d,p,q}$ be a Lagrangian sphere. Then L is compactly-supported Hamiltonian isotopic to a sphere obtained by applying a word of symplectomorphisms generated by $\tau_1, \dots, \tau_{d-1}$, and $\tau_{p,q}$, to a standard sphere.*

For a more precise statement, see Theorem 5.3.1.

Remark 1.2.3. This is a natural extension of Wu's result [57] on the A_n Milnor fibres. In particular, the only symplectic knotting of Lagrangian spheres comes from twists about the components of the Lagrangian skeleton of $B_{d,p,q}$.

1.3 Outline of the proof

At its core, the idea of the proof is to use the theory of Lefschetz fibrations and matching cycles to construct Lagrangian isotopies of spheres. So, we begin with a brief recap on the key features of Lefschetz fibrations that are relevant to us.

1.3.1 What is a Lefschetz fibration?

Let M be an 4-dimensional oriented manifold, and S an oriented surface. In this paper, a Lefschetz fibration will mean the following (taken from [55, Definition 3.18]):

Definition 1.3.1. A Lefschetz fibration is a smooth map $\pi : M \to S$ with finitely many critical points M_{crit}, and, for each $p \in M_{\text{crit}}$, there exist complex coordinate charts at $p \in M$ and $\pi(p) \in S$ agreeing with the orientations, such that, in these coordinates, we have

$$\pi(z_1, z_2) = z_1^2 + z_2^2. \tag{1.3.1}$$

Remark 1.3.2. 1. In words, this can be interpreted as saying that π is a fibre bundle away from a (real) codimension 2 subset of singular fibres, each of which is modelled on the A_1 node. The fibres are two dimensional surfaces, the critical fibres being singular. We will often assume that each fibre has *at most 1 singularity.*

2. In general, a Lefschetz fibration isn't actually a fibration in the topological sense. Indeed, the most basic example

$$\pi : \mathbb{C}^2 \to \mathbb{C} : \pi(z_1, z_2) = z_1^2 + z_2^2$$

reveals that the fibres are not all homotopy equivalent. The smooth fibres $\pi^{-1}(z)$ $(z \neq 0)$ are homeomorphic to $\mathbb{C}^\times = \mathbb{C}\backslash\{0\}$, whilst $\pi^{-1}(0)$ is the wedge sum of two planes $\mathbb{C} \wedge \mathbb{C}$.

To introduce symplectic geometry to the story, we consider the situation where the ambient space M is equipped with a closed 2-form Ω that restricts to a positively oriented area form on the (smooth parts of the) fibres of π. This implies that each fibre $M_z = \pi^{-1}(z)$ is a symplectic manifold with symplectic form $\Omega|_{M_z}$. This extra structure picks out a connection on the tangent bundle TM by taking the symplectic orthogonal complement to the tangent spaces of the fibres: for $p \in M_z \backslash M_{\text{crit}}$ define the *horizontal space* $H(p)$ to be

$$H(p) := T_p M_z^\Omega := \{v \in T_p M \mid \Omega(v, w) = 0, \, \forall w \in T_p M_z\}.$$

Given an embedded smooth path $\gamma : [0, 1] \to S\backslash S_{\text{crit}}$, we define the parallel transport along γ as with any fibre bundle with connection (see [31, Chapter II], for example). Explicitly, since $\mathrm{d}\pi(p)|_{H_p} : H_p \to T_{\pi(p)} S$ is a linear isomorphism, we can define the (unique) vector field $\tilde{X}$ (defined on the pullback $\gamma^*\pi$) by the condition

$$\mathrm{d}\pi(p)|_{H_p}(\tilde{X}(p)) = \frac{\mathrm{d}\gamma}{\mathrm{d}t}(\pi(p)).$$

The parallel transport of γ is defined to be the map $\tau_\gamma : M_{\gamma(0)} \to M_{\gamma(1)}$ which integrates the vector field $\tilde{X}$.

Remark 1.3.3. Of course, we need some condition to tell us that τ_γ is well-defined. For example, it is enough to assume that π has closed fibres, as we'll make use of in this thesis. For other cases we'll justify the existence of τ_γ as and when we need. A reference on Lefschetz fibrations in the *exact* symplectic setting is [47, Part III].

1.3.1.1 Vanishing and matching cycles

With this in hand, we can define the notions of vanishing paths, cycles, and thimbles. A *vanishing path* is an embedded path $\gamma : [0,1] \to S$ such that $\gamma^{-1}(S_{\text{crit}}) = \{1\}$. Let $p_1 \in M_{\gamma(1)}$, be a critical point and consider the parallel transport for the restriction $\gamma_s := \gamma|_{[0,s]}$, $s < 1$. Define the map $\tau_1 : M_{\gamma(0)} \to M_{\gamma(1)}$ from the smooth fibre $M_{\gamma(0)}$ to the singular one $M_{\gamma(1)}$ by

$$\tau_1 := \lim_{s \to 1} \tau_{\gamma_s}.$$

This is well-defined and continuous [18, Lemma 1.2]. Therefore, we can define the *vanishing cycle* $V_\gamma \subset M_{\gamma(0)}$ to be the locus of points that are transported to the singular point p_1:

$$V_\gamma = \{p \in M_{\gamma(0)} \mid \tau_1(p) = p_1 \in M_{\gamma(1)} \cap M_{\text{crit}}\}.$$

Define also the *vanishing thimble* associated to γ to be

$$\Delta_\gamma := \bigcup_{s \in [0,1]} \tau_{\gamma_s}(V_\gamma).$$

That is, Δ_γ is the trace of the vanishing cycle under the parallel transport maps τ_{γ_s}. It is an embedded 2-ball, as one can prove by expressing Δ_γ as the stable manifold of a hyperbolic vector field [46, Lemma 1.13]. Moreover, it satisfies[7] $\Omega|_{T\Delta_\gamma} \equiv 0$. The vanishing cycle V_γ is the boundary of Δ_γ, which is therefore a Lagrangian[8] S^1 in $M_{\gamma(0)}$.

Similarly there are matching paths and cycles. A *matching path* $\gamma : [-1,1] \to S$ is an embedded path such that $\gamma^{-1}(S_{\text{crit}}) = \{-1, 1\}$. Therefore, the restrictions $\gamma^+ := \gamma|_{[0,1]}$, and $\gamma^- := \gamma|_{[-1,0]}$ define two vanishing paths, for which we form the vanishing thimbles $\Delta_{\gamma^\pm}$. If the corresponding vanishing cycles $V_{\gamma^\pm}$ are equal, we

---

[7] We avoid saying the word *Lagrangian* here, since we aren't assuming that the 2-form Ω is symplectic on the ambient space M.

[8] Of course, saying Lagrangian here is redundant, since any 1-dimensional submanifold of a 2-dimensional symplectic manifold is Lagrangian. However, in higher dimensions this is still true and thus non-trivial.

can glue the two thimbles together to form a Lagrangian 2-sphere Σ_γ called the *matching cycle*. Of course, the condition that $V_{\gamma^+} = V_{\gamma^-}$ is not always satisfied, so one needs to be slightly more careful to define Σ_γ in general. Provided that $V_{\gamma^\pm}$ are symplectically isotopic in $M_{\gamma(0)}$, these complications can be surmounted, but we delay talking about them for now. For the full construction in the exact case, see [47, 16g].

1.3.1.2 Matching cycles and Lagrangian isotopies

Consider an isotopy of matching paths $\gamma_s : [-1, 1] \to S$. By finiteness of the set of critical values S_{crit}, this is an isotopy rel the end points: $\gamma_s(\pm 1) = \gamma_0(\pm 1)$. Such an isotopy gives rise to a Lagrangian isotopy of the corresponding matching cycles. That is, Σ_{γ_s} is a Lagrangian isotopy from Σ_{γ_0} to Σ_{γ_1}. Therefore, we have a map

$$\{\text{matching paths}\}/\text{isotopy} \to \{\text{Lagrangian spheres}\}/\text{Lagrangian isotopy}. \quad (1.3.2)$$

In general, this map is neither injective nor surjective, meaning that two non-isotopic matching paths can give rise to Lagrangian isotopic matching cycles, and not every Lagrangian sphere need be realised as the matching cycle of a matching path [47, Example 16.12]. The grand plan of the proof of Theorem 1.2.2 is to show that, in the case of $B_{d,p,q}$, there exists a Lefschetz fibration[9] such that the above map is a *bijection*. Most of the work involved is proving surjectivity, for which we employ the theory of J-holomorphic curves. We state this as a theorem, as it is the main technical achievement of this thesis:

Theorem 1.3.4 (Corollary 4.3.3)**.** *There exists a genus 0 Lefschetz fibration defined on an open dense subset of $B_{d,p,q}$ such that any Lagrangian sphere L is Lagrangian isotopic to a matching cycle.*

The main theorem (Theorem 1.2.2) follows from this and a computation involving the theory of mapping class groups of surfaces, see Chapter 5.

[9] A Lefschetz fibration *of sorts*. See Section 2.5.

Remark 1.3.5. Contrast the above with work of Auroux, Muñoz, and Presas [3] where they prove a similar theorem. However, they construct a Lefschetz fibration for each Lagrangian sphere, whereas we construct a *single* Lefschetz fibration for which all Lagrangian spheres are isotopic to matching cycles.

1.3.2 What is a J-holomorphic curve?

Most of the technical drive behind the proof comes from the theory of J-holomorphic curves in symplectic manifolds.

Definition 1.3.6. Let M be a smooth manifold. An *almost complex structure* on M is an endomorphism of the tangent bundle $J : TM \to TM$ such that $J^2 = -1$. Suppose we have a Riemann surface (Σ, j) and a map $u : \Sigma \to M$. We say that u is a (j, J)*-holomorphic curve* (or simply J-holomorphic if the complex structure on the domain is understood) if it satisfies the differential equation

$$J \circ \mathrm{d}u = \mathrm{d}u \circ j. \tag{1.3.3}$$

The classic example of a J-holomorphic curve is a holomorphic curve in a complex manifold. They have seen fantastic applications in symplectic topology — especially in dimension 4 — since Gromov originally introduced them in a groundbreaking paper [22] in 1985.

We'll use J-holomorphic curves to construct Lefschetz fibrations on $B_{d,p,q}$ that are somehow adapted to the Lagrangian sphere $L \subset B_{d,p,q}$. Eventually, we end up with a Lefschetz fibration for which L is Lagrangian isotopic to a matching cycle, proving surjectivity of the map in Equation (1.3.2).

The basic idea is as follows. Given an almost complex structure J, we consider a family of J-holomorphic curves, called a *moduli space*, written $\overline{\mathcal{M}}(J)$, comprised of an open dense subset of smooth curves $\mathcal{M}(J)$, and finitely many singular curves $\partial\overline{\mathcal{M}}(J) := \overline{\mathcal{M}}(J)\backslash\mathcal{M}(J)$. The singular curves come in two types: genus zero nodal curves with exactly two components, and a unique *exotic* curve which generally has many genus 0 components. On the complement of the exotic curve, the natural map $\pi_J : B_{d,p,q} \to \overline{\mathcal{M}}(J)$ defined by mapping a point to the unique curve it lies on is a Lefschetz fibration. Therefore, we seek an almost complex structure J — adapted to L in some way — such that L is Lagrangian isotopic to a matching cycle of π_J. We will construct such a J through a technique called neck stretching (the details of which are outlined in Section 1.3.3). This is not quite enough, however, since we want a *single* Lefschetz fibration whose matching cycles exhaust the set of Lagrangian isotopy classes of spheres. We resolve this via a basic argument involving the connectivity of the space of almost complex structures. Indeed, a 1-parameter family of almost complex structures J_s yields a 1-parameter family of Lefschetz fibrations $\pi_s := \pi_{J_s}$ and corresponding Lagrangian isotopies between their matching cycles. With this in mind, fix a reference almost complex structure J_{ref} and its corresponding Lefschetz fibration $\pi_{\text{ref}} := \pi_{J_{\text{ref}}}$. Then we can choose a homotopy J_s from $J_0 = J$ to $J_1 = J_{\text{ref}}$ to obtain a Lagrangian isotopy from L to a matching cycle of π_{ref}.

A few comments are in order about the above sketch. Firstly, we discuss the smoothness of π_J and $\overline{\mathcal{M}}(J)$. The subset $\mathcal{M}(J)$ is a smooth manifold under certain conditions. Proving this is a long and technical argument involving applying the implicit function theorem in infinite dimensions. A comprehensive reference on the case of closed J-holomorphic curves is [37] (in particular, see Chapter 3 therein for the proof of smoothness of the moduli space). Fortunately for us, in dimension 4 there are exceptionally powerful *automatic transversality* results [25, 52] which guarantee that the moduli space of smooth curves $\mathcal{M}(J)$ is a smooth manifold under simple conditions. For example, in the case of a closed genus 0 curve $u : S^2 \to M$, a neighbourhood of u is a smooth manifold if

$$c_1(u) = c_1(u^*TM) > 0,$$

where $c_1(u^*TM)$ is the first Chern number[10] of u^*TM.

One can assign a topology to $\overline{\mathcal{M}}(J)$ which agrees with the manifold topology on the subset $\mathcal{M}(J)$. Indeed, a remarkable insight of the paper [22] demonstrated that

[10]For an embedded J-holomorphic curve $u : \Sigma \to M$ with $u_*[\Sigma] = A \in H_2(M;\mathbb{Z})$ the adjunction

there is a well-behaved compactness theory for J-holomorphic curves in the presence of a symplectic form which *tames* J, in the sense that, for all $v \in TM$ non-zero,

$$\omega(v, Jv) > 0. \tag{1.3.4}$$

The full details of the theory we tacitly use here can be found in [37, Chapters 4–6]. To sketch what can happen, take a sequence of closed J-holomorphic curves $u_\nu : \Sigma \to M$ with uniformly bounded symplectic area,[11] that is, there exists $C > 0$ such that

$$E(u_\nu) := \int_\Sigma u_\nu^* \omega < C.$$

Suppose further that the L^∞-norm of the derivatives is unbounded:

$$\sup_\nu \|\mathrm{d}u_\nu\|_{L^\infty} = \infty.$$

Then u_ν has a subsequence that converges to a *singular* J-holomorphic curve. This type of convergence is known as *bubbling*, since a finite portion of the energy $E(u_\nu)$ gets concentrated to a point as $\nu \to \infty$, and (after reparametrising) one sees another J-holomorphic curve emerging at this point. In this thesis, we will largely deal with the simplest case where the limiting singular curve is a nodal curve with two genus 0 components. The example of the family of smooth conics $(xy = \epsilon z^2) \subset \mathbb{CP}^2$ degenerating to the singular one $(xy = 0)$ as $\epsilon \to 0$ illustrates the picture well.[12]

Under this notion of convergence, $\overline{\mathcal{M}}(J)$ will be a (compact) metric space, but in general it is no longer a manifold. However, in the simple cases considered here, the

formula [37, Theorem 2.6.4] states that

$$c_1(u) = A \cdot A + \chi(\Sigma).$$

[11] This uniform bound will always be satisfied for us, since we will only consider sequences of curves that represent the same homology class.

[12] The parametrisation $u_\epsilon : \mathbb{C} \to \mathbb{CP}^2 : u_\epsilon(w) = [w^2 : \epsilon : w]$ converges C^∞_{loc} to the line $u_0^+(w) = [w : 0 : 1]$. In this example, the bubbling happens at the point $w = 0$. Using the map $\phi_\epsilon : B_{\frac{1}{\epsilon}}(0) \to \mathbb{C} : \phi_\epsilon(w) = \epsilon w$, we reparametrise u_ϵ to see the bubble forming: $u_\epsilon \circ \phi_\epsilon$ converges to $u_0^-(w) = [0 : 1 : w]$.

complement of the exotic curve can be equipped with a smooth structure that agrees with that given by the implicit function theorem on the locus of smooth curves.

Secondly, we need to justify the claim that the maps $\pi_J : B_{d,p,q} \to \overline{\mathcal{M}}(J)$ are actually Lefschetz fibrations. The singular curves with two components that appear in $\overline{\mathcal{M}}(J)$ form the critical fibres of the map $\pi_J : B_{d,p,q} \to \overline{\mathcal{M}}(J)$. This claim implicitly depends on showing that the local coordinate form in Equation (1.3.1) holds at the nodal points of these curves, which follows from an argument of Wendl and Lisi sketched in [55, Appendix A]. We discuss it here in Section 2.5.

Thirdly, we outline why 1-parameter families of almost complex structures J_s give rise to 1-parameter families of Lefschetz fibrations π_s. This amounts to showing that, as we vary the almost complex structure J, the structure of the moduli spaces $\overline{\mathcal{M}}(J_s)$ is essentially unchanged. In practise it requires showing that no further bubbling occurs as we vary J, which is done in Section 2.4.2. This analysis is heavily reliant on two big results: (1) automatic tranversality, which ensures that when we deform J there are still solutions to Equation (1.3.3); and, (2) positivity of intersections, which gives strong geometric constraints on J-holomorphic curves. Positivity of intersections was first proved in full generality by McDuff [36]. The statement relevant to us is:

Theorem 1.3.7 (Positivity of intersections)**.** *Let $u : \Sigma \to M$, and $v : \Sigma' \to M$ be two closed J-holomorphic curves such that $u^{-1}(\operatorname{im} v)$ contains no non-empty open set. Then the intersection number $u \cdot v = u_*[\Sigma] \cdot v_*[\Sigma']$ is non-negative, and it equals 0 if, and only if, u and v are disjoint.*

Finally, we briefly discuss the motivation for using neck stretching to construct an almost complex structure adapted to L, and specifically, why one would expect this to produce a Lefschetz fibration with a matching cycle isotopic to L. The starting point is Weinstein's Lagrangian neighbourhood theorem [51]. Recall the canonical 1-form λ_{can} defined on the cotangent bundle of any smooth manifold M. Given coordinates $(q_1, \ldots, q_n)$ on M, let $(p_1, \ldots, p_n, q_1, \ldots, q_n)$ be the induced coordinates on T^*M, so that the point $(p_1, \ldots, p_n, q_1, \ldots, q_n)$ represents the covector $\sum_{i=1}^{n} p_i \mathrm{d}q_i$.

In these coordinates, λ_{can} has the following form:

$$\lambda_{\text{can}} = \sum_{i=1}^{n} p_i \mathrm{d}q_i.$$

We define the canonical symplectic form on T^*M by

$$\omega_{\text{can}} = \mathrm{d}\lambda_{\text{can}}. \tag{1.3.5}$$

Notice that the zero-section $0_M \subset T^*M$ is Lagrangian with respect to ω_{can}. Weinstein's Lagrangian neighbourhood theorem says that any Lagrangian $L \subset (N, \omega)$ diffeomorphic to M in some symplectic manifold (N, ω) admits a neighbourhood symplectomorphic to a neighbourhood of $0_M \subset (T^*M, \omega_{\text{can}})$. This means we can understand the symplectic geometry locally in a neighbourhood of a Lagrangian simply by considering its cotangent bundle.

The following example acts as a local model for a Lefschetz fibration in a neighbourhood of a Lagrangian sphere:

Example 1.3.8. Let $Q = (z_1^2 + z_2^2 + z_3^2 = 1) \subset \mathbb{C}^3$ be the affine quadric equipped with the restriction of the standard symplectic form on $\mathbb{C}^3$. This is symplectomorphic to $(T^*S^2, \omega_{\text{can}})$ [42, Lemma 18.1] via a symplectomorphism that identifies the zero-section with the real locus of Q. The projection $\pi_Q : Q \to \mathbb{C} : \pi(z) = z_3$ is a genus 0 Lefschetz fibration with exactly two singular fibres $\pi_Q^{-1}(\pm 1)$. Moreover, the matching cycle of the path $\gamma : [-1, 1] \to \mathbb{C} : \gamma(t) = t$ is exactly the real locus $\Re Q$. See Figure 1.2 for a cartoon of π_Q.

Weinstein's Lagrangian neighbourhood theorem allows us to compare Lefschetz fibrations on $B_{d,p,q}$ to π_Q by restricting them to a neighbourhood of a Lagrangian sphere L. Starting with an initial almost complex structure J, neck stretching yields a specific deformation J_t of J such that, for $t \gg 0$, the J_t-holomorphic curves that pass through a fixed neighbourhood of L closely resemble the fibres of π_Q. As L is identified with the real locus $\Re Q$ in the above model, and $\Re Q$ is a matching cycle of π_Q, one might hope that, in some sense, L is close to being a matching cycle of π_{J_t} for large t. We formalise this in Chapter 4 and prove that t can be chosen large enough so that L is Lagrangian isotopic to a matching cycle of π_{J_t}.

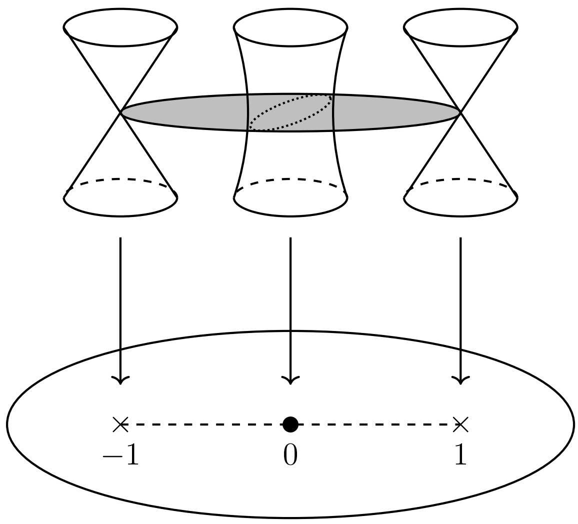

Figure 1.2: A cartoon of the Lefschetz fibration $\pi_Q : Q \to \mathbb{C}$ on the affine quadric Q. The shaded sphere is the real locus, which is the matching cycle corresponding to the dashed matching path drawn in the base. It intersects the regular central fibre in the vanishing cycle, indicated in the figure by the dotted circle.

1.3.3 What is neck stretching?

Neck stretching, or the splitting construction, is a process by which we alter an almost complex structure in a neighbourhood of a particular hypersurface of a symplectic manifold. We refer to [14, §1.3], [5, §3.4], or [9, §2.7] for the full construction, as we will only need a special case of it here. Note that the Lagrangian isotopy theorems mentioned earlier [23, 15] crucially rely on it.

Consider a Lagrangian L of a closed symplectic manifold (X, ω) and equip L with a Riemannian metric. This allows us to define the sphere bundle of covectors of a fixed length:

$$T_R^*L := \{v \in T^*L \mid |v| = R\}.$$

Along with the restriction of the canonical 1-form λ_{can}, the pair $(T_R^*L, \lambda_{\text{can}})$ is an example of a contact manifold. For us, this will mean the following:

Definition 1.3.9. Let M be a manifold of dimension $2n - 1$. A *contact form*

$\lambda \in \Omega^1(M)$ is a 1-form satisfying

$$\lambda \wedge \mathrm{d}\lambda^{n-1} \neq 0.$$

The distribution $\xi = \ker \lambda$ is called the *contact structure* and the pair (M, ξ) is a *contact manifold*. Notice that the pair $(\xi, \mathrm{d}\lambda)$ forms a symplectic vector bundle. That is, $\mathrm{d}\lambda$ restricts to a non-degenerate form on each contact plane ξ_p.

The manifold $\mathbb{R} \times M$ equipped with the 2-form $\mathrm{d}(e^r\lambda)$ is symplectic and the pair $(\mathbb{R} \times M, \mathrm{d}(e^r\lambda))$ is called the *symplectisation* of (M, λ).

Remark 1.3.10.

1. We give this general definition to condense the notation slightly, although the reader should keep in mind that we will only be interested in the case where $(M, \lambda) = (T^*_R L, \lambda_{\text{can}})$ as above.

2. For an example of neck stretching around a Lagrangian submanifold other than a sphere see the work of Dimitroglou-Rizell, Ivrii, and Goodman [11] where they prove the nearby Lagrangian conjecture for T^*T^2.Examples of stretching around non-orientable Lagrangian surfaces can be found in [39, 16].

3. The map

$$\Phi : \mathbb{R} \times T^*_R L \to T^*L \backslash 0_L : \Phi(r, v) = e^r v$$

is an *exact symplectomorphism*. That is, it satisfies,

$$\Phi^*\lambda_{\text{can}} = e^r \lambda_{\text{can}}|_{T^*_R L}.$$

Therefore, by the Weinstein Lagrangian neighbourhood theorem, some portion $[-\epsilon, \epsilon] \times T^*_R L$ symplectically embeds into M. We'll call this the *neck region*, or simply the *neck*.

On symplectisations, we can define a special class of almost complex structures, but to do so, we need to take the taming condition (1.3.4) one step further:

Definition 1.3.11. An almost complex structure J on a symplectic vector bundle $(E, \omega) \to X$ is called *compatible* with ω when it satisfies the taming condition of

Equation (1.3.4) and ω is J-invariant: for all $u, v \in E$,

$$\omega(Ju, Jv) = \omega(u, v). \tag{1.3.6}$$

This is equivalent to saying that the tensor field g defined by

$$g(u, v) := \omega(u, Jv)$$

is a bundle metric on E.

On $\mathbb{R} \times M$ there are two special vector fields: ∂_r given by the $\mathbb{R}$-coordinate, and the *Reeb vector field* R_λ defined by

$$\lambda(R_\lambda) \equiv 1, \text{ and } \iota_{R_\lambda} \mathrm{d}\lambda \equiv 0.$$

Definition 1.3.12. We say that an almost complex structure J on a symplectisation $\mathbb{R} \times M$ is *cylindrical* when

1. $J\partial_r = R_\lambda$,
2. J is invariant under $\mathbb{R}$-translation, and
3. the restriction $J|_\xi$ of J to the contact structure ξ is a compatible almost complex structure on the symplectic vector bundle $(\xi, \mathrm{d}\lambda) \to M$.

Since the neck region is symplectomorphic to part of a symplectisation, we can consider the class of almost complex structures J on X that restrict to a cylindrical almost complex structure J_M on the neck $[-\epsilon, \epsilon] \times M$. These are sometimes called almost complex structures that are *adapted* to the hypersurface $\{0\} \times M \subset X$, but we will often abuse language and call them cylindrical. The process of neck stretching is to start with one of these adapted almost complex structures and replace the neck with larger and larger portions $[-t - \epsilon, \epsilon] \times M$ of $\mathbb{R} \times M$ along with an almost complex structure that is cylindrical along this longer neck. More precisely, following [9, §2.7] excise the neck region to obtain a compact manifold[13] $Y := X \backslash (-\epsilon, \epsilon) \times M$

[13] In [9] this is denoted by $\bar{X}_0$.

with boundary $\partial Y = M_- \sqcup M_+$, where $M_\pm = \{\pm\epsilon\} \times M \subset X$. Form the *stretched* manifold X_t by

$$X_t := Y \cup_{M_- \sqcup M_+} [-t-\epsilon, \epsilon] \times M,$$

which is diffeomorphic to X. Define the almost complex structure J_t on X_t by

$$J_t = \begin{cases} J, & \text{on } Y \\ J_M, & \text{on } [-t-\epsilon, \epsilon] \times M. \end{cases} \tag{1.3.7}$$

The family of almost complex structures J_t is called a *neck stretch* of J.

A neck stretch J_t has no well-defined limit on X as $t \to \infty$. However, due to their $\mathbb{R}$-invariant nature on the neck, they can be seen to converge (on compact subsets) to almost complex structures on either the symplectisation $\mathbb{R} \times M$, or the *completed*[14] symplectic manifold $\hat{Y} = Y \cup_{M_-} [-\epsilon, \infty) \times M \cup_{M_+} (-\infty, \epsilon] \times M$. Taking a monotonic sequence $t_k \to \infty$ and the corresponding neck stretch $J_k := J_{t_k}$, we consider sequences $f_k : \Sigma \to X_k$ of J_k-holomorphic curves with uniformly bounded energy.[15] To make sense of a limit of f_k as $k \to \infty$ we need to recall the notion of a *holomorphic building*. Essentially, this is a map $F : \Sigma^* \to X^*$ from a (potentially disconnected and punctured) Riemann surface $\Sigma^* = \bigsqcup_{\nu=0}^{N} \Sigma^{(\nu)}$ to the manifold

$$X^* = \hat{Y} \sqcup \bigsqcup_{\nu=1}^{N} \mathbb{R} \times M,$$

where the restriction $F^{(\nu)} = F|_{\Sigma^{(\nu)}}$ to each *level* maps into

$$X^{(\nu)} = \begin{cases} \hat{X}, & \text{if } \nu = 0, N+1 \\ \mathbb{R} \times M, & \text{if } \nu = 1, \dots, N. \end{cases}$$

[14]The completion of a symplectic manifold W with contact-type boundary (M, λ) is obtained by gluing on infinite half-cylinders of the form $[0, \infty) \times M$ or $(-\infty, 0] \times M$. The distinction between these cases comes from the fact that each connected component M' of M has a collar neighbourhood symplectomorphic to $((-\epsilon, 0] \times M', \mathrm{d}(e^{\pm r}\lambda))$. The sign of ∞ in the intervals agrees with the sign of r in $\mathrm{d}(e^{\pm r}\lambda)$.

[15]The definition of energy is a little technical in general [5, §6.1]. However, in our case, as we are neck stretching around a contact-type hypersurface in a closed manifold, the energy bound will be automatically satisfied by considering sequences of homologous curves.

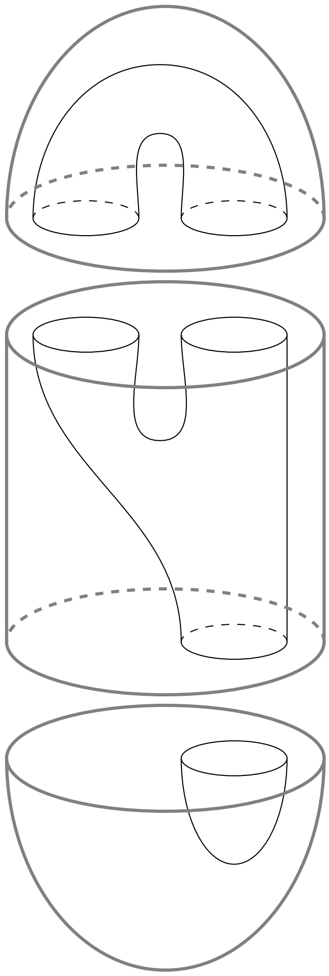

Figure 1.3: An example of a genus 1 holomorphic building with 3 levels: a top and bottom, and one symplectisation level in between.

The limits of the neck stretch sequence J_t fit together into an almost complex structure J^* on X^*, with respect to which F is J^*-holomorphic. That is, $F^{(\nu)}$ is a $J^{(\nu)}$-holomorphic curve in $X^{(\nu)}$. One should imagine the curves f_k becoming stretched so much as $k \to \infty$ that they "break" into multiple curves in the limit. See Figure 1.3 for a picture illustrating the level structure of a holomorphic building. See also [5, Figure 11] and [9, Figure 1] for more sophisticated cartoons of what happens.

We will only consider the case where M is a separating hypersurface in X, meaning that $Y = X \backslash (-\epsilon, \epsilon) \times M$ has two connected components[16] $Y_\pm$. This means that $X^{(0)} = \hat{Y}$ decomposes into two components $X_\pm^{(0)}$, which are the completions of $Y_\pm$. The 0-level of a holomorphic building then naturally splits into two components called the *top* and *bottom* levels: $\Sigma^{(0)} = \Sigma_+^{(0)} \sqcup \Sigma_-^{(0)}$ and $F_\pm^{(0)} := F|_{\Sigma_\pm^{(0)}}$ maps into $X_\pm^{(0)}$. As the completions of the symplectic manifolds $Y_\pm$, $X_\pm^{(0)}$ come equipped with symplectic forms $\omega^\pm$ satisfying:

$$(X_+^{(0)}, \omega^+) \cong (X \backslash L, \omega) \text{ and } (X_-^{(0)}, \omega^-) \cong (T^*L, \mathrm{d}\lambda_{\mathrm{can}}).$$

Moreover, the stretched manifolds X_t can be equipped with natural symplectic structures [9, Example 2.4]:

$$\omega_t = \begin{cases} \omega, & \text{on } Y_+, \\ \mathrm{d}(e^r \lambda), & \text{on } [-t-\epsilon, \epsilon] \times M, \\ e^{-t}\omega, & \text{on } Y_-, \end{cases}$$

which converge (on compact subsets and after rescaling by e^t in the case of $X_-^{(0)}$) to the familiar symplectic forms ω on $X_+^{(0)} \cong X \backslash L$ and $\mathrm{d}\lambda_{\mathrm{can}}$ on $X_-^{(0)} \cong T^*L$.

The power of this technique will come from our understanding of J-holomorphic curves in $X_-^{(0)} \cong T^*L \cong T^*S^2$. It follows from the techniques of Hind [23] and automatic transversality [52] that any cylindrical almost complex structure on T^*S^2 admits two transverse foliations by J-holomorphic planes (once punctured spheres). These are analogous to the transversely intersecting foliations of $S^2 \times S^2$ by horizontal $[S^2 \times \{\text{point}\}]$ and vertical $[\{\text{point}\} \times S^2]$ spheres. The relevant facts are collected in Section 3.3.

Using these foliations, and intersection theory for punctured J-holomorphic curves [50], we deduce what happens to the curves in the moduli spaces $\overline{\mathcal{M}}(J_t)$ as $t \to \infty$. This is the purpose of Section 3.4. They give rise to a J-holomorphic foliation of T^*S^2 by cylinders (copies of $\mathbb{C}^\times = \mathbb{C} \backslash \{0\}$) in the limit, whose leaves form the fibres of a Lefschetz fibration akin to π_Q of Example 1.3.8. The crucial

[16]The sign of $Y_\pm$ corresponds to that of the boundary $\partial Y_\pm = \{\pm\epsilon\} \times M$.

consequence of this is that our Lagrangian sphere L will be a matching cycle of the limiting Lefschetz fibration. We then apply a uniform convergence argument to show that, for sufficiently large neck stretches (sufficiently large t), L is Lagrangian isotopic to a matching cycle of one of the Lefschetz fibrations $\pi : B_{d,p,q} \to \overline{\mathcal{M}}(J_t)$. This is covered in Chapter 4.

1.3.4 Compactifying $B_{d,p,q}$

The previous discussion is almost enough to prove Theorem 1.3.4. However, we have oversimplified some details for the sake of narrative clarity. Perhaps the most glaring omission is that the domain of the Lefschetz fibration $\pi : B_{d,p,q} \to \overline{\mathcal{M}}(J)$ is non-compact, but in Section 1.3.3 we quietly assumed that the ambient manifold X was *closed*. The construction of the map π actually goes as follows. First, in Section 2.2, we find a suitable compactification $X_{d,p,q}$ (of a suitable subset) of $B_{d,p,q}$ whose symplectic geometry reflects enough of that of $B_{d,p,q}$ to be able to construct the Lagrangian isotopy in $X_{d,p,q}$ and the pull this back to $B_{d,p,q}$. Then we consider the moduli space of genus 0 curves living in a distinguished homology class $F \in H_2(X_{d,p,q}; \mathbb{Z})$ satisfying $F^2 = F \cdot F = 0$, and show that it has the nice compactification and invariance (under deformations of J) properties described in Section 1.3.2. This is carried out in Section 2.4.

The purpose of compactifying is to make use of the simpler theory of closed J-holomorphic curves in closed symplectic manifolds. Compactifying gives a method to control the bubbling that occurs as we vary J, and ultimately prove that this is well behaved (Corollary 2.4.16). Otherwise, working directly with $B_{d,p,q}$ would force one to work with punctured curves. This approach would need an alternative argument to control bubbling/breaking.[17]

The method of producing the compactification $X_{d,p,q}$ comes from the almost

[17] Although, the author believes this is achievable and even desirable, in fact, since this would avoid annoying arguments like Lemma 4.3.2. However, the compactness theory for punctured curves is more complicated.

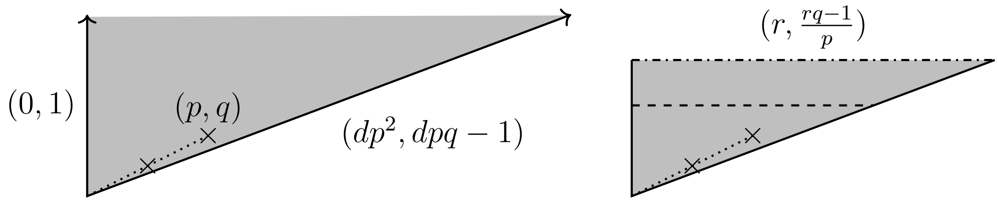

Figure 1.4: The symplectic cut performed to "chop off" the non-compact end of $B_{d,p,q}$. The cut is indicated by the dash-dot line which has primitive integer direction $\left(r, \frac{rq-1}{p}\right)$, where r is the unique integer $0 < r < p$ such that $rq \equiv 1 \mod p$. The preimage of the area below the horizontal dashed line is a compact symplectic submanifold with boundary whose symplectic completion is symplectomorphic to $B_{d,p,q}$. Drawn here is the most basic example $(d,p,q) = (2,2,1)$. However, this is not indicative of the general case: when $q > 1$ the dash dot line corresponding to the symplectic cut will have positive gradient. See Figure 1.5 for another example.

toric structure on $B_{d,p,q}$. Using a technique called symplectic cut (originally due to Lerman [33]) one can essentially "chop off" the non-compact end of $B_{d,p,q}$ to obtain a closed symplectic manifold (see Figure 1.4). Doing so introduces a number of singularities, which must be resolved to produce a smooth manifold. All this is done in Lemma 2.2.1. The requisite knowledge is presented lucidly in [17].

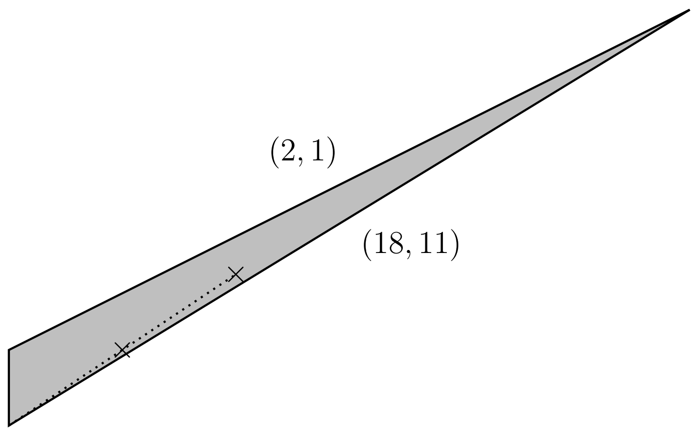

Figure 1.5: Another example of $B_{d,p,q}$ after cutting. Here $(d, p, q) = (2, 3, 2)$ is drawn. As is clear, the pictures quickly become unwieldy in this form. See Figure 2.7 for an integral affine transformation of the fundamental action domain of $B_{d,p,q}$ that behaves better under increasing the values (d, p, q).

Chapter 2

A compactification of $B_{d,p,q}$ and its Lefschetz fibrations

This chapter achieves two goals. Firstly, we use an almost toric structure on $B_{d,p,q}$ to construct a compactification $X = X_{d,p,q}$ via a series of symplectic cuts (see Lemma 2.2.1). Secondly, we show that for every almost complex structure J on X satisfying some constraints (but importantly including any almost complex structure arising from neck stretching) there exists a J-holomorphic foliation of X whose leaves form a Lefschetz fibration (Proposition 2.5.1) on an open dense subset.

2.1 Spiel on toric base diagrams

In this chapter we use a little of the theory of (almost) toric geometry. As mentioned in the introduction, the requisite knowledge can be found in [17] and the references therein. We state the main definition of a *Delzant polytope* here for convenience (taken from Definition 3.5 in [17]).

Definition 2.1.1. A rational convex polytope P (which we will call a *moment polytope* or simply a *polytope*) is a subset of $\mathbb{R}^n$ defined as the intersection of a finite collection of half spaces $\{x \in \mathbb{R}^n \mid \sum_{i=1}^n \alpha_i x_i \leq b\}$ with $\alpha_i \in \mathbb{Z}$ and $b \in \mathbb{R}$. We say that P is a *Delzant polytope* if it is a convex rational polytope such that every point on a k-dimensional facet has a neighbourhood that is integral affine isomorphic (an isomorphism of the form $Ax + b$ with $A \in SL_n(\mathbb{Z})$ and $b \in \mathbb{R}^n$) to a neighbourhood of the origin in the polytope $[0, \infty)^{n-k} \times \mathbb{R}^k$. A vertex of a polytope is called *Delzant* if the germ of the polytope at that vertex is Delzant.

Remark 2.1.2.

1. The Delzant condition is important since it means that the symplectic manifold corresponding to the polytope is smooth.

2. Importantly, as we only work with 4-dimensional symplectic manifolds (corresponding to $n = 2$ in the above definition), a vertex v of a polygon is Delzant if, and only if,[1] the primitive integer vectors u and w pointing along the edges that meet at v form a matrix $u \wedge w$ with determinant ± 1.

