

Surface-fibrations, four-manifolds, and symplectic Floer homology

Timothy Perutz

Imperial College, University of London

A thesis presented for the degree of Doctor of Philosophy
of the University of London

2005

Abstract

It is known that any smooth, closed, oriented four-manifold which is not negative-definite supports, after blowing up a finite set of points, a ‘broken fibration’. This is a singular surface-fibration of a special kind; its critical points are isolated points and circles. We construct an invariant of broken fibrations and begin a programme to study its properties. It has the same format as the Seiberg-Witten invariant of the underlying four-manifold, and is equal to it in some cases; we conjecture that this is always so. The invariant fits into a field theory for cobordisms equipped with broken fibrations over surfaces with boundary. The formal structure of this theory is very similar to that of a version of monopole Floer homology. The construction is purely symplectic and generalises that of the Donaldson-Smith standard surface count for Lefschetz fibrations. It involves moduli spaces of pseudoholomorphic sections of relative Hilbert schemes of points on the fibres, with Lagrangian boundary conditions which arise by a vanishing-cycle construction.

Broken fibrations form one part of what may be termed ‘near-symplectic geometry’. The other part concerns ‘near-symplectic forms’. The first chapter of the thesis reviews and develops the properties of these forms and their relation with broken fibrations.

Declaration

The material presented in this thesis is the author’s own, except where it appears with attribution to others.

Contents

Introduction	5
Geometric structures on four-manifolds	5
Near-symplectic topology, gauge theory, and pseudoholomorphic curves	9
1 Near-symplectic geometry	1
1.1 Near-symplectic forms	1
1.1.1 Topology associated with near-symplectic forms	13
1.2 Near-symplectic broken fibrations	23
2 Symmetric products and locally Hamiltonian fibrations	34
2.1 Relative symmetric products	34
2.1.1 Cohomology of relative symmetric products	36
2.2 Locally Hamiltonian fibrations	38
2.3 The vortex equations	41
2.3.1 Moduli spaces of vortices	41
2.3.2 The Kähler class on the vortex moduli space	46
2.3.3 Families of vortex moduli spaces	51
3 Invariants from pseudoholomorphic sections	58
3.1 Pseudoholomorphic sections and Lagrangian boundary conditions	59
3.1.1 Transversality	61
3.1.2 Compactness	64
3.1.3 Orientations	67
3.1.4 Invariance	70
3.2 Field theories	70
3.2.1 The theory in outline	72
3.2.2 Floer homology in the fibre-monotone case	78
3.2.3 Canonical relative gradings	80

3.2.4	Quantum cap product	81
3.2.5	Open-closed invariants	81
3.2.6	Weakly monotone fibres	85
4	Lagrangian matching invariants	89
4.1	Floer homology for relative symmetric products	90
4.2	Lagrangian matching invariants	98
4.2.1	Defining the invariants	103
4.2.2	Some simple calculations	107
4.3	Geometry of the compactified families	111
4.3.1	The Jacobian of a compact nodal curve	112
4.3.2	Compactifying the Picard family	114
4.3.3	The Hilbert scheme of r points on a nodal curve	117
4.3.4	The relative Hilbert scheme	121
4.4	Constructing the Lagrangian boundary conditions	125
4.4.1	Vanishing cycles as Lagrangian correspondences	125
4.4.2	Some auxiliary spaces	127
4.5	Floer homology of the Jacobian families	133
4.6	An exact triangle?	135
A	Symmetric products of non-singular complex curves	138
B	Topology of mapping tori	141

Introduction

Geometric structures on four-manifolds

Two-forms and almost complex structures

A possible approach to understanding the four-dimensional smooth category, with its enigmas, is to begin with a restricted class of manifolds, sufficiently tractable that at least some of the mysteries can be resolved, and then progressively to enlarge the class. A good way to do this is via triples (X, ω, J) , where the closed four-manifold X has been equipped with a closed two-form $\omega \in Z_X^2$ and an almost complex structure $J \in \text{End}(TX)$. The first class to consider is of triples which make X a compact Kähler surface.

General four-manifolds do not support almost complex structures, and we shall consider situations where J is defined only on an dense open subset $U \subset X$. We always insist that, where it is defined, it tames ω :

$$\omega_x(v, Jv) > 0 \quad \text{for every non-zero } v \in T_x U,$$

so that $\omega|_U$ is symplectic. Usually we also impose the compatibility relation

$$\omega_x(Jv, v') + \omega_x(v, Jv') = 0,$$

which means that $g := \omega(\cdot, J\cdot)$ is a Riemannian metric. Then $(U, \omega|_U, J)$ is, by definition, an almost Kähler manifold. Global almost Kähler structures are much more plentiful than the Kähler ones, where J is required to be integrable and ω of type $(1, 1)$; many compact symplectic four-manifolds, even simply connected ones, do not support integrable complex structures. However, the property of supporting a symplectic structure is an even subtler one (to give a simple example, a necessary but insufficient condition is that, in any connected sum decomposition, one of the summands has $b_2^+ = 0$).

We can relax the almost-Kähler condition, allowing ω to be a **near-symplectic form**. This means that at each point x either $\omega_x^2 > 0$ or $\omega_x = 0$, and that along its zero-locus $Z_\omega = \omega^{-1}(0)$, ω satisfies the transversality condition that $\nabla_x \omega: T_x X \rightarrow \Lambda^2 T_x^* X$ has rank

3 (which is the largest possible). This implies that Z_ω is an embedded 1-submanifold. We then ask that $(X \setminus Z_\omega, \omega|_{(X \setminus Z_\omega)}, J)$ be almost Kähler, and that the associated metric g extends to a metric on X (J will then *not* extend).

It has been known for several years that *near-symplectic structures* (X, ω, J) , *compatible with a chosen orientation of X , exist if and only if there exists a class $c \in H^2(X; \mathbb{R})$ with $c^2[X] > 0$* . This is a consequence of Hodge theory, together with a non-trivial transversality argument. If c_g is the harmonic representative of such a c , for the Riemannian metric g , then its self-dual part $\omega_{c,g} = c_g^+$ is a near-symplectic form for generic g .

Pencils and fibrations

A **topological Lefschetz pencil** (X, B, π) on X is given by a discrete subset $B \subset X$ and a smooth map $\pi: X \setminus B \rightarrow S$ to a smooth oriented surface S such that the fibres of π have compact closures in X . Each point $b \in B$ must have a neighbourhood U_b such that the map $U_b \setminus \{b\} \rightarrow \pi(U_b \setminus \{b\})$ is equivalent, as a germ of maps of oriented manifolds, to the map $\mathbb{C}^2 \setminus \{0\} \rightarrow \mathbb{C}P^1$, $(z_1, z_2) \mapsto (z_1 : z_2)$. Each point $p \in X \setminus B$ has a neighbourhood U_p such that $\pi: (U_p, p) \rightarrow (\pi(U_p), \pi(p))$ is equivalent to one of two maps $(\mathbb{C}^2, 0) \rightarrow (\mathbb{C}, 0)$ —either $(z_1, z_2) \mapsto z_1 z_2$, or $(z_1, z_2) \mapsto z_1$. It is a **topological Lefschetz fibration** when $B = \emptyset$.

There is a simple device which converts a Lefschetz pencil over S^2 into a Lefschetz fibration: one blows up X along B (using local complex coordinates near B as in the model above) and composes π with the blow-down map $\sigma: \widehat{X} \rightarrow X$. This has a unique extension to a Lefschetz fibration $\pi \circ \sigma: \widehat{X} \rightarrow S^2$, the exceptional divisors appearing as sections.

In the preprint [3], Auroux, Donaldson and Katzarkov introduce the following generalisation of the notion of topological Lefschetz pencil. I have adopted Ivan Smith's punning suggestion for its name.

Definition 0.0.1. A **broken pencil** (X, B, π) consists of a smooth oriented 4-manifold X , a 0-submanifold $B \subset X$ and a smooth map $\pi: X \setminus B \rightarrow S$ to an oriented surface, such that

- The fibres of π have compact closures in X .
- The critical set $\text{crit}(\pi) \subset X \setminus B$ is a disjoint union of a 1-submanifold Z and a discrete set of points D . Each point $b \in B$ must have an open neighbourhood U_b such that the map $U_b \setminus \{b\} \rightarrow \pi(U_b \setminus \{b\})$ is equivalent, as a germ of maps of oriented manifolds, to the map $\mathbb{C}^2 \setminus \{0\} \rightarrow \mathbb{C}P^1$, $(z_1, z_2) \mapsto (z_1 : z_2)$. Each point $p \in D$ has an open neighbourhood U_p such that $\pi: (U_p, p) \rightarrow (\pi(U_p), \pi(p))$ is equivalent to the map $(\mathbb{C}^2, 0) \rightarrow (\mathbb{C}, 0)$, $(z_1, z_2) \mapsto z_1 z_2$.
- For any component $Z_i \subset Z$, $\pi|_{Z_i}$ is injective; and for any other component Z_j , either $\pi(Z_i) = \pi(Z_j)$ or $\pi(Z_i) \cap \pi(Z_j) = \emptyset$.

- Each point $p \in Z$ has a neighbourhood U_p such that $\pi: (U_p, p) \rightarrow (\pi(U_p), \pi(p))$ is equivalent to a map

$$\mathbb{R}^3 \times \mathbb{R} \rightarrow \mathbb{R}^2, \quad (x, t) \mapsto (q(x), t),$$

where $q: \mathbb{R}^3 \rightarrow \mathbb{R}$ is a non-degenerate quadratic form of signature $(2, 1)$.

A **broken fibration** is a pair (X, π) such that (X, \emptyset, π) is a broken pencil.

In the above, ‘equivalence’ is of smooth germs, respecting the orientations of X and S . So the last point means that there is an orientation-preserving diffeomorphism Ψ from U_p to an open subset of \mathbb{R}^4 , with $\Psi(p) = 0$, and an orientation-preserving diffeomorphism ψ from an open subset neighbourhood of $\pi(p)$ in S to an open subset of \mathbb{R}^2 , with $\psi(\pi(p)) = 0$, such that the equation $\psi \circ f \circ \Psi^{-1}(x, t) = (q(x), t)$ holds on some neighbourhood of $0 \in \mathbb{R}^4$.

The subset $\pi(Z) \subset S$ is then a 1-submanifold. In a neighbourhood of Z_i , the map is equivalent—as germ, preserving all the orientations—to one of two model maps $\mathbb{R}^3 \times S^1 \rightarrow \mathbb{R} \times S^1$ which we shall write down later.

The relation of broken pencils to near-symplectic forms is established in [3]. It is, roughly speaking, that of topological Lefschetz pencils to symplectic forms. We briefly recall part of that relationship. It runs in two directions. An *integral* symplectic form ω is the curvature of a hermitian line bundle $L \rightarrow X$, and there exist asymptotically holomorphic sections s_k^0, s_k^1 of $L^{\otimes k}$ such that $(s_k^0 : s_k^1)$ defines a Lefschetz pencil over \mathbb{CP}^1 when $k \gg 0$; its fibres represent the homology class dual to $k[\omega]$. If ω is not an integral class, then one can still find a sequence of pencils over S^2 with symplectic fibres. In the opposite direction, Lefschetz pencils (subject to the condition that some class in H^2 evaluates positively on all components of the fibres) carry symplectic forms, canonical up to deformation, with symplectic fibres.

For the near-symplectic generalisation, we shall concentrate on the non-integral case, referring to [3] for the integral one from which it is derived, and for the proof. To formulate it, we need a notion of compatibility for a near-symplectic form $\omega \in Z_X^2$ and a broken pencil (X, B, π) . This requires that Z , the 1-dimensional part of $\text{crit}(\pi)$, coincides with $\omega^{-1}(0)$ and that ω is non-degenerate on $\ker(D_x \pi)$ for every regular point $x \in X \setminus B$, and also some subtler conditions: There are natural orientations for $\omega^{-1}(0)$ and for Z , and these must match. Both π and ω determine normal sub-bundles $N_{Z/X} \subset TX|_Z$, carrying quadratic forms S_π and S_ω (both defined up to conformal equivalence); these data must coincide.

Theorem 0.0.2 (Auroux-Donaldson-Katzarkov). *Let X be a closed, oriented four-manifold.*

(a) *Suppose that $\omega \in Z_X^2$ is a near-symplectic form with zero-set Z . Then there exists a broken pencil (X, B, π) over S^2 , compatible with ω , and enjoying two extra topological properties: $\pi(Z)$ is connected, and π is **directional**.*

(b) *Let (X, B, π) be a broken pencil. Assume that there is a class $c \in H^2(X; \mathbb{R})$ which evaluates positively on any homology class represented by the closure of a component of a regular fibre. Then there is a distinguished deformation-class of near-symplectic forms on X , representing c and compatible with the pencil.*

We explain the term ‘directional’. Let Γ be a component of $\pi(Z)$ (in the context of the theorem, there is only one). Let $\gamma: [-\epsilon, \epsilon] \rightarrow S$ be a short, embedded path crossing $\pi(Z)$ transversely (say $\gamma(0) \in \pi(Z)$). Then $W := \pi^{-1}(\text{im}(\gamma))$ is a 3-manifold with boundary, and $\gamma^{-1} \circ \pi: W \rightarrow [-\epsilon, \epsilon]$ has Morse critical points, each of index 1 or 2. We say Γ is directional if all of the indices are equal to 1 or all are equal to 2. It is clear that this property is independent of the choice of path γ . The fibration is directional if every component Γ is directional.

This theorem opens up a multitude of questions. For example:

- Is there a set of standard operations that can be used to relate different broken pencils on the same four-manifold? Is there a direct topological proof of the theorem?
- Do broken pencils shed light on the minimal genus problem? In particular, is it true that an embedded surface with transverse, locally positive intersections with the fibres minimises genus within its homology class?
- Do broken pencils with $b_2^+ > 1$ have Seiberg-Witten simple type?

There is also a generic question: how does one explain some chosen aspect of four-manifold topology in terms of broken pencils? One such question asks for a geometric (more particularly, a pseudoholomorphic) formulation of Seiberg-Witten theory on a broken fibration. In this thesis we begin a programme to address this problem. We should remark that it has already been solved in the case of Lefschetz pencils on closed 4-manifolds: Donaldson and Smith [9] define a ‘standard surface count’ as a Gromov invariant for pseudoholomorphic sections of a bundle of symmetric products associated to the Lefschetz fibration; Usher [52] shows that it is equal to the Gromov invariant of the underlying symplectic 4-manifold; and Taubes (in a sequence of papers running to some 400 pages) shows that the Gromov invariant is the Seiberg-Witten invariant of the underlying 4-manifold.

We extend the Donaldson-Smith invariant in two directions. One is to fit it into a TQFT for fibrations over compact surfaces with boundary. The second—which is the main innovation in the thesis—is to extend the theory to handle broken fibrations. It is not at all obvious how one should so, and our approach is in some ways an artificial one, in that it relies on a trick to produce Lagrangian boundary conditions for the pseudoholomorphic curves. Constructing invariants from Lagrangian boundary conditions involves, in general, ferocious technical problems; in our case the technicalities are quite mild, at least under certain numerical hypotheses.

Near-symplectic topology, gauge theory and pseudoholomorphic curves

Geometric interpretations. In algebraic topology, a question one asks about a generalised homology theory is which geometric objects (submanifolds, vector bundles, etc.) naturally represent homology classes. The next question is whether the theory can be built using those objects alone. Analogous questions arise in gauge theory (for example, Donaldson or Seiberg-Witten theory of 4-manifolds); one looks for interpretations which reflect particular geometric structures on the 4-manifold.

Such interpretations are well-known on Kähler surfaces. The instanton invariants can be computed using moduli spaces of stable holomorphic vector bundles, and the Seiberg-Witten invariants \mathcal{SW}_X from spaces of smooth divisors. As Taubes demonstrated, the latter remains true on symplectic four-manifolds: \mathcal{SW}_X is equal to a Gromov-Witten invariant which counts pseudoholomorphic curves. It can be computed using only smoothly embedded curves and their unramified coverings.

Seiberg-Witten theory on non-symplectic four-manifolds is not so well understood. However, Taubes has given a geometric realisation of solutions to the Seiberg-Witten equations on a four-manifold equipped with a near-symplectic form ω : from a family of solutions of the equations (with a perturbation term $r\omega$, and r tending to infinity), one can produce, roughly speaking, a pseudoholomorphic curve C in X with interior in $X \setminus Z_\omega$ and boundary Z . A useful framework here is to choose a contact form α on the boundary of a tubular neighbourhood of Z_ω , with $\omega = d\alpha$. The ends of $X \setminus Z_\omega$ are then *concave* ends of the symplectisation of α . A better description of C is as a finite-area pseudoholomorphic curve in $X \setminus Z_\omega$, asymptotic to periodic Reeb trajectories of α . Objects of this kind are studied in symplectic field theory, and they have a Fredholm deformation theory. Taubes' programme is aimed at constructing a near-symplectic invariant from their moduli spaces and comparing it to \mathcal{SW} .

Though pseudoholomorphic curves of the kinds considered by Taubes are not used directly in this thesis, they do *motivate* our main construction. Our approach is, loosely speaking, parallel to his: in place of contact geometry and Reeb boundary conditions, we use fibred geometry and Lagrangian boundary conditions.

Let $\text{Spin}^c(X)$ denote the $H^2(X; \mathbb{Z})$ -torsor of (isomorphism classes of) Spin^c -structures on the oriented four-manifold X . When X is given a structure of broken fibration, with $F \in H_2(X; \mathbb{Z})$ the class of a regular fibre, we write

$$\text{Spin}^c(X)_n = \{\mathfrak{s} \in \text{Spin}^c(X) : \langle c_1(\mathfrak{s}), F \rangle = 2n, (*)\},$$

where $(*)$ is the condition that for any connected component Σ of any regular fibre, one has $\langle c_1(\mathfrak{s}), [\Sigma] \rangle \geq \chi(\Sigma)$.

Construction 0.0.3. *To any broken fibration (X, π) over S^2 , directional, with $\pi(Z)$ connected, and with connected fibres, we associate an invariant $\mathcal{L}_{(X, \pi)}$, the **Lagrangian matching invariant**. This is a map*

$$\bigcup_{n \geq 1} \text{Spin}^c(X)_n \rightarrow \mathbb{A}(X), \quad \mathfrak{s} \mapsto \mathcal{L}_{(X, \pi)}(\mathfrak{s}).$$

Here $\mathbb{A}(X)$ is the graded abelian group underlying the graded ring

$$\mathbb{Z}[U] \otimes_{\mathbb{Z}} \Lambda^* H^1(X; \mathbb{Z}), \quad \deg(U) = 2.$$

The element $\mathcal{L}_{(X, \pi)}(\mathfrak{s})$ is homogeneous of degree $d(\mathfrak{s}) = (c_1(\mathfrak{s})^2 - 2e(X) - 3\sigma(X))/4$.¹ It is invariant under isotopies of π , and equivariant under isomorphisms $(X, \pi, Z) \cong (X', \pi', Z')$.

As well as $n \geq 1$, one can also let n take certain negative values (detailed in Chapter 4). I expect that the case $n = 0$ can be included, by means of arguments which we will sketch in the text. The hypothesis that the fibres be connected can also be dropped, at the price of imposing further restrictions on the Spin^c -structures in the domain.

The format of this invariant is familiar from Seiberg-Witten theory: assuming $b_2^+(X) > 1$, the Seiberg-Witten invariant is a map $\mathcal{SW}_X: \text{Spin}^c(X) \rightarrow \mathbb{A}(X)$, and $\mathcal{SW}_X(\mathfrak{s})$ is homogeneous of degree $d(\mathfrak{s})$. The overall sign of \mathcal{SW}_X depends on a choice of ‘homology orientation’, an orientation for the vector space $H^1(X; \mathbb{R}) \oplus H^+(X; \mathbb{R})$ (the second term is any maximal positive-definite subspace of $H^2(X; \mathbb{R})$). A compact almost complex manifold has a canonical homology orientation. There is also a canonical homology orientation for a four-manifold with a near-symplectic form (or broken fibration), because it is the union of a symplectic manifold with contact boundary, to which one can apply Kronheimer and Mrowka’s recipe [21], and a universal piece (a tubular neighbourhood of the zero-set of the form).

Conjecture 0.0.4. $\mathcal{L}_{X, \pi} = \varepsilon \cdot \mathcal{SW}_X$ on $\text{Spin}^c(X)_d$, $d \geq 1$ for a universal sign $\varepsilon \in \{\pm 1\}$. In particular, $\mathcal{L}_{X, \pi}$ depends only on X and not on π .

When $b_2^+(X) = 1$, we calculate \mathcal{SW}_X in the ‘Taubes chamber’ of a compatible near symplectic-form.

The conjecture holds for Lefschetz fibrations on symplectic manifolds, but that is not new. In Chapter 4 we will give further examples where the conjecture holds, for certain kinds of fibrations where the underlying four-manifolds are (i) $S^1 \times M^3$, for any \mathbb{Z} -homology- $(S^1 \times S^2)$ M (here the invariants are equivalent to the Alexander polynomial of M); (ii) connected sums (here the invariants vanish).

Assuming the conjecture is true, the Lagrangian matching invariant gives a symplectic interpretation for $\mathcal{SW}_X(\mathfrak{s})$ for most Spin^c -structures (all if we can plug the gap $d = 0$),

¹ $d(\mathfrak{s}) \in \mathbb{Z}$ because $c_1(\mathfrak{s})$, a characteristic element, satisfies $c_1(\mathfrak{s})^2 \equiv \sigma \pmod{8}$.

providing that one can find a broken fibration with connected fibres.² For in computing \mathcal{SW}_X one can use two basic properties of Seiberg-Witten invariants, ‘conjugation invariance’ and the blow-up formula. Conjugation invariance says that

$$\mathcal{SW}_X(\mathfrak{s}) = (-1)^{\frac{1}{8}(c_1(\mathfrak{s})^2 - \sigma) + 1 + b_1 - b_2^+} \mathcal{SW}_X(\bar{\mathfrak{s}}).$$

Now $c_1(\bar{\mathfrak{s}}) = -c_1(\mathfrak{s})$, and consequently there are only certain values of d for which the analogous statement for Lagrangian matching invariant can be formulated. On the other hand, if the conjecture is correct it tells us that, provided all the fibres of π are connected, the Seiberg-Witten invariant is determined by its restriction to $\bigcup_{d \geq 0} \text{Spin}^c(X)_d$.

The blow-up formula implies that we can compute the invariants of the manifold underlying a broken pencil from those of the broken fibration obtained by blowing up the basepoints.

Idea of the construction. When (X, π) is a *Lefschetz* fibration over a closed surface S , $\mathcal{L}_{X, \pi}$ is (by definition) its Donaldson-Smith invariant. We briefly recall how that is defined. If π is a submersion, the r -fold relative symmetric product $\text{Sym}_S^r(X)$ (the quotient of the r -fold fibre product by the symmetric group) has a structure of smooth manifold, for each $r > 0$. This fails when π has nodal fibres, but there is a natural ‘resolution of singularities’ $\text{Hilb}_S^r(X) \rightarrow \text{Sym}_S^r(X)$, the ‘relative Hilbert scheme of r points’, which is not only smooth but also supports a canonical deformation-class of symplectic structures. The set of homotopy classes of sections of $\text{Hilb}_S^r(X) \rightarrow S$ may (usually) be identified with the set of classes $\beta \in H_2(X; \mathbb{Z})$ whose intersection number with the fibre is r . Now $\mathcal{L}_{X, \pi}(\beta \cdot \mathfrak{s}_{\text{can}})$ is defined as a Gromov invariant, ‘counting’ pseudoholomorphic sections of the Hilbert scheme in the class β . For reasons of technical convenience we assume that $r \geq g - 1$, where g is the fibre-genus.

In the Donaldson-Smith theory, $\text{Hilb}_S^r(X)$ serves as an amenable compactification of the relative symmetric product over the smooth locus; its geometric structure does not greatly affect things. In extending the construction to broken fibrations, relative Hilbert schemes also play a crucial role, and one which directly involves the structure of the special fibres. This is curious in that there is no obvious reason why Hilbert schemes should have a role here at all.

To simplify the exposition, suppose that we have a near-symplectic fibration (X, π) over S^2 , with just one singular circle Z mapping to an ‘equator’ $\pi(Z)$. Thicken $\pi(Z)$ to a narrow, closed annulus T (the ‘tropics’); thus $S^2 \setminus \text{int}(T) = N \sqcup S$, the northern and southern hemispheres. Above them we have $X^N = \pi^{-1}(N)$ (where the fibres have genus g , say), $X^S = \pi^{-1}(S)$ (fibre-genus $g-1$). We can consider the Hilbert schemes $\text{Hilb}_N^r(X^N)$ and $\text{Hilb}_S^{r-1}(X^S)$, which can be made into symplectic manifolds with boundary. The respective

²Determining whether these exist for general four-manifolds is an interesting problem.

boundaries are $Y_r^N := \text{Sym}_{S^1}^r(Y^N)$ and $Y_{r-1}^S := \text{Sym}_{S^1}^{r-1}(Y^S)$, where $Y^N = \partial X^N$ and $Y^S = \partial X^S$.

The construction hinges on finding a middle-dimensional sub-fibre bundle \mathcal{Q} of the fibre product $Y_r^N \times_{S^1} Y_{r-1}^S \rightarrow S^1$, with the properties that (i) projection $\mathcal{Q} \rightarrow Y_r^N$ is an embedding; (ii) projection $\mathcal{Q} \rightarrow Y_{r-1}^S$ is a trivial S^1 -bundle; (iii) there are symplectic forms Ω_N on $\text{Hilb}_N^r(X^N)$ and Ω_S on $\text{Hilb}_S^{r-1}(X^S)$ such that \mathcal{Q} is isotropic with respect to $-(\Omega_N|_{Y_r^N}) \oplus +(\Omega_S|_{Y_{r-1}^S})$. We then define $\mathcal{L}_{X,\pi}$ as an ‘open Gromov invariant’ for pairs (u_N, u_S) of pseudoholomorphic sections of $\text{Hilb}_N^r(X^N)$ and $\text{Hilb}_S^{r-1}(X^S)$ such that the boundary values $(u_N|_{\partial N}, u_S|_{\partial S})$ lie in the ‘Lagrangian boundary condition’ \mathcal{Q} . We treat separately the different homotopy-classes of pairs of sections; these correspond to elements of $\text{Spin}^c(X)_d$, where $d = 1 - g + r = 1 - (g - 1) + (r - 1)$.

When $r = 1$, the Lagrangian boundary condition $\mathcal{Q} := \mathcal{Q} \subset Y^N$ can be thought of as a vanishing cycle for the near-symplectic singularity Z ; it is a torus or Klein bottle formed from circles in the fibres of $Y^N \rightarrow S^1$ which shrink to points of Z .³ For higher r , \mathcal{Q} is constructed as an S^1 -family of vanishing cycles for an S^1 -family of Hilbert schemes of elementary Lefschetz fibrations over disks. At an intuitive level (or at the level of homotopy classes), pairs (u_N, u_S) with boundary on \mathcal{Q} should be thought of as surfaces $\bar{u}_N \subset X_N$ and \bar{u}_S , where \bar{u}_N has one boundary component on \mathcal{Q} . The remaining $(r - 1)$ boundary components are supposed to match up with those of \bar{u}_S , in the sense that they can be joined via cylinders in $\pi^{-1}(T)$. This explains the name of the invariant, which is intended to be suggestive rather than literal.

Twisted monopole Floer homology. The Seiberg-Witten monopole Floer homology theory developed by Kronheimer and Mrowka (the most detailed exposition presently available is [22]) gives a functor on the following cobordism category \mathcal{C} . An object is a triple (Y, \mathfrak{t}, a) of closed, oriented 3-manifold, Spin^c -structure and cohomology class $a \in H^2(Y; \mathbb{R})$. A morphism from $(Y_0, a_0, \mathfrak{t}_0)$ to $(Y_1, a_1, \mathfrak{t}_1)$ is a diffeomorphism class of cobordisms between them; a cobordism here consists of a compact, oriented 4-manifold, X , Spin^c -structure \mathfrak{s} , and $b \in H^2(X; \mathbb{R})$, together with a diffeomorphism from $\partial X \rightarrow \bar{Y}_0 \cup Y_1$ which respects the Spin^c -structures and classes). The target category is that of modules over the Novikov field $\Lambda_{\mathbb{Z}/2}$.⁴

Definition 0.0.5. For any commutative ring R , the **Novikov ring** Λ_R is the ring of formal series $\sum_{\lambda \in \mathbb{R}} c_\lambda t^\lambda$, $c_\lambda \in R$, where for any $\lambda_0 \in \mathbb{R}$ there are only finitely many $\lambda < \lambda_0$ with $c_\lambda \neq 0$. (When R is a field, Λ_R is known to be a field as well.)

Thus one has modules $HM_\bullet(Y, a, \mathfrak{t})$ and homomorphisms $HM_\bullet(X, b, \mathfrak{s})$ between them.

³I am indebted to Paul Seidel for the idea that this surface could serve as a boundary condition for sections of X^N .

⁴The orientation issues controlling lifts to $\Lambda_{\mathbb{Z}}$ are not addressed in [22].

To be precise, HM_\bullet is Kronheimer-Mrowka's ' HM -to', in the twisted version where the Chern-Simons-Dirac functional is deformed so that the curvature term F_A is replaced by $F_A - ia_{\text{harm}}$. Some of the properties of this theory are as follows.

- $HM_\bullet(Y, a, \mathfrak{t})$ is 'graded': when $c_1(\mathfrak{t})$ is not torsion, there is a decomposition

$$HM_\bullet(Y, a, \mathfrak{t}) = \bigoplus_{j \in J(Y, \mathfrak{t})} HM_j(Y, a)$$

where $J(Y, \mathfrak{t})$ is the set of homotopy classes of oriented 2-plane fields underlying \mathfrak{t} . This is a transitive \mathbb{Z} -set with stabiliser $\langle c_1(\mathfrak{t}), H_2(Y; \mathbb{Z}) \rangle \subset 2\mathbb{Z}$. When $c_1(\mathfrak{t})$ is torsion, $HM_\bullet(Y, a, \mathfrak{t})$ is a completion of the above direct sum.

- The homomorphism $HM_\bullet(X, b, \mathfrak{s}) : HM_\bullet(Y_0, a_0, \mathfrak{t}_0) \rightarrow HM_\bullet(Y_1, a_1, \mathfrak{t}_1)$ is also graded: it maps $HM_{j_0}(Y_0, a_0, \mathfrak{t}_0)$ to $HM_{j_1}(Y_1, a_1, \mathfrak{t}_1)$, where j_1 is characterised by the existence of an almost complex structure J on X which underlies \mathfrak{s} , and which preserves—and is positive on—2-plane fields representing j_0 and j_1 .
- There is a canonical labelling of elements of the sets $J(Y, \mathfrak{t})$ as 'even' and 'odd' in such a way that the theory—modules and homomorphisms—becomes mod 2 graded.
- Each module $HM_\bullet(Y, a, \mathfrak{t})$ has a distinguished endomorphism U of degree -2 . These endomorphisms commute with the cobordism maps.

Symplectic Floer homology of symmetric products.

The cobordism category \mathcal{NS} of near-symplectic broken fibrations has as its objects quadruples (Y, π, σ, γ) , where Y is a closed, orientable 3-manifold, $\pi : Y \rightarrow S^1$ a smooth fibre bundle, $\sigma \in Z_Y^2$ a closed 2-form which is non-degenerate on the fibres, and $\gamma \in H_1(Y; \mathbb{Z})$. The morphisms from $(Y_0, \pi_0, \sigma_0, \gamma_0)$ to $(Y_1, \pi_1, \sigma_1, \gamma_1)$ are isomorphism classes of near-symplectic cobordisms. A cobordism consists of a near-symplectic fibration (X, π, ω) over a compact surface S , together with a class $\beta \in H'_2 \subset H_2(X, \partial X \cup Z_\omega; \mathbb{Z})$ (see below) and an isomorphism of these data, restricted to ∂S , with $(Y_0, \pi_0, \sigma_0, \gamma_0)^t \cup (Y_1, \pi_1, \sigma_1, \gamma_1)$. Here the superscript t means that we pull back the data by the map $x \mapsto -x$ on $S^1 = \mathbb{R}/\mathbb{Z}$, and we take the boundary $\partial\beta \in H_1(\partial X; \mathbb{Z})$. There is a technical condition on cobordisms which we impose in order to get a well-defined composition: the 'Hamiltonian curvature' of ω with respect to π (i.e. the curvature of the symplectic connection determined by ω) should be zero near ∂X .

A class $\beta \in H_2(X, \partial X \cup Z_\omega; \mathbb{Z})$ lies in H'_2 if $\delta : H_2(X, \partial X \cup Z_\omega; \mathbb{Z}) \rightarrow H_1(Z_\omega; \mathbb{Z})$ maps β to $[Z_\omega]$.

The subcategory \mathcal{NS}_d is obtained from \mathcal{NS} by imposing an extra condition on the homology classes γ and β : we require $\gamma \in H_1^{(d)}(Y; \mathbb{Z}) \subset H_1(Y; \mathbb{Z})$ and $\beta \in H_2^{(d)} \subset H'_2$.

These subsets consist of classes having intersection number $d + \frac{1}{2}\chi(\Sigma)$ with every regular fibre Σ , and non-negative intersection with every irreducible component of every fibre.

$\mathcal{NS}_d^{\text{conn}}$ is the subcategory of \mathcal{NS}_d where the fibrations have connected fibres.

For each $d \geq 1$, we shall construct a functor HF_\bullet from $\mathcal{NS}_d^{\text{conn}}$ to $\Lambda_{\mathbb{F}}$ -modules, for an arbitrary field \mathbb{F} . (We use fields rather than working over $\Lambda_{\mathbb{Z}}$ to avoid a minor complication involving possible Tor-terms in a Künneth formula.) As in the closed case, there is a partial extension to the case where there are disconnected fibres.

Reformulating the data. Before we state the main properties of HF_\bullet , we note that the data involved in \mathcal{NS}_d can be reformulated so as to become comparable to monopole Floer theory data. This goes as follows.

(a) Given (Y, π, σ) : there is an identification of $H^2(Y; \mathbb{Z})$ -torsors $H_1(Y; \mathbb{Z}) \cong \text{Spin}^c(Y)$. This sends γ to $\mathfrak{t}_\gamma := \text{PD}(\gamma) \cdot \mathfrak{t}_{\text{can}}$, where $\mathfrak{t}_{\text{can}}$ is the Spin^c -structure induced by the vertical tangent bundle $T^\vee Y$ (considered as oriented 2-plane field).

(b) Given (X, π_S, ω) : note that $H'_2 \subset H_2(X, \partial X \cup Z; \mathbb{Z})$ is an $H^2(X; \mathbb{Z})$ -torsor ($c \in H^2(X; \mathbb{Z})$ acts by addition of its Poincaré-Lefschetz dual) and is naturally identified with $\text{Spin}^c(X)$. The correspondence sends $\beta \in H'_2$ to $\mathfrak{s}_\beta := \text{PD}(\beta) \cdot \mathfrak{s}_{\text{can}}$. Here $\mathfrak{s}_{\text{can}} \in \text{Spin}^c(X \setminus Z)$ is induced by an almost complex structure on $X \setminus Z$ compatible with ω ; $\text{PD}(\beta) \cdot \mathfrak{s}_{\text{can}} \in \text{Spin}^c(X \setminus Z)$ extends uniquely to $\text{Spin}^c(X)$ (that it extends is a consequence of the fact that β maps to $[Z_\omega] \in H_1(Z_\omega X; \mathbb{Z})$ —see the discussion of ‘Taubes’ map’ in Chapter 1). Restriction of Spin^c -structures to ∂X corresponds to the boundary map $H'_2 \rightarrow H_1(\partial X; \mathbb{Z})$.

Under this correspondence, $H_2^{(d)}$ is mapped to $\text{Spin}^c(X)_d$.

(c) A loop $\tilde{\gamma}: S^1 \rightarrow Y$, positively transverse to the fibres of π and representing the class $\gamma \in H_1(Y; \mathbb{Z})$, determines a lift of \mathfrak{t}_γ to $J(Y, \mathfrak{t}_\gamma)$. This is represented by an oriented 2-plane field $\xi_{\tilde{\gamma}}$ obtained from $T^\vee Y$ by applying an automorphism of TY supported in a tubular neighbourhood $S^1 \times D^2$ of $\text{im}(\tilde{\gamma})$. The automorphism sends $(u, v) \in T_t S^1 \times T_x D^2 = T_t S^1 \times \mathbb{C}$ to $(u, e^{\pi i(|x|-1)} v)$.

Now we return to the properties of HF_\bullet .

- The module $HF_\bullet(Y, \pi, \sigma, \gamma)$ is *relatively* graded by the cyclic group $\mathbb{Z}/2(r+1-g)$, where $r = \gamma \cdot [\Sigma]$. This is isomorphic as \mathbb{Z} -set to $J(Y, \mathfrak{t}_\gamma)$. Indeed, $HF_\bullet(Y, \pi, \sigma, \gamma)$ is the homology of a complex on the free module $\Lambda_{\mathbb{Z}} \mathcal{H}(Y, \sigma, \gamma)$, where the generating set $\mathcal{H}(Y, \sigma, \gamma)$ has a \mathbb{Z} -action with stabiliser $2(r+1-g)$.
- $HF_\bullet(X, \pi_S, \omega, \beta)$ has a definite degree (called zero). Choose some $\nu_0 \in \mathcal{H}(Y_0, \sigma_0, \gamma_0)$; it defines a loop $\tilde{\gamma}$ and hence an element $j_0 := \xi_{\tilde{\gamma}}$ in $J(Y_0)$. Now $\nu_1 \in \mathcal{H}(Y_1, \sigma_1, \gamma_1)$

has the same degree as ν_1 if its induced class $j_1 \in J(Y_1)$ is related to j_0 via an almost complex structure representing \mathfrak{s}_β , as above.

- The subdivision of $J(Y, \mathfrak{t}_\gamma)$ into even and odd parts induces a splitting

$$HF_\bullet(Y, \pi, \sigma, \gamma) = HF_{\text{even}}(Y, \pi, \sigma, \mathfrak{t}_\gamma) \oplus HF_{\text{odd}}(Y, \pi, \sigma, \mathfrak{t}_\gamma)$$

which is respected by the cobordism maps.

- $HF_\bullet(Y, \pi, \sigma, \gamma)$ carries a canonical endomorphism U of degree -2 , and the cobordism maps commute with the U -maps. In fact, U is part of an action on $HF_\bullet(Y, \pi, \sigma, \gamma)$ by the integral quantum cohomology ring $QH^*(\text{Sym}^r(\Sigma))$, where $\Sigma = \pi^{-1}(0)$ and $r = \langle c_1(\mathfrak{t}), [\Sigma] \rangle$.

We define a functor $\mathcal{F} : \mathcal{NS}_d^{\text{conn}} \rightarrow \mathcal{C}$, by sending (Y, π, σ, γ) to $(Y, \mathfrak{t}_\gamma, \tau[\sigma] + 4\pi \text{PD}(\gamma))$, where $\tau = 4\pi \langle c_1(\mathfrak{t}), [\Sigma] \rangle / \int_\Sigma \sigma$ for a fibre Σ . The morphism $(X, \pi_S, \omega, \beta)$ is sent to $(X, \mathfrak{s}_\beta, \tau[\omega] + 4\pi \text{PD}(\beta))$, with τ defined similarly.