Let $p, q \in \mathbb{Z}$ be such that $\gcd(p, q) = 1$, then write $0 < r < p$ for the unique integer such that $qr \equiv 1 \mod p$. Consider the singular toric manifold with the non-Delzant moment polygon, $\Pi(p, q)$, shown in Figure 2.1. As in [17, §4.5], we can use

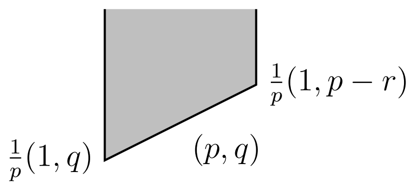

Figure 2.1: The moment polygon $\Pi(p, q)$.

symplectic cuts to resolve the cyclic quotient singularities present in the symplectic orbifold corresponding to the polygon $\Pi(p, q)$. The minimal resolution of $\Pi(p, q)$ is a smooth toric manifold with moment polygon sketched in Figure 2.2.

Denote the continued fractions $\frac{p}{q} = [x_1, \ldots, x_m]$ and $\frac{p}{p-r} = [y_1, \ldots, y_k]$. We use

[1] Note that, for n-dimensional facets (interior points) the Delzant condition is trivial, and for $n-1$ dimensional facets (faces) the condition is automatically satisfied since they lie on hypersurfaces of the form $\{x \in \mathbb{R}^n \mid \sum_{i=1}^{n} \alpha_i x_i = b\}$. Therefore, checking the vertices of a 2-dimensional polygon is enough to verify whether it is Delzant or not.

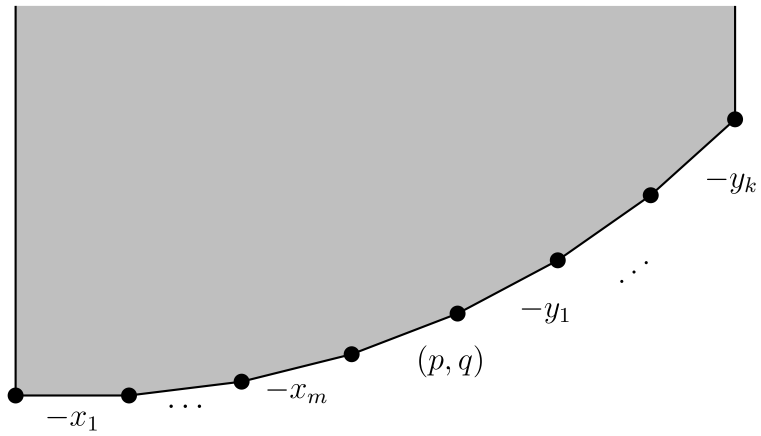

Figure 2.2: The minimal resolution of $\Pi(p, q)$. The vertices are marked for clarity. The labels $-x_i$, and $-y_j$ represent the self-intersection numbers of the symplectic spheres that live above the corresponding edges. The label (p, q) indicates the primitive integer direction of its edge.

the Hirzebruch-Jung convention for continued fractions, meaning:

$$[x_1, \ldots, x_m] := x_1 - \cfrac{1}{x_2 - \cfrac{1}{\ddots - \frac{1}{x_m}}}.$$

Recall that a zero continued fraction (ZCF) is one which evaluates to zero:

$$[a_1, \ldots, a_l] = 0.$$

Blowing up $[a_1, \ldots, a_l]$ corresponds to replacing it with any one of the following:

$$[1, a_1 + 1, \ldots, a_l],\ [a_1, \ldots, a_i + 1, 1, a_{i+1} + 1, \ldots, a_l],\ \text{or}\ [a_1, \ldots, a_l + 1, 1],$$

all of which are ZCFs themselves [17, Example 9.10]. The reverse operation is called *blow down*, and any ZCF admits a blow down to $[1, 1]$ [17, Lemma 9.11]. The following lemma collects some facts that we will use when constructing the compactification.

Lemma 2.1.3. *1. The continued fraction $\chi = [x_1, \ldots, x_m, 1, y_1, \ldots, y_k]$ is a ZCF. Then, since $x_i, y_i \geq 2$, there is a* unique *sequence of blow ups from the ZCF $[1, 1]$ to χ.*

2. *The symplectic sphere corresponding to the edge marked (p,q) in the polygon of Figure 2.2 has self intersection -1.*

Proof. (1). The fact that $\chi = 0$ follows from a simple computation using the relation $qr \equiv 1 \mod n$. Since

$$\chi = x_1 - \cfrac{1}{x_2 - \cfrac{1}{\ddots - \cfrac{1}{1-[y_1,\dots,y_k]^{-1}}}},$$

we calculate that

$$1 - [y_1,\dots,y_k]^{-1} = 1 - \frac{p-r}{p} = \frac{r}{p}.$$

Since the continued fraction of $\frac{p}{r}$ is just the reverse of that of $\frac{p}{q}$, we find that

$$\frac{1}{1-[y_1,\dots,y_k]^{-1}} = \frac{p}{r} = [x_m,\dots,x_1],$$

which we use to see

$$\chi = x_1 - \cfrac{1}{x_2 - \cfrac{1}{\ddots - \cfrac{1}{x_m-[x_m,\dots,x_1]}}}.$$

Since

$$x_m - [x_m,\dots,x_1] = \frac{1}{[x_{m-1},\dots,x_1]},$$

it follows that $\chi = 0$.

As is well-known, see for example [17, Lemma 9.11], every ZCF can be obtained from $[1,1]$ via iterated blow up, and the uniqueness follows from the observation that χ has a unique entry equal to 1. Therefore, the initial blow up must be $[1,1] \to [2,1,2]$, and subsequent blow ups must be adjacent to the unique 1 entry in the ZCF.[2]

(2). We prove this by direct computation. Consider the part of the polygon shown in Figure 2.3(a). Recall that each of the vertices were the result of resolving a singularity corresponding to the polygon $\pi(p,a)$ in Figure 2.3(b) for some $0 < a < p$. In particular, let $A \in SL_2(\mathbb{Z})$ be the unique matrix that maps the vertex of $\pi(p,a)$ to the $\frac{1}{p}(1,p-r)$ vertex in Figure 2.1. That is, A satisfies

$$(0,1)A = (-p,-q), \qquad \text{and} \qquad (0,1)A^{-1} = (p,a).$$

[2] For example, $[2,1,2] \to [3,1,2,2]$ is permissible, whilst $[2,1,2] \to [1,3,1,2]$ isn't.

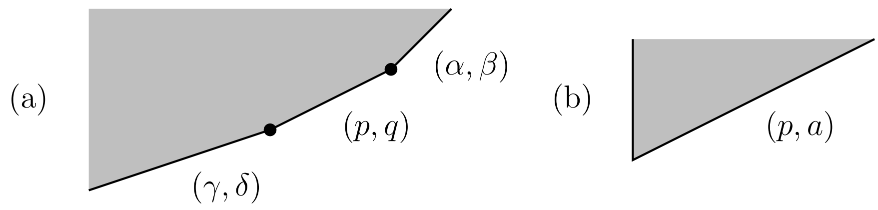

Figure 2.3: (a) A neighbourhood of the edge (p, q) in the polygon of Figure 2.2. (b) the moment polygon $\pi(p, a)$.

Then, since $(\alpha, \beta) = (1, 0)A$ we have that[3] $0 < \alpha = a < p$. A similar argument shows that $0 < \gamma < p$.

Both vertices in Figure 2.3(a) are Delzant, which means

$$\begin{aligned} -\delta p + \gamma q &= \det \begin{pmatrix} p & q \\ -\gamma & -\delta \end{pmatrix} = 1, \text{ and} \\ -\alpha q + \beta p &= \det \begin{pmatrix} \alpha & \beta \\ -p & -q \end{pmatrix} = 1. \end{aligned} \tag{2.1.1}$$

Reducing modulo p yields $\gamma \equiv -\alpha \mod p$, which, combined with the fact that $0 < \alpha, \gamma < p$, means $\gamma = p - \alpha$. Lemma 3.20 of [17], asserts that the self intersection of the (p, q)-edge is given by

$$\det \begin{pmatrix} -\gamma & -\delta \\ \alpha & \beta \end{pmatrix} = -\gamma\beta + \delta\alpha.$$

Therefore, in view of equation (2.1.1), we have that

$$-\gamma\beta + \delta\alpha = \frac{-\gamma - \alpha}{p} = -1,$$

which completes the proof. □

[3] Indeed,

$$A^{-1} = \begin{pmatrix} * & * \\ p & a \end{pmatrix}$$

so that

$$\alpha = (1, 0)A \begin{pmatrix} 1 \\ 0 \end{pmatrix} = (1, 0) \begin{pmatrix} a & * \\ -p & * \end{pmatrix} \begin{pmatrix} 1 \\ 0 \end{pmatrix} = a.$$

Remark 2.1.4. A more conceptual way to prove statement (2) above is to use the fact that there is a unique series of blow ups

$$[1,1] \to [x_1, \dots, x_m, 1, y_1, \dots, y_k]$$

and perform the corresponding symplectic blow ups as in Figure 2.4. Indeed, Example 9.10 and Corollary 9.13 of [17] show that the combinatorics of these two procedures are the same. Thus, we obtain two families of toric moment polygons,

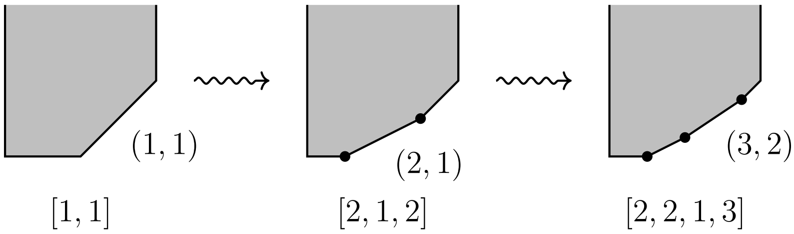

Figure 2.4: An example of blowing up the toric manifold associated to the ZCF $[1,1]$. The symplectic cuts made are labelled with the primitive integer direction they point in. The vertices are marked for clarity.

one from resolving all the polygons $\Pi(p,q)$, and the other from blowing up ZCFs. Note in particular that the second family of polygons will necessarily have a unique edge with self intersection -1. An inductive argument on the length of the blow-ups of the ZCFs shows that these families coincide, from which it follows that the (p,q) edge in Figure 2.2 must be a -1-curve, since all the other edges are of self intersection at most -2.

2.2 Compactification construction

The Fubini-Study form ω_{FS} on $\mathbb{CP}^2$ can be defined as the symplectic reduction of $(\mathbb{C}^3, \omega_{\mathbb{C}^3})$ with respect to the Hamiltonian $\frac{1}{2}|z|^2$ at the regular level $(\frac{1}{2}|z|^2 = 1)$ [17, Example 4.9].[4] It has a Hamiltonian torus action with moment polygon

[4]Equivalently, in each of the usual affine coordinate patches on $\mathbb{CP}^2$ we have

$$\omega_{\text{FS}} = i\partial\bar{\partial}\log(1 + |z_1|^2 + |z_2|^2).$$

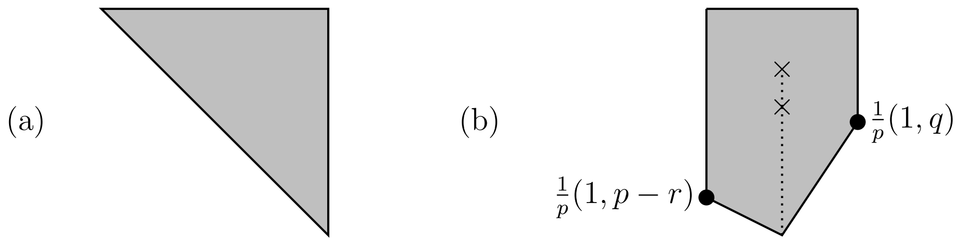

Figure 2.5: (a) The moment polygon of $(\mathbb{CP}^2, \omega_{\mathrm{FS}})$. (b) The almost toric base diagram of the compactification X (modulo resolving the marked singularities as in Figure 2.2).

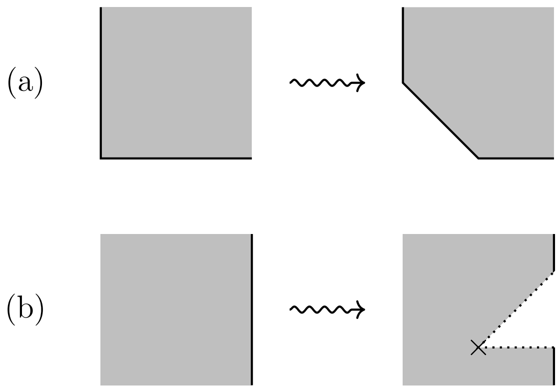

Figure 2.6: (a) A symplectic cut at a Delzant vertex. (b) A non-toric blow up at an edge of a moment polygon.

Figure 2.5(a). The next result shows that we can perform a series of alterations to this polygon, called symplectic cuts and non-toric blow ups, to obtain a symplectic compactification of $B_{d,p,q}$ determined up to symplectic deformation[5] by the polygon Figure 2.5(b). Performing a symplectic cut to a Delzant vertex in a moment polygon is characterised by Figure 2.6(a). The non-toric blow up transforms a polygon as in Figure 2.6(b). Both of these operations correspond to the symplectic blow up. The reader may consult Sections 4.4 and 9.1 of [17] for further details.

This definition ensures that $\int_H \omega_{\mathrm{FS}} = 2\pi$ for the class H of a projective line.

[5] The reason this is not "up to symplectomorphism" is because of the freedom of choice over the affine lengths of the edges introduced by resolving the singularities. However, this is largely inconsequential.

Lemma 2.2.1. *There exists a number*[6] $\mu > 0$ *and a symplectic compactification* $(X, \omega) = (X_{d,p,q}, \omega_{d,p,q})$ *of* $B_{d,p,q}$ *such that* (X, ω) *is a blow up of* $(\mathbb{CP}^2, \mu\omega_{\mathrm{FS}})$*. The blow ups compute a basis* $\{H, S, E_i, \mathcal{E}_j \mid 0 \le i \le n, 1 \le j \le d\}$ *of* $H_2(X)$ *such that the following holds:*

1. H *is the class of a line in* $\mathbb{CP}^2$*,* S *and* E_i *are exceptional curves corresponding to symplectic cuts, and* $\mathcal{E}_j$ *are exceptional curves corresponding to non-toric blow ups;*

2. $n = m + k - 1$ *is equal to the number of blow ups*

$$[1, 1] \to [x_1, \ldots, x_m, 1, y_1, \ldots, y_k]$$

determined by Lemma 2.1.3;

3. *the* ω*-areas of the non-toric curves* $\mathcal{E}_j$ *are all equal, that is, there exists* $l > 0$ *such that, for each* $1 \le j \le d$*,*

$$\omega(\mathcal{E}_j) = l;$$

and,

4. *define the* homology class of a smooth fibre *by* $F := H - S$*, then* $\omega(F) = 2l$*, and* $\omega(F) > \omega(E_n)$*.*

Proof. Consider the almost toric base diagram of $B_{d,p,q}$ shown in Figure 2.7(a).[7] Choose $A \in SL_2(\mathbb{Z})$ such that $(0, 1)A = (p, q)$. Writing $A = \begin{pmatrix} r & s \\ p & q \end{pmatrix}$, the condition $\det A = 1$ ensures that $qr \equiv 1 \mod p$ and we may choose A such that $0 < r < p$. Applying A^{-1} to the base diagram then yields Figure 2.7(b). We make a series of symplectic cuts, shown as dash dot lines in Figure 2.8: first a horizontal one at

[6] The actual value of μ is irrelevant, but one can calculate it to be $\sqrt{2}$ times the Euclidean distance of the diagonal edge from the top-right vertex in the polygon shown in Figure 2.5(a).

[7] Note that the figures drawn here are for the case $(d, p, q) = (2, 2, 1)$ but the process works for all triples.

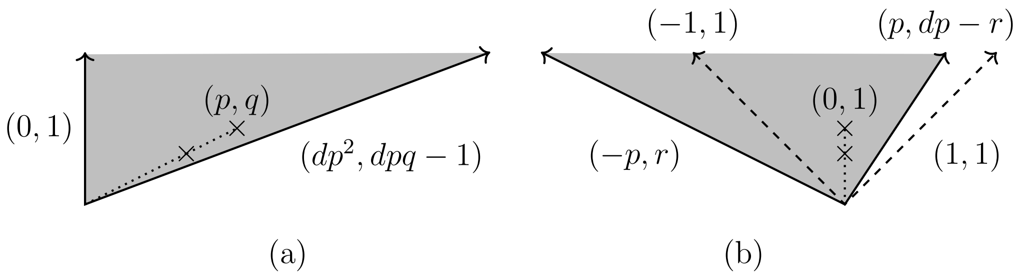

Figure 2.7: The figure from left to right shows the effect of the transformation A^{-1}. The dashed rays in (b) are included to show the extremes that the solid rays cannot cross since $0 < r < p$ and $d > 1$.

some level above the focus-focus critical values; and then vertical cuts on either side. Doing so introduces two cyclic quotient singularities, which are marked in the figure. Taking the minimal resolution yields the desired compactification X.

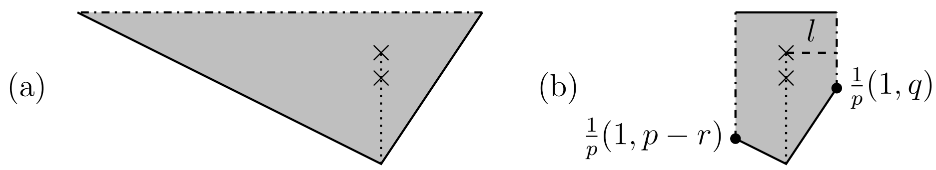

Figure 2.8: Performing symplectic cuts (dash dot lines) to compactify $B_{d,p,q}$. After performing the vertical cuts, we obtain a manifold with two cyclic quotient singularities, which are marked as bullets in the figure, along with their type. The dashed line labelled l indicates the affine displacement of the right hand vertical cut from the monodromy eigenline.

We now perform a sequence of blow-downs to the base diagram of X terminating at the moment polygon of $\mathbb{CP}^2$ (Figure 2.10(b)). This proves that X is a blow-up of $\mathbb{CP}^2$ and the sequence of blow-downs we perform picks out the claimed basis of $H_2(X)$. To this end, we rotate the branch cut by 90° anticlockwise (see [17, Section 7.2]), to obtain[8] Figure 2.9(a). The d bites in the diagram correspond to the exceptional loci of d non-toric blow ups. In particular, they represent a collection

[8]Note that we continue to draw the base diagram of the singular manifold (Figure 2.8(b)) prior to taking the minimal resolution. This is to make the pictures easier to draw and interpret.

of d disjoint symplectic -1-spheres. Moreover, each of them has ω-area equal to l, since the focus-focus critical values are all the same affine length l from the right-hand vertical edge. Label the homology classes of these spheres by $\mathcal{E}_j$ and blow them down to obtain Figure 2.9(b). Observe that the singularities of the resultant

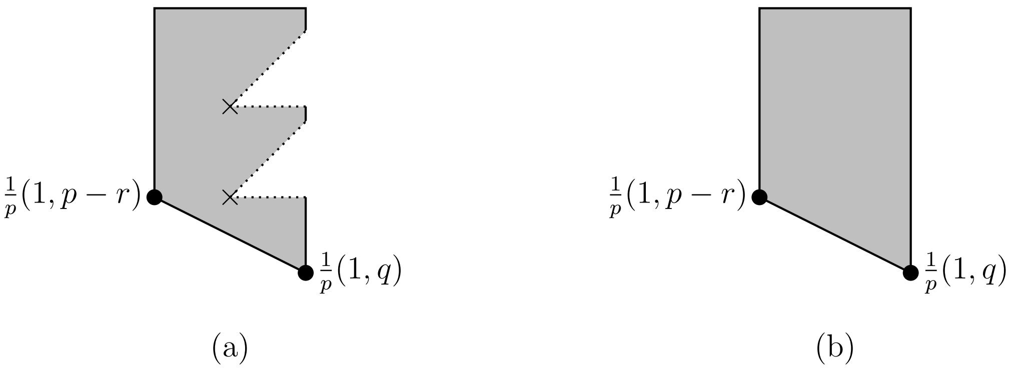

Figure 2.9: Non-toric blow downs.

manifold are modelled on a reflection about the vertical axis of those of[9] $\Pi(p, r)$. Therefore, the minimal resolution is given by (the reflection of) Figure 2.2. Then Lemma 2.1.3 implies that there is a unique sequence of toric blow downs to the polygon in Figure 2.10(a), which is the moment polygon of $\mathbb{CP}^2\#2\overline{\mathbb{CP}^2}$. Reversing this process, that is blowing up instead of down, yields the desired sequence of blow ups from $(\mathbb{CP}^2, \mu\omega_{\text{FS}})$ to (X, ω). Labelling the homology classes of the exceptional loci of the toric blow ups as S and E_i, and the non-toric ones as $\mathcal{E}_j$, we obtain the claimed basis

$$\{H, S, E_i, \mathcal{E}_j \mid 0 \leq i \leq n, 1 \leq j \leq d\}.$$

We have established properties (1)–(3), so it remains to show (4). The equality $\omega(H - S) = 2l$ can be seen by examining Figure 2.8(b): $\omega(H - S)$ is equal to (2π times) the affine length of the horizontal edge and the vertical cuts can be made equidistant from the monodromy eigenline. The inequality $\omega(H - S) > \omega(E_n)$

Drawing the resolved manifold would involve long chains of edges as in Figure 2.2, the precise directions and lengths of which add nothing to the proof, so we omit them.

[9]This is not a typo; reflection in the vertical acts as inverse modulo p on the singularity: $\frac{1}{p}(1, q) \mapsto \frac{1}{p}(1, r)$.

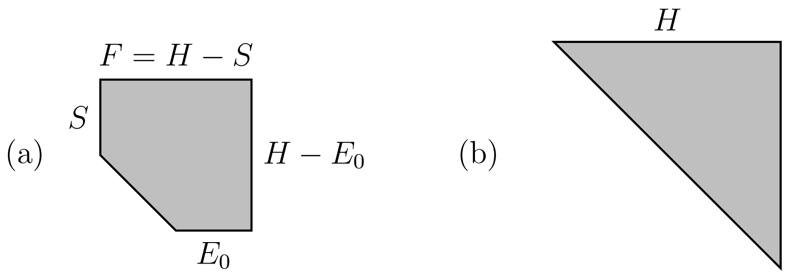

Figure 2.10: (a) The moment polygon of the toric manifold obtained by blowing down $X_{d,p,q}$. The edges are labelled with the corresponding homology classes. (b) The moment polygon of $\mathbb{CP}^2$, obtained by blowing down the edges labelled S and E_0.

follows from considering Figure 2.11, which is the same polygon as in Figure 2.9(b).

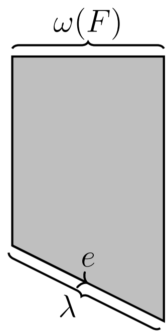

Figure 2.11: The moment polygon of Figure 2.9(b) with affine lengths indicated. The edge e points in the $(p, -r)$ direction. Since the edge corresponding to E_n is obtained by symplectically cutting the bottom edge, we have the inequality $\omega(E_n) < \lambda$.

The edge corresponding to E_n is obtained by symplectic cutting the bottom edge e of Figure 2.11, so its affine length $\omega(E_n)$ is strictly smaller than that of e, which we denote by λ. Since e points in the direction $(p, -r)$, we must have that $\lambda p = \omega(H - S)$, since the top edge is horizontal. As $p > 1$, we obtain that $\lambda < \omega(H - S)$. This completes the proof. □

Remark 2.2.2. The toric boundary $D \subset X$ is a *symplectic divisor* of (X, ω), meaning that is a union of (real) codimension 2 symplectic submanifolds of (X, ω). It is a cycle of transversely intersecting spheres, with each intersection point being positive. We represent D by the dual intersection graph of its components in Figure 2.12(a). Removing the components $\{S, F, S' := H - E_0 - \sum_{j=1}^{d} \mathcal{E}_j\}$ yields a subgraph

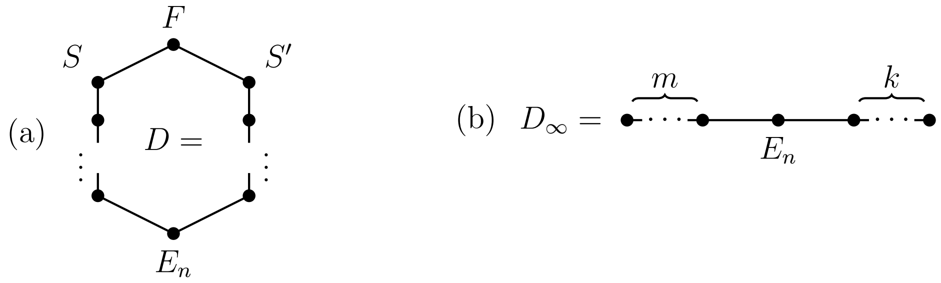

Figure 2.12: (a) The dual intersection graph of the divisor $D \subset X$. Vertices are labelled with the respective homology classes of the component symplectic spheres. The integers m and k correspond to the lengths of the continued fractions of $\frac{p}{p-r}$ and $\frac{p}{q}$ respectively, as in Lemma 2.1.3. (b) The subgraph consisting of spheres introduced by resolving the singularities in the compactification process.

representing the part of the toric boundary obtained by resolving the singularities in Figure 2.9, which we will call D_∞.

2.3 Lagrangian spheres in $B_{d,p,q}$

We classify the homology classes in $H_2(B_{d,p,q}; \mathbb{Z}) \cong \mathbb{Z}^{d-1}$ that support Lagrangian spheres. The point of this is to show that, in the compactification X, any Lagrangian sphere $L \subset B_{d,p,q} \subset X$ is homologous to $e - e'$ where $e, e' \in H_2(X; \mathbb{Z})$ are classes represented by J-holomorphic -1-curves. This will be important for our neck stretching analysis performed in Section 3.

The homotopy type of $B_{d,p,q}$ is a (p, q)-pinwheel wedged with $d - 1$ spheres [17, Lemma 7.11]. Moreover, the standard A_{d-1} configuration of Lagrangian spheres is a generating set of the second homology:

$$H_2(B_{d,p,q}) = H_2(B_{d,p,q}; \mathbb{Z}) = \mathbb{Z}\langle L_i : 1 \leq i \leq d-1 \rangle,$$

Therefore, the intersection form with respect to this basis is

$$Q := \begin{pmatrix} -2 & 1 & 0 & \cdots & \cdots & \cdots & 0 \\ 1 & -2 & 1 & 0 & \cdots & \cdots & 0 \\ 0 & 1 & -2 & 1 & \cdots & \cdots & 0 \\ \vdots & \vdots & \ddots & \ddots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & \cdots & 1 & -2 & 1 \\ 0 & 0 & \cdots & \cdots & 0 & 1 & -2 \end{pmatrix}. \tag{2.3.1}$$

Lemma 2.3.1. *The set $\mathcal{L} \subset H_2(B_{d,p,q})$ of homology classes represented by Lagrangian spheres in $B_{d,p,q}$ is given by*

$$\mathcal{L} = \{\pm(L_i + L_{i+1} + \ldots + L_j) \mid 1 \leq i \leq j \leq d-1\}. \tag{2.3.2}$$

Proof. Put $\mathcal{L}' = \{\pm(L_i + L_{i+1} + \ldots + L_j) \mid 1 \leq i \leq j \leq d-1\}$. Recall that the Picard-Lefschetz formula states that, for the 4-dimensional Dehn twist τ_L associated to a Lagrangian 2-sphere L, the action of τ_L on homology is given by

$$(\tau_L)_*(A) = \begin{cases} A + (A \cdot L)L, & \text{if } A \in H_2, \\ A, & \text{if } A \in H_k,\ k \neq 2, \end{cases}$$

see [48] for example. Since τ_L is a symplectomorphism, the image of any Lagrangian submanifold under τ_L is again Lagrangian. Therefore, by iterating the calculation

$$(\tau_{L_i})_* L_{i+1} = L_{i+1} + (L_i \cdot L_{i+1})L_i = L_{i+1} + L_i,$$

we find that $\mathcal{L}' \subset \mathcal{L}$.

On the other hand, Weinstein's Lagrangian tubular neighbourhood theorem [51] implies that any Lagrangian 2-sphere has self-intersection -2, which reduces the reverse inclusion $\mathcal{L} \subset \mathcal{L}'$ to a matter of algebra. Indeed, consider the homology class $L = \sum_{i=1}^{d-1} a_i L_i$ of a Lagrangian sphere in $B_{d,p,q}$. Then, using Equation (2.3.1), we have that

$$-2 = L^2 = -2\sum_{i=1}^{d-1} a_i^2 + 2\sum_{i=1}^{d-2} a_i a_{i+1}.$$

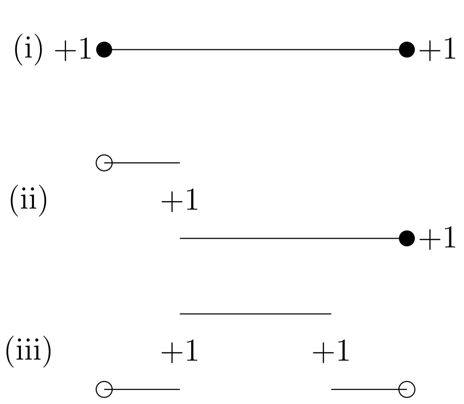

(i) $a_1 = \pm 1,\ a_{d-1} = \pm 1$
$\pm(L_1 + \ldots + L_{d-1})$

(ii) $a_i - a_{i+1} = 1,\ a_{d-1} = -1$
$-L_{i+1} - \ldots - L_{d-1}$

(iii) $a_{i-1} - a_i = -1,\ a_j - a_{j+1} = 1$
$L_i + \ldots + L_j$

Figure 2.13: These diagrams demonstrate how the coefficients a_i contribute to the sum in Equation (2.3.3). The left circle represents the value of a_1, and the right that of a_{d-1}; an open circle is 0, whilst a filled one is ± 1. Adjacent coefficients (a_i and a_{i+1}) with the same value are joined with a straight line. As indicated, a contribution of $+1$ is made to the sum any time the horizontal line jumps or a circle is filled. The labels to the right of each diagram detail the non-zero terms of the sum and the homology class they represent.

Since $(a_i - a_{i+1})^2 = a_i^2 - 2a_ia_{i+1} + a_{i+1}^2$, we can rearrange the above to obtain

$$a_1^2 + \sum_{i=1}^{d-2}(a_i - a_{i+1})^2 + a_{d-1}^2 = 2. \tag{2.3.3}$$

The left-hand side of this equation is a sum of square integers, so it follows that exactly two of the terms a_1, $(a_i - a_{i+1})$, or a_{d-1} have modulus equal to 1, whilst the remainder are zero. From this we deduce that $\mathcal{L} \subset \mathcal{L}'$ by a case-by-case analysis (although the reader may find Figure 2.13 alone sufficiently convincing).

We must have that $|a_1| \leq 1$ and that, for some $k \geq 0$, the first k terms of the sum $\mathcal{S} := \sum_{i=1}^{d-2}(a_i - a_{i+1})^2$ are zero, implying that $a_1 = \ldots = a_{k+1}$. There are three cases (which Figure 2.13 exemplifies): either (i) $k = d-2$ which implies that $a_1 = \ldots = a_{d-1} = \pm 1$, or $k < d-2$ and $\mathcal{S}$ has either (ii) one, or (iii) two non-zero terms. In the case (ii), there exists $i \geq 1$ such that

$$a_1 a_{d-1} = 0, \qquad a_1 = \ldots = a_i,$$
$$|a_1 + a_{d-1}| = 1, \text{ and } a_{i+1} = \ldots = a_{d-1}.$$

On the other hand, for (iii) there exist $1 < i \leq j < d-1$ such that

$$a_1 = \ldots = a_{i-1} = 0 = a_{j+1} = \ldots = a_{d-1}, \text{ and } a_i = \ldots = a_j = \pm 1.$$

These equations imply that $L \in \mathcal{L}'$ and hence $\mathcal{L} = \mathcal{L}'$. $\square$

Remark 2.3.2. The $\mathbb{Z}$-linear map $\phi : H_2(B_{d,p,q}) \to \mathbb{R}^d$ generated by $\phi(L_i) = e_i - e_{i+1}$, where $e_1, \ldots, e_d$ is the standard Euclidean basis of $\mathbb{R}^d$, identifies $H_2(B_{d,p,q})$ with the A_{d-1} lattice in $\mathbb{R}^d$. In this context, the calculation of the homology classes of Lagrangian spheres above is equivalent to computing the roots of A_{d-1}.

Lemma 2.3.3. *The construction of $X = X_{d,p,q}$ induces an injection $H_2(B_{d,p,q}) \hookrightarrow H_2(X)$ sending the generators L_i of $H_2(B_{d,p,q})$ to $\mathcal{E}_i - \mathcal{E}_{i+1} \in H_2(X)$ (up to sign).*

Proof. Zooming in to Figure 2.6(b), one can see that each sphere L_i has a neighbourhood $N_i \subset X_{d,p,q}$ that looks like that shown in Figure 2.14. The second

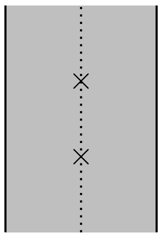

Figure 2.14: A neighbourhood of one of the standard Lagrangian spheres $L_i \subset X$.

homology of N_i is freely generated by $F, \mathcal{E}_i, \mathcal{E}_{i+1} \in H_2(X)$, so we must have

$$L_i = \lambda F + \mu_i \mathcal{E}_i + \mu_{i+1} \mathcal{E}_{i+1}.$$

Note that the left edge of the diagram in Figure 2.14 is part of the S sphere, which is a non-zero class S_i in $H_2(N_i, \partial N_i)$ satisfying

$$F \cdot S_i = 1, \qquad \text{and} \qquad \mathcal{E}_j \cdot S_i = 0,$$

for $j = i, i+1$. Therefore, since L_i is disjoint from S_i, we find that $\lambda = 0$. Now, L_i must be a -2 class, and so

$$-2 = L_i^2 = -\mu_i^2 - \mu_{i+1}^2,$$

which implies that $|\mu_i| = 1 = |\mu_{i+1}|$. Finally, the Lagrangian condition ensures that

$$0 = \omega(L_i) = \mu_i \omega(\mathcal{E}_i) + \mu_{i+1}\omega(\mathcal{E}_{i+1}) = l(\mu_i + \mu_{i+1}),$$

which shows that the coefficients μ_i and μ_{i+1} must have opposite sign. Hence, the result follows. $\square$

Combining the above with Lemma 2.3.1, we have proved the following:

Corollary 2.3.4. *Any Lagrangian sphere $L \subset B_{d,p,q} \subset X$ has homology class of the form*

$$[L] = \mathcal{E}_i - \mathcal{E}_j,$$

for some $1 \leq i \neq j \leq d$.

2.4 Moduli spaces of fibre curves

In this section, we examine the J-holomorphic curves of X that represent the homology class $F = H - S$. First, we define the central object of study. Given an almost complex structure J on X and a complex structure j on S^2, define the *moduli space of genus 0 curves in the homology class F*, $\mathcal{M}_{0,0}(X, F; J)$, to be the set[10] of maps $u : (S^2, j) \to X$ satisfying

$$\mathrm{d}u \circ j = J \circ \mathrm{d}u \qquad \text{and} \qquad [u] := u_*[S^2] = F,$$

modulo the action of the reparametrisation group $\mathrm{Aut}(S^2, j) \cong PSL_2(\mathbb{C})$. The present goal is to prove that (after removing a single *exotic* stable curve) the forgetful map

$$\overline{\mathcal{M}}_{0,1}(X, F; J) \to \overline{\mathcal{M}}_{0,0}(X, F; J) \tag{2.4.1}$$

gives rise to a Lefschetz fibration for all almost complex structures J in some suitable subset $\mathcal{J}(\Delta)$ of the space of compatible structures $\mathcal{J}(X, \omega)$, defined as follows.

[10]The notation used here is adopted from [37]. We'll often suppress X in the notation $\mathcal{M}_{0,0}(X, F; J)$.

Definition 2.4.1. Let Δ be a symplectic divisor in (X,ω), and define

$$\mathcal{J}(\Delta) := \left\{ J \in \mathcal{J}(X,\omega) \;\middle|\; \begin{array}{l} \text{each irreducible component of } \Delta \text{ ad-} \\ \text{mits a } J\text{-holomorphic representative} \end{array} \right\}.$$

The reason we're interested in this space of almost complex structures is that the Lagrangian spheres we consider live in the complement of a divisor in X. Under neck stretching, the almost complex structure only changes in a small neighbourhood of the Lagrangian, and so, supposing the initial almost complex structure J_0 is a member of $\mathcal{J}(\Delta)$, then J_t will also be for all $t \geq 0$.

Consider the divisor D' obtained by excising the single component corresponding to the edge of the toric boundary of X with homology class E_n. That is, D' is given by the subgraph of the dual intersection graph of D (Figure 2.12(a)) consisting of every vertex except the bottom one. We are primarily concerned with the case where $\Delta = D'$ in the above definition. We will show that, for each $J \in \mathcal{J}(D')$, there exists a unique J-holomorphic curve in the class E_n. Note that this does not necessarily mean that $J \in \mathcal{J}(D)$, since the aforementioned curve may not have the same image as the E_n-component[11] of D. Indeed, *a priori* the Lagrangian sphere L may intersect the E_n-component of D, whereas Corollary 3.4.11 shows that, for a sufficiently long neck stretch J_t, the unique J_t-holomorphic curve of class E_n is disjoint from L.

The existence of the J-holomorphic curve in the class E_n will then be used to show that, for all $J \in \mathcal{J}(D')$, the curves in the Gromov compactification $\overline{\mathcal{M}}_{0,0}(F;J)$ come in the following types:

- smooth curves of class F, called a *smooth fibre curves*;
- nodal curves with exactly 2 components with homology classes $\mathcal{E}_j, F - \mathcal{E}_j$, which will correspond to the Lefschetz critical fibres; and
- a single *exotic* nodal curve $\mathbf{u}^J_\infty$ with $n+2$ genus 0 components (where n is

[11] Although the reader will lose little by ignoring this fact.

as in Lemma 2.2.1), all, except one (the component covering the unique J-holomorphic curve in the class E_n), of which cover those of $D_\infty \cap D'$.

Moreover, the unique J-holomorphic curves in the classes S and S' are sections of the fibration (2.4.1) — see Corollary 2.4.16 and Proposition 2.5.1 for further details.

2.4.1 Almost complex structures on symplectic divisors

That $\mathcal{J}(D') \neq \emptyset$, follows from straightforward extensions of results on the symplectic neighbourhood theorem. See Appendix A.1 for the detailed construction. To apply the results therein, we need to check that the components of D' intersect symplectically orthogonally. However, this follows from standard facts on toric geometry, see [17, §3.2] for example. Indeed, each intersection point between components of D is modelled on a Delzant corner of the moment polytope of X, which is fibred symplectomorphic to $(\mathbb{C}^2, \omega_{\mathbb{C}^2})$ with its usual toric structure. The components of D' correspond to the coordinate planes in $\mathbb{C}^2$, which intersect symplectically orthogonally. See Figure 2.15 for a cartoon picture of what happens.

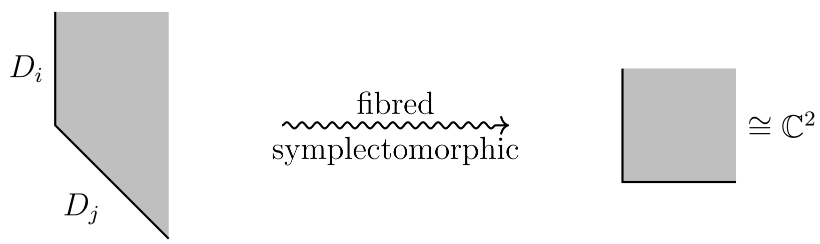

Figure 2.15: Intersecting components of the divisor D' do so symplectically orthogonally, since any Delzant corner of a moment polytope is integral affine isomorphic to the standard one. The fibred symplectomorphism sends the pieces D_i and D_j to the coordinate planes in $\mathbb{C}^2$.

2.4.2 Non-bubbling under variations of J

Choose $J \in \mathcal{J}(D')$. In this section we show that the class E_n admits a J-holomorphic representative. To this end, recall that the intersection pairing on $H_2(X) = H_2(X;\mathbb{Z})$ is non-degenerate. Given any subset $\mathcal{A} \subset H_2(X)$ we denote the orthogonal complement (with respect to the intersection pairing) of the submodule $\mathbb{Z}\langle\mathcal{A}\rangle$ it generates by $\mathcal{A}^\perp$. By non-degeneracy, $\dim \mathcal{A}^\perp = \operatorname{codim} \mathbb{Z}\langle\mathcal{A}\rangle$. This first result is a basic application of positivity of intersections (Theorem 1.3.7).