Conjecture 0.0.6. *There is an isomorphism of functors $HM_\bullet \circ \mathcal{F} \cong HF_\bullet$.*

The monopole Floer functor has two companions, \widehat{HM}_\bullet and \overline{HM}_\bullet , and the three fit into natural exact triangles. However, $\overline{HM}_\bullet(Y, a, \mathfrak{t})$, which is defined using the reducible solutions, is non-zero only when $2\pi c_1(\mathfrak{t}) = a$. This implies that $\overline{HM}_\bullet \circ \mathcal{F}$ is zero (the equation $\tau[\sigma] + 4\pi \text{PD}(\gamma) = 2\pi c_1(\mathfrak{t}_\gamma)$ has no solution with $\gamma \cdot [\Sigma] \geq \max(g-1, 1)$).

As we develop the theory based on relative symmetric products, we shall also keep track of a simpler theory using Picard (or Jacobian) varieties of the fibres. (From the gauge-theoretic perspective of Chapter 2, the degree r Picard variety is a ‘degeneration’ of the r th symmetric product: the symmetric product appears as a regular fibre of a moment map, the Jacobian as a critical fibre.) It seems likely that the Picard theory is equivalent to $\overline{HM}_\bullet(Y, 2\pi c_1(\mathfrak{t}), \mathfrak{t})$ —for instance, the homology modules are in both cases determined up to isomorphism by $H^*(Y; \mathbb{Z})$ and $c_1(\mathfrak{t})$ —but I have not yet tried to formulate a precise conjecture.

Further related theories. The Heegaard Floer homology HF^+ of Ozsváth and Szabó (which is the only one of the theories we could compare to HF_\bullet for which the foundations have appeared, in full, in print) is conjecturally isomorphic to the twisted monopole theory $HM_\bullet(Y, 2\pi c_1(\mathfrak{t}), \mathfrak{t})$ (see [27]). I do not see a direct connection between HF_\bullet and HF^+ . There is, however, a filtration on the differentials of HF_\bullet which seems formally similar to the one defining HF^+ , coming from intersections with the diagonal in the relative symmetric product, and this might repay closer attention.

One other theory worth mentioning is the periodic Floer homology PFH which is being built by Hutchings (see [19]). This is intended to be the home of relative Gromov-Taubes

invariants, ultimately including ones for near-symplectic forms. Again, it will conjecturally be isomorphic to a twisted monopole theory. Generalising Usher’s proof that the Donaldson-Smith invariant equals the Gromov invariant to an isomorphism between HF_\bullet and PFH appears to be a challenging task, but not hopelessly so.

With an array of potentially equivalent (or closely related) theories, it is highly desirable to have a method of comparing them based on their formal rather than analytic properties. I hope that near-symplectic fibrations, which are composed of simple pieces—elementary Lefschetz and near-symplectic fibrations over annuli and trivial fibrations over disks—may help in finding one. This remains to be seen.

Acknowledgements

I am enormously grateful to my supervisor, Simon Donaldson, for his encouragement, mathematical ideas, advice on doing mathematics, and considerable patience during the course of this project. It was a privilege to hear his ideas on broken fibrations at an early stage in their development, and I thank him for leading me into an area which I have found fascinating.

The work presented here builds on that of several mathematicians. Besides the impetus provided by the work of Auroux, Donaldson and Katzarkov [3] and Taubes [50], there is crucial technical input from papers by Donaldson and Smith [8, 49]; Salamon [38]; Seidel [46]; Seidel and Smith [47]. Paul Seidel has been a particular influence, both in print (his papers have been my guide to symplectic Floer homology) and in person: he has pointed me in the right direction on several occasions. It was Ivan Smith’s suggestion to look for Lagrangian correspondences. My thanks to them, and also to Denis Auroux, Kevin Costello, Joel Fine, Michael Thaddeus and Richard Thomas for their helpful comments.

Chapter 1

Near-symplectic geometry

Closed, transverse self-dual 2-forms on conformal 4-manifolds have been studied for some years. They play a central role in Taubes’ ongoing programme, outlined in [50], which aims to generalise the equality ‘SW = Gr’ to non-symplectic manifolds. In their preprint [3], Auroux, Donaldson and Katzarkov give an intrinsic formulation of the condition on a closed 2-form that it be transverse self-dual for some conformal structure; they call such a form near-symplectic. These forms are the subject of the first part of this chapter. This is partly a review of known facts (we give references as we proceed), but we also aim to dispel the fog surrounding certain claims for which proofs have not appeared in the literature.

The second part of the chapter is motivated by the main discovery of [3]—that certain singular surface-fibrations are topological counterparts to near-symplectic forms. Our account develops some of the basic properties of these fibrations.

1.1 Near-symplectic forms

Definition 1.1.1. A 2-form ω on an oriented 4-manifold X is called **near-positive** if at each point $x \in X$, either

1. $\omega_x^2 > 0$, or
2. $\omega_x = 0$, and $\text{im}(\nabla_x \omega) \subset \Lambda^2 T_x^* X$ is a maximal (that is, 3-dimensional) positive-definite subspace for the wedge-square quadratic form.

It is **near-symplectic** if, in addition, $d\omega = 0$.

To amplify: when $\omega_x = 0$, the gradient $\nabla_x \omega : T_x X \rightarrow \Lambda^2 T_x^* X$ is intrinsic—as at any zero of a smooth section of a vector bundle. If $\omega^2 \geq 0$ in a neighbourhood of x then $(\nabla_x \omega)(v)^2 \geq 0$

for all $v \in T_x X$. The wedge-product quadratic form $\Lambda^2 T_x^* X \otimes \Lambda^2 T_x^* X \rightarrow \Lambda^4 T_x^* X \rightarrow \mathbb{R}^4$ is defined up to a positive scalar, and has signature $(3, 3)$.

We write Z_ω for the zero-set $\{x \in X : \omega_x = 0\}$ of the form ω .

Example 1.1.2. On $\mathbb{R}^4 = \mathbb{R} \times \mathbb{R}^3$, with positively-oriented coordinates (t, x_1, x_2, x_3) , introduce the 2-forms

$$\begin{aligned}\beta_1 &= dt \wedge dx_1 + dx_2 \wedge dx_3, \\ \beta_2 &= dt \wedge dx_2 + dx_3 \wedge dx_1, \\ \beta_3 &= dt \wedge dx_3 + dx_1 \wedge dx_2.\end{aligned}$$

Now define $\Theta \in \Omega_{\mathbb{R}^4}^2$ by

$$\Theta = x_1 \beta_1 + x_2 \beta_2 - 2x_3 \beta_3,$$

so that $\Theta^2 = 2(x_1^2 + x_2^2 + 4x_3^2) dt \wedge dx_1 \wedge dx_2 \wedge dx_3$. The form Θ is near-symplectic with zero-set $Z_\Theta = \{x_1 = x_2 = x_3 = 0\}$. This example should be considered prototypical, for reasons that will emerge.

Lemma 1.1.3. *If $\omega \in \Omega_X^2$ is a near-positive form then its zero-set $Z_\omega \subset X$ is a smooth 1-dimensional submanifold.*

Proof. Take $z \in Z_\omega$. Working over a small ball $B \ni z$, choose a 3-plane subbundle $E \subset \Lambda^2 T^* B$ such that E_z is complementary to $\text{im}(\nabla_z \omega)$. Project $\omega|_B$ to a section $\bar{\omega}$ of $(\Lambda^2 T^* B)/E$. Then $\nabla_z \bar{\omega}$ is surjective, so, shrinking B if necessary, $\bar{\omega}$ vanishes transversely along a 1-submanifold $\bar{Z} \subset B$. Moreover, $T_x \bar{Z} = \ker(\nabla_x \omega)$ for $x \in \bar{Z}$, so ω is constant along \bar{Z} ; hence $\bar{Z} = Z_\omega \cap B$. □

Near-positive forms and almost complex structures

The condition for a 2-form to be near-positive prescribes how it meets the strata of $\Lambda^2 T^* X$. The group $\text{GL}(4, \mathbb{R})$ has precisely three orbits $\{O_i\}_{i=0,2,4}$ on the vector space $\Lambda^2(\mathbb{R}^4)^\vee$: O_i is the set of skew forms of rank i . Thus O_4 is a 6-dimensional open orbit diffeomorphic to $\text{GL}(4, \mathbb{R})/\text{Sp}(4, \mathbb{R})$. The stabiliser of the form $dx_1 \wedge dx_2$ comprises the matrices

$$\begin{pmatrix} A & B \\ 0 & D \end{pmatrix}, \quad A \in \text{Sp}(2, \mathbb{R})$$

so O_2 is 5-dimensional; O_0 is a point. Generic behaviour for a 2-form on a 4-manifold is, therefore, that it should be symplectic on the complement of a codimension-1 submanifold, and nowhere zero.

The conformal group $\mathbb{R}^+ \cdot \mathrm{SO}(4) \subset \mathrm{GL}(4, \mathbb{R})$ acts on the 3-dimensional subspace $\Lambda^+ \subset \Lambda^2(\mathbb{R}^4)^\vee$ of self-dual forms. It has just two orbits. In the presence of a conformal structure on a 4-manifold, a generic self-dual 2-form is non-degenerate on the complement of a codimension-3 submanifold on which it is zero.

Near-positive forms evidently behave quite unlike generic 2-forms. We will see that they are the same thing as transverse self-dual forms for arbitrary metrics.

Three preliminaries. (1) A choice of orientation and conformal structure on a four-manifold X is the same thing as a reduction of the structure group of TX from $\mathrm{GL}(4, \mathbb{R})$ to $\mathbb{R}^+ \mathrm{SO}(4)$. An $\mathbb{R}^+ \mathrm{SO}(4)$ -structure defines a Hodge-star operator and a splitting $\Lambda_X^2 = \Lambda_X^+ \oplus \Lambda_X^-$ of the bundle $\Lambda^2(X)$ into self-dual and anti-self-dual sub-bundles. Write Ω_X^\pm for the spaces of sections $\Gamma(X, \Lambda_X^\pm)$.

(2) We recall an algebraic relation between self-duality and complex structures in four dimensions. Let (V, ω) be a four-dimensional symplectic vector space. Suppose that ω is self-dual with respect to a conformal structure $[g]$; then the map $J \in \mathrm{End}(V)$ defined by

$$g(\cdot, \cdot) = \omega(\cdot, J\cdot) \tag{1.1}$$

satisfies $J^2 = -\lambda \mathrm{id}_V$, with $\lambda > 0$. Conversely, if J is a complex structure such that (1.1) defines an inner product g , then $*_g \omega = \omega$.

These assertions are clear when the triple (ω, J, g) is the standard one (ω_0, J_0, g_0) . Given g such that $*_g \omega = \omega$, choose an orthonormal basis for V . Then $\omega = (A^{-1})^* \omega_0$ for some $A \in \mathbb{R}^+ \mathrm{SO}(4)$, which implies that $J = A J_0 A^\top$ and $J^2 = -\det(A) \mathrm{id}_V$. For the converse, observe that $(V, J, g - i\omega)$ is a 2-dimensional hermitian vector space. By choosing a unitary basis we identify it with the standard one.

(3) On a 4-dimensional vector space V with given orientation $o_V \in \Lambda^4 V$, the map which sends a conformal structure $[g]$ to the subspace $\Lambda_{[g]}^+ \subset \Lambda^2 V$ establishes a bijection between conformal structures and 3-dimensional subspaces $\Lambda^+ \subset \Lambda^2 V$ which are positive-definite for the wedge-square quadratic form. For instance, one recovers Λ^- as the annihilator of Λ^+ . See Donaldson and Kronheimer [8, p.8].

Proposition 1.1.4. (a) Let X be a 4-manifold with chosen $\mathbb{R}^+ \mathrm{SO}(4)$ -structure. Let $\omega \in \Omega_X^+$ be a transverse section of Λ^+ . Then ω is near-positive.

(b) Let ω be a near-positive form on an oriented 4-manifold X . The following sets are in canonical bijection:

\mathcal{C} : conformal structures making ω self-dual;

\mathcal{S} : positive-definite 3-plane sub-bundles $\Lambda^+ \subset \Lambda_X^2$ having ω as section;

\mathcal{J} : almost complex structures J on $X \setminus Z_\omega$ such that $g_J(\cdot, \cdot) := \omega(\cdot, J\cdot)$ defines a metric whose conformal class extends to one on X .

(c) *These sets are non-empty, and contractible in the C^∞ topologies.*

Proof. (a) Self-duality implies that $\omega_x^2 = |\omega_x|^2 \text{vol}_x$, so that ω is symplectic where non-zero. As ω is a transverse section of a rank 3 vector bundle, its zero set is a codimension 3 submanifold of X ; hence it is near-symplectic.

(b) The map which sends a conformal structure $[g]$ to the bundle $\Lambda_{[g]}^+$ gives $\mathcal{C} \leftrightarrow \mathcal{S}$, by **3** above.

A metric g such that $*_g \omega = \omega$, gives rise (as in **2**) to an endomorphism of $T(X \setminus Z_\omega)$ which, after rescaling, is an almost complex structure in \mathcal{J} . Conversely, mapping $J \in \mathcal{J}$ to g_J , the self-duality equation $*_{g_J} \omega = \omega$ holds on $X \setminus Z_\omega$ and hence on X . This gives $\mathcal{C} \leftrightarrow \mathcal{J}$.

(c) To see that the sets are non-empty, take a tubular neighbourhood N of Z_ω , and extend the bundle $\text{im}(\nabla \omega) \rightarrow Z_\omega$ to a positive sub-bundle $\Lambda_0^+ \subset \Lambda_N^2$. We then obtain a canonical splitting $\Lambda^2 = \Lambda_0^+ \oplus \Lambda_0^-$ and a conformal structure $[g_0]$. Write $\omega = \omega_0^+ + \omega_0^-$, with $\omega_0^\pm \in \Gamma(N, \Lambda_0^\pm)$. Each map $m: \Lambda_0^+ \rightarrow \Lambda_0^-$, of pointwise norm < 1 , gives a new positive sub-bundle $\Lambda^+ = \text{graph}(m)$ and a new decomposition $\omega = \omega^+ + \omega^-$. With respect to $[g_0]$, we have $\omega^- = \eta + m^* \eta$, where $\eta = (1 - mm^*)^{-1}(\omega_0^- - m\omega_0^+)$. Thus to make ω self-dual, it suffices to find a map m , of norm < 1 , zero along Z_ω , such that $m\omega_0^+ = \omega_0^-$. This is clearly possible in a neighbourhood of Z_ω .

To find a global extension of our conformal structure we can switch to the almost complex structure viewpoint: by a well-known argument from symplectic geometry, for any compatible almost complex structure I over $U \subset X$, and any open set U' with $\overline{U'} \subset U$, there is a global J which extends $I|_{U'}$.

For contractibility of \mathcal{S} , take $[g] \in \mathcal{C}$. Then elements of \mathcal{S} correspond bijectively to graphs of maps $m: \Lambda_{[g]}^+ \rightarrow \Lambda_{[g]}^-$, of pointwise norm < 1 , such that $m(\omega) = 0$ and $m|_{Z_\omega} = 0$. Thus \mathcal{S} deformation-retracts onto $\{\Lambda_{[g]}^+\}$ by $\text{graph}(m) \mapsto \text{graph}(tm)$, $t \in [0, 1]$.

Since the bijections with \mathcal{C} (topologised as a quotient of Met_X) and with \mathcal{J} are homeomorphisms, those spaces are likewise contractible. □

Existence of near-symplectic forms; examples

Let X be a compact, connected, oriented 4-manifold. Introduce the space

$$P_X = \{(g, \omega) \in \text{Met}_X \times Z_X^2 : *_g \omega = \omega, \omega \neq 0\},$$

where the metrics and closed 2-forms are of class C^k , for some chosen integer $k \geq 2$, and we use the C^k topology. There is an open subspace P_X^* of pairs (g, ω) such that $\omega \in \Gamma(X, \Lambda_g^+)$ vanishes transversely. Write $q: P_X \rightarrow Z_X^2$ for the projection map. Notice that for $(g, \omega) \in P_X$, ω is g -harmonic, and so, by unique continuation, the set $\omega^{-1}(0)$ has empty interior. Its complement is then open and dense, and on it $\omega_x^2 > 0$.

The open subspace $q(P_X^*) \subset \text{im}(q)$ is the space of near-symplectic forms.

For proofs of the following result, first proved by Taubes, see LeBrun [26] or Honda [17].

Theorem 1.1.5. P_X^* is dense in P_X .

One derives this from transversality of a ‘universal’ section. The proof of transversality involves an asymptotic expansion for certain Green currents. There is a parametric version:

Theorem 1.1.6. Let $[0, 1] \ni t \mapsto \gamma(t) = (g_t, \omega_t)$ be a continuous path in P_X . Assume that $\gamma(j) \in P_X^*$ for $j \in \{0, 1\}$. Then γ is homotopic rel $\{0, 1\}$ to a C^k path $t \mapsto \gamma'(t) = (g'_t, \omega'_t)$ which vanishes transversely. Thus the set $Z = \{(t, x) \in [0, 1] \times X : \omega'_t(x) = 0\}$ is a submanifold of $[0, 1] \times X$ realising an embedded cobordism between the 1-manifolds $\{0\} \times Z_{\omega_0}$ and $\{1\} \times Z_{\omega_1}$.

The transversality argument (used to prove density of P_X^*) extends readily to paths $[0, 1] \rightarrow P_X$ (cf. Honda [17]); thus, using the Sard-Smale theorem, one can perturb γ to a path γ' with transverse zeros.

Corollary 1.1.7. (a) There exists a C^k near-symplectic form on X if and only if $b_2^+(X) > 0$. (b) $\text{im}(q) \subset Z_X^2$ has exactly two path-components if $b_2^+(X) = 1$ and one if $b_2^+(X) > 1$.

(c) Any two C^k near-symplectic forms ω_0, ω_1 in the same path-component can be joined by a C^k path $t \mapsto \omega_t$ in $\text{im}(q)$ which vanishes transversely, so that the set $Z = \{(t, z) : \omega_t(z) = 0\} \subset [0, 1] \times X$ is a cobordism from $\{0\} \times Z_{\omega_0}$ to $\{1\} \times Z_{\omega_1}$.

Proof. (a) Any near-symplectic form ω satisfies $[\omega]^2 = \int_X \omega^2 > 0$; conversely, if $b_2^+ > 0$ the density theorem tells us that C^k near-symplectic forms exist.

(b) This follows from the observation that path-components of $\text{im}(q)$ correspond exactly to those of $H := \{c \in H^2(X; \mathbb{R}) : c^2 > 0\}$. To see the correspondence, take $(\omega_0, g_0), (\omega_1, g_1) \in P_X$ and a path in H , say $t \mapsto c_t$, with $[\omega_0] = c_0$ and $[\omega_1] = c_1$. We then join g_0 to g_1 by a path of metrics $t \mapsto g_t$, and let ω_t be the self-dual projection of the g_t -harmonic representative of c_t , whereupon $t \mapsto (g_t, \omega_t)$ is a path in P_X .

(c) The theorem on paths in P_X tells us that we can make the path ω_t transverse when considered as a section Ω_0 of the appropriate 3-plane subbundle $\Lambda^+ \rightarrow [0, 1] \times X$ of $\text{pr}_2^* \Lambda_X^2$. \square

Remark 1.1.8. (1) In (a) one can replace C^k by C^∞ . One way to see this is to take a C^k near-symplectic form ω , and tubular neighbourhoods $N \subset\subset N'$ of Z_ω . There is a perturbation, within the space of closed forms, to a form which equals ω on N and which is C^∞ (but still symplectic) on $X \setminus N'$. One can then replace $\omega|_{N'}$ by a smooth form, e.g. using Honda’s lemma 1.1.18 which we discuss below.

(2) The transversality arguments can presumably be pushed further. It seems likely that one can arrange that the projection $\pi: Z \rightarrow [0, 1]$ is a Morse function. Then at points (t, z)

with $\omega_t(z) = 0$, $\nabla_z \omega_t$ would have rank ≥ 2 (the critical points of π would be the points $(t, z) \in Z$ where the rank equals 2). We do not pursue this direction here, but we will see some examples of paths of this type.

Gay and Kirby [14] have now proved the existence of near-symplectic forms on arbitrary closed 4-manifolds with $b_2^+ > 1$ without relying on the Hodge theorem. Their essential tool is contact geometry, specifically Eliashberg's theorem that overtwisted contact structures representing homotopic distributions are isotopic. The zero-circles of the forms which arise are all 'even' in the sense of Definition 1.1.17 below, and there is a formula for their number.

Example 1.1.9 (self-dual forms on Euclidean \mathbb{R}^4). Taking (t, x_1, x_2, x_3) as coordinates on \mathbb{R}^4 , we consider 2-forms $\theta \in \Omega_{\mathbb{R}^4}^2$, self-dual for the Euclidean metric g_0 , with $Z_\theta = \{x_1 = x_2 = x_3 = 0\}$. The forms $(\beta_1, \beta_2, \beta_3)$ from Example 1.1.2 give a basis for the self-dual forms over $C_{\mathbb{R}^4}^\infty$. Write $\theta = \sum \theta_i \beta_i$, where $\theta_i(t, 0, 0, 0) \equiv 0$. Then, for $z \in Z_\theta$,

$$d\theta(z) = \left(\sum_i \partial_i \theta_i \right) dx_1 \wedge dx_2 \wedge dx_3 + \sum_{i < j} (\partial_j \theta_i - \partial_i \theta_j) dt \wedge dx_i \wedge dx_j,$$

since $\partial_t \theta_i(z) = 0$. Consequently $d\theta(z) = 0$ if and only if the matrix

$$\left(\frac{\partial \theta_i}{\partial x_j}(z) \right)_{1 \leq i, j \leq 3}$$

is symmetric and trace-free. The condition for near-positivity (that is, for transversality of θ as section of Λ^+) is that the matrix be non-singular.

The form Θ of Example (1.1.2) indeed has matrix of partial derivatives which is symmetric, trace-free and non-singular.

Later we will reformulate this trivial example in more intrinsic terms.

Example 1.1.10. Two examples of near-symplectic forms with compact zero-set arise through symmetries of Θ . These symmetries are

$$\Theta(t, x) = \Theta(t - 1, x) = \Theta(t, \sigma x),$$

where $\sigma(x_1, x_2, x_3) = (x_1, -x_2, -x_3)$. By the first of these, Θ descends to a form Θ_{ev} on $S^1 \times \mathbb{R}^3$. By the second, it descends to a form Θ_{odd} on the mapping torus $T(\sigma) = (\mathbb{R} \times \mathbb{R}^3)/(t - 1, x) \sim (t, \sigma x)$. Its zero-set is the zero-section of the mapping torus.

Since $\sigma^2 = \text{id}$, there is a double covering $S^1 \times \mathbb{R}^3 = T(\text{id}) \rightarrow T(\sigma)$, $(t, x) \mapsto [2t, 2x]$. Under this map, Θ_{odd} pulls back to $8\Theta_{\text{ev}}$.

Example 1.1.11. There is an example of Luttinger and Simpson [28] of a family of closed forms $\{\omega_t\}_{t \in [-r, r]}$ on a ball in \mathbb{R}^4 , whose zero-set Z_{ω_t} is a circle for $t > 0$, a point p at $t = 0$, and empty when $t < 0$. They are near-symplectic when $t \neq 0$, and $\nabla_p \omega_0$ has rank 2.

Here is such a family (only marginally different from theirs). Using oriented coordinates (x_1, x_2, x_3, x_4) , and writing dx_{ij} for $dx_i \wedge dx_j$, define

$$\zeta = x_2(dx_{12} + dx_{34}) - x_4(dx_{14} + dx_{23}), \quad (1.2)$$

$$\eta_t = \frac{1}{2}(x_1^2 + x_3^2 - t)(dx_{13} + dx_{42}) + x_1x_2dx_{41} + x_3x_4dx_{32}. \quad (1.3)$$

Both these forms are closed, and $\zeta^2 \geq 0$. Fix a small constant ϵ , and consider the family $\omega(\epsilon, t) = \zeta + \epsilon\eta_t$ over a compact t -interval $[-r, r]$.

To be precise, consider a small neighbourhood U of the disk $D_r = \{(x_1, 0, x_3, 0) : x_1^2 + x_3^2 \leq r\}$. Let $Z = \{(t; x_1, 0, x_3, 0) \in [-r, r] \times U : x_1^2 + x_3^2 = t\}$. Thus $\omega_x(\epsilon, t) = 0$ when $(t, x) \in Z$. Moreover, $\omega_x(\epsilon, t)^2 > 0$ if $(t, x) \in ([-r, r] \times D_r) \setminus Z$, and there is a neighbourhood N of Z such that, for each $(t, x) \in N \setminus Z$, $\omega_x(\epsilon, t)^2 > 0$ provided $\epsilon = \epsilon(t, x)$ is small enough. By compactness of $[-r, r] \times D_r$, we can find a constant ϵ such that $\omega_x(\epsilon, t)^2 > 0$ for $(t, x) \in U \setminus Z$ (here we shrink U if required).

It is easy to check, by writing down partial derivatives along the zero-set, that $\nabla_x \omega(\epsilon, t)$ satisfies the rank conditions when $(t, x) \in Z$.

Example 1.1.12. One can vary the last example by redefining η_t :

$$\eta_t = \frac{1}{2}(x_1^2 - x_3^2 - t)(dx_{13} + dx_{42}) + x_1x_2dx_{41} - x_3x_4dx_{32}. \quad (1.4)$$

Again, $t \in [-r, r]$ and ϵ is a sufficiently small constant. Then the forms $\zeta + \epsilon\eta_t$ are near-symplectic for $t \neq 0$ on a suitable neighbourhood U of D_r , but the zero-set is now $\{(x_1, 0, x_3, 0) : x_1^2 - x_3^2 = t\}$. Thus as t passes through zero, the zero-set undergoes a surgery. At $t = 0$, $\nabla_p(\zeta + \epsilon\eta_0)$ has rank 2 at the isolated zero p .

Example 1.1.13. The following observation was pointed out to me by Donaldson. Note first that

$$\Lambda^+(\mathbb{R}^4) \otimes_{\mathbb{R}} \mathbb{C} = \Lambda^{2,0}(\mathbb{C}^2) \oplus \mathbb{C} \cdot \omega_{\mathbb{C}^2} \oplus \Lambda^{0,2}(\mathbb{C}^2), \quad \Lambda^-(\mathbb{R}^4) \otimes_{\mathbb{R}} \mathbb{C} = \Lambda_0^{1,1}(\mathbb{C}^2).$$

If f is a real function then the $(1,1)$ -form $dd^c f$ has self-dual component $\frac{1}{2}(\Delta f)\omega_{\mathbb{C}^2}$.¹ Here Δ is the Laplacian $4\sum_{i=1}^2 \frac{\partial^2}{\partial z_i \partial \bar{z}_i} = \sum_j \frac{\partial^2}{\partial x_j^2}$. Thus f is anti-self-dual iff f is harmonic (i.e. $\Delta f = 0$). We can regard anti-self-dual forms on \mathbb{C}^2 as self-dual forms on $\overline{\mathbb{C}^2}$, and so obtain a linear map

$$\{\text{harmonic functions on } U \subset \mathbb{R}^4\} \rightarrow \{\text{closed self-dual forms on } U \subset \overline{\mathbb{C}^2}\}.$$

This is potentially useful as a source for near-symplectic forms, including ones with many zero-circles; for example, it can be applied to harmonic polynomials.

¹Our normalisation for the d^c -operator $\Omega_X^0 \rightarrow \Omega_X^1$ on an almost complex manifold is that $d^c f = df \circ j = i(\partial - \bar{\partial})$.

Example 1.1.14. The following example is similar to, but different from, the previous one. Work on \mathbb{C}^2 , with coordinates $z_1 = x_1 + ix_2$, $z_2 = x_3 + it$. Recall that a smooth function $f: \mathbb{C}^2 \rightarrow \mathbb{R}$ is plurisubharmonic if $-dd^c f$ is non-negative, in that $-dd^c f(v, jv) \geq 0$, for all $v \in T\mathbb{C}^2$, where j denotes the complex structure. Suitable plurisubharmonic functions f give rise to near-symplectic forms $dd^c f$. Indeed, if there is a 1-submanifold $Z \subset \mathbb{C}^2$ such that $dd^c f(z) = 0$ for $z \in Z$, and such that all non-zero tangent vectors $v \in T(\mathbb{C}^2 \setminus Z)$ satisfy $-dd^c f(v, jv) > 0$, then $\omega = dd^c f$ is near-symplectic as soon as it satisfies the transversality condition. The latter is a condition on the third derivatives of f .

For example, the standard 2-form Θ can be expressed as

$$\Theta = dd^c (x_3(2x_3^2/3 - x_1^2/2 - x_2^2/2)). \quad (1.5)$$

Example 1.1.15. When (M, g) is a compact, oriented Riemannian three-manifold, and $\alpha \in \Omega_M^1$ is a harmonic 1-form, the 2-form

$$*\alpha + dt \wedge \alpha \in \Omega_{S^1 \times M}^2$$

is closed and self-dual with respect to the product metric $dt^2 \oplus g$ on $S^1 \times M$. Its zero-set is $S^1 \times \alpha^{-1}(0)$, and it will vanish transversely as soon as α does.

A theorem of Calabi says that, for a circle-valued Morse function $f: M \rightarrow S^1$, the 1-form df is harmonic for some metric if and only if the critical points of f have index 1 or 2, and any two regular points p_0, p_1 are connected by a smooth path p_t with $df(\dot{p}_t) > 0$. One can use such a function to find an S^1 -invariant near-symplectic form on $S^1 \times M$, either by invoking Calabi's result, or else by applying a Gompf argument (as in [3]) to the *broken fibration* $\text{id} \times f: S^1 \times M \rightarrow S^1 \times S^1$.

Geometry near the zero-circles

Taubes [51] has established the following picture of the differential geometry of a closed, transverse, g -self-dual form ω near its zero-set $Z = Z_\omega$. Since we are working with a fixed metric, we need not distinguish the intrinsic normal bundle $N_{Z/X}$ from the metric complement $(TZ)^\perp$.

(a) $\nabla\omega$ defines a vector bundle isomorphism $N_{Z/X} \rightarrow \Lambda^+|Z$. In particular, there is a universal procedure for orienting Z (as for the zero-manifold of any section of an oriented vector bundle). Following Taubes, we adopt the convention which gives Z the *co-orientation* o_N that makes $\nabla\omega$ orientation-reversing, and the *orientation* o_Z such that $o_Z \wedge o_N = o_X|Z$.

On \mathbb{R}^4 , with oriented coordinates (t, x_1, x_2, x_3) , $(\beta_1, \beta_2, \beta_3)$ is an oriented basis for Λ^+ . The standard form Θ has $\nabla\Theta = \text{diag}(1, 1, -2)$ with respect to the bases $(\partial_{x_1}, \partial_{x_2}, \partial_{x_3})$ and $(\beta_1, \beta_2, \beta_3)$. Thus the convention makes the vector field $-\partial_t$ along Z positive.

(b) Let z (resp. ζ) denote the unit-length oriented vector field (resp. 1-form) on Z . The interior product map

$$\Lambda^+|Z \rightarrow N_{Z/X}^\vee, \quad \eta \mapsto \iota(z)\eta$$

is a bundle isomorphism; its inverse is $\alpha \mapsto (\zeta \wedge \alpha) + *(\zeta \wedge \alpha)$.

(c) We have a sequence

$$N_{Z/X} \xrightarrow{\nabla\omega} \Lambda^+|Z \xrightarrow{\iota(z)} N_{Z/X}^\vee \xrightarrow{\text{metric}} N_{Z/X},$$

whose composite we denote by $S_{\omega,g} \in \text{End}(N_{Z/X})$. By working in Gaussian normal coordinates along Z (in which the formulae resemble those of (1.1.9)) one finds that $S_{\omega,g}$ is a *self-adjoint, trace-free automorphism*.

(d) It follows that, at each point of Z , $S_{\omega,g}$ has a basis of eigenvectors, and that two of its eigenvalues are positive and one negative. The positive and negative eigenspaces trace out eigenbundles L_g^+ and L_g^- (here L_g^- has rank 1 and is obviously locally trivial; L_g^+ is its orthogonal complement). Thus $N_{Z/X} = L_g^+ \oplus L_g^-$.

The auxiliary choice of metric is actually irrelevant. Whilst there are many conformal classes $[g]$ which make ω self-dual, they are all the same along Z , because for $z \in Z$, they satisfy $\text{im}(\nabla_z \omega) = \Lambda_{[g]}^+$. The sub-bundle $TZ^\perp \subset TX|Z$ is therefore intrinsic to ω .

We condense these points into a metric-free statement:

Proposition 1.1.16. *A near-symplectic form ω determines*

1. *a canonical orientation for the 1-manifold Z_ω ;*
2. *a canonical embedding of the intrinsic normal bundle $N_{Z_\omega/X}$ as a sub-bundle of $TX|Z$ complementary to TZ ;*
3. *a conformal class of quadratic forms S_ω on $N_{Z_\omega/X}$, of signature $(2,1)$;*
4. *a vector bundle splitting $N_{Z_\omega/X} = L^+ \oplus L^-$, such that S_ω is positive-definite on L^+ and negative-definite on L^- .*

A trivial 3-plane bundle over S^1 admits just two topological splittings.

Definition 1.1.17. The **parity** function of ω is the function

$$\varepsilon_\omega : \pi_0(Z_\omega) \rightarrow \mathbb{Z}/2, \quad Z_l \mapsto \langle w_1(L^-), [Z_l] \rangle.$$

This function is invariant in the sense that $\varepsilon_{\psi^*\omega}(\psi^{-1}(Z_l)) = \varepsilon_\omega(Z_l)$ under oriented diffeomorphisms ψ ; also, for an isotopy ω_t of near-symplectic forms with constant zero-set, one has $\varepsilon_{\omega_1} = \varepsilon_{\omega_0}$.

Our model forms $\Theta_{\text{ev}} \in \Omega_{S^1 \times \mathbb{R}^3}^2$ and $\Theta_{\text{odd}} \in \Omega_{T(\sigma)}^2$ each vanish along a circle, and the parities are those suggested by the notation.

Let $D_{\text{ev}}(r) \subset S^1 \times \mathbb{R}^3$ and $D_{\text{odd}}(r) \subset T(\sigma)$ be the open disk sub-bundles (over S^1) of radius r .

Lemma 1.1.18 (Honda). *Let $\omega \in \Omega_X^2$ be near-symplectic, Z a compact component of Z_ω , and U a neighbourhood of Z . There exists a smooth family $\{\omega_s\}_{s \in [0,1]}$ of near-symplectic forms, with fixed zero-set, such that $\omega_s(p) = \omega(p)$ if $s = 0$ or $p \in X \setminus U$, and where ω_1 is standard in the following sense: there is (according to the parity of Z) an oriented embedding $i: (D_{\text{ev/odd}}(r), S^1) \rightarrow (X, Z)$, for some r , such that $i^*\omega_1 = \Theta_{\text{ev/odd}}$.*

Proof. Honda's proof, which we follow in its essentials, amounts to a careful application of Moser's argument on $X \setminus Z$.

1. We may assume that X is a tubular neighbourhood $S^1 \times D^3$ of Z . Take a metric g for which ω is self-dual, and let t be a positively-oriented, constant-speed circle coordinate along Z , identifying it with the circle \mathbb{R}/\mathbb{Z} .

Suppose that Z has even parity. We can then find an orthonormal frame (e_1, e_2, e_3) for $N_{Z/X}$, such that $L_g^+ = \text{span}(e_1, e_2)$ and $L_g^- = \text{span}(e_3)$.

The metric and the choice of e_i give rise to normal coordinates (t, x_1, x_2, x_3) near Z , identifying a neighbourhood with $S^1 \times D^3(r)$. The three basic 2-forms β_i are defined as in Example 1.1.2, referring to these coordinates. In the basis (e_1, e_2, e_3) , the linear map $S_{\omega,g}(t)$ is represented by a trace-free symmetric matrix $S(t) = S^+(t) \oplus S^-(t)$, where $S^+(t)$ is 2×2 and positive-definite, and $S^-(t) < 0$. Regarding x and β as column vectors, we have an expansion near Z :

$$\begin{aligned} \omega(t, x) &= x \cdot S(t)\beta + O(|x|^2) \\ &= (x_1 \ x_2) S^+(t) \begin{pmatrix} \beta_1 \\ \beta_2 \end{pmatrix} + x_3 S^-(t)\beta_3 + O(|x|^2). \end{aligned}$$

The advantage of paying attention to the eigenbundle splitting is that one can then apply a convexity argument. Define $\theta_s = (1-s)\omega + s(x_1\beta_1 + x_2\beta_2 - 2x_3\beta_3)$. On a small neighbourhood of Z , this is a family of near-symplectic forms, with common zero-set Z .

If Z has odd parity then we can find orthonormal bases $(e_1(t), e_2(t), e_3(t))$ such that $L_g^+ = \text{span}(e_1, e_2)$ and $L_g^- = \text{span}(e_3)$, but with non-trivial monodromy around Z . Indeed, we can assume that

$$(e_1(1), e_2(1), e_3(1)) = (e_1(0), -e_2(0), -e_3(0)).$$

By exponentiating the e_i we get 'coordinates' $[t, x_1, x_2, x_3]$ on the mapping torus $T(\sigma)$, and we can proceed to write down a family θ_s in the same way as before.

2. On the total space of a vector bundle with a chosen metric one has the following explicit Poincaré lemma. Let $p: V \rightarrow M$ be the bundle projection, $i: M \rightarrow V$ the zero-section. Let $h_t: V \rightarrow V$ be the fibrewise dilation $x \mapsto tx$, and R the canonical vertical

vector field $\sum x_j \partial_j$ (here (x_j) is any basis for a fibre of V^*). The map

$$H: \Omega_V^* \rightarrow \Omega_V^{*-1}; \quad \alpha \mapsto \int_0^1 h_t^*(\iota(R)\alpha) t^{-1} dt$$

satisfies $\alpha - p^* i^* \alpha = d(H\alpha) + Hd\alpha$. Notice that if α vanishes along $i(M)$ then $H(\alpha)$ vanishes to second order along $i(M)$.

Applying this to a neighbourhood of the zero-section in $N_{Z/X}$, we find that $\eta_s := H(\theta_s)$ vanishes to second order along Z , and satisfies $d\eta_s = \theta_s$.

3. Let $U' \subset U$ be a neighbourhood of Z on which θ_s is defined and near-symplectic. Take a cutoff function $\chi = \chi(|x|)$, supported in U' and identically 1 near Z , and introduce the vector fields v_s on $X \setminus Z$ defined by $\iota(v_s)\theta_s + \chi\eta_s = 0$.