Lemma 2.4.2. *Let $A \in H_2(X)$ be any homology class represented by a nodal J-holomorphic sphere, that is, there exists $\mathcal{A} \subset H_2(X)$ such that $A = \sum_{A_\alpha \in \mathcal{A}} A_\alpha$ where each $A_\alpha \in \mathcal{A}$ has a J-holomorphic representative. Let $C \subset D$ be a connected subgraph*[12] *of D and suppose that $A \cdot C_i = 0$ for each J-holomorphic component C_i of C, and that, for some i, there exists $k_i \in \mathbb{N}$ such that $k_i[C_i] \in \mathcal{A}$. Then, for all i, there exist positive integers k_i such that $k_i[C_i] \in \mathcal{A}$.*

Proof. We induct on the number, m, of irreducible components of C. The base case $m = 1$ is trivial, so assume the result holds for some $m > 0$. In other words, $C = \bigcup_{1\le i\le m+1}(C_i)$ and, for $1 \le i \le m$, there exist positive integers $k_i > 0$ such that $k_i[C_i] \in \mathcal{A}$. Note that, by construction, C_{m+1} has at most two neighbours in C, each of which it intersects exactly once positively. Therefore, we have that

$$C_{m+1} \cdot \sum_{i=1}^{m} k_i C_i > 0.$$

Combining this with the assumption that $A \cdot C_{m+1} = 0$, we obtain

$$\left(A - \sum_{i=1}^{m} k_i[C_i]\right) \cdot C_{m+1} < 0.$$

Positivity of intersections then implies that one of the remaining components $A_\alpha \in \mathcal{A}\backslash\{k_i[C_i] \mid 1 \le i \le m\}$ covers C_{m+1}. Thus, there exists $k_{m+1} > 0$ such that $k_{m+1}[C_{m+1}] \in \mathcal{A}$. □

[12] In the sense of the dual intersection graph of D.

Lemma 2.4.3. *For any $J \in \mathcal{J}(D')$, any stable curve of class E_n (if it exists) is smooth: $\overline{\mathcal{M}}_{0,0}(E_n; J) = \mathcal{M}_{0,0}(E_n; J)$.*

Proof. Consider the connected subgraph $C \subset D$ consisting of D minus E_n and its neighbours — that is, C consists of all the vertices except the bottom three in Figure 2.12. Said another way, C is a the maximal subgraph none of whose components intersect E_n. Denote the set of homology classes of the components of C as $\mathcal{C}$.

Suppose that $\mathbf{u} = (u_\alpha) \in \overline{\mathcal{M}}_{0,0}(E_n; J)$ is a stable J-holomorphic curve in the class E_n, and that, for some components u_α of $\mathbf{u}$ and C_i of C, we have $A_\alpha := [u_\alpha] = k_i[C_i]$. Since $F \in \mathcal{C}$, and, for all $c \in \mathcal{C}$, $E_n \cdot c = 0$, we may apply Lemma 2.4.2 to deduce that $A_\beta = k_\beta F$ for some component u_β of $\mathbf{u}$, which contradicts $\omega(F) > \omega(E_n) \geq \omega(A_\beta)$. Therefore, for all α, $c \in \mathcal{C}$, and $k \in \mathbb{N}$, we have $A_\alpha \neq kc$. Then, positivity of intersections and $E_n \in \mathcal{C}^\perp$ imply that $A_\alpha \in \mathcal{C}^\perp$ for all α. Indeed, $0 = E_n \cdot c = \sum_\alpha A_\alpha \cdot c$ and $A_\alpha \cdot c \geq 0$.

A simple calculation shows that $\mathcal{C}$ is a rank $n+2$ linearly independent set, and so, since $\dim H_2(X) = n + d + 3$, to explicitly describe A_α it suffices to find a rank $d+1$ linearly independent set in $\mathcal{C}^\perp$. Let E^-, E^+ denote the neighbours of $E_n \subset D$ and note that $E^+ - E^- \in \mathcal{C}^\perp$, since they both intersect E_n once positively. Then it is easily checked that

$$\mathcal{C}^\perp = \mathbb{Z}\langle E_n, E^+ - E^-, \mathcal{E}_j - \mathcal{E}_{j+1} : 1 \leq j < d \rangle.$$

Recall that during the compactification process of Lemma 2.2.1, the curves $E^\pm$ were introduced through symplectic cutting the toric boundary during the resolution of singularities procedure. Since we have choice over the affine lengths of these cuts, they can be made equal. That is, $\omega(E^+) = \omega(E^-)$. Furthermore, by Lemma 2.2.1, $\omega(\mathcal{E}_j) = l$ does not depend on j, so we find that, for some integers $\lambda_n, \lambda_\pm, \mu_j \in \mathbb{Z}$,

$$\omega(A_\alpha) = \omega\left(\lambda_n E_n + \lambda_\pm(E^+ - E^-) + \sum_{j=1}^{d-1} \mu_j(\mathcal{E}_j - \mathcal{E}_{j+1})\right) = \lambda_n \omega(E_n).$$

That is, the ω-area of u_α is an integer multiple of $\omega(E_n)$. Since

$$\sum_\alpha \omega(A_\alpha) = \omega(E_n),$$

this is only possible if it is in fact equal to $\omega(E_n)$. Since non-constant J-holomorphic curves have positive area, this implies that $A_\alpha = E_n$ and hence, $\mathbf{u} = u_\alpha$ is a smooth curve. □

Corollary 2.4.4. *For any $J \in \mathcal{J}(D')$, the class E_n admits a J-holomorphic representative.*

Proof. Let $\mathcal{J}(D', E_n)$ denote the subset of $\mathcal{J}(D')$ of almost complex structures that admit a J-holomorphic representative of E_n. Section 2.4.1 says that $\mathcal{J}(D) \subset \mathcal{J}(D', E_n)$ is non-empty, and so, since $c_1(E_n) > 0$, we may apply automatic transversality [25] to show that $\mathcal{J}(D', E_n)$ is open in $\mathcal{J}(D')$. Moreover, by Lemma 2.4.3, $\mathcal{J}(D', E_n)$ is also a closed subset. Indeed, pick a sequence $J_\nu \in \mathcal{J}(D', E_n)$ converging to $J \in \mathcal{J}(D')$ and a corresponding sequence of curves $u_\nu \in \mathcal{M}_{0,0}(E_n; J_\nu)$. Gromov compactness [37, Theorem 5.3.1] ensures that there is a convergent subsequence $u_\nu \to \mathbf{u}$, for some $\mathbf{u} \in \overline{\mathcal{M}}_{0,0}(E_n; J)$. However, Lemma 2.4.3 states that $\mathbf{u}$ must actually be a smooth curve, and so we find that $J \in \mathcal{J}(D', E_n)$. Thus, since $\mathcal{J}(D')$ is connected, we must have that $\mathcal{J}(D', E_n) = \mathcal{J}(D')$. □

Remark 2.4.5. Corollary 2.4.4 is a type of *non-bubbling result*. We can paraphrase it by saying that "J-holomorphic curves in the class E_n do not bubble under deformations of $J \in \mathcal{J}(D')$."

2.4.3 The universal J-holomorphic curve

For $J \in \mathcal{J}(D')$, let E_J be the image of the unique J-holomorphic curve in the homology class E_n. Recall that the symplectic divisor $D_\infty \subset D$ is that arising from the resolution of singularities carried out in Lemma 2.2.1. By replacing the E_n-component of D_∞ with E_J we obtain a J-holomorphic divisor $D_{\infty,J}$. More precisely this is $(D_\infty \cap D') \cup E_J$. See Figure 2.16 for a sketch of how D_∞ and $D_{\infty,J}$ differ in

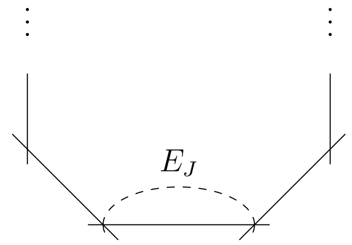

Figure 2.16: Part of the divisors D_∞ and $D_{\infty,J}$. The only difference between the two is the central component: $D_{\infty,J}$ contains the unique J-holomorphic curve E_J in the class E_n, whereas the E_n-component of D_∞ is part of the toric boundary divisor of X.

a neighbourhood of E_J. Excising this divisor we obtain $X_J := X \backslash D_{\infty,J}$, which we will show in Section 2.5 is the open dense subset of X on which there exists a genus 0 J-holomorphic Lefschetz fibration.

Following Wendl [55, §7.3.3], we define the *universal J-holomorphic curve* to be the forgetful map

$$\pi_J : \overline{\mathcal{M}}_{0,1}(X_J, F; J) \to \overline{\mathcal{M}}_{0,0}(X_J, F; J).$$

Our goal is to prove that these moduli spaces can be endowed with smooth structures that realise π_J as a smooth Lefschetz fibration. The only notable difference between our situation and that in [55] is that the target manifold X_J is not closed. However, due to the specifics of our situation, we can work around this and show that the required results of [55, §7.3.4] hold here.

We prove the following results first under a mild genericity condition $J \in \mathcal{J}_{\mathrm{reg}}(D')$ explained in Definition 2.4.9 below, and then extend them to all $J \in \mathcal{J}(D')$ by an automatic transversality and non-bubbling argument as in Corollary 2.4.4.

Remark 2.4.6. The reason we invoke genericity in the first place is that it gives us a convenient argument to deduce the structure of $\overline{\mathcal{M}}_{0,0}(X_J, F; J)$. One may wonder why we don't just assume genericity henceforth and skip the automatic transversality argument. The reason is that in Chapter 3 we will have to modify our almost complex structures a few times, and knowing the results of this section apply to all $J \in \mathcal{J}(D')$ makes those arguments more straightforward.

The reason that genericity is convenient for describing the structure of $\overline{\mathcal{M}}_{0,0}(X_J, F; J)$ is that, as we harness in the proof of Lemma 2.4.11, it allows us to deduce that non-smooth stable curves in $\overline{\mathcal{M}}_{0,0}(X_J, F; J)$ are nodal curves with exactly two index 0 components intersecting each other exactly once transversely. This in turn facilitates the application of the gluing theorem (see for example [55, Corollary 2.39]) to deduce that each of these nodal curves has a neighbourhood homeomorphic to a 2-disc.

The result describing non-smooth stable curves in $\overline{\mathcal{M}}_{0,0}(X_J, F; J)$ is crucial. When we drop the genericity assumption, before Lemma 2.4.15, we need to reprove it, in essence. This is where automatic transversality (combined with non-bubbling results for certain index 0 curves) enters and tells us that we may still apply gluing to understand the structure of $\overline{\mathcal{M}}_{0,0}(X_J, F; J)$ for non-generic J.

Lemma 2.4.7 (*c.f.* Lemma 7.43 of [55])**.** *The moduli space $\overline{\mathcal{M}}_{0,0}(X_J, F; J)$ is homeomorphic to the complex plane $\mathbb{C}$.*

Lemma 2.4.8 (*c.f.* Lemma 7.45 of [55])**.** *The moduli spaces $\overline{\mathcal{M}}_{0,1}(X_J, F; J)$ and $\overline{\mathcal{M}}_{0,0}(X_J, F; J)$ admit smooth structures, making them open manifolds of dimension 4 and 2 respectively, such that π_J is a Lefschetz fibration with genus zero fibres and exactly one critical point in each singular fibre.*

Definition 2.4.9. Fix $J_0 \in \mathcal{J}(D')$ and denote $D_{J_0} := D' \cup E_{J_0}$. Let $U_{J_0} = X \backslash D_{J_0} \subset X_{J_0}$ and consider the subset of almost complex structures $\mathcal{J}_{\text{reg}}(U_{J_0}, J_0) \subset \mathcal{J}(D')$ defined to be those that satisfy

$$J|_{D_{J_0}} = J|_{X \backslash U_{J_0}} = J_0|_{X \backslash U_{J_0}} = J_0|_{D_{J_0}},$$

and every J-holomorphic curve that maps an injective point[13] into U_{J_0} is Fredholm regular.[14] This is a Baire subset[15] [53, Theorem 4.8] (see also [56, Theorem A.4]

[13] An injective point of a curve $u : \Sigma \to X$ is one where $u^{-1}(u(z)) = \{z\}$ and $du(z)$ is injective.

[14] A curve u is Fredholm regular when the linearisation of the Cauchy-Riemann operator at u is surjective. Said another way, this means the moduli space near u is cut out transversely and is thus a smooth manifold.

[15] A Baire subset is one which contains a countable intersection of open dense sets.

and Remark 3.2.3 of [37]) and so we can choose a small perturbation of J_0 that lives in $\mathcal{J}_{\text{reg}}(U_{J_0}, J_0)$. We shall abuse notation and denote this set by $\mathcal{J}_{\text{reg}}(D')$ since the choice of J_0 does not really matter.

Remark 2.4.10. One should think of the set U_{J_0} above as the set in which almost complex structures in $\mathcal{J}_{\text{reg}}(D')$ are allowed to vary. Note that, since $J|_{D_{\infty,J_0}} = J_0|_{D_{\infty,J_0}}$ we have that $D_{\infty,J} = D_{\infty,J_0}$ and thus $X_J = X_{J_0}$.

We first prove a result that restricts the form of the non-smooth stable curves in $\overline{\mathcal{M}}_{0,0}(X_J, F; J)$.

Lemma 2.4.11. *For every generic $J \in \mathcal{J}_{\text{reg}}(D')$, every non-smooth stable curve $\mathbf{u} \in \overline{\mathcal{M}}_{0,0}(X_J, F; J)$ is a nodal curve with exactly two transversely intersecting components $\mathbf{u} = (u_1, u_2)$ satisfying $([u_1], [u_2]) = (\mathcal{E}_j, F - \mathcal{E}_j)$ for some $1 \leq j \leq d$. In particular, there are no multiple covers.*

Proof. Let A_α be the homology class of an irreducible non-constant component u_α of $\mathbf{u} = (u_\alpha)_\alpha$. Since u_α is disjoint from D_{∞,J_0}, positivity of intersections implies that $A_\alpha \cdot D_\beta = 0$ for each irreducible component D_β of D_{∞,J_0}. We write this condition as $A_\alpha \in D_{\infty,J_0}^\perp$. Since the components of D_{∞,J_0} form a rank $n+2$ linearly independent set, and $\dim H_2(X) = n + d + 3$, we find that $D_{\infty,J_0}^\perp = \mathbb{Z}\langle F, \mathcal{E}_j : 1 \leq j \leq d\rangle$. Thus we may write, for some integers $\lambda, \mu_j \in \mathbb{Z}$,

$$A_\alpha = \lambda F + \sum_{j=1}^{d} \mu_j \mathcal{E}_j. \tag{2.4.2}$$

We claim that $c_1(A_\alpha) > 0$ for all non-constant components of $\mathbf{u}$. Indeed, this will follow from the Fredholm regularity condition of $J \in \mathcal{J}_{\text{reg}}(D')$, provided that each non-constant component passes through U_{J_0}. Now, by the open mapping theorem, a J-holomorphic map $u_\alpha : S^2 \to D_{J_0}$ is either constant, or covers a component of D_{J_0}. However, since $A_\alpha \in D_{\infty,J_0}^\perp$ and the only component of $D_{J_0} \backslash D_{\infty,J_0} = \overset{S}{\bullet}\!\!-\!\!\!-\!\!\overset{F}{\bullet}\!\!-\!\!\!-\!\!\overset{S'}{\bullet}$ disjoint from D_{∞,J_0} is the F-component, the latter case would imply that $A_\alpha = F$, which is only possible if $\mathbf{u}$ is smooth. Therefore, in the non-smooth case, we conclude that every component u_α passes through U_{J_0}. However, *a priori*, u_α may not be simple, so we pass to its underlying simple curve, say u'_α. This satisfies the condition that it maps an injective point into U_{J_0}, and so, it is Fredholm regular, implying that $c_1(u'_\alpha) > 0$, and thus $c_1(u_\alpha) = c_1(A_\alpha) > 0$.

Combining the above with the arguments of Proposition 4.8, Lemma 4.12, and §4.3 of [55], we deduce that $\mathbf{u}$ is a nodal curve with exactly two embedded components of index 0 intersecting transversely in a single node. Let us relabel the components as $\mathbf{u} = (u_1, u_2)$. Since the index of a curve is given by

$$\mathrm{ind}(u) = 2c_1([u]) - 2 \geq 0, \tag{2.4.3}$$

we find that $c_1(A_i) = 1$ for both components (u_1, u_2) of $\mathbf{u}$. As u_i is embedded, the adjunction formula [37, Theorem 2.6.4] says

$$c_1(A_i) = A_i^2 + 2,$$

implying that $A_i^2 = -1$ and so $\sum_j \mu_j^2 = 1$, from which we deduce that exactly one μ_j is non-zero, and furthermore, this coefficient is ± 1. The equation

$$2\lambda + \mu_j = c_1(A_i) = 1$$

ensures that either $A_i = \mathcal{E}_j$, or, $A_i = F - \mathcal{E}_j$. Finally, the condition

$$\sum_{i=1}^{2} A_i = F$$

implies that the other component u_{i+1} of $\mathbf{u}$ satisfies $[u_{i+1}] = A_{i+1} = F - A_i$, from which the result follows. $\square$

Remark 2.4.12. We call the types of nodal curves arising from Lemma 2.4.11 *curves of type* $(\mathcal{E}_j, F - \mathcal{E}_j)$.

Proof of Lemma 2.4.7 assuming genericity. A topological manifold structure on $\overline{\mathcal{M}}_{0,0}(X_J, F; J)$ is constructed via a gluing argument, such as that in the proof of Lemma 7.43 in [55]. Our application of gluing is valid since Lemma 2.4.11 showed that the only non-smooth stable curves are nodal curves with exactly two transversely intersecting components.

It remains to show that $\overline{\mathcal{M}}_{0,0}(X_J, F; J)$ is homeomorphic to $\mathbb{C}$. To this end, recall the evaluation map

$$\mathrm{ev} : \overline{\mathcal{M}}_{0,1}(X_J, F; J) \to X_J : \mathrm{ev}([(\mathbf{u}, x)]) = \mathbf{u}(x).$$

In our case, ev is bijective, which follows by a standard foliation and compactness argument (for example, combine [55, Proposition 2.53], with an argument such as the end of the proof of Theorem 1.16 of [56]). Moreover, since it's a proper map,[16] and X_J is locally compact and Hausdorff, it follows that ev is a homeomorphism onto X_J. Now fix a parametrisation $v : S^2 \to X$ of the divisor component of class S, and remove the unique point[17] from the domain that maps to $D_{\infty,J}$. Then, since $S \cdot F = 1$, the composition $\pi_J \circ \mathrm{ev}^{-1} \circ v : \mathbb{C} \to \overline{\mathcal{M}}_{0,0}(X_J, F; J)$ is the required homeomorphism. □

Proof of Lemma 2.4.8 assuming genericity. The smooth structures on the moduli spaces are constructed through the usual implicit function theorem argument [37, §3] on the set of Fredholm regular curves, combined with the gluing argument at the nodes given in [55, Lemma 7.45]. The same result there also shows that π_J also has a Lefschetz fibration structure at the nodal points. Since these arguments are inherently local, they also apply to our situation of the non-compact target manifold X_J. □

Corollary 2.4.13. *For generic $J \in \mathcal{J}_{\mathrm{reg}}(D')$, the moduli space $\overline{\mathcal{M}}_{0,1}(X_J, F; J)$ is homeomorphic to $Y\#k\mathbb{CP}^2$, where Y is a ruled surface over $\overline{\mathcal{M}}_{0,0}(X_J, F; J) \cong \mathbb{C}$, and $k \geq 0$ is the number of nodal fibres. In particular, the Euler characteristic satisfies*

$$\chi(\overline{\mathcal{M}}_{0,1}(X_J, F; J)) = k + 2. \tag{2.4.4}$$

Proof. The homeomorphism claim is Corollary 7.46 of [55], and the Euler characteristic formula then follows from a calculation using elementary properties of χ.

[16] Preimages of compact sets are compact.

[17] This is the point $S \cap D_{\infty,J}$.

Specifically, Y is fibred by 2-spheres over $\mathbb{C}$, and so

$$\chi(Y) = \chi(S^2)\chi(\mathbb{C}) = 2,$$

and connect summing with $\mathbb{CP}^2$ increases χ by 1. $\square$

Corollary 2.4.14. *For every generic $J \in \mathcal{J}_{\mathrm{reg}}(D')$, there are exactly d non-smooth stable curves in $\overline{\mathcal{M}}_{0,0}(X_J, F; J)$ given exactly by the curves of type $(\mathcal{E}_j, F - \mathcal{E}_j)$ for all $1 \leq j \leq d$.*

Proof. Since the evaluation map $\mathrm{ev} : \overline{\mathcal{M}}_{0,1}(X_J, F; J) \to X_J$ is a homeomorphism and, by Corollary 2.4.13, we have that

$$\#\text{nodal curves} + 2 = \chi(\overline{\mathcal{M}}_{0,1}(X_J, F; J)) = \chi(X_J) = d + 2,$$

and so there must be exactly d nodal curves in X_J, all of which are of type $(\mathcal{E}_j, F - \mathcal{E}_j)$ by Lemma 2.4.11. Since a nodal curve of this type is unique, that is, for each $1 \leq j \leq d$, there is at most one nodal curve of type $(\mathcal{E}_j, F - \mathcal{E}_j)$, then there is exactly one for each integer j. $\square$

It is at this moment that we drop the genericity condition. The next result is a generalisation of Lemma 2.4.11, and it will be used to show that all of the previous results that relied on genericity continue to hold.

Lemma 2.4.15. *For every (not necessarily generic) $J \in \mathcal{J}(D')$, a non-smooth stable curve in $\overline{\mathcal{M}}_{0,0}(X_J, F; J)$ must be of type $(\mathcal{E}_j, F - \mathcal{E}_j)$ for some $1 \leq j \leq d$.*

Proof. The argument is similar to that of Corollary 2.4.4. We will show that the set of almost complex structures $\mathcal{J}(D', A)$ that support a J-holomorphic curve in the class $A \in H_2(X)$, where A is one of the classes $\mathcal{E}_j, F - \mathcal{E}_j$, is equal to the whole space $\mathcal{J}(D')$ by automatic transversality and non-bubbling results. The fact that $\mathcal{J}(D', A)$ is non-empty follows immediately from Corollary 2.4.14, so we need only prove the non-bubbling results.

Suppose that $J \in \mathcal{J}(D')$, assume that $A = \mathcal{E}_j$, and let $\mathbf{u} \in \overline{\mathcal{M}}_{0,0}(\mathcal{E}_j; J)$ be a stable curve of class $\mathcal{E}_j$. Choose a non-constant component u_α of $\mathbf{u}$ and note, as in the proof of Lemma 2.4.3, that Lemma 2.4.2 ensures that

$$A_\alpha = [u_\alpha] \in (D_{\infty,J} \cup S)^\perp.$$

Since $D_\infty \cup S$ forms a rank $n+3$ linearly independent set in $H_2(X)$, we find that $(D_{\infty,J} \cup S)^\perp = \mathbb{Z}\langle \mathcal{E}_k : 1 \leq k \leq d \rangle$, and so, for some integers $\mu_k \in \mathbb{Z}$

$$A_\alpha = \sum_{k=1}^{d} \mu_k \mathcal{E}_k.$$

Recall that the areas of the classes $\mathcal{E}_k$ are all equal to l. Thus, we have that $\omega(A_\alpha)$ is an integer multiple of l:

$$\omega(A_\alpha) = l \sum_{k=1}^{d} \mu_k.$$

However, this is only possible if $\omega(A_\alpha) = l$, and so we conclude that $\mathbf{u} = u_\alpha$ is actually a smooth curve.

Repeating the above argument with $\mathcal{E}_j$ replaced with $F - \mathcal{E}_j$ and S replaced with S' shows that $F - \mathcal{E}_j$ curves also don't undergo bubbling. Hence, by the argument of Corollary 2.4.4, $\mathcal{J}(D', A) = \mathcal{J}(D')$. Summarising, we have shown that the classes $\mathcal{E}_j, F - \mathcal{E}_j$ are J-holomorphic for all $J \in \mathcal{J}(D')$. It remains to show that any non-smooth stable curve must be of type $(\mathcal{E}_j, F - \mathcal{E}_j)$.

Suppose that we have a stable curve $\mathbf{u} \in \overline{\mathcal{M}}_{0,0}(X_J, F; J)$ that is not of type $(\mathcal{E}_j, F - \mathcal{E}_j)$ for all $1 \leq j \leq d$. Then, since the classes $\mathcal{E}_j$ and $F - \mathcal{E}_j$ are J-holomorphic, a simple positivity of intersections argument, akin to Lemma 2.4.2, implies that $\mathbf{u}$ has no components in common with $\mathcal{E}_j$ and $F - \mathcal{E}_j$. In particular, for any component u_α of $\mathbf{u}$, we have that

$$A_\alpha \in \left(D_{\infty,J} \cup \bigcup_{j=1}^{d} \mathcal{E}_j \right)^\perp = \mathbb{Z}\langle F \rangle,$$

and so $\mathbf{u}$ must actually be a smooth curve. This completes the proof. $\square$

Corollary 2.4.16. *For every $J \in \mathcal{J}(D')$ the compactifying curves in $\overline{\mathcal{M}}_{0,0}(X, F; J)$, that is, the elements of $\overline{\mathcal{M}}_{0,0}(X, F; J) \backslash \mathcal{M}_{0,0}(X, F; J)$, are given exactly by:*

- *for each* $1 \leq j \leq d$*, a nodal curve of type* $(\mathcal{E}_j, F - \mathcal{E}_j)$*, and*
- *a stable curve* $\mathbf{u}_\infty^J$ *whose components cover those of* $D_{\infty,J}$*. We call this the* exotic curve *in* $\overline{\mathcal{M}}_{0,0}(X, F; J)$*.*

In particular, $\overline{\mathcal{M}}_{0,0}(X, F; J)$ *is homeomorphic to* S^2*.*

Proof. To prove that the curves of type $(\mathcal{E}_j, F-\mathcal{E}_j)$ do indeed exist, we wish to repeat the argument of Corollary 2.4.14. We achieve this by noting that Lemmata 2.4.7 and 2.4.8 continue to hold in the non-generic case, since the only way they depend on genericity is to prove that non-smooth stable curves consist only of nodal curves with exactly two components intersecting transversely, but Lemma 2.4.15 shows that this continues to hold in the non-generic case. Hence, the result of Corollary 2.4.14 continues to hold.

To see that the stable curve $\mathbf{u}_\infty^J \in \overline{\mathcal{M}}_{0,0}(X, F; J)$ exists, we use the fact that $\mathrm{ev} : \overline{\mathcal{M}}_{0,1}(X_J, F; J) \to X_J$ is a homeomorphism to pick a sequence of smooth curves $u_\nu \in \mathcal{M}_{0,0}(X, F; J)$ passing through points $x_\nu \in X$ converging to $x \in D_{\infty,J}$. Gromov compactness allows us to extract a convergent subsequence $u_\nu \to \mathbf{u}_\infty^J$. Since $F^2 = 0$, the image of $\mathbf{u}_\infty^J$ must be disjoint from all the curves in $\overline{\mathcal{M}}_{0,1}(X_J, F; J)$, so it follows that $\operatorname{im} \mathbf{u}_\infty^J \subset D_{\infty,J}$. Unique continuation [37, §2.3] of J-holomorphic curves implies that each of the non-constant components of $\mathbf{u}_\infty^J$ have the same image as a component of $D_{\infty,J}$. Hence, $\mathbf{u}_\infty^J$ covers $D_{\infty,J}$.

Finally, the Gromov topology on $\overline{\mathcal{M}}_{0,0}(X, F; J)$ realises it as the one point compactification of $\overline{\mathcal{M}}_{0,0}(X_J, F; J) \cong \mathbb{C}$. Whence we obtain that $\overline{\mathcal{M}}_{0,0}(X, F; J)$ is homeomorphic to S^2. □

2.5 The Lefschetz fibrations $\pi_J : X_J \to \mathbb{C}$

The results of Section 2.4.3 allow us to construct Lefschetz fibrations on $X_J = X \backslash D_{\infty,J}$ for all $J \in \mathcal{J}(D')$. This section proves the following proposition.

Proposition 2.5.1. *For every $J \in \mathcal{J}(D')$, there exists a Lefschetz fibration $\pi_J : X_J \to \mathbb{C}$ which has smooth fibres in the class F, and exactly d nodal fibres of type $(\mathcal{E}_j, F - \mathcal{E}_j)$.*

This almost immediately follows from §2.4.3, however there are a few smoothness technicalities to check. The aim is to show that the map

$$X_J \to \overline{\mathcal{M}}_{0,0}(X_J, F; J) : x \mapsto \text{the curve passing through } x$$

is a smooth Lefschetz fibration. Precisely, this map is equal to $\pi_J \circ \mathrm{ev}^{-1}$, where $\pi_J : \overline{\mathcal{M}}_{0,1}(X_J, F; J) \to \overline{\mathcal{M}}_{0,0}(X_J, F; J)$ is the universal curve from the previous section. We abuse notation and denote this composition by π_J also, as we won't make further reference to the universal curve. A detailed proof of the following can be found in Lemma 6.29 of [34]. We recall it here as our subsequent arguments crucially depend on it.

Lemma 2.5.2. *There exists a unique smooth structure on $\overline{\mathcal{M}}_{0,0}(X_J, F; J)$ such that the map $\pi_J : X_J \to \overline{\mathcal{M}}_{0,0}(X_J, F; J)$ is smooth except perhaps at images in X_J of the nodes of nodal curves in $\overline{\mathcal{M}}_{0,0}(X_J, F; J)$.*

Proof. We construct a smooth atlas on $\overline{\mathcal{M}}_{0,0}(X_J, F; J)$ making use of the smoothness and embeddedness properties of the curves in this moduli space. Choose a point $p \in X_J$ that is *not* the image of a node, and fix an embedded open 2-disc D_p passing through p that is transverse to the tangent space of the curve $u_p \in \overline{\mathcal{M}}_{0,0}(X_J, F; J)$ passing through p. After possibly shrinking D_p, we may assume that nearby curves in some open neighbourhood $U \subset \overline{\mathcal{M}}_{0,0}(X_J, F; J)$ of u_p intersect D_p exactly once, transversely. Therefore, we have a homeomorphism $U \to D_p$. We consider $U \to D_p$ to be a smooth chart by smoothly identifying D_p with the open unit disc in $\mathbb{C}$. We take the atlas to be the maximal one containing the collection of all such charts, all of which are smoothly compatible since the foliation $\overline{\mathcal{M}}_{0,0}(X_J, F; J)$ is smooth. The claims about smoothness of π_J and uniqueness of the corresponding atlas follow by construction. □

Remark 2.5.3. As in Lemma 2.4.7, $\overline{\mathcal{M}}_{0,0}(X_J, F; J)$ equipped with the smooth structure of Lemma 2.5.2 is diffeomorphic to the complex plane $\mathbb{C}$. Indeed, as a 2-dimensional topological manifold, there is a unique smooth structure up to diffeomorphism. However, it is not clear that these smooth structures are *compatible* with one another. Indeed, smoothness of gluing maps is a subtle question in general [55, Remark 2.40] [1, Remark 6.30], and one should be careful when discussing smooth structures on spaces of stable J-holomorphic curves.

Next, we want to show that π_J is smooth over the nodal points. Firstly, we note a result sketched in the appendix of [55], which explains how moduli spaces of square zero spheres that degenerate to a single curve with exactly one transverse node "look like" Lefschetz singular fibres.

Lemma 2.5.4 (Corollary A.3 of [55])**.** *Each nodal point in X_J admits a complex coordinate chart in which the leaves of the foliation $\overline{\mathcal{M}}_{0,0}(X_J, F; J)$ are identified with the fibres of the map $\pi_0 : \mathbb{C}^2 \to \mathbb{C} : \pi_0(z_1, z_2) = z_1 z_2$.*

Corollary 2.5.5. *The map $\pi_J : X_J \to \overline{\mathcal{M}}_{0,0}(X_J, F; J)$ is a smooth Lefschetz fibration. The nodal fibres are given exactly by the d nodal curves of type $(\mathcal{E}_j, F - \mathcal{E}_j)$.*

Proof. Let $V \subset \mathbb{C}$ be an open neighbourhood of $0 \in \mathbb{C}$ and choose a local section $s : V \to \mathbb{C}^2$ of π_0. Note that $\operatorname{im} s$ is necessarily disjoint from the unique nodal point $0 \in \mathbb{C}^2$. Denote one of the coordinate charts given by Lemma 2.5.4 by $\psi : U \to \mathbb{C}^2$. The fact that s is a section and ψ maps curves in $\overline{\mathcal{M}}_{0,0}(X_J, F; J)$ to fibres of π_0 ensures that the composition $\phi := \pi_J \circ \psi^{-1} \circ s : V \to \overline{\mathcal{M}}_{0,0}(X_J, F; J)$ is one of the charts of the atlas constructed in Lemma 2.5.2. Again, as ψ maps curves to fibres, we have that

$$\pi_J = \phi \circ \pi_0 \circ \psi, \tag{2.5.1}$$

that is, the following diagram commutes:

$$\begin{tikzcd} U \arrow[r, "\psi"] \arrow[d, "\pi_J"] & \mathbb{C}^2 \arrow[d, "\pi_0"] \\ \overline{\mathcal{M}}_{0,0}(X_J, F; J) & V \arrow[l, "\phi"] \end{tikzcd} . \tag{2.5.2}$$

Since the all the maps on the right-hand-side of (2.5.1) are smooth everywhere, we must have that π_J is smooth at the node. Moreover, by the commutativity of (2.5.2), the local charts ψ and ϕ give π_J the structure of a Lefschetz singularity at the node. The claim about the nodal fibres follows directly from Corollary 2.4.16. □

Proof of Proposition 2.5.1. All that remains to do is identify the moduli spaces $\overline{\mathcal{M}}_{0,0}(X_J, F; J)$ with $\mathbb{C}$ for all $J \in \mathcal{J}(D')$. In fact, we consider the subset $\mathcal{J}' \subset \mathcal{J}(D')$ of almost complex structures that are fixed along $TS \subset TX|_S$. Then, similarly to the proof of Lemma 2.4.11, we can fix a parametrisation $v : S^2 \to X$ of the S curve that is J-holomorphic for all $J \in \mathcal{J}'$ and such that $v(0)$ is in the F-component of D and $v(\infty) = S \cap D_{\infty,J}$. The condition $S \cdot F = 1$ ensures that the composition $\pi_J \circ v$ is a homeomorphism, and moreover, it is a diffeomorphism by construction of the smooth structure on $\overline{\mathcal{M}}_{0,0}(X_J, F; J)$ since the transverse disk D_p can be taken to be $v(S^2 \backslash \{\infty\})$. This completes the proof. □

Remark 2.5.6.

1. One can visualise the map $\pi_J : X_J \to \mathbb{C}$ as a Riemann sphere with $d + 2$ marked points. Indeed, for each $J \in \mathcal{J}(D')$, denote that curves of type $(\mathcal{E}_j, F - \mathcal{E}_j)$ by $\mathbf{u}_j^J$. Then we define a marked surface by

$$(S^2, (0, \infty, v^{-1}(\mathbf{u}_1^J), \ldots, v^{-1}(\mathbf{u}_d^J))).$$

   As J varies, the d marked points corresponding to Lefschetz singular fibres move around, whilst the points 0 and ∞ remain fixed.

2. The domain X_J of π_J depends on J. However, as we shall see later, the Lagrangian isotopy from L to a matching cycle of π_J occurs away from the curve $\mathbf{u}_\infty^J$. Therefore, by excising a suitable tubular neighbourhood of $\mathbf{u}_\infty^J$, we can fix a subset of X on which all of the maps π_J are defined (see the discussion preceding Corollary 4.3.3).

Chapter 3

Constructing a foliation adapted to a Lagrangian sphere

In this chapter we apply neck stretching to produce a foliation on T^*S^2 by J-holomorphic cylinders with exactly two singular leaves. Crucially, we show that the smooth leaves intersecting the zero section do so along smooth circles. This then leads to a J-holomorphic Lefschetz fibration on T^*S^2 for which the zero section is a matching cycle.

First, we recall some further details on the neck stretching setup. We then construct an almost complex structure $J_{T^*S^2}$ on T^*S^2 that is suitable for neck stretching, and satisfies the property that it is *anti-invariant* under fibre-wise multiplication by -1, which is the property that powers the intersection along circles argument.

The intersection theory for punctured J-holomorphic curves is reviewed in Section 3.2, as well as explaining how this works in the case of Morse-Bott degenerate asymptotics. This theory is then used in Section 3.3 to prove some basic facts about the foliations of T^*S^2 by J-holomorphic planes constructed by Hind [23].

The bulk of the chapter takes places in Section 3.4 where we apply intersection theory to analyse the holomorphic buildings that arise as limits of curves under neck stretching. Finally, Section 3.5 uses the limit analysis to construct a neck stretching sequence J_k such that all the curves in $\overline{\mathcal{M}}_{0,0}(F; J_k)$ that pass through the Weinstein neighbourhood of L converge to their respective holomorphic buildings as $k \to \infty$.

3.1 Neck-stretching, holomorphic buildings, and SFT compactness

Recall the neck stretching setup described in Section 1.3.3. Specifically, with $(M, \lambda) = (T^*_R L, \lambda_{\text{can}})$, a portion of the symplectisation $(-\epsilon, \epsilon) \times M \subset (\mathbb{R} \times M, \mathrm{d}(\mathrm{e}^r \lambda))$ symplectically embeds into X. Up to rescaling by the Liouville flow, we can assume that $R = 1$. We write $Y = X \backslash (-\epsilon, \epsilon) \times M$ as before, and form the stretched manifolds by inserting longer and longer necks:

$$X_t := Y \cup_{M_- \sqcup M_+} [-t - \epsilon, \epsilon] \times M.$$

These manifolds are diffeomorphic to $X_0 = X$ via a diffeomorphism that extends one of the form $(-t - \epsilon, \epsilon) \cong (-\epsilon, \epsilon)$. Since the Lagrangian sphere L is disjoint from the divisor $D' \subset X$, a cylindrical almost complex structure defined on the Weinstein neighbourhood $T^*_{\leq 1} L \subset X$ of L, can be extended to an element of[1] $\mathcal{J}(D')$. Since neck stretching only alters almost complex structures in the neck region, we have that $J_t \in \mathcal{J}(D')$ for all $t \geq 0$. Therefore, the results of Section 2.4 apply to J_t.

Throughout this chapter we'll use the rich and technical theory of punctured J-holomorphic curves and Symplectic Field Theory (SFT). We give an overview of the main relevant definitions, without getting too bogged down in the details.[2] The notation used here is that of [9], but the reader may also consult references such as [5, 14, 54] for different perspectives.

[1] Choose a tubular neighbourhood of D' disjoint from $T^*_{\leq 1} L$ and define J here, then extend arbitrarily everywhere else.

[2] The reader is cautioned that some important (but inconsequential for the narrative presented here) technical details are ignored here, and they can assume that any time a reference to the literature is made, the author is glossing over something not strictly relevant to the narrative here.

Consider a non-compact symplectic manifold (W, ω) with ends modelled on the half-symplectisations $(-\infty, 0] \times M$ or $[0, \infty) \times M$ of a contact manifold (M, λ). Given a compatible almost complex structure J that is cylindrical over the ends, we consider *punctured J-holomorphic curves* $f : \dot{\Sigma} \to W$ whose domain is a Riemann surface of finite type. A remarkable property of these curves is that, if they have *finite energy* (which we will not define here, see instead [9, 5]) then each of their punctures is asymptotic to a Reeb orbit of the contact manifold (M, λ) [28, 27, 40]. Our main examples will be the manifolds and curves that arise from neck stretching. In particular, $X_+^{(0)} \cong X \backslash L$ is a symplectic manifold with a *concave* cylindrical end modelled on $(-\infty, 0] \times T_1^* L$, and $X_-^{(0)} \cong T^* L$ is one with a *convex* end modelled on $[0, \infty) \times T_1^* L$. The punctured curves in these manifolds will form the components of a holomorphic building obtained through a compactness theorem described below. Another important example of a punctured curve is that of a trivial cylinder over a Reeb orbit:

Example 3.1.1. Given a Reeb orbit γ of period T in a contact manifold (M, λ), the trivial cylinder over γ is the map $u_\gamma : \mathbb{R} \times S^1 \to \mathbb{R} \times M$ defined by

$$u_\gamma(s, t) = (Ts, \gamma(t)).$$

Any cylindrical almost complex structure J on $\mathbb{R} \times M$ realises u_γ as a punctured J-holomorphic curve with exactly one positive and one negative puncture, both asymptotic to γ.