The family $\{v_s\}_{s \in [0,1]}$ generates a flow $\{\phi_s\}_{s \in [0,1]}$ on $X \setminus Z$, supported inside U . The reason is that $\nabla \dot{\theta}_s(u)$ is non-zero for all $0 \neq u \in N_{Z/X}$, while $\nabla \eta_s(u) = 0$, so that $|v_s(x)| \leq k|x|$ near Z , for some constant k . A trajectory x_s , defined on some interval $[0, s']$, satisfies $\frac{d}{ds}(\log |x_s|) \geq -k$; integrating over $[0, s']$, we obtain $|x_{s'}| \geq e^{-ks'}|x_0|$. This shows that the trajectory stays inside $X \setminus Z$, and therefore that ϕ_s is defined. We can extend ϕ_s to a self-homeomorphism of X by putting $\phi_s|_Z = \text{id}_Z$. The extension, though Lipschitz along Z , is perhaps not smooth.

Finally we can set $\omega_s = \phi_s^* \omega$ on $X \setminus Z$. Moser's argument shows that $\omega_s = \theta_s$ near Z and $\omega_s = \omega$ outside U . Thus ω_s extends smoothly over Z , giving an isotopy with the required properties. □

Remark 1.1.19. (a) As the proof makes clear, it is also true that any two near-symplectic forms, say $\omega \in \Omega_X^2$ and $\omega' \in \Omega_{X'}^2$, with zero-circles Z and Z' of the same parity, are diffeomorphic over deleted neighbourhoods $N \setminus Z$ and $N' \setminus Z'$, by a diffeomorphism that extends to a homeomorphism $N \cong N'$.

(b) Let ω and ω_0 be near-symplectic on X , with the same oriented zero-set Z ; take $p \in Z$. Then for any small neighbourhood $N_p \ni p$, there is a family $\{\omega_t\}_{t \in [0,1]}$ of near-symplectic forms with fixed zero-set, such that $\omega_t = \omega_0$ outside N_p and $\omega_1 = \omega$ near p . One can prove this by similar arguments to those in the proof of Honda's lemma. The first step is to twist ω_0 , via a family of diffeomorphisms supported near p , so as to match the splittings of $(N_{Z/X})_p$.

Theorem 1.1.20. Suppose that the connected, oriented four-manifold X carries a near-symplectic form whose zero-set has n components, $n \geq 2$. Then it carries a near-symplectic form, equal to the original one on the complement of a ball, whose zero-set has $n - 1$ components.

Corollary 1.1.21. Any closed, connected, oriented 4-manifold X , not negative-definite, carries a near-symplectic form whose zero-set has at most one path-component.

Proof. Let θ be near-symplectic, and Z_0, Z_1 distinct components of $Z = Z_\theta$. We begin by explaining how we intend to perform surgery on Z , but don't yet pay attention to θ . Take points $p_0 \in Z_0$ and $p_1 \in Z_1$, and a smooth path $\gamma: [0, 1] \rightarrow X$ from p_0 to p_1 , transverse to Z , with $\gamma^{-1}(Z) = \{0, 1\}$. For any $r > 0$, there is a neighbourhood U of $\text{im}(\gamma)$ and a diffeomorphism $\psi: U \rightarrow D(0; 2r) \subset \mathbb{R}^4$, such that

$$\begin{aligned}\psi(Z_i \cap U) &= \{(x_1, 0, x_3, 0) \in D(0; 2r) : x_1^2 - x_3^2 = r^2, (-1)^i x_1 > 0\}, \\ \psi(\text{im}(\gamma)) &= [-r, r] \times \{(0, 0, 0)\}.\end{aligned}$$

Introduce the family $C_t = \{(x_1, 0, x_3, 0) \in D(0; 2r) : x_1^2 - x_3^2 = r^2(1 - 2t)\}$ for $t \in [0, 1]$. There is an isotopy $t \mapsto \phi_t \in \text{Diff}(D(0; 2r))$ with $\phi_0 = \text{id}$, such that $\phi_t|_{D(0; r)} = \text{id}_{D(0; r)}$ and $\phi_t(C_0 \cap [D(0; 2r) - D(0; 3r/2)]) = C_t \cap [D(0; 2r) - D(0; 3r/2)]$. We replace X by $X' := (X \setminus U) \cup_{\partial D(0; 2r)} \overline{D(0; 2r)}$, where the attaching map is $\psi^{-1}\phi_1|_{\partial D(0; 2r)}$. Inside X' is the 1-manifold $Z' = (Z \cap (X \setminus U)) \cup C_1 \subset X'$. We have $X' \cong X$ and $|\pi_0(Z')| = |\pi_0(Z)| - 1$.

Realising such a surgery on zero-sets is not particularly difficult but a little laborious.

1. In Example 1.1.12, we found closed 2-forms $\zeta + \epsilon\eta_t$ on $D(0; 2r)$ vanishing precisely along C_t . For some $r > 0$, and some $\epsilon = \epsilon(r)$, the forms are near-symplectic when $|t| \leq r$. Fix such parameters (r, ϵ) , and write $\omega_s = \zeta + \epsilon\eta_{r(1-2s)} \in \Omega_{D(0; 2r)}^2$, where $s \in [0, 1]$.

After perturbing θ near $\{p_0, p_1\}$ as in Remark 1.1.19 (b), we may choose γ and ψ so that $\theta_0 := \psi^{-1*}\theta$ coincides with ω_0 on a neighbourhood of their common zero-set $C_0 \subset D(0; 2r)$.

2. Let $P = \psi(\text{im } \gamma) = [-r, r] \times \{(0, 0, 0)\} \subset D(0; 2r)$. Choose linearly independent vector fields $(v_1 = \partial_{x_1}, v_2, v_3, v_4)$ along P such that the matrix $(\theta_0(v_i, v_j))_{ij}$ takes the form

$$\lambda(t) \begin{pmatrix} J & 0 \\ 0 & J \end{pmatrix}, \quad J = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}, \quad \lambda(t) \geq 0.$$

Choose also linearly independent vector fields $(v'_1 = \partial_{x_1}, v'_2, v'_3, v'_4)$ along P such that the matrix $(\omega_0(v'_i, v'_j))$ takes the form

$$\lambda'(t) \begin{pmatrix} J & 0 \\ 0 & J \end{pmatrix}, \quad \lambda'(t) \geq 0.$$

We may assume that, near ∂P (where θ_0 and ω_0 coincide), these two frames are equal: $(v'_1, v'_2, v'_3, v'_4) = (v_1, v_2, v_3, v_4)$. Then there is an open neighbourhood $N \supset P$ and a diffeomorphism $\Phi: N \rightarrow N$, such that $\Phi(x) = x$ if $x \in P$ or x lies close enough to ∂P , and such that, along P , $D\Phi(v_i) = v'_i$. Define

$$\theta_s = (1 - s)\theta_0 + s\Phi^*\omega_0 \in \Omega_N^2, \quad s \in [0, 1].$$

There is a smaller open neighbourhood $N' \supset P$ on which these forms are near-symplectic.

3. The function $s \mapsto \theta_s \in \Omega_{N'}^2$, has constant derivative $d\theta_s/ds = \Phi^*\omega_0 - \theta_0$. We may write $d\theta_s/ds = d\alpha$, where $\alpha \in \Omega_{N'}^1$, vanishes near ∂P . Take an open set $N'' \supset P$

with $\overline{N''} \subset N'$, and a cutoff function $\chi: D(0; 2r) \rightarrow [0, 1]$ such that $\chi(N'') = \{1\}$ and $\chi(D(0; 2r) \setminus N') = \{0\}$. Define a family of vector fields v_s on $D(0; 2r)$, supported on N'' , by $\iota(v_s)\theta_s + \chi\alpha = 0$. These fields are zero near the (common) zero-set of the θ_s . The family $\{v_s\}$ generates a flow $s \mapsto F_s \in \text{Diff}(D(0; 2r))$, and the near-symplectic form $F_1^*\theta_0$ coincides with θ_0 outside N' and with $\theta_1 = \Phi^*\omega_0$ in N'' .

4. From the previous step we make the following conclusion. There is a near-symplectic form $\theta \in \Omega^2(X)$, isotopic to the one we started with; an open set $U'' \subset X$, diffeomorphic to a 4-ball, such that $U'' \cap \theta^{-1}(0)$ consists of two disjoint arcs; and a diffeomorphism $\psi': N'' \rightarrow U''$, such that $\psi'^*\theta = \omega_0$.

We are now ready to carry out the surgery. On N'' , we can perturb the family ω_t , isotoping the zero-sets, to a new family $\omega'_t \in \Omega_{N''}^2$ of near-symplectic forms, where $\omega'_0 = \omega_0$. These are to be such that $\omega'_t|_{N'''} = \omega_t|_{N'''}$, where N''' is a yet-smaller neighbourhood of P , and $\omega'_t|_V = \omega_0|_V$, where $V \subset N''$ is a small neighbourhood of $\omega_0^{-1}(0) \cap \partial N''$ (we may assume that $\partial N''$ is a 3-sphere, and that $\omega_0^{-1}(0) \cap \partial N''$ is a set of four points).

Then, by another application of Moser's argument, we can find an isotopy $[0, 1] \ni t \mapsto \phi_t \in \text{Diff}(\overline{D(0; 2r)})$, such that $\phi_0 = \text{id}$; $\phi_t|_V = \text{id}_V$; and ϕ_t preserves a shell $D(0; 2r) \setminus D(0; 2r - \delta)$ and on it satisfies $\phi_t^*\omega'_t = \omega_0$.

Finally we excise U'' from X , and glue it back in via the diffeomorphism ϕ_1 of its boundary. This carries a well-defined symplectic form, equal to $(\psi'^{-1})^*\omega_1$ on U'' , and realises the surgery $(X, Z) \rightarrow (X', Z')$ of zero-sets.

□

Remark 1.1.22. *The claim that one can reduce to one the number of zero-circles is not new, but a detailed proof has not appeared. It is mentioned in Luttinger-Simpson's preprint [28], where the reader is referred to work in preparation, and in [50], where Taubes briefly indicates Luttinger's argument. This seems to be somewhat different from ours in that it uses a model where one circle slides onto another and fuses with it in a process analogous (though not precisely in the sense of 1.1.15) to cancellations of Morse critical points in three dimensions.*

1.1.1 Topology associated with near-symplectic forms

We will need a number of results in 3- and 4-dimensional topology, mostly standard, which we collect here in readiness to applying them to near-symplectic forms and, later, broken fibrations.

2-plane fields on 3-manifolds

The set $J(Y)$ of homotopy classes of smooth, oriented 2-plane fields on an oriented 3-manifold Y is non-empty because the Euler class $e(TY)$ is zero. Once TY has been trivi-

alised, it is identified with the free homotopy set $[Y, S^2]$. Traditionally one uses obstruction theory to describe the structure of such a set, but an elegant modern approach, which does not require a trivialisation of TY , is to exploit the classification of Spin^c -structures. We draw on Kronheimer and Mrowka's exposition [21]. On a Riemannian 3-manifold, a Spin^c -structure $\tilde{\mathfrak{t}}$ is specified by a Hermitian \mathbb{C}^2 -bundle $\mathbb{S} \rightarrow Y$ and a map $\rho: T^*Y \rightarrow \text{End}(\mathbb{S})$ satisfying $\rho(\alpha)^2 = -|\alpha|^2 \text{id}$. In this section we will fussily write $\tilde{\mathfrak{t}}$ for a Spin^c -structure, and \mathfrak{t} for its isomorphism class. Now, the key observation from [21] is this:

- Oriented 2-plane fields are in bijection with isomorphism classes of pairs $(\tilde{\mathfrak{t}}, \Psi)$, where $\Psi \in \Gamma(Y, \mathbb{S})$ is a unit-length spinor.

The 2-plane field associated with $(\tilde{\mathfrak{t}}, \Psi)$ is $\ker(\alpha)$, the kernel of the unique 1-form α with $|\alpha|^2 = 1$ such that $\rho(\alpha)$ has $\mathbb{C}\Psi$ and Ψ^\perp as $\pm i$ -eigenspaces. Conversely, an oriented 2-plane field ξ inherits a hermitian structure from the metric on Y . There is an induced structure of hermitian complex vector bundle on the real vector bundle $\mathbb{S} := T^*Y \oplus \epsilon^1 = \xi^* \oplus (\xi^*)^\perp \oplus \epsilon^1$. This bundle has an obvious unit-length section Ψ . The Clifford map ρ is given, at each point $y \in Y$, by a map $L \oplus \mathbb{R} \rightarrow \text{U}(L \oplus \mathbb{C})$, where L is a 1-dimensional hermitian \mathbb{C} -vector space; as such, it can be described using the quaternions \mathbb{H} . Right multiplication by j makes \mathbb{H} a complex vector space, with complex subspaces $\langle i, k \rangle$ and $\langle 1, j \rangle$. Left multiplication gives a map from $\langle i, k \rangle \oplus \langle j \rangle$ to the unitary endomorphisms of $\mathbb{H} = \mathbb{C}^2$, and this is the model for the Clifford map. It is easy to check that the procedures described here are mutually inverse.

Using this bijection it becomes rather straightforward to describe $J(Y)$:

Coarse decomposition. This is by isomorphism classes of Spin^c -structures, and is slightly finer than that by Euler class:

$$J(Y) = \bigcup_{\mathfrak{t} \in \text{Spin}^c(Y)} J(Y, \mathfrak{t}).$$

An isomorphism class of Spin^c -structures on Y is the same thing as a ‘homology class’, of nowhere-vanishing vector fields, that is, a homotopy class on the complement of a 3-ball; $J(Y, \mathfrak{t})$ is the set of homotopy classes determining the same homology class \mathfrak{t} . For a pair of 2-plane fields (ξ_0, ξ_1) , representing $(\mathfrak{t}_0, \mathfrak{t}_1)$, the difference $c_1(\mathfrak{t}_0) - c_1(\mathfrak{t}_1) \in H^2(Y; \mathbb{Z})$ is the classical obstruction to homotopy of ξ_0 and ξ_1 over the 2-skeleton (which lies in $H^2(Y; \pi_2(S^2)) = H^2(Y; \mathbb{Z})$).

Fine decomposition. When Y is connected, each $J(Y, \mathfrak{t})$ is a transitive \mathbb{Z} -set. The stabiliser is the image of $H^1(Y) \rightarrow \mathbb{Z}$, $a \mapsto (a \cup c_1(\mathfrak{s}))[Y] \in 2\mathbb{Z}$.

The action $j \mapsto j[n]$ of $n \in \mathbb{Z}$ is induced by an automorphism of TY , supported over a small ball in Y : take a map $(B^3, \partial B^3) \rightarrow (\text{SO}(3), 1)$ of degree $2n$, and transfer it to TY via an embedding $B^3 \hookrightarrow Y$.

We may represent $j_0, j_1 \in J(Y, \mathfrak{t})$ by spinors Ψ_0, Ψ_1 , where $\Psi_1 = \theta\Psi_0$ for a function $\theta: Y \rightarrow S^1$. There is no obstruction to homotopy of Ψ_0 and Ψ_1 over the 2-skeleton of Y , since $\pi_2(S^3) = 0$. We define the difference class $\delta(j_0, j_1) \in \mathbb{Z}$ to be the relative Euler number

$$\delta(j_0, j_1) = e(\mathbb{S} \times [0, 1]; \Psi_0, \Psi_1)[X \times [0, 1], X \times \{0, 1\}].$$

This is the obstruction to homotopy of Ψ_0 and Ψ_1 over the 3-skeleton. Thus if the Euler number is zero then Ψ_0 is homotopic to Ψ_1 and $j_0 = j_1$.

A calculation in a model case gives that $\delta(j, j[n]) = n$, so the \mathbb{Z} -action is transitive.

We have $e(\mathbb{S} \times [0, 1]; \Psi_0, \theta\Psi_0) = \pm(c_1(\mathbb{S}) \cup [\theta])[Y]$, where $[\theta]$ is the class of θ in $H^1(Y; \mathbb{Z})$. This can be seen using the linear homotopy $(1-t)\Psi_0 + t\theta\Psi_0'$, where Ψ_0' is a unit-length spinor homotopic to Ψ_0 such that $\Psi_0 \wedge \Psi_0'$ vanishes transversely. Its zero set is concentrated at $t = 1/2$; there, it is the intersection of the zero-set of $\Psi_0 \wedge \Psi_0'$ (representing $c_1(\mathbb{S})$) and a level set of θ .

This shows that the stabiliser is as claimed.

In particular:

- $J(S^3) \cong \mathbb{Z}$ as \mathbb{Z} -sets.
- $J(S^1 \times S^2) = \bigcup_{n \in \mathbb{Z}} J(S^1 \times S^2, \mathfrak{t}_n)$, where \mathfrak{t}_n is the Spin^c -structure characterised by $\langle c_1(\mathfrak{t}_n), [\{t\} \times S^2] \rangle = 2n$. As \mathbb{Z} -sets,

$$J(S^1 \times S^2, \mathfrak{t}_n) \cong \mathbb{Z}/(2n). \quad (1.6)$$

Almost complex structures on punctured 4-manifolds

The basic reference for the following is the paper of Hirzebruch-Hopf [15]; useful modern sources are [21] and [14]. Let X be a smooth, closed, oriented 4-manifold.

(a) $w_2(TX) \in H^2(X; \mathbb{Z}/2)$ has a lift to $H^2(X; \mathbb{Z})$.

In general, a class $x \in H^2(X; \mathbb{Z}/2)$ has an integral lift iff it annihilates the mod 2 reduction of any torsion element in $H^2(X; \mathbb{Z})$ [15]. This is true for w_2 by the Wu formula $w_2 \cup y = y^2$. A lift c thus represents a characteristic element in the unimodular lattice $H^2(X; \mathbb{Z})/\text{tors.}$: $c \cup x \equiv x^2 \pmod{2}$ for all $x \in H^2(X; \mathbb{Z})$.

(b) Choose a set of disjointly embedded four-balls B_1, \dots, B_N in X . For each integral lift c of $w_2(TX)$ there is an almost complex structure J_0 over $X_0 := X \setminus \bigcup_{i=1}^N B_i$ with $c_1(X_0, J_0) = c|_{X_0}$. Indeed, the choice of c determines J_0 over the 2-skeleton, and it then extends unobstructedly over the 3-skeleton. Conversely, the Chern class of any such almost complex structure extends to a characteristic element in $H^2(X; \mathbb{Z})$, and J_0 extends to a stably almost complex structure over X , or equivalently to a Spin^c -structure. The obstruction

to extending J_0 to a true almost complex structure lies in $H^4(X; \pi_3(S^2)) = H^4(X; \mathbb{Z}) = \mathbb{Z}$. It can be computed as a sum of Hopf invariants $h_i(J_0)$: let ν_i be an outward-pointing normal vector field to ∂B_i such that $J_0\nu_i$ is tangent to ∂B_i . After trivialising TB_i , we can regard $J_0\nu_i$ as a map $\partial B_i \rightarrow S^2$, and define $h_i(J_0)$ to be its class in $\pi_3(S^2) = \mathbb{Z}$.

(c)

$$\frac{1}{4} (c^2 - 2e(X) - 3\sigma(X)) = - \sum_{i=1}^N h_i(J_0). \quad (1.7)$$

Proof of (1.7). The formula is true when J_0 extends to an almost complex structure on X . Hence, given $I \in J(S^3)$ (thought of as a complex structure on $TS^3 \oplus \varepsilon^1$), any extension of (S^3, I) to an almost complex 4-manifold (M, J_0) with boundary (S^3, I) will give the same number $c_1(M, J_0)^2 - 2e(M) - 3\sigma(M)$. In particular, we can choose M to have precisely N components. Since the left-hand side is additive under disjoint union, we deduce that it suffices to prove the formula when $N = 1$.

Identify $J(S^3)$ with \mathbb{Z} in such a way that $I \mapsto 0 \in \mathbb{Z}$ when I extends over B^4 . Then both sides of (1.7) define homomorphisms $J(S^3) = \mathbb{Z} \rightarrow \mathbb{Z}$ (to see that the left-hand side is additive, observe that one can replace a ball with Hopf invariant $h = h_1 + h_2$ by two disjoint balls with Hopf invariants h_1, h_2). The standard almost complex structure on \mathbb{R}^4 does not extend to S^4 ; its Hopf invariant at ∞ is $+1$. Since (1.7) holds for $1 \in J(S^3)$ it is generally true. □

As c is characteristic, $c^2 - \sigma(X) \equiv 0 \pmod{8}$, and hence

$$1 - b_1(X) + b_2^+(X) \equiv \sum_{i=1}^N h_i(J_0) \pmod{2}. \quad (1.8)$$

(d) Almost complex structures on a Riemannian 4-manifold have a spinor interpretation:

- There is a correspondence between orthogonal almost complex structures and isomorphism classes of pairs (\tilde{s}, Φ) of Spin^c -structure and unit-length positive spinor $\Phi \in \Gamma(\mathbb{S}^+)$.

The Spin^c structure \tilde{s} is specified by hermitian \mathbb{C}^2 -bundles $\mathbb{S}^+, \mathbb{S}^-$ with a map $\rho: T^*X \rightarrow \text{Hom}(\mathbb{S}^+, \mathbb{S}^-)$ such that $\rho^*(\alpha)\rho(\alpha) = |\alpha|^2 \text{id}_{\mathbb{S}^+}$.

Take a compact, oriented 4-manifold X_0 , with non-empty boundary in each component, and introduce the set $J(Y, \mathfrak{s})$ of homotopy classes of almost complex structures on X_0 underlying $\mathfrak{s} \in \text{Spin}^c(X_0)$. This set admits an action by $H_1(X_0, \partial X_0; \mathbb{Z})$. For $a \in J(Y, \mathfrak{s})$, and an oriented path γ with boundary in (and transverse to) ∂X_0 , define $a[\gamma]$ by choosing

a map $(B^3, \partial B^3) \rightarrow (\mathrm{SO}(3), 1)$ of degree 2, and using this to induce an automorphism α of TX_0 supported in a tubular neighbourhood $B^3 \times [0, 1]$ of $\mathrm{im}(\gamma)$. Choose a representative J for a , and let $a[\gamma]$ be the class of α^*J . Note that this does still represent \mathfrak{s} (the difference between the two Spin^c -structures a class in $H^2(X_0; \mathbb{Z})$; this is zero because every 2-cycle is homologous to one disjoint from γ).

Remark 1.1.23. A class $a \in J(X_0, \mathfrak{s})$ preserves a unique homotopy class of 2-plane fields over the boundary ∂X_0 . Thus there is a ‘restriction’ map $J(X_0, \mathfrak{s}) \rightarrow J(\partial X_0, \mathfrak{s}|_{\partial X_0})$. The group actions are compatible: $\partial(a[\gamma]) = (a|_{\partial X_0})[\partial\gamma]$, where $[\partial\gamma] \in H_0(\partial X_0; \mathbb{Z})$.

Lemma 1.1.24. The $H_1(X_0, \partial X_0; \mathbb{Z})$ -action on $J(Y, \mathfrak{s})$ is transitive. Its stabiliser is the image of the homomorphism $H_3(X_0, \mathbb{Z}) \rightarrow H_1(X_0; \mathbb{Z})$, $x \mapsto c_1(\mathfrak{s}) \cap x$.

Proof. We can represent orthogonal almost complex structures J_0, J_1 in the class \mathfrak{s} by spinors Ψ_0 and $\Psi_1 = \theta\Psi_0$, for a function $\theta: X_0 \rightarrow S^1$. Homotopy of Ψ_0 and Ψ_1 is detected by the relative Euler class of $\mathbb{S}^+ \times [0, 1] \rightarrow X_0 \times [0, 1]$:

$$e(\mathbb{S}^+ \times [0, 1]; \Psi_0, \Psi_1) \in H_1(X_0 \times [0, 1], \partial X_0 \times [0, 1]; \mathbb{Z}). \quad (1.9)$$

Again, this is because homotopy over the 2-skeleton is unobstructed, whilst the relative Euler class is the obstruction to homotopy over the 3-skeleton. We put $\delta(\Psi_0, \Psi_1) = i(e)$, where $i: H_1(X_0 \times [0, 1], \partial X_0 \times [0, 1]; \mathbb{Z}) \rightarrow H_1(X_0, \partial X_0; \mathbb{Z})$ is the obvious isomorphism. Thus $\delta(\Psi_0, \Psi_1)$ is represented by the projection to X_0 of the 1-cycle of zeroes of a section of $\mathbb{S}^+ \times [0, 1]$ extending the Ψ_i .

Now the lemma follows from two simple computations. First, if $a \in J(X, \mathfrak{s})$ is represented by $(\tilde{\mathfrak{s}}, \Psi_0)$ then $a[\gamma]$ is represented by a pair $(\tilde{\mathfrak{s}}, \Psi_1)$ with $\delta(\Psi_0, \Psi_1) = [\gamma]$. For, if $\alpha \in \mathrm{Aut}(TX)$ induces the action by γ , then (using a standard model for α) the zeroes of $(1-t)\Psi_0 + t\alpha^*\Psi_0$ occur only at $t = 1/2$ and along γ . Hence $J(X_0, \mathfrak{s})$ is a transitive $H_1(X_0, \partial X_0; \mathbb{Z})$ -set. Second, the function θ represents a class $[\theta] \in H^1(X_0, \partial X_0; \mathbb{Z})$, and

$$\delta(\Psi_0, \theta\Psi_0) = \pm c_1(\mathfrak{s}) \cap \mathrm{PD}[\theta] \in H_1(X_0; \mathbb{Z}) \quad (\text{universal sign}).$$

This can be seen using the linear homotopy $(1-t)\Psi_0 + t\theta\Psi'_0$, where Ψ' is a unit-length spinor homotopic to Ψ such that $\Psi_0 \wedge \Psi'_0$ vanishes transversely. Its zero set is concentrated at $t = 1/2$; there, it is the intersection of the zero-set of $\Psi_0 \wedge \Psi'_0$ (representing $c_1(\mathbb{S}^+)$) and a level set of θ . Thus the stabiliser is as claimed. \square

Example 1.1.25. Choose disjoint embeddings $B^4 \hookrightarrow S^4$ and $S^1 \times B^3 \hookrightarrow S^4$, and let W be the closure of the complement of their images. We have

$$J(W) = \bigcup_{n \in \mathbb{Z}} J(W, \mathfrak{s}_n),$$

where $\mathfrak{s}_n \in \text{Spin}^c(W)$ is the unique element which restricts to \mathfrak{t}_n on $S^1 \times S^2$; and $H_1(W, \partial W; \mathbb{Z}) = \mathbb{Z}$, where $1 \in \mathbb{Z}$ corresponds to a path γ running from S^3 to $S^1 \times S^2$. Each $J(W, \mathfrak{s}_n)$ is acted upon transitively by \mathbb{Z} , and because $H^1(W; \mathbb{Z}) = 0$ this action is free.

There are ‘restriction’ maps $r_1: J(W, \mathfrak{s}_n) \rightarrow J(S^3)$ and $r_2: J(W, \mathfrak{s}_n) \rightarrow J(S^1 \times S^2, \mathfrak{t}_n)$ as in Remark 1.1.23. Taking orientations into account, we see that

$$r_1(1 + a) = -1 + r_1(a), \quad r_2(1 + a) = 1 + r_2(a).$$

The sets $J(S^3)$ and $J(S^1 \times S^2, \mathfrak{t}_n)$ have distinguished basepoints, represented respectively by the boundary of the standard complex structure on B^4 , and by an S^1 -invariant complex structure. Thus to complete the description it is necessary only to give one pair $(r_1(a), r_2(a)) \in \mathbb{Z} \times \mathbb{Z}/(2n)$. The case relevant to near-symplectic geometry is $n = -1$, and in this case we will see that there is an a with $(r_1(a), r_2(a)) = (0, 0) \in \mathbb{Z} \times \mathbb{Z}/2$.

Elementary topology of near-symplectic 4-manifolds

One more topological lemma, due to Hirzebruch-Hopf: *the Euler class $e(\Lambda_X^+) \in H^3(X; \mathbb{Z})$ is zero.* For there is a standard homomorphism $\lambda^+: \text{SO}(4) \rightarrow \text{SO}(3)$ such that $\Lambda^+ = \lambda_*^+(TX)$, and a calculation using characters of representations gives that $w_2(\lambda_*^+P) = w_2(P)$ for any principal $\text{SO}(4)$ -bundle P . The Euler class e of $\text{SO}(3)$ -bundles equals $\beta \circ w_2$, where $\beta: H_2(-; \mathbb{Z}/2) \rightarrow H_3(-; \mathbb{Z})$ is the Bockstein operation. Since $w_2(TX)$ admits an integral lift, we have $e(\Lambda^+) = \beta w_2(TX) = 0$.

Proposition 1.1.26. *For any near-positive form $\omega \in \Omega_X^2$, there exists a compact orientable surface-with-boundary Σ and an embedding $i: \Sigma \rightarrow X$ with $i^{-1}(Z_\omega) = \partial\Sigma$.*

Proof. As the zero-set of a transverse section of Λ_g^+ (for suitable g), Z_ω represents the Poincaré dual to the Euler class $e(\Lambda^+)$, which is zero. Consequently there exists a class $a \in H_2(X, \mathbb{Z})$ with $\delta a = [Z_\omega] \in H_1(Z_\omega)$, and this is dual to a line bundle $L_a \rightarrow X \setminus Z_\omega$.

Take a closed tubular neighbourhood $N \cong \coprod_{l \in \pi_0(Z_\omega)} (S^1 \times \overline{D^3})$ of Z_ω , and let $S_l = \{\text{pt.}\} \times \partial D^3$ be a linking 2-sphere for the l th component. Then $\langle c_1(L_a), [S_l] \rangle = (\delta a)[Z_l] = 1$, and this characterises $L_a|_N$ up to isomorphism. There is a transverse section s' of $L_a|_N$ such that, in the l th component of N , s' has zero set $S^1 \times (0, 1] \times \{(0, 0)\} \subset S^1 \times \overline{D^3}$ (to find such a section one could start from a section of the degree 1 line bundle over S^2 , vanishing at one point).

Extend s' to a smooth section of L_a . A perturbation supported away from Z_ω then gives a transverse section s . We define $i: \Sigma \rightarrow X$ to be the inclusion $\overline{s^{-1}(0)} \hookrightarrow X$, which clearly has the required properties. □

A near-positive form ω determines a homotopy-class of almost complex structures on $X \setminus Z$, ($Z = Z_\omega$) and hence a canonical Spin^c -structure $\mathfrak{s}_\omega \in \text{Spin}^c(X \setminus Z)$ with spinor

bundle $\Lambda^{*,0}T(X \setminus Z)$. This Spin^c -structure sometimes extends over Z and sometimes not. In the near-symplectic case it does not extend. Indeed,

$$\langle c_1(\mathfrak{s}_\omega), [S_l] \rangle = -2$$

for each linking 2-sphere S_l of Z , where we use the pairing of homology and cohomology on $X \setminus Z$ (Taubes [51]). To orient S_l here, think of it as the boundary of a 3-ball in a fibre of $N_{Z/X}$, which is already oriented; orient $S^2 \subset \mathbb{R}^3$ as usual, so that $\partial_1 \wedge \partial_2$ is an orientation at $(0, 0, -1)$.

Taubes' map. There is a useful map

$$\tau: \text{Spin}^c(X) \rightarrow H_2(X, Z; \mathbb{Z}). \quad (1.10)$$

For $\mathfrak{s} \in \text{Spin}^c(X)$, let $L_{\mathfrak{s}}$ be the unique line bundle on $X \setminus Z$ with $\mathfrak{s}|(X \setminus Z) \cong \mathfrak{s}_\omega \otimes L_{\mathfrak{s}}$, and define $\tau(\mathfrak{s})$ as the class Poincaré-Lefschetz dual to $c_1(L_{\mathfrak{s}})$.

Note (i) that $H^2(X; \mathbb{Z}) \rightarrow H^2(X \setminus Z; \mathbb{Z})$ (or equivalently $H_2(X; \mathbb{Z}) \rightarrow H_2(X, Z; \mathbb{Z})$) is injective; and (ii) that the cokernel of $H_2(X; \mathbb{Z}) \rightarrow H_2(X, Z; \mathbb{Z})$ is isomorphic to $H_1(Z; \mathbb{Z})$ via the exact sequence of the pair. The relation

$$\tau(\mathfrak{s} \otimes L) = \tau(\mathfrak{s}) + \text{PD}(c_1(L)|(X \setminus Z)),$$

with (i), shows that τ is injective. The formula

$$0 = \langle c_1(\mathfrak{s}|(X \setminus Z)), [S_i] \rangle = -2 + 2\langle c_1(L_{\mathfrak{s}}), [S_i] \rangle,$$

with (ii), implies that $\text{im}(\tau)$ is the set of classes whose boundary in $H_1(Z; \mathbb{Z})$ is $[Z]$. In particular, a surface with boundary Z represents a class in $\text{im}(\tau)$.

As mentioned in the introduction, Taubes proves that when Z is the zero-set of a near-symplectic form and $\text{SW}_X(\mathfrak{s}) \neq 0$, there is a 'pseudoholomorphic subvariety' C which carries a fundamental class $[C] \in H_2(X, Z; \mathbb{Z})$, with $\delta[C] = [Z]$. Its relation to \mathfrak{s} is that $[C] = \tau(\mathfrak{s})$.

The number of even circles

The following result is stated without proof by Taubes, who attributes it to R. Gompf.

Proposition 1.1.27 (Gompf). *Let $\omega \in \Omega_X^2$ be a near-symplectic form on a closed 4-manifold X . For $i \in \mathbb{Z}/2$, let n_i denote the number of components of Z_ω for which $\varepsilon_\omega = i$. Then*

$$n_0 \equiv 1 - b_1(X) + b_2^+(X) \pmod{2}.$$

Hence n_0 is even when X admits an almost complex structure.

Proof. The proof is in five steps; the fourth relies on a computation which we carry out afterwards. Choose a Riemannian metric making ω self-dual.

1. Write $\beta(X)$ for the integer $1 - b_1(X) + b_2^+(X)$. As we have seen, when one has disjointly embedded closed 4-balls $B_1, \dots, B_N \subset X$, and a unit-length section of Λ^+ over $X \setminus \bigcup \text{int}(B^4)$ with Hopf invariants h_1, \dots, h_N , the relation

$$\beta(X) \equiv h_1 + \dots + h_N \pmod{2}$$

holds. (Previously we phrased this in terms of almost complex structures.)

2. Let X' denote the manifold obtained by surgery along an embedded, framed, oriented 1-submanifold γ . This means that we excise an open tubular neighbourhood of γ to obtain a manifold X_0 ; the framing gives (up to homotopy) an identification of ∂X_0 with a number of copies of $S^1 \times S^2$. To each boundary component we attach the complement of a standard copy of $B^3 \times S^1 \subset S^4$ (this is diffeomorphic to $S^2 \times B^2$). Then

$$\beta(X') = \beta(X) + 1.$$

Indeed, $\beta = (e + \sigma)/2$. As a cobordism invariant, the signature σ is unaffected by surgery, while a pair of Mayer-Vietoris sequences yields

$$e(X') - e(X) = e(S^2 \times B^2) - e(B^3 \times S^1) = 2 - 0.$$

3. We now carry out surgery along Z_ω . Along each component, there are two homotopically distinct framings, since $\pi_1(\text{SO}(3)) = \mathbb{Z}/2$, but either will do for our purposes. Note that $\beta(X') = \beta(X) + n_0 + n_1$.

4. We now have a manifold

$$X' = X_0 \cup \bigcup_{i=1}^N X'_i,$$

with standard pieces X'_i , and a non-vanishing 2-form $\omega|_{X_0}$ on X_0 . By restricting ω to $\partial X'_i$ we obtain a class in $J(S^1 \times S^2)$. Because of our choice of framing, this class is the odd or even element of $J(S^1 \times S^2, \mathfrak{t}_1) = \mathbb{Z}/2$ according to the parities of the original circles.

In Example 1.1.25, we claimed that the element $a \in J(S^1 \times S^2, \mathfrak{t}_{-1})$ extends over the standard manifold W to give an element $b \in J(S^3) = \mathbb{Z}$ if and only if $a = b \pmod{2}$. We will make good on this claim in a separate calculation that follows this proof.

5. The previous steps together give the result: the form $\omega|_{X_0}$ extends to a unit-length section of $\Lambda^+(X' \setminus \bigcup_{i=1}^{n_0+n_1} B_i)$ (where B_i is a small ball in X'_i). The sum of the Hopf invariants is $n_1 \pmod{2}$, and we have $\beta(X) + n_0 + n_1 \equiv n_1 \pmod{2}$.

□

Lemma 1.1.28. Θ_{ev} , restricted to $\{|x| = 1\}$ and considered as a map $S^1 \times S^2 \rightarrow \mathbb{R}^3 \setminus \{0\}$, extends to a map $\overline{B^2} \times S^2 \rightarrow \mathbb{R}^3 \setminus \{0\}$.

Proof. A tubular neighbourhood of a circle $S^1 \subset S^4$ gives a decomposition $S^4 = (S^1 \times \overline{B^3}) \cup_{S^1 \times S^2} (\overline{B^2} \times S^2)$. What follows is an explicit model for this decomposition.

Consider the function $f : \mathbb{C}^2 \rightarrow [0, \infty)$, $f(z_1, z_2) = (|z_1| - 2)^2 + |z_2|^2$. The pair $(f^{-1}[0, 1], f^{-1}(0))$ is diffeomorphic to $(S^1 \times \overline{B^3}, S^1 \times \{0\})$, via the map

$$\alpha : (z_1, z_2) \mapsto (|z_1|^{-1} z_1; |z_1| - 2, \operatorname{Re}(z_2), \operatorname{Im}(z_2)).$$

The region $f^{-1}[1, \infty)$ is diffeomorphic to $(\overline{B^2} \times S^2) \setminus \{(0, 0; 1, 0, 0)\}$ (to see this, observe that the map

$$f^{-1}[1, \infty) \rightarrow S^2; \quad (z_1, z_2) \mapsto f^{-1/2}(|z_1| - 2, \operatorname{Re}(z_2), \operatorname{Im}(z_2))$$

is a submersion, hence a disk-bundle). The pullback $\alpha^* \Theta_{\text{ev}}$ is given by $(z_1, z_2) \mapsto (-2(|z_1| - 2), z_2)$. Define

$$F : \mathbb{C}^2 \setminus f^{-1}(0) \rightarrow \mathbb{R}^3 \setminus \{0\}; \quad (z_1, z_2) \mapsto f(z_1, z_2)^{-1/2}(-2(|z_1| - 2), z_2).$$

Then $F = \alpha^* \Theta_{\text{ev}}$ on $f^{-1}(1)$. We consider the map $F/|F|$, restricted to a sphere $S_R^3 = \{|z_1|^2 + |z_2|^2 = R^2\}$, $R \gg 0$. This map is not surjective ($(1, 0, 0)$ is not in its image). But a map $S^3 \rightarrow S^2$ which is not surjective has trivial Hopf invariant, and hence extends over the four-ball.

□

This shows that (in the notation of Example 1.1.25), there exists $j \in J(W)$ with $r_1(j) = 0 \in \mathbb{Z}/2$, $r_2(j) = 0 \in \mathbb{Z}$, as required.