Recall from Section 1.3.3 the split manifold X^* and the almost complex structure J^* to which the neck stretching sequence J_t converges (on compact subsets). A J^*-holomorphic building in X^* is a J^*-holomorphic map $F : \Sigma^* \to X^*$ where Σ^* is a (potentially disconnected and non-compact) Riemann surface. Each level $F^{(\nu)} : \Sigma^{(\nu)} \to X^{(\nu)}$ is a finite-energy punctured curve, with each puncture asymptotic to a Reeb orbit in (M, λ). The set of punctures Γ is partitioned into positive and negative sets $\Gamma^\pm$ determined by whether the asymptote corresponding to a puncture lives in either the convex or concave ideal boundary of $X^{(\nu)}$ respectively. Crucially, the asymptotics in adjoining levels of a J^*-holomorphic building agree. That is, the ν level connects to the $(\nu+1)$ level, and the set of positive asymptotics of $F^{(\nu)}$ *is equal to* the set of negative asymptotics of $F^{(\nu+1)}$. Amongst other things, this means we can glue Σ^* together into a compact Riemann surface $\bar{\Sigma}$ and F correspondingly glues together into a continuous map $\bar{F}:\bar{\Sigma}\to X$, whence we obtain the homology class $[\bar{F}]\in H_2(X)$ of the building.[3]

A *stable* J^*-holomorphic building is one where none of the levels $F^{(\nu)}$ restrict to constant maps on any sphere component with fewer than 3 nodal points or punctures, and moreover, that no level $F^{(\nu)}$ is comprised solely of a union of trivial cylinders over Reeb orbits. We will produce stable J^*-holomorphic buildings by using the SFT compactness theorem [5, Theorem 10.3] [9, Theorem 1.1]:

Theorem 3.1.2 (SFT compactness)**.** *Given a neck stretching family J_t and a sequence of J_k-holomorphic curves $f_k:\Sigma\to X_k$ with uniformly bounded energy, there exists a subsequence converging to a stable J^*-holomorphic building F.*

The nature of the convergence is technical (see Definition 2.7 of [9] for the convention used here in full) and we note just a couple of its properties. For each level ν, up to reparametrisation, there exists a sequence of translations of the maps f_k in the neck region so that they converge to the level $F^{(\nu)}$ in the C^∞_{loc} topology. Additionally, the convergence ensures that of the homology classes:

$$[f_k]\to[\bar{F}]\in H_2(X).$$

In particular, if $[f_k]=A$ is independent of k, then $[\bar{F}]=A$.

Remark 3.1.3. The neck stretching setup allows us to think about points in the Weinstein neighbourhood $x\in Y_-=T^*_{\leq 1}L$ as either living in X, or in the completion $X_-^{(0)}\cong T^*L$. This is useful for phrasing statements like: given a $J^{(0)}$-holomorphic curve $F^{(0)}$, that forms part of a J^*-holomorphic building F, passing through $x\in X_-^{(0)}$, there exists a sequence of *closed* J_k-holomorphic curves f_k^x passing through x that converges to F. This phrasing will be used repeatedly in Chapter 4.

[3] Sometimes we will drop the bar notation and just write $[F]$ when no confusion is possible.

3.1.1 An almost complex structure adapted to a Lagrangian sphere

In this section we construct a family of $SO(3)$-invariant, ω_{can}-compatible almost complex structures on T^*S^2 that can be taken to be cylindrical on the complement of an arbitrarily small neighbourhood of the zero section. The construction is similar to Lemma 4.12 of [45], however we make it explicit here to ensure the property of anti-invariance under fibre-wise multiplication by -1.

First we recall a result taken from Seidel's thesis [42]. Let $\pi : \mathbb{C}^3 \to \mathbb{C}$ be the map $\pi(z) = z_1^2 + z_2^2 + z_3^2$ and denote by E the restriction of π to the real axis $\mathbb{R} \subset \mathbb{C}$. Equip E with the restriction of the standard Kähler structure $(\Omega = \omega_{\mathbb{C}^3}, i)$ on $\mathbb{C}^3$.

Lemma 3.1.4 ([42, Lemma 18.1])**.** *The restriction of* (E, Ω) *to* $(0, 1]$ *is isomorphic to the trivial symplectic fibre bundle* $(0, 1] \times (T^*S^2, \omega_{\text{can}})$*. An explicit isomorphism is given by:*

$$\Psi : z = x + \mathrm{i}y \mapsto (\pi(z), |x|^{-1}x, -|x|y).$$

Here we have employed the usual embedding $T^*S^2 = \{(q, p) \in \mathbb{R}^3 \times \mathbb{R}^3 \mid |q| = 1, \langle q, p\rangle = 0\}$, and the fact that $\omega_{\text{can}} = \sum_{i=1}^3 \mathrm{d}p_i \wedge \mathrm{d}q_i$. Under this identification, fibre-wise multiplication by -1 $(q, p) \mapsto (q, -p)$ is identified with complex conjugation.

Remark 3.1.5. In fact, Ψ restricts to an *exact* symplectomorphism $\Psi_s : (E_s, -y_i \mathrm{d}x_i) \to (\{s\} \times T^*S^2, p_i \mathrm{d}q_i)$ on each fibre $E_s = \pi^{-1}(s)$. This is critical for our situation since cylindrical almost complex structures depend on the Liouville structures present. Note also that Ψ is equivariant with respect to the diagonal Hamiltonian $SO(3)$-action on $\mathbb{R}^3 \times \mathbb{R}^3 = \mathbb{C}^3$.

The almost complex structure induced by the standard one i is not cylindrical with respect to the contact-type hypersurface $T_1^*S^2$, therefore, we explicitly construct one that is.

Proposition 3.1.6. *There exists a compatible almost complex structure* $J_{T^*S^2}$ *on* T^*S^2 *that is* anti-invariant *under the anti-symplectomorphism* φ *of fibre-wise* *multiplication by* -1*, that is,* $\varphi_* J_{T^*S^2} = -J_{T^*S^2}$*. Moreover,* $J_{T^*S^2}$ *can be made cylindrical outside of an arbitrarily small neighbourhood of* 0_{S^2}*.*

Proof. Consider the manifold $W := ([0,1] \times T^*S^2)\backslash(\{0\} \times 0_{S^2})$. Observe that Ψ extends to a isomorphism $E|_{[0,1]}\backslash\{0\} \to W$ that is exact on the fibres. Pushing forward i via Ψ_s yields a family of almost complex structures J_s on T^*S^2, defined only away from the zero-section in the case of J_0. Each of these is anti-invariant under φ and J_0 is cylindrical on $(T^*S^2\backslash 0_{S^2}, \lambda_{\text{can}}) = (\mathbb{R} \times T_1^*S^2, e^r\lambda_{\text{can}}|_{T_1^*S^2})$.

Choose a smooth function $\rho : \mathbb{R} \to \mathbb{R}$ that satisfies:

$$\begin{aligned} \rho|_{(-\infty,\frac{1}{4}]} &= 1 \\ \rho|_{[\frac{3}{4},\infty)} &= 0 \\ \rho'|_{(\frac{1}{4},\frac{3}{4})} &< 0. \end{aligned}$$

Then $J_{T^*S^2}(q,p) := J_{\rho(|p|)}(q,p)$ is the desired almost complex structure. By altering ρ the last claim of the result follows. ☐

3.2 Intersection theory of punctured curves

Intersection theory for punctured J-holomorphic curves was first laid out by Siefring in [50]. In that paper, Siefring assumes a certain non-degeneracy condition on the contact form. However, as we explain in this section, some of Siefring's results apply more widely (with few alterations) to the so-called Morse-Bott degenerate case considered in this work. We discuss these degeneracy conditions below in Section 3.2.1, along with the choices that need to be made to define the intersection number in the degenerate case.

3.2.1 Asymptotic constraints

When considering punctured J-holomorphic curves in a symplectic manifold (W, ω, J) with a cylindrical end modelled on a contact manifold (M, λ), there is a non-degeneracy condition one can ask of a closed Reeb orbit of λ. Specifically, let

$\gamma : [0, T] \to M$ be tangent to the Reeb vector field R_λ so that $\gamma(0) = \gamma(T)$. Consider the contact plane over $\gamma(0)$: $\xi_{\gamma(0)} = \ker \lambda(\gamma(0))$. The Reeb flow $\phi_t^{R_\lambda}$ preserves $\ker \lambda$, and so we can make the following definition: γ is *non-degenerate* when the restriction $\mathrm{d}(\phi_T^{R_\lambda})(\gamma(0))|_{\xi_{\gamma(0)}}$ has no eigenvalue equal to 1. We say that λ is non-degenerate when all of its closed Reeb orbits are non-degenerate.

A problem with this property is that many natural contact forms possessing symmetry are degenerate. Instead may consider the *Morse-Bott degenerate* condition, the full definition of which can be found in Definition 1.7 of [4]. The important property for us is that Morse-Bott degenerate Reeb orbits come in families of orbifolds of orbits sharing the same period T. We consider the especially simple case of $(M, \lambda) = (T_1^*S^2, \lambda_{\text{can}})$, which has the property that every simple[4] Reeb orbit has the same period, and $M \cong \mathbb{RP}^3$ is smoothly foliated by them. This implies that the quotient of M by the S^1 action of the Reeb flow is a smooth 2-dimensional manifold (diffeomorphic to S^2).

Associated to any Reeb orbit γ is a differential operator called an *asymptotic operator* $\mathbf{A}_\gamma$. Somewhat counter intuitively, knowing the precise definition will not be that useful for what follows.[5] The reader may consult [54, §3] for the precise definition along with justifications for the statements we make here. What is more useful for us is a basic understanding of the spectrum $\sigma(\mathbf{A}_\gamma)$: it is a discrete set composed exclusively of real eigenvalues. In dimension 4 (as we consider here) this property furnishes the definition of a certain winding number function defined on $\sigma(\mathbf{A}_\gamma)$ — see Equation (3.2.2).

The non-degeneracy condition of γ can now be characterised via its asymptotic operator as: $0 \notin \sigma(\mathbf{A}_\gamma)$. On the other hand, the Morse-Bott degenerate operators we handle here satisfy $\dim \ker(\mathbf{A}_\gamma) = 2$, which leads to problems in defining the

[4]A simple Reeb orbit is one with minimal period.

[5]However, we note that in a unitary trivialisation of the bundle $\gamma^*\xi$, one can always write

$$\mathbf{A}_\gamma = -J_0\partial_t - S(t),$$

where $S(t)$ is a loop of symmetric matrices and J_0 is the standard complex structure on $\mathbb{C}^n$.

intersection number and virtual dimensions of moduli spaces of curves. The root of the problem is that the conventional definition of the Conley-Zehnder index μ_{CZ} via the spectral flow [10] is undefined. This is resolved by perturbing the asymptotic operators by a small number $\epsilon \in \mathbb{R}\backslash\{0\}$ to make them non-degenerate and then computing the Conley-Zehnder index of the perturbed operator. The sign of ϵ corresponds[6] to whether the orbit is considered as *constrained* or *unconstrained* in the moduli problem. That is, whether or not we consider families of curves whose asymptotic orbits are allowed to move freely.

To make this correspondence precise, let $\dot{\Sigma} = \Sigma\backslash\Gamma$ be a punctured Riemann surface, where $\Gamma = \Gamma^+ \cup \Gamma^-$ is the finite set of punctures partitioned into positive and negative subsets. Suppose that (W, ω) has convex and concave ends modelled on (M_+, λ_+) and (M_-, λ_-) respectively, and let $f : \dot{\Sigma} \to W$ be a punctured curve asymptotic to a Reeb orbit γ at $z \in \Gamma$. Fix $\epsilon \in \mathbb{R}\backslash\{0\}$ so that

$$[-|\epsilon|, |\epsilon|] \cap \sigma(\mathbf{A}_\gamma)\backslash\{0\} = \emptyset, \tag{3.2.1}$$

then γ is constrained if the signs of the puncture z and ϵ agree, otherwise, it is unconstrained. We shall write $\gamma + \epsilon$ to denote the orbit subject to the perturbation ϵ, which corresponds to the asymptotic operator $\mathbf{A}_\gamma + \epsilon$.

Remark 3.2.1. In the sequel we will always assume (without always saying so) that ϵ (or often δ) has been chosen small enough so that Equation (3.2.1) holds for a fixed finite collection of asymptotic operators. These collections arise in situations when one considers curves with multiple punctures, or, more generally, a finite collection of punctured curves. When considering the asymptotic operators associated to the orbits of a punctured curve, we will often write $\mathbf{A}_z$ to denote the operator associated to the asymptotic orbit of the puncture $z \in \Gamma$. Furthermore, when it is clear from the context, we sometimes just write $\mathbf{A}$.

[6]The relationship between the sign of the perturbation and the geometric idea of whether an orbit is constrained or not is explained in §3.2 of [52]. Specifically, the reader is directed to the discussion of the splitting $T_u\mathcal{B} = W^{1,p,\delta}_\Lambda(u^*TW) \oplus V_\Gamma \oplus X_\Gamma$, and the proof of Proposition 3.7 therein.

Before proceeding further, we remark on some differences in approaches to defining the Conley-Zehnder index in the literature.

Remark 3.2.2. As already described, one can define an index in the Morse-Bott degenerate case by making a choice $\mathbf{c}_z \in \mathbb{R}\backslash\{0\}$ for each asymptotic operator $\mathbf{A}_z$ and compute the index μ_{CZ} of the perturbed *non-degenerate* operator $\mathbf{A}_z + \mathbf{c}_z$, as in Equation (3.5) of [52]; or, one can compute the Robbin-Salamon index μ_{RS} [41] of the linearised Reeb flow as done in [4]. It turns out that the latter can be realised as a special case of the former. Indeed, as in [4, §5], denote the linearised Reeb flow by $\Psi^{\pm}$ and let $\mathbf{A}_{\pm}$ be the asymptotic operator with parallel transport $\Psi^{\pm}$. Then

$$\begin{aligned}\mu^{\tau}_{\mathrm{CZ}}(\mathbf{A}_{\pm} \pm \delta) &= \mu_{\mathrm{RS}}(\Psi^{\pm'}) \\ &= \mu_{\mathrm{RS}}(\Psi^{\pm}) \mp \frac{1}{2}\dim\ker(\Psi^{\pm} - I) \\ &= \mu_{\mathrm{RS}}(\Psi^{\pm}) \mp \frac{1}{2}\dim\ker(\mathbf{A}_{\pm}),\end{aligned}$$

where $\Psi^{\pm'}$ is the perturbed parallel transport defined in §5.2.3 of [4]. Note that the left-hand side of the first line is what appears in the index formulas of [52], and the last line appears in the index formula of $\bar{\partial}_1$ of [4, §5].

In the sequel, we shall often use the notation $\mu_{\mathrm{RS}}(\gamma)$ to represent the Robbin-Salamon index of the linearised Reeb flow around the Reeb orbit γ. That is, $\mu_{\mathrm{RS}}(\gamma) = \mu_{\mathrm{RS}}(\Psi^{\pm})$.

3.2.2 Definitions and key results

Now let $g : \dot{\Sigma}' \to W$ be another J-holomorphic curve. In defining the intersection number between f and g, it is important to be precise about the orbit constraints one considers. Here, we will usually consider the case where *every* orbit of f and g is unconstrained, although, we will also consider the opposite case, where every orbit is constrained. To keep the notation light, we will denote these numbers as $i_U(f, g)$ and $i_C(f, g)$ respectively. At times where it is clear from the context which intersection number we are handling, we shall sometimes refer to them both as the *Siefring intersection number*, due to Siefring's original work [50].

The numbers $i_U(f,g)$ and $i_C(f,g)$ are generalisations of the usual intersection product for closed curves. See [56, §4] for a concise account of the non-degenerate case, and [52, §4.1] for the corresponding Morse-Bott statements that we use here. Moreover, there is a further generalisation: that of the intersection number of two holomorphic buildings. The mantra that we shall justify in this section is that the intersection theory of buildings arising in the splitting (or neck stretching) construction behaves almost exactly like that of closed curves.

Before stating the main definitions, we recall some terminology. We adopt the conventions of Wendl, which differ slightly from Siefring's; see Remark 4.7 of [56] for a note on the differences. Given a Reeb orbit γ let τ denote a unitary trivialisation of the contact bundle $\gamma^*\xi \to S^1$ over it.[7] As is customary, whenever a trivialisation over a multiple cover is required, we use the pullback of τ under the covering map $S^1 \to S^1 : \theta \mapsto k\theta$ and denote this by τ^k (although sometimes we abuse notation and just use τ). Similarly, we write γ^k for the pullback of γ under the k-fold covering map, and we say that γ^k is a k-fold cover of γ.

The reader may consult Sections 3 and 4 of the excellent book [56], or the original paper [26] for full details on the following statements about asymptotic winding numbers. See also Section 3.1.3 of [50] for a summary of the situation. Relative to τ, there is a well-defined winding number function $\mathrm{wind}^\tau_{\mathbf{A}} : \sigma(\mathbf{A}) \to \mathbb{Z}$. The extremal winding numbers $\alpha^\tau_\pm$ are defined as follows:[8]

$$\begin{aligned} \alpha^\tau_+(\mathbf{A}) &= \min\{\mathrm{wind}^\tau_{\mathbf{A}}(\lambda) \mid \lambda \in (0,+\infty) \cap \sigma(\mathbf{A})\} \\ \alpha^\tau_-(\mathbf{A}) &= \max\{\mathrm{wind}^\tau_{\mathbf{A}}(\lambda) \mid \lambda \in (-\infty,0) \cap \sigma(\mathbf{A})\}. \end{aligned} \tag{3.2.2}$$

We will also write $\alpha^\tau_\pm(\gamma)$ to mean $\alpha^\tau_\pm(\mathbf{A}_\gamma)$. Now suppose that γ is a simple Reeb orbit, then we define the numbers

$$\Omega^\tau_\pm(\gamma^k+\epsilon, \gamma^m+\epsilon') = km \min\left\{\frac{\mp\alpha^\tau_\mp(\gamma^k+\epsilon)}{k}, \frac{\mp\alpha^\tau_\mp(\gamma^m+\epsilon')}{m}\right\}.$$

This definition is extended to all pairs of Reeb orbits by asserting that

$$\Omega^\tau_\pm(\gamma+\epsilon, \gamma'+\epsilon') = 0$$

[7] Occasionally we will also use Φ instead of τ.

[8] Observe that these definitions are sound, no matter the degeneracy of $\mathbf{A}$.

if γ and γ' are not covers of the same simple orbit.

The final ingredients needed to state the main definitions are the relative intersection number $f \bullet_\tau g$ of two punctured curves, and the relative first Chern number $c_1^\tau(E)$ associated to a complex bundle $E \to \dot{\Sigma}$ over a punctured curve. The former is defined to be the usual algebraic count of intersections between f and the push-off of g in the direction determined by τ near infinity, and the latter is defined via algebraic counts of zeros of sections of E that are non-zero and constant near infinity. For precise definitions see Sections 4.2 and 3.4 of [56] respectively. For a punctured curve $f : \dot{\Sigma} \to W$ we'll often use the shorthand $c_1^\tau(f) := c_1^\tau(f^*TW)$.

Definition 3.2.3. 1. Let $f : \dot{\Sigma} \to W$ and $g : \dot{\Sigma}' \to W$ be punctured holomorphic curves in an almost complex manifold with cylindrical ends W, then for any $\delta > 0$ sufficiently small, their constrained and unconstrained intersection numbers are defined as[9]

$$i_C(f, g) := f \bullet_\tau g - \sum_{(z,w)\in\Gamma^\pm\times(\Gamma')^\pm} \Omega_\pm^\tau(\gamma_z \pm \delta, \gamma_w \pm \delta),$$
$$i_U(f, g) := f \bullet_\tau g - \sum_{(z,w)\in\Gamma^\pm\times(\Gamma')^\pm} \Omega_\pm^\tau(\gamma_z \mp \delta, \gamma_w \mp \delta).$$

2. Let $F = (F^{(0)}, \ldots, F^{(N)})$ and $G = (G^{(0)}, \ldots, G^{(N)})$ be holomorphic buildings in a split manifold, then we define their intersection number to be

$$i(F, G) := \sum_{i=0}^{N} F^{(i)} \bullet_\tau G^{(i)}.$$

3. The normal Chern number of a curve f is defined as

$$c_N(f) = c_1^\tau(f) - \chi(\dot{\Sigma}) \pm \sum_{z\in\Gamma^\pm} \alpha_\mp^\tau(\gamma_z \mp \delta).$$

[9] In this equation and in those that follow, the notation $\pm$ indicates that the sum appears twice, once with a plus sign, and again with a minus sign. For example,

$$\sum_{(z,w)\in\Gamma^\pm\times(\Gamma')^\pm} \Omega_\pm^\tau(\gamma_z \pm \delta, \gamma_w \pm \delta) = \sum_{(z,w)\in\Gamma^+\times(\Gamma')^+} \Omega_+^\tau(\gamma_z + \delta, \gamma_w + \delta) + \sum_{(z,w)\in\Gamma^-\times(\Gamma')^-} \Omega_-^\tau(\gamma_z - \delta, \gamma_w - \delta).$$

These definitions are sound since the dependence on the trivialisation τ cancels in each of the equations. In particular see Appendix C.5 of [56] for the justification of this claim for intersections of buildings.

We now state some key results that will be crucial in our analysis of holomorphic buildings arising from the neck stretching process. First, the adjunction inequality, which is an immediate corollary of the adjunction formula given in [52, §4].

Theorem 3.2.4 (Adjunction inequality.). *Let $f : \dot{\Sigma} \to W$ be a punctured curve in a manifold with cylindrical ends. Then,*

$$i_U(f, f) \geq c_N(f).$$

Next, we have a statement of continuity of the intersection products with respect to the topology of the moduli spaces of holomorphic buildings. This can be derived from Lemma 5.7 of [50].

Theorem 3.2.5. *If f_k and g_k are sequences of closed curves in (X_k, J_k) converging to holomorphic buildings F and G respectively under neck stretching, then $i(F, G) = f_k \cdot g_k$ for k large.*

Finally, the following statement is our justification of the mantra that holomorphic buildings in split manifolds behave like closed holomorphic curves. The proof is given in Section 3.2.3. No novel techniques are involved, so it is only included for completeness.

Proposition 3.2.6. *In the Morse-Bott setting considered here (that is, where the asymptotic operators associated to punctures of curves satisfy* $\dim \ker \mathbf{A} = 2$*), the following holds:*

1. *The intersection product for holomorphic buildings is additive over the levels with respect to $i_U(\cdot, \cdot)$: for $F = (F^{(0)}, \dots, F^{(N)})$ and $G = (G^{(0)}, \dots, G^{(N)})$ we have*
$$i(F, G) = \sum_{i=0}^{N} i_U(F^{(i)}, G^{(i)}).$$

2. *The unconstrained intersection number $i_U(\cdot,\cdot)$ is linear with respect to multiple covers. That is, if f is a d-fold cover of f', then $i_U(f,g) = d(i_U(f',g))$.*

3.2.3 Proof of Proposition 3.2.6

We prove Part 1, and Part 2 will be picked up along the way as a corollary of our techniques. The key is to understand how a common orbit γ of the buildings F and G contributes to the sum $\sum_{i=0}^{N} i_U(F^{(i)}, G^{(i)})$ as both a positive orbit for $(F^{(i)}, G^{(i)})$ and a negative one for $(F^{(i+1)}, G^{(i+1)})$. Such an orbit is called a *breaking orbit*, and the corresponding pair punctures (z,w) are called a *breaking pair.*

Now,

$$i(F,G) - \sum_{i=0}^{N} i_U(F^{(i)}, G^{(i)}) = \sum_{\substack{(z,w) \\ \text{breaking pairs}}} \mathrm{Br}(\gamma_z + \delta, \gamma_w + \delta),$$

where

$$\mathrm{Br}(\gamma_z + \epsilon, \gamma_w + \epsilon) := \Omega_+^\tau(\gamma_z - \epsilon, \gamma_w - \epsilon) + \Omega_-^\tau(\gamma_z + \epsilon, \gamma_w + \epsilon).$$

is the *breaking contribution* of the breaking pair (z,w). Henceforth we suppress the ϵ notation in $\mathrm{Br}(\gamma + \epsilon, \gamma' + \epsilon)$ since it is independent of the number $\epsilon > 0$ chosen, provided that it is small enough. It suffices to show that $\mathrm{Br}(\gamma, \gamma') = 0$ for all unconstrained orbits γ and γ' in Morse-Bott degenerate families.

Given an asymptotic operator $\mathbf{A} = \mathbf{A}_\gamma : \Gamma(\gamma^*\xi) \to \Gamma(\gamma^*\xi)$ of a Reeb orbit γ, denote its k-fold cover by $\mathbf{A}^k = \mathbf{A}_{\gamma^k}$. Consider the map

$$\ker(\mathbf{A}) \to \ker(\mathbf{A}^k) : f \mapsto f_k,$$

sending a section $f \in \ker(\mathbf{A})$ to its pullback f_k under the k-fold covering map $S^1 \to S^1 : \theta \mapsto k\theta$. The next lemma shows that it is an linear isomorphism.

Lemma 3.2.7. *Suppose that $\mathbf{A}$ is the asymptotic operator of a hermitian line bundle*[10] *$E \to S^1$. Then the map $\ker(\mathbf{A}) \to \ker(\mathbf{A}^k) : f \mapsto f_k$ is a linear isomorphism.*

[10] Which is always the case for us, since we deal only with 4-dimensional manifolds.

Proof. The map is well-defined since, if f is an eigenfunction of $\mathbf{A}$ with eigenvalue λ, then f_k is one of $\mathbf{A}^k$ with eigenvalue $k\lambda$. Injectivity is easily seen, and surjectivity follows from the fact that the covering multiplicity of an eigenfunction g of $\mathbf{A}^k$ is given by [50, Lemma 3.2]

$$\mathrm{cov}(g) = \gcd(\mathrm{wind}(\tau^k \circ g), k),$$

and that the winding function is constant on eigenspaces [26, Lemma 3.4]. Explicitly, take a non-zero $f \in \ker(\mathbf{A})$, then $f_k \in \ker(\mathbf{A}^k)$ is non-trivial, so

$$\mathrm{wind}(\tau^k \circ g) = \mathrm{wind}(\tau^k \circ f_k) = k\mathrm{wind}(\tau \circ f),$$

and thus $\mathrm{cov}(g) = k$. $\square$

Lemma 3.2.8. *Suppose that γ is a simple Morse-Bott degenerate Reeb orbit. Then, for all integers $k, m > 0$, we have that* $\mathrm{Br}(\gamma^k, \gamma^m) = 0$.

Proof. Recall the definition

$$\Omega^\tau_\pm(\gamma^k + \epsilon, \gamma^m + \epsilon') = km \min\left\{\frac{\mp\alpha^\tau_\mp(\gamma^k + \epsilon)}{k}, \frac{\mp\alpha^\tau_\mp(\gamma^m + \epsilon')}{m}\right\}.$$

Thus,

$$\mathrm{Br}(\gamma^k, \gamma^m) = \min\{m\alpha^\tau_+(\gamma^k + \epsilon), k\alpha^\tau_+(\gamma^m + \epsilon)\} - \max\{m\alpha^\tau_-(\gamma^k - \epsilon), k\alpha^\tau_-(\gamma^m - \epsilon)\},$$

and the result then follows from the claim that there exists an integer l such that, for all $k > 0$, $\alpha^\tau_+(\gamma^k + \epsilon) = kl = \alpha^\tau_-(\gamma^k - \epsilon)$. Indeed, from this it follows that all the terms involved in the min/max functions above are equal to kml.

To prove the claim, recall that $\mathrm{wind}^\tau_{\mathbf{A}} : \sigma(\mathbf{A}) \to \mathbb{Z}$ is monotone increasing and attains each value in $\mathbb{Z}$ exactly twice (accounting for multiplicities of eigenvalues). Therefore, since $\dim\ker(\mathbf{A}^k) = 2$, and $\delta > 0$ is chosen such that $\delta < |\lambda|$ for each non-zero eigenvalue $\lambda \in \sigma(\mathbf{A}^k)$, it is clear from the definitions that (compare Figure 3.1)

$$\alpha^\tau_+(\mathbf{A}^k + \delta) = \mathrm{wind}^\tau_{\mathbf{A}^k}(0) = \alpha^\tau_-(\mathbf{A}^k - \delta).$$

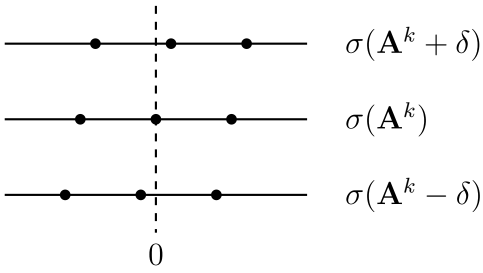

Figure 3.1: The effect of perturbing on the spectrum of an asymptotic operator $\mathbf{A}$. The bullets represent elements of the spectrum.

The previous lemma ensures that any non-zero $f \in \ker(\mathbf{A}^k)$ is a k-fold cover of a non-zero element of $\ker(\mathbf{A})$, and so, $\mathrm{wind}^\tau_{\mathbf{A}^k}(0) = k\mathrm{wind}^\tau_{\mathbf{A}}(0)$. Setting $l = \mathrm{wind}^\tau_{\mathbf{A}}(0)$ completes the proof. $\square$

Remark 3.2.9. Note that we can extract a proof of Part 2 of Proposition 3.2.6 from the proof of the claim above through the observation that it shows that[11]

$$\alpha^\tau_\pm(\mathbf{A}^k \pm \delta) = k\alpha^\tau_\pm(\mathbf{A} \pm \delta).$$

Indeed, this is because $\mathrm{wind}^\tau_{\mathbf{A}}(0) = \alpha^\tau_\pm(\mathbf{A} \pm \delta)$. Therefore, both components $f \bullet_\tau g$ and $\Omega^\tau_\pm(\gamma^k, \gamma^l)$ scale linearly with the covering number. More precisely, $(df) \bullet_\tau g = d(f \bullet_\tau g)$ and $\Omega^\tau_\pm(\gamma^{dk}, \gamma^{dl}) = d\Omega^\tau_\pm(\gamma^k, \gamma^l)$. Hence, $i_U(df, g) = di_U(f, g)$.

3.2.4 Positivity of intersections for curves with degenerate asymptotics

Siefring's paper [50] essentially shows that we have positivity of intersections for punctured curves with Morse-Bott degenerate asymptotics.

Theorem 3.2.10 (Positivity of intersections of punctured curves)**.** *Suppose that f, g are punctured curves such that $f^{-1}(\mathrm{im}\, g)$ does not contain a non-empty open*

[11] Note the distinction to the non-degenerate case, where this is only an inequality in general, *cf.* [56, Proposition C.2]. See also [52, §4.2] for a discussion on covering relations in the degenerate case, in particular Proposition 4.6 therein, of which this result is a special case.

set. Then $i_*(f,g) \geq 0$, where the notation i_* indicates either the constrained or unconstrained intersection number. Moreover, equality occurs if, and only if, f and g are disjoint and remain so after homotopy.

We sketch the proof, which essentially depends on two facts. The first is local positivity of intersections, which is the same story as the case of closed curves, covered in detail in [37], for example. The second is the asymptotic representation for a J-holomorphic half-cylinder. This topic has been covered extensively in the literature, for example [28, 56] cover the non-degenerate case, and [27, 4, 5] cover the Morse-Bott degenerate case.

Despite the non-degeneracy assertion made in [50], only minor modifications are required to prove Theorem 3.2.10. The technical results on winding numbers that power the relevant positivity of intersections results are isolated in Section 3.1.4 of [50]. These in turn depend on two key ingredients: the asymptotic representative of the difference between two J-holomorphic half-cylinders (Theorem 3.6 of [50][12]); and the properties of the winding function[13] $\mathrm{wind}^\tau_{\mathbf{A}}$ (Lemma 3.1 of [50]), which is a summary of the properties proved in Lemmata 3.4–3.7 of [26]. The latter of these is valid irrespective of whether $\mathbf{A}$ is degenerate or not, and the former, as Siefring remarks in Section 2.2 of [49], only depends on the exponential decay of a half-cylinder. Therefore, one can apply the decay results of[14] [27] and [4, §3.3] (see also [5, Appendix A], or [40]) to conclude that Theorem 3.6 of [50] holds for punctured J-holomorphic curves with Morse-Bott asymptotics.

To apply the results in [50], one must be careful about exactly what is meant by the Conley-Zehnder index that appears there.[15] The properties of $\alpha^\tau_\pm$, ensure that

$$\alpha^\tau_\pm(\mathbf{A} \mp \delta) = \alpha^\tau_\pm(\mathbf{A}),$$

[12] Which is stated and proved as Theorem 2.2 in [49].

[13] Denoted by $w(\lambda, [\Phi])$ in [50].

[14] Note that in [27] the manifold of unparametrised Reeb orbits is assumed to be a circle. The proof for more general manifolds of orbits (including our case of interest, the 2-sphere) can be found in either [4, §3.3], or [5, Appendix A], or [40].

[15] Denoted by μ^Φ in [50].

provided that $\delta > 0$ is sufficiently small. Therefore, if one interprets the definition of μ^{τ}_{CZ} given in [50] literally,[16] then the constructions therein[17] give the constrained intersection number $i_C(\cdot,\cdot)$. Therefore, positivity of intersections holds for i_C. A special case of Proposition 4.11 of [52] is the inequality

$$i_U(u,v) \geq i_C(u,v),$$

which implies positivity of intersections for i_U too.

3.3 J-holomorphic planes in T^*S^2

Hind proved[18] that, for any cylindrical almost complex structure J on T^*S^2, there are two transverse foliations $\mathcal{M}_\pm$ by finite-energy J-holomorphic planes. They are distinguished by the sign of the intersection with the zero-section $0_{S^2} \subset T^*S^2$:

$$P_\pm \in \mathcal{M}_\pm \Rightarrow P_\pm \cdot 0_{S^2} = \pm 1.$$

[16] This is exactly Definition 3.9 of [26]. In particular, in our situation where the Morse-Bott manifold is the whole contact manifold, and so $\dim \ker \mathbf{A} = 2$, we have that the parity $p(\gamma^k)$ (of the perturbed operator $A \pm \epsilon$) is 1. Moreover, comparing the notation of Siefring and Wendl, we have, for an integer $k \in \mathbb{Z}\backslash\{0\}$,

$$\alpha^{\Phi}(\gamma^k) = \pm\alpha^{\Phi}_{\mp}(\gamma^{|k|}) = \pm\alpha^{\Phi}_{\mp}(\gamma^{|k|} \pm \epsilon),$$

where the leftmost sign agrees with the distinction of γ as a positive or negative orbit. It then follows that

$$\mu^{\Phi}(\gamma^k) = 2\alpha^{\Phi}(\gamma^k) + p(\gamma^k) = \pm(2\alpha^{\Phi}_{\mp}(\gamma^{|k|}) \pm 1) = \pm\mu^{\Phi}_{\mathrm{CZ}}(\gamma^{|k|} \pm \epsilon),$$

where the final equality follows from equation (2.3) of [52].

[17] To directly compare the notation that appears in [50] with that used here, one can work through the definitions to show that

$$m_z m_w \max\left\{\frac{\alpha^{\Phi}(\gamma_z^{m_z})}{|m_z|}, \frac{\alpha^{\Phi}(\gamma_w^{m_w})}{|m_w|}\right\} = -\Omega^{\Phi}_{\pm}(\gamma_z^{|m_z|}, \gamma_w^{|m_w|}).$$

The numbers on the left hand side are used to define the intersection in [50], and those on the right are used to define the constrained intersection in [52].

[18] The paper [23] combined with results on automatic transversality.

As remarked by Evans [15, §6.4], two planes $P_\pm \in \mathcal{M}_\pm$ of opposite parity intersect positively exactly once if, and only if, they have distinct asymptotic orbits. Otherwise, they are disjoint. However, the Siefring intersection "sees" this intersection, even for planes with a common orbit:

Lemma 3.3.1. *1. The unconstrained intersection of two planes of opposite parity is equal to 1: $i_U(P_+, P_-) = 1$.*

2. Let $P_\pm, Q_\pm \in \mathcal{M}_\pm$ be two planes of the same parity. Then $i_U(P_\pm, Q_\pm) = 0$.

Proof. 1. As the Siefring intersection is homotopy invariant, we can perturb the plane P_+ to one nearby, say P'_+ in its moduli space, guaranteeing that P'_+ and P_- have distinct orbits. Then $i_U(P'_+, P_-) = P'_+ \cdot P_- = 1$, as already noted.

2. As above, we can homotope one of the planes so that it has distinct asymptotic to the other, ensuring that the asymptotic contribution to the intersection product is zero and thus $i_U(P_\pm, Q_\pm) = P_\pm \cdot Q_\pm = 0$, where the last equality follows since planes of the same parity form a foliation.

□

Proposition 3.3.2. *Let $u : \dot{\Sigma} \to T^*S^2$ be a punctured J-holomorphic curve. Then exactly one of the following is true:*

1. u is a cover of a plane in $\mathcal{M}_\pm$; or,

2. u intersects both types of plane positively. More precisely, $i_U(u, P_\pm) > 0$.

Proof. Assume the first scenario is false, that is, u is not a cover of any $P_\pm \in \mathcal{M}_\pm$. Then, Theorem 3.2.10 implies that $i_U(u, P_\pm) \geq 0$. Note that this inequality must be strict since $\mathcal{M}_\pm$ is a foliation, and so u intersects one of the planes $P_\pm$ in at least one point. □

Corollary 3.3.3. *Any unconstrained intersection of punctured J-holomorphic curves in T^*S^2 is non-negative.*

3.4 Analysis of holomorphic limit buildings

3.4.1 Consequences of neck stretching about a Lagrangian sphere

The following result can be found in essential form in [23, p.315]. See also an expanded proof in [15, Lemma 7.5]. It is a simple consequence of a curve in a symplectisation having positive dλ-energy and the topology of $\mathbb{RP}^3$. We recall the argument here as we will need a modified version of it later.

Lemma 3.4.1. *Suppose that* $G = (G^{(0)}, \dots, G^{(N)}) : \Sigma^* \to X^*$ *is a genus 0 holomorphic building such that* $\Sigma_+^{(0)}$ *is connected and* $G_+^{(0)}$ *has only simple punctures. Then* $N = 0$*, that is,* G *has no non-trivial symplectisation levels.*

Proof. Let $g = (a, u) : \dot{\Sigma} \to \mathbb{R} \times T_1^* L$ be a connected component of the symplectisation level $G^{(N)}$ — the one connected to the top level $G_+^{(0)}$. The maximum principle implies that g has at least one positive puncture. On the other hand, since the genus[19] of G is 0 and $G_+^{(0)}$ is connected, g has *at most*, and therefore exactly, one positive puncture γ. Moreover, it must have at least one negative puncture, as otherwise $\dot{\Sigma}$ would be the complex plane and thus g would represent a contraction of a simple Reeb orbit, which represents the non-trivial element in $\pi_1(\mathbb{RP}^3) \cong \mathbb{Z}_2$.

Next we consider the $E_{\mathrm{d}\lambda}$-energy of g, defined in [5, §5.3] as[20]

$$E_{\mathrm{d}\lambda}(g) = \int_{\dot{\Sigma}} u^* \mathrm{d}\lambda.$$

By Lemmata 5.4 and 5.16 of [5] we have that this energy is non-negative and equal to the difference between the sum of the periods of the positive punctures and the sum of those of the negative ones. Thus, since γ is simple and all simple Reeb orbits in $(T_1^* S^2, \lambda_{\mathrm{can}})$ have the same period, there is at most one negative puncture, and therefore, exactly one. Whence we obtain $E_{\mathrm{d}\lambda}(g) = 0$. Theorem 6.11 of [26] then

[19] The genus of a holomorphic building $G : \Sigma^* \to X^*$ is the genus of the surface $\bar{\Sigma}$ obtained by gluing Σ^* together along its punctures. We will always deal with genus 0 buildings.

[20] In the notation of that paper we have $\omega = \mathrm{d}\lambda$.

implies that g is a reparametrisation of a trivial cylinder. Repeating this argument shows that every symplectisation level is composed solely of trivial cylinders, a situation which is ruled out by the stability condition. Hence the result is proved. □

The index formula of a punctured curve will be crucial in the proofs of this section. The index[21] of a curve is the expected dimension of the moduli space it lives in, and thus it is often also called the virtual dimension. Let $u : \dot{\Sigma} \to W$ be a punctured J-holomorphic curve with $k_\pm$ positive/negative punctures mapping into a symplectic manifold with cylindrical ends W. The index of u is given by[22]

$$\mathrm{ind}(u) = 2c_1^\tau(u) - \chi(\dot{\Sigma}) \pm \sum_{i=1}^{k_\pm} \mu_{\mathrm{CZ}}^\tau(\mathbf{A}_{z_i} \mp \delta).$$

In the notation of [4] and [23], this becomes[23]

$$\mathrm{ind}(u) = 2c_1^\tau(u) - \chi(\dot{\Sigma}) \pm \sum_{i=1}^{k_\pm} \left(\mu_{\mathrm{RS}}(\gamma_i^\pm) \pm \frac{1}{2}\dim(\gamma_i^+)\right),$$

where $\dim(\gamma)$ is the dimension of the moduli space of *unparametrised* Reeb orbits that γ lives in. In our situation, we deal only with genus zero curves with asymptotics that live in 2-dimensional families. Fixing τ to be the trivialisation that appears in Lemma 7 of [23], then we have $\mu_{\mathrm{RS}}(\gamma_i^\pm) = 2\,\mathrm{cov}(\gamma_i^\pm)$, so the index formula reduces to:

$$\mathrm{ind}(u) = 2\left(k_+ + k_- - 1 + c_1^\tau(u) \pm \sum_{i=1}^{k_\pm} \mathrm{cov}(\gamma_i^\pm)\right). \tag{3.4.1}$$

We primarily deal with the cases where one of $k_\pm$ is zero, which simplifies the formula further. In particular, in the case where u maps into the top level $X_+^{(0)}$, we have $k_+ = 0$ and so, combined with the fact that $\mathrm{cov}(\gamma) \geq 1$, we obtain the

[21] The name index comes from the fact that it is equal to the Fredholm index of a certain operator derived from the Cauchy-Riemann equation (see [4] for example).