Remark 1.1.29. *It follows from the proof of the proposition above that the distributions ξ_{Θ_0} and $\xi_{\Theta'}$ (or ξ_{Θ_1}) on $S^1 \times S^2$ are not homotopic. This means that we have two ways to distinguish the two types of zero-circle: by orientability of the splitting of the normal bundle, or by the homotopy class of the foliation on the boundary of a tubular neighbourhood.*

Contact structures

Contact geometry plays an important role in Taubes' programme, and in the work of Gay and Kirby [14]. Here we briefly acknowledge its presence.

Consider first the standard near-symplectic Θ on $\mathbb{R} \times \mathbb{R}^3 = \mathbb{C}^2$ (cf. Example 1.1.2). We saw in Example 1.1.14 that $\Theta = dd^c f$, where $f = x_3(x_3^2/6 - r^2/2)$ (where $r^2 = x_1^2 + x_2^2$). Let $\alpha = d^c f$:

$$\alpha(t, x) = \frac{1}{2}(2x_3^2 - r^2)dt + x_3(x_2 dx_1 - x_1 dx_2) = \frac{1}{2}(2x_3^2 - r^2)dt + (x_3 r^2)d\theta.$$

This form, as observed by Honda, satisfies $\alpha \wedge d\alpha \wedge d(|x|^2) \neq 0$, which implies that its restriction to $\mathbb{R} \times S^2 = \{|x|^2 = 1\}$ satisfies the contact condition $\alpha \wedge d\alpha \neq 0$. We can rewrite

the contact form as

$$\alpha|_{\mathbb{R} \times S^2} = \left(\frac{3}{2}r^2 - 1\right) dt - r^2 \sqrt{1 - r^2} d\theta.$$

The contact distribution $\xi = \ker(\alpha)$ coincides with $T(S^1 \times S^2) \cap J_0 T(S^1 \times S^2)$, where J_0 is the complex structure determined by Θ and g_0 .

The vector field

$$\sigma(t, x) = (1 - q\rho^{-2})x_1\partial_1 + (1 - q\rho^{-2})x_2\partial_2 + (1 + 2q\rho^{-2})x_3\partial_3$$

satisfies $\iota(\sigma)\Theta = \alpha$, so that $\mathcal{L}_\sigma\Theta = \Theta$, i.e. σ is a Liouville vector field.

The contact structure ξ is *overtwisted*. In fact the equator $x_3 = 0$ in $\{0\} \times S^2$ bounds an overtwisted disk:

$$\bar{D}(1) \rightarrow \mathbb{R} \times S^2, \quad (re^{i\theta}) \mapsto (\epsilon(1 - r^2), r \cos \theta, r \sin \theta, \sqrt{1 - r^2}).$$

Here $\epsilon > 0$ is a (small) constant. Under this embedding, α pulls back to $-2\epsilon r(3r^2/2 - 1)dr - r^2\sqrt{1 - r^2}d\theta$, from which one sees that the characteristic foliation of \bar{D} is spanned by the vector field $-2\epsilon r(3r^2/2 - 1)\partial_\theta + r^2\sqrt{1 - r^2}\partial_r$. The foliation is singular just at the origin (where \bar{D} is tangent to ξ , and the foliation has a focus); and \bar{D} is a closed leaf (a limit cycle of the foliation). By definition, this is an overtwisted disk.

There are likewise forms α_{ev} , α_{odd} such that $\Theta_{\text{ev/odd}} = d\alpha_{\text{ev/odd}}$, arising by the same gluing operations as those for $\Theta_{\text{ev/odd}}$. These are contact structures on $S^1 \times S^2$, evidently overtwisted. By applying Honda's lemma, we get contact structures for general near-symplectic forms, and we have:

Proposition 1.1.30. *If $\omega \in Z_X^2$ is a near-symplectic form, N an open tubular neighbourhood of Z_ω , then $X \setminus N$ is a symplectic manifold with a concave contact-type boundary $(\partial N, \alpha)$. The contact structure is overtwisted and represents the homotopy class ξ_ω . The ends of $X \setminus Z_\omega$ can be identified with the concave end of the symplectisation of α .*

By a famous theorem of Eliashberg, an overtwisted contact structure on a 3-manifold Y is determined up to isotopy (hence up to isomorphism) by its class in $J(Y)$.

Remark 1.1.31. *The Spin^c -structure induced by the oriented 2-plane field ξ on ∂N is the same as the restriction from $N \setminus Z$ of the one induced by a compatible almost complex structure—indeed, an almost complex structure compatible with the symplectisation structure is also compatible with ω . From this we can give a proof (different from the one in [51]) that $\langle e(\xi_\omega), [S] \rangle = -2$, where S is any linking 2-sphere of Z . This goes as follows.*

By Honda's lemma, we may assume that ω is one of the standard forms Θ_{ev} or Θ_{odd} . Moreover, since the statement is local near a chosen point of Z , there is no loss in considering the local model given by the form Θ on \mathbb{R}^4 . The distribution ξ is homotopic to the

tangent distribution to the fibres of $(t, x) \mapsto (t, x_1^2 + x_2^2 - 2x_3^2)$. Its restriction to the sphere $S = \{t = 0, |x| = 1\}$ has as section the rotational vector field $x_1\partial_2 - x_2\partial_1$. This section has two zeros, $(0, 0, 0, \pm 1)$, both transverse, and positive with respect to the standard orientation of $S^2 \subset \mathbb{R}^3$ (after all $e(S^2) = +2$). But for us, the coordinates (x_1, x_2, x_3) are anti-oriented, so we must insert a minus sign.

1.2 Near-symplectic broken fibrations

The model map Q and 2-form Θ

Our aim is to describe the interplay between near-symplectic forms and certain singular fibrations. We do so first in a model case, involving the map

$$q: \mathbb{R}^3 \rightarrow \mathbb{R}; \quad (x_1, x_2, x_3) \mapsto \frac{1}{2}(x_1^2 + x_2^2) - x_3^2,$$

and the standard self-dual 2-form Θ on Euclidean \mathbb{R}^4 (see Example 1.1.2). Notice that

$$\Theta = dx_0 \wedge dq + *_4(dx_0 \wedge dq) = dx_0 \wedge dq + *_3dq.$$

Our model singularity will be the germ at the origin of the map

$$Q = q \times \text{id}_{\mathbb{R}}: \mathbb{R} \times \mathbb{R}^3 \rightarrow \mathbb{R} \times \mathbb{R}. \quad (1.11)$$

Here are some elementary properties of Q and of Θ .

- (a) The critical locus of Q is the set $Z = \{x_1 = x_2 = x_3 = 0\}$, which is just Z_{Θ} .
- (b) $Q^{-1}(s, t)$ is a hyperboloid of one sheet (diffeomorphic to $S^1 \times D^1$) when $s > 0$, a hyperboloid of two sheets (diffeomorphic to $D^2 \times S^0$) when $s < 0$, and a cone when $s = 0$.
- (c) $Q^{-1}([-\epsilon, \epsilon] \times \{t\})$ is a smooth 3-manifold with boundary, realising a (non-compact) cobordism from $Q^{-1}(-\epsilon, t)$ to $Q^{-1}(\epsilon, t)$. This is a standard model for surgery along a circle—or, going in the opposite direction, a pair of points—in an orientable surface.
- (d) Let (r, θ) be polar coordinates in the (x_1, x_2) -plane. Then, as Taubes observes, $(\theta, x_3r^2, x_0, q) \in S^1 \times \mathbb{R}^3$ give Darboux-type coordinates for the region $\{r \neq 0\}$ in \mathbb{R}^4 : we have

$$\Theta = dx_0 \wedge dq + d\theta \wedge d(x_3r^2).$$

These coordinates are orthogonal with respect to g_0 ; the coframe

$$\left(dx_0, \frac{dq}{\sqrt{r^2 + 4x_3^2}}, \frac{d\theta}{r}, \frac{d(x_3r^2)}{r\sqrt{r^2 + 4x_3^2}} \right)$$

is orthonormal. The dual orthonormal frame (v_0, v_1, v_2, v_3) is given by

$$\left(\partial_0, \frac{x_1\partial_1 + x_2\partial_2 - 2x_3\partial_3}{\sqrt{r^2 + 4x_3^2}}, \frac{x_1\partial_2 - x_2\partial_1}{r}, \frac{2x_3(x_1\partial_1 + x_2\partial_2) + r^2\partial_3}{r\sqrt{r^2 + 4x_3^2}} \right).$$

(e) If J_0 is the almost complex structure on $\mathbb{R}^4 \setminus Z$ determined by the equation $g_0(u, v) = \Theta(u, J_0 v)$ then, by the previous item, J_0 preserves the 2-plane distributions $\text{span}\{v_0, v_1\}$ and $\text{span}\{v_2, v_3\}$. Specifically, $J_0(v_0) = (r^2 + 4x_3^2)^{-1/2}v_1$, and $J(v_2) = (r^2 + 4x_3^2)^{-1/2}v_3$.

(f) At a point $x \in \mathbb{R}^4 \setminus Z$ the tangent space to the fibre of Q is

$$\ker(D_x Q) = \text{span}\{x_1 \partial_2 - x_2 \partial_1, 2x_3(x_1 \partial_1 + x_2 \partial_2) + r^2 \partial_3\}.$$

As we have just seen, this is a J_0 -complex line, so the fibres are J_0 -holomorphic curves. The function $Q : \mathbb{R}^4 \setminus Z \rightarrow \mathbb{C}$ is even holomorphic: $DQ \circ J = i \circ DQ$.

(g) The spaces $\mathcal{H}_{\theta_0, c} = \{(\theta, x_3 r^2) = (\theta_0, c)\}$ ($c \neq 0$) and $\mathcal{H}_{\theta_0, 0} = \{\theta = \theta_0, x_3 = 0\}$, together with $\{r = 0\}$, foliate $X \setminus Z$ (the foliation is singular along Z). The leaves are transverse to the fibres of Q ; indeed, the tangent space to the leaf through x is the annihilator with respect to Θ_x of $\ker(D_x Q)$. The leaves are also J_0 -curves.

There is a unique smooth section $s_{\theta_0, c}$ of $Q : \mathbb{R}^4 \rightarrow \mathbb{R}^2$ whose image is $\mathcal{H}_{\theta_0, c}$ ($c \neq 0$); it avoids Z . There is also a unique section of Q over $[0, \infty) \times \mathbb{R}$ whose image is $\mathcal{H}_{\theta_0, 0}$ (a surface with Z as its boundary).

It is a short step up from Θ to the two model forms Θ_{ev} and Θ_{odd} . In either case, Q is compatible with the relevant gluing map, and can therefore be promoted to maps $Q_{\text{ev}} : \mathbb{R}^3 \times S^1 \rightarrow \mathbb{R} \times S^1$ or $Q_{\text{odd}} : T(\sigma) \rightarrow \mathbb{R} \times S^1$. The discussion above then goes over essentially unchanged.

Indefinite quadratic singularities

Following Auroux, Donaldson and Katzarkov, we make the following definition.

Definition 1.2.1. Let X^4 and S^2 be oriented manifolds, $Z \subset X$ a 1-submanifold. A smooth map $\pi : X \rightarrow S$ has **indefinite quadratic singularities** along Z if, near any point of $z \in Z$, it is locally equivalent to the map $Q : \mathbb{R}^4 \rightarrow \mathbb{R}^2$.

Local equivalence means that there are charts $\psi : (\mathbb{R} \times \mathbb{R}^3, \mathbb{R} \times \{0\}, 0) \rightarrow (X, Z, z)$ and $\chi : (\mathbb{R} \times \mathbb{R}, \mathbb{R} \times \{0\}, 0) \rightarrow (S, \pi(Z), \pi(z))$, respecting the orientations of X and S , and a neighbourhood $U \ni 0$ in \mathbb{R}^4 , such that $(\chi^{-1} \circ \pi \circ \psi)|U = Q|U$.

Thus $\pi(Z)$ is a 1-manifold, and $\pi_Z : Z \rightarrow \pi(Z)$ a covering map.

Proposition 1.2.2. *A map π with indefinite quadratic singularities along Z determines*

1. *a canonical orientation for the 1-manifold Z ;*
2. *a canonical embedding of the intrinsic normal bundle $N_{Z/X}$ as a sub-bundle of $TX|_Z$ complementary to TZ ;*
3. *a conformal class of quadratic forms S_π on $N_{Z/X}$, of signature $(2, 1)$; in particular,*

4. an isotopy class of vector bundle splittings $N_{Z/X} = L^+ \oplus L^-$, such that S_π is positive-definite on L^+ and negative-definite on L^- .

Proof. The orientation is defined by taking $z \in Z$, and a short path $\gamma: (-\epsilon, \epsilon) \rightarrow S$ with $\gamma(0) = z$, transverse to $\pi(Z)$. For any small ball $U_z \ni z$, the sets $U_z \cap \pi^{-1}(\gamma(t))$ have one component when $t < 0$ and two when $t > 0$, or vice versa. Thus the segment $\pi(Z \cap U_z) \subset S$ has a ‘connected side’ and a ‘disconnected side’. We orient Z by declaring that $v \in T_z Z$ is positive when $j(\pi_* v)$ points into the disconnected side of $\pi(Z \cap U_z)$ (here $j \in \text{End}(TS)$ is a complex structure compatible with the orientation of S).

The normal sub-bundle is the 3-plane bundle $\ker(D\pi) \rightarrow Z$ spanned by the tangent cones of the fibres of π . This sub-bundle carries an intrinsic Hessian quadratic form $D^2\pi: \ker(D\pi) \otimes \ker(D\pi) \rightarrow \pi^* N_{\pi(Z)/S}$ (here $N_{\pi(Z)/S}$ is the intrinsic normal bundle). Choosing a trivialisation of $N_{\pi(Z)/S}$ of the correct sign one obtains an \mathbb{R} -valued form of signature $(2, 1)$. □

We can now formulate a notion of compatibility between near-symplectic forms and maps with indefinite quadratic singularities.²

Definition 1.2.3. Suppose that $\pi: X \rightarrow S$ has indefinite quadratic singularities along Z , and that $\omega \in \Omega_X^2$ is a near-symplectic form with zero-set $Z_\omega = Z$. We say that π is **compatible with ω along Z** if (a) the orientations of Z induced by ω and by π agree; (b) the normal sub-bundle $N_{Z_\omega/X} \subset TX|_Z$ determined by ω coincides with $\ker(D_z \pi)$; and (c) the conformal quadratic forms S_ω and $D^2\pi$ on this common normal bundle are equal.

Any two compatible forms are linearly homotopic near Z . One can turn this round to see that compatible near-symplectic forms always exist on some neighbourhood U_Z of Z , because one can patch together forms defined on balls. Conversely, given ω one can find (using S_ω) a compatible map π on a neighbourhood U_Z .

The following lemma shows that one can find almost complex structures compatible with all the available structure simultaneously.

Lemma 1.2.4. *Suppose that π is compatible with ω along Z . Then there is a neighbourhood U of Z and a metric h on U , smooth on $U \setminus Z$ and continuous everywhere, such that $*_h \omega = \omega$ and, for all $x \in U \setminus Z$, the 2-plane $V_x = \ker(D_x \pi)$ is orthogonal to $H_x := \text{Ann}_\omega(V_x)$. The associated almost complex structure J_h preserves both these distributions, and in particular, ω is positive on the oriented 2-planes V_x .*

Proof. Here we choose to work with coordinates which make π look standard. It is not difficult to see that there is a tubular neighbourhood $S^1 \times D^3$, with coordinates (t, x_1, x_2, x_3) ,

²As a (stronger) alternative, one could require that the pair (π, ω) should be locally equivalent to (Q, Θ) .

such that

$$q = \sum_{1 \leq i, j \leq 3} q_{ij}(t) x_i x_j$$

where q is π followed by projection on the normal coordinate along $\pi(Z)$. The matrix $q_{ij}(t)$ is symmetric and trace-free, with signature $(2, 1)$. Then, because of the compatibility between ω and π , we have

$$\omega = \sum_{1 \leq i, j \leq 3} q_{ij}(t) x_i \beta_j + O(|x|^2)$$

(here β_j is as in Example 1.1.2). Write ω_0 for the leading order part of ω . There is a metric g_0 on U_z such that $*_{g_0} \omega_0 = \omega_0$ and for which J_{g_0} preserves the vertical distribution V on $U_z \setminus Z$.

The forms ω , ω_0 give different horizontal distributions on $U_z \setminus Z$, $H = \text{Ann}_\omega(V)$, and $H_0 = \text{Ann}_{\omega_0}(V)$. Moreover, H is the graph of a bundle map $\rho \in \text{Hom}(H_0, V)$. Define a metric g by declaring $g|V = g_0|V$, $g(V, H) = 0$, and, for $h, h' \in H_0$, $g(h + \rho h, h' + \rho h') = g_0(h, h')$; put $g_z = (g_0)_z$ for $z \in Z$. This defines a continuous metric because $\rho_x \rightarrow 0$ as $x \rightarrow Z$.

The metric g meets the requirement $V \perp H$, but ω need not be g -self-dual. However, by rescaling g on the planes H (and leaving it alone on V) we obtain a new metric h which meets both these conditions. The conformal scaling factor tends to 1 at Z because ω_0 is g_0 -self-dual. Thus h is still continuous.

The planes V_x are ω_x -positive when x is close to Z because this is so for ω_0 . The statement about J_h now follows formally.

□

Definition 1.2.5. A **broken fibration** is a smooth, proper map $\pi: X \rightarrow S$ from an oriented 4-manifold to an oriented 2-manifold, such that

- The critical set X^{crit} is the union of submanifolds Z and C of dimensions 1 and 0; moreover, $\pi(Z)$ and $\pi(C)$ are disjoint submanifolds of S ;
- π has indefinite quadratic singularities along Z ;
- π maps each component of Z diffeomorphically onto its image;
- the map $\pi: X \setminus Z \rightarrow S \setminus \pi(Z)$ is a Lefschetz fibration. That is, its germ at any point in C is equivalent to the germ at $0 \in \mathbb{C}^2$ of $(z_1, z_2) \mapsto z_1 z_2$.

A **near-symplectic broken fibration** (X, π, ω) is a broken fibration together with a near-symplectic form $\omega \in \Omega_X^2$ which is compatible with π along Z , and which restricts positively to the planes $\ker(D_x \pi)$ at all regular points x .

Theorem 1.2.6 ([3]). *When X is closed, a singular Lefschetz fibration (X, π) lifts to a near-symplectic broken fibration (X, π, ω) if and only if there is a class $c \in H^2(X; \mathbb{R})$ which satisfies $\langle c, u \rangle > 0$, where $u \in H_2(X; \mathbb{R})$ ranges over classes represented by components of regular fibres of π .*

The much deeper result of [3] concerns the existence of broken fibrations (or rather, pencils) compatible with a given near-symplectic form. Now that we have given the relevant definitions, we refer back to the Introduction for the statement.

Fibred vanishing cycles

From a near-symplectic broken fibration (X, π, ω) one obtains a connection on the non-singular locus X^* —i.e., a field of subspaces $T_x^h(X^*) \subset T_x(X^*)$ complementary to $\ker(d\pi)_x$ —by letting $T_x^h X$ be the annihilator with respect to ω_x of $T_x^v X = \ker(d\pi)_x$.

We will need some notation for the parallel transport maps associated with this connection. Write

$$m_\gamma: X_{\gamma(a)} \rightarrow X_{\gamma(b)}$$

for the symplectomorphism induced by a path $\gamma: [a, b] \rightarrow S^*$. Also, write

$$m_\Gamma: X_{\Gamma(S^1 \times \{a\})} \rightarrow X_{\Gamma(S^1 \times \{b\})}$$

for the map of fibre bundles induced by a family of paths parametrised by the circle, $\Gamma: S^1 \times [a, b] \rightarrow S^*$.

We will discuss the local geometry of such connections in Chapter 2. Notice that, in the case of (Q, Θ) , the connection is *flat*: in point (\mathbf{g}) at the beginning of Section 1.2, we exhibited a horizontal foliation.

Theorem 1.2.7. *Let (X, π, ω) be a near-symplectic broken fibration, and Z_0 a component of Z . Then there exists an embedded, connected, Lagrangian surface $V \subset X \setminus \text{crit}(\pi)$, such that $\pi(V)$ is a smooth circle in S , parallel to $\pi(Z_0)$ and on its connected side, with $\pi|_V: V \rightarrow \pi(V)$ a fibre bundle. Moreover, to the data (X, π, ω, Z_0) one can associate a distinguished deformation-class of such objects.*

Such a surface, as an S^1 -bundle over S^1 , is diffeomorphic to the torus or Klein bottle. A representative of the special deformation-class is called a **fibred vanishing cycle**. In this case, we get a torus when the parity of Z_0 is even, a Klein bottle when it is odd.

Proof. This follows Seidel's treatment of vanishing cycles of Lefschetz critical points in [46].

We assume for the time being that every critical point lying over $\pi(Z_0)$ is in fact in Z_0 . Then we can choose an embedding $\alpha: (S^1 \times [-1, 1], S^1 \times \{0\}) \rightarrow (S, \pi(Z_i))$, such that

$\alpha(S^1 \times [-1, 0))$ is the connected side of $\pi(Z_0)$ and π has no critical points in $\pi^{-1}(\text{im } \alpha)$. Replace (X, π, ω) by their pullbacks by α , so that the base is $S^1 \times [-1, 1]$.

For $\epsilon \in [0, 1)$ and $t \in S^1$, let $\gamma(t, \epsilon)$ denote the path $[-1, -\epsilon] \rightarrow S^1 \times [-1, 1]$, $s \mapsto (t, s)$. Define

$$W_\alpha = \{x \in X_{(t,s)} : s \in [-1, 0), \lim_{\epsilon \rightarrow 0^+} m_{\gamma(t,\epsilon)}(x) \in Z_0\} \cup Z_0.$$

(This is to be read as ‘the limit exists as ϵ decreases to 0, and is a point of Z_0 ’.) Our candidate for fibred vanishing cycle is $W_\alpha \cap \pi^{-1}(S^1 \times \{-1\})$.

To understand the symplectic flow defining W_α , we show that its generating vector field, which is the horizontal lift $\tilde{\partial}_s$, can be expressed in the form λu for a non-vanishing vector field u which extends over Z to a Morse-Bott vector field with an index 1 zero along Z , and a function $\lambda > 0$.

Let $f = -\text{pr}_1 \circ \pi : X \rightarrow S^1$. The equation $df = \iota(v_f)\omega$ defines a symplectic vector field v_f on the region $X \setminus Z_0$ where ω is symplectic (v_f will become singular along Z_0). Since f factors through π , this is a horizontal field: $v_f(x) \in T_x^h(X)$. Since f is constant along trajectories of v_f , $D\pi(v_f(x))$ is always tangent to the ray $\{-f(x)\} \times [-1, 1]$. This shows that $v_f = \lambda \tilde{\partial}_s$, where the function λ is non-vanishing since f has no critical points on $X \setminus Z_0$. Actually $\lambda > 0$.

Consider yet another vector field, $\nabla_g q$. Here $q = \text{pr}_2 \circ \pi : X \rightarrow \mathbb{R}$, and we will decide on the metric g as we proceed. The advantage of $\nabla_g q$ is that it evidently has Z as index 1, non-degenerate critical manifold. In the model (Θ, Q) —though not in general—we can arrange (by using the flat metric g_0) that it be proportional to v_f .

Now, locally we can write $q = (x_1^2 + x_2^2 - 2x_3^2)/2$ and $\omega = \Theta + O(x_1^2 + x_2^2 + 4x_3^2)$. If we choose the euclidean metric then we have explicit formulae for orthonormal frames. Using these we see that, whilst a unit-length Θ -horizontal vector field e will not be ω -horizontal, its vertical component $e_v(x)$ goes to zero in C^1 as x approaches a point of Z_0 .

Write $\nabla q = (\nabla q)_1 + (\nabla q)_2$ where $(\nabla q)_1(x) \in \mathbb{R}\tilde{\partial}_s(x)$ and $(\nabla q)_2(x) \in \mathbb{R}\tilde{\partial}_t(x) \oplus T_x^\vee X$. We see from the previous paragraph that, as x approaches a point of Z_0 , we have that $|(\nabla q)_2(x)|/|(\nabla q)_1(x)| \rightarrow 0$, and moreover, the Hessian of $(\nabla q)_1$ remains non-degenerate on the normal bundle to Z_0 . Thus $(\nabla q)_1$ is a positive multiple of v_f over a deleted neighbourhood $N(Z_0) \setminus Z_0$, and is still Morse-Bott.

We now define u to be any vector field which is a positive multiple of v_f over $X \setminus Z_0$, and which is equal to $(\nabla q)_1$ near Z_0 . Observe that W_α is the stable manifold of Z_0 under u . As such, it is the total space of a 2-disk bundle $l : W_\alpha \rightarrow Z$, and its structure group is reduced canonically to $O(2)$. Its boundary $V_\alpha := \partial W_\alpha$ is a circle-bundle over Z .

It is clear that V_α is a surface, and a circle-bundle over the circle $\{-1\} \times S^1$. To see that it is Lagrangian, consider the maps $m_\epsilon : X_{S^1 \times \{-1\}} \rightarrow X_{S^1 \times \{-\epsilon\}}$ given by parallel transport

along the rays $\{t\} \times [-1, -\epsilon]$. These maps are symplectic in the sense that

$$m_\epsilon^*(\omega|(S^1 \times \{-\epsilon\})) = \omega|(S^1 \times \{-1\})$$

(see Chapter 2). The maps $m_\epsilon|_{V_\alpha}$ have a (pointwise) limit $m_0: V_\alpha \rightarrow Z_0$ as $\epsilon \rightarrow 0^+$. For each $x \in V_\alpha$, $|\omega(m_\epsilon)(x)| \rightarrow 0$, while $|D_x m_\epsilon|$ remains bounded. It follows that $\omega|_{V_\alpha} = m_0^*(\omega|_Z) = 0$.

Thus, to each embedding α is associated the fibred vanishing cycle V_α . Any other choice of embedding, α' say, can be deformed to α , and thus $V_{\alpha'}$ can be deformed to V_α through fibred vanishing cycles.

It remains to remove the assumption on Z_0 . To do so, notice that in the general case the set W_α is still defined. There may no longer be parallel transport maps defined globally on the fibres of π , but the argument can be carried out in a small neighbourhood of W_α which contains no other critical points. The only thing which becomes more complicated is the notation. \square

Surgery on symplectic surface-automorphisms

When $\phi \in \text{Aut}(\Sigma, \omega)$ leaves invariant a circle $L \subset \Sigma$, and preserves the foliation by parallel circles of a neighbourhood of L , one can perform surgery along L in a way which is compatible with ϕ . The result is a self-diffeomorphism $\sigma_L \phi$ of the new surface $\sigma_L \Sigma$; moreover $\sigma_L \phi$ preserves a symplectic structure obtained from ω by a ‘surgery’ procedure. Further, one gets a self-diffeomorphism of the standard cobordism from Σ to $\sigma_L \Sigma$, and its mapping torus is a near-symplectic broken fibration over an annulus.

Construction 1.2.8. *Let L be a compact, embedded 1-manifold in the symplectic surface (Σ, ω) . Suppose that $\phi \in \text{Aut}(\Sigma, \omega)$, $\phi(L) = L$. Then one can construct:*

- *an oriented 3-manifold-with-boundary C ;*
- *a Morse function $q: C \rightarrow [-1, 1]$, with $\partial C = q^{-1}\{1, -1\}$, whose critical points have index 1 and are in bijection with $\pi_0(L)$;*
- *an oriented diffeomorphism $F: q^{-1}(1) \cong \Sigma$;*
- *a diffeomorphism $\Phi: C \rightarrow C$, extending $F^{-1} \circ \phi \circ F$;*
- *on the mapping torus $X = \text{T}(\Phi)$, a 2-form Ω which makes the map $\text{id}_{S^1} \times q: X \rightarrow S^1 \times [-1, 1]$ into a near-symplectic broken fibration, and which extends the natural form induced by ω_ϕ on $q^{-1}(S^1 \times \{1\})$.*

The construction is canonical up to isotopy. That is, any two choices are related by a path through data of the same kind, forming smooth families in the obvious senses.

Remark 1.2.9. *The components of the surface swept out by L give a set of vanishing cycles for the near-symplectic broken fibration.*

1. The first step is to find an ‘adapted’ neighbourhood N_L of L in Σ . There are several ways to express what is meant by adapted here; one is that there should be a Hamiltonian S^1 -action on N_L whose moment map $\mu: N(L) \rightarrow \mathbb{R}$ satisfies $L = \mu^{-1}(0)$, such that $\mu \circ \phi = \pm \mu$ (the sign need only be locally constant).

The existence of an adapted neighbourhood is a consequence of the Lagrangian neighbourhood theorem, explained (in slightly different terms) in Chapter 2, Lemma 2.2.5.

The condition that ϕ preserves L and (anti)commutes with the Hamiltonian circle-action implies that it takes the following form. Each component L_i has an annular neighbourhood which may be identified with $[-\delta, \delta] \times S^1$, with standard coordinates (h, θ) . The symplectic form is $d\theta \wedge dh$, and the moment map is $\mu(h, \theta) = mh$, for some $m \in \mathbb{Z}$. The coordinates are unique up to component-wise rotations and inversions $(h, \theta) \mapsto (-h, -\theta)$ (of course one may also decrease δ). The automorphism ϕ permutes the L_i ; suppose that $\phi(L_i) = L_j$. One can choose embeddings of this type, $e_i, e_j: [-\delta, \delta] \times S^1 \rightarrow \Sigma$, near L_i and L_j , in such a way that ϕ becomes a (multiple) Dehn twist:

$$e_j^{-1} \phi \circ e_i: (h, \theta) \mapsto (h, \theta + f(h)),$$

where $f: [-\delta, \delta] \rightarrow \mathbb{R}$ takes constant, integer values near $\pm\delta$.

2. We use the notation of the paragraph headed ‘The model map Q and 2-form Θ ’ at the beginning of Section 1.2. For $\delta > 0$, define

$$S_*(\delta) = \{x \in \mathbb{R}^3 : |q(x)| \leq 1, |x_3|r^2 \leq \delta\};$$

denote by $S_b(\delta)$ the fibre over b of $q: S_*(\delta) \rightarrow [-1, 1]$. The space $S_*(\delta)$ is a smooth 3-manifold with corners: there is a ‘vertical’ boundary $S_{-1}(\delta) \cup S_1(\delta)$ and a ‘horizontal’ boundary $\{x \in S_* : |x_3|r^2 = \delta\}$, meeting in codimension-2 corners. Each $S_b(\delta)$ is a surface with boundary, except for $S_0(\delta)$, which has a singular point.

We identify the annulus $S^1 \times [-\delta, \delta]$ with $S_1(\delta)$ using the coordinates θ and $h = x_3r^2$.

Shrinking it if necessary, we may take the neighbourhood N_L to be the image of a (standardising) embedding $e: \coprod_{i=1}^n S^1 \times [-\delta, \delta] \hookrightarrow \Sigma$.

Let $\Sigma' = \Sigma \setminus e(\coprod S^1 \times (-\delta/2, \delta/2))$ and define

$$C = (\Sigma' \times [-1, 1]) \cup_O \coprod_{i=1}^n S_*(\delta),$$

where the pieces are glued together along the overlap region

$$O = \coprod_{i=1}^n S^1 \times ([-\delta, -\delta/2] \cup [\delta/2, \delta]) \times [-1, 1].$$

Here O embeds in $\Sigma' \times [-1, 1]$ using $e \times \text{id}_{[-1, 1]}$, and in $S_*(\delta)$ using the coordinates $(\theta, x_3 r^2, q)$.

As announced, C is a smooth, orientable 3-manifold. We choose the orientation compatible with $dx_1 \wedge dx_2 \wedge dx_3$ on $S_*(\delta)$. There is a unique function $C \rightarrow [-1, 1]$, given by q on $S_*(\delta)$ and by the projection on $\Sigma' \times [-1, 1]$. We still denote this function by q , and its fibres by Σ_b . This function is locally constant on ∂C , and $q^{-1}(1) = \Sigma' \cup N_L = \Sigma$.

3. We now construct the diffeomorphism Φ . On $\Sigma' \times [-1, 1]$ it is simply $\phi \times \text{id}_{[-1, 1]}$. Recall the local model

$$(\theta, h) \mapsto (\theta + f(h), h)$$

on $S^1 \times [-\delta, \delta]$. We extend this to $S_*(\delta)$: write $w = x_1 + ix_2$ and use the coordinates (w, x_3) . Set

$$\Phi(w, x_3) = (e^{2\pi i f(x_3 |w|^2)} w, x_3).$$

This is plainly a smooth map; it preserves $|w|^2$ and x_3 , hence also h and q , which implies that it is a self-diffeomorphism of $S_*(\delta)$ mapping the fibres S_b to themselves. On $S_1(\delta) = S^1 \times [-\delta, \delta]$ it matches with the model. In the gluing region this model matches with $\phi \times \text{id}_{[-1, 1]}$. We thus obtain a diffeomorphism $\Phi: C \rightarrow C$, preserving the fibres Σ_b .

The next step is to construct a suitable 2-form on the mapping torus $X = T(\Phi)$. One obtains a well-defined, closed 2-form $\beta \in Z_C^2$ by gluing together the forms

$$\text{pr}_1^* \omega \in Z_{\Sigma' \times [-1, 1]}^2, \quad x_1 dx_{2,3} + x_2 dx_{3,1} - 2x_3 dx_{1,2} \in Z_{S_*(\delta)}^2$$

(the latter is just $d\theta \wedge dh$ on the region where $r \neq 0$). Using the function $q \in C^\infty(C)$, put

$$\Omega = \beta_\Phi + dq \wedge dt \in Z_X^2.$$

This form has the right properties: it restricts to ω_ϕ over $S^1 \times \{1\}$ and is positive on the fibres of $\pi = \text{id}_{S^1} \times q$. Thus (X, π, Ω) is a near-symplectic broken fibration and $e(S^1 \times \{0\})$ sweeps out a vanishing cycle.

Uniqueness up to isotopy is apparent from the construction. The only choice involved is of the standardising tubular neighbourhood of L ; this is clearly unique up to isotopy (to be precise, this is so only once we orient L ; but reversing the orientation does not affect the outcome of the construction).

Remark 1.2.10. *The construction can be carried out starting simply from a surface Σ , a number of embedded loops L_1, \dots, L_k , and a mapping class $\gamma \in \pi_0 \text{Diff}^+(\Sigma)$ which preserves the isotopy class of each L_i . For one can choose an area form ω , and a representative $\phi \in \text{Aut}(\Sigma, \omega)$ such that $\phi(L_i) = L_i$ for each i .*

This seems a convenient place to insert two simple topological lemmas. Let $Y = \pi^{-1}(S^1 \times \{1\}) = T(\phi)$, and $\bar{Y} = \pi^{-1}(S^1 \times \{-1\}) = T(\sigma_L \phi)$.

Definition 1.2.11. If M is an module over a commutative ring, and $f \in \text{End}(M)$, we write

$$M^f = \ker(\text{id} - f), \quad M_f = \text{coker}(\text{id} - f).$$

The kernel and cokernel are, respectively, the modules of **invariants** and of **coinvariants** of f .

Lemma 1.2.12. *The restriction map $i_Y^*: H^2(X; \mathbb{Z}) \rightarrow H^2(\bar{Y}; \mathbb{Z})$ is surjective. Its kernel is isomorphic to the subgroup $\langle \text{PD}[L_1], \dots, \text{PD}[L_k] \rangle$ of the group of coinvariants $H^1(\Sigma; \mathbb{Z})_\phi$.*

Proof. Radial symplectic flow defines a deformation retraction $r: X \rightarrow Y_0 := \pi^{-1}(S^1 \times \{0\})$ (thus $r|_Y$ is the map which collapses the fibred vanishing cycle Q to Z). Surjectivity of i_Y^* follows from that of $(r \circ i_Y)^*$. But that can be seen by the five-lemma, since—regarding Y_0 as a mapping torus $T_{\Sigma_0}(\phi_0)$ —there is a commutative diagram with exact rows

$$\begin{array}{ccccccc} 0 & \longrightarrow & H^1(\Sigma_0; \mathbb{Z})_{\phi_0} & \longrightarrow & H^2(T_{\Sigma_0}(\phi_0); \mathbb{Z}) & \longrightarrow & H^2(\Sigma_0; \mathbb{Z})^{\phi_0} \longrightarrow 0 \\ & & \downarrow & & \downarrow i_Y^* r^* & & \downarrow \cong \\ 0 & \longrightarrow & H^1(\sigma_L \Sigma; \mathbb{Z})_{\sigma_L \phi} & \longrightarrow & H^2(T(\sigma_L \phi); \mathbb{Z}) & \longrightarrow & H^2(\sigma_L \Sigma; \mathbb{Z})^{\sigma_L \phi} \longrightarrow 0. \end{array}$$

(cf. Appendix B). The first vertical map, between spaces of coinvariants, is surjective because $H^1(\Sigma_0; \mathbb{Z}) \rightarrow H^1(\sigma_L \Sigma; \mathbb{Z})$ is. The third vertical map is an isomorphism because both $H^2(\Sigma_0; \mathbb{Z})$ and $H^2(\sigma_L \Sigma; \mathbb{Z})$ are free abelian on the set of components of $\sigma_L \Sigma$.

Since r^* is an isomorphism, $\ker(i_Y^*) \cong \ker(i_Y^* r^*)$. The latter is isomorphic to the kernel of the left-hand vertical map, and in turn to $\langle \text{PD}[L_1], \dots, \text{PD}[L_k] \rangle \subset H^1(\Sigma; \mathbb{Z})_\phi$. □

Lemma 1.2.13. *For each $\mathfrak{s} \in \text{Spin}^c(X)$, the class $\tau(\mathfrak{s}) \in H_2(X, Z; \mathbb{Z})$ from (1.10) satisfies*

$$\tau(\mathfrak{s}) \cap ([\Sigma] - [\sigma_L \Sigma]) = \frac{1}{2}(\chi(\sigma_L \Sigma) - \chi(\Sigma)).$$

(Here \cap can be interpreted in a number of equivalent ways, one of which is as the product $H_2(X, Z; \mathbb{Z}) \times H_2(X \setminus Z; \mathbb{Z}) \rightarrow H^2(X \setminus Z; \mathbb{Z}) \times H_2(X \setminus Z; \mathbb{Z}) \rightarrow \mathbb{Z}$.)

For example, as one crosses a circle of critical values which is the image of just one circle of critical points, passing from genus g to genus $g - 1$ fibres, the intersection with the fibre of any class $\tau(\mathfrak{s})$ decreases by 1.

This is fairly clear, and can be deduced from the previous lemma, but here is a direct proof. We have

$$\begin{aligned} 2\tau(\mathfrak{s}) \cap ([\Sigma] - [\sigma_L \Sigma]) &= 2\langle c_1(L_{\mathfrak{s}}), [\Sigma] - [\sigma_L \Sigma] \rangle \\ &= \langle c_1(\mathfrak{s})|_{X \setminus Z} - c_1(X \setminus Z, \omega), [\Sigma] - [\sigma_L \Sigma] \rangle \\ &= -\langle c_1(X \setminus Z, \omega), [\Sigma] - [\sigma_L \Sigma] \rangle. \end{aligned}$$

But the adjunction formula in $X \setminus Z$ implies that Σ has Euler characteristic $\chi(\Sigma) = \langle c_1(X \setminus Z, \omega), [\Sigma] \rangle$; similarly for $\sigma_L \Sigma$.