[22] Observe that this is the *unconstrained* index. Moreover, in the case that u is a closed curve, it reduces to the index formula given in Equation (2.4.3).

[23] See Remark 3.2.2.

inequality:[24]

$$\operatorname{ind}(u) \leq 2(c_1^\tau(u) - 1). \tag{3.4.2}$$

Lemma 3.4.2. *Let $F = (F^{(0)}, \dots, F^{(N)})$ be a holomorphic building. Then the relative first Chern numbers of every level except the top one vanish. In particular, the relative first Chern number of the top level equals the first Chern number of the homology class $[\bar{F}]$:*

$$c_1^\tau(F_+^{(0)}) = c_1([\bar{F}]). \tag{3.4.3}$$

Proof. The chosen trivialisation τ extends to a global trivialisation of $TX^{(i)}$ and $TX_-^{(0)}$ for the symplectisation levels $X^{(i)} = \mathbb{R} \times T_1^*S^2$ and bottom level $X_-^{(0)} = T^*S^2$. Indeed, for the symplectisation levels we have the global splitting

$$T(\mathbb{R} \times T_1^*S^2) = \langle Z, R \rangle \oplus \xi,$$

where Z and R are the Liouville and Reeb vector fields respectively. For $X_-^{(0)} = T^*S^2$ the extension is facilitated by the global splitting of $T(T^*S^2)$ into vertical and horizontal Lagrangian planes [23, Lemma 7]. Therefore, we have $c_1^\tau(u) = 0$ for any punctured curve mapping into $X^{(i)}$ or $X_-^{(0)}$. Combining this with the fact[25] that, for a holomorphic building $F = (F^{(0)}, \dots, F^{(N)})$, we have

$$\sum_{i=0}^{N} c_1^\tau(F^{(i)}) = c_1([\bar{F}]),$$

we obtain the result. □

Remark 3.4.3. Compare the above discussion with Section 6.3 of [15], where Evans gives a proof using complex geometry and the compactification of T^*S^2 to the projective quadric surface.

We now turn to recording a genericity result. This will be achieved by perturbing the almost complex structure $J_+^{(0)}$ on the top level $X_+^{(0)} \cong X \setminus L$ in a suitable open set.

[24] We caution the reader that this inequality is *only* valid with respect to the fixed trivialisation τ.

[25] This is a consequence of gluing the levels $F^{(i)}$ back together to obtain the cycle $[\bar{F}]$ in X.

This set must be chosen carefully though, as an arbitrary perturbation would either destroy the $c_1 \leq 0$ curves in the divisor $D' \subset X$, or the cylindrical nature of $J_+^{(0)}$ in the neck region $(-\infty, \epsilon) \times T_1^* L$. Let V be the open set $X_+^{(0)} \setminus ((-\infty, \epsilon] \times T_1^* L \cup D')$. We are then free to perturb $J_+^{(0)}$ in the set V to make it generic. This means that any $J_+^{(0)}$-holomorphic curve u mapping an injective point into V is Fredholm regular, and thus $\mathrm{ind}(u) \geq 0$.

Lemma 3.4.4. *Let $J_+^{(0)}$ be as above and $u : \dot{\Sigma} \to X_+^{(0)}$ be a somewhere injective $J_+^{(0)}$-holomorphic curve that is either closed and not contained in D' or has at least one (negative) puncture. Then u intersects $V \subseteq X_+^{(0)}$ and is thus Fredholm regular.*

Proof. First we deal with the closed case. Since the cylindrical part $(-\infty, \epsilon) \times T_1^* L$ of $X_+^{(0)}$ is an exact symplectic manifold, it contains no non-constant closed J-holomorphic curves. Therefore, u must pass through V.

For the punctured case, a similar argument works. Suppose that the result is false. Then the image of u is contained entirely in $(-\infty, \epsilon) \times T_1^* L$, which is an exact, cylindrical symplectic manifold with positive *and* negative boundary. However, u has no positive punctures, which contradicts the maximum principle. Therefore, u must intersect V. □

Corollary 3.4.5. *Let u be a closed or punctured J-holomorphic curve in $X_+^{(0)}$ that is not contained in the divisor D'. Then,*

$$c_1^{\tau}(u) \geq 1. \tag{3.4.4}$$

Proof. If u is not simple, then by Theorems 2.35 and 6.34 of [54], we may pass to its underlying simple curve $\tilde{u}$. Lemma 3.4.4 ensures that $\tilde{u}$ is Fredholm regular, and so

$$\mathrm{ind}(\tilde{u}) \geq 0.$$

The inequality (3.4.2) then implies that $c_1^{\tau}(\tilde{u}) \geq 1$, so the result follows by linearity of c_1^{τ} with respect to covers. □

3.4.2 Analysis of buildings: limits of $c_1 = 1$ curves

In this section we consider sequences of closed curves in a homology class $E \in H_2(X)$ satisfying $E^2 = -1$ (which is equivalent to $c_1(E) = 1$ for embedded genus zero curves, by the adjunction formula). Recall that in homology, $L = \mathcal{E}_i - \mathcal{E}_j$ for some $1 \leq i \neq j \leq d$. Since the indices i and j won't play a role in what follows, we do away with them and define[26]

$$\begin{aligned} \mathcal{E}^+ &= \mathcal{E}_j, & & \mathcal{G}^+ = F - \mathcal{E}^+, \\ & & \text{and} & \\ \mathcal{E}^- &= \mathcal{E}_i, & & \mathcal{G}^- = F - \mathcal{E}^-. \end{aligned} \tag{3.4.5}$$

The main result of this section, proved in Propositions 3.4.10 and 3.4.16, is:

Proposition 3.4.6. *Let $E \in H_2(X)$ satisfy $c_1(E) = 1$ and suppose we have a sequence $e_k \in \mathcal{M}_{0,0}(X, E; J_k)$ of embedded J-holomorphic curves converging to a limit building F under neck stretching about L. Then $F = (F_+^{(0)}, F_-^{(0)})$ has no symplectisation levels, and moreover $F_-^{(0)}$ is non-empty if, and only if, $E \cdot L = E \cdot (\mathcal{E}^- - \mathcal{E}^+) \neq 0$, in which case $F_-^{(0)}$ consists of a union of planes of the same parity in the moduli spaces $\mathcal{M}_\pm$ in T^*S^2. The parity of the planes coincides with the sign of the intersection $E \cdot L$.*

Furthermore, in the case $E \in \{\mathcal{E}^\pm, \mathcal{G}^\pm\}$, the buildings $F_{\mathcal{E}^\pm}$ arising from limits of curves in the classes $\mathcal{E}^\pm$ have identical top levels. The analogous result holds for $\mathcal{G}^\pm$. We write these relationships as

$$(F_{\mathcal{E}^+})_+^{(0)} = (F_{\mathcal{E}^-})_+^{(0)} \text{ and } (F_{\mathcal{G}^+})_+^{(0)} = (F_{\mathcal{G}^-})_+^{(0)}.$$

On the other hand, the four planes $(F_{\mathcal{E}^\pm})_-^{(0)}$, $(F_{\mathcal{G}^\pm})_-^{(0)}$ in T^*S^2 are pairwise distinct.

Remark 3.4.7. 1. In the case $F_-^{(0)} = \emptyset$ in the above result, this means that $F_+^{(0)}$ is actually a *closed* holomorphic curve in $X_+^{(0)}$.

[26] This notation is chosen so that $\mathcal{E}^\pm \cdot L = \pm 1$. Although, note that $\mathcal{G}^\pm \cdot L = \mp 1$, which might seem confusing, however the author believes that the other definition of $\mathcal{G}^\pm = F - \mathcal{E}^\mp$ required to achieve $\mathcal{G}^\pm \cdot L = \pm 1$ is worse.

2. In a later section, we will show that the two pairs of planes $(F_{\mathcal{E}^\pm})^{(0)}_- \cup (F_{\mathcal{G}^\pm})^{(0)}_-$ are the nodal fibres of a Lefschetz fibration on T^*S^2 such that $L = 0_{S^2}$ is a matching cycle.

Lemma 3.4.8. *Suppose that $F = (F^{(0)}, \dots, F^{(N)})$ is the limiting building of a sequence of curves in the moduli spaces $\mathcal{M}_{0,0}(X, E; J_k)$, where $E \in H_2(X; \mathbb{Z})$ satisfies $c_1(E) = 1$. Then $F^{(0)}_+$ has connected domain,* i.e., *it consists of exactly one component, and each of its punctures is simple.*

Proof. If $F^{(0)}_+$ had more than one component then at least one would satisfy $c_1^\tau \leq 0$, since

$$\sum_{\substack{\text{components } f \\ \text{of } F^{(0)}_+}} c_1^\tau(f) = c_1^\tau(F^{(0)}_+) = c_1([\bar{F}]) = c_1(E) = 1,$$

by Lemma 3.4.2. However, this is impossible by Corollary 3.4.5. Therefore, $F^{(0)}_+$ consists of exactly one component which satisfies $c_1^\tau(F^{(0)}_+) = 1$. Since $F^{(0)}_+$ is Fredholm regular by Lemma 3.4.4, we must have $\operatorname{ind}(F^{(0)}_+) = 0$ by (3.4.2). This implies that $k_- - \sum_{i=1}^{k_-} \operatorname{cov}(\gamma_i^-) = 0$, which implies that the punctures of $F^{(0)}_+$ are simple. ☐

Hence, by Lemma 3.4.1, $F = (F^{(0)}_+, F^{(0)}_-)$ is a holomorphic building with only a top and bottom level. Therefore, by Theorem 3.2.5 and Proposition 3.2.6 we have

$$i_U(F^{(0)}_+, F^{(0)}_+) + i_U(F^{(0)}_-, F^{(0)}_-) = i(F, F) = E^2 = -1. \tag{3.4.6}$$

The adjunction formula implies that $i_U(F^{(0)}_+, F^{(0)}_+) \geq c_N(F^{(0)}_+)$, and so we obtain the inequality

$$i_U(F^{(0)}_-, F^{(0)}_-) \leq -1 - c_N(F^{(0)}_+). \tag{3.4.7}$$

Lemma 3.4.9. *The normal Chern number $c_N(F^{(0)}_+) = -1$, and thus*

$$i_U(F^{(0)}_-, F^{(0)}_-) \leq 0. \tag{3.4.8}$$

Proof. By definition,

$$c_N(u) = c_1^\tau(u^*TW) - \chi(\dot{\Sigma}) \pm \sum_{z \in \Gamma^\pm} \alpha^\tau_\mp(\gamma_z \mp \delta).$$

We have the formulas [52, §3.2]

$$2\alpha^\tau_\pm(\gamma+\epsilon) = \mu^\tau_{\mathrm{CZ}}(\gamma+\epsilon) \pm p(\gamma+\epsilon),$$
$$\mu^\tau_{\mathrm{CZ}}(\gamma\mp\delta) = 2\operatorname{cov}(\gamma) \pm 1,$$

where $p(\gamma+\epsilon) := \alpha^\tau_+(\gamma+\epsilon) - \alpha^\tau_-(\gamma+\epsilon)$ is called the *parity*. From these we obtain

$$\alpha^\tau_\mp(\gamma\mp\delta) = \operatorname{cov}(\gamma).$$

Therefore, for a curve with only simple asymptotics,

$$c_N(u) = c^\tau_1(u^*TW) - (2-k_+-k_-) \pm k_\pm = c^\tau_1(u^*TW) + 2(k_+-1). \tag{3.4.9}$$

In particular, for $u = F^{(0)}_+$ we have $k_+ = 0$ and $c^\tau_1(F^{(0)}_+) = 1$, resulting in $c_N(F^{(0)}_+) = -1$. The inequality (3.4.8) then follows from (3.4.7). □

Proposition 3.4.10. *Let $F = (F^{(0)}_+, F^{(0)}_-)$ be the limit building of a sequence of curves in the moduli spaces $\mathcal{M}_{0,0}(X, E; J_k)$, where $E \in H_2(X;\mathbb{Z})$ satisfies $c_1(E) = 1$ and $E \cdot L = \iota$. Then $F^{(0)}_+$ is a sphere with exactly $|\iota|$ punctures and*

$$F^{(0)}_- \text{ consists of exactly } \begin{cases} \iota \quad P_+ \text{ planes}, & \text{if } \iota \geq 0, \text{ or} \\ |\iota| \quad P_- \text{ planes}, & \text{if } \iota < 0. \end{cases}$$

Corollary 3.4.11. *Suppose that the sequence in the above proposition is given by the unique curves in $\mathcal{M}_{0,0}(X, E_n; J_k)$. Then $F^{(0)}_-$ is empty and $F^{(0)}_+$ is a smooth closed curve in $X^{(0)}_+$.*

Remark 3.4.12. We paraphrase this result by saying that the E_n curve disjoins from L under neck stretching.

Corollary 3.4.13. *Suppose that E is one of the classes in Equation (3.4.5). Then $F^{(0)}_-$ consists of exactly one plane in the moduli spaces $\mathcal{M}_\pm$ of planes in T^*S^2 defined in Section 3.3. The sign of the plane $F^{(0)}_-$ is the same as that of $E \cdot L$.*

Remark 3.4.14. It is important (albeit obvious) to note that the planes in T^*S^2 arising as limits of curves in the classes $\mathcal{E}^\pm$ and $\mathcal{G}^\pm$ have opposite parity. That is, one is in $\mathcal{M}_+$ and the other in $\mathcal{M}_-$.

Proof of Proposition 3.4.10. The genus g of F is 0, and so each component of $F_-^{(0)}$ is a genus 0 punctured curve. Moreover, as $F_+^{(0)}$ has only 1 component, the topology of $F_-^{(0)}$ must be a union of discs — anything else would increase g. Due to the classification of simple J-holomorphic planes of T^*S^2 (Proposition 3.3.2), we deduce that $F_-^{(0)}$ is a union of $P_\pm$ planes. We write this as[27] $F_-^{(0)} = m_+P_+ + m_-P_- := \bigsqcup_{m_+} P_+ \sqcup \bigsqcup_{m_-} P_-$, for some integers $m_\pm \geq 0$. The results of Lemma 3.3.1 and additivity of the Siefring intersection over disjoint unions [50, Proposition 4.3(3)] then imply that

$$i_U(F_-^{(0)}, F_-^{(0)}) = 2m_+m_- \geq 0.$$

Combining this with Lemma 3.4.9 yields that $i_U(F_-^{(0)}, F_-^{(0)}) = 0$, and thus at least one of $m_\pm$ is zero. The equation

$$m_+ - m_- = F_-^{(0)} \cdot L = E \cdot L = \iota$$

then completes the proof. $\square$

Remark 3.4.15. Observe that, through repeated applications of the SFT compactness theorem, we can produce a single neck-stretching sequence $J_k \to J_\infty$ such that, for every $E \in \{E_n, \mathcal{E}^+, \mathcal{E}^-, \mathcal{G}^+, \mathcal{G}^-\}$, the curves in the moduli spaces $\mathcal{M}_{0,0}(E; J_k)$ converge to J^*-holomorphic buildings as $k \to \infty$. Indeed, fix a particular class E and sequences $J_k \to J_\infty$ and $u_k \in \mathcal{M}_{0,0}(E; J_k)$. The compactness theorem produces a subsequence k such that u_k converges. We can then apply SFT compactness to a sequence $v_k \in \mathcal{M}_{0,0}(E'; J_k)$ for some $E' \neq E$ to obtain another subsequence, and so on and so forth.

This idea can be generalised to include sequences of curves in classes other than the exceptional ones used above, provided that the sequences of curves are somehow fixed. For example, we could make a point constraint. When dealing with the classes

[27] Note that the $\sqcup$ notation only means that the domain of $F_-^{(0)}$ is a disjoint union of Riemann surfaces, its image need not be a disjoint union of planes. Moreover, the use of the plus sign in the equation $F_-^{(0)} = m_+P_+ + m_-P_-$ is meant to reflect the fact that the unconstrained intersection product is additive over unions.

E above there was no need to do this since the moduli spaces $\mathcal{M}_{0,0}(E;J)$ are just single points.

We shall use this principle repeatedly in the sequel, although we will not always make reference to it to save cluttering the narrative. Suffice to say that if we need to compare two J^*-holomorphic buildings arising from sequences in the classes $\alpha, \beta \in H_2(X)$, then this will be implicitly done with respect to a single neck-stretching sequence for which both the corresponding sequences of curves converge.

We now turn to deriving relations between the limits of curves with homology classes in $\{\mathcal{E}^+, \mathcal{E}^-, \mathcal{G}^+, \mathcal{G}^-\}$.

Proposition 3.4.16. *The holomorphic limit buildings* $F_{\mathcal{E}^+} = ((F_{\mathcal{E}^+})_+^{(0)}, (F_{\mathcal{E}^+})_-^{(0)})$, $F_{\mathcal{E}^-} = ((F_{\mathcal{E}^-})_+^{(0)}, (F_{\mathcal{E}^-})_-^{(0)})$, $F_{\mathcal{G}^+} = ((F_{\mathcal{G}^+})_+^{(0)}, (F_{\mathcal{G}^+})_-^{(0)})$, *and* $F_{\mathcal{G}^-} = ((F_{\mathcal{G}^-})_+^{(0)}, (F_{\mathcal{G}^-})_-^{(0)})$ *satisfy*

$$(F_{\mathcal{E}^+})_+^{(0)} = (F_{\mathcal{E}^-})_+^{(0)} \qquad \textit{and} \qquad (F_{\mathcal{G}^+})_+^{(0)} = (F_{\mathcal{G}^-})_+^{(0)}.$$

Proof. Corollary 3.4.13 and Proposition 3.2.6 imply that

$$0 = \mathcal{E}^- \cdot \mathcal{E}^+ = i_U((F_{\mathcal{E}^-})_+^{(0)}, (F_{\mathcal{E}^+})_+^{(0)}) + i_U((F_{\mathcal{E}^-})_-^{(0)}, (F_{\mathcal{E}^+})_-^{(0)}) = i_U((F_{\mathcal{E}^-})_+^{(0)}, (F_{\mathcal{E}^+})_+^{(0)}) + 1.$$

Therefore, by positivity of intersections and the fact that

$$c_1^\tau((F_{\mathcal{E}^-})_+^{(0)}) = c_1^\tau((F_{\mathcal{E}^+})_+^{(0)}) = 1,$$

we must have that $(F_{\mathcal{E}^-})_+^{(0)}$ is a reparametrisation of $(F_{\mathcal{E}^+})_+^{(0)}$, that is, $(F_{\mathcal{E}^-})_+^{(0)} = (F_{\mathcal{E}^+})_+^{(0)}$. Similarly, $(F_{\mathcal{G}^-})_+^{(0)} = (F_{\mathcal{G}^+})_+^{(0)}$. □

Recall the so-called Morse-Bott contribution to the intersection number [52, §4]:

$$i^{\pm}_{\mathrm{MB}}(\gamma + \epsilon, \gamma' + \epsilon') := \Omega^\tau_\pm(\gamma, \gamma') - \Omega^\tau_\pm(\gamma + \epsilon, \gamma' + \epsilon').$$

A simple calculation shows that this integer is non-negative and that it satisfies

$$i_U(u, v) = i_C(u, v) + \sum_{(z,w) \in \Gamma^\pm \times (\Gamma')^\pm} i^{\pm}_{\mathrm{MB}}(\gamma_z \mp \delta, \gamma_w \mp \delta).$$

Lemma 3.4.17. *Let γ and γ' denote the asymptotic orbits of $(F_{\mathcal{E}^+})_+^{(0)}$ and $(F_{\mathcal{G}^+})_+^{(0)}$ respectively. Then $\gamma' \neq \gamma$, that is, they are geometrically distinct orbits.*

Proof. Plugging $(u,v) = ((F_{\mathcal{E}^+})_+^{(0)}, (F_{\mathcal{G}^+})_+^{(0)})$ into the above formula yields

$$i_U((F_{\mathcal{E}^+})_+^{(0)}, (F_{\mathcal{G}^+})_+^{(0)}) = i_C((F_{\mathcal{E}^+})_+^{(0)}, (F_{\mathcal{G}^+})_+^{(0)}) + i_{\mathrm{MB}}^-(\gamma+\delta, \gamma'+\delta).$$

Recall that each fibration π_J has two distinguished J-holomorphic sections, one of class S, and the other $S' = H - E_0 - \sum_{j=1}^d \mathcal{E}_j$, and that $\mathcal{E}^+$ intersects S' and not S, and vice versa for $\mathcal{G}^+$. It is then trivial that $(F_{\mathcal{E}^+})_+^{(0)}$ and $(F_{\mathcal{G}^+})_+^{(0)}$ are geometrically distinct, and so positivity of intersections gives $i_C((F_{\mathcal{E}^+})_+^{(0)}, (F_{\mathcal{G}^+})_+^{(0)}) \geq 0$ and thus

$$i_U((F_{\mathcal{E}^+})_+^{(0)}, (F_{\mathcal{G}^+})_+^{(0)}) \geq i_{\mathrm{MB}}^-(\gamma+\delta, \gamma'+\delta).$$

Observe that $[\mathcal{E}^+]\cdot[\mathcal{G}^+] = 1$, and, since $(F_{\mathcal{E}^+})_-^{(0)}$ and $(F_{\mathcal{G}^+})_-^{(0)}$ are planes of opposite parity in T^*S^2, $i_U((F_{\mathcal{E}^+})_-^{(0)}, (F_{\mathcal{G}^+})_-^{(0)}) = 1$. The additivity of the unconstrained intersection product then implies that

$$i_U((F_{\mathcal{E}^+})_+^{(0)}, (F_{\mathcal{G}^+})_+^{(0)}) = 0,$$

Combining this with the above inequality, we obtain $i_{\mathrm{MB}}^-(\gamma+\delta, \gamma'+\delta) = 0$. Therefore, the result follows from the next lemma. □

Lemma 3.4.18. *For a Morse-Bott degenerate asymptote γ satisfying $\dim\ker \mathbf{A}_\gamma = 2$, we have*

$$i_{\mathrm{MB}}^\pm(\gamma^k \mp \delta, \gamma^l \mp \delta) = \min\{l,k\} > 0.$$

Proof. This is a simple consequence of the definitions and the properties of asymptotic operators satisfying $\dim\ker \mathbf{A} = 2$. Indeed,

$$\begin{aligned} i_{\mathrm{MB}}^\pm(\gamma^k \mp \delta, \gamma^l \mp \delta) &= \Omega_\pm^\tau(\gamma^k, \gamma^l) - \Omega_\pm^\tau(\gamma^k \mp \delta, \gamma^l \mp \delta) \\ &= \min\{\mp l\alpha_\mp^\tau(\gamma^k), \mp k\alpha_\mp^\tau(\gamma^l)\} - \min\{\mp l\alpha_\mp^\tau(\gamma^k \mp \delta), \mp k\alpha_\mp^\tau(\gamma^l \mp \delta)\}. \end{aligned}$$

Moreover, by the properties of the extremal winding numbers $\alpha_\pm^\tau$, we have (*c.f.* the proof of Lemma 3.2.8 and Remark 3.2.9)

$$\alpha_\pm^{\tau^k}(\gamma^k) = k\mathrm{wind}_{\mathbf{A}_\gamma}^\tau(0) \pm 1 \qquad \text{and} \qquad \alpha_\pm^{\tau^k}(\gamma^k \pm \delta) = k\mathrm{wind}_{\mathbf{A}_\gamma}^\tau(0).$$

Thus the result follows. □

Since $\gamma' \neq \gamma$, and there is a unique $\mathcal{M}_\pm$ plane asymptotic to each orbit, we obtain the following:

Corollary 3.4.19. *The planes* $\{(F_{\mathcal{E}^+})^{(0)}_-, (F_{\mathcal{G}^+})^{(0)}_-, (F_{\mathcal{E}^-})^{(0)}_-, (F_{\mathcal{G}^-})^{(0)}_-\}$ *are pairwise geometrically distinct. Specifically, the planes with common parity* $(F_{\mathcal{E}^+})^{(0)}_-, (F_{\mathcal{G}^-})^{(0)}_- \in \mathcal{M}_+$ *and* $(F_{\mathcal{G}^+})^{(0)}_-, (F_{\mathcal{E}^-})^{(0)}_- \in \mathcal{M}_-$ *are distinct.*

3.4.3 Controlling the asymptotic orbits

The aim of this section is to show that the almost complex structure on X that undergoes neck stretching can be chosen so that the top level of the limiting building of the sequence of $\mathcal{G}^+$ curves, $(F_{\mathcal{G}^+})^{(0)}_+$, has asymptotic Reeb orbit γ' equal to the image of the asymptote γ of $(F_{\mathcal{E}^+})^{(0)}_+$ under fibre-wise multiplication by -1. This is the property that will power our argument that the $J^{(0)}_-$-holomorphic cylinders in T^*S^2 (obtained via neck stretching in Section 3.4.4) that intersect the zero section do so along circles.

As mentioned in Section 3.1.1, fibre-wise multiplication by -1 in T^*S^2 corresponds to complex conjugation under the $SO(3)$-equivariant isomorphism $T^*S^2 \cong (z_1^2 + z_2^2 + z_3^2 = 1)$. We say these asymptotes are conjugate and write this as $\gamma' = \overline{\gamma}$. We first make an explicit construction to force the top level curves $(F_{\mathcal{E}^+})^{(0)}_+$ and $(F_{\mathcal{G}^+})^{(0)}_+$ to have conjugate asymptotics. Recall the almost-Kähler structure $(T^*S^2, \omega_{\text{can}}, J_{T^*S^2})$, where $J_{T^*S^2}$ is any of the compatible, cylindrical almost complex structures constructed in Section 3.1.1, and that the natural $SO(3)$ action by rotations is by exact almost-Kähler isometries. Note in particular that Section 3.1.1 allows us to produce a compatible almost complex structure that is cylindrical outside of an arbitrarily small neighbourhood of the zero section. We'll use this property to alter $J^{(0)}_+$ in the neck region of $X^{(0)}_+$.

Denote by M the unit cotangent bundle $T_1^*S^2$. Since both conjugation and the Hamiltonian $SO(3)$ action preserve the length function $|p|$ on T^*S^2, they restrict to $\mathbb{R}$-equivariant maps on the symplectisation $\mathbb{R} \times M$. Moreover, the $SO(3)$ action is by exact symplectomorphisms, so it descends to give an action on the quotient

M/S^1 of the unit cotangent bundle by the Reeb flow. Recall that $M/S^1 \cong S^2$ and, under this identification, the action of $SO(3)$ is the usual one by rotations.

We introduce some notation for convenience: as $X_+^{(0)}$ is a manifold with the cylindrical end $(-\infty, \epsilon) \times M$, it makes sense to define the subsets, for all $r_0 < \epsilon$,

$$X_{\leq r_0} := (-\infty, r_0] \times M, \text{ and,}$$
$$X_{\geq r_0} := X_+^{(0)} \backslash (-\infty, r_0) \times M.$$

Since $(F_{\mathcal{E}^+})_+^{(0)}$ and $(F_{\mathcal{G}^+})_+^{(0)}$ are asymptotic to the distinct Reeb orbits $\gamma, \gamma' \in M/S^1$, there exists a number $r_1 \ll 0$ and an open neighbourhood (which is a lift via the composition $X_{\leq r_1} \to M \to M/S^1$ of a neighbourhood of $\gamma \in M/S^1$) $U \subset X_{\leq r_1}$ of the half cylinder

$$C_{\mathcal{E}^+} := \mathrm{im}(F_{\mathcal{E}^+})_+^{(0)} \cap X_{\leq r_1}$$

such that $\mathrm{im}(F_{\mathcal{G}^+})_+^{(0)}$ is disjoint from U. We are now ready to state and prove the main result of this section.

Lemma 3.4.20. *There exists a compatible almost complex structure J on $X_+^{(0)}$ and a number $r_2 < r_1$ that satisfy the following:*

1. *J agrees with $J_+^{(0)}$ on the complement of $[r_2, r_1] \times M$,* i.e. *on $X_{\geq r_1} \cup X_{\leq r_2}$, and on the neighbourhood U of $C_{\mathcal{E}^+}$. In particular, J is cylindrical on $X_{\leq r_2}$.*

2. *there exists a J-holomorphic plane asymptotic to $\bar{\gamma}$ in the same relative homology class as $(F_{\mathcal{G}^+})_+^{(0)}$.*

Proof. The half-cylinder $C_{\mathcal{E}^+}$ is asymptotic to γ, whilst $C_{\mathcal{G}^+} := \mathrm{im}(F_{\mathcal{G}^+})_+^{(0)} \cap X_{\leq r_1}$ is asymptotic to $\gamma' \neq \gamma$. The asymptotic convergence of $C_{\mathcal{E}^+}$ and $C_{\mathcal{G}^+}$ implies that their projections to the manifold of Reeb orbits M/S^1 yield disjoint closed neighbourhoods N_γ and $N_{\gamma'}$ of $\gamma, \gamma' \in M/S^1$. Choose a path $\gamma_t \in M/S^1$ disjoint from N_γ such that $\gamma_0 = \gamma'$ and $\gamma_1 = \bar{\gamma}$ and pick a corresponding 1-parameter subgroup $R_t \in SO(3)$ such that $R_t(\gamma') = \gamma_t$. Note that the asymptotic convergence of $C_{\mathcal{E}^+}$ and $C_{\mathcal{G}^+}$ ensures that $r_1 \ll 0$ can be chosen large enough to ensure that $R_t(N_{\gamma'})$ is disjoint from N_γ for all $t \in [0, 1]$.

Now choose $r_2 < r_1$ and a diffeomorphism $\psi : [r_2, r_1] \to [0, 1]$ such that $\psi' \equiv -1$ near the ends. Consider the diffeomorphism

$$R : X_{\leq r_1} \to X_{\leq r_1} : R(r, x) = \begin{cases} (r, R_{\psi(r)}(x)), & \text{if } r \in [r_2, r_1], \\ (r, R_1(x)), & \text{if } r < r_2, \end{cases}$$

and define the *symplectic* half-cylinder $C'_{\mathcal{G}^+} := R(C_{\mathcal{G}^+})$. Note that $C'_{\mathcal{G}^+}$ is disjoint from the neighbourhood U of $C_{\mathcal{E}^+}$. We glue $C'_{\mathcal{G}^+}$ to $\operatorname{im}(F_{\mathcal{G}^+})^{(0)}_+ \cap X_{\geq r_1}$ to obtain a symplectic plane $P_{\mathcal{G}^+}$ asymptotic to $\bar{\gamma}$. Linear interpolation from R to the identity map yields a (relative) homology between $\operatorname{im}(F_{\mathcal{G}^+})^{(0)}_+$ and $P_{\mathcal{G}^+}$.

Since $C'_{\mathcal{G}^+}$ is disjoint from U, we may choose a compatible almost complex structure J on $[r_2, r_1] \times M$ that makes $C'_{\mathcal{G}^+}$ holomorphic, and agrees with $J^{(0)}_+$ on U. Moreover, it can be chosen to agree with $J^{(0)}_+$ on $\{r_1\} \times M \cup \{r_2\} \times M$ since $R_0 = \mathrm{id}$ and R_1 is holomorphic with respect to $J^{(0)}_+$. Extending J by $J^{(0)}_+$ on the complement of $[r_2, r_1] \times M$ completes the construction. $\square$

Since $(X^{(0)}_+, \omega^+)$ is symplectomorphic to $(X \backslash L, \omega)$, we can push forward our new almost complex structure J to $X \backslash L$ and perform neck stretching to this new J. What this amounts to is stretching the neck around the contact-type hypersurface $T^*_{e^{r_2}} S^2 \subset X$. To make this completely precise, we need to extend J over L. However, this is easily achieved by altering the bump function used in the construction of $J_{T^*S^2}$ in Section 3.1.1. The result is a new compatible almost complex structure on X that is adapted to the hypersurface $T^*_{e^{r_2}} S^2$, is equal to our original almost complex structure on the set $X_{\geq r_1}$, and, under stretching the neck, converges to the almost complex structure of Lemma 3.4.20 on $X^{(0)}_+$. To avoid cluttering the notation, we shall continue to refer to the new split almost complex structure obtained on $X^* = (X^{(0)}, \ldots, X^{(N)})$ as J^*, as we will no longer make reference to the old one.

Remark 3.4.21. Note that the new stretched almost complex structure obtained on $X^{(0)}_- = T^*L$ may be different to the original one, since we have *increased* the portion of the neck $\mathbb{R} \times T^*_1 S^2 \cong T^*L \backslash L$ on which $J^{(0)}_-$ is cylindrical. This amounts to pushing forward the old $J^{(0)}_-$ by the Liouville flow for some negative time.

Ultimately, though, this is of no importance, since the results of Section 3.3 are valid for whatever cylindrical almost complex structure we choose on T^*S^2.

Since we've changed our neck stretching sequence J_k, morally we need to redo the limit analysis of the previous section. Indeed, recall the open set $V = X_+^{(0)} \setminus (((-\infty, \epsilon] \times T_1^* L) \cup D')$ from Section 3.4.1 in which we perturbed to ensure genericity. The almost complex structure $J_+^{(0)}$ constructed here satisfies a form of symmetry, which perturbing would destroy. Therefore, *a priori*, $J_+^{(0)}$-holomorphic curves don't satisfy the Fredholm regularity result of Lemma 3.4.4.

Note that we *can* perturb $J_+^{(0)}$ on the complement of the images of $(F_{\mathcal{E}^+})_+^{(0)}$ and $P_{\mathcal{G}^+}$ in V. Moreover, since a punctured $J_+^{(0)}$-holomorphic curve in $X_+^{(0)}$ either passes through this complement, or has identical image to one of $(F_{\mathcal{E}^+})_+^{(0)}$ or $P_{\mathcal{G}^+}$, we can apply Lemma 3.4.4 to every punctured curve in $X_+^{(0)}$ except $(F_{\mathcal{E}^+})_+^{(0)}$ and $P_{\mathcal{G}^+}$. Thus, if we can show that $(F_{\mathcal{E}^+})_+^{(0)}$ and $P_{\mathcal{G}^+}$ are Fredholm regular, then we can apply Lemma 3.4.4 exactly as in Section 3.4.2.

As a result, a convergent sequence of curves $e_k \in \mathcal{M}_{0,0}(\mathcal{E}^+; J_k)$ converges to $(F_{\mathcal{E}^+})_+^{(0)}$ in the top level, and similarly, for curves in $\mathcal{M}_{0,0}(\mathcal{G}^+; J_k)$, we obtain that they converge to the constructed plane $P_{\mathcal{G}^+}$.

Lemma 3.4.22. *The $J_+^{(0)}$-holomorphic planes $(F_{\mathcal{E}^+})_+^{(0)}$ and $P_{\mathcal{G}^+}$ are Fredholm regular.*

Proof. The case of $(F_{\mathcal{E}^+})_+^{(0)}$ is trivial since $J_+^{(0)}$ was unchanged in a neighbourhood of this curve. For $P_{\mathcal{G}^+}$ we apply automatic transversality [52, Theorem 1]. Since $P_{\mathcal{G}^+}$ is embedded, the automatic transversality condition is:

$$\operatorname{ind}(P_{\mathcal{G}^+}) > c_N(P_{\mathcal{G}^+}).$$

We compute that $\operatorname{ind}(P_{\mathcal{G}^+}) = 0$, which follows from Equation (3.4.1), Lemma 3.4.2, and that $P_{\mathcal{G}^+}$ has a single simply covered asymptote. On the other hand,

$$c_N(P_{\mathcal{G}^+}) = -1$$

follows as in Lemma 3.4.9. □

Proposition 3.4.23. *Let J_k be a neck stretching sequence converging to J^*, and $e_k \in \mathcal{M}_{0,0}(\mathcal{E}^+; J_k)$ a corresponding sequence converging to a J^*-holomorphic building $F = (F^{(0)}, \dots, F^{(N)})$. Then the top level $F_+^{(0)}$ consists only of the plane $(F_{\mathcal{E}^+})_+^{(0)}$ obtained in Proposition 3.4.16, there are no symplectisation levels, and the bottom level consists of exactly one plane in $\mathcal{M}_+$.*

Similarly, a convergent sequence $g_k \in \mathcal{M}_{0,0}(\mathcal{G}^+; J_k)$ converges to a building $G = (G^{(0)}, \dots, G^{(N)})$ whose top level $G_+^{(0)}$ consists only on the plane $P_{\mathcal{G}^+}$ constructed in Lemma 3.4.20. As above, the bottom level then consists of exactly one plane in $\mathcal{M}_-$. In particular, $F_+^{(0)}$ and $G_+^{(0)}$ have conjugate asymptotes.

Proof. Since $J_+^{(0)}$ is unchanged in a neighbourhood of $(F_{\mathcal{E}^+})_+^{(0)}$, there is nothing to check for the claimed convergence $e_k \to F$. We prove that $G_+^{(0)} = P_{\mathcal{G}^+}$. Note that $P_{\mathcal{G}^+}$ can be completed into a J^*-holomorphic building G' by adding in the unique $\mathcal{M}_-$ plane in T^*S^2 with asymptotic orbit $\bar{\gamma}$. Then, both G and G' are J^*-holomorphic buildings of the same homology class: $[\bar{G}] = \mathcal{G}^+ = [\bar{G}']$. Thus, they must intersect somewhere, since $(\mathcal{G}^+)^2 = -1$. If $G_+^{(0)}$ is geometrically distinct from $(G')_+^{(0)} = P_{\mathcal{G}^+}$, then positivity of intersections would force there to be a negative intersection in the bottom level, contradicting Corollary 3.3.3. Hence, $G_+^{(0)} = P_{\mathcal{G}^+}$ and the result follows. □

We have constructed two J^*-holomorphic buildings F and G in the homology classes $\mathcal{E}^+$ and $\mathcal{G}^+$, whose asymptotic orbits are conjugate. To free up the notation, we shall continue to denote these buildings by $F_{\mathcal{E}^+}$ and $F_{\mathcal{G}^+}$ respectively. Similarly we obtain the J^*-holomorphic buildings $F_{\mathcal{E}^-}$ and $F_{\mathcal{G}^-}$, which satisfy

$$(F_{\mathcal{E}^-})_+^{(0)} = (F_{\mathcal{E}^+})_+^{(0)} \text{ and } (F_{\mathcal{G}^-})_+^{(0)} = (F_{\mathcal{G}^+})_+^{(0)},$$

by Proposition 3.4.16, and thus they have the same pair of conjugate asymptotes $\{\gamma, \bar{\gamma}\}$.

Corollary 3.4.24. *The unique point of intersection of the planes $(F_{\mathcal{E}^\pm})_-^{(0)} \in \mathcal{M}_\pm$ and $(F_{\mathcal{G}^\pm})_-^{(0)} \in \mathcal{M}_\mp$ lies on the zero section.*

Proof. Conjugation sends $\mathcal{M}_+$ planes to $\mathcal{M}_-$ and vice versa. Therefore, since there is a unique $\mathcal{M}_\pm$ plane asymptotic to each simple Reeb orbit and the asymptotics of $(F_{\mathcal{E}^\pm})^{(0)}_-$ and $(F_{\mathcal{G}^\pm})^{(0)}_-$ are conjugate, we have that

$$\overline{(F_{\mathcal{E}^\pm})^{(0)}_-} = (F_{\mathcal{G}^\pm})^{(0)}_-.$$

Moreover, the zero section is the fixed locus of conjugation, so, since each plane intersects the zero section, the claim follows. □

Remark 3.4.25. This result partially justifies the claim made in Remark 3.4.7 that the bottom levels of the buildings $F_{\mathcal{E}^\pm} \cup F_{\mathcal{G}^\pm}$ form the nodal fibres of a Lefschetz fibration on T^*L such that L is a matching cycle.