Separating model for $X_1 \# X_2$

We refer to [3] for interesting ways of modifying broken fibrations corresponding to the operations $X \rightarrow X \# \mathbb{C}P^2$ and $X \rightarrow X \# \overline{\mathbb{C}P^2}$. We add here a less sophisticated operation.

Suppose given broken fibrations $\pi_i: X_i \rightarrow S$, $i = 1, 2$, over the same base, a point $s \in S$ which is a regular value both of π_1 and of π_2 , with $\pi_i^{-1}(s)$ connected in both cases, and a small, closed coordinate disk $(D, 0) \hookrightarrow (S, s)$. There is then a broken fibration $\pi: X \rightarrow S$ such that

- $\pi^{-1}(S \setminus D) = \pi_1^{-1}(S \setminus D) \sqcup \pi_2^{-1}(S \setminus D)$, and on this region π is the disjoint union of π_1 and π_2 ;
- $X^{\text{crit}} \cap \pi^{-1}(D)$ is a circle;
- s is a regular value and the fibre $X_s = \pi^{-1}(s)$ is connected.

These conditions determine X up to diffeomorphism supported over D , and we will see that $X \cong X_1 \# X_2$. To construct this broken fibration, take an standard near-symplectic singularity over the annulus $D \setminus D(0; \epsilon)$, with fibres $S^1 \times [-1, 1]$ over the outer boundary and $D^2 \times S^0$ over the inner. Glue the inner boundary to a trivial fibration over $\overline{D(0; \epsilon)}$ to obtain a broken fibration $Y \rightarrow D$. It has a ‘horizontal’ boundary $S^0 \times S^1 \times D$, and a ‘vertical’ boundary $S^0 \times D^2 \times S^1$; their union is diffeomorphic to $S^0 \times S^3$ with codimension-2 corners along $S^0 \times (S^1 \times S^1)$.

Now symplectically trivialise $X_i|_D$, say $\pi_i^{-1}(D) \cong \Sigma_i \times D$, choose small disks $D_i \subset \Sigma_i$, and excise $D_i \times D \subset X_i$; one can then glue in Y in an obvious way, forming the manifold X and an induced map $\pi: X \rightarrow S$. The fibre $\pi^{-1}(s)$ is $\Sigma_1 \# \Sigma_2$, the connected sum along the D_i .

To see that X is the connected sum $X_1 \# X_2$, we need a suitable identification of Y with $S^3 \times [-1, 1]$. Let $W \subset Y$ be the union of the circle $\text{crit}(Y \rightarrow D)$ and the ‘vanishing cycles’ in the fibres of Y lying over points in the open disk $D' \subset D$ bounded by the circle of critical values. Then W is a smooth submanifold, diffeomorphic to S^3 (indeed, the map $W \xrightarrow{\pi} D' \rightarrow \bar{D}'/\partial\bar{D}' = S^2$ corresponds to the Hopf fibration). W is the union of two copies of $S^1 \times D^2$, lying over a central disk and over an annulus. By considering these two portions separately, one can see that Y is a trivial $[-1, 1]$ -fibre bundle over W ; this gives the required identification with $S^3 \times [-1, 1]$.

Chapter 2

Symmetric products and locally Hamiltonian fibrations

This chapter develops technical material needed for the Floer theory later in the thesis. The setting is that of symplectic fibrations, or more precisely, locally Hamiltonian fibrations: fibre bundles $X \rightarrow S$ equipped with closed, fibrewise-nondegenerate 2-forms Ω . When $X \rightarrow S$ is a bundle of compact Riemann surfaces, one can form its relative symmetric products $\text{Sym}_S^r(X) \rightarrow S$. The main aim of the chapter is to make $\text{Sym}_S^r(X)$ into a locally Hamiltonian fibration, with Kähler fibres, in a functorial manner. (Strictly speaking, to make the construction functorial one should keep track of a hermitian line bundle with connection over X .)

In the case where the base S is a circle, such a method was established by Salamon [38]. The crucial point is to think of the symmetric products as moduli spaces of solutions to the vortex equations. Our forms generalise Salamon's, and we also use the vortex equations, but the method of construction is somewhat different. The one important advantage of our approach is that it facilitates a description of the cohomology class of the closed 2-form in terms of the class $[\Omega] \in H^2(X; \mathbb{R})$. This will be useful in setting up both Floer homology groups and the Lagrangian matching invariant. We will not make any essential use of the geometry of these forms over bases other than S^1 . However, the greater generality may have its uses; it may on occasion be useful to have control over the Hamiltonian curvature.

2.1 Relative symmetric products

Suppose given a smooth fibre bundle $\pi: X \rightarrow S$, whose fibre is a compact orientable surface Σ . There are two constructions we can make without imposing any further structure. These

are to form the Picard bundle

$$\mathrm{Pic}_S^0(X) \rightarrow S,$$

which is a smooth fibre bundle, and the symmetric product bundles

$$\mathrm{Sym}_S^r(X) \rightarrow X, \quad r = 1, 2, \dots,$$

which at this stage are just topological fibre bundles.

The Picard bundle. As a set, $\mathrm{Pic}_S^0(X)$ consists of pairs (s, c) with $s \in S$ and $c \in \mathrm{Pic}^0(X_s) = H^1(X_s; \mathbb{R})/H^1(X_s; \mathbb{Z})$. The fibre bundle $X \rightarrow S$ has a principal $\mathrm{Diff}^+(\Sigma)$ -bundle P_X , consisting of pairs (s, f_s) where $s \in S$ and $f_s: X_s \cong \Sigma$. The Picard bundle $\mathrm{Pic}_S^0(X)$ is the associated bundle $P_X \times_{\mathrm{Diff}^+(\Sigma)} \mathrm{Pic}^0(\Sigma)$. This gives it a topology and a differentiable structure.

Symmetric product bundles. The r th symmetric product bundle, or relative symmetric product, $\mathrm{Sym}_S^r(X) \rightarrow S$, is defined to be the quotient by the symmetric group S_r of the fibre product

$$X_S^{\times r} = \{(x_1, \dots, x_r) \in X^{\times r} : \pi(x_1) = \dots = \pi(x_r)\}.$$

The fibre $\mathrm{Sym}^r(\Sigma)$ is a topological manifold, but not yet a smooth one.

Background material on (absolute) symmetric products is provided in Appendix A. We will draw on it freely.

Vertical complex structures. Further structure appears when we choose an oriented complex structure J_v on the vertical tangent bundle $T^v X = \ker(D\pi) \rightarrow X$. This is a contractible choice, because on an oriented vector bundle of rank two, complex structures correspond to fibrewise metrics modulo functions on the base.

A complex structure on Σ gives a group isomorphism

$$\mathrm{Pic}^0(\Sigma) \cong H^1(\Sigma, \mathcal{O})/H^1(\Sigma, \mathbb{Z}),$$

induced by the composition

$$H^1(\Sigma; \mathbb{R}) \hookrightarrow H^1(\Sigma; \mathbb{C}) = H^{1,0}(\Sigma) \oplus H^{0,1}(\Sigma) \rightarrow H^{0,1}(\Sigma) \cong H^1(\Sigma, \mathcal{O}).$$

Thus the fibres of the Picard bundle are now complex tori. A small variant is to consider $\mathrm{Pic}^r(\Sigma)$, the subgroup in $H^1(\Sigma, \mathcal{O}^\times)$ of degree r line bundles ($r \in \mathbb{Z}$), and its associated bundle $\mathrm{Pic}_S^r(X)$.

The symmetric product $\mathrm{Sym}^r(\Sigma)$ is also now a complex manifold; its charts are given by the elementary symmetric functions in those of Σ . The Abel-Jacobi map

$$\mathbf{aj}: \mathrm{Sym}^r(\Sigma) \rightarrow \mathrm{Pic}^r(\Sigma)$$

assigns to the divisor D the class of the holomorphic line bundle $\mathcal{O}(D)$ on Σ .

The complex structure J_v also induces a structure of smooth manifold on $\mathrm{Sym}_S^r(X) \rightarrow S$, making it a smooth fibre bundle over S . The Abel-Jacobi map globalises to a smooth and fibrewise-holomorphic mapping $\mathbf{aj}: \mathrm{Sym}_S^r(X) \rightarrow \mathrm{Pic}_S^r(X)$ over S . The smooth atlas on the relative symmetric product is generated by charts which are obtained by fibrewise application of the elementary symmetric functions to ‘restricted charts’

$$\Psi: D^2 \times U \rightarrow X.$$

This means that there is a chart $\psi: U \rightarrow S$ such that $\mathrm{pr}_2 \circ \Psi = \psi$, and more importantly that $\Psi: D^2 \times \{s\} \rightarrow X_{\psi(s)}$ is a holomorphic embedding, for each $s \in U$. As observed by Donaldson and Smith [9], the existence of such charts is a consequence of the parametrized Riemann mapping theorem. Different choices J_v^0, J_v^1 yield diffeomorphic smooth structures on $\mathrm{Sym}_S^r(X)$, as one can see by considering the relative symmetric product of $X \times [0, 1] \rightarrow S \times [0, 1]$, equipped with an interpolating family J_v^t .

Since the space \mathcal{J}_Σ of all positively oriented complex structures on Σ is contractible, the relative symmetric products of the trivial bundles $\mathcal{J}_\Sigma \times \Sigma \rightarrow \mathcal{J}_\Sigma$ are themselves trivial. Thus $\mathrm{Sym}_{\mathcal{J}_\Sigma}^r(\mathcal{J}_\Sigma \times \Sigma)$ admits a flat connection, which shows that there is a coherent system of diffeomorphisms between the $(\mathrm{Sym}^r(\Sigma), j)$. It seems, however, that there is no canonical choice of flat connection. We will use a connection which is canonical, but not flat.

2.1.1 Cohomology of relative symmetric products

There is a sequence of natural operations sending homology classes on the relative symmetric product $\mathrm{Sym}_S^r(X)$ of $X \rightarrow S$, where S is a manifold with boundary, to classes on X . These, along with formally identical constructions in slightly different settings, will appear at several points in the thesis. They come about via the universal (or tautological) divisor

$$\Delta^{\mathrm{univ}} \subset \mathrm{Sym}_S^r(X) \times_S X.$$

This carries a codimension-2 homology class

$$[\Delta^{\mathrm{univ}}] \in H_*(\mathrm{Sym}_S^r(X) \times_S X, \mathrm{Sym}_{\partial S}^r(\partial X) \times_{\partial S} \partial X; \mathbb{Z}),$$

Poincaré-Lefschetz dual to a class $\delta \in H^2(\mathrm{Sym}_S^r(X) \times_S X; \mathbb{Z})$. Thus when $X \rightarrow S$ is a holomorphic fibration, $\delta = c_1(\mathcal{O}(\Delta))$. Using the projection maps

$$\mathrm{Sym}_S^r(X) \xleftarrow{p_1} \mathrm{Sym}_S^r(X) \times_S X \xrightarrow{p_2} X$$

and iterated cap products, define, for each $k \geq 0$, the map

$$H_*(\mathrm{Sym}_S^r(X); \mathbb{Z}) \rightarrow H_{*+2-2k}(X; \mathbb{Z}); \quad u \mapsto u_{(k)} = p_{2*}((p_1^! u) \cap [\Delta^{\mathrm{univ}}]^k). \quad (2.1)$$

Similarly, one has

$$H^*(X; \mathbb{Z}) \rightarrow H^{*+2k-2}(\mathrm{Sym}_S^r(X); \mathbb{Z}); \quad c \mapsto c^{(k)} = p_{1!}((p_2^*c) \cup \delta^k). \quad (2.2)$$

These operations evidently behave in a natural way under base-change (i.e. pulling back by $S' \rightarrow S$). Unravelling the definitions, one finds that they are adjoint in the sense that

$$\langle c^{(k)}, u \rangle = \langle c, u_{(k)} \rangle.$$

Lemma 2.1.1. *We have*

$$c_1(T^\vee \mathrm{Sym}_S^r(X)) = \frac{1}{2}(c^{(1)} + 1^{(2)}), \quad c = c_1(T^\vee X).$$

Proof. We apply the Grothendieck-Riemann-Roch theorem to the map p_1 (cf. Smith [49] and the book [2] for calculations in the same vein.) Since we are dealing with a merely differentiable (not complex) family of complex manifolds, what we actually invoke is the Atiyah-Singer family index theorem; this tells us that the GRR *formula* remains valid.

The holomorphic tangent sheaf \mathcal{T}^ver of the relative symmetric product has an intrinsic description (we abbreviate Δ^{univ} to Δ):

$$\mathcal{T}^\mathrm{ver}(\mathrm{Sym}_S^r(X)) \cong p_{1*}(\mathcal{O}_\Delta(\Delta)).$$

This is an immediate generalisation of an isomorphism proved in [2]. A vertical holomorphic tangent vector at $(s, D = \sum n_i p_i)$ looks, in local coordinates z_i on X_s centred on the points p_i , like $\sum a_i \frac{\partial}{\partial z_i}$; to this one assigns the ‘Laurent tail’ $\sum a_i z_i^{-n_i} \in H^0(X_s, \mathcal{O}_D(D))$. This mapping is coordinate-independent, and globalises to a sheaf isomorphism.

GRR for p_1 says that

$$\mathrm{ch}(p_{1!} \mathcal{O}_\Delta(\Delta)) = p_{1!}(\mathrm{td}(p_2^* \mathcal{T}^\mathrm{ver} X) \cdot \mathrm{ch}(\mathcal{O}_\Delta(\Delta))),$$

since $p_2^*(\mathcal{T}^\mathrm{ver} X)$ is the tangent sheaf of $\mathrm{Sym}_S^r(X) \times_S X$ relative to $\mathrm{Sym}_S^r(X)$. But $R^i p_{1*} \mathcal{O}_\Delta(\Delta)$ is zero for $i > 0$, so

$$\mathrm{ch}(\mathcal{T}^\mathrm{ver}(\mathrm{Sym}_S^r(X))) = p_{1!}(p_2^* \mathrm{td}(\mathcal{T}^\mathrm{ver} X) \cdot \mathrm{ch}(\mathcal{O}_\Delta(\Delta)))$$

We now input

$$\mathrm{td}(\mathcal{T}^\mathrm{ver} X) = 1 + \frac{1}{2}c + \frac{1}{12}c^2 + \dots, \quad \mathrm{ch}(\mathcal{O}_\Delta(\Delta)) = e^\delta - 1.$$

Writing c also for p_2^*c , we obtain

$$\begin{aligned} \mathrm{ch}(\mathcal{T}^\mathrm{ver} \mathrm{Sym}_S^r X) &= p_{1!} \left(\delta + \frac{1}{2}(c\delta + \delta^2) + \left(\frac{1}{12}c^2\delta + \frac{1}{4}c\delta^2 + \frac{1}{6}\delta^3 \right) + \dots \right) \\ &= 1^{(1)} + \frac{1}{2}(c^{(1)} + 1^{(2)}) + \left(\frac{1}{12}(c^2)^{(1)} + \frac{1}{4}c^{(2)} + \frac{1}{6}1^{(3)} \right) + \dots \end{aligned}$$

Since $c_1 = \mathrm{ch}_1$, the result follows. In general, ch_k will be a sum of the terms $(c^i)^{(k+1-i)}$, with rational coefficients that could be computed from the Bernoulli numbers. \square

2.2 Locally Hamiltonian fibrations

Here we collect some standard material concerning fibre bundles equipped with closed, fibre-symplectic 2-forms.

Definition 2.2.1. (a) A **locally Hamiltonian fibration** (X, π, Ω) is a smooth fibre bundle $\pi: X \rightarrow S$ over a manifold S , with a closed 2-form $\Omega \in Z_X^2$ on the total space, such that $\Omega_s := \Omega|_{\pi^{-1}(s)}$ is non-degenerate on $X_s := \pi^{-1}(s)$, for each $s \in S$. Sometimes, when the map is clear from the context, we will just write (X, Ω) .

(b) When (X, π, Ω_0) and (X, π, Ω_1) are both locally Hamiltonian fibrations, they are called **isotopic** if there exists a closed form $\Omega \in Z_{[0,1] \times X}^2$ making $([0, 1] \times X, \pi \times \text{id}, \Omega)$ locally Hamiltonian, with $\Omega|_{\{i\} \times X} = \Omega_i$.

(c) Locally Hamiltonian fibrations are called **equivalent** if they are related under the equivalence relation generated by isotopy and 2-form-preserving bundle isomorphism.

Local geometry

A locally Hamiltonian fibration (X, π, Ω) defines a subbundle $T^h X$ of TX complementary to $T_x^\vee X = \ker(D_x \pi)$: the fibre $T_x^h X \subset T_x X$ is the Ω_x -annihilator of $T_x^\vee X$. The horizontal subbundle $T^h X$ defines a connection on $X \rightarrow S$. The connection has the usual paraphernalia. Horizontal lifting

$$TS \rightarrow T^h X; \quad v \mapsto \tilde{v}$$

gives rise (when the horizontal vector field can be integrated, for instance when π is proper) to parallel transport maps

$$m_\gamma: X_{\gamma(a)} \rightarrow X_{\gamma(b)}, \quad \gamma: [a, b] \rightarrow S,$$

satisfying a composition rule. If one fixes a basepoint $b \in S$ and a diffeomorphism $X_b \rightarrow M$, then a fibre bundle $X \rightarrow S$ gives rise to a principal $\text{Diff}(M)$ -bundle $P_X \rightarrow S$ of frames for the fibres of $X \rightarrow S$. The connection is then encoded in a 1-form

$$A(\Omega) \in \Omega^1(P_X; \Gamma(TM)).$$

Locally on the base, in a fixed trivialisation, it may be thought of as a 1-form on S with values in $\Gamma(TM)$.

We do not discuss the curvature of the connection here, since it will not play any role in the thesis.

The first part of the following standard lemma (compare e.g. [31, 46]) gives the crucial property of locally Hamiltonian fibrations.

Lemma 2.2.2. *Assume that the parallel transport maps are defined for all paths in S ; this is true if π is proper.*

1. The parallel transport maps are symplectomorphisms.
2. Choose an identification $\iota : (M, \omega) \rightarrow (X_s, \Omega_s)$. Let D^n denote the open unit disk with coordinates x_1, \dots, x_n , and let $\chi : (D^n, 0) \rightarrow (S, s)$ be a chart. There is then a trivialisation $\tilde{\chi} : D \times M \rightarrow \chi^* X$, with $\tilde{\chi}|_{\{0\}} \times M = \iota$, such that

$$\tilde{\chi}^* \Omega = \omega + \sum_i \eta_i \wedge dx_i + \sum_{i < j} \zeta_{ij} dx_i \wedge dx_j. \quad (2.3)$$

for forms $\eta_i \in \Omega_M^1$ and $\zeta_{ij} \in \Omega_M^0$ (these are not intrinsic to Ω). Any two such trivialisations differ by a map $(D^n, 0) \rightarrow (\text{Aut}(M, \omega), \text{id}_M)$.

3. The form

$$A = \sum_i \eta_i \wedge dx_i \in \Omega^1(D, Z^1(M))$$

appearing in (2.3) has the property that, for tangent vectors $(h, v) \in T_s D \times T_p M$, the relation $A(h) = \iota(v)\omega$ holds iff $(h, v) \in T^h(D \times M)$. This means that A is ω -dual to the connection form $A(\Omega)$: writing $\iota(\xi_i)\omega = \eta_i$, we have

$$A(\Omega) = \sum_i dx_i \otimes \xi_i.$$

Proof. (1) is Moser's lemma in geometric form. There is no loss in considering a trivial bundle $[0, 1] \times M \rightarrow [0, 1]$, with 2-form $\Omega = \omega_t + \eta_t \wedge dt$. Then $\tilde{\partial}_t = \partial_t + v_t$, where $\iota(v_t)\omega_t = \eta_t$. Since $d\Omega = 0$, we have $d\eta_t + \dot{\omega}_t = 0$. The flow $\{m_t\}_{t \in [0, 1]}$ of v_t on M is then symplectic: $m_t^* \omega_t = \omega_0$. But this flow is the parallel transport $m_{[0, t]}$.

In (2), one can construct $\tilde{\chi}$ using radial parallel transport. (3) is easily checked. \square

The same argument as (1) gives the following 'Moser lemma for families':

Lemma 2.2.3. *Isotopic proper locally Hamiltonian fibrations are isomorphic: in $X \times [0, 1] \rightarrow S \times [0, 1]$, parallel transport along the paths $[0, 1] \times \{s\}$ in $[0, 1] \times S$ induces an isomorphism between $X \times \{0\}$ and $X \times \{1\}$.*

Locally Hamiltonian fibrations over $S^1 = \mathbb{R}/\mathbb{Z}$

Given (M, ω) and $\phi \in \text{Aut}(M, \omega)$, the mapping torus $(T(\phi), \omega_\phi)$ is a flat locally Hamiltonian fibration over S^1 . Here $T(\phi) = \mathbb{R} \times M / \sim$, where $(t, x) \sim (t - 1, \phi(x))$, and ω_ϕ denotes the form induced by the pullback of ω to $M \times \mathbb{R}$. Conversely, for any locally Hamiltonian fibration (X, π, Ω) over S^1 , the map

$$X_t \ni x \mapsto (t, (m_{[0, t]})^{-1}x) \in \mathbb{R} \times X_0$$

induces an isomorphism of (X, π, Ω) with the mapping torus of its monodromy $m_{[0,1]} \in \text{Aut}(X_0, \Omega_0)$.

The topology of mapping tori is reviewed in Appendix B. This includes a description of the de Rham complex.

Consider a family of closed 1-forms $\{\eta_t\}_{t \in \mathbb{R}}$ satisfying $\phi^* \eta_t = \eta_{t+1}$, so that $\eta_t \wedge dt$ descends to a closed 2-form on $T(\phi)$.

Lemma 2.2.4. *The following are equivalent:*

1. *the locally Hamiltonian fibration $(T(\phi), \omega_\phi + \eta_t \wedge dt)$ is isotopic to $(T(\phi), \omega_\phi)$;*
2. *$\eta_t \wedge dt$ is an exact 2-form on $T(\phi)$;*
3. *the class $\int_0^1 [\eta_t] dt \in H^1(M; \mathbb{R})$ lies in the image of $\text{id} - \phi^*$.*

Proof. (1) \Leftrightarrow (2): straightforward.

(2) \Rightarrow (3): say $\eta_t \wedge dt = d(\theta_t + h_t dt)$, where $\theta_t \in Z^1(M)$ and h_t are ϕ -periodic. Then $\eta_t = dh_t - \partial_t \theta_t$, and $\int_0^1 [\eta_t] dt = [\theta_0] - [\theta_1] = (\text{id} - \phi^*)[\theta_0]$.

(3) \Rightarrow (2): if $\int_0^1 [\eta_t] dt \in \text{im}(\text{id} - \phi^*)$ then $(\int_0^1 \eta_t dt, 0)$ is a coboundary in $\text{cone}(\text{id} - \phi^*)$, cf. Appendix B. This means that $dt \wedge \eta_t$ is exact on $T(\phi)$. □

The vector fields dual to the 1-forms η_t generate a symplectic isotopy $\{\psi_t\}_{t \in \mathbb{R}}$, and the monodromy of $(T(\phi), \omega + \eta_t \wedge dt)$ is $\phi \circ \psi_1$. Thus $(T(\phi), \omega + \eta_t \wedge dt) \cong (T(\phi \circ \psi_1), \omega_{\phi \circ \psi_1})$. We have, by definition,

$$\int_0^1 [\eta_t] dt = \text{Flux}(\{\psi_t\}) \in H^1(M; \mathbb{R}).$$

When $\phi \in \text{Aut}_0(M, \omega)$, the lemma specialises to the well-known fact that, for a symplectic isotopy ψ_t starting at $\psi_0 = \text{id}_M$, the map ψ_1 is Hamiltonian if and only if $\text{Flux}(\{\psi_t\}) = 0$.

We also make a note of a version of the Lagrangian neighbourhood theorem:

Lemma 2.2.5. *Let (X, π, Ω) be a compact locally Hamiltonian fibration over S^1 , with fibres of dimension $2n$, and suppose that $Q \subset X$ is a smooth sub-bundle, with fibres of dimension n , such that $\Omega|_Q = 0$. Then there is a neighbourhood N of Q , and a diffeomorphism Ψ from a neighbourhood of the zero-section in $T_v^*Q := \{\lambda \in T^*Q : \lambda(T^h X) = 0\}$ to N , such that*

- $\pi \circ \Psi$ is the natural projection $T_v^*Q \rightarrow B$;
- Ψ acts as the identity along the zero-section $Q \subset T_v^*Q$;
- $\Psi^* \Omega$ equals the restriction to T_v^*Q of the canonical two-form on T^*Q .

We ignore the question of whether this is true for more general bases than S^1 . As stated, it is an easy consequence of the usual Lagrangian neighbourhood theorem. We may assume that $X = T(\phi)$ and $\Omega = \omega_\phi$, where $\phi \in \text{Aut}(M, \omega)$. The sub-bundle Q is preserved by parallel transport, and is thus the image in $T(\phi)$ of $\mathbb{R} \times L \subset \mathbb{R} \times M$, where $L \subset M$ is a Lagrangian submanifold. Now, $L \subset M$ has a neighbourhood symplectomorphic to a neighbourhood of the zero-section in T^*L . Thus we can replace (M, L) by (T^*L, L) , and ω by the canonical two-form $d\lambda_{\text{can}}$. In this case the conclusion is clear.

Topological interpretation. The topological interpretation of locally Hamiltonian fibrations is surprisingly subtle, and here we just remark briefly on it. Let (X, π, Ω) be a locally Hamiltonian fibration equipped with an identification $i: (M, \omega) \rightarrow (X_s, \Omega_s)$. Then Ω determines a reduction of the structure group of π from $\text{Diff}^+(M)$ to $\text{Aut}(M, \omega)$ —this is the topological content of Lemma 2.2.2, (1)—and $X \rightarrow S$ becomes a symplectic bundle. It also determines a reduction of the structure group of γ^*X to the Hamiltonian group $\text{Ham}(M, \omega)$, whenever $\gamma: S' \rightarrow S$ is a smooth map from a simply connected space S' (this is the reason for the name ‘locally Hamiltonian’).

There are partial converses to these statements.

- A symplectic fibration $(M, \omega) \hookrightarrow X \rightarrow S$ underlies a locally Hamiltonian fibration iff the class $[\omega] \in H^2(M; \mathbb{R})$ extends to a class on X . (This always holds when $S = S^1$.)
- A symplectic fibration $(M, \omega) \hookrightarrow X \rightarrow S$ admits a reduction from $\text{Aut}(M, \omega)$ to $\text{Ham}(M, \omega)$ iff the symplectic class extends to X and $\gamma^*X \rightarrow S^1$ is a trivial symplectic bundle for every loop $\gamma: S^1 \rightarrow S$.

The first statement is implied by Thurston’s patching argument. For proofs of the second, see [31, 24].

Very recent work by McDuff [30] improves on this: there is a topological group $\text{Ham}^s(M, \omega)$ (homotopy-equivalent to a CW complex) and a continuous, injective homomorphism

$$\text{Ham}^s(M, \omega) \hookrightarrow \text{Aut}(M, \omega)$$

with the property that a symplectic fibration with typical fibre (M, ω) admits a reduction to $\text{Ham}^s(M, \omega)$ iff it underlies a locally Hamiltonian fibration.

2.3 The vortex equations

2.3.1 Moduli spaces of vortices

Fix a closed Riemann surface (Σ, j) , a Kähler form $\omega \in \Omega_{\Sigma}^{1,1}$, and a hermitian line bundle $(L, |\cdot|)$ over Σ , of degree $r > 0$, and introduce the following notation:

- $\mathcal{C}(L, |\cdot|)$, or $\mathcal{C}(L)$, is the space of $U(1)$ -connections (an affine space modelled on the imaginary 1-forms $i\Omega_\Sigma^1$).
- $\mathcal{G} = U(1)^\Sigma$ is the gauge group.
- $i\Omega_\Sigma^0$ is the Lie algebra of \mathcal{G} . The pairing

$$\Omega_\Sigma^0 \otimes i\Omega_\Sigma^0 \rightarrow \mathbb{R}, \quad f \otimes ig \mapsto \int_\Sigma fg \omega$$

embeds Ω_Σ^0 into the dual of the Lie algebra. We consider moment maps for Hamiltonian \mathcal{G} -actions as maps into Ω_Σ^0 .

Connections, sections and gauge transformations are by default C^∞ , and the spaces are given their C^∞ topologies. We also need

- \mathcal{A}_1^2 , the space of pairs (A, ψ) of unitary connection A on L and section ψ of L , each of Sobolev class L_1^2 (an L_1^2 connection differs from a smooth one by an L_1^2 form).
- $L_2^2(\Sigma, U(1))$, the group of gauge transformations of Sobolev class L_2^2 . Note that a map $\Sigma \rightarrow \mathbb{C}$ of class L_2^2 is continuous, hence has a pointwise norm.

Action of the gauge group

Action on connections. The conformal structure j induces a Kähler structure on $\mathcal{C}(L)$. Its 2-form is independent of j ; it is

$$(a_1, a_2) \mapsto \int_\Sigma ia_1 \wedge ia_2, \quad a_1, a_2 \in i\Omega_\Sigma^1. \quad (2.4)$$

The complex structure is the Hodge star $a \mapsto *_j a$. The 2-form is exact, equal to dp , where

$$p(A; a) = \frac{1}{2} \int_\Sigma ia \wedge i(A_0 - A), \quad \text{any } A_0 \in \mathcal{C}(L). \quad (2.5)$$

The action of the gauge group \mathcal{G} on $\mathcal{C}(L)$ is Hamiltonian, with (equivariant) moment map

$$\mathcal{C}(L) \rightarrow i\Omega_\Sigma^0; \quad A \mapsto *iF_A. \quad (2.6)$$

Action on sections. The symplectic form ω induces a Kähler structure on $\Omega_\Sigma^0(L)$, with 2-form

$$(\phi_1, \phi_2) \mapsto \int_\Sigma \text{Im} \langle \phi_1, \phi_2 \rangle \omega, \quad \phi_1, \phi_2 \in \Omega_\Sigma^0(L) \quad (2.7)$$

and complex structure $\phi \mapsto i\phi$. The 2-form is equal to dq , where

$$q(\psi; \phi) = \frac{1}{2} \int_\Sigma \text{Im} \langle \psi, \phi \rangle \omega.$$

The gauge-action on $\Omega_\Sigma^0(L)$ is again Hamiltonian; the moment map is

$$\Omega_\Sigma^0(L) \rightarrow \Omega_\Sigma^0; \quad \psi \mapsto \frac{1}{2}|\psi|^2. \quad (2.8)$$

Action on pairs. The manifold $\mathcal{C}(L) \times \Omega_\Sigma^0(L)$ carries a product Kähler structure σ , which depends on both j and ω . The moment map μ for the diagonal \mathcal{G} -action is the sum of the moment maps of the factors,

$$\mu(A, \psi) = *iF_A + \frac{1}{2}|\psi|^2 \in \Omega_\Sigma^0. \quad (2.9)$$

The Chern-Weil formula gives some basic information about this moment map:

$$\frac{1}{2\pi} \int_\Sigma \tau \omega \begin{cases} < r : & \mu^{-1}(\tau) = \emptyset; \\ = r : & \mu(A, \psi) = \tau \Rightarrow \psi \equiv 0; \\ > r : & \mu(A, \psi) = \tau \Rightarrow \psi \neq 0. \end{cases}$$

In fact, μ is submersive at (A, ψ) precisely when $\psi \neq 0$, which is also the locus on which the gauge-action is free. When $\int_\Sigma \tau \omega > 2\pi r$, the free gauge-action on $\mu^{-1}(\tau)$ admits local slices (see below), so the Kähler quotient $\mu^{-1}(\tau)/\mathcal{G}$ is a Kähler manifold.

The vortex equations

The **vortex equations** with parameter τ are the following coupled equations for a pair $(A, \psi) \in \mathcal{C}(L) \times \Omega_\Sigma^0$:

$$\bar{\partial}_A \psi = 0 \quad \text{in } \Omega_\Sigma^{0,1}(L), \quad (2.10)$$

$$\mu(A, \psi) = \tau \quad \text{in } \Omega_\Sigma^0. \quad (2.11)$$

Individually, we will refer to them as the Cauchy-Riemann equation and the moment map equation. The space of solutions $\mathcal{U}_\Sigma(L, \tau)$ is invariant under \mathcal{G} , and the quotient space $\text{Vor}(L, \tau) = \mathcal{U}_\Sigma(L, \tau)/\mathcal{G}$ is called the **vortex moduli space**. The fundamental results about $\text{Vor}(L, \tau)$ are as follows.

Theorem 2.3.1. *Assume $\int \tau \omega > 2\pi r$.*

(a) *The space $\text{Vor}(L, \tau)$ is a finite-dimensional, complex—therefore smooth and Kähler—submanifold of $\mu^{-1}(\tau)/\mathcal{G}$.*

(b) *The map*

$$Z: \text{Vor}(L, \tau) \rightarrow \text{Sym}^r(\Sigma), \quad [A, \psi] \mapsto \psi^{-1}(0)$$

is an isomorphism of complex manifolds.

The unitary connection A induces a holomorphic structure on L ; a local section is holomorphic iff it lies in $\ker \bar{\partial}_A$. By means of the holomorphic structure, one attaches multiplicities to points of $\psi^{-1}(0)$, so that ψ has r zeros in all. This makes sense of Z .

(a) is proved by an elliptic regularity argument. As for (b), the statement that Z is bijective is an existence and uniqueness theorem for solutions to the vortex equations. This is the heart of the theorem, and various proofs are known (see Jaffe and Taubes [20] and García-Prada [13] for two of them).

Addendum 2.3.2. *When $\int \tau\omega = 2\pi r$, the moduli space*

$$\text{Vor}(L, \tau) = \{(A, 0) : iF_A = \tau\omega\}/\mathcal{G}$$

is a finite-dimensional, complex—therefore smooth and Kähler—submanifold of $\mu^{-1}(\tau)/\mathcal{G}$. The map

$$\text{Vor}(L, \tau) \rightarrow \text{Pic}^L(\Sigma); \quad [A, 0] \mapsto L_A$$

is an isomorphism of complex manifolds.

Smoothness of the moduli space

This is a standard application of elliptic theory. The tangent space to the affine space \mathcal{A}_1^2 is the space of pairs (a, ϕ) , where a is an imaginary 1-form, ϕ a section, both of class L_1^2 . Using weak derivatives, define

$$v_\tau(A, \psi) = (\bar{\partial}_A \psi, *iF_A + \frac{1}{2}|\psi|^2 - \tau).$$

Supposing $v_\tau(A, \psi) = 0$, set

$$\mathcal{N}_{A, \psi, \epsilon} = \{(a, \phi) : v_\tau(A + a, \psi + \phi) = 0, |(a, \phi)|_{L_1^2} < \epsilon, (2.12)\}.$$

Here (2.12) is the gauge-fixing equation

$$d^*(ia) + \text{Im}\langle \psi, \phi \rangle = 0, \tag{2.12}$$

which says that (a, ϕ) is orthogonal to the gauge-orbit of (A, ψ) . Note that its left-hand side is gauge-equivariant. The linearisations of the two vortex equations at the solution (A, ψ) are

$$\bar{\partial}_A \phi + a^{0,1}\psi = 0, \quad *ida + \text{Re}\langle \psi, \phi \rangle = 0. \tag{2.13}$$

The second of these and (2.12) are real and imaginary parts of the single equation

$$\bar{\partial}^*(a^{0,1}) - \frac{1}{2}\langle \psi, \phi \rangle = 0. \tag{2.14}$$

Hence the space of solutions to equations (2.13, 2.12) is the kernel of the \mathbb{C} -linear differential operator

$$D_{(A,\psi)} : (a, \phi) \mapsto (\bar{\partial}_A \phi + a^{0,1} \psi, \bar{\partial}^*(a^{0,1}) - \frac{1}{2} \langle \psi, \phi \rangle). \quad (2.15)$$

$D_{(A,\psi)}$ is a compact perturbation of the Fredholm operator $(a, \phi) \mapsto (\bar{\partial}_A \phi, \bar{\partial}^*(a^{0,1}))$, which has (complex) index $(r+1-g) - (1-g) = r$, and is hence also Fredholm of index r . It is surjective, as can be seen by computing $D_{(A,\psi)}^* D_{(A,\psi)}$ (see [38]). Hence its kernel has constant rank r .

For small ϵ , the map $(a, \phi) \mapsto [A + a, \psi + \phi]$ gives an embedding of $\mathcal{N}_{A,\psi,\epsilon}$ as a neighbourhood of $[A, \psi]$ in the moduli space $\text{Vor}(L, \tau)$ (elliptic regularity shows that solutions of $v_\tau = 0$ are smooth). On the other hand, an implicit function theorem argument identifies $\mathcal{N}_{A,\psi,\epsilon}$ with $\ker D_{(A,\psi)}$, showing that $\text{Vor}(L, \tau)$ is a smooth manifold of \mathcal{B}^* . Since its tangent spaces are complex linear, we conclude that $\text{Vor}(L, \tau)$ is indeed an r -dimensional complex submanifold of \mathcal{B}^* .

Relations with complex geometry

Certain aspects of the geometry of $\text{Sym}^r(\Sigma)$ are visible on the vortex moduli space—for example, the Abel-Jacobi map. There is a commutative diagram

$$\begin{array}{ccc} \text{Vor}(L, \tau) & \xrightarrow{[A,\psi] \mapsto L_A} & \text{Pic}^L(\Sigma) \\ \downarrow z & & \downarrow \\ \text{Sym}^r(\Sigma) & \xrightarrow{\mathbf{aj}} & \text{Pic}^r(\Sigma) \end{array}$$

in which the vertical arrows are isomorphisms of complex manifolds. On the right are two realisations of the Picard torus: $\text{Pic}^L(\Sigma)$, the space of holomorphic structures on L , and $\text{Pic}^r(\Sigma)$, the affine subspace $c_1^{-1}(r)$ of $H^1(\Sigma, \mathcal{O}^\times)$.

Another such aspect concerns Chern forms and the Poincaré-Lelong formula. The Chern form $c_1(L, |\cdot|)$ of a hermitian holomorphic line bundle is the closed, integral $(1, 1)$ -form given locally by $i\pi^{-1} \partial \bar{\partial} \log(|s|)$ for any *non-vanishing* holomorphic section s . The Chern form of L_A is $\frac{i}{2\pi} F_A$. For any holomorphic section ψ of a holomorphic, hermitian line bundle, $u := \log |\psi|^2$ lies in $L^1(\Sigma)$. According to the Poincaré-Lelong formula, there is an equality of closed integral *currents*

$$\frac{i}{2\pi} \partial \bar{\partial} u + \delta_{Z(\psi)} = c_1(L, |\cdot|).$$

When this formula is applied to a pair $(A, \psi) \in \mathcal{U}_\Sigma(L, \tau)$, the curvature form on the right can be rewritten using (2.11), to give the equation

$$\frac{i}{2\pi} \partial \bar{\partial} u + \delta_Z = \frac{1}{2\pi} (\tau - \frac{1}{2} e^u) \omega, \quad u = \log |\psi|^2. \quad (2.16)$$

This equation can also be obtained directly. Indeed, if we trivialise L over a neighbourhood of a point where ψ is non-zero, (2.10) may be rewritten as $A^{0,1} = -\bar{\partial}(\log \psi)$. The Dirac-delta correction δ_Z can be inferred from Taylor expansions of ψ .

In the approach of [20], equation (2.16) is fundamental. There it is shown, using variational methods, that for any $Z \in \text{Sym}^r(\Sigma)$ there exists a unique solution $u = u_Z$, and that e^{u_Z} is smooth.

2.3.2 The Kähler class on the vortex moduli space

As we have seen, the moduli space $\text{Vor}(L, \tau)$ is, in a canonical way, a Kähler manifold. Pulling back its Kähler form σ_τ by Z^{-1} , we obtain a Kähler form on $\text{Sym}^r(\Sigma)$, which we also call σ_τ . The target of this section is to prove a formula for its cohomology class. Under the standard isomorphism $H^2(\text{Sym}^r(\Sigma); \mathbb{Z}) \cong H_0(\Sigma; \mathbb{Z}) \oplus \Lambda^2 H_1(\Sigma; \mathbb{Z})$, η maps to $[\text{pt}] \in H_0(\Sigma; \mathbb{Z})$ and θ to the intersection form $\cap \in \Lambda^2 H^1(\Sigma; \mathbb{Z})$; cf. Appendix A.

Theorem 2.3.3. *The following equation holds in $H^2(\text{Sym}^r(\Sigma); \mathbb{R})$:*

$$\frac{1}{2\pi}[\sigma_\tau] = \left(\int_{\Sigma} \tau \omega \right) \eta + 2\pi(\theta - r\eta).$$

Our method is to exhibit connections on two line bundles over the orbit space of irreducible pairs, $\mathcal{A}^*/\mathcal{G}$. The Chern classes of these line bundles restrict to η and $\theta - r\eta$ on $\text{Vor}(L, \tau)$, while the appropriate linear combination of their curvature forms restricts exactly to the form σ_τ .

Remark 2.3.4. *The formula specialises, when $\tau \equiv 1/2$, to one which appears in the mathematical physics literature (Manton-Nasir [32]). That work, while rigorous in itself, relies on a local expansion of the Kähler form [39] which deserves more thorough analytic treatment.*

Cohomology of the orbit space

We write \mathcal{A}^* for the space of pairs $(A, \psi) \in \mathcal{C}(L) \times \Omega_{\Sigma}^0$ with ψ not identically zero, \mathcal{B}^* for the orbit space $\mathcal{A}^*/\mathcal{G}$, and $i: \text{Vor}(L, \tau) \rightarrow \mathcal{B}^*$ for the inclusion.

Lemma 2.3.5. *i induces a surjection on cohomology, and an isomorphism on $H^{\leq 2}$.*

Proof. One can use slant products to set up maps

$$\begin{aligned} \mu_1: H_*(\Sigma; \mathbb{Z}) &\rightarrow H^{2-*}(\mathcal{B}^*; \mathbb{Z}), & h &\mapsto c_1(\mathcal{L}_1)/h, \\ \mu_2: H_*(\Sigma; \mathbb{Z}) &\rightarrow H^{2-*}(\text{Sym}^r(\Sigma); \mathbb{Z}), & h &\mapsto c_1(\mathcal{L}_2)/h. \end{aligned}$$

Here the line bundle $\mathcal{L}_1 \rightarrow \mathcal{B}^* \times \Sigma$ is $\mathcal{L}_1 = \widetilde{\mathcal{L}_1}/\mathcal{G}$, where the equivariant line bundle $\widetilde{\mathcal{L}_1} \rightarrow \mathcal{A}^* \times \Sigma$ is the pullback of $L \rightarrow \Sigma$; and $\mathcal{L}_2 \rightarrow \text{Sym}^r(\Sigma) \times \Sigma$ is the topological line bundle corresponding to the universal divisor $\Delta^{\text{univ}} \subset \text{Sym}^r(\Sigma) \times \Sigma$.

These maps extend uniquely to homomorphisms of graded rings

$$\begin{aligned}\Lambda^* H_1(\Sigma; \mathbb{Z}) \otimes_{\mathbb{Z}} \mathbb{Z}[H_0(\Sigma)] &\rightarrow H^*(\mathcal{B}^*; \mathbb{Z}), \\ \Lambda^* H_1(\Sigma; \mathbb{Z}) \otimes_{\mathbb{Z}} \mathbb{Z}[H_0(\Sigma)] &\rightarrow H^*(\text{Sym}^r(\Sigma); \mathbb{Z}),\end{aligned}$$

since the graded ring on the left is freely generated by $H_0(\Sigma; p\mathbb{Z}) \oplus H_1(\Sigma; \mathbb{Z})$. The grading is such that $H_i(\Sigma; \mathbb{Z})$ has degree $2 - i$. The first of these two maps is an isomorphism (see Atiyah and Bott [1, pp. 539–545]). The second is surjective, since the image of μ_2 contains $H^1(\text{Sym}^r(\Sigma); \mathbb{Z})$ and the class η , and these generate the cohomology ring.

To prove the lemma it suffices to show that $i^* \circ \mu_1 = Z^* \circ \mu_2$. This follows from the fact that $(i \times 1)^* \mathcal{L}_1$ is isomorphic to $(Z \times 1)^* \mathcal{L}_2$. To see that these bundles are isomorphic, observe that the former has a tautological section which vanishes precisely along Δ^{univ} . \square

It is convenient to have some notation to hand for integral (co)homology classes on Σ . Take a standard symplectic basis for the homology,

$$\begin{aligned}e_0 &\in H_0(\Sigma), \\ \{\alpha_i, \beta_j\}_{1 \leq i, j \leq g} &\in H_1(\Sigma), \\ e_2 = [\Sigma] &\in H_2(\Sigma),\end{aligned}$$

and the dual basis for cohomology,

$$\begin{aligned}e^0 = 1 &\in H^0(\Sigma), \\ \{\alpha^i, \beta^j\}_{1 \leq i, j \leq g} &\in H^1(\Sigma), \\ e^2 &\in H^2(\Sigma).\end{aligned}$$

Put

$$\tilde{\eta} = \mu_1(e_0), \quad \tilde{\theta} = \sum_{i=1}^g \mu_1(\alpha_i) \cup \mu_1(\beta_i). \quad (2.17)$$

Lemma 2.3.6. $c_1(\mathcal{L}_1)^2/e_2 = 2r\tilde{\eta} - 2\tilde{\theta}$ in $H^2(\mathcal{B}^*; \mathbb{Z})$.

Proof. The group $H^2(\mathcal{B}^* \times \Sigma; \mathbb{Z})$ is the direct sum of its Künneth components $H^0(\mathcal{B}^*; \mathbb{Z}) \otimes H^2(\Sigma; \mathbb{Z})$, $H^1(\mathcal{B}^*; \mathbb{Z}) \otimes H^1(\Sigma; \mathbb{Z})$ and $H^2(\mathcal{B}^*; \mathbb{Z}) \otimes H^0(\Sigma; \mathbb{Z})$. The Chern class $c_1(\mathcal{L}_1)$ is *tautologically* the sum of

$$\begin{aligned}\mu_1(e_2) \otimes e^2 &\in H^0(\mathcal{B}^*; \mathbb{Z}) \otimes H^2(\Sigma; \mathbb{Z}), \\ \sum_{i=1}^g (\mu_1(\alpha_i) \otimes \alpha^i + \mu_1(\beta_i) \otimes \beta^i) &\in H^1(\mathcal{B}^*; \mathbb{Z}) \otimes H^1(\Sigma; \mathbb{Z}), \\ \mu_1(e_0) \otimes e^0 &\in H^2(\mathcal{B}^*; \mathbb{Z}) \otimes H^0(\Sigma; \mathbb{Z}).\end{aligned}$$

Let us call these terms A , B and C respectively. Note that $A = r.1 \otimes e^2$ (by definition of \mathcal{L}_1) and $C = \tilde{\eta} \otimes e^0$. The Künneth isomorphism is compatible with cup products, providing that one uses the *graded* tensor product of graded rings. Thus $A \cup C = r\tilde{\eta} \otimes e^2 = C \cup A$, and

$$B^2 = \left(\sum_{i=1}^g \mu_1(\alpha_i) \cup \mu_1(\beta_i) - \mu_1(\beta_i) \cup \mu_1(\alpha_i) \right) \otimes e^2 = -2\tilde{\theta} \otimes e^2.$$

Hence $c_1(\mathcal{L}_1)^2/e_2 = 2(r\tilde{\eta} - \tilde{\theta})$.

□

A connection on \mathcal{L}_1

We now write down a certain connection ∇ on \mathcal{L}_1 , and compute its curvature. This calculation is modelled on that of Donaldson and Kronheimer [8, p. 195]. We will use the curvature form, together with its wedge-square, to construct a closed 2-form on \mathcal{B}^* , representing a known cohomology class, whose restriction to $\text{Vor}(L, \tau)$ is σ_τ .

The connection ∇ is built, using the quotient construction explained in [8], from

- a certain unitary, \mathcal{G} -invariant connection $\hat{\nabla}$ on the line bundle $\text{pr}_2^*L \rightarrow \mathcal{A}^* \times \Sigma$; and
- a certain connection Γ on the principal \mathcal{G} -bundle $\mathcal{A}^* \rightarrow \mathcal{B}^*$, pulled back to $\mathcal{B}^* \times \Sigma$.

The connection $\hat{\nabla}$ is trivial in the \mathcal{A}^* -directions and tautological in the Σ -directions. To amplify: a section of pr_2^*L is a map $s: \mathcal{A}^* \times \Sigma \rightarrow L$ with $s(A, \psi, x) \in L_x$, and at the point (A, ψ, x) ,

$$\hat{\nabla}_{(a, \phi, v)} s = d_{A, v}(s|_{\{A, \psi\} \times \Sigma})(x) + \left(\frac{d}{dt} \Big|_{t=0} s(A + ta, \psi + t\phi, x) \right)(x). \quad (2.18)$$

The curvature of $\hat{\nabla}$ is given by

$$\begin{aligned} F_{\hat{\nabla}}((0, 0, u), (0, 0, v)) &= F_A(u, v), \\ F_{\hat{\nabla}}((a, \phi, 0), (0, 0, v)) &= \langle a, v \rangle, \\ F_{\hat{\nabla}}((a, \phi, 0), (a', \phi', 0)) &= 0. \end{aligned} \quad (2.19)$$

We can obtain a connection on $\mathcal{A}^* \rightarrow \mathcal{B}^*$ from our gauge-fixing condition: the horizontal space over $[A, \psi]$ is the kernel of the linear operator

$$(a, \phi) \mapsto d^*(ia) - \text{Im}\langle \psi, \phi \rangle.$$

To write down the connection 1-form Γ , we need the Green's operator G_ψ associated to the Laplacian

$$\Delta_\psi = d^*d + |\psi|^2: \quad \Omega_\Sigma^0 \rightarrow \Omega_\Sigma^0.$$

Δ_ψ is surjective (since d^*d maps onto the functions of mean-value zero), inducing an isomorphism of $\ker(\Delta_\psi)^\perp$ with Ω_Σ^0 ; its inverse is G_ψ .

Lemma 2.3.7. *The connection 1-form Γ is given by*

$$\Gamma_{(A,\psi)}(a, \phi) = iG_\psi(d^*ia - \text{Im}\langle\psi, \phi\rangle) \in i\Omega_\Sigma^0.$$

Proof. This form has the correct kernel, so to justify the assertion one simply observes that it is invariant under \mathcal{G} :

$$\Gamma_{(A,\psi,x)}(-df, f\psi, 0) = f, \quad f \in i\Omega_\Sigma^0.$$

□

In accordance with the general pattern explained in [8], the curvature of the quotient connection ∇ on $\mathcal{L}_1 \rightarrow \mathcal{B}^* \times \Sigma$ is given by

$$\begin{aligned} F_\nabla((0, 0, u), (0, 0, v)) &= F_A(u, v), \\ F_\nabla((a, \phi, 0), (0, 0, v)) &= \langle a, v \rangle, \\ F_\nabla((a_1, \phi_1, 0), (a_2, \phi_2, 0)) &= 2iG_\psi(d^*ib - \text{Im}\langle\psi, \chi\rangle). \end{aligned} \tag{2.20}$$

Here (a_1, ϕ_1) and (a_2, ϕ_2) are vector fields on \mathcal{B}^* which are *horizontal* with respect to Γ ; their Lie bracket is (b, χ) .

Note. *From now on, the (a_j, ϕ_j) are assumed horizontal.*

Lemma 2.3.8. *We have $d^*(ib) - \text{Im}\langle\psi, \chi\rangle = -\text{Im}\langle\phi_1, \phi_2\rangle$.*

Proof. Denote the pair $(A + ta_1, \psi + t\phi_1)$ by c_t . Then, at (A, ψ) ,

$$\begin{aligned} b &= \frac{1}{t}(a_2(c_0) - a_2(c_t)) + o(t), \\ \chi &= \frac{1}{t}(\phi_2(c_0) - \phi_2(c_t)) + o(t) \end{aligned}$$

as $t \rightarrow 0$. But at c_t , $d^*ia_2 = \text{Im}\langle\psi + t\phi_1, \phi_2(c_t)\rangle$, and from this one obtains

$$d^*(ib) - \text{Im}\langle\psi, \chi\rangle = -\lim_{t \rightarrow 0} \text{Im}\langle\phi_1(c_0), \phi_2(c_t)\rangle = -\text{Im}\langle\phi_1(c_0), \phi_2(c_0)\rangle.$$

□

2-forms as curvature integrals

We are now in a position to write down closed 2-forms representing $c_1(\mathcal{L}_1)/e_0$ and for $c_1(\mathcal{L}_1)^2/e_2$ in de Rham cohomology. The class $c_1(\mathcal{L}_1)$ has the Chern-Weil representative $iF_\nabla/2\pi$, so

$$c_1(\mathcal{L}_1)/e_0 = \left[\frac{1}{2\pi} \int_\Sigma iF_\nabla \wedge \omega_0 \right], \quad \text{where } \int_\Sigma \omega_0 = 1. \tag{2.21}$$

Explicitly, this representative for $c_1(\mathcal{L}_1)/e_0$ is the 2-form

$$((a_1, \phi_1), (a_2, \phi_2)) \mapsto \frac{1}{\pi[\omega]} \int_{\Sigma} G_{\psi}(\text{Im}\langle \phi_1, \phi_2 \rangle) \omega. \quad (2.22)$$

Similarly,

$$c_1(\mathcal{L}_1)^2/e_2 = \left[\frac{1}{4\pi^2} \int_{\Sigma} iF_{\nabla} \wedge iF_{\nabla} \right]. \quad (2.23)$$

This integral involves the product of the first and third curvature terms, and the square of the second term. So $c_1(\mathcal{L}_1)^2/e_2$ has the representative

$$((a_1, \phi_1), (a_2, \phi_2)) \mapsto \frac{1}{\pi^2} \int_{\Sigma} G_{\psi}(\text{Im}\langle \phi_1, \phi_2 \rangle) iF_A - \frac{1}{2\pi^2} \int_{\Sigma} ia_1 \wedge ia_2. \quad (2.24)$$

Notice the appearance of an expression familiar from (2.4) as the second term.

At this point we impose the moment map equation, restricting these forms and classes to the locus where $\mu(A, \psi) = \tau$. On that locus, the class

$$4\pi^2(\tilde{\theta} - r\tilde{\eta}) + 2\pi[\tau\omega]\tilde{\eta} = 2\pi(-\pi c_1(\mathcal{L}_1)^2/e_2 + [\tau\omega]c_1(\mathcal{L}_1)/e_0) \quad (2.25)$$

is represented by the form

$$\begin{aligned} & \int_{\Sigma} ia_1 \wedge ia_2 + 2 \int_{\Sigma} G_{\psi}(\text{Im}\langle \phi_1, \phi_2 \rangle) (\tau\omega - iF_A) \\ &= \int_{\Sigma} ia_1 \wedge ia_2 + \int_{\Sigma} G_{\psi}(\text{Im}\langle \phi_1, \phi_2 \rangle) |\psi|^2 \omega \\ &= \int_{\Sigma} ia_1 \wedge ia_2 + \int_{\Sigma} \text{Im}\langle \phi_1, \phi_2 \rangle \omega \\ &= \sigma((a_1, \phi_1), (a_2, \phi_2)). \end{aligned} \quad (2.26)$$

(Recall that σ is our standard Kähler form on \mathcal{A}^*). The penultimate equality uses the observation that, because the Laplacian of a function f has mean value zero,

$$\int_{\Sigma} f\omega = \int_{\Sigma} \Delta_{\psi} G_{\psi}(f) \omega = \int_{\Sigma} |\psi|^2 G_{\psi}(f) \omega.$$

Proof of Theorem 2.3.3. What we have just found is that the class $2\pi([\tau\omega]\tilde{\eta} + 2\pi(\tilde{\theta} - r\tilde{\eta}))$ on \mathcal{B}^* , restricted to $\mu^{-1}(\tau)/\mathcal{G}$, is equal to $[\sigma_{\tau}]$. Restricting further to the vortex moduli space, we find that the class of our preferred Kähler form is

$$2\pi([\tau\omega]\tilde{\eta} + 2\pi(\tilde{\theta} - r\tilde{\eta}))|_{\text{Vor}(L, \tau)} \in H^2(\text{Vor}(L, \tau); \mathbb{R}).$$

Hence, pulling back by Z , we find that the class of our Kähler form on $\text{Sym}^r(\Sigma)$ is $2\pi([\tau\omega]\eta + 2\pi(\theta - r\eta))$, which is the formula we have been working towards. \square

Remark 2.3.9. *The Duistermaat-Heckman formula [10] for the variation of cohomology of symplectic quotients gives another proof that the cohomology class $[\sigma_\tau]$ varies linearly with τ —provided that τ is a constant function—and computes the slope.¹*

Suppose that one has a Hamiltonian S^1 -action on (M, ω) , with moment map $\mu: M \rightarrow \mathfrak{t}^*$. Here $\mathfrak{t} = \text{Lie}(S^1)$. Identify \mathfrak{t}^* with \mathbb{R} so that the lattice dual to $\exp^{-1}(1) \subset \mathfrak{t}$ corresponds to $\mathbb{Z} \subset \mathbb{R}$. Suppose that μ is proper, and that its restriction to $\mu^{-1}(\ell)$ is submersive, for some open interval $\ell \subset \mathbb{R}$. The family of symplectic quotients $(M_t, \omega_t)_{t \in \ell}$ is then a trivial fibre bundle, and a trivialisation gives an identification of the cohomology of M_t with that of a fixed fibre M_s . The identification is canonical, hence $\{[\omega_t]\}_{t \in \ell}$ can be considered as a family of classes on M_s . Suppose that S^1 acts freely on $\mu^{-1}(s)$, so that $\mu^{-1}(s) \rightarrow M_s$ is a principal circle-bundle, with Chern class $c \in H^2(M_s; \mathbb{R})$. The Duistermaat-Heckman formula says that

$$\frac{d}{dt}[\omega_t] = 2\pi c. \quad (2.27)$$

We apply this with $M = \bigcup_{\tau \in \mathbb{R}} \mathcal{U}_\Sigma(L, \tau) / \mathcal{G}_0$, where \mathcal{G}_0 is the based gauge group $\{u: \Sigma \rightarrow U(1) : u(x) = 1\}$, $x \in \Sigma$ an arbitrary basepoint, and $\tau \in \mathbb{R}$ stands for a constant function on Σ . The circle acts by constant gauge transformations. We take $\ell = (2\pi r / \int \omega, \infty)$; the Chern class c of $M_\tau \rightarrow \text{Vor}(L, \tau)$ is η . Formula 2.27 gives

$$\frac{d}{d\tau}[\sigma_\tau] = 2\pi \left(\int_\Sigma \omega \right) \eta,$$

which is consistent with our result. One can formally recover the constant term $4\pi^2(\theta - r\eta)$ by specialising to the degenerate parameter $\tau = 2\pi r / \int \omega$ (for which the formula $[\sigma_\tau] = 4\pi^2\theta$ is easily verified); however, justifying this formal manipulation would need further thought.

Since Duistermaat and Heckman's proof identifies the variation in the symplectic forms with the curvature of a connection on $\mu^{-1}(s) \rightarrow M_s$, the two methods are not very different.

2.3.3 Families of vortex moduli spaces

Construction of the vortex fibration

(a) Suppose that $X \rightarrow S$ is a smooth, proper fibre bundle of relative dimension 2, where X and S are connected and oriented. Let $L \rightarrow X$ be a principal $U(1)$ -bundle, and assume that $L|_{X_s} \rightarrow X_s$ has degree $r > 0$.

We can consider $X \rightarrow S$ as a *fibration*; the typical fibre is some closed, oriented surface Σ , and the structure group $\text{Diff}^+(\Sigma)$. Putting $P = L|_{X_s}$, we can consider the composite map $L \rightarrow X \rightarrow S$ as a fibration with typical fibre P and structure group $\text{Diff}_P^+(\Sigma)$. The latter is the group of pairs (\tilde{g}, g) , where $\tilde{g} \in \text{Aut}(P)$ is an automorphism covering $g \in \text{Diff}^+(\Sigma)$, so it is an extension of $\text{Diff}^+(\Sigma)$ by the gauge group.

¹This was pointed out to me by Michael Thaddeus.

There are natural left actions of $\text{Diff}_P^+(\Sigma)$ on the space of connections $\mathcal{C}(P)$ and on the space of sections $\Omega_\Sigma^0(P)$. These arise through the covariance of connections and of sections; representing a connection by its 1-form $A \in \Omega_P^1$, we have

$$\tilde{g}.A = \tilde{g}^{-1*}A; \quad \tilde{g}.\psi = \tilde{g} \circ \psi \circ g^{-1}.$$

One can then form the associated fibrations

$$L \times_{\text{Diff}_P^+(\Sigma)} \mathcal{C}(P) \rightarrow S, \quad L \times_{\text{Diff}_P^+(\Sigma)} \Omega_\Sigma^0(P) \rightarrow S,$$

with structure group $\text{Diff}_P^+(\Sigma)$. The underlying fibre bundles can be thought of as the bundles of connections (resp. sections) along the fibres of $X \rightarrow S$:

$$\begin{aligned} L \times_{\text{Diff}_P^+(\Sigma)} \mathcal{C}(P) &\cong \{(s, A) : s \in S, A \in \mathcal{C}_{X_s}(L_s)\}, \\ L \times_{\text{Diff}_P^+(\Sigma)} \Omega_\Sigma^0(P) &\cong \{(s, \psi) : s \in S, \psi \in \Omega_{X_s}^0(L_s)\}. \end{aligned}$$

The first of these has the special property that it is a *symplectic* fibration: its structure group is reduced to $\text{Aut}(\mathcal{C}(P), dp)$ (cf. Equation (2.5)).

Other fibrations can be derived from these basic ones. The space

$$\mathcal{A}(P) = \mathcal{C}(P) \times \Omega_\Sigma^0(P \times_{U(1)} \mathbb{C}),$$

comprising pairs (A, ψ) where ψ is a section of the line bundle associated with P , is also a $\text{Diff}_P^+(\Sigma)$ -space (the action is the diagonal one), and so is

$$\mathcal{B}(P) = \mathcal{A}(P)/\mathcal{G}$$

(because \mathcal{G} acts on $\mathcal{A}(P)$ as a subgroup of $\text{Diff}_P^+(\Sigma)$). The associated fibrations are

$$\begin{aligned} \mathcal{A}_{X/S}(L) &:= L \times_{\text{Diff}_P^+(\Sigma)} \mathcal{A}(P), \\ \mathcal{B}_{X/S}(L) &:= L \times_{\text{Diff}_P^+(\Sigma)} \mathcal{B}(P). \end{aligned}$$

(b) Suppose now that $X \rightarrow S$ is itself a symplectic fibration, i.e. that its structure group is reduced to $\text{Aut}(\Sigma, \omega)$ for some area form ω . Then the structure group $L \rightarrow S$ is reduced to $\text{Aut}_P(\Sigma, \omega)$, the group of pairs (\tilde{f}, f) with $f^*\omega = \omega$, and $\mathcal{A}_{X/S}(L) \rightarrow S$ is again a symplectic fibration. Note that $P \times_{U(1)} \mathbb{C}$ is a *hermitian* line bundle, so our formula for the symplectic form on $\mathcal{A}(P)$ makes sense.

The moment map² $\mu: \mathcal{A}(P) \rightarrow \Omega_\Sigma^2$, $(A, \psi) \mapsto iF_A + |\psi|^2\omega/2$, generalises to a bundle-map

$$\mu_{X/S}: \mathcal{A}_{X/S}(L) \rightarrow L \times_{\text{Diff}_P^+(\Sigma)} \Omega_\Sigma^2$$

²At this point we choose to use 2-forms (not 0-forms) as values for the moment map.

over S . Since $X \rightarrow S$ is a symplectic fibration, the bundle $L \times_{\text{Diff}_P^+(\Sigma)} \Omega_\Sigma^2$ of fibrewise 2-forms has a preferred section ω , giving rise to the sub-bundles

$$\mu_{X/S}^{-1}(\tau\omega) \subset \mathcal{A}_{X/S}(L), \quad \pi(\mu_{X/S}^{-1}(\tau\omega)) \subset \mathcal{B}_{X/S}(L).$$

Here π is the projection $\mathcal{A}_{X/S}(L) \rightarrow \mathcal{B}_{X/S}(L)$, so $\pi(\mu_{X/S}^{-1}(\tau\omega)) \rightarrow S$ is an $\text{Aut}_P(\Sigma, \omega)$ -bundle with typical fibre $\mu^{-1}(\tau\omega)/\mathcal{G}$. It is also a symplectic fibration.

(c) We now impose a fibred version of the Cauchy-Riemann equation. This differs from what we have done so far in that it cannot be expressed in terms of associated bundles.

Let $\{j_s \in \mathcal{J}(X_s, \omega_s)\}_{s \in S}$ be a smooth family of complex structures, compatible with the symplectic forms. Then one has a space $\mathcal{M}_{X/S}(L) \subset \mathcal{A}_{X/S}(L)$ of triples (s, A, ψ) satisfying $\bar{\partial}_{j_s, A} \psi = 0$.

We can finally define the **vortex fibration** $\text{Vor}_{X/S}(L, \tau) \rightarrow S$. It has total space

$$\text{Vor}_{X/S}(L, \tau) = \pi(\mu_{X/S}^{-1}(\tau\omega) \cap \mathcal{M}_{X/S}(L)), \quad (2.28)$$

and it maps to S in the obvious way. The fibre over s can be identified with the vortex moduli space $\text{Vor}_{X_s}(L|X_s, \tau)$, and so with $\text{Sym}^r(X_s)$.

Lemma 2.3.10. *The space $\text{Vor}_{X/S}(L, \tau)$ has a structure of smooth manifold which makes the projection $p: \text{Vor}_{X/S}(L, \tau) \rightarrow S$ a smooth submersion, hence a fibre bundle.*

Proof. The linearisation of the defining equations for $\text{Vor}_{X/S}(L, \tau)$, and the fibrewise gauge-fixing condition, define an \mathbb{R} -linear operator $D_{(s, A, \psi)}$:

$$D_{(s, A, \psi)}(v, a, \phi) = D_{A, \psi}(a, \phi) + P(v), \quad v \in T_s S. \quad (2.29)$$

Here P is the 0th-order operator $P(v) = \frac{1}{2}i(d_A \psi) \circ \frac{\partial j}{\partial v}$. The operator $D_{(s, A, \psi)}$ is thus Fredholm, of real index $2r + \dim(S)$, and surjective (since $D_{(A, \psi)}$ is). The kernel of $D_{(s, A, \psi)}$ is the putative tangent space to $\text{Vor}_{X/S}(L, \tau)$ at (s, A, ψ) , and the projection $\pi: \ker D_{(s, A, \psi)} \rightarrow T_s S$ is putatively the derivative of p . Note that π is surjective, because its kernel is exactly $\ker D_{A, \psi}$, which we know has dimension $2r$.

Now the same standard elliptic theory which we sketched above here gives smoothness of the vortex fibration and of the map p . □

Line bundles and cohomology operations. Let $\widetilde{\mathcal{L}}_1 \rightarrow \mathcal{A}_{X/S}^* \times_S X$ be the pullback of the line bundle $L \rightarrow X$. It is an equivariant line bundle under the fibrewise gauge-action, and so descends to a line bundle

$$\mathcal{L}_1 \rightarrow \mathcal{B}_{X/S}^* \times_S X.$$

The universal divisor $\Delta^{\text{univ}} \subset \text{Sym}_S^r(X) \times_S X$ corresponds to a unique line bundle

$$\mathcal{L}_2 \rightarrow \text{Sym}_S^r(X) \times_S X.$$

Lemma 2.3.11. *There is a natural isomorphism*

$$(i \times 1)^* \mathcal{L}_1 \rightarrow Z^* \mathcal{L}_2,$$

where i is the inclusion of $\text{Vor}_{X/S}(L, \tau)$ in $\mathcal{B}_{X/S}^*$, and Z the natural isomorphism of $\text{Vor}_{X/S}(L, \tau)$ with $\text{Sym}_S^r X$.

Proof. The section

$$([A, \psi], x) \mapsto [\psi(x)]$$

of $(i \times 1)^* \mathcal{L}_1$ vanishes precisely along $Z^{-1}(\Delta)$.

□

Using these two line bundles one can construct operations

$$\begin{aligned} H^*(X) &\rightarrow H^{*+2k-2}(\mathcal{B}_{X/S}^*(L)), & c &\mapsto \tilde{c}^{(k)}, \\ H^*(X) &\rightarrow H^{*+2k-2}(\text{Sym}_S^r(X)), & c &\mapsto c^{(k)}. \end{aligned}$$

defined for arbitrary coefficient rings. The second of these earlier was discussed at the beginning of this chapter. Introduce the projections

$$\begin{aligned} \mathcal{B}_{X/S}^* &\xleftarrow{p_1} \mathcal{B}_{X/S}^* \times_S X && \xrightarrow{p_2} X, \\ \text{Sym}_S^r(X) &\xleftarrow{p_1} \text{Sym}_S^r(X) \times_S X && \xrightarrow{p_2} X, \end{aligned}$$

and set

$$\tilde{c}^{(k)} = p_{1!}(c_1(\mathcal{L}_1)^k \cup p_2^* c), \tag{2.30}$$

$$c^{(k)} = p_{1!}(c_1(\mathcal{L}_2)^k \cup p_2^* c). \tag{2.31}$$

Because of the relation between \mathcal{L}_1 and \mathcal{L}_2 , we have $i^* \tilde{c}^{(k)} = Z^* c^{(k)}$.

Associated fibrations as locally Hamiltonian fibrations

In the previous section, we constructed various associated fibrations within the category of *symplectic* fibrations—fibre bundles with symplectic forms on the fibres. Our present task is to refine these constructions to the category of *locally Hamiltonian* fibrations. The vortex fibration will then become locally Hamiltonian by restricting a closed 2-form defined on a larger space. The cleanest way that I have found to do this is to ‘reverse-engineer’ our cohomology calculation for the vortex moduli space. This goes as follows.

We need a fibrewise-equivariant connection $\hat{\nabla}$ on the bundle $\tilde{\mathcal{L}}_1 \rightarrow \mathcal{A}_{X/S}^* \times_S X$. To obtain one, choose a connection A_{ref} on $L \rightarrow X$. We define $\hat{\nabla}$ to be the unique connection which restricts to the natural one (2.18) on each fibre over S , and which is given by A_{ref} on $T^h X$.

In conjunction with the fibrewise gauge-fixing condition, $\hat{\nabla}$ defines a quotient connection ∇ on $\mathcal{L}_1 \rightarrow \mathcal{B}_{X/S}^* \times_S X$.

We now define a 2-form $\tilde{f}(\tau\Omega, L)$ on $\mathcal{B}_{X/S}^*$:

$$\tilde{f}(\tau\Omega, L) = 2\pi \int_{X/S} iF_{\nabla} \wedge (\tau\Omega - \frac{\pi}{2} iF_{\nabla}) \in Z_{\mathcal{B}_{X/S}^*}^2. \quad (2.32)$$

This is a closed form because F_{∇} and Ω are, and the fibre-integral of a closed form is closed.

Bearing in mind that integration along the fibre corresponds to the cohomology push-forward, we can read off the cohomology class of $\tilde{f}(\tau\Omega, L)$:

$$[\tilde{f}(\tau\Omega, L)] = 2\pi([\tau\Omega]^{(1)} - \pi\tilde{1}^{(2)}).$$

Forming $\tilde{f}(\tau\Omega, L)$ is obviously compatible with restricting the base S . Thus, by our earlier calculations, the form $i^*\tilde{f}(\tau\Omega, L)$ on the vortex bundle restricts to the preferred Kähler form on each fibre.

We define $f(\Omega, \tau, L)$ to be the restriction of $i^*\tilde{f}(\tau\Omega, L)$ to the vortex bundle. Thus

$$[f(\tau\Omega, L)] = 2\pi([\tau\Omega]^{(1)} - \pi 1^{(2)}).$$

The following theorem is now an immediate consequence of what we have done.

Theorem 2.3.12. *Let (X, π, Ω) be a proper, locally Hamiltonian surface-fibration over a manifold S , and r a positive integer. Choose*

- *an Ω -positive complex structure j on $T^v X = \ker D\pi$;*
- *a hermitian line bundle $(L, |\cdot|)$ over X such that $L|_{X_s}$ has degree r for each $s \in S$, and an affine connection A_{ref} on L ;*
- *a real parameter τ . We require $\tau > 2\pi r/a$, where a is the symplectic area of a fibre.*

There is a procedure which associates to these data a locally Hamiltonian structure $f(\Omega, \tau, L)$ on the relative symmetric product $\text{Sym}_S^r(X; j)$ (that is, on $\text{Sym}_S^r(X)$ with its smooth structure induced by j). This procedure is compatible with restriction of the base S . The form $f(\Omega, \tau, L)$ restricts to a Kähler form on each fibre $\text{Sym}^r(X_s)$. Its cohomology class is

$$[f(\Omega, \tau, L)] = 2\pi \left(\tau[\Omega]^{(1)} - \pi 1^{(2)} \right) \in H^2(\text{Sym}_S^r(X); \mathbb{R}).$$

It is interesting that the class $[f(\Omega, \tau, L)]$ does not depend on L .

Suppose we wish to make $\text{Sym}_S^r(X)$ into a locally Hamiltonian fibration starting from the fibre bundle $X \rightarrow S$ alone. To do so, the *essential* choices we must make are those of the cohomology class $[\Omega] \in H^2(X; \mathbb{R})$ and the parameter τ , up to rescalings preserving the product $\tau[\Omega]$.³ This can be made precise as follows.

Corollary 2.3.13. *Fix a proper surface-bundle $\pi: X \rightarrow S$. Choose two sets of data*

$$(\Omega_0, j_0, L_0, |\cdot|_0, A_{\text{ref},0}, \tau_0), \quad (\Omega_1, j_1, L_1, |\cdot|_1, A_{\text{ref},1}, \tau_1)$$

as above, and suppose that $[\tau_0\Omega_0] = [\tau_1\Omega_1] \in H^2(X; \mathbb{R})$. Then the locally Hamiltonian fibrations

$$(\text{Sym}_S^r(X; j_0), \pi_r, f(\Omega, L_0, \tau_0)), \quad (\text{Sym}_S^r(X; j_1), \pi_r, f(\Omega, L_1, \tau))$$

are equivalent.

Proof. Because $\tau_0\Omega_0$ and $\tau_1\Omega_1$ represent the same cohomology class, the locally Hamiltonian fibrations $(X, \pi, \tau_0\Omega_0)$ and $(X, \pi, \tau_1\Omega_1)$ are isotopic: an isotopy is given by the form $\tau_0\Omega_0 + d(t\beta) \in Z^2(X \times [0, 1])$, where $\tau_1\Omega_1 - \tau_0\Omega_0 = d\beta$. This restricts to the slice $X \times \{t\}$ as $(1-t)\tau_0\Omega_0 + t\tau_1\Omega_1$, and hence is positive on the fibres $X_{s,t}$ of $X \times [0, 1] \rightarrow S \times [0, 1]$. We can give $X \times [0, 1]$ a vertical complex structure J by choosing a path j_t between the given ones.

In the case that $L_1 = L_0$, there is a hermitian line bundle $(L, |\cdot|)$ with connection over $X \times [0, 1]$ which restricts to $(L_i, |\cdot|_i)$ on the ends. The form

$$f(\tau_0\Omega_0 + d(t\beta), L, 1)$$

on $\text{Sym}_{S \times [0,1]}^r(X \times [0, 1]; J)$ restricts on the ends to $f(\tau_i\Omega_i, L_i, 1)$, and it follows that $(\text{Sym}_S^r(X; j_0), \pi_r, f(\Omega, L_0, \tau_0))$ is equivalent to $(\text{Sym}_S^r(X; j_1), \pi_r, f(\Omega, L_1, \tau_1))$.

It remains to show that changing the line bundle does not affect things, and for this we may assume that $\Omega_0 = \Omega_1$ (write Ω for this single form) and $j_0 = j_1$. By the previous corollary, we can write $f(\Omega, L_0, \tau) - f(\Omega, L_1, \tau) = d\gamma$. Then, since $f(\Omega, L_0, \tau)$ and $f(\Omega, L_1, \tau)$ are both Kähler, the form $f(\Omega, L_0, \tau) + d(t\gamma)$ on $\text{Sym}_S^r(X) \times [0, 1]$ gives an isotopy between them.

□

Explicit formulae. One can give explicit formulae for $f(\Omega, \tau, L)$ in terms of local coordinates on S . These will not be needed elsewhere in the thesis, but it is interesting to see how $f(\Omega, \tau, L)$ reflects the geometry of the original locally Hamiltonian fibration.

³When the base is S^1 , this is the choice which may affect Floer homology.

Consider a trivial bundle $\Sigma \times U \rightarrow U$, where U is a ball about $0 \in \mathbb{R}^n$, with coordinates x_i on U . We work in a symplectic trivialisation, writing the closed 2-form Ω in the form of (2.3):

$$\Omega = \omega + \sum_i \eta_x^i \wedge dx_i + \sum_{i < j} \zeta_x^{ij} dx_i \wedge dx_j. \quad (2.33)$$

Thus the conclusions of Lemma (2.2.2) hold.