3.4.4 Analysis of buildings: limits of $c_1 = 2$ curves

In this section we prove (Proposition 3.4.30) that limits of curves in the class of a fibre $H - S \in H_2(X)$ have bottom level (if non-empty) either a smooth cylinder, or one of the pairs of nodal planes $(F_{\mathcal{E}^\pm})^{(0)}_- + (F_{\mathcal{G}^\pm})^{(0)}_-$. First, we need a basic result on unconstrained intersections of trivial cylinders in symplectisations.[28]

Lemma 3.4.26. *Let γ be a simply covered Morse-Bott degenerate Reeb orbit living in a positive dimensional orbifold N of unparametrised orbits of a contact manifold M. Then the unconstrained self intersection of the corresponding trivial cylinder $u_\gamma : \mathbb{R} \times S^1 \to \mathbb{R} \times M$ is zero:*

$$i_U(u_\gamma, u_\gamma) = 0.$$

Proof. Since $\dim(N) > 0$, we can homotope u_γ to $u_{\tilde{\gamma}}$ for some $\tilde{\gamma} \neq \gamma$ in the same family N. Since i_U is homotopy invariant, we obtain the result, as u_γ and $u_{\tilde{\gamma}}$ are disjoint and asymptotic to distinct orbits. □

[28]Compare with the non-degenerate case where

$$u_\gamma * u_\gamma = \begin{cases} -1, & \text{if } \gamma \text{ is odd} \\ 0, & \text{if } \gamma \text{ is even.} \end{cases}$$

Remark 3.4.27. In our situation, all simple Reeb orbits live in the same manifold $N \cong S^2$ of unparametrised orbits, and so the above result combined with Proposition 3.2.6(2) yields, for any integers $k, l > 0$ and simple orbits γ and γ',

$$i_U(u_{\gamma^k}, u_{(\gamma')^l}) = 0.$$

From this and positivity of intersections (Theorem 3.2.10), we deduce that any curve in a symplectisation level intersects a trivial cylinder non-negatively.

Let $F = (F^{(0)}, \dots, F^{(N)})$ denote a J^*-holomorphic building arising as the limit of a sequence of fibre curves. If the bottom level $F_-^{(0)}$ is non-empty, the following lemma shows that its top-level $F_+^{(0)}$ has to be the disjoint union of the limits that arose in Proposition 3.4.23.

Lemma 3.4.28. *Let $F = (F^{(0)}, \dots, F^{(N)})$ be the limit of a sequence of curves in the class $[\bar{F}] = H - S \in H_2(X)$ of a fibre. Suppose that $F_-^{(0)} \neq \emptyset$. Then $F_+^{(0)}$ is the disjoint union of the planes $(F_{\mathcal{E}^+})_+^{(0)}$ and $(F_{\mathcal{G}^+})_+^{(0)}$, where $F_{\mathcal{E}^+}$ and $F_{\mathcal{G}^+}$ are the holomorphic buildings that arose in section 3.4.3. That is, $F_+^{(0)} = (F_{\mathcal{E}^+})_+^{(0)} + (F_{\mathcal{G}^+})_+^{(0)} = (F_{\mathcal{E}^+})_+^{(0)} \sqcup (F_{\mathcal{G}^+})_+^{(0)}$ is the disjoint union of two planes with conjugate asymptotics $\{\gamma, \bar{\gamma}\}$.*

Proof. Since $F_-^{(0)} \neq \emptyset$, we can use Proposition 3.3.2 to analyse how $F_-^{(0)}$ intersects the planes in $\mathcal{M}_\pm$. Note that $F_-^{(0)}$ cannot possibly be composed exclusively of (covers of) planes in only one of the families $\mathcal{M}_+$ or $\mathcal{M}_-$, as otherwise the intersection of $F_-^{(0)}$ with L would be non-zero, contradicting $[\bar{F}] \cdot L = 0$. Thus, $F_-^{(0)}$ intersects both $(F_{\mathcal{E}^+})_-^{(0)} \in \mathcal{M}_+$ and $(F_{\mathcal{G}^+})_-^{(0)} \in \mathcal{M}_-$ positively. That is

$$i_U(F_-^{(0)}, (F_{\mathcal{E}^+})_-^{(0)}) > 0 \qquad \text{and} \qquad i_U(F_-^{(0)}, (F_{\mathcal{G}^+})_-^{(0)}) > 0.$$

Now let $\digamma = (\digamma_+^{(0)}, \digamma_-^{(0)})$ denote either of the buildings $F_{\mathcal{E}^+}$ or $F_{\mathcal{G}^+}$. *A priori*, the building F may have non-trivial symplectisation levels, so to make sense of the intersection number of the buildings $i(F, \digamma)$ we may need to extend $\digamma$ by *trivial* symplectisation levels.[29] With this understood, Theorem 3.2.5, Proposition 3.2.6,

[29]See Appendix C.5 of [56] for a full discussion of why this operation is well defined.

$i_U(F_-^{(0)}, \boldsymbol{F}_-^{(0)}) > 0$ and $[\bar{F}] \cdot [\bar{\boldsymbol{F}}] = 0$ imply that

$$i_U(F_+^{(0)}, \boldsymbol{F}_+^{(0)}) + \sum_{i=1}^{N} i_U(F^{(i)}, \boldsymbol{F}^{(i)}) < 0.$$

Furthermore, Remark 3.4.27 implies that, for each $1 \leq i \leq N$, $i_U(F^{(i)}, \boldsymbol{F}^{(i)}) \geq 0$, and thus $i_U(F_+^{(0)}, \boldsymbol{F}_+^{(0)}) < 0$. Therefore, by positivity of intersections, at least one component of $F_+^{(0)}$ covers $\boldsymbol{F}_+^{(0)}$.

Summarising, we have proved that some components of $F_+^{(0)}$ cover both $(F_{\mathcal{E}^+})_+^{(0)}$ *and* $(F_{\mathcal{G}^+})_+^{(0)}$, and so there exist positive integers $m, n > 0$ such that

$$F_+^{(0)} = m(F_{\mathcal{E}^+})_+^{(0)} + n(F_{\mathcal{G}^+})_+^{(0)} + (\text{other terms}).$$

However, this already exhausts the total ω^+-area. Indeed, by Corollary 2.11 of [9], we have that

$$\begin{aligned}
\omega([\bar{F}]) &= \int_{\Sigma_+^{(0)}} (F_+^{(0)})^* \omega^+ \\
&\geq m \int_{\mathbb{C}} ((F_{\mathcal{E}^+})_+^{(0)})^* \omega^+ + n \int_{\mathbb{C}} ((F_{\mathcal{G}^+})_+^{(0)})^* \omega^+ \\
&= m\omega([\mathcal{E}^+]) + n\omega([\bar{F}] - [\mathcal{E}^+]) \\
&= n\omega([\bar{F}]) + (m-n)\omega([\mathcal{E}^+]) \\
&= m\omega([\bar{F}]) + (n-m)(\omega([\bar{F}]) - \omega([\mathcal{E}^+])).
\end{aligned}$$

Recall that Lemma 2.2.1 states that $\omega([\bar{F}]) \geq \omega([\mathcal{E}^+]) = l$, and so, in either of the cases $n \geq m$, or $m \geq n$, we obtain

$$\omega([\bar{F}]) \geq m\omega([\mathcal{E}^+]) + n\omega([\bar{F}] - [\mathcal{E}^+]) \geq \omega([\bar{F}]),$$

and thus $n = m = 1$. Moreover, this implies that there can be no other terms in the expression for $F_+^{(0)}$. That is, $F_+^{(0)} = (F_{\mathcal{E}^+})_+^{(0)} + (F_{\mathcal{G}^+})_+^{(0)}$, and so the result is proved. □

Lemma 3.4.29. *Let $F = (F^{(0)}, \ldots, F^{(N)})$ be the limit of a sequence of curves in the class of a fibre $[\bar{F}] = H - S \in H_2(X)$ such that $F_-^{(0)} \neq \emptyset$. Then $N = 0$, that is, $F = (F_+^{(0)}, F_-^{(0)})$ consists only of a top and bottom level.*

Proof. Lemma 3.4.28 shows that $F_+^{(0)} = (F_{\mathcal{E}^+})_+^{(0)} + (F_{\mathcal{G}^+})_+^{(0)}$ and so any non-trivial component f of $F^{(N)}$ (the level adjacent to the top level $F_+^{(0)}$) must have exactly two positive punctures. Indeed, if it only had one, then the argument of Lemma 3.4.1 would imply that it is a trivial cylinder, and thus $F^{(N)}$ would be a union of trivial cylinders, which is ruled out by the stability condition. Therefore, f has exactly two positive punctures: γ and $\bar{\gamma}$ of $(F_{\mathcal{E}^+})_+^{(0)}$ and $(F_{\mathcal{G}^+})_+^{(0)}$.

Since $F_+^{(0)} = (F_{\mathcal{E}^+})_+^{(0)} + (F_{\mathcal{G}^+})_+^{(0)}$, we have that $i_U(F_+^{(0)}, (F_{\mathcal{E}^+})_+^{(0)}) = -1$, and so, to balance the equation $i(F, \mathcal{E}^+) = [\bar{F}] \cdot [\mathcal{E}^+] = 0$, there must be a positive intersection between F and $F_{\mathcal{E}^+}$ in a different level. The previous paragraph implies that this positive intersection is eaten up by the non-trivial level $F^{(N)}$, which then forces $F_-^{(0)}$ to take on an illegal form. Indeed, extend $F_{\mathcal{E}^+}$ by trivial intermediate levels, then, as $(F_{\mathcal{E}^+})^{(N)}$ and f are geometrically distinct, their intersection number is bounded below by the Morse-Bott contributions from their asymptotic orbits, which are computed in Lemma 3.4.18:

$$i_U(f, (F_{\mathcal{E}^+})^{(N)}) \geq i_{\mathrm{MB}}^+(\gamma, \gamma) = 1.$$

Combining this with Remark 3.4.27, which says that any further intersections appearing in intermediate levels are non-negative, we obtain that

$$\begin{aligned} 0 = i(F, \mathcal{E}^+) &\geq i_U(F_+^{(0)}, (F_{\mathcal{E}^+})_+^{(0)}) + i_U(f, (F_{\mathcal{E}^+})^{(N)}) + i_U(F_-^{(0)}, (F_{\mathcal{E}^+})_-^{(0)}) \\ &\geq i_U(F_+^{(0)}, (F_{\mathcal{E}^+})_+^{(0)}) + 1 + i_U(F_-^{(0)}, (F_{\mathcal{E}^+})_-^{(0)}) \\ &= i_U(F_-^{(0)}, (F_{\mathcal{E}^+})_-^{(0)}), \end{aligned}$$

and therefore $i_U(F_-^{(0)}, (F_{\mathcal{E}^+})_-^{(0)}) = 0$. Since $F_-^{(0)} \neq \emptyset$, Lemma 3.3.1 implies that $F_-^{(0)}$ must consist of covers of J-holomorphic planes of the same parity as $(F_{\mathcal{E}^+})_-^{(0)}$, which contradicts $[\bar{F}] \cdot L = 0$. Hence, there are no non-trivial symplectisation levels, and so the result is proved. □

The following is the main result of the analysis of $c_1 = 2$ buildings.

Proposition 3.4.30. *Let F be a limiting building of a sequence of curves in the class $[\bar{F}] = H - S \in H_2(X)$ of a fibre such that $F_-^{(0)} \neq \emptyset$. Then $F = (F_+^{(0)}, F_-^{(0)})$ where $F_+^{(0)} = (F_{\mathcal{E}^+})_+^{(0)} + (F_{\mathcal{G}^+})_+^{(0)}$ and $F_-^{(0)}$ is either*

1. *one of the pairs of nodal curves $(F_{\mathcal{E}^+})_-^{(0)} + (F_{\mathcal{G}^+})_-^{(0)}$ or $(F_{\mathcal{E}^-})_-^{(0)} + (F_{\mathcal{G}^-})_-^{(0)}$, with $F_{\mathcal{E}^+}, F_{\mathcal{G}^+}, F_{\mathcal{E}^-}, F_{\mathcal{G}^-}$ as in section 3.4.3; or,*
2. *a smooth cylinder with the same asymptotes $\{\gamma, \bar{\gamma}\}$ as $(F_{\mathcal{E}^+})_+^{(0)} + (F_{\mathcal{G}^+})_+^{(0)}$.*

Proof. The previous results show that F is a building with only a top and bottom level, and that $F_+^{(0)} = (F_{\mathcal{E}^+})_+^{(0)} + (F_{\mathcal{G}^+})_+^{(0)}$. This tells us that the positive asymptotics of $F_-^{(0)}$ are exactly γ and $\bar{\gamma}$ of $(F_{\mathcal{E}^+})_+^{(0)}$ and $(F_{\mathcal{G}^+})_+^{(0)}$. Thus, $F_-^{(0)}$ is a connected genus 0 holomorphic curve with exactly two positive punctures. Therefore, either it is a smooth cylinder, or it has nodes forming a chain of closed spheres connecting two planes. However, there are no closed holomorphic spheres in T^*L since this is an exact symplectic manifold, and so there can be at most one node.

Observe that a nodal pair $P_1 + P_2$ of J-holomorphic planes satisfying $(P_1 + P_2) \cdot L = 0$ and having asymptotic orbits $\{\gamma, \bar{\gamma}\}$ must be exactly one of the pairs $(F_{\mathcal{E}^+})_-^{(0)} + (F_{\mathcal{G}^+})_-^{(0)}$ or $(F_{\mathcal{E}^-})_-^{(0)} + (F_{\mathcal{G}^-})_-^{(0)}$. Therefore, either we are in case (1), or $F_-^{(0)}$ passes through a point in T^*L not contained in the images of the curves $(F_{\mathcal{E}^+})_-^{(0)}, (F_{\mathcal{G}^+})_-^{(0)}, (F_{\mathcal{E}^-})_-^{(0)}$, or $(F_{\mathcal{G}^-})_-^{(0)}$, and is thus a smooth cylinder. □

3.5 Constructing the foliation

In this section we use the limit analysis of section 3.4 to construct a $J_{T^*S^2}$-holomorphic foliation of T^*S^2 by cylinders. The process also picks out a particular neck stretching sequence of almost complex structures J_k such that all the curves of interest converge. First, suppose that we have a neck stretching sequence J_k, and a countable collection of sequences (f_k^m) of J_k-holomorphic curves with uniformly bounded energy. One then applies a diagonal argument, a countable generalisation of that discussed in Remark 3.4.15, using the SFT compactness theorem to extract a subsequence J_k such that, for all m, the sequences (f_k^m) converge to J^*-holomorphic buildings as $k \to \infty$.

We apply this process to the following sequences. Recall the -1-classes $\{\mathcal{E}_j \mid 1 \leq j \leq d\}$ defined in Lemma 2.2.1, where classes of the form $\mathcal{E}_i - \mathcal{E}_j$ support the Lagrangian spheres in $B_{d,p,q}$. The critical points of the Lefschetz fibrations $\pi_{J_k} : X_{J_k} \to \mathbb{C}$ correspond exactly to the unique intersection points of $\mathcal{E}_j \cdot (F - \mathcal{E}_j)$. This allows us to partition the set of critical points into so-called *relevant* and *irrelevant* sets determined by whether the intersection $\mathcal{E}_j \cdot L$ is non-zero or not. That is, the relevant critical points correspond to the classes $\{\mathcal{E}^+, \mathcal{E}^-\}$ as defined in (3.4.5), and the irrelevant ones correspond to the remaining $\mathcal{E}_j$ classes. Recall also the class E_n, which represents (the underlying simple curve of) a component of the exotic curve $\mathbf{u}_\infty^J$ (see Corollary 2.4.16). Even though curves in this class do not correspond to Lefschetz critical points, we shall also call them irrelevant. With this in mind, for each $1 \leq j \leq d$, take the unique sequences of J_k-holomorphic -1-curves

$$
\begin{aligned}
e_k^j &\in \mathcal{M}_{0,0}(\mathcal{E}_j; J_k), \\
g_k^j &\in \mathcal{M}_{0,0}(F - \mathcal{E}_j; J_k), \\
\varepsilon_k &\in \mathcal{M}_{0,0}(E_n; J_k),
\end{aligned}
$$

along with a countable number of sequences of fibre curves

$$
f_k^m \in \mathcal{M}_{0,0}(F; J_k)
$$

determined by point constraints for some fixed dense set $\{x^m \in Y_-\}$ in the Weinstein neighbourhood Y_- of L. We extract a subsequence J_k for which all these sequences converge to the J^*-holomorphic buildings F^m.

Note that Proposition 3.4.10 implies that, for sufficiently large k, the images of the *irrelevant* curves stay bounded away from L. Moreover, a further diagonal argument ensures that the convergence of all curves is monotonic, which will be useful in Section 4.1.

To upgrade the dense set $\{(F^m)_-^{(0)}\}$ of $J_{T^*S^2}$-holomorphic curves in T^*S^2 to a foliation we apply a bubbling argument, which is inspired by [29, §6].[30] The idea is choose $x \in T^*S^2$ and to take a subsequence of (x^m) converging to x and analyse the corresponding limits of the curves $(F^m)_-^{(0)}$ under SFT compactness. However, this time the SFT compactness theorem is that relating to manifolds with cylindrical ends [5, Theorem 10.2]. To that end, denote the moduli space[31] of $J_-^{(0)}$-holomorphic cylinders with fixed positive asymptotes $\{\gamma, \bar{\gamma}\}$ by $\mathcal{M}_{\mathrm{cyl}}$. We first compute the self-intersection of a cylinder in $\mathcal{M}_{\mathrm{cyl}}$:

Lemma 3.5.1. *The $J_{T^*S^2}$-holomorphic cylinders in $\mathcal{M}_{\mathrm{cyl}}$ have zero* constrained *self-intersection. In other words, for a building in the class $[\bar{F}] = H - S$ arising from Proposition 3.4.30 we have*

$$i_U(F_-^{(0)}, F_-^{(0)}) = 2 \qquad \textit{and} \qquad i_C(F_-^{(0)}, F_-^{(0)}) = 0.$$

Proof. Since $F_+^{(0)} = (F_{\mathcal{E}^+})_+^{(0)} + (F_{\mathcal{G}^+})_+^{(0)}$ and

$$i_U((F_{\mathcal{E}^+})_+^{(0)}, (F_{\mathcal{E}^+})_+^{(0)}) = -1 = i_U((F_{\mathcal{G}^+})_+^{(0)}, (F_{\mathcal{G}^+})_+^{(0)}),$$

we have that

$$\begin{aligned} 0 &= [\bar{F}]^2 \\ &= i_U((F_{\mathcal{E}^+})_+^{(0)}, (F_{\mathcal{E}^+})_+^{(0)}) + i_U((F_{\mathcal{G}^+})_+^{(0)}, (F_{\mathcal{G}^+})_+^{(0)}) + i_U(F_-^{(0)}, F_-^{(0)}) \\ &= i_U(F_-^{(0)}, F_-^{(0)}) - 2, \end{aligned}$$

and so, $i_U(F_-^{(0)}, F_-^{(0)}) = 2$. Then, since

$$i_U(F_-^{(0)}, F_-^{(0)}) = i_C(F_-^{(0)}, F_-^{(0)}) + i_{\mathrm{MB}}^+(\gamma, \gamma) + i_{\mathrm{MB}}^+(\bar{\gamma}, \bar{\gamma}) = i_C(F_-^{(0)}, F_-^{(0)}) + 2,$$

we obtain the result. □

[30] The argument of [29, §6] is much more complicated than what is required for our situation.

[31] More precisely, this is the moduli space of genus 0 curves with exactly two positive asymptotes that represent the homology class

$$(0, 1, 1) \in H_2(T^*S^2, \gamma \sqcup \bar{\gamma}) \cong H_2(T^*S^2) \oplus H_1(\gamma \sqcup \bar{\gamma}) \cong \langle 0_{S^2}, \gamma, \bar{\gamma} \rangle \cong \mathbb{Z}^3.$$

Applying the SFT compactness theorem to sequences of curves in $\mathcal{M}_{\text{cyl}}$ yields holomorphic buildings of height $k_-|1|k_+$, as in Section 8 of [5]. Since T^*S^2 is a manifold with no negative cylindrical ends, we have that $k_- = 0$, so the resultant buildings have a *main level* $F^{(0)} : \Sigma^{(0)} \to T^*S^2$ and k_+ *upper levels* $F^{(\nu)} : \Sigma^{(\nu)} \to \mathbb{R} \times T_1^*S^2$.

Lemma 3.5.2. *A holomorphic building with non-empty main level in the SFT compactification of $\mathcal{M}_{\text{cyl}}$ has no non-trivial upper levels, and so, is given by a curve in T^*S^2. Moreover, this curve must be exactly one of those in Proposition 3.4.30.*

Proof. Denote the building by $F = (F^{(0)}, \ldots, F^{(k_+)})$. Its upper-most level $F^{(k_+)} : \Sigma^{(k_+)} \to \mathbb{R} \times T_1^*S^2$ is a $J^{(k_+)}$-holomorphic curve in a symplectisation with positive asymptotes given by $\{\gamma, \bar{\gamma}\}$. If $F^{(k_+)}$ is not the union of trivial cylinders $u_\gamma \sqcup u_{\bar{\gamma}}$, then it must intersect them positively:

$$i_U(F^{(k_+)}, u_\gamma \sqcup u_{\bar{\gamma}}) \geq i^+_{\text{MB}}(\gamma, \gamma) + i^+_{\text{MB}}(\bar{\gamma}, \bar{\gamma}) = 2.$$

Let $P^\gamma_\pm \in \mathcal{M}_\pm$ denote the unique $J_{T^*S^2}$-holomorphic plane asymptotic to γ. Since F is the limit of cylinders in $\mathcal{M}_{\text{cyl}}$, we have that (extending $P^\gamma_\pm + P^{\bar{\gamma}}_\mp$ by trivial upper levels)

$$2 = i(F, P^\gamma_\pm + P^{\bar{\gamma}}_\mp) \geq i_U(F^{(0)}, P^\gamma_\pm + P^{\bar{\gamma}}_\mp) + i_U(F^{(k_+)}, u_\gamma \sqcup u_{\bar{\gamma}}),$$

which implies that

$$i_U(F^{(0)}, P^\gamma_\pm + P^{\bar{\gamma}}_\mp) \leq 0.$$

In light of Corollary 3.3.3 and Proposition 3.3.2, this yields

$$i_U(F^{(0)}, P^\gamma_\pm) = 0 = i_U(F^{(0)}, P^{\bar{\gamma}}_\pm),$$

which is only possible if $F^{(0)}$ is simultaneously a cover of all four planes $P^\gamma_\pm$, $P^{\bar{\gamma}}_\pm$. However, since these planes are distinct, this is impossible.

We have shown that F has no non-trivial upper levels, and so $F^{(0)}$ is a curve in T^*S^2 with exactly two asymptotes $\{\gamma, \bar{\gamma}\}$. Therefore, we can apply Proposition 3.4.30 to complete the proof. □

We write $\overline{\mathcal{M}}_{\text{cyl}}$ to denote the subset of the SFT compactification of $\mathcal{M}_{\text{cyl}}$ consisting of buildings with non-empty main level. This is a topological surface homeomorphic to the complex plane $\mathbb{C}$. Moreover, just as in Section 2.5, $\overline{\mathcal{M}}_{\text{cyl}}$ can be equipped with a smooth structure (indeed the unique one) that makes the natural map

$$\pi_{T^*S^2} : T^*S^2 \to \overline{\mathcal{M}}_{\text{cyl}} : x \mapsto \text{the curve passing through } x$$

into a Lefschetz fibration with exactly two critical points corresponding to the intersection points of the two pairs of planes $P^{\gamma}_{\pm} + P^{\bar{\gamma}}_{\mp}$.

The next goal is to prove that L is fibred by circles in the foliation $\overline{\mathcal{M}}_{\text{cyl}}$, and thus the image of L under $\pi_{T^*S^2}$ is a smooth path.

Proposition 3.5.3. *Let $u : (\mathbb{C}^{\times}, j) \to T^*S^2$ be a smooth, properly embedded*[32] *$J_{T^*S^2}$-holomorphic cylinder with conjugate asymptotic Reeb orbits $\{\gamma, \bar{\gamma}\}$, as in scenario (2) of Proposition 3.4.30, that intersects the zero section $L \subset T^*L$. Then this intersection is along a circle contained in L.*

Proof. The construction of the almost complex structure $J_{T^*S^2}$ ensures that it is anti-invariant under the action of conjugation on T^*L. That is,

$$\overline{J_{T^*S^2}} = -J_{T^*S^2}.$$

Composing with conjugation gives a $(-j, J_{T^*S^2})$-holomorphic cylinder $\bar{u}$. We aim to show that $i_C(u, \bar{u}) = 0$. To this end, since $\overline{\mathcal{M}}_{\text{cyl}} \cong \mathbb{C}$, we can choose a homotopy from u to one of the nodal cylinders $P^{\gamma}_{+} + P^{\bar{\gamma}}_{-}$, where $P^{\gamma}_{+} = (F_{\mathcal{E}^+})^{(0)}_{-}$ and $P^{\bar{\gamma}}_{-} = (F_{\mathcal{E}^+})^{(0)}_{-}$ are the $\mathcal{M}_{\pm}$ planes asymptotic to γ and $\bar{\gamma}$ respectively. Then, as conjugation sends this pair of planes to itself, we can use the homotopy invariance, additivity, and

[32] All of the curves in $\mathcal{M}_{\text{cyl}}$ are properly embedded. Properness follows from their asymptotic behaviour, and embeddedness follows from the adjunction formula [52, §4.1].

symmetry of i_C to compute:

$$\begin{aligned}
i_C(u,\bar{u}) &= i_C(P_+^\gamma + P_-^{\bar{\gamma}}, \overline{P_+^\gamma + P_-^{\bar{\gamma}}}) \\
&= i_C(P_+^\gamma + P_-^{\bar{\gamma}}, P_+^\gamma + P_-^{\bar{\gamma}}) \\
&= i_C(P_+^\gamma, P_+^\gamma) + 2i_C(P_+^\gamma, P_-^{\bar{\gamma}}) + i_C(P_-^{\bar{\gamma}}, P_-^{\bar{\gamma}}) \\
&= -1 + 2 - 1 = 0.
\end{aligned}$$

However, since the zero section L is the fixed locus of conjugation, there is necessarily an intersection between u and $\bar{u}$. In view of the above, we deduce that u and $\bar{u}$ have the same image — they are geometrically indistinct. This implies that conjugation restricts to a j-anti-holomorphic involution of $(\mathbb{C}^\times, j)$ with non-empty fixed locus. Moreover, it's compact since u is proper. The fixed locus of such an involution is diffeomorphic to a circle, as is proved in the next lemma. As u is an embedding, this completes the proof. □

The following result is surely well known, and we include the proof only for completeness.

Lemma 3.5.4. *The fixed locus of an anti-holomorphic involution ι of $(\mathbb{C}^\times, j)$ is either empty, or non-empty and diffeomorphic to S^1, or two copies of $\mathbb{R}$.*

Proof. The action of the diffeomorphism group of $\mathbb{C}^\times$ is transitive on complex structures, so we can assume that $j = i$ is the standard complex structure. Identify $(\mathbb{C}^\times, i)$ with $(\mathbb{CP}^1 \backslash \{0, \infty\}, i)$, so that we have an anti-holomorphic involution ι of $\mathbb{CP}^1$ that preserves the set $\{0, \infty\}$. Then $z \mapsto \iota(\bar{z})$ is holomorphic, and preserves $\{0, \infty\}$, so it must be of the form

$$\iota(\bar{z}) = \begin{cases} \alpha z, & \text{if } \iota(0) = 0, \\ \alpha / z, & \text{if } \iota(\infty) = \infty, \end{cases}$$

for some $\alpha \in \mathbb{C}^\times$. In the first case $\iota(z) = \alpha \bar{z}$, if there is a fixed point $z_0 \in \mathbb{C}^\times$, then we have that $\alpha = z_0 / \bar{z}_0$, and

$$\operatorname{fix}(\iota|_{\mathbb{C}^\times}) = \{ r z_0 \mid r \in \mathbb{R} \backslash \{0\} \} \cong \mathbb{R} \sqcup \mathbb{R}.$$

On the other hand, if $\iota(z) = \alpha/\bar{z}$ has a non-zero fixed point $z \in \mathbb{C}^\times$, then

$$\alpha = z\bar{z} = |z|^2$$

must be a positive real number, and

$$\mathrm{fix}(\iota|_{\mathbb{C}^\times}) = (|z|^2 = \alpha) \cong S^1.$$

□

The next result is the pay-off for the work of this section.

Corollary 3.5.5. *The zero section $L \subset T^*S^2$ is a matching cycle of the Lefschetz fibration $\pi_{T^*S^2}$.*

Proof. Since L is fibred by circles, it lives submersively over its projection by $\pi_{T^*S^2}$, which is a smooth embedded path $p : [-1,1] \to \overline{\mathcal{M}}_{\mathrm{cyl}}$ joining the two critical values of $\pi_{T^*S^2}$. Denote the symplectic parallel transport of this fibration by ϕ_t. By Lemma 1.17 of [18] L is the trace under ϕ_t of its intersection with the fibre over $p(0)$, that is,

$$L = \bigcup_{t \in [-1,1]} \phi_t \left(L \cap \pi^{-1}_{T^*S^2}(p(0)) \right).$$

Since the critical points of $\pi_{T^*S^2}$ lie on L, we conclude that it is indeed a matching cycle. □

Chapter 4

Isotopy to a matching cycle

In this chapter we use the neck stretching analysis of Chapter 3 to show that a matching cycle Σ_k of of π_{J_k} converges to the matching cycle of $\pi_{T^*S^2}$. In turn, this yields a Lagrangian isotopy from L to Σ_k for sufficiently large k.

Before discussing convergence of matching cycles, we need to understand how to construct a sequence of matching paths for π_{J_k} that converges (in some sense) to the matching path of Corollary 3.5.5. This is the subject of Section 4.1. Once this is done, Section 4.2 shows that the convergence of the matching cycles essentially follows from smooth dependence of ODEs on their defining vector field and initial condition. Finally, Section 4.3 completes the proof that every Lagrangian sphere $L \subset B_{d,p,q} \subset X$ is Lagrangian isotopic to a matching cycle of a fixed Lefschetz fibration π_{ref}.

4.1 Convergence of matching paths

Let $x \in T^*S^2$ and recall from Section 3.5 that there is a unique sequence of curves $f_k^x \in \overline{\mathcal{M}}_{0,0}(X, F; J_k)$ that converge to a $J_{T^*S^2}$-holomorphic building passing through x. In this chapter, we'll primarily be interested in the convergence to the bottom levels in the Weinstein neighbourhood Y_- of L.

Before we can talk about convergence of matching cycles of the fibrations π_{J_k} :

$X_{J_k} \to \overline{\mathcal{M}}_{0,0}(X_{J_k}, F; J_k)$ to the zero section matching cycle of $\pi_{T^*S^2} : T^*S^2 \to \overline{\mathcal{M}}_{\mathrm{cyl}}$, we need to make sense of what it means for matching paths in the bases $\overline{\mathcal{M}}_{0,0}(X_{J_k}, F; J_k)$ to converge to the matching path of Corollary 3.5.5. The global charts of $\overline{\mathcal{M}}_{0,0}(X_{J_k}, F; J_k)$ constructed in Section 2.5 are not fit for this purpose, since the neck stretching analysis of Section 3.1 shows that, as $k \to \infty$, the subset $\overline{\mathcal{M}}_L(J_k) \subset \overline{\mathcal{M}}_{0,0}(X_{J_k}, F; J_k)$ of curves that pass through Y_- shrinks to a point. We construct a new chart of $\overline{\mathcal{M}}_{0,0}(X_{J_k}, F; J_k)$ with image contained in $\overline{\mathcal{M}}_L(J_k)$ using one of the $J_{T^*S^2}$-holomorphic $\mathcal{M}_+$ planes in T^*S^2. Choose a simple Reeb orbit γ_0 that is neither γ nor $\bar{\gamma}$. Fix the unique plane $P_+^{\gamma_0} \in \mathcal{M}_+$ asymptotic to γ_0 along with a parametrisation $u : \mathbb{C} \to T^*S^2$ of it.

Lemma 4.1.1. *The plane $P_+^{\gamma_0}$ intersects each curve in $\overline{\mathcal{M}}_{\mathrm{cyl}}$ exactly once, transversely. In particular, for each $C \in \overline{\mathcal{M}}_{\mathrm{cyl}}$,*

$$i_U(P_+^{\gamma_0}, C) = i_C(P_+^{\gamma_0}, C) = 1.$$

Therefore, the composition $\pi_{T^*S^2} \circ u : \mathbb{C} \to \overline{\mathcal{M}}_{\mathrm{cyl}}$ is a global chart.

Proof. The intersection claim follows easily from homotopy invariance of i_C and the fact that $\gamma_0 \neq \gamma, \bar{\gamma}$ (see Section 3.3)

$$i_C(P_+^{\gamma_0}, C) = i_C(P_+^{\gamma_0}, P_+^{\gamma} + P_-^{\bar{\gamma}}) = i_C(P_+^{\gamma_0}, P_-^{\bar{\gamma}}) = 1.$$

By positivity of intersections, there is exactly one positive, transverse intersection as claimed. The fact that this provides a global chart follows by definition of the smooth structure on $\overline{\mathcal{M}}_{\mathrm{cyl}}$. □

Remark 4.1.2. Since $PSL(2, \mathbb{C})$ acts transitively on the configuration space of 3 points on $\mathbb{CP}^1$, we can reparametrise u to ensure that the two critical values of $\pi_{T^*S^2}$ correspond to the points ± 1.

Recall that the neck stretching setup gives, for each k, an almost complex embedding $\iota_k : (T^*_{\leq e^k} S^2, J_{T^*S^2}) \to (X, J_k)$ with image the Weinstein neighbourhood $Y_- = T^*_{\leq 1} L$ of L. This allows us to embed larger and larger portions of the plane

$P_+^{\gamma_0}$ as the neck gets longer. The idea is that, eventually, the curves in $\overline{\mathcal{M}}_L(J_k)$ get close enough to their limiting bottom level curves in $\overline{\mathcal{M}}_{\mathrm{cyl}}$ so that the transverse intersection property of the previous lemma holds for them too.

Lemma 4.1.3. *There exists an integer $K > 0$ such that for all $k \geq K$ every curve in $\overline{\mathcal{M}}_L(J_k)$ intersects $P_+^{\gamma_0}$ exactly once transversely. Therefore, the parametrisation u of $P_+^{\gamma_0}$ facilitates a chart of the moduli spaces $\overline{\mathcal{M}}_{0,0}(X_{J_k}, F; J_k)$ for all $k \geq K$. Moreover, K can be chosen large enough so that there are only the two* relevant *critical values of the maps π_{J_k} contained in the image of this chart. That is, the two that correspond to the nodal curves $P_\pm^{\gamma} + P_\mp^{\bar{\gamma}}$.*

Proof. Fix $x \in Y_-$, then there is a unique sequence $f_k^x \in \overline{\mathcal{M}}_L(J_k)$ of fibre curves that pass through x. We first prove that the result holds for this sequence. That is, that there exists $K_x > 0$ such that, for all $k \geq K_x$, f_k^x transversely intersects $P_+^{\gamma_0}$ exactly once. Recall from Section 3.5 that the sequence J_k was chosen to ensure that each of the sequences f_k^x converges[1] to a J-holomorphic building $F^x : \Sigma^* \to X^*$, where the bottom level $(F^x)_-^{(0)}$ is the element of $\overline{\mathcal{M}}_{\mathrm{cyl}}$ passing through x. The nature of the convergence $f_k^x \to F^x$ implies that, for any $\epsilon > 0$, there exists $K_{x,\epsilon}$ such that (up to reparametrisation) f_k^x is ϵ-close to $(F^x)_-^{(0)}$ in the C^1_{loc} topology.

Now, $(F^x)_-^{(0)}$ and $P_+^{\gamma_0}$ intersect exactly once, but since they are non-compact, we need to be slightly careful before claiming that all nearby curves satisfy the same property. Fortunately, they are non-compact in a very controlled way, due to the asymptotic convergence to the Reeb orbits. The convergence $f_k^x \to F^x$ guarantees,[2] for all sufficiently large k, that the curves f_k^x are bounded away from the cylindrical end of $P_+^{\gamma_0}$. Therefore, the only intersections between f_k^x and $P_+^{\gamma_0}$ can occur in a compact set. Thus, by standard differential topology, any such map transversely

[1] Note that some of the sequences f_k^x necessarily contain nodal curves, since there are two nodal curves in $\overline{\mathcal{M}}_{\mathrm{cyl}}$. We have not defined exactly what it means for such a sequence (consisting of potentially non-smooth J_k-holomorphic curves) to converge to a holomorphic building. However, since we are only interested in the convergence *away from the nodes*, we will not go into the details.

[2] See Definition 2.7(d) of [9].

intersects $P_+^{\gamma_0}$ in exactly the same number of points as $(F^x)_-^{(0)}$ — that is, exactly once. Therefore, we deduce that there exists $\epsilon_x > 0$, and $K_x = K_{x,\epsilon_x} > 0$, such that, for all $k \geq K_x$, f_k^x is ϵ_x-close to $(F^x)_-^{(0)}$ *and* bounded away from the cylindrical end of $P_+^{\gamma_0}$, from which it follows that f_k^x intersects $P_+^{\gamma_0}$ exactly once, transversely.

We turn to showing that the result continues to hold as stated. This is a consequence of the local description of the moduli spaces $\overline{\mathcal{M}}_L(J_k)$ and compactness of the Weinstein neighbourhood Y_-. First, we show that if the transverse intersection property holds for a curve $f_{K_x}^x \in \overline{\mathcal{M}}_L(J_{K_x})$, then it holds for all nearby curves $f \in \overline{\mathcal{M}}_L(J_{K_x})$. This is achieved by finding a tubular neighbourhood N_x of $f_{K_x}^x$ such that the intersection $P_+^{\gamma_0} \cap N_x$ is connected and C^1-close to a fibre of N_x. The claim then follows from standard differential topology.

If the curve $f_{K_x}^x$ is nodal, then the tubular neighbourhood N_x is constructed via the gluing map. The nodal points of curves in $\overline{\mathcal{M}}_L(J_k)$ converge to the nodal points of the curves in $\overline{\mathcal{M}}_{\mathrm{cyl}}$, and since $P_+^{\gamma_0}$ intersects the curves in $\overline{\mathcal{M}}_{\mathrm{cyl}}$ away from the nodes, we can increase K_x so that either $f_{K_x}^x$ is no longer nodal, or the node is bounded away from $P_+^{\gamma_0}$. Note that, away from nodal points the gluing map is a diffeomorphism, and this is where the unique transverse intersection $f_{K_x}^x \cap P_+^{\gamma_0}$ occurs. By varying the gluing parameter, we obtain arbitrarily small tubular neighbourhoods N_x of $f_{K_x}^x$.

On the other hand, if $f_{K_x}^x$ is smooth we use the local description of the moduli space of smooth curves $\mathcal{M}_{0,0}(X, F; J_{K_x})$ to express nearby curves as sections of the (trivial) normal bundle $\nu f_{K_x}^x$. Thus, we obtain the tubular neighbourhood N_x. In both the smooth and the nodal cases, $f_{K_x}^x$ is bounded away from the cylindrical part of $P_+^{\gamma_0}$, so we can choose the tubular neighbourhood N_x small enough so that the intersection $N_x \cap P_+^{\gamma_0}$ is connected. In both the smooth and the nodal cases, this implies that the intersection property continues to hold for all sufficiently nearby curves $f \in \overline{\mathcal{M}}_L(J_{K_x})$. In particular, there exists a number $\delta_x > 0$ such that, for all $f \in \overline{\mathcal{M}}_L(J_{K_x})$ satisfying[3] $d(f, f_{K_x}^x) < \delta_x$, f intersects $P_+^{\gamma_0}$ exactly once transversely.

[3]This metric refers to any that induces the manifold topology on $\overline{\mathcal{M}}_L(J_{K_x})$.

In summary, so far we have proved that, for all $x \in Y_-$, there exists $K_x > 0$ and $\delta_x > 0$ such that, for all $y \in Y_-$, and $f^y_{K_x} \in \overline{\mathcal{M}}_L(J_{K_x})$ satisfying $d(f^y_{K_x}, f^x_{K_x}) < \delta_x$, we have that $f^y_{K_x}$ intersects $P^{\gamma_0}_+$ exactly once, transversely. Furthermore, since the convergence $f^y_k \to F^y$ is C^1-monotonic, we can upgrade the statement to hold for all $k \geq K_x$. Specifically, for all $k \geq K_x$ and $y \in Y_-$ such that $d(f^y_{K_x}, f^x_{K_x}) < \delta_x$, we have that f^y_k intersects $P^{\gamma_0}_+$ as desired. We now harness compactness to turn this into a global statement.

For a given $x \in Y_-$, the curves near $f^x_{K_x}$ form a local foliation around x. Therefore, we have that

$$U_x := \bigcup_{\substack{f \in \overline{\mathcal{M}}_L(J_{K_x}) \\ d(f, f^x_{K_x}) < \delta_x}} \operatorname{im} f \cap Y_-$$

is a neighbourhood of x. Thus, we obtain a cover of Y_-. Since it's compact, we can pass to a finite subcover $U_{x_1}, \ldots, U_{x_k}$ and define

$$K := \max\{K_{x_1}, \ldots, K_{x_k}\} < \infty.$$

Then, as any point $x \in Y_-$ is necessarily contained in some U_{x_i}, we have that $d(f^x_{K_{x_i}}, f^{x_i}_{K_{x_i}}) < \delta_{x_i}$, which implies that f^x_k intersects $P^{\gamma_0}_+$ exactly once, transversely, for all $k \geq K \geq K_{x_i}$. As any curve $f \in \overline{\mathcal{M}}_L(J_k)$ is equal to f^x_k for some $x \in Y_-$, this completes the proof.