The tangent space to $\mathcal{B}_{X/S}^* \times_S X$ at a point $([A, \Psi], x)$ is $T_{[A, \psi]} \mathcal{B}^* \oplus T_x^v X \oplus T_x^h X$, and there is a corresponding decomposition of the curvature of ∇ which we write as

$$F_\nabla = F_{\mathcal{B}\mathcal{B}} + F_{\mathcal{B}v} + F_{vv} + F_{hh} + F_{\mathcal{B}h} + F_{vh}.$$

Examining the definition of ∇ one sees that $F_{\mathcal{B}h}$ and F_{vh} are zero, and

$$F_{hh}(u_1, u_2) = F_{A_{\text{ref}}}(u_1, u_2), \quad u_1, u_2 \in T_x^h X.$$

(Recall that A_{ref} is a chosen connection on L .) We write (2.33) schematically as $\Omega = \Omega_{vv} + \Omega_{vh} + \Omega_{hh}$, and $\tilde{f}(\tau\Omega, L)$ as

$$\tilde{f}(\tau\Omega, L) = \tilde{f}(\tau\Omega, L)_{\mathcal{B}\mathcal{B}} + \tilde{f}(\tau\Omega, L)_{\mathcal{B}h} + \tilde{f}(\tau\Omega, L)_{hh}.$$

We can represent forms on $\mathcal{B}_{X/S}^* \times_S X$ as forms on $\mathcal{A}_{X/S}^* \times_S X$ which are invariant under the fibrewise gauge-actions. We have

$$\begin{aligned} \tilde{f}(\tau\Omega, L)_{\mathcal{B}\mathcal{B}} &= \int_\Sigma i a_1 \wedge i a_2 + 2 \int_\Sigma G_\psi(\text{Im}\langle \phi_1, \phi_2 \rangle)(\tau\omega - iF_A) \\ \tilde{f}(\tau\Omega, L)_{\mathcal{B}h} &= 2\pi \int_\Sigma iF_{\mathcal{B}v} \wedge \tau\Omega_{vh} = 2\pi \sum_j \left(\int_\Sigma \tau i a \wedge \eta^j \right) dx_j, \\ \tilde{f}(\tau\Omega, L)_{hh} &= 2\pi \int_\Sigma (iF_{vv} \wedge \tau\Omega_{hh} + iF_{hh} \wedge \tau\Omega_{vv} - \pi iF_{hh} \wedge iF_{vv}) \\ &= 2\pi \sum_{j < k} \left(\int_\Sigma \tau iF_A \zeta^{jk} \right) dx_j dx_k + 2\pi^2 iF_{A_{\text{ref}}} \int_\Sigma (\tau\omega - iF_A). \end{aligned}$$

The restriction to the locus where the moment map equation is satisfied is then

$$\sigma + 2\pi \sum_j \left(\int_\Sigma \tau i a \wedge \eta^j \right) dx_j + 2\pi \sum_{j < k} \left(\int_\Sigma \tau \left(\tau - \frac{|\psi|^2}{2} \right) \omega \zeta^{jk} \right) dx_j dx_k + \pi^2 iF_{A_{\text{ref}}} \int_\Sigma |\psi|^2 \omega. \quad (2.34)$$

at a point $([A, \psi], s)$. When the base is S^1 , only the first two terms are non-zero. One can verify directly that, in this case, the Hamiltonian connection is exactly the same as Salamon's [38].

Chapter 3

Invariants from pseudoholomorphic sections

This chapter gives an exposition of those parts of the theory of pseudoholomorphic curves which we will need to construct a field theory for near-symplectic broken fibrations. The analytic foundations on which these rest are relatively simple; we make no use of virtual cycles or Kuranishi structures. As a result, all our invariants take integer (not merely rational) values.

We describe invariants of two kinds (and eventually combine them). The first are ‘open Gromov invariants’, defined from moduli spaces of pseudoholomorphic sections of locally Hamiltonian fibrations over the disk (or, more generally, over surfaces with boundary), with Lagrangian boundary conditions. The ‘Lagrangian matching invariants’ for broken fibrations on closed 4-manifolds fit into this framework.

After that, we discuss Floer homology for symplectic automorphisms, stressing the field-theoretic character of the theory and its relation to the open Gromov invariants. In doing so we risk overplaying formal aspects; on the other hand, it is precisely these which are relevant to our programme of studying fibred 3- and 4-manifolds via pseudoholomorphic sections of relative symmetric products.

Using the 2-forms constructed in Chapter 2 we obtain invariants for mapping classes of surfaces, generalising those of Seidel [45] which already distinguish the trivial class.

Because of our avoidance of the virtual class machinery, we will not obtain invariants of r -fold symmetric products for every r (there are too many bubbles when r is a little less than $g - 1$). Appealing to virtual techniques does yield Floer homology over \mathbb{Q} for every r ; however, to obtain a Lagrangian matching invariant one would need to give serious thought to the possible failure (due to boundary bubbling) of the fundamental chain to be a cycle.

Bibliographic note. The idea of presenting symplectic Floer homology as a topological field theory for surfaces coupled to (non-trivial) symplectic fibre bundles appears in Seidel’s thesis [42]. A more general open-closed field theory is described in [44], while the role of open Gromov invariants—as distinct from relative invariants with values in Lagrangian Floer homology groups—is brought out in [46]. My exposition owes much to these papers.

Lalonde [23] has constructed a slightly different field theory for Hamiltonian fibrations over Riemann surfaces by allowing variations in the complex structure on the base. It seems likely that Lalonde’s work generalises in a straightforward way to *locally* Hamiltonian fibrations.

3.1 Pseudoholomorphic sections and Lagrangian boundary conditions

Definition 3.1.1. A **simple quadratic fibration** $(X^{2n+2}, \pi, \Omega, J_0, j_0)$ consists of a surjective, proper map $\pi: X \rightarrow S$ from an oriented manifold to an oriented surface; a closed 2-form $\Omega \in Z_X^2$; and the germ (J_0, j_0) of a pair of integrable complex structures near X^{crit} and $\pi(X^{\text{crit}})$ such that $D\pi \circ J_0 = j_0 \circ D\pi$. The requirements are as follows.

- the critical set $X^{\text{crit}} \subset X$ must be a $(2n-2)$ -dimensional submanifold, hence a complex submanifold with respect to J .
- On the complex normal bundle $N = N_{X^{\text{crit}}/X} \rightarrow X^{\text{crit}}$, the Hessian $D_x^2\pi: N_x \otimes N_x \rightarrow \mathbb{C}$ is non-degenerate as complex quadratic form. (It is automatic that $D_x\pi = 0$ for any $x \in X^{\text{crit}}$, since π is holomorphic there.)
- The form Ω must be positive on the oriented vector space $\ker(D_x\pi) \cong \mathbb{R}^{2n}$, for all regular points x .
- On the neighbourhood of X^{crit} where J is defined, Ω is a Kähler form.

We allow X and S to have boundary, provided that $\partial X = \pi^{-1}(\partial S)$ and $\pi(X^{\text{crit}}) \subset \text{int}(S)$. We write S^* for the set of regular values, $S \setminus \pi(X^{\text{crit}})$, and X^* for $\pi^{-1}(S^*)$.

Simple quadratic fibrations may be thought of as objects of ‘complex Morse-Bott theory’. The key examples for us will be the relative Hilbert scheme $\text{Hilb}_S^r(X)$ of r points on a four-dimensional Lefschetz fibration $X \rightarrow S$, and also the compactified Picard family $\mathcal{P}_S^r(X)$.

Definition 3.1.2. A **Lagrangian boundary condition** for the simple quadratic fibration $(X^{2n+2}, \pi, \Omega, J_0, j_0)$ over compact surface-with-boundary S is a submanifold $Q^{n+1} \subset X|\partial S$ such that (i) $\pi|Q: Q \rightarrow \partial S$ is a proper, surjective submersion, and (ii) $\Omega|Q = 0$.

The 2-form Ω need not be symplectic. However, if one chooses a positive area form $\beta \in \Omega_S^2$ then, for all $c \gg 0$, $\Omega + c\pi^*\beta$ is a symplectic form, and Q a Lagrangian submanifold. For any homotopy class h of sections $(S, \partial S) \rightarrow (X, Q)$, there is a moduli space \mathcal{M}_h of pseudoholomorphic sections representing h , which under certain hypotheses is smooth and compact, and from this one can extract invariants. Our Lagrangian matching invariant will be of this kind.

Remark 3.1.3. *It may seem arbitrary that we keep (Ω, J_0, j_0) as part of the data defining a simple quadratic fibration, but not a globally-defined pair (J, j) . The rationale (besides consistency with Seidel's work on Lefschetz fibrations) is that in Floer theory it is safe to discard choices from contractible spaces, but one should keep track of choices from merely path-connected ones \mathcal{F} , bearing in mind that a homotopically non-trivial map $S^n \rightarrow \mathcal{F}$ may give rise to a potentially interesting Floer-theoretic invariant.*

Almost complex structures adapted to simple quadratic fibrations

Denote by $\mathcal{J}(E)$ the space of all orientation-compatible, C^∞ complex structures on the oriented real vector bundle $E \rightarrow M$, with its C^∞ topology. If (E, σ) is a symplectic vector bundle then the subspace $\mathcal{J}(E, \sigma) \subset \mathcal{J}(E)$ of compatible complex structures—those $J \in \mathcal{J}(E)$ such that the formula $g_x(v_1, v_2) := \sigma(v_1, Jv_2)$ defines a metric on the vector bundle E —is contractible.

Definition 3.1.4. Let $(X, \pi, \Omega, J_0, j_0)$ be a simple quadratic fibration over S . A pair $(J, j) \in \mathcal{J}(TX) \times \mathcal{J}(TS)$ is **adapted** to the fibration if

1. π is (J, j) -holomorphic, i.e. $D\pi \circ J = j \circ D\pi$;
2. J is compatible with Ω on the smooth fibres, i.e. $J|_{T^\vee X^*} \in \mathcal{J}(T^\vee X^*, \Omega|_{T^\vee X^*})$;
3. J extends the germ J_0 , and j extends j_0 .

Trying not to make the notation *too* cumbersome, we write $\mathcal{J}(X, \pi)$ for the space of adapted pairs, and $\mathcal{J}(X, \pi, j) = \{J : (J, j) \in \mathcal{J}(X, \pi)\}$. With respect to the splitting $TX^* = T^\vee X^* \oplus \pi^*TS^*$ defined by the symplectic connection (or, in fact, any other connection), a pair $(J, j) \in \mathcal{J}(X, \pi)$ has a block decomposition over X^* of shape

$$J = \begin{pmatrix} J^{vv} & J^{vh} \\ 0 & j \end{pmatrix}, \quad J^{vv} \circ J^{vh} + J^{vh} \circ j = 0. \quad (3.1)$$

With J^{vv} and j fixed, J^{vh} is just a \mathbb{C} -antilinear homomorphism:

$$J^{vh} \in \Gamma \operatorname{Hom}^{0,1}(T^h X^*, T^\vee X^*).$$

Lemma 3.1.5. *The spaces $\mathcal{J}(X, \pi, j)$ are contractible.*

Proof. To see that $\mathcal{J}(X, \pi, j) \neq \emptyset$, fix a closed neighbourhood U of $\text{crit}(\pi)$ on which J_0 can be defined; let $U^* = U \setminus X^{\text{crit}}$. One can extend $J_0|_{T^\vee U^*}$ to an element $J^{vv} \in \mathcal{J}(T^\vee X^*, \Omega|_{T^\vee X^*})$. Over U^* , $J^{vv} \oplus j$ differs from J_0 by an antilinear homomorphism, which can be extended to one defined over X^* . This gives rise to a complex structure J of the right sort.

Now consider the restriction map $r: \mathcal{J}(X, \pi, j) \rightarrow \mathcal{J}(T^\vee X^*, \Omega|_{T^\vee X^*})$. Its fibres are affine spaces modelled on the vector spaces $\Gamma_c \text{Hom}^{0,1}(T^h X^*, T^\vee X^*)$, where Γ_c means sections supported outside X^{crit} . The map r admits a section s , and $\mathcal{J}(X, \pi, j)$ deformation-retracts to the contractible space $\text{im}(s)$, hence is contractible. \square

An easy matrix calculation, linearising (3.1), shows that $\mathcal{J}(X, \pi)$ has formal tangent spaces $T_J \mathcal{J}(X, \pi, j)$ which fit into short exact sequences

$$0 \rightarrow \text{Hom}^{0,1}((\pi^* TS, j), (T^\vee X, J^{vv})) \rightarrow T_J \mathcal{J}(X, \pi, j) \rightarrow T_{J^{vv}} \mathcal{J}(T^\vee X) \rightarrow 0.$$

$\mathcal{J}(X, \pi, j)$ can indeed be made into a smooth Fréchet manifold, which fibres smoothly over $\mathcal{J}(T^\vee X)$; the fibres are spaces of sections of affine spaces.

Note that, for any rapidly-decreasing sequence ϵ of positive reals, there is a dense Banach submanifold $\mathcal{J}(X, \pi, j)_\epsilon \subset \mathcal{J}(X, \pi, j)$; its tangent vectors are bundle maps with finite ‘ C_ϵ^∞ -norm’ (see [40]).

3.1.1 Transversality

Write $\mathcal{S}(X)$ for the space of C^∞ sections of π , and $\mathcal{S}(X, Q)$ for the subspace of sections which map ∂S into Q . Fix $(J, j) \in \mathcal{J}(X, \pi, \Omega)$. The moduli space of (j, J) -pseudoholomorphic sections with boundary in Q is the space

$$\mathcal{M}_{J,j}(X, Q) = \{u \in \mathcal{S}(X, Q) : J \circ (Du) = (Du) \circ j\} \subset \mathcal{S}(X, Q),$$

and it is this that we wish to analyse, under the assumption (in force from now on) that S is compact.

Remark 3.1.6. *Infinite-dimensional transversality theory only works for Banach manifolds. We will try to keep track of the relevant Banach spaces to the extent that we do not make false statements. In the first place, the space of sections $\mathcal{S}(X)$ should be enlarged so that it has tangent spaces $L_k^p(\gamma^* T^\vee X)$, where $kp > \dim(S)$; in the present case L_1^p will do, for any $p > 2$. The enlargement $\mathcal{S}_k^p(X)$ is the smallest set of sections closed under exponentiation (whenever defined) of such vector fields along smooth sections, with respect to some chosen fibre metric on $T^\vee X \rightarrow X$. This is a Banach manifold, and it is contained in the space of continuous sections.*

Transversality for sections

The first point to make is that a C^1 section of π *cannot intersect* X^{crit} . Thus the presence of critical points makes no difference to transversality theory for moduli spaces of sections (it does affect their compactifications).

Let $\mathcal{E}_{J,j} \rightarrow \mathcal{S}(X)$ be the natural infinite-rank vector bundle with fibres $(\mathcal{E}_{J,j})_u = \Gamma(\text{Hom}^{0,1}(TS, u^*T^\vee X))$. The space $\mathcal{M}_{J,j}(X)$ of pseudoholomorphic sections of π is the zero-set of the section

$$\bar{\partial}_{J,j} = \frac{1}{2}(D + J \circ D \circ j) \in \Gamma(\mathcal{E}_{J,j}).$$

Allowing J to vary within $\mathcal{J}(X, \pi, j)$ one gets a ‘universal’ vector bundle

$$\mathcal{E}_j \rightarrow \mathcal{S}(X) \times \mathcal{J}(X, \pi; j).$$

This has a section $\bar{\partial}_j$, whose zero-set is $\mathcal{M}_j(X) = \bigcup_J \mathcal{M}_{J,j}(X)$. Introduce a torsion-free connection ∇ on $T(X \setminus X^{\text{crit}})$ which extends a connection on the vertical subbundle. The linearisation of $\bar{\partial}_{J,j}$ at $u \in \mathcal{M}_{J,j}(X)$ is a linear map

$$D_u: \Gamma(u^*T^\vee X) \rightarrow (\mathcal{E}_{J,j})_u = \Gamma(\text{Hom}^{0,1}(TS, u^*T^\vee X)).$$

Explicitly,

$$D_u(v) = \frac{1}{2}(\mathcal{L}_{\tilde{v}}J) \circ Du \circ j \tag{3.2}$$

$$\begin{aligned} &= \frac{1}{2}(-\nabla_{J\bullet}v + J\nabla_{\bullet}v + \nabla_v J) \circ Du \circ j \\ &= \frac{1}{2}((u^*\nabla v) + J \circ (u^*\nabla)v \circ j) + \frac{1}{2}(\nabla_v J) \circ Du \circ j \\ &= (u^*\nabla)^{0,1}(v) + \frac{1}{2}(\nabla_v J) \circ Du \circ j, \end{aligned} \tag{3.3}$$

where \tilde{v} is an extension of v to a vertical vector field defined on a neighbourhood of $\text{im}(u)$. The linearisation of $\bar{\partial}_j$ is given by $D_{u,j}^{\text{univ}}: (v, Y) \mapsto D_u v + \frac{1}{2}Y \circ Du \circ j$.

We have not yet imposed the boundary condition. To do so, we should restrict $\mathcal{E}_{J,j}$ to $\mathcal{S}(X, Q)$, and consider its subbundle $\mathcal{E}_{J,j}^Q \rightarrow \mathcal{S}(X, Q)$ with fibres $\Gamma \text{Hom}_Q^{0,1}(TS, u^*T^\vee X)$: here $\text{Hom}_Q^{0,1}(TS, u^*T^\vee X)$ denotes the (j, J^{vv}) -antilinear homomorphisms $TS \rightarrow u^*T^\vee X$ which carry $T(\partial S)$ to $u^*T^\vee Q$. Then we have

$$D_u: \Gamma(u^*T^\vee X) \rightarrow (\mathcal{E}_{J,j}^Q)_u;$$

similar remarks apply to the universal version.

The linearised operators extend continuously to maps between the Banach completions of their domains and targets. Thus we have

$$D_u: L_1^p(u^*T^\vee X, (\partial u)^*T^\vee Q) \rightarrow L^p(\text{Hom}_Q^{0,1}(TS, u^*T^\vee X), \tag{3.4}$$

$$D_{u,j}^{\text{univ}}: L_1^p(u^*T^\vee X, (\partial u)^*T^\vee Q) \oplus C_\epsilon^\infty(T_J \mathcal{J}(X, \pi, j)) \rightarrow L^p(\text{Hom}_Q^{0,1}(TS, u^*T^\vee X). \tag{3.5}$$

Now the crucial point is that D_u is Fredholm (we will come back to this point in the next paragraph). It follows that so too is $D_{u,J}^{\text{univ}}$. The latter is also *surjective* (one makes use of a unique continuation theorem in proving this), whence $\mathcal{M}_j(X, Q)$ is a smooth Banach submanifold of $\mathcal{S}_1^p(X, Q) \times \mathcal{J}(X, \pi; j)_\epsilon$.

Write $\mathcal{J}^{\text{reg}}(X, \pi; j)$ for the space of **regular** almost complex structures: those with the property that D_u is onto for every $u \in \mathcal{M}_{J,j}(X, Q)$. Their importance is that when J is regular, $\mathcal{M}_{J,j}(X, Q)$ is a smooth manifold of local dimension $\text{ind } D_u$. The key transversality statement is that

$$\mathcal{J}^{\text{reg}}(X, \pi, j) \text{ is } C^\infty\text{-dense in } \mathcal{J}(X, \pi, j).$$

The standard way to prove it is to consider the projection $\text{pr}_2: \mathcal{M}_j(X, Q) \rightarrow \mathcal{J}(X, \pi; j)_\epsilon$. Its derivative at (u, J) has kernel $\ker(D_u)$ and closed image $\{Y : Y \circ Du \circ j \in \text{im}(D_u)\}$. The cokernel is thus identified with that of D_u . Since it has Fredholm derivative, the Sard-Smale theorem applies to pr_2 . What it says is that the regular values of pr_2 form a Baire-dense subset of $\mathcal{J}(X, \pi; j)_\epsilon$; but the regular values are just the regular almost complex structures (of class C_ϵ^∞). A Baire-dense subset of $\mathcal{J}(X, \pi; j)_\epsilon$ is still a dense subset of $\mathcal{J}(X, \pi; j)$, which gives the result.

Intersection with cycles in fibres. Mark a finite set of points $\{s_i\}_{i \in I}$ in S , manifolds Z_i , and smooth maps $\zeta_i: Z_i \rightarrow X_{s_i}$. There is an evaluation map

$$\text{ev}_I^J = \prod_{i \in I} \text{ev}_i: \mathcal{M}_{J,j}(X, Q) \rightarrow \prod_i X_{s_i}.$$

The arguments that establish the surjectivity of $D_{u,J}^{\text{univ}}$ extend easily to give a ‘transversality of evaluation’ lemma (see [46]) which says that, for any $J \in \mathcal{J}(X, \pi; j)$, and any neighbourhood U of $\{s_i : i \in I\}$, there exist arbitrarily small perturbations J_t of J , supported in $\pi^{-1}(U)$, such that $J_t \in \mathcal{J}^{\text{reg}}(X, \pi; j)$ and $\text{ev}_I^{J_t}$ is transverse to $\zeta = \prod_i \zeta_i$. Similar remarks apply to marked points $s'_k \in \partial S$ and maps $\zeta'_k: Z'_k \rightarrow Q_{s'_k}$.

Fredholm theory. The linearised operator $D_u: \Gamma(u^*T^\vee X) \rightarrow (\mathcal{E}_{J,j})_u$ does not make sense for arbitrary $u \in \mathcal{S}(X, Q)$ until one chooses a connection on $\mathcal{E}_{J,j}$. Then, instead of (3.3), one has

$$D_u v = \nabla^{0,1} v + a v \tag{3.6}$$

for some connection ∇ on the complex vector bundle $E = u^*T^\vee X$, and some bundle map $a: E \rightarrow \text{Hom}^{0,1}(TS, E)$. In other words, D_u is a *Cauchy-Riemann operator* on E . Over the boundary of S we have a real vector subbundle $F = (\partial u)^*T^\vee Q \subset E|_{\partial S}$ which is middle-dimensional and totally real. We have a linear space $L_1^p(E, F)$ of sections with boundary in

F , and an \mathbb{R} -linear operator

$$L: L_1^p(E, F) \rightarrow L^p(\text{Hom}^{0,1}((TS, T\partial S), (E, F))).$$

Data (E, F, L) of this sort constitute a **Cauchy-Riemann operator with totally real boundary condition**. The fundamental fact which we quote is that any such Cauchy-Riemann operator is Fredholm, with index given by the Riemann-Roch formula for surfaces with boundary. In the case at hand this says that

$$\text{ind}(D_u) = n\chi(S) + \mu_Q([u]), \quad (3.7)$$

where $\mu_Q: H_2(X, Q; \mathbb{Z}) \rightarrow \mathbb{Z}$ is the Maslov index homomorphism, which can be defined as follows. First, consider the subgroup K of classes $a \in H_2(X, Q; \mathbb{Z})$ with $\langle w_1(TQ), \partial a \rangle = 0$. Represent $a \in K$ by a map $u: (\Sigma, \partial\Sigma) \rightarrow (X, L)$ from an oriented surface. Trivialise $(\partial u)^* T^\vee Q$ (possible since $a \in K$), and extend this to a symplectic trivialisation t of $(\partial u)^* T^\vee X$. The Maslov index is twice the relative Chern number:

$$\mu_Q(a) = 2c_1(u^* TX; t)[\Sigma, \partial\Sigma].$$

This gives a well-defined homomorphism $K \rightarrow \mathbb{Z}$. For general $a \in H_2(X, Q; \mathbb{Z})$, put $\mu_Q(a) = \frac{1}{2}\mu_Q(2a)$; this is legitimate since $2a \in K$.

3.1.2 Compactness

The **action** of a section $u \in \mathcal{S}(X, Q)$ is

$$\mathcal{A}(u) = \int_S u^* \Omega. \quad (3.8)$$

Choosing a positive area form $\beta \in \Omega_S^2$ and a pair (J, j) is adapted to π , for all $c \gg 0$ the form $\Omega + c\pi^*\beta$ is J -positive. The symplectic area $\int_S u^*(\Omega + c\beta)$ is obviously equal to $\mathcal{A}(u) + c \text{vol}(S)$.

Gromov's compactness theorem (the relevant version is that of Ye [53]) says that any sequence u_1, u_2, \dots in $\mathcal{M}_{J,j}(X, Q)$ of bounded symplectic area—equivalently, bounded action—has a subsequence which converges, in Gromov's topology, to a pseudoholomorphic curve v with bubbles and boundary bubbles.

The domain of v is a compact, nodal complex curve with nodal boundary, (\bar{S}, \bar{j}) . It can be obtained from (S, j) by attaching trees of spheres at various interior points r_1, \dots, r_j , and trees of disks (attached to one another along their boundaries) at points $r'_1, \dots, r'_k \in \partial S$. The principal component of v is a pseudoholomorphic map $v^{\text{prin}}: S \rightarrow X$. Convergence in Gromov's topology means that on $S_0 := S \setminus \{r_1, \dots, r_j, r'_1, \dots, r'_k\}$ the sequence converges to v^{prin} ; on compact subsets this convergence is uniform in any C^k -topology. Near

$\{r_1, \dots, r_j, r'_1, \dots, r'_k\}$, convergence involves reparametrisations of the u_i , and we decline to give the details, though we remark that both the index and action behave well under Gromov limits.

Clearly $v^{\text{prin}}|_{S_0}$ is still a section, and it follows that the differentiable map v^{prin} is a section too. In particular, $\text{im}(v^{\text{prin}}) \cap X^{\text{crit}} = \emptyset$. The boundaries of the disks are of course mapped into Q . Each bubble or boundary-bubble component must be mapped by $\pi \circ v$ to a point (by the maximum principle, $\pi \circ v$ is constant or surjective on each component; by the definition of Gromov convergence, the latter can't arise by bubbling).

Hence to compactify $\{u \in \mathcal{M}_{J,j}(X, Q) : \mathcal{A}(u) \leq \lambda\}$ one need only consider curves whose components are (i) sections; (ii) bubbles in regular fibres; (iii) bubbles in singular fibres; (iv) boundary bubbles.

Fibred monotonicity. When S is closed, the moduli space of pseudoholomorphic sections always carries a rational fundamental cycle, at least in the virtual sense. In the open case this need not be so, because of two ‘anomalies’. It may be that there is a rational fundamental chain but that it is not a cycle, because of boundary bubbles in virtual codimension one; and it may be that the moduli space is not orientable, in which case the virtual methods fail. Various hypotheses can be imposed to prevent such difficulties from arising. Ours will be a simple one: fibred monotonicity. In our applications, which will involve symmetric products $\text{Sym}^r(\Sigma)$, we will be able to fulfil this hypothesis when $r \geq g(\Sigma)$.

It is useful to be able to include the case $r = g(\Sigma) - 1$ as well, but in that case the fibres are only ‘weakly monotone’. At the end of the chapter we will give a sketch of how the theory can be extended to that case using the methods of Hofer and Salamon [16].

Definition 3.1.7. (a) A closed Lagrangian submanifold L in a compact symplectic manifold (M, ω) is called **monotone** if there exists $c > 0$ such that

$$\int_D u^* \omega = c \mu(u)$$

for every smooth map $u: (\bar{D}, \partial \bar{D}) \rightarrow (M, L)$. **(b)** A simple quadratic fibration (X, π, Ω) , with Lagrangian boundary condition Q , is **fibre-monotone** if Q_s is monotone in $(X_s, \Omega|_{X_s})$ for each $s \in \partial S$.

For a Lagrangian submanifold $L \subset M$, let $\mu_{\min}(L) \geq 0$ denote the minimal Maslov index. Let $c_{\min}(M) \geq 0$ be the minimal Chern number of M . That is,

$$\text{im}(\mu_L: \pi_2(M, L) \rightarrow \mathbb{Z}) = \mu_{\min}(L)\mathbb{Z}, \quad \text{im}(c_1: \pi_2(M) \rightarrow \mathbb{Z}) = c_{\min}(M)\mathbb{Z}.$$

Since the Maslov index of a disk whose boundary is mapped to a point is twice its Chern number, $\mu_{\min}(L) \leq 2c_{\min}(M)$. For a Lagrangian boundary condition Q , let $\mu_{\min}(Q) = \gcd_{s \in \partial S} \mu_{\min}(Q_s)$.

Theorem 3.1.8. *Let Q be a Lagrangian boundary condition for the simple quadratic fibration (X^{2n+2}, π) over S , and fix $h \in \pi_0 \mathcal{S}(X, Q)$. Fix also finite sets of points $\{s_a\}_{a \in A}$ in S^* , $\{s'_b\}_{b \in B}$ in ∂S and smooth maps $\zeta_a: Z_a \rightarrow X_{s_a}$, $\zeta'_b: Z'_b \rightarrow X_{s'_b}$ from closed, oriented manifolds Z_a and Z'_b . Suppose that*

1. *the Lagrangian boundary condition $Q \subset \partial X$ is fibre-monotone;*
2. *the virtual dimension $d(h) := n\chi(S) + \mu_Q(h)$ satisfies*

$$d(h) - \sum_{a \in A} (2n - \dim(Z_a)) - \sum_{b \in B} (n - \dim(Z'_b)) - \mu_{\min}(Q) < 0; \quad (3.9)$$

3. *the normal bundle $N_{X^{\text{crit}}/X}$ is topologically trivial;*
4. *for any $(J, j) \in \mathcal{J}(X, \pi)$ and any holomorphic sphere $\beta: S^2 \rightarrow X$ with $\pi \circ \beta$ constant, $\beta^{-1}(X^{\text{crit}})$ is never a single point.*

Then for any extension of j_0 to $j \in \mathcal{J}(TS)$ and for a dense set of $J \in \mathcal{J}^{\text{reg}}(X, \pi, j)$, the moduli space $\mathcal{M}_{J,j}(X, Q)_h$ is a compact manifold of dimension $d(h) - \sum_a (2n - \dim(Z_a)) - \sum_b (n - \dim(Z'_b))$ when this number is non-negative, and empty when it is negative.

The conditions here are of course tailored to our applications; many variations are possible to deal with differing situations.

Proof. Write $Z = \prod_a Z_a$, and $\zeta = \prod_a \zeta_a: Z \rightarrow \prod_a X_{s_a}$. Similarly define Z' , ζ' . By the remarks on transversality of evaluation, there is a dense subset of $\mathcal{J}^{\text{reg}}(X, \pi, j)$ whose members J have the property that the evaluation map

$$\text{ev} = \text{ev}_A \times \text{ev}'_B: \mathcal{M}_{J,j}(X, Q) \rightarrow \prod_a X_{s_a} \times \prod_b X_{s'_b}$$

is transverse to $\zeta \times \zeta'$. This means that there is a smooth moduli space

$$\mathcal{M} = \mathcal{M}_{J,j}(X, Q) \times_{(\text{ev}, \zeta, \times \zeta')} (Z \times Z').$$

Any sequence in \mathcal{M} , for which the sections all lie in h , has a subsequence with a Gromov limit comprising a principal component u_{prin} , non-constant bubbles $\{u_i\}_{i \in I}$, and non-constant boundary bubbles $\{u'_k\}_{k \in K}$. These satisfy

$$d(h) = \text{ind}(u_{\text{prin}}) + 2 \sum_{i \in I} c_1(u_i) + \sum_{k \in K} \mu(u'_k). \quad (3.10)$$

By fibred monotonicity, the Maslov indices $\mu(u'_k)$ are positive multiples of $\mu_{\min}(Q)$, hence $\geq \mu_{\min}(Q)$. The contribution of each term $c_1(\beta_i)$ from a bubble in a regular fibre is also positive, at least $\mu_{\min}(Q)$. Let us temporarily grant that this is also true for bubbles

in singular fibres. Then if $I \cup K \neq \emptyset$ we will have $\text{ind}(u_{\text{prin}}) - \sum_a (2n - \dim(Z_a)) - \sum_b (n - \dim(Z'_b)) < 0$. But for regular almost complex structures, this number gives the dimension near u_{prin} of \mathcal{M} , which cannot be negative.

We now consider bubbles β in a singular fibre X_s . If $\beta^{-1}(X^{\text{crit}}) = \emptyset$ then β is homotopic to a sphere in a regular fibre. The same is true when $\beta^{-1}(X^{\text{crit}}) = S^2$, because of the triviality of the normal bundle. An argument due to Donaldson and Smith [9, Lemma A.11] shows that bubbles which intersect X^{crit} but do not map into it are also impossible: suppose that β is a component of a limit of pseudoholomorphic sections, $\beta(z) \in X^{\text{crit}}$; let $D_z \subset S$ be a small disk centred at z . Then $\beta(D_z)$ lies in one of the two branches of the normal crossing singularity of X_s at $\beta(z)$, and intersects the other branch positively. But $\beta|_{D_z}$ is the limit of local sections of π —maps which do not hit X^{crit} , and which have precisely one, transverse intersection with X_s . Hence $z = \beta^{-1}(X^{\text{crit}})$, which contradicts our hypotheses. \square

3.1.3 Orientations

Orientability of a Fredholm moduli space \mathcal{M} means triviality of the determinant line bundle $\text{Det} \rightarrow \mathcal{M}$. Recall that this is a line bundle with fibres $\Lambda^{\max} \ker(F_x) \otimes \Lambda^{\max}(\text{coker}(F_x))^\vee$, where F_x is the linear Fredholm operator associated with $x \in \mathcal{M}$.

In the case of the moduli space $\mathcal{M}_S(M, L)$ of pseudoholomorphic maps $(S, \partial S) \rightarrow (M, L)$ from a compact Riemann surface to an almost complex manifold M^{2n} with totally real submanifold L^n , de Silva [48] shows that, for a loop $\gamma: S^1 \rightarrow \mathcal{M}_S(M, L)$, one has

$$\langle w_1(\gamma^* \text{Det}), [S^1] \rangle = \langle (\partial\gamma)^* w_2(TL), [S^1 \times \partial S] \rangle, \quad (3.11)$$

where $\partial\gamma: S^1 \times \partial S \rightarrow L$ gives the boundary values of γ .

Fukaya, Oh, Ohta and Ono [12] show that one gets an orientation for $\mathcal{M}_S(M, L)$ by giving a ‘relative spin structure’, that is, an oriented vector bundle $\xi \rightarrow M$, with $w_2(\xi|L) = w_2(TL)$, and a spin structure on $\xi|L \oplus TL$. This is not quite enough for our purposes, because it does not allow non-orientable Lagrangian boundary conditions. We formulate the following definition.

Definition 3.1.9. A **relative pin structure** for $Q \subset \partial X$ is

- a stable oriented vector bundle $\xi \rightarrow X$, with $w_2(\xi|Q) = w_2(T^\vee Q)$; and
- a stable pin structure on $\xi|Q \oplus T^\vee Q$.

Recall that $\text{Pin}(n)$ is the double cover of $\text{O}(n)$ which is non-trivial over each of the two components, and that a pin structure on a real n -plane bundle ζ over Z is a homotopy class of homotopy-liftings $Z \rightarrow B\text{Pin}(n)$ of the classifying map $Z \rightarrow B\text{O}(n)$. The inclusions $\text{O}(n) \rightarrow \text{O}(n+1)$ lift to the pin groups, and this gives the meaning of ‘stable’.

The classifying space $B\text{Pin}(n)$ is path-connected with $\pi_1 = \mathbb{Z}/2$, $\pi_2 = \pi_3 = 0$. This leads to the following observation, which explains the relevance of pin structures.

Lemma 3.1.10. *A stable pin structure on ζ determines*

1. *an isomorphism $a^* \det(\zeta) \cong \iota$ for any map $a: S^1 \rightarrow Z$: here $\iota \rightarrow S^1$ is the trivial line bundle ε if ζ is orientable, the Möbius bundle μ if not; and*
2. *a stable trivialisation of $b^*(\zeta \oplus \det(\zeta))$ for any two-complex Σ and any map $b: \Sigma \rightarrow Z$.*

Proposition 3.1.11. *Take a simple quadratic fibration $X \rightarrow \bar{D}$, with Lagrangian boundary condition $Q \subset \partial X$. A relative pin structure for (X, Q) induces an orientation of the moduli space $\mathcal{M}(X, Q)$.*

Remark 3.1.12. *This is still true when the disk \bar{D} is replaced by a general base surface S .*

We precede the proof with a discussion of the space of Cauchy-Riemann operators with totally real boundary condition over the disk. Let E be the trivial bundle $\mathbb{C}^n \times \bar{D} \rightarrow \bar{D}$. The Grassmannian of conjugation-invariant, n -dimensional real subspaces of \mathbb{C}^n is $U(n)/O(n)$, and any smooth map $f: \partial \bar{D} \rightarrow U(n)/O(n)$ sweeps out a totally real sub-bundle $F(f) \subset \mathbb{C}^n \times \partial \bar{D}$ over $\partial \bar{D}$. By choosing a connection ∇ on E and a bundle map a , as in (3.6), we get a Fredholm Cauchy-Riemann problem for $(E, F(f))$. Triples (f, ∇, a) form a space $\mathcal{CR}(n)$, which maps to the loop space $L(U(n)/O(n))$ with contractible fibres. It has components $\mathcal{CR}_j(n)$, indexed by $j \in H_1(U(n)/O(n)) = \mathbb{Z}$. As far as index theory is concerned, we may as well work with $L(U(n)/O(n))$ rather than $\mathcal{CR}(n)$.

We are interested in the determinant line bundle $\text{Det} \rightarrow L(U(n)/O(n))$. The so-called ‘family index anomaly’ is the fact that this is a non-trivial bundle. However, its pullback to either of the spaces, $L(U(n)/O(n))^\sim$ or $L(U(n)/O(n))^\smile$ is trivial. These are defined as follows.

- An element of $L(U(n)/O(n))^\sim$ is a pair (f, λ) , where $f \in L(U(n)/O(n))$ and λ is a homotopy class of stable isomorphisms $F(f) \simeq \varepsilon^n$ or $F(f) \simeq \varepsilon^{n-1} \oplus \mu$. (For a given f , only one of these two possibilities can occur.)
- An element of $L(U(n)/O(n))^\smile$ is a pair (f, h) , where $f \in L(U(n)/O(n))$ and h is a homotopy $f \simeq l^j$ for some j and a fixed loop $l: S^1 \rightarrow U(n)/O(n)$ of Maslov index 1.

The natural map $L(U(n)/O(n))^\sim \rightarrow L(U(n)/O(n))$ is a two-fold covering. But each of the two components of $L(U(n)/O(n))$ has fundamental group $\mathbb{Z}/2$, and it is not difficult to see that $\pi_0 L(U(n)/O(n))^\sim = \mathbb{Z}/2$. Hence the covering is universal over each component, and the pullback of Det to $L(U(n)/O(n))^\sim$ is orientable.

It is clear that the pullback of Det to $L(U(n)/O(n))^\smile$ is orientable. To specify an orientation it is enough to orient the determinant lines of the standard Cauchy-Riemann operators $\bar{\partial}$ on $(E, F(l^j))$. Here is a recipe for doing this:

- We can take $F(l^j)_z = z^{j/2}\mathbb{R} \oplus \mathbb{R}^{n-1}$. Then, for $j \geq 0$, $\bar{\partial}^{-1}(0)$ consists of n -tuples $(c_1 f; c_2, \dots, c_n)$, where $f(z) = (e^{2i\theta_1}z - e^{-2i\theta_1}) \dots (e^{2i\theta_j}z - e^{-2i\theta_j})$ for constants $e^{i\theta_p} \in S^1$ and $c_q \in \mathbb{R}$. That *every* solution is of this form is an application of the maximum principle. To go between j and $-j$, replace z by z^{-1} .
- These moduli spaces are regular: they are manifolds of the correct dimension $j+n$, so it suffices to check that $\ker(D\bar{\partial})$ has constant rank $j+n$. Having done that, we only have to orient their tangent bundles: we give $T(\bar{\partial}^{-1}(0)) = T(\mathbb{T}^j \times \mathbb{R}^n)$ its standard orientation.

There is a natural map

$$p: L(\mathrm{U}(n)/\mathrm{O}(n))^\sim \rightarrow L(\mathrm{U}(n)/\mathrm{O}(n))^\sim,$$

over $L(\mathrm{U}(n)/\mathrm{O}(n))$. Thus there is at most one way to orient $\mathrm{Det} \rightarrow L(\mathrm{U}(n)/\mathrm{O}(n))^\sim$ so that its pullback to $L(\mathrm{U}(n)/\mathrm{O}(n))^\sim$ has the orientation we have specified. We just need to check that the orientation on $L(\mathrm{U}(n)/\mathrm{O}(n))^\sim$ really does descend to $L(\mathrm{U}(n)/\mathrm{O}(n))^\sim$. The fibres of $L(\mathrm{U}(n)/\mathrm{O}(n))^\sim \rightarrow L(\mathrm{U}(n)/\mathrm{O}(n))$ are homotopy-equivalent to $\Omega^2(\mathrm{U}(n)/\mathrm{O}(n))$, and the effect of p on the fibre over $f \in L(\mathrm{U}(n)/\mathrm{O}(n))$ is that it collapses the two components of $\Omega^2(\mathrm{U}(n)/\mathrm{O}(n))$ to the two points lying over f in $L(\mathrm{U}(n)/\mathrm{O}(n))^\sim$. Thus p has path-connected fibres, and the orientation descends.

Proof of 3.1.11. The Cauchy-Riemann problem associated with our moduli space $\mathcal{M}(X, Q)$ is classified by a homotopy class of maps $g: \mathcal{M}(X, Q) \rightarrow L(\mathrm{U}(n)/\mathrm{O}(n))$. To prove the proposition it will suffice to show that a relative pin structure determines a lift $\tilde{g}: \mathcal{M}(X, Q) \rightarrow L(\mathrm{U}(n)/\mathrm{O}(n))^\sim$. Indeed, it is enough to lift g over loops in \mathcal{M} .

We argue roughly as in [12, p. 191]. Take a loop in $\mathcal{M}(X, Q)$, considered as a map $u: (S^1 \times D, S^1 \times \partial D) \rightarrow (X, Q)$. Write v for the induced map $S^1 \times \partial D \rightarrow Q$. The bundle $u^*T^\vee X \rightarrow S^1 \times D$ is trivialised by the orientation of $T^\vee X$, so to lift the loop u to $L(\mathrm{U}(n)/\mathrm{O}(n))^\sim$ we need to produce a stable isomorphism $v^*T^\vee Q \simeq \mathrm{pr}_2^* \iota$, where as before, $\iota \rightarrow S^1$ is either the trivial bundle ε or the Möbius bundle μ .

The pin structure induces a stable trivialisation of $v^*(\xi \oplus T^\vee Q \oplus \det(\xi \oplus T^\vee Q))$. On the other hand, the orientation of ξ induces a stable trivialisation of $u^*\xi$, which restricts to one of $v^*\xi$. Taking this into account, we are left with a stable trivialisation of $v^*(T^\vee Q \oplus \det(T^\vee Q))$.

The pin structure also gives us, for each $t \in S^1$, an isomorphism $v_t^* \det(\xi T^\vee Q) \cong \iota$ (where $v_t = v(t, \cdot)$), and hence also an isomorphism $v_t^*(\det T^\vee Q) \cong \iota$; allowing t to vary we get $v^*(\det T^\vee Q) \cong \mathrm{pr}_2^* \iota$. Combining this with the stable trivialisation of $v^*(T^\vee Q \oplus \det(T^\vee Q))$ we obtain a stable isomorphism $v^*T^\vee Q \simeq \mathrm{pr}_2^* \iota$, as desired.

□

3.1.4 Invariance

We now give the parametric version of Theorem 3.1.8. All the data involved in the construction can be allowed to move in smooth families. To begin with, *fix* the simple quadratic fibration $(X, \pi, \Omega, J_0, j_0)$, Lagrangian boundary condition Q , and homotopy class h .

Each j and each $J \in \mathcal{J}^{\text{reg}}(X, \pi; j)$ gives a moduli space $\mathcal{M}_{J,j}(X, Q;)_h$ which we know to be a smooth, oriented manifold; after a small perturbation of J , the fibre product $\mathcal{M}_{J,j}(X, Q;)_h \times_{\text{ev}, \zeta \times \zeta'} Z \times Z'$ with a family of smooth oriented cycles in the fibres (as considered above) is also smooth and oriented. It carries a smooth boundary-evaluation map

$$\text{ev}(J, j): \mathcal{M}_{J,j}(X, Q)_h \rightarrow \mathcal{S}(Q)$$

to the space of sections of Q .

Theorem 3.1.13. *Assume that the hypotheses of Theorem 3.1.8 hold, and that the left-hand side of (3.9) is $< -(k+1)$, for some $k \geq 0$. Then the oriented bordism class of the map $\text{ev}(J, j)$ is independent of (J, j) . Moreover, any smooth map $S^k \rightarrow \mathcal{J}^{\text{reg}}(X, \pi)$ induces a bundle \mathcal{M}_{S^k} over S^k and a map $\text{ev}: \mathcal{M}_{S^k} \rightarrow \mathcal{S}(Q)$ which extends to a map to $\mathcal{S}(Q)$ from a bundle over B^{k+1} .*

Proof. One considers parametrised moduli spaces. The space $\mathcal{J}(X, \pi)$ is contractible; thus any map $f: S^k \rightarrow \mathcal{J}^{\text{reg}}(X, \pi)$ extends to a map $F: B^{k+1} \rightarrow \mathcal{J}(X, \pi)$. The transversality theory extends in a standard and straightforward way, showing that after a small homotopy of F , fixing f , one gets a smooth moduli space of the appropriate dimension which inherits an orientation from that on B^{k+1} . The numerical conditions imply that compactness goes through as before, and this gives the result. \square

Likewise, the bordism class of the evaluation map is unchanged under deformations of Ω , J_0 , j_0 , the points s_a , s'_b and the cycles ζ_a and ζ'_b .

One can also allow deformations of the data (X, π, Q) , this results in a ‘dynamic’ bordism class (that is, one with varying target).

3.2 Field theories

Field theories, in their mathematical incarnation, are functors from cobordism categories to linear tensor categories, such that disjoint union corresponds to tensor product and orientation-reversal to dualisation.

In the theory we describe, the cobordism category has as its objects locally Hamiltonian fibrations over closed, oriented 1-manifolds, and as its morphisms locally Hamiltonian fibrations over compact, oriented 2-manifolds with boundary; the fibres are supposed to be diffeomorphic to some fixed compact symplectic manifold. The target category is that of

$\mathbb{Z}/2$ -graded modules over the universal Novikov ring $\Lambda_{\mathbb{Z}}$, and the functor, HF_* is a version of symplectic Floer homology.

Cobordism category of locally Hamiltonian fibrations

(1a) The $(1+1)$ -dimensional cobordism category has as its objects closed, oriented, smooth 1-manifolds. A morphism from Z_0 to Z_1 is an oriented-diffeomorphism class of cobordisms, where a cobordism is a compact, oriented, smooth 2-manifold S and a diffeomorphism $\partial S \rightarrow -Z_0 \amalg Z_1$. Composing two cobordisms involves choosing boundary collars, but the small resulting ambiguity disappears upon passing to the diffeomorphism class.

(2a) For each symplectic manifold (M, ω) , define a category $\mathcal{F}(M, \omega)$:

- An object of \mathcal{F} is a proper locally Hamiltonian fibration (Y, π, σ) over a closed, oriented 1-manifold Z , such that every fibre $(X_s, \sigma|_{X_s})$ is symplectomorphic to (M, ω) .
- A morphism from (Y_-, π_-, σ_-) to (Y_+, π_+, σ_+) is an isomorphism class of **cobordisms** between them.

A ‘cobordism’ is a locally Hamiltonian fibration (X, π, Ω) over a compact oriented surface S , with fibres symplectomorphic to (M, ω) , flat near ∂S , together with an isomorphism from $(X, \pi, \Omega)|_{\partial S}$ to the disjoint union of $(Y_-, \pi_-, \sigma_-)^{\text{op}}$ and (Y_+, π_+, σ_+) (here the superscript $^{\text{op}}$ means that the data are pulled back by $S^1 \rightarrow S^1, t \mapsto -t$).

The flatness condition facilitates composition of cobordisms. A flat locally Hamiltonian fibration over $Z \times [0, 1)$ (or over $Z \times (-1, 0]$) is canonically isomorphic to the pullback of its restriction to $Z \times \{0\}$. To compose cobordisms, choose boundary collars over which the Hamiltonian connection is flat. One can then glue the surfaces along these collars, and glue the fibrations along their restrictions to the collars.

(1b) As a small variation on the definitions above, one can work with manifolds with cylindrical ends instead of boundaries. This is what is required for the technical construction (though not necessarily the presentation) of Floer theory.

The cylindrical cobordism category has the same objects as that of (1a), but a morphism is now a diffeomorphism class of **cylindrical cobordisms**. A cylindrical cobordism from Z_- to Z_+ is an oriented surface S , together with disjoint, proper, oriented embeddings $e_-: (-\infty, 0] \times Z_- \rightarrow S$ and $e_+: [0, \infty) \times Z_+ \rightarrow S$, such that $S \setminus (\text{im}(e_-) \cup \text{im}(e_+))$ has compact closure. Composition can then be defined by some universal rule such as the following: define $g: (-1, 1) \rightarrow [0, \infty) \times (-\infty, 0]$ by $t \mapsto (-2^{-1} + (1-t)^{-1}, 2^{-1} - (1+t)^{-1})$; glue an incoming end e_+ to an outgoing end e'_- using $g \times \text{id}_{S^1}$.

(2b) $\mathcal{F}(M, \omega)_{\text{cyl}}$ is the cylindrical-ends version of $\mathcal{F}(M, \omega)$, and has the same objects. A morphism from (Y_-, π_-, σ_-) to (Y_+, π_+, σ_+) is an isomorphism class of cylindrical cobordisms between them. A cylindrical cobordism is a locally Hamiltonian fibration (X, π, Ω) over an oriented surface, together with embeddings $e_-: (-\infty, 0] \times Z_- \rightarrow S$ and $e_+: [0, \infty) \times Z_+ \rightarrow S$ as in (1b), and lifts $\hat{e}_-: (-\infty, 0] \times Y_- \rightarrow X$, $\hat{e}_+: [0, \infty) \times Y_+ \rightarrow X$. These must satisfy $\hat{e}_\pm^* \Omega = \text{pr}_2^* \sigma_\pm$, that is, they must be embeddings of locally Hamiltonian fibrations.

Structure of Floer homology

Floer homology for automorphisms of monotone (or weakly monotone) symplectic manifolds is defined over the universal Novikov ring Λ_R , as defined in the introduction, over an arbitrary commutative ring R . For us, usually $R = \mathbb{Z}$.

The purpose of Novikov rings in Floer theory is to keep track of the areas of holomorphic sections with fixed asymptotic values. To put this another way, the coefficients record the periods of the action functional. If one can arrange that it has no periods (that is, it is an exact 1-form) then one can work over $\mathbb{Z}/2$ or \mathbb{Z} .

Floer homology is a covariant functor from $\mathcal{F}(M, \omega)$ to $\mathbb{Z}/2$ -graded $\Lambda_{\mathbb{Z}/2}$ -modules, where (M, ω) is a weakly monotone compact symplectic manifold. It sends an object (Y, π, σ) to the module

$$HF_*(Y, \sigma) = HF_0(Y, \sigma) \oplus HF_1(Y, \sigma).$$

There is a canonical splitting into ‘topological sectors’, i.e. components of the space of sections $\mathcal{S}(Y)$,

$$HF_*(Y, \sigma) = \bigoplus_{\gamma \in \pi_0 \mathcal{S}(Y)} HF_*(Y, \sigma)_\gamma,$$

and a similar decomposition of the map associated to a cobordism.

3.2.1 The theory in outline

Full and self-contained accounts of the foundations of *Hamiltonian* Floer homology in the monotone case can be found e.g. in Schwarz [40]. There is no comparable source for Floer homology of general symplectic automorphisms (that is, of locally Hamiltonian fibrations), so it is useful to note that statements made in the Hamiltonian theory often generalise immediately to the broader context. This applies to the linear analysis in its entirety, as it takes place on vector bundles over surfaces, and not in the fibre bundles themselves.

Critical points of the action functional

The Floer homology of an object (Y, π, σ) of $\mathcal{F}(M, \omega)_{\text{cyl}}$ is the Morse-Novikov homology of the action 1-form $\mathcal{A}_{Y, \sigma} \in \Omega^1(\mathcal{S}(Y))$ on the space of sections $\mathcal{S}(Y)$:

$$\mathcal{A}_{Y, \sigma}(\gamma; \xi) = \int_Z \sigma(\dot{\gamma}, \xi), \quad \xi \in C_Z^\infty(\gamma^* T^\vee Y). \quad (3.12)$$

The set of zeros of $\mathcal{A}_{Y, \sigma}$ coincides with the set $\mathcal{H}(Y, \sigma)$ of horizontal sections defined by the Hamiltonian connection:

$$\mathcal{H}(Y, \sigma) = \{\nu: Z \rightarrow Y : \pi \circ \nu = \text{id}_Z, \text{im } D_z \nu \subset T_{\nu(z)}^{\text{h}} X\}. \quad (3.13)$$

One should again enlarge the space of sections $\mathcal{S}(Y)$ to a Banach manifold, namely $\mathcal{S}_1^2(Y)$, which has tangent spaces $L_1^2(\gamma^* T^\vee Y)$. The transversality condition for a zero ν of $\mathcal{A}_{Y, \sigma}$ is surjectivity of the operator

$$\xi \mapsto \nabla_{\partial_t} \xi; \quad L_1^2(\nu^* T^\vee Y) \rightarrow L^2(\nu^* T^\vee Y),$$

where ∇ is the intrinsic connection along ν . After fixing basepoints $z_i \in Z$, one in each component, we can consider the linear holonomy maps $L_{\nu, i} \in \text{End}(T_{\nu(z_i)}^\vee Y)$. Transversality of ν is equivalent in turn to the invertibility of the linear maps $\text{id} - L_{\nu, i} \in \text{End}(T_{\nu(z_i)}^\vee Y)$. This is the condition that $\nu(z)$ is a non-degenerate fixed point of the monodromy $\phi \in \text{Aut}(Y_z, \sigma|_{Y_z})$. It follows that when Y is **non-degenerate**, meaning that every $\nu \in \mathcal{H}(Y, \sigma)$ is transverse, the set $\mathcal{H}(Y, \sigma)$ is *finite* (as $Y_z \cong M$ is compact).

For a non-degenerate object (Y, π, σ) , set

$$CF(Y, \sigma) = \bigoplus_{\nu \in \mathcal{H}(Y, \sigma)} \Lambda_R,$$

for some chosen ring R , and denote the generators of this free Λ_R -module by $\langle \nu \rangle$. This is the module underlying Floer's (co)chain complex. The definition does not yet use the orientation of Z .

Each $\nu \in \mathcal{H}(Y, \sigma)$ has a **Lefschetz number** $l_\nu \in \mathbb{Z}/2$, defined when Z is connected to be 0 if $\det(\text{id} - L_\nu) > 0$ and 1 if $\det(\text{id} - L_\nu) < 0$, and in general by summing the Lefschetz numbers over components of Z . These give the $\mathbb{Z}/2$ -grading of the complex, $CF = CF_0 \oplus CF_1$, with CF_i generated by $\{\nu : l_\nu = i\}$. It follows that when $Z = \coprod Z_i$ there are canonical $\mathbb{Z}/2$ -graded isomorphisms $CF(Y) \cong \bigotimes_i CF(Y|_{Z_i})$.

Remark 3.2.1. For a mapping torus (M_ϕ, ω_ϕ) , over S^1 , monodromy gives a bijection between $\mathcal{H}(M_\phi, \omega_\phi)$ and the fixed-point set $\text{Fix}(\phi)$. Non-degeneracy for horizontal sections corresponds to non-degeneracy for fixed points, and likewise for Lefschetz numbers.

Cylindrical almost complex structures. Take a locally Hamiltonian fibration (Y, π, σ) over S^1 . A compatible vertical almost complex structure $J^v \in \mathcal{J}(T^v Y, \sigma)$ extends uniquely to an almost complex structure J_0 on $TY \oplus \mathbb{R}$ such that $J(\tilde{\partial}_t) = 1 \in \mathbb{R}$ (here ∂_t is the unit length vector field on $S^1 = \mathbb{R}/\mathbb{Z}$ and $\tilde{\partial}_t$ its horizontal lift).

Since $T(Y \times \mathbb{R}) = TY \oplus \mathbb{R}$, J_0 induces a translation-invariant complex structure on $T(Y \times \mathbb{R})$. It is still denoted J_0 , and called **cylindrical**. If we regard $Y \times \mathbb{R}$ as a locally Hamiltonian fibration over the cylinder $S^1 \times \mathbb{R}$, with 2-form $\text{pr}_1^* \sigma$, then J_0 preserves both vertical and horizontal subbundles, and the projection to $S^1 \times \mathbb{R} = \mathbb{C}/i\mathbb{Z}$ is holomorphic.

Now consider a cylindrical cobordism $(X, \pi, \Omega; e_\pm, \hat{e}_\pm)$ from (Y_-, π_-, σ_-) to (Y_+, π_+, σ_+) in the category $\mathcal{F}(M, \omega)_{\text{cyl}}$. Fix cylindrical almost complex structures J_\pm on the ends, together with a complex structure j on S such that $e_-^* j$ (resp. $e_+^* j$) is equal to i on the complement of a compact subset in $(-\infty, 0] \times S^1$ (resp. $[0, \infty) \times S^1$). There is then a contractible subspace $\mathcal{J}(X, \pi; j, J_\pm) \subset \mathcal{J}(X, \pi; j)$ comprising almost complex structures J which are **adapted to J_\pm** . This means that J satisfies $\hat{e}_\pm^* J = J_\pm$ on $(-\infty, T] \times Y_-$ and $[T, \infty) \times Y_+$, for some $T \geq 0$.

Finite-action pseudoholomorphic sections

Let (Y_\pm, σ_\pm) be objects of $\mathcal{F}(M, \omega)_{\text{cyl}}$, and $(X, \pi, \Omega, e_\pm, \hat{e}_\pm)$ a morphism between them. Fix vertical complex structures $J_\pm \in \mathcal{J}(T^v Y_\pm)$, and j on the base S . For each $J \in \mathcal{J}(X, \pi; j, J_\pm)$ we have a moduli space

$$\mathcal{M}_{J,j}(X, \pi, \Omega) \subset \mathcal{S}(X)$$

of finite-action pseudoholomorphic sections, i.e. of C^∞ sections $u: S \rightarrow X$ satisfying

$$J \circ (Du) = (Du) \circ j; \tag{3.14}$$

$$\int_S u^* \Omega < \infty. \tag{3.15}$$

The interpretation of (3.15) is that the integral converges absolutely on each cylindrical end. Such sections behave well when the objects (Y_\pm, σ_\pm) are non-degenerate. In fact, under the assumption of non-degeneracy, any $u \in \mathcal{M}_{J,j}(X, \pi, \Omega)$ has the following asymptotic behaviour:

- As $s \rightarrow -\infty$, the loops $(\hat{e}_-^* u)(s, \cdot)$ converge pointwise towards a horizontal section $\nu_- \in \mathcal{H}(Y_-, \sigma_-)$.
- The convergence is exponentially fast with respect to the Riemannian metric $g_- = \sigma_-(\cdot, J_- \cdot)$ on Y_- .
- The loops $(\hat{e}_+^* u)(s, \cdot)$ converge in similar fashion to some $\nu_+ \in \mathcal{H}(Y_+, \sigma_+)$.

The transversality theory discussed earlier in the chapter applies, using very similar arguments, to the moduli spaces of finite-energy pseudoholomorphic sections. The one important difference is that the linearised operators D_u are now *cylindrical Cauchy-Riemann operators*. We will discuss these below. Again, these are Fredholm operators (in a suitable Sobolev setting) and by repeating the earlier argument involving the Sard-Smale theorem we conclude that

$$\mathcal{J}^{\text{reg}}(X, \pi; J^\pm j) \text{ is } C^\infty\text{-dense in } \mathcal{J}(X, \pi; J^\pm, j).$$

Here $\mathcal{J}^{\text{reg}}(X, \pi; J^\pm, j)$ is the subspace for which the operators D_u are surjective for all $u \in \mathcal{M}_{J,j}(X, \pi, \Omega)$.

Cylindrical Cauchy-Riemann operators

Model case: trivialised vector bundles over the cylinder. The 2-point-compactified real line $\overline{\mathbb{R}} = \{-\infty\} \cup \mathbb{R} \cup \{\infty\}$ becomes a smooth manifold upon declaring the bijection $\overline{\mathbb{R}} \rightarrow [-1, 1]$, $t \mapsto t(1+t^2)^{-1/2}$, $\pm\infty \mapsto \pm 1$, to be a diffeomorphism. Let $\overline{Z} := \overline{\mathbb{R}} \times S^1$. We work with the trivial rank n hermitian vector bundle ξ over \overline{Z} . The hermitian structure $h_0 = g_0 + i\omega_0$ is the standard one. Generally, when we have objects defined over \overline{Z} , we will denote their restrictions to the circles $\{\pm\infty\} \times S^1$ by adding superscripts \pm .

A *real, self-adjoint* section $A = A(s, t)$ of the (trivial) endomorphism bundle $\text{End}(\xi)$ determines a linear operator

$$L_A: L_1^2(\xi) \rightarrow L^2(\xi); \quad L_A u = \partial_s u + i\partial_t u + A \circ u.$$

Here L^2 and L_1^2 are defined using the hermitian metric h_0 . This operator is asymptotic at $\pm\infty$ to a pair of operators L_A^\pm ,

$$(L_A^\pm v)(t) = i\dot{v}(t) + A^\pm(t)v(t), \quad v \in L_1^2(\xi|_{\{\pm\infty\} \times S^1}),$$

both self-adjoint in L^2 . The loop $A^+(t)$ determines a path $\gamma = \gamma_{A^+}: [0, 1] \mapsto \text{End}(\mathbb{R}^{2n})$, by

$$\dot{\gamma}(t) = iA(t) \circ \gamma(t), \quad \gamma(0) = \text{id}.$$

This is a Hamiltonian flow; thus $\text{im}(\gamma) \subset \text{Sp}(\mathbb{R}^{2n})$. We say that A^+ is **non-degenerate** if $\ker(L_A^+) = 0$. When this holds, $\gamma(1) \in \text{Sp}^*(\mathbb{R}^{2n})$, that is, it does not have 1 as an eigenvalue.

Making the same definitions for L^- , we put

$$\mathcal{A}_n = \{L_A : A^+ \text{ and } A^- \text{ are non-degenerate}\} \subset \mathcal{L}(L_1^2(\xi), L^2(\xi)).$$

Once one knows that its elements are continuous, one can give \mathcal{A}_n the operator-norm topology. The crucial facts about \mathcal{A}_n are these:

- Elements of \mathcal{A}_n are (real) Fredholm operators with index

$$\text{ind}(L_A) = \text{CZ}(\gamma_{A^+}) - \text{CZ}(\gamma_{A^-}), \quad (3.16)$$

where CZ is the Conley-Zehnder index for paths in $\text{Sp}(\mathbb{R}^{2n})$ starting at $\text{id}_{\mathbb{R}^{2n}}$ and ending in $\text{Sp}^*(\mathbb{R}^{2n})$. The parity $(-1)^{\text{ind}(L_A)}$ is the product of the signs of the non-zero real numbers $\det(\text{id} - \gamma_{A^\pm}(1))$.

- For $L \in \mathcal{A}_n$, put $\mathcal{A}_n(L) = \{K \in \mathcal{F}_n : K^\pm = L^\pm\}$. This is a contractible subspace; in particular, the restriction of the determinant index bundle, $\text{Det}|_{\mathcal{A}_n(L)}$, is trivial.

Let $\mathcal{A}_n^{ac} \subset \mathcal{A}_n$ denote the subspace comprising L_A where A is asymptotically constant, i.e. independent of $s \in \mathbb{R}$ for $|s| \gg 0$; it is a deformation-retract of \mathcal{A}_n .

There is a simple procedure for gluing a pair $L_1, L_2 \in \mathcal{A}_n^{ac}$, providing $L_1^+ = L_2^-$, to obtain an element $L_1 \#_\rho L_2 \in \mathcal{A}_n^{ac}$ depending on a translation parameter $\rho \gg 0$. The index is additive under gluing, and there is in fact a natural class of isomorphisms

$$\text{Det}(L_1) \otimes \text{Det}(L_2) \rightarrow \text{Det}(L_1 \#_\rho L_2).$$

These satisfy an associativity rule. Given orientations $\mathfrak{o}(L_i)$ for the lines $\text{Det}(L_i)$ ($i = 1, 2$) one gets an orientation $\mathfrak{o}(L_1) \wedge \mathfrak{o}(L_2) = \mathfrak{o}(L_1 \#_\rho L_2)$ for $\text{Det}(L_1 \#_\rho L_2)$, and—under obvious hypotheses— $(\mathfrak{o}(L_1) \wedge \mathfrak{o}(L_2)) \wedge \mathfrak{o}(L_3) = \mathfrak{o}(L_1) \wedge (\mathfrak{o}(L_2) \wedge \mathfrak{o}(L_3))$.

Definition 3.2.2. A **linear coherent orientation** is an orientation $L \mapsto \mathfrak{o}(L)$ for $\text{Det} \rightarrow \mathcal{A}_n$ with the property that

$$\mathfrak{o}(L_1 \#_\rho L_2) = \mathfrak{o}(L_1) \wedge \mathfrak{o}(L_2)$$

for all glueable pairs (L_1, L_2) of asymptotically constant operators.

Linear coherent orientations exist (but are far from unique). The starting point is to orient each L for which L^+ is some fixed operator l ; this can be done arbitrarily, except that the determinant of the pullback of l to \overline{Z} has a canonical orientation (since L is an isomorphism), and this should be adopted. The requirement of coherence then determines orientations for L with $L^- = l$, then for $L_1 \#_\rho L_2$ with $L_1^+ = \ell = L_2^-$, and thereby for all L .

Surfaces with cylindrical ends. We draw here on the discussion by Eliashberg, Givental and Hofer [11]. Define a category \mathcal{C} as follows. An object of \mathcal{C} is an integer $m \geq 0$ and an m -tuple (A_1, \dots, A_m) , where each $A_i : S^1 \rightarrow \text{End}(\mathbb{R}^{n_i})$ is a non-degenerate loop of self-adjoint operators. A morphism from (A_1^-, \dots, A_l^-) to (A_1^+, \dots, A_m^+) is, roughly, a Cauchy-Riemann operator on a Riemann surface with cylindrical ends, asymptotic to the operators $L_{A_i^\pm}$. Precisely, it is an isomorphism class of data consisting of: a Riemann surface (S, j) with cylindrical ends

$$e_- : (-\infty, 0] \times S^1 \times \{1, \dots, l\} \rightarrow S, \quad e_+ : [0, \infty) \times S^1 \times \{1, \dots, m\};$$

a hermitian vector bundle (E, h) over S ; hermitian trivialisations

$$\Phi_- : e_-^* E \rightarrow \prod_{i=1}^l S^1 \times \mathbb{C}^{n_i^-}, \quad \Phi_+ : e_+^* E \rightarrow \prod_{i=1}^m S^1 \times \mathbb{C}^{n_i^+};$$

and a **cylindrical Cauchy-Riemann operator**

$$L : L_1^2(E) \rightarrow L^2(\text{Hom}^{0,1}(TS, E)), \quad Lu(x) = \nabla h + i(\nabla u) \circ j + Au.$$

The Sobolev spaces L^2 and L_1^2 are defined using the trivialisations Φ_\pm ; ∇ is any connection, and $A : E \rightarrow \text{Hom}^{0,1}(TS, E)$ a bundle map. We require that $\Phi_-^* \circ e_-^*(L)$ should be asymptotic to (A_1^-, \dots, A_l^-) , and similarly at the positive ends.

There is an obvious gluing procedure for *asymptotically constant* Cauchy-Riemann operators. The composition rule for morphisms in \mathcal{C} is the unique one which is consistent with this satisfying the rules for a category.

\mathcal{C} is a topological category: the morphism spaces have natural topologies in which a component is represented by an isomorphism class of pairs (oriented topological surface, complex vector bundle).

The Cauchy-Riemann operators involved in \mathcal{C} are again Fredholm. This can be seen in terms of a ‘locality’ principle, proved in [7]; when one has parametrices for elliptic operators on the sets in an open cover of a manifold, one can patch them to obtain a global parametrix. Over each morphism space there is therefore a determinant line bundle Det .

Definition 3.2.3. A **coherent orientation system** is the assignment of an orientation of Det over each morphism space, in a way which is compatible with composition.

Coherent orientation systems exist in profusion, as Floer and Hofer showed. Moreover, one can impose three further compatibility requirements: a simple compatibility property for direct sums; a more subtle ‘cut-and-paste’ property (see [11]); and the requirement that the determinant of a complex-linear operator over a closed Riemann surface is assigned its complex orientation. A coherent orientation system for cylinders extends uniquely to a coherent orientation system satisfying these properties.

Canonical coherent orientations. We return now to the category $\mathcal{F}(M, \omega)_{\text{cyl}}$, and to the problem of orienting its moduli spaces of finite energy pseudoholomorphic sections. The usual procedure at this point is to make an arbitrary choice of a coherent orientation system, which can then be pulled back to $\mathcal{F}(M, \omega)_{\text{cyl}}$.¹ There is, however, a canonical choice (this may be known to experts, but I am unaware of any reference).

¹Or, in situations where one has a proposed isomorphism of some Floer homology group with, say, a classical homology group, one makes the unique choice which will validate it.

Take an object (Y, π, σ) of $\mathcal{F}(M, \omega)_{\text{cyl}}$. The fibre product $(Y \times_{S^1} Y, \pi \times_{S^1} \pi, \sigma \oplus -\sigma)$ is a locally Hamiltonian fibration over S^1 , and it contains two Lagrangian subbundles: one is the diagonal Δ , and the other is $\Gamma := \{(y, my) : y \in Y\}$, the graph of the monodromy m of (Y, σ) . An argument similar to the one we gave in the Lagrangian situation above (or, in fact, a reduction to that situation) shows that to give coherent orientations for the spaces of asymptotically horizontal sections of $Y \times \mathbb{R}$, it suffices to give

- a vector bundle $\xi \rightarrow Y \times_{S^1} Y$; and
- spin structures on $\xi|_{\Delta} \oplus T^\vee \Delta \rightarrow \Delta$ and $\xi|_{\Gamma} \oplus T^\vee \Gamma \rightarrow \Gamma$.

Take $\xi = \text{pr}_1^* T^\vee Y$. Both $T^\vee \Gamma \oplus T^\vee \Gamma \rightarrow \Gamma$ and $T^\vee \Delta \oplus T^\vee \Delta$ are isomorphic, via pr_1 , to $T^\vee Y \oplus T^\vee Y \rightarrow Y$. The structure group of $T^\vee Y$ is $\text{SO}(n)$, where n is the rank, and that of $T^\vee Y \oplus T^\vee Y$ is reduced to the diagonal subgroup $\text{SO}(n) \subset \text{SO}(2n)$. But, by looking at its action on π_1 , one sees that the diagonal inclusion factors uniquely through a homomorphism to $\text{Spin}(2n)$, and this gives a canonical spin structure on $T^\vee Y \oplus T^\vee Y$.

3.2.2 Floer homology in the fibre-monotone case

Compactness. The Gromov-Floer compactness theorem says that a sequence $\{u_n\}_n$ in $\mathcal{M}_J(\nu_- |X| \nu_+)$ (the component of the moduli space where the sections approach ν_\pm at $\pm\infty$), of bounded action $\mathcal{A}(u_n)$, converges in the *Gromov-Floer topology* to a broken section plus bubbles. Again, we decline to define convergence in this topology, but simply describe the limit points. A broken section comprises finite sequences $(\nu_-^0, \dots, \nu_-^l)$ and $(\nu_+^0, \dots, \nu_+^m)$ in $\mathcal{H}(Y, \sigma)$, with $\nu_- = \nu_-^0$ and $\nu_+ = \nu_+^0$, together with an element $w = (w_-^0, \dots, w_-^{l-1}; w_{\text{prin}}; w_+^{m-1}, \dots, w_+^0)$ of the product

$$\prod_{p=0}^{l-1} \mathcal{M}_{J,-}(\nu_-^p, \nu_-^{p+1}) \times \mathcal{M}_{J,j}(\nu_-^l |X| \nu_+^m) \times \prod_{q=0}^{m-1} \mathcal{M}_{J,+}(\nu_+^{m-q}, \nu_+^{m-q-1}).$$

Only w_{prin} is allowed to be a constant section. A Gromov-Floer limit consists of a broken section together with trees of holomorphic spheres in fibres, rooted at points in the images of components of w . The limit of a sequence u_n with constant action λ has total action λ (‘total action’ means the sum of the actions of all components—broken section and bubbles) and likewise for indices.

A similar result holds for sequences in $\mathcal{M}_J(\nu^-, \nu^+)$; one just omits the ‘principal component’ w_{prin} , and insists that at most one component of a broken section can be constant.

Fibred monotonicity implies that sequence of sections in low-dimensional, regular moduli spaces do not have bubbles in their limits. Each bubble contributes at least $2c_{\min} > 0$ to the total index, while for regular almost complex structures each component of a broken section makes a non-negative contribution. Hence a limit point of $\mathcal{M}_J(\nu_-, \nu_+)_k$ with $k < 2c_{\min}$ is a

broken section; if $k = 2c_{\min}$ then a limit point is either a broken section or a constant section with one bubble. A limit point of $\mathcal{M}_{J,j}(\nu_-|X|\nu_+)_k$ with $k < 2c_{\min}$ is a broken section.

Floer homology. We now review the definition of Floer homology. The Floer complex $(CF(Y, \sigma), \partial)$ is initially defined when (Y, σ) is non-degenerate. Once one has established that Hamiltonian deformations induce canonical isomorphisms on homology $HF(Y, \sigma) := H(CF(Y, \sigma), \partial)$, one can define $HF(Y', \sigma')$ in general by perturbing (Y', σ') to be non-degenerate.

Assume that (Y, σ) is non-degenerate and that $J \in \mathcal{J}^{\text{reg}}(Y, \sigma)$. Because J is translation-invariant, there is a reparametrisation action by \mathbb{R} on the moduli space of trajectories $\mathcal{M}_J(\nu_-, \nu_+)_k$ of fixed index k . This action is free—except on the constant trajectories c_ν —because of the exponential convergence of trajectories. Thus $\mathcal{M}_J(\nu_-, \nu_+)_k/\mathbb{R}$ is a manifold of dimension $k - 1$, assuming $k > 0$; when $k = 0$ it is identified with $\mathcal{H}(Y, \sigma)$, and when $k < 0$ it is empty.

The action functional $\mathcal{A}: \mathcal{M}_J(\nu_-, \nu_+)_1 \rightarrow \mathbb{R}$ decomposes the moduli space as the union of subspaces $\mathcal{M}_J(\nu_-, \nu_+)_{1,\lambda} = \mathcal{A}^{-1}(\lambda)$. Gromov-Floer compactness implies that for any $\lambda_0 \in \mathbb{R}$, the space $\bigcup_{\lambda < \lambda_0} \mathcal{M}_J(\nu_-, \nu_+)_{1,\lambda}/\mathbb{R}$ is compact in any C^k topology. As a 0-manifold, it is then a finite set, and we define $n_j(\nu_-, \nu_+)_\lambda \in \mathbb{Z}/2$ by counting its points modulo 2. We arrange these numbers in a formal series

$$n_J(\nu_-, \nu_+) = \sum_{\lambda \in \mathbb{R}} n_j(\nu_-, \nu_+)_\lambda t^\lambda.$$

This enables us to define the Floer complex over $\Lambda_{\mathbb{Z}/2}$: the boundary operator ∂ on $CF(Y, \sigma) = \bigoplus_\nu \Lambda_{\mathbb{Z}/2}$ is defined by

$$\partial\langle\nu_-\rangle = \sum_{\nu_+ \in \mathcal{H}(Y, \sigma)} n_J(\nu_-, \nu_+) \langle\nu_+\rangle.$$

The fact that $\partial \circ \partial = 0$ is a consequence of a famous argument involving the boundary of the compactification of the index 2 trajectory spaces. There is one loophole which needs to be closed: when $c_{\min} = 1$ and $\nu_- = \nu_+$, a limit of index 2 trajectories could contain a bubble. This can be done using a transversality argument for pseudoholomorphic spheres, the conclusion of which is that, for generic J , the spheres of Chern number 1 sweep out a compact submanifold of Y of codimension $2n - (2c_1 + 2n - 6 + 2) = 2$, which can be made transverse to the horizontal sections, hence disjoint from them (we may assume $n \geq 2$). Such arguments were developed systematically by Hofer and Salamon; a brief discussion appears at the end of the chapter.

The chain map associated to a cobordism (X, π, Ω) from (Y_-, σ_-) to (Y_+, σ_+) is defined from the moduli spaces $\mathcal{M}_J(\nu_-|X|\nu_+)_0$, where $J \in \mathcal{J}^{\text{reg}}(X, \pi; j, J_\pm)$. Here there is no \mathbb{R} -action, and the constant-energy subspaces $\mathcal{M}_J(\nu_-|X|\nu_+)_0$ are already zero-dimensional

and compact. By counting (modulo 2) the points at fixed energy, one gets coefficients $n_J(\nu_-|X|\nu_+) \in \Lambda_{\mathbb{Z}/2}$, which go into the definition of a homomorphism $\Phi_X: CF_{J_-}(Y_-, \sigma_-) \rightarrow CF_{J_+}(Y_+, \sigma_+)$:

$$\Phi_X \langle \nu_- \rangle = \sum_{\nu_+ \in \mathcal{H}(Y_+, \sigma_+)} n_J(\nu_-|X|\nu_+) \langle \nu_+ \rangle.$$

That this is a chain map follows from the structure of the compactification of the index 1 moduli spaces. A family X_t of cobordisms between the same ends induces chain homotopic maps, by an argument with parametrised moduli spaces. The gluing property of the moduli spaces implies that, on homology, the relation $\Phi_{X_1 \circ X_2} = \Phi_{X_1} \circ \Phi_{X_2}$ holds. Crucial examples of cobordisms are deformations of the locally Hamiltonian fibration (Y, σ) and of the complex structures (J, j) . A chain homotopy argument shows that the induced ‘continuation maps’ induce isomorphisms on homology, and that these isomorphisms depend only on the modules they map between. These arguments are common to elliptic Morse homology theories, and to apply them one just needs a theory with good transversality, gluing, and (crucially) compactness properties. Taking orientations into account, they are also valid over $\Lambda_{\mathbb{Z}}$. Here $CF(Y, \sigma) = \bigoplus_{\nu \in \mathcal