The claim about the critical values follows from the discussion on choosing the sequence J_k in Section 3.5. □

As a result, for sufficiently large k, we can choose a sequence of matching paths $p_k : [-1, 1] \to \overline{\mathcal{M}}_{0,0}(X_{J_k}, F; J_k)$ joining the two relevant critical points, such that, in the coordinates of the charts given by Lemma 4.1.3, p_k converges to the matching path $p : [-1, 1] \to \mathbb{C}$ of the Lefschetz fibration $\pi_{T^*S^2}$ associated to L. Our next goal is to show that we can choose k large enough so that the matching cycle of π_{J_k} associated to p_k is Lagrangian isotopic to L.

4.2 Convergence of the parallel transport

4.2.1 The parallel transport of $\pi_{T^*S^2} : T^*S^2 \to \overline{\mathcal{M}}_{\mathrm{cyl}}$

The problem of symplectic parallel transport involves lifting a vector field defined over a compact path in the base to a horizontal vector field in the total space and integrating it. This amounts to solving an ODE over a compact family of fibres of the Lefschetz fibration $\pi_{T^*S^2}$. However, since these fibres are non-compact, we need to justify why solutions to this ODE exist for all time.

More precisely, fix an embedded path $\gamma : [0,1] \to \overline{\mathcal{M}}_{\mathrm{cyl}}$ whose image avoids the critical values of $\pi_{T^*S^2}$. Taking symplectic orthogonal complements of the tangent spaces to the regular fibres of $\pi_{T^*S^2}$ yields a field of *horizontal* planes $H \subset TT^*S^2$, such that the restriction $\mathrm{d}\pi_{T^*S^2}|_H : H \to T\overline{\mathcal{M}}_{\mathrm{cyl}}$ is an isomorphism. Therefore, there exists a unique horizontal lift $\tilde{X}$ of the vector field $X = \frac{\mathrm{d}}{\mathrm{d}t}p$. This defines an ODE on the space $\pi^{-1}_{T^*S^2}(\operatorname{im}\gamma)$.

Lemma 4.2.1. *The symplectic parallel transport of the Lefschetz fibration $\pi_{T^*S^2} : T^*S^2 \to \overline{\mathcal{M}}_{\mathrm{cyl}}$ is well-defined.*

Proof. The curves in the moduli space $\overline{\mathcal{M}}_{\mathrm{cyl}}$ are finite energy $J_{T^*S^2}$-holomorphic curves and so, by the results of Hofer, Wysocki, and Zehnder [27, Theorem 1.3], Bourgeois [4, §3.3], and Mora-Donato [40, Proposition 1.2], they satisfy exponential convergence to their asymptotic Reeb cylinders. The fixed path $\gamma : [0,1] \to \overline{\mathcal{M}}_{\mathrm{cyl}}$ yields a compact family $[0,1] \cong \kappa \subset \overline{\mathcal{M}}_{\mathrm{cyl}}$ of such curves, and so, there exists a compact Liouville subdomain of T^*S^2 outside of which, all the curves in κ satisfy the exponential convergence estimates.

Fix the standard Riemannian metric on T^*S^2 coming from the embedding

$$T^*S^2 = \{(p,q) \in \mathbb{R}^3 \times \mathbb{R}^3 \mid |q| = 1,\ \langle p,q\rangle = 0\},$$

and consider the Hamiltonian $H : T^*S^2 \to \mathbb{R} : H(p,q) = \frac{1}{2}|p|^2$. The ω_{can}-orthogonal complement of the tangent space of a Reeb cylinder is a contact plane, which is contained in $\ker \mathrm{d}H$. Therefore, the exponential convergence of the curves implies that the horizontal spaces converge exponentially to the contact planes. In particular, the horizontal lift $\tilde{X}$ satisfies

$$|L_{\tilde{X}}H| < Ce^{-ds},$$

where $C > 0$ and $d > 0$ are constants, and $s \in [s_0, \infty)$ is the Liouville coordinate on T^*S^2. An integral curve α of $\tilde{X}$ escapes to infinity if, and only if, $|H(\alpha)| \to \infty$. Therefore, the above estimate shows that this is impossible, as we are integrating over the compact set $[0, 1]$. □

4.2.2 The parallel transport of $\pi_{J_k} : X_{J_k} \to \mathbb{C}$

Recall from Lemma 4.1.3 that, for large enough k, we can identify the bases of the Lefschetz fibrations $\pi_{J_k}|_{Y_-}$ and $\pi_{T^*S^2}|_{Y_-}$ with a bounded disc $D \subset \mathbb{C}$.

Lemma 4.2.2. *Let $\gamma_k : [0, 1] \to D$ be a sequence of embedded paths converging C^1 to $\gamma : [0, 1] \to D$. Suppose that each path avoids the critical locus of the corresponding Lefschetz fibration, then the parallel transport ϕ_{γ_k} of $\pi_{J_k}|_{Y_-}$ over γ_k is C^0-close to that of $\pi_{T^*S^2}|_{Y_-}$ over γ.*

Proof. Since we are dealing with compact families of curves, so one can show that the fibres of $\pi_{J_k}|_{Y_-}$ living over γ_k converge uniformly to those of $\pi_{T^*S^2}|_{Y_-}$ over γ. It follows that the ODE defining the parallel transport ϕ_γ of $\pi_{T^*S^2}$ can be arbitrarily well approximated by that defining ϕ_{γ_k} by increasing k. Therefore, by smooth dependence of ODEs on the initial condition and the vector field that defines it (see [12, Appendix B] for example) we obtain the result. □

Lemma 4.2.3. *Let $\gamma_k : [-1, 1] \to D$ be a sequence of matching paths converging to the matching path $\gamma : [-1, 1] \to D$ corresponding to the zero section matching cycle of $\pi_{T^*S^2}$. Then, for sufficiently large k, the matching cycle Σ_k of $\pi_{J_k}|_{Y_-}$ associated to γ_k is C^0-close to the zero section L.*

Proof. Since the critical points of $\pi_{J_k}|_{Y_-}$ converge to those of $\pi_{T^*S^2}|_{Y_-}$, smooth dependence of ODEs implies that we have C^0 convergence of the vanishing thimbles associated to the paths $\gamma_k^- := \gamma_k|_{[-1,0]}$ and $\gamma_k^+ := \gamma_k|_{[0,1]}$ in a small neighbourhood of the critical points. Therefore, we can apply Lemma 4.2.2 to the restrictions $\gamma_k^-|_{[-1+\epsilon,0]}$ and $\gamma_k^+|_{[0,1-\epsilon]}$ for some small $\epsilon > 0$ to deduce that the vanishing thimbles of π_{J_k} over the paths $\gamma_k^\pm$ can be made C^0-close to those of $\pi_{T^*S^2}$ over $\gamma^\pm$ by increasing k.

Choose k large enough so that the vanishing thimbles are very close to those of $\pi_{T^*S^2}$ and in particular, are contained in the Weinstein neighbourhood Y_-. To form the matching cycle a deformation of the symplectic structure on Y_- is made to account for the fact that the vanishing cycles over $\gamma_k(0)$ may not agree (see [47, Lemma 15.3]). A Moser-type argument ([47, Lemma 7.1]) is then used to map the resulting sphere back to the original symplectic structure. Since the vanishing thimbles themselves are C^0-close, the deformation (and resulting Moser isotopy) can be made so that the matching sphere Σ_k remains C^0-close to L. This completes the proof. □

4.3 Constructing the isotopy

Theorem 4.3.1. *There exists k sufficiently large such that L is Lagrangian isotopic to a matching cycle of π_{J_k}.*

Proof. Lemma 4.2.3 shows that the relevant matching cycle Σ_k is a Lagrangian sphere contained in a neighbourhood of L that is symplectomorphic to $(T^*_{\leq e^{R_k}}S^2, e^{-R_k}\omega_{\text{can}})$ for some positive number $R_k > 0$. Therefore, we can apply Hind's theorem [24, Theorem 18] on uniqueness of Lagrangian spheres in T^*S^2 to obtain a Lagrangian isotopy from Σ_k to L. If necessary, by re-scaling by the Liouville flow, we can ensure that the isotopy is supported in $T^*_{\leq e^{R_k}}S^2$. Finally, since this neighbourhood of L symplectically embeds into X, we obtain the isotopy in X as desired. □

In summary, we have found a Lefschetz fibration $\pi_{J_k} : X_{J_k} \to \mathbb{C}$ and a Lagrangian isotopy supported in a small neighbourhood of L taking L to a matching cycle Σ_k.

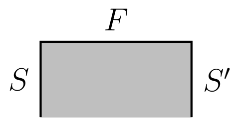

Figure 4.1: The neighbourhood N of the fibre component of D' in which we fix the almost complex structures to be equal to the usual product complex structure coming from $S^2 \times S^2$.

The next task is to "undo the neck stretch" and find another Lagrangian isotopy taking Σ_k to a matching cycle of some fixed Lefschetz fibration with respect to which we will do the remaining computations.

We use the fact that the space $\mathcal{J}(D')$ of compatible almost complex structures for which each component of the divisor $D' \subset X$ is J-holomorphic is connected. In fact, we restrict ourselves even further to the subset of $\mathcal{J}(D')$ of almost complex structures that are *fixed* in a neighbourhood of the F-component of D'. In more detail, consider a neighbourhood N of this component of the form shown in Figure 4.1. Remark A.1.4 ensures that we can choose[4] $J_{\text{ref}} \in \mathcal{J}(D')$ so that it agrees with the product almost complex structure coming from $S^2 \times S^2$ in N. Let $U := X \backslash N$ and denote the subset of $\mathcal{J}(D')$ of almost complex structures that agree with J_{ref} in $N = X \backslash U$ by $\mathcal{J}(U, J_{\text{ref}})$. That is,

$$\mathcal{J}(U, J_{\text{ref}}) := \{J \in \mathcal{J}(D') \mid J|_{X\backslash U} = J|_N = J_{\text{ref}}|_N = J_{\text{ref}}|_{X\backslash U}\}.$$

This space is connected, which follows from Sévennec's argument (see [2, Proposition 1.1.6] for example). Observe that, for a suitable choice of neighbourhood N, all the stretched almost complex structures J_k are contained in $\mathcal{J}(U, J_{\text{ref}})$.

The idea is to use connectedness of $\mathcal{J}(U, J_{\text{ref}})$ to choose a path J_t of almost complex structures from $J_0 = J_k$ to $J_1 = J_{\text{ref}}$. Applying the results of Section 2.4 yields a path of Lefschetz fibrations $\pi_t = \pi_{J_t} : X_{J_t} \to \mathbb{C}$ from $\pi_0 = \pi_{J_k}$ to $\pi_1 = \pi_{J_{\text{ref}}}$. This path of fibrations induces a path in the configuration space $C(\mathbb{C}, d)$ of d points in $\mathbb{C}$ corresponding to the positions of the critical values of π_t (see Remark 2.5.6(1)).

[4] In fact, in Section 5.3 we shall explicitly construct a specific J_{ref}.

We note two features of this path. Firstly, recall that the point at infinity in the bases $\overline{\mathcal{M}}_{0,0}(X_{J_t}, F; J_t) \cong \mathbb{C}$ corresponds to the exotic stable curve $\mathbf{u}_\infty^{J_t}$. Therefore, the points must remain in a bounded subset by the fact that Lefschetz critical fibres never intersect $\mathbf{u}_\infty^{J_t}$. Secondly, since the almost complex structures J_t are fixed on N and N is foliated by smooth J_t-holomorphic F-curves the points also remain bounded away from $0 \in \mathbb{C}$. As a result, this path in $C(\mathbb{C}, d)$ is really a path in $C(A, d)$ where $A \cong [0,1] \times S^1$ is a compact annulus. Now, choose an isotopy of matching paths $\gamma_t : [-1,1] \to A$ for π_t such that $\gamma_0 = \gamma_k$ is the matching path for Σ_k from Lemma 4.2.3. Forming the corresponding matching cycles L_t yields a Lagrangian isotopy from $L_0 = \Sigma_k$ to the matching cycle of π_{ref} over γ_1. This isotopy is disjoint from open neighbourhoods of $\mathbf{u}_\infty^{J_t}$ and the F-component of the divisor $D' \subset X$. However, *a priori* it may pass through the sections S and S'. This is a problem since we want the Lagrangian isotopy to be supported in a subset of X that is symplectomorphic to a subset of $B_{d,p,q}$. The point of the next result is to avoid this behaviour by explicitly altering the Lefschetz fibrations π_t.

The curves in $\overline{\mathcal{M}}_{0,0}(X, F; J)$ form singular foliations $\mathscr{F}_t$ on X. The singular leaves are exactly the singular Lefschetz fibres of π_t, and one *exotic* leaf given by $\mathbf{u}_\infty^{J_t}$. In small tubular neighbourhoods of the J_t-holomorphic sections S and S', these foliations are smooth with leaves given by symplectic 2-discs. We denote these by $\mathscr{D}_t$.

Lemma 4.3.2. *There exists an isotopy $\mathscr{D}_{s,t}$ of foliations such that $\mathscr{D}_{0,t} = \mathscr{D}_t$ and $\mathscr{D}_{1,t}$ are foliations by symplectic 2-discs intersecting the sections S and S' positively and symplectically orthogonally. The isotopy can be chosen so that it is invariant in s outside of an arbitrarily small neighbourhood of each section. As a result, we obtain singular foliations $\mathscr{F}_{s,t}$ of X such that, excluding a single exotic leaf, $\mathscr{F}_{s,t}$ forms a* symplectic Lefschetz fibration[5] *on X.*

Proof. Put $S_1 = S$ and $S_2 = S'$. Note that near a section S_i, the fibrations

[5] That is, Lefschetz fibrations whose fibres are only symplectic submanifolds, and not necessarily J-holomorphic.

π_t are foliations $\mathscr{D}_t$ of symplectic 2-discs whose leaves intersect S_i transversely and positively. Therefore, we can use a variation of an argument of Gompf [21, Lemma 2.3] to construct an isotopy of foliations $\mathscr{D}_{s,t}$ such that: (1) the point of intersection of a leaf and S_i is invariant in s; and (2) outside of a small neighbourhood of S_i, the leaves are also invariant in s. Furthermore, the leaves of $\mathscr{D}_{1,t}$ intersect S_i symplectically orthogonally. Consequently, we can glue the symplectic discs in $\mathscr{D}_{s,t}$ to the leaves of $\mathscr{F}_t$ to obtain new foliations $\mathscr{F}_{s,t}$. By construction, these are constant in s in a complement of small neighbourhoods of the sections S_i, which implies that (except for the one exotic fibre corresponding to the $\mathbf{u}_\infty^{J_t}$ curve) their leaves form Lefschetz fibrations, since the existence of a Lefschetz chart at a critical point is a local condition. This completes the proof. □

Let X_t denote X with the exotic leaf of $\mathscr{F}_{1,t}$ excised. Write $\pi_{1,t} : X_t \to \mathbb{C}$ for the corresponding Lefschetz fibration. Define the reference Lefschetz fibration by $\pi_{\text{ref}} := \pi_{1,1}$, and observe that it does not depend on the Lagrangian sphere L. In addition to the tubular neighbourhood N defined earlier, choose a tubular neighbourhood foliated by leaves of $\mathscr{F}_{1,1}$ of the exotic leaf, then define $\mathring{X} \subset U \subset X$ to be the excision of this as well as the sections S and S'. The restriction $\pi_{\text{ref}}|_{\mathring{X}}$ is a Lefschetz fibration over the compact annulus.

Corollary 4.3.3. *There exists a Lagrangian isotopy from L to a matching cycle of π_{ref} that is supported in $\mathring{X}$.*

Proof. Theorem 4.3.1 shows that it suffices to find such an isotopy from the matching cycle Σ_k to one of π_{ref}. Observe that Σ_k is still a matching cycle for the deformed fibration $\pi_{1,0}$ since none of the structure changes near Σ_k throughout the isotopy $\mathscr{F}_{s,t}$. Therefore, we can construct the Lagrangian isotopy using the matching paths γ_t defined earlier, except this time we form the matching cycles with respect to the maps $\pi_{1,t}$. To see that this has support as claimed, note that the orthogonality condition ensures that the parallel transport maps of $\pi_{1,t}$ preserve the sections S_i. Therefore, they preserve their complements too. Since Σ_k lives in this complement, the claim follows from this and the discussion preceding Lemma 4.3.2. $\square$

Chapter 5

Mapping class groups of surfaces and symplectomorphisms

In this chapter we harness the theory of mapping class groups of surfaces to improve Corollary 4.3.3 to the main result Theorem 1.2.2. Corollary 4.3.3 tells us that any Lagrangian sphere $L \subset B_{d,p,q}$ is Lagrangian isotopic to a matching cycle of the Lefschetz fibration π_{ref}. To convert this into an isotopy statement phrased in terms of Dehn twists, we need to understand their relation to (isotopy classes of) matching paths. This will be facilitated by the fact that the natural action of the mapping class group of the d-punctured annulus $A^{\times d} := [0,1] \times S^1 \backslash \{d \text{ points}\}$ is transitive on matching paths, and so, it will suffice to understand particularly simple matching paths.

5.1 The mapping class group of the punctured annulus

The mapping class group of a surface is a classical object with plenty of rich theory. Farb and Margalit's book [20] — in particular, Sections 1–4 — contains everything we will use here. We recall the main definition.

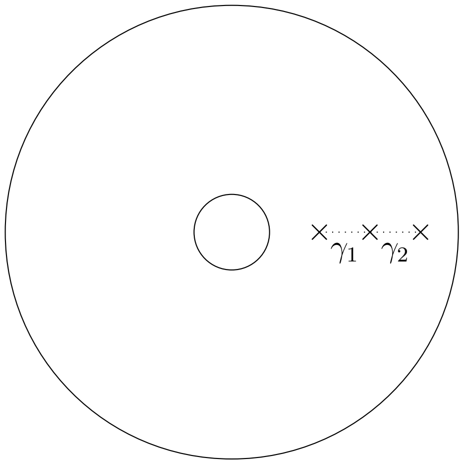

Figure 5.1: The punctured annulus and the matching paths γ_i.

Definition 5.1.1. Let S be an oriented surface (possibly with boundary and punctures). The *mapping class group of* S, $\mathrm{Mod}(S)$, is defined to be the group of connected components of the group of orientation preserving diffeomorphisms[1] of S that fix the boundary point wise. That is,

$$\mathrm{Mod}(S) := \pi_0(\mathrm{Diff}^+(S, \partial S)).$$

Recall from Section 4.3, the restriction of the Lefschetz fibration π_{ref} to $\mathring{X}$ has base a d-punctured annulus A. Therefore, we seek to understand the mapping class group $\mathrm{Mod}(A)$. In Appendix A.2 we compute a presentation of $\mathrm{Mod}(A)$:

Example 5.1.2 (Appendix A.2)**.** The mapping class group of the punctured annulus is isomorphic to the direct product of the *annular braid group* with a copy of $\mathbb{Z}$. This is finitely generated, and each generator is one of three types: the Dehn twist T about the central boundary, the so-called central twist τ, and half-twists σ_i around the straight line paths γ_i joining the adjacent i and $(i+1)$th punctures. Explicitly,

[1] In fact, as Section 1.4.2 of [20] explains, we are free to consider homeomorphisms or diffeomorphisms interchangeably.

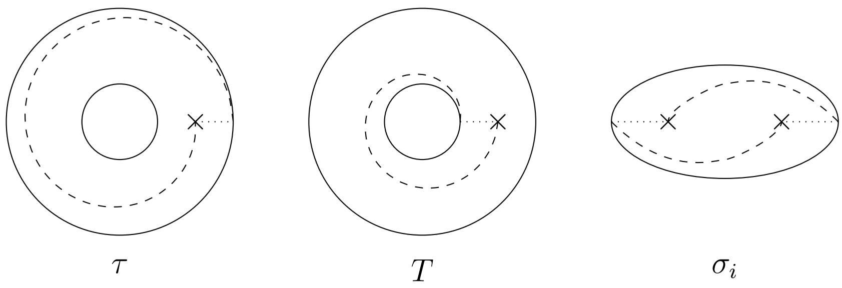

Figure 5.2: Local characterisations of each generator of Mod(A). In each case, dotted lines map to dashed.

the presentation is:

$$\mathrm{Mod}(A) = \left\langle \tau, \sigma_1, \ldots, \sigma_{n-1} \;\middle|\; \begin{array}{c} (\tau\sigma_1)^2 = (\sigma_1\tau)^2, \\ \tau\sigma_i = \sigma_i\tau \ \forall i > 1, \\ \sigma_i\sigma_{i+1}\sigma_i = \sigma_{i+1}\sigma_i\sigma_{i+1}, \\ \sigma_i\sigma_j = \sigma_j\sigma_i \ \forall |i-j| > 1 \end{array} \right\rangle \times \langle T \rangle.$$

The behaviour of each generator is characterised by the diagrams in Figure 5.2.

Since any Lagrangian sphere $L \subset B_{d,p,q}$ is Lagrangian isotopic to a matching cycle Σ_γ fibred over the matching path γ, the following basic fact will be crucial in proving Theorem 1.2.2.

Fact. The mapping class group Mod(A) acts transitively on the set of isotopy classes of matching paths $\mathcal{P}$.

Remark 5.1.3. In fact, this is true if A is replaced by any punctured oriented surface, as can be proved by a cut-and-paste argument appealing to the classification of surfaces, as the author learned from the MathOverflow answer [35].

Therefore, for some mapping class $\rho \in \mathrm{Mod}(A)$, $\rho(\gamma)$ is isotopic to the so-called standard matching path γ_1 (see Figure 5.1). It follows that the matching cycle Σ_γ is Lagrangian isotopic to $\Sigma_{\rho(\gamma_1)}$. So to prove Theorem 1.2.2, we seek to a correspondence between symplectomorphisms of $B_{d,p,q}$ and mapping classes in Mod(A).

That the half-twists σ_i correspond to generalised Dehn twists about the Lagrangian spheres Σ_{γ_i} is well known (and proved in [47, 16h]). The Dehn twist $T \in \mathrm{Mod}(A)$ acts trivially on any class $[\gamma] \in \mathcal{P}$ since no matching path has an end point lying on the central boundary component. So, our task is to understand to what symplectomorphism the central twist τ corresponds. More precisely, is there a symplectomorphism ϕ of $B_{d,p,q}$ such that $\phi(\Sigma_{\gamma_1})$ is Lagrangian isotopic to $\Sigma_{\tau(\gamma_1)}$?

5.2 Some symplectomorphisms of $B_{d,p,q}$

The aim of this section is to prove the following result.

Proposition 5.2.1. *There exists a compactly-supported symplectomorphism $\tau_{p,q}$ of $B_{d,p,q}$ arising from the symplectic monodromy of the $\frac{1}{p^2}(1, pq-1)$ singularity. The central twist $\tau \in \mathrm{Mod}(A)$ corresponds to $\tau_{p,q}$, that is, $\tau_{p,q}(\Sigma_{\gamma_1})$ is Lagrangian isotopic to $\Sigma_{\tau(\gamma_1)}$. Moreover, no iterate $\Sigma_{\tau^k(\gamma_1)}$, for $k \neq 0$, is Lagrangian isotopic to Σ_{γ_1}.*

As a corollary of this, we find that $\tau_{p,q}$ has infinite order in the symplectic mapping class group $\pi_0(\mathrm{Symp}_c(B_{d,p,q}))$.

We adapt Seidel's approach [45, §4.c] to the symplectic monodromy to suit our situation. We work in $\mathbb{C}^3$ with coordinates $z = (z_1, z_2, z_3)$. Let ψ be a cut-off function[2] satisfying

$$\psi(t^2) = \begin{cases} 1, & \text{if } t \leq \frac{1}{3} \\ 0, & \text{if } t \geq \frac{2}{3}. \end{cases}$$

and define

$$\tilde{M}_w := \{(z_1, z_2, z_3) \in \mathbb{C}^3 \mid |z| \leq 1,\ z_1 z_2 = z_3^{dp} + \psi(|z|^2)w\}$$

Note that, for some $\epsilon > 0$ sufficiently small, the manifolds $\tilde{M}_w$ with $0 < |w| \leq \epsilon$ are smooth symplectic manifolds diffeomorphic to the Milnor fibre $(z_1 z_2 = z_3^{dp}) \cap (|z| \leq 1)$ of the A_{dp-1} singularity (see Lemmata 4.9 and 4.10 of [45]). The $\tilde{M}_w$ are invariant

[2]The purpose of which is to ensure that the symplectic parallel transport maps used in the definition of monodromy are well defined.

under the $\frac{1}{p}(1, -1, q)$ action on $\mathbb{C}^3$, and so (being a subgroup of the unitary group $U(3)$) we can take the quotient to obtain symplectic manifolds M_w diffeomorphic to $B_{d,p,q}$.

The Milnor fibration associated to the singular point $0 \in M_0$ is defined to be restriction of the projection $\mathbb{C}^3 \times S^1 \to S^1$:

$$\pi : M := \bigcup_{|w|=\epsilon} M_w \times \{w\} \to S^1,$$

Pulling back the standard symplectic form $\omega_{\mathbb{C}^3}$ to M yields a closed 2-form Ω whose restriction to each fibre is symplectic. Therefore, we can define the symplectic parallel transport of this fibration. Write $\omega = \Omega|_{M_\epsilon}$.

Definition 5.2.2. Winding once (anticlockwise) around the base S^1 yields the *symplectic monodromy map* $f \in \mathrm{Aut}(M_\epsilon, \partial M_\epsilon, \omega)$.

Lemma 5.2.3. *In the case $d = 1$, the monodromy, which we now denote by $\tau_{p,q}$, of $B_{p,q}$ induces the central twist $\tau \in \mathrm{Mod}(A)$ in the mapping class group of the punctured annulus A. More precisely, let γ and $\gamma' = \tau(\gamma)$ be the dotted and dashed (respectively) vanishing paths in Figure 5.2(τ). Form the associated vanishing thimbles D_γ and $D_{\gamma'}$. Then $\tau_{p,q}(D_\gamma)$ is Lagrangian isotopic to $D_{\gamma'} = D_{\tau(\gamma)}$.*

Proof. Consider the map $\varpi : M \to \mathbb{C} : \varpi(z, w) = z_3^p$. We have the following diagram of maps:

$$\begin{array}{ccc} M & \xrightarrow{\pi} & S^1 \\ \downarrow{\scriptstyle \varpi} & & \\ \mathbb{C} & & \end{array}$$

Restricting ϖ to each fibre M_w of π yields a Lefschetz fibration $\varpi_w := \varpi|M_w$ (except on the exotic fibre $\varpi_w^{-1}(0)$). The core idea of this proof is thinking about how the unique Lefschetz critical value $x_w \in \mathbb{C}^\times$ of ϖ_w winds around the origin as we go around the base S^1 of π.

For $t \in [0, 1]$ choose a smooth family γ_t of vanishing paths such that $\gamma_0 = \gamma$, $\gamma_1 = \gamma' = \tau(\gamma_0)$, $\gamma_t(1) = x_w$, and γ_t restricted to an interval of the form $[0, b]$ for

$b < 1$ agrees with $\gamma_0|_{[0,b]}$. Let $\tilde{X}$ be the horizontal lift (with respect to π and Ω) of the vector field $X(t) = 2\pi i e^{2\pi i t}$ and consider its flow μ_t. Note that μ_t maps M_w to $M_{e^{2\pi i t}w}$ and $\mu_1|_{M_\epsilon} = \tau_{p,q}$ by definition. Form the vanishing thimbles $D_{\gamma_t} \subset M_{e^{2\pi i t}w}$ and consider the Lagrangian isotopy

$$L_t := \mu_{-t}(D_{\gamma_t}),$$

which is contained in (M_ϵ, ω). Moreover, $L_0 = D_{\gamma_0}$ and $L_1 = \mu_{-1}(D_{\gamma_1}) = \tau_{p,q}^{-1}(D_{\tau(\gamma_0)})$. Therefore, it follows that $\tau_{p,q}(D_\gamma) \simeq D_{\gamma'}$, as desired. $\square$

Remark 5.2.4. Consider the circle action σ_t with weights $(1, dpq-1, q)$ on $\mathbb{C}^3$:

$$\sigma_t(z_1, z_2, z_3) = (e^{2\pi i t} z_1, e^{2\pi i (dpq-1)t} z_2, e^{2\pi i q t} z_3).$$

In the quotient $\mathbb{C}^3/\frac{1}{p}(1,-1,q)$ this descends to a circle action that factors through $\sigma_{t/p}$. The restriction of $\sigma_{t/p}$ to ∂M_w gives the boundary ∂M_w the structure of a Seifert fibration. Consider the Hamiltonian $H \in C^\infty(M_w, \mathbb{R})$:

$$H(z) = \pi(|z_1|^2 + (dpq-1)|z_2|^2 + q|z_3|^2).$$

Its time -1 flow ϕ^H_{-1} is a representative of the boundary Dehn twist on M_w (induced by the Seifert structure on ∂M_w). Since the Milnor fibration $\pi : M \to S^1$ is equivariant with respect to the circle actions $(\sigma_{t/p}, e^{2\pi d q i})$ on M and $e^{2\pi d q i}$ on S^1, one can apply an argument of Seidel [45, Lemma 4.16] (with $\sigma_{t/p}$ in place of σ_t and $\beta = dq$) to show that

$$[\phi^H_{-1}] = [f]^{dq}$$

in the symplectic mapping class group of (M_ϵ, ω).

We can view the $B_{p,q}$ monodromy explicitly inside $B_{d,p,q}$ as follows. Consider the singular manifold

$$\tilde{N} = \{(z_1, z_2, z_3) \in \mathbb{C}^3 \mid z_1 z_2 = z_3^p \prod_{i=2}^{d} (z_3^p - a_i)\},$$

and its quotient $N = \tilde{N}/\frac{1}{p}(1,-1,q)$. This is a partial smoothing of the $\frac{1}{dp^2}(1,dpq-1)$ singularity. The unique singular point $0 \in N$ is of type $\frac{1}{p^2}(1,pq-1)$. Similarly to above, we form the manifolds

$$\tilde{N}_t := \{(z_1,z_2,z_3) \in \mathbb{C}^3 \mid |z| \leq 1,\ z_1 z_2 = (z_3^p - \psi(|z|^2)e^{2\pi i t}a_1)\prod_{i=2}^{d}(z_3^p - \psi(|z|^2)a_i)\}$$

where $0 < a_1 < \ldots < a_d$ are sufficiently small real numbers, along with their quotients $N_t = \tilde{N}_t/\frac{1}{p}(1,-1,q)$. The family N_t is a smoothing of the singular manifold N with a unique $\frac{1}{p^2}(1,pq-1)$ singularity. Therefore the monodromy of N_t is the $B_{p,q}$ monodromy. Observe that, $N_t \cong B_{d,p,q}$. As in the proof of Lemma 5.2.3, projecting to z_3^p yields a Lefschetz fibration (up to the p-covered fibre over 0) $\varpi_t : N_t \to \mathbb{C}$, and the same proof shows that the monodromy of N_t induces the central twist $\tau \in \mathrm{Mod}(A)$. Since we primarily work with N_0, we abuse notation and continue to write $\varpi : N_0 \to \mathbb{C}$ for $\varpi_0 : N_0 \to \mathbb{C}$.

Recall the matching path γ_1 shown in Figure 5.1.

Lemma 5.2.5. *The $B_{p,q}$ symplectic monodromy $\tau_{p,q}$ acts non-trivially on $\pi_2(B_{d,p,q})$. Moreover, it acts with order p on the matching cycle Σ_{γ_1}.*

Proof. The quotient map $\tilde{c} : \tilde{N}_0 \to N_0$ is the (degree p) universal cover of $N_0 \cong B_{d,p,q}$. Consider the Lefschetz fibration fibration $\tilde{\varpi} : \tilde{N}_0 \to \mathbb{C} : \tilde{\varpi}(z_1,z_2,z_3) = z_3$ and the canonical p-to-1 branched cover $c : \mathbb{C} \to \mathbb{C} : c(x) = x^p$. Observe that the diagram

$$\begin{tikzcd} \tilde{N}_0 \arrow[r, "\tilde{c}"] \arrow[d, "\tilde{\varpi}"] & N_0 \arrow[d, "\varpi"] \\ \mathbb{C} \arrow[r, "c"] & \mathbb{C} \end{tikzcd} \tag{5.2.1}$$

commutes and realises the Lefschetz fibration $\tilde{\varpi} : \tilde{N}_0 \to \mathbb{C}$ as the pullback under c of $\varpi : N_0 \to \mathbb{C}$. Combined with the fact that $\tilde{c}$ respects the symplectic structures, this implies that $\tilde{c}$ commutes with the parallel transport of $\tilde{\varpi}$ and ϖ. In particular, given a matching path γ of ϖ and a lift $\tilde{\gamma}$ via c to a matching path for $\tilde{\varpi}$, the matching cycle $\Sigma_{\tilde{\gamma}} \subset \tilde{N}_0$ will be a lift of $\Sigma_\gamma \subset N_0$ with respect to $\tilde{c}$.

Since $\tilde{c}$ is the universal cover, it induces isomorphisms on homotopy groups $\pi_i(\tilde{N}_0) \cong \pi_i(N_0)$ for all $i > 1$. Moreover, as $\tilde{N}_0$ is simply connected, the Hurewicz

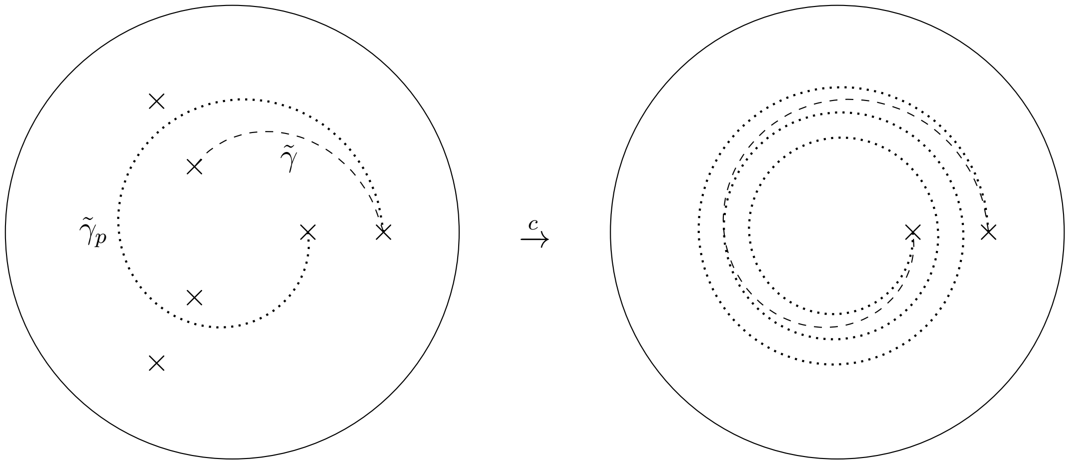

Figure 5.3: The dashed path $\tilde{\gamma}$ is a lift of the matching path $\tau(\gamma_1)$. The dotted path $\tilde{\gamma}_p$ is a lift of the path $\tau^p(\gamma_1)$. The example drawn here is $d = 2$ and $p = 3$.

theorem implies that $\pi_2(\tilde{N}_0) \cong H_2(\tilde{N}_0; \mathbb{Z}) \cong \mathbb{Z}^{dp-1}$. Therefore, the calculation of $(\tau_{p,q})_*[\Sigma_{\gamma_1}] = [\Sigma_{\tau(\gamma_1)}] \in \pi_2(N_0)$ reduces to computing the homology class of $\Sigma_{\tilde{\gamma}}$, where $\tilde{\gamma}$ lifts $\tau(\gamma_1)$.

To this end, we fix a basis of $H_2(\tilde{N}_0; \mathbb{Z})$ to be the collection of matching cycles corresponding to the following matching paths: the $p-1$ arcs, $\mu_j : [0,1] \to \mathbb{C}$: $\mu_j(s) = \sqrt[p]{a_1} e^{2\pi i(j+s)/p}$, for $0 \leq j < p-1$; and the $(d-1)p$ paths $\mu_{j,k} : [0,1] \to \mathbb{C}$, for $0 \leq j < p$ and $1 \leq k < d$, where $\mu_{0,k}$ is the straight path between the points $\sqrt[p]{a_k}$ and $\sqrt[p]{a_{k+1}}$, and $\mu_{j,k} = e^{2\pi i j/p} \mu_{0,k}$. The illustration in Figure 5.4 should dispel any confusion.

We have that $\Sigma_{\tilde{\gamma}}$ is isotopic to the Dehn twist $\tau_{\Sigma_{\mu_0}}(\Sigma_{\mu_{0,1}})$ of $\Sigma_{\mu_{0,1}}$ about Σ_{μ_0}, so we can use the Picard-Lefschetz formula to calculate the homology class:

$$[\tau_{\Sigma_{\mu_0}}(\Sigma_{\mu_{0,1}})] = [\Sigma_{\mu_{0,1}}] + ([\Sigma_{\mu_0}] \cdot [\Sigma_{\mu_{0,1}}])[\Sigma_{\mu_0}] = [\Sigma_{\mu_{0,1}}] + [\Sigma_{\mu_0}].$$

This proves that $(\tau_{p,q})_*$ acts non-trivially on $H_2(\tilde{N}_0; \mathbb{Z}) \cong \pi_2(B_{d,p,q})$.

To show that $(\tau_{p,q}^p)_*[\Sigma_{\gamma_1}] = [\Sigma_{\gamma_1}]$ it suffices to check that $[\Sigma_{\tilde{\gamma}_p}] = [\Sigma_{\mu_{0,1}}]$, where $\tilde{\gamma}_p$ is the lift of $\tau^p(\gamma_1)$ drawn in Figure 5.3. Denote by μ the matching path obtained by successively half-twisting μ_0 along μ_i for $i > 0$ (see Figure 5.4(b) for an example

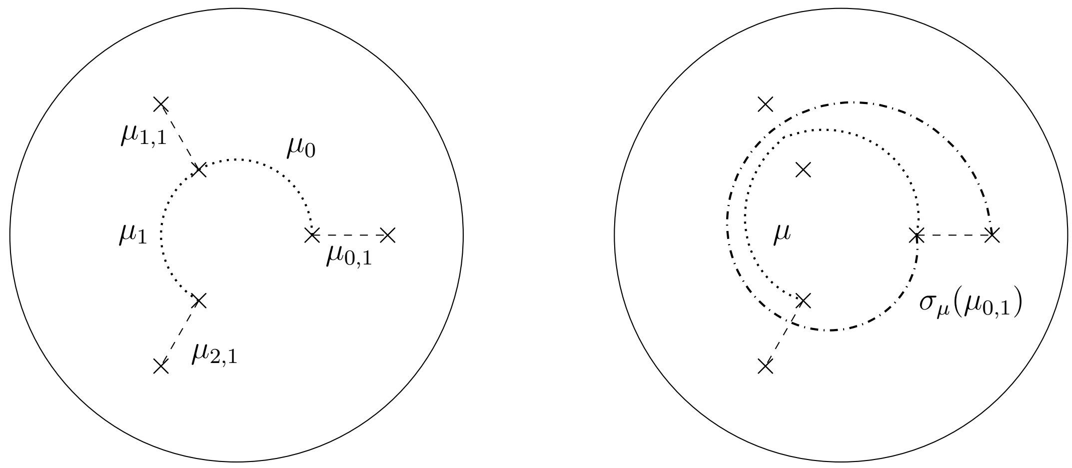

(a) (b)

Figure 5.4: (a) The matching paths corresponding to the chosen basis of $H_2(\tilde{N}_0;\mathbb{Z})$. (b) An A_3-configuration of matching paths: $\mu_{0,1}$ (dashed), $\mu = \sigma_{\mu_1}(\mu_0)$ (dotted), and $\mu_{2,1}$ (dashed); along with the matching path $\tilde{\gamma}_p = \sigma_\mu(\mu_{0,1})$ (dash dot). The case drawn here has $d = 2$, and $p = 3$.

with $d = 2$, and $p = 3$):

$$\mu := \sigma_{\mu_{p-2}}\sigma_{\mu_{p-3}}\cdots\sigma_{\mu_1}(\mu_0).$$

Then, $\tilde{\gamma}_p$ is isotopic to $\sigma_\mu^2(\mu_{0,1})$, which implies that the matching cycle $\Sigma_{\tilde{\gamma}_p}$ is Lagrangian isotopic to the Dehn twisted sphere $\tau^2_{\Sigma_\mu}(\Sigma_{\mu_{0,1}})$, which is *smoothly* isotopic to $\Sigma_{\mu_{0,1}}$. This implies that $[\Sigma_{\tilde{\gamma}_p}] = [\Sigma_{\mu_{0,1}}]$, which completes the proof. □

The previous lemma shows that $\tau^p_{p,q}(\Sigma_{\gamma_1}) \simeq \Sigma_{\tau^p(\gamma_1)}$ is homotopic to Σ_{γ_1}. However, they are not Lagrangian isotopic as we now show.

Proposition 5.2.6. *The matching cycle $\tau^p_{p,q}(\Sigma_{\gamma_1}) \simeq \Sigma_{\tau^p(\gamma_1)}$ is* not *Lagrangian isotopic to Σ_{γ_1}.*

Proof. Suppose that the result is false. Then, denoting $L = \Sigma_{\gamma_1}$ and $L' = \Sigma_{\tau^p(\gamma_1)}$, there exists a *Hamiltonian* isotopy $\phi^H_t \in \mathrm{Ham}(N_0,\omega)$ satisfying

$$\phi^H_1(L) = L'.$$

Then H lifts to the Hamiltonian $\tilde{H} := H \circ \tilde{c} : \tilde{N}_0 \to \mathbb{R}$, whose flow $\phi_t^{\tilde{H}}$ covers that of H:

$$\tilde{c} \circ \phi_t^{\tilde{H}} = \phi_t^H.$$

Therefore, by choosing the lift $\tilde{L} := \Sigma_{\mu_{0,1}}$ of L we must have that $\tilde{L}' := \phi_1^{\tilde{H}}(\tilde{L})$ is a lift of L'. Moreover, by the diagram in Equation (5.2.1), $\tilde{\varpi}(\tilde{L}')$ is (the image of) a lift of the matching path $\tau^p(\gamma_1)$. For topological reasons, this lift is exactly $\tilde{\gamma}_p$ as shown in Figure 5.3.

Picking up where we left off in the proof of Lemma 5.2.5, we deduce that $\tilde{L}'$ is Lagrangian isotopic to $\tau^2_{\Sigma_\mu}(\Sigma_{\mu_{0,1}})$. As noted, this is smoothly isotopic to $\tilde{L} = \Sigma_{\mu_{0,1}}$, but a famous result of Seidel [44] shows that it is not Lagrangian isotopic. Indeed, an A_3-configuration of Lagrangian spheres one can use to apply this result is given by the collection of matching cycles corresponding to the matching paths $\mu_{0,1}$, $\mu = \sigma_{\mu_{p-2}}\sigma_{\mu_{p-3}} \cdots \sigma_{\mu_1}(\mu_0)$, and $\mu_{p-1,1}$. Applying Seidel's theorem shows that the Lagrangian Floer cohomology $HF(\tilde{L}', \Sigma_{\mu_{p-1,1}})$ is non-zero,[3] so $\tilde{L}'$ cannot be Hamiltonian isotoped to $\tilde{L}$ since $HF(\tilde{L}, \Sigma_{\mu_{p-1,1}}) = 0$. This yields a contradiction, and hence our initial assumption that $L = \Sigma_{\gamma_1}$ and $L' = \Sigma_{\tau^p(\gamma_1)}$ were Lagrangian isotopic is false. □

Corollary 5.2.7. *The $B_{p,q}$ monodromy has infinite order in the symplectic mapping class group of $B_{d,p,q}$.*

Proof. A basic extension of the above proof shows that $L_k := \Sigma_{\tau^{kp}(\gamma_1)}$ is not Lagrangian isotopic to $L_0 = \Sigma_{\gamma_1}$ for all $k \neq 0$. Since $L_k \simeq \tau_{p,q}^{kp}(L_0)$ this proves the result. □

Remark 5.2.8. This is not a surprising result; it parallels the story of weighted homogeneous hypersurface singularities. What would be interesting is investigating whether it is truly a symplectic phenomenon or not. That is, is the monodromy of infinite order in the *smooth* mapping class group?

[3] Note how in Figure 5.4 the matching path $\sigma_\mu(\mu_{0,1})$ intersects $\mu_{p-1,1}$, whereas $\mu_{0,1}$ doesn't.

Note that in the case $(d, p, q) = (n+1, 1, 1)$ — which is the A_n du Val singularity — this is truly a symplectic phenomenon by Brieskorn's simultaneous resolution [6]. On the other hand, recent work of Konno, Lin, Mukherjee, and Muñoz-Echániz [32] shows that the monodromy diffeomorphism of the Milnor fibration for every weighted homogeneous hypersurface singularity *excluding* the ADE singularities has infinite order in the smooth mapping class group. Their theorem does not apply to the $B_{d,p,q}$ case since $b^+(B_{d,p,q}) = 0$ and $\pi_1(B_{d,p,q}) \neq 1$.

5.3 Proof of the main theorem

Consider the subgroup G of the symplectic mapping class group $\pi_0(\mathrm{Symp}_c(B_{d,p,q}))$ generated by the $d-1$ Dehn twists about the standard spheres $L_i = \Sigma_{\gamma_i}$ and the $\frac{1}{p^2}(1, pq-1)$ symplectic monodromy $\tau_{p,q}$. We now state the main theorem in its most precise form:

Theorem 5.3.1. *For every Lagrangian sphere $L \subset B_{d,p,q}$ there exists $\phi \in G$ such that L is Lagrangian isotopic to $\phi(L_1)$.*

We'll do this by combining the main result of Chapter 4 with a proof that the Lefschetz fibration $\varpi : N_0 \to \mathbb{C}$ in the previous section compactifies in a suitable sense to π_{ref}. We show that we can choose the reference almost complex structure J_{ref} on $X = X_{d,p,q}$ so that the fibres of the map $\pi_{J_{\mathrm{ref}}} : X \to \mathbb{C}$ are compactifications of those of $\varpi : N_0 \to \mathbb{C}$.[4] The upshot of this is that all the results proved in Section 5.2 transfer immediately to π_{ref}.

Lemma 5.3.2. *There exists a symplectic embedding $\iota : N_0 \hookrightarrow X = X_{d,p,q}$ and an almost complex structure $J_{\mathrm{ref}} \in \mathcal{J}(D')$ on X such that the fibres of $\varpi : N_0 \to \mathbb{C}$ compactify to J_{ref}-holomorphic fibres of the map $\pi_{J_{\mathrm{ref}}} : X \to \mathbb{C}$.*

Proof. The existence of this embedding will follow from showing that N_0 has an almost toric structure with base diagram isomorphic to that shown in Figure 1.1.

[4] Recall that π_{ref} is the deformation of $\pi_{J_{\mathrm{ref}}}$ under the Gompf argument of Lemma 4.3.2.

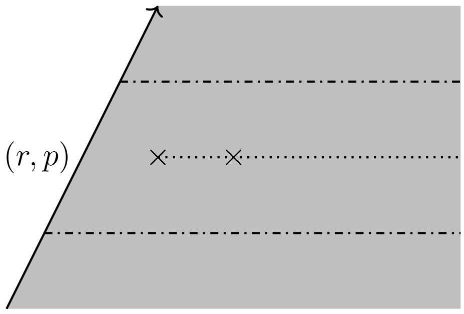

(a): N_0

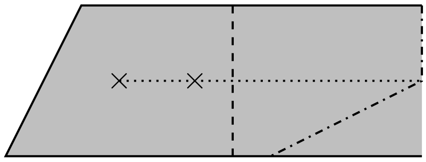

(b): $\bar{N}_0$

Figure 5.5: A fundamental domain derived from the Hamiltonian system $\mathbf{H} : N_0 \to \mathbb{R}^2$. Here r is the unique integer $0 < r < p$ such that $rq \equiv 1 \mod p$. Cuts made are shown in dash dot.

Indeed, the compactification X was constructed using only the almost toric base diagram. After this, we then investigate how the almost complex structure J on N_0, inherited from its embedding in $\mathbb{C}^3$, and the fibres of ϖ behave with respect to the Hamiltonian $H : N_0 \to \mathbb{R}$ used in the symplectic cut construction.

As in Lemma 7.2 of [17], the Hamiltonian system $\mathbf{H} = (|z_3|^2, \frac{1}{2}(|z_1|^2 - |z_2|^2))$ on N_0 has a fundamental domain whose image under action coordinates is that shown in Figure 5.5(a). In particular, the y-coordinate is the Hamiltonian $H(z) = \frac{1}{2}(|z_1|^2 - |z_2|^2)$ which generates the circle action $t \cdot (z_1, z_2, z_3) = (e^{it}z_1, e^{-it}z_2, z_3)$.

After applying a $SL_2(\mathbb{Z})$ transformation to make the toric boundary vertical and rotating the branch cut 180° clockwise, one obtains exactly the fundamental domain of Figure 1.1. Under this correspondence, the vertical cuts made in Figure 2.8(b) (which correspond to the J-holomorphic sections of the Lefschetz fibrations $\pi_J : X \to \mathbb{C}$) translate to horizontal cuts in the above fundamental domain. This means that the symplectic cut construction is done with respect to the Hamiltonian H. Recall that a fibre $\varpi^{-1}(w)$ is given by the equation

$$z_1 z_2 = \prod_{i=1}^{d} (w - \psi(|z|^2)a_i),$$

which is preserved by the Hamiltonian flow ϕ_t^H. Therefore, the intersection of a fibre with a regular level $H^{-1}(r)$ is given by an orbit of ϕ_t^H. This implies that, after taking the symplectic cuts, the fibres of ϖ become symplectic spheres in X.

As for the almost complex structure, note that in the region where the cut is made, N_0 is a complex submanifold of $\mathbb{C}^3/\frac{1}{p}(1,-1,q)$. Indeed, $N_0 \cap (\frac{2}{3} \leq |z|^2 \leq 1) = M_0 \cap (\frac{2}{3} \leq |z|^2 \leq 1)$ and M_0 is holomorphic. Therefore, as the circle action ϕ_t^H is a subgroup of the $U(3)$ action, this implies that J is invariant under ϕ_t^H and thus descends to give an almost complex structure $\bar{J}$ on the cut manifold $\bar{N}_0$ shown in Figure 5.5(b). To produce an almost complex structure on X we need to make a cut corresponding to the horizontal cut shown in Figure 2.8(a). However, *a priori*, the symplectic sphere introduced by this cut may not be $\bar{J}$-holomorphic. Therefore, we define J_{ref} as follows. Fix a closed subset of $\bar{N}_0$ containing all of the focus-focus critical points of the form indicated by the vertical dashed line in Figure 5.5(b), and define J_{ref} to be $\bar{J}$ here. Then, in a neighbourhood of the vertical cut passing through the branch cut define J_{ref} to be that coming from the standard product complex structure in $S^2 \times S^2$ (compare Section 4.3 and Figure 4.1). Finally, on the remaining region pick J_{ref} arbitrarily so that it makes the horizontal toric boundary J_{ref}-holomorphic (which is possible by Lemma A.1.3).[5]

The final claim that the fibres of ϖ compactify to fibres of $\pi_{J_{\text{ref}}}$ is equivalent to saying that the compactified spheres live in the homology class $F = H - S$ (in the basis computed in Lemma 2.2.1). Let $A \in H_2(X)$ be the homology class of the compactified fibres. The claim follows from the fact that $A \cdot E_i = 0$, $A \cdot \mathcal{E}_j = 0$, $A \cdot F = 0$, and $A \cdot S = 1$. $\square$

Corollary 5.3.3. *The symplectic embedding $\iota : N_0 \to X$ induces a commutative*

[5] Of course, one also needs to handle resolving the singularities introduced by the horizontal cuts, as in Lemma 2.1.3. However, this can be done so that the resulting resolution loci are J_{ref}-holomorphic without affecting the property that the fibres of $\varpi : N_0 \to \mathbb{C}$ away from the exotic fibre over 0 compactify to J_{ref}-holomorphic spheres in X.

square

$$\begin{tikzcd} N_0 \arrow[r, "\iota"] \arrow[d, "\varpi"] & X \arrow[d, "\pi_{J_{\mathrm{ref}}}"] \\ \mathbb{C} \arrow[r] & \mathbb{C} \end{tikzcd}.$$

Furthermore, this induces a bijection of sets of isotopy classes of matching paths of ϖ *and* $\pi_{J_{\mathrm{ref}}}$.

Proof. The existence of the commutative square is essentially just a rephrasing of Lemma 5.3.2. Notice that the restriction $\pi_{J_{\mathrm{ref}}}|_{X\setminus \operatorname{im}\iota}$ consists only of regular fibres by the choice of closed set on which J_{ref} agrees with the almost complex structure $\bar{J}$ coming from symplectic cutting N_0. This induces the claimed bijection of isotopy classes of matching paths, which, said another way, means that every matching cycle of $\pi_{J_{\mathrm{ref}}}$ is Lagrangian isotopic to a matching cycle of ϖ. □

Proof of Theorem 5.3.1. The symplectic completion of N_0 is symplectomorphic to $B_{d,p,q}$. Therefore, by applying the negative Liouville flow, we can assume that $L \subset N_0$ and thus, after compactifying, $L \subset X$. Therefore, by Corollary 4.3.3, L is Lagrangian isotopic to a matching cycle of π_{ref} in $\mathring{X}$. The results of Section 5.2 and Corollary 5.3.3 imply that this matching cycle is Lagrangian isotopic to a composition of Dehn twists and the $B_{p,q}$ symplectic monodromy applied to the matching cycle $L_0 = \Sigma_{\gamma_1}$. Finally, since the isotopy is supported in $\mathring{X}$, we can find a contact-type hypersurface M in $X\backslash D'$ such that the isotopy is contained in the *interior* $X_M \subset X\backslash D'$ of M and so that the completion of X_M is symplectomorphic to $B_{d,p,q}$. See Figure 5.6 for an example. Thereby we view the isotopy as taking place in $B_{d,p,q}$.

□

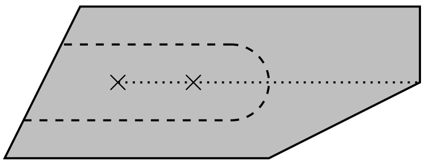

Figure 5.6: The dashed curve represents the image of the contact-type hypersurface $M \subset X \backslash D'$ chosen to ensure that the completion of the interior is symplectomorphic to $B_{d,p,q}$. We can ensure that the Lagrangian isotopy is contained in the interior of this hypersurface by altering the curve so that it doesn't touch any of the boundary edges except the leftmost one.

Appendix A

Constructions used in the text

A.1 Almost complex structures on symplectic divisors

We construct the almost complex structures necessary for the arguments of Section 2.4. Let (M, ω) be a closed symplectic 4-manifold and $S \subset M$ a symplectic divisor. In particular, throughout the whole of this section, we will assume the following:

- the components S_i of S are closed, embedded symplectic submanifolds;
- intersecting components S_i and S_j do so symplectically orthogonally;[1] and,
- there are no triple intersections, *i.e.* for distinct components S_i, S_j, S_k, we have $S_i \cap S_j \cap S_k = \emptyset$.

We will construct a compatible almost complex structure on M that realises each component of S as a J-holomorphic submanifold. The idea is to first construct J at the points of intersection, then extend in neighbourhoods of each component,

[1] This condition is not strictly necessary, but it makes the proofs simpler and it is satisfied in the situation we handle in this paper.

and finally to the whole of M. All the arguments presented here are variations on standard material, see for example [38, Chapters 2–3] and [8, Part III].

First consider the simplest situation where $S = S_1 \cup S_2$ is composed of two symplectic surfaces that intersect symplectically orthogonally exactly once. Note that, by the symplectic neighbourhood theorem [38, Theorem 3.4.10], a neighbourhood of S_1 in M is isomorphic to a neighbourhood of the zero section in the symplectic normal bundle νS_1 of S_1. Thus, we may choose a symplectic trivialisation of νS_1 in a neighbourhood of the unique intersection point $\{x\} = S_1 \cap S_2$. In other words, we have an embedding $\phi : U \to M$ of a neighbourhood U of $0 \in (\mathbb{R}^4, \omega_0)$, where $\omega_0 = \begin{pmatrix} 0 & I \\ -I & 0 \end{pmatrix}$ is the standard symplectic form, such that $\phi(\mathbb{R}^2 \times \{0\}) \subset S_1$. Furthermore, the symplectic orthogonality condition between S_1 and S_2 ensures that $\mathrm{d}\phi(0)(\{0\} \times \mathbb{R}^2) = T_x S_2$. The next step is to adjust the map ϕ to a symplectic embedding that maps the symplectic planes $\mathbb{R}^2 \times \{0\}$ and $\{0\} \times \mathbb{R}^2$ into S_1 and S_2 respectively, thus symplectically identifying the intersection $S_1 \cap S_2$ with the standard intersection of coordinate planes in $\mathbb{R}^4$. This is achieved by the following result.

Lemma A.1.1. *Let $W \subset (\mathbb{R}^{2n}, \omega_0) = (\mathbb{C}^n, \omega_0)$ be an embedded symplectic submanifold of the standard symplectic vector space that intersects the coordinate plane $\mathbb{C}^k \times \{0\}$ exactly once symplectically orthogonally at $0 \in \mathbb{C}^n$. Then there exists a symplectomorphism of a neighbourhood of $0 \in \mathbb{C}^n$ that maps the planes $\mathbb{C}^k \times \{0\}$ and $\{0\} \times \mathbb{C}^{n-k}$ onto $\mathbb{C}^k \times \{0\}$ and W respectively.*

Proof. Choose a symplectic embedding $i : B^{2(n-k)} \subset \mathbb{C}^{n-k} \to \mathbb{C}$ of a ball parametrising W and satisfying $i(0) = 0$. In the following, we will denote the planes $\mathbb{C}^k \times \{0\}$ and $\{0\} \times \mathbb{C}^{n-k}$ by $\mathbb{C}_{z_1}$ and $\mathbb{C}_{z_2}$ respectively. The symplectic neighbourhood theorem allows us to view a neighbourhood of the image of i as $i^*\nu W$. Note that the normal fibre $\nu_0 W$ over 0 is exactly $\mathbb{C}_{z_1}$. Moreover, the symplectic form ω_0 defines a connection $H \subset Ti^*\nu W$ on $i^*\nu W$ by taking the ω_0 complement of the vertical distribution. This implies that the symplectic form ω_0 splits over $i^*\nu W$ into vertical and horizontal components: $\omega_0 = \omega' \oplus \omega''$. Let $\langle X_j, Y_j \rangle$ be the standard symplectic basis of $T_0\mathbb{C}_{z_1} = \mathbb{C}_{z_1} = (T_0W)^\omega$ and extend this to a symplectic frame of $i^*\nu W$ via parallel transport. Using complex coordinates $z_1 = (x_1 + iy_1, \ldots, x_k + iy_k)$, define a map $\varphi : B^{2n} \subset \mathbb{C}^n = \mathbb{C}_{z_1} \times \mathbb{C}_{z_2} \to i^*\nu W$ by

$$\varphi(z_1, z_2) = \left(\sum_{j=1}^{k} x_j X_j(z_2) + y_j Y_j(z_2), z_2 \right).$$

Note that φ is a bundle isomorphism $\mathbb{C}^k \times B^{2(n-k)} \cong i^*\nu W$. Moreover, since $\nabla X_j = 0$ and $\nabla Y_j = 0$, φ pulls back the connection H on $i^*\nu W$ to the canonical one $TB^{2(n-k)} \subset T\mathbb{C}^k \oplus TB^{2(n-k)}$ on $\mathbb{C}^k \times B^{2(n-k)}$. Consequently, φ respects the splitting of $i^*\nu W$ induced by H, satisfies $\varphi^*\omega' = \omega_{\mathbb{C}^k}$, and $\varphi^*\omega'' = \omega_{B^{2(n-k)}}$, which yields

$$\varphi^*\omega_0 = \varphi^*(\omega' \oplus \omega'') = \omega_{\mathbb{C}^k} \oplus \omega_{B^{2(n-k)}} = \omega_0.$$

Hence, φ is the required symplectomorphism. □

Corollary A.1.2. *Let x be an intersection point of a symplectic divisor $S \subset (M, \omega)$, then there exists an integrable compatible almost complex structure on a neighbourhood of x that preserves each of the pieces of the divisor that intersect at x.*

Proof. The assumptions on the divisor S ensure that intersections between components are isolated and satisfy the assumptions of the local result of Lemma A.1.1. Therefore, we can push forward the standard complex structure i on $\mathbb{C}^2$ via a chart given by Lemma A.1.1. Since i preserves the coordinate planes, the result follows. □

We are now a position to choose a compatible almost complex structure J in a neighbourhood of all the intersection points of the divisor $S \subset (M, \omega)$. Then next step is to extend this to each of the components $S_i \subset S$ so that each S_i is J-holomorphic.

Lemma A.1.3. *Let $W \subset (M, \omega)$ be an embedded symplectic submanifold.*[2] *Then, by the symplectic neighbourhood theorem, W has a neighbourhood symplectomorphic*

[2] Here, W and M may be of any dimensions.

to its symplectic normal bundle $(\nu W, \omega)$. *Let* $H \subset T\nu W$ *be the canonical symplectic connection induced by* ω*, and suppose that we are given a compatible almost complex structure that splits* $J = J_H \oplus J_V$ *according to the connection* H*, and is defined over an open neighbourhood* $U \subset \nu W$ *of an open set of* W*. Then, for any open subset* $U' \subset \overline{U'} \subset U$*, there exists a compatible almost complex structure* $J' \in \mathcal{J}(T\nu W, \omega)$ *that agrees with* J *on* U' *and preserves the splitting* $T\nu W = H \oplus V$*. In particular,* W *is a* J'*-holomorphic submanifold.*

Proof. The non-degeneracy of the symplectic form ω splits the bundle $T\nu W$ into horizontal and vertical components (H, ω_H) and (V, ω_V), that is,

$$(T\nu W, \omega) = (H, \omega_H) \oplus (V, \omega_V).$$

The hypothesis on J says that it decomposes into two compatible almost complex structures $J_H \in \mathcal{J}(H|_U, \omega_H)$ and $J_V \in \mathcal{J}(V|_U, \omega_V)$. Therefore, to extend J to all of νW, it suffices to extend both J_H and J_V. This follows from the non-emptiness and contractibility of the space of compatible almost complex structures on any symplectic vector bundle, see [38, Proposition 2.6.4] for example. The non-emptiness and contractibility of $\mathcal{J}(H, \omega_H)$ and $\mathcal{J}(V, \omega_V)$, ensure that we can find global sections $J'_H \in \mathcal{J}(H, \omega_H)$ and $J'_V \in \mathcal{J}(V, \omega_V)$ that agree with J_H and J_V over $\overline{U'}$. Finally, since $J' := J'_H \oplus J'_V$ preserves the horizontal bundle H, and, for any point $x \in W$, $H_{(x,0)} = T_x W$, we have that W is a J'-holomorphic submanifold. This completes the proof. □

Remark A.1.4. Suppose that the normal bundle $(\nu W, \omega)$ above is trivial, *i.e.* isomorphic to $(W \times \mathbb{R}^{2m}, \omega_W \oplus \omega_{\mathbb{R}^{2m}})$. Then the horizontal distribution is simply $W \times \{x\} \subset W \times \mathbb{R}^{2m}$. Therefore, the extended almost complex structure J' preserves those subspaces, making the canonical sections $s_x : W \to W \times \mathbb{R}^{2m} : s_x(w) = (w, x)$ J'-holomorphic.

Returning to the situation of Section 2.4, since a neighbourhood of the F-component of the divisor D is symplectomorphic to a neighbourhood of a horizontal sphere in $S^2 \times S^2$, we can choose J near the points of intersection with the sections to be that coming from the corresponding intersections of horizontal and vertical spheres in $S^2 \times S^2$. Hence, the above implies that we can choose J near the F-component of the divisor D so that nearby curves in $\mathcal{M}_{0,0}(F; J)$ all intersect the sections ω-orthogonally.

Corollary A.1.5. *Given a symplectic divisor S in a symplectic 4-manifold (M, ω), there exists a compatible almost complex structure $J \in \mathcal{J}(M, \omega)$ that realises each component S_i of S as a J-holomorphic curve.*

Proof. Corollary A.1.2 and Lemma A.1.3 yield a compatible almost complex structure J' defined in a neighbourhood of S such that each component S_i is J'-holomorphic. Indeed, the set U' in the statement of Lemma A.1.3 can be chosen to be a shrunken neighbourhood of that given by Corollary A.1.2. Therefore, all that remains to do is extend J' to all of M, which again follows from [38, Proposition 2.6.4]. □

A.2 The mapping class group of the punctured annulus

The following argument is based on Adrien Brochier's MathOverflow answer [7]. Let S be a topological surface, and let $C(S, n)$ denote the configuration space of n points on S. Denote the 2-disc by D and recall that the braid group Br_n on n strands is defined to be the fundamental group of $C(D, n)$. Similarly, denoting the compact annulus by A, the *annular* braid group on n strands $Br_n(A)$ is defined to be $\pi_1(C(A, n))$.

Lemma A.2.1. *The annular braid group has the following presentation:*

$$Br_n(A) = \left\langle \tau, \sigma_1, \ldots, \sigma_{n-1} \;\middle|\; \begin{array}{c} (\tau\sigma_1)^2 = (\sigma_1\tau)^2, \\ \tau\sigma_i = \sigma_i\tau \;\forall i > 1, \\ \sigma_i\sigma_{i+1}\sigma_i = \sigma_{i+1}\sigma_i\sigma_{i+1}, \\ \sigma_i\sigma_j = \sigma_j\sigma_i \;\forall |i-j| > 1 \end{array} \right\rangle. \tag{A.2.1}$$

Proof. Recall the Artin presentation of Br_{n+1}:

$$Br_{n+1} = \left\langle \sigma_0, \ldots, \sigma_{n-1} \;\middle|\; \begin{array}{c} \sigma_i\sigma_{i+1}\sigma_i = \sigma_{i+1}\sigma_i\sigma_{i+1}, \\ \sigma_i\sigma_j = \sigma_j\sigma_i \ \forall|i-j| > 1 \end{array} \right\rangle. \tag{A.2.2}$$

By "filling in" the central hole of A with a punctured disc, we construct an injective homomorphism $i : Br_n(A) \to Br_{n+1}$ sending each annular braid to the corresponding planar one which fixes the 0th strand.[3] Consider the map $\phi : Br_{n+1} \to \Sigma_{n+1} = \mathrm{Sym}\{0, \ldots, n\}$ sending each braid to its permutation of the punctures. The image of i is $\phi^{-1}(\mathrm{Stab}(0))$, which has generators $\langle \sigma_0^2, \sigma_1, \ldots, \sigma_{n-1}\rangle$. Setting $\tau = \sigma_0^2$, we obtain the claimed generating set for $Br_n(A)$. The relation $(\tau\sigma_1)^2 = (\sigma_1\tau)^2$ follows from the braid relation $\sigma_0\sigma_1\sigma_0 = \sigma_1\sigma_0\sigma_1$. □

Denote the mapping class group of a surface S as $\mathrm{Mod}(S)$, interpreted as in [20, §2]. As with $\mathrm{Mod}(D^{\times n})$, where $D^{\times n}$ is the n-punctured disc $D^{\times n} = D \backslash \{n \text{ points}\}$, $\mathrm{Mod}(A^{\times n})$ is closely related to its braid group $Br_n(A)$.

Proposition A.2.2. *Let T_α denote the isotopy class of the Dehn twist about a simple closed curve α near the central boundary of $A^{\times n}$. Then*

$$\mathrm{Mod}(A^{\times n}) \cong Br_n(A) \times \langle T_\alpha \rangle = Br_n(A) \times \mathbb{Z}. \tag{A.2.3}$$

Proof. The result follows from the Birman exact sequence [20, Theorem 9.1]:

$$1 \longrightarrow \pi_1(C(A,n)) = Br_n(A) \longrightarrow \mathrm{Mod}(A^{\times n}) \longrightarrow \mathrm{Mod}(A) \longrightarrow 1. \tag{A.2.4}$$

We may apply this theorem since the identity component of $\mathrm{Homeo}^+(A, \partial A)$ is contractible [13, Theorem 1(D)], and so in particular, $\pi_1(\mathrm{Homeo}^+(A, \partial A)) = 1$. Since $\mathrm{Mod}(A) = \langle T_\alpha \rangle = \mathbb{Z}$ is free, this sequence splits and we obtain the result. □

[3] The strand associated with the central puncture in D.

References

[1] Mohammed Abouzaid, Mark McLean, and Ivan Smith. *Complex cobordism, Hamiltonian loops and global Kuranishi charts*. 2021. arXiv: 2110.14320v2 [math.SG].

[2] Michèle Audin. “Symplectic and almost complex manifolds”. In: *Holomorphic curves in symplectic geometry*. Vol. 117. Progr. Math. With an appendix by P. Gauduchon. Birkhäuser, Basel, 1994, pp. 41–74. DOI: 10.1007/978-3-0348-8508-9_3.

[3] Denis Auroux, Vicente Muñoz, and Francisco Presas. “Lagrangian submanifolds and Lefschetz pencils”. In: *J. Symplectic Geom.* 3.2 (2005), pp. 171–219. ISSN: 1527-5256.

[4] Frédéric Bourgeois. *A Morse-Bott approach to contact homology*. Thesis (Ph.D.)–Stanford University. ProQuest LLC, Ann Arbor, MI, 2002, p. 123. ISBN: 978-0493-62828-8.

[5] Frédéric Bourgeois et al. “Compactness Results in Symplectic Field Theory”. In: *Geometry & Topology* 7.2 (2003), pp. 799–888.

[6] Egbert Brieskorn. “Die Auflösung der rationalen Singularitäten holomorpher Abbildungen”. In: *Math. Ann.* 178 (1968), pp. 255–270. ISSN: 0025-5831. DOI: 10.1007/BF01352140.

[7] Adrien Brochier. *Mapping class group of n-punctured annulus*. MathOverflow. URL: https://mathoverflow.net/q/488830.

[8] Ana Cannas da Silva. *Lectures on symplectic geometry*. Vol. 1764. Lecture Notes in Mathematics. Springer-Verlag, Berlin, 2001, pp. xii+217. ISBN: 3-540-42195-5. DOI: 10.1007/978-3-540-45330-7.

[9] Kai Cieliebak and Klaus Mohnke. "Compactness for Punctured Holomorphic Curves". In: *Journal of Symplectic Geometry* 3.4 (2005), pp. 589–654.

[10] Charles C. Conley and Eduard J. Zehnder. "An index theory for periodic solutions of a Hamiltonian system". In: *Geometric dynamics (Rio de Janeiro, 1981)*. Vol. 1007. Lecture Notes in Math. Springer, Berlin, 1983, pp. 132–145. DOI: 10.1007/BFb0061415.

[11] Georgios Dimitroglou Rizell, Elizabeth Goodman, and Alexander Ivrii. "Lagrangian Isotopy of Tori in $S^2 \times S^2$ and $\mathbb{CP}^2 \times \mathbb{CP}^2$". In: *Geometric and Functional Analysis* 26 (2016), pp. 1297–1358.

[12] Johannes Jisse Duistermaat and Johan A. C. Kolk. *Lie groups*. Universitext. Springer-Verlag, Berlin, 2000, pp. viii+344. ISBN: 3-540-15293-8. DOI: 10.1007/978-3-642-56936-4.

[13] C. J. Earle and A. Schatz. "Teichmüller theory for surfaces with boundary". In: *J. Differential Geometry* 4 (1970), pp. 169–185. ISSN: 0022-040X.

[14] Yakov Eliashberg, Alexander Givental, and Helmut Hofer. "Introduction to Symplectic Field Theory". In: *Visions in Mathematics: GAFA 2000 Special Volume, Part II* (2010), pp. 560–673.

[15] Jonathan David Evans. "Lagrangian spheres in del Pezzo surfaces". In: *Journal of Topology* 3.1 (2010), pp. 181–227.

[16] Jonathan David Evans. "A Lagrangian Klein bottle you can't squeeze". In: *J. Fixed Point Theory Appl.* 24.2 (2022), Paper No. 47, 10. ISSN: 1661-7738. DOI: 10.1007/s11784-022-00945-w.

[17] Jonathan David Evans. *Lectures on Lagrangian torus fibrations*. Vol. 105. London Mathematical Society Student Texts. Cambridge University Press, Cambridge, 2023, pp. xiii+225. ISBN: 978-1-009-37262-6; 978-1-009-37263-3. DOI: 10.1017/9781009372671.

[18] Jonathan David Evans. *KIAS Lectures on Symplectic Aspects of Degenerations*. 2024. arXiv: 2403.03519 [math.SG].

[19] Jonathan David Evans and Ivan Smith. "Markov numbers and Lagrangian cell complexes in the complex projective plane". In: *Geom. Topol.* 22.2 (2018), pp. 1143–1180. ISSN: 1465-3060. DOI: 10.2140/gt.2018.22.1143.

[20] Benson Farb and Dan Margalit. *A primer on mapping class groups*. Vol. 49. Princeton Mathematical Series. Princeton University Press, Princeton, NJ, 2012, pp. xiv+472. ISBN: 978-0-691-14794-9.

[21] Robert E. Gompf. "A new construction of symplectic manifolds". In: *Ann. of Math. (2)* 142.3 (1995), pp. 527–595. ISSN: 0003-486X. DOI: 10.2307/2118554.

[22] Mikhael Gromov. "Pseudo holomorphic curves in symplectic manifolds". In: *Invent. Math.* 82.2 (1985), pp. 307–347. ISSN: 0020-9910. DOI: 10.1007/BF01388806.

[23] Richard Hind. "Lagrangian spheres in $S^2 \times S^2$". In: *Geometric & Functional Analysis GAFA* 14.2 (2004), pp. 303–318.

[24] Richard Hind. "Lagrangian unknottedness in Stein surfaces". In: *Asian J. Math.* 16.1 (2012), pp. 1–36. ISSN: 1093-6106. DOI: 10.4310/AJM.2012.v16.n1.a1.

[25] Helmut Hofer, Véronique Lizan, and Jean-Claude Sikorav. "On genericity for holomorphic curves in four-dimensional almost-complex manifolds". In: *J. Geom. Anal.* 7.1 (1997), pp. 149–159. ISSN: 1050-6926. DOI: 10.1007/BF02921708.

[26] Helmut Hofer, Krzysztof Wysocki, and Eduard Zehnder. "Properties of pseudoholomorphic curves in symplectisations II: Embedding controls and algebraic invariants". In: *Geom. Funct. Anal.* 5.2 (1995), pp. 270–328. ISSN: 1016-443X. DOI: 10.1007/BF01895669.

[27] Helmut Hofer, Krzysztof Wysocki, and Eduard Zehnder. "Properties of pseudoholomorphic curves in symplectisation IV: Asymptotics with degeneracies". In: *Contact and symplectic geometry (Cambridge, 1994)*. Vol. 8. Publ. Newton Inst. Cambridge Univ. Press, Cambridge, 1996, pp. 78–117.

[28] Helmut Hofer, Krzysztof Wysocki, and Eduard Zehnder. "Properties of pseudoholomorphic curves in symplectisations I: Asymptotics". In: *Ann. Inst. H. Poincaré C Anal. Non Linéaire* 13.3 (1996), pp. 337–379. ISSN: 0294-1449. DOI: 10.1016/s0294-1449(16)30108-1.

[29] Helmut Hofer, Krzysztof Wysocki, and Eduard Zehnder. "Finite energy foliations of tight three-spheres and Hamiltonian dynamics". In: *Ann. of Math. (2)* 157.1 (2003), pp. 125–255. ISSN: 0003-486X. DOI: 10.4007/annals.2003.157.125.

[30] Tatyana Khodorovskiy. *Symplectic Rational Blow-up*. 2013. arXiv: 1303.2581 [math.SG].

[31] Shoshichi Kobayashi and Katsumi Nomizu. *Foundations of differential geometry. Vol I*. Interscience Publishers (a division of John Wiley & Sons, Inc.), New York-London, 1963, pp. xi+329.

[32] Hokuto Konno et al. *The monodromy diffeomorphism of weighted singularities and Seiberg–Witten theory*. 2024. arXiv: 2411.12202 [math.GT].

[33] Eugene Lerman. "Symplectic cuts". In: *Math. Res. Lett.* 2.3 (1995), pp. 247–258. ISSN: 1073-2780. DOI: 10.4310/MRL.1995.v2.n3.a2.

[34] Samuel Lisi, Jeremy Van Horn-Morris, and Chris Wendl. *On Symplectic Fillings of Spinal Open Book Decompositions II: Holomorphic Curves and Classification*. 2020. arXiv: 2010.16330 [math.SG].

[35] Bruno Martelli. *Is the action of the mapping class group transitive on embedded arcs?* MathOverflow. eprint: https://mathoverflow.net/q/49964.

[36] Dusa McDuff. "The local behaviour of holomorphic curves in almost complex 4-manifolds". In: *J. Differential Geom.* 34.1 (1991), pp. 143–164. ISSN: 0022-040X.

[37] Dusa McDuff and Dietmar Salamon. *J-holomorphic Curves and Symplectic Topology*. eng. Second edition. Providence, Rhode Island: American Mathematical Society, 2012. ISBN: 9780821887462.

[38] Dusa McDuff and Dietmar Salamon. *Introduction to symplectic topology*. Third. Oxford Graduate Texts in Mathematics. Oxford University Press, Oxford, 2017, pp. xi+623. ISBN: 978-0-19-879490-5; 978-0-19-879489-9. DOI: 10.1093/oso/9780198794899.001.0001.

[39] Klaus Mohnke. *How to (symplectically) thread the eye of a (Lagrangian) needle*. 2003. arXiv: math/0106139 [math.SG].

[40] Eduardo Mora-Donato. *Pseudoholomorphic cylinders in symplectisations*. Thesis (Ph.D.)–New York University. ProQuest LLC, Ann Arbor, MI, 2003, p. 101. ISBN: 978-0496-37151-8.

[41] Joel Robbin and Dietmar Salamon. "The Maslov index for paths". In: *Topology* 32.4 (1993), pp. 827–844. ISSN: 0040-9383. DOI: 10.1016/0040-9383(93)90052-W.

[42] Paul Seidel. "Floer homology and the symplectic isotopy problem". PhD thesis. University of Oxford, 1997.

[43] Paul Seidel. *Symplectic automorphisms of T^*S^2*. 1998. arXiv: math/9803084 [math.DG].

[44] Paul Seidel. "Lagrangian two-spheres can be symplectically knotted". In: *J. Differential Geom.* 52.1 (1999), pp. 145–171. ISSN: 0022-040X.

[45] Paul Seidel. “Graded Lagrangian submanifolds”. In: *Bull. Soc. Math. France* 128.1 (2000), pp. 103–149. ISSN: 0037-9484.

[46] Paul Seidel. “A long exact sequence for symplectic Floer cohomology”. In: *Topology* 42.5 (2003), pp. 1003–1063. ISSN: 0040-9383. DOI: 10.1016/S0040-9383(02)00028-9.

[47] Paul Seidel. *Fukaya categories and Picard-Lefschetz theory*. Zurich Lectures in Advanced Mathematics. European Mathematical Society (EMS), Zürich, 2008, pp. viii+326. ISBN: 978-3-03719-063-0. DOI: 10.4171/063.

[48] Paul Seidel. “Lectures on four-dimensional Dehn twists”. In: *Symplectic 4-manifolds and algebraic surfaces*. Vol. 1938. Lecture Notes in Math. Springer, Berlin, 2008, pp. 231–267. DOI: 10.1007/978-3-540-78279-7_4.

[49] Richard Siefring. “Relative asymptotic behavior of pseudoholomorphic half-cylinders”. In: *Comm. Pure Appl. Math.* 61.12 (2008), pp. 1631–1684. ISSN: 0010-3640. DOI: 10.1002/cpa.20224.

[50] Richard Siefring. “Intersection theory of punctured pseudoholomorphic curves”. In: *Geom. Topol.* 15.4 (2011), pp. 2351–2457. ISSN: 1465-3060. DOI: 10.2140/gt.2011.15.2351.

[51] Alan Weinstein. “Symplectic manifolds and their Lagrangian submanifolds”. In: *Advances in Math.* 6 (1971), 329–346 (1971). ISSN: 0001-8708. DOI: 10.1016/0001-8708(71)90020-X.

[52] Chris Wendl. “Automatic transversality and orbifolds of punctured holomorphic curves in dimension four”. In: *Commentarii Mathematici Helvetici* 85.2 (2010), pp. 347–407.

[53] Chris Wendl. *Lectures on Holomorphic Curves in Symplectic and Contact Geometry*. 2014. arXiv: 1011.1690v2 [math.SG].

[54] Chris Wendl. *Lectures on Symplectic Field Theory*. 2016. arXiv: 1612.01009 [math.SG].

[55] Chris Wendl. *Holomorphic Curves in Low Dimensions. From symplectic ruled surfaces to planar contact manifolds.* Vol. 2216. Lecture Notes in Mathematics. Springer, Cham, 2018, pp. xiii+292. ISBN: 978-3-319-91369-8; 978-3-319-91371-1. DOI: 10.1007/978-3-319-91371-1.

[56] Chris Wendl. *Lectures on Contact 3-Manifolds, Holomorphic Curves and Intersection Theory.* First. Cambridge University Press, 2020. ISBN: 978-1-108-49740-4. DOI: 10.1017/9781108608954.

[57] Weiwei Wu. "Exact Lagrangians in A_n-surface singularities". In: *Math. Ann.* 359.1-2 (2014), pp. 153–168. ISSN: 0025-5831. DOI: 10.1007/s00208-013-0993-3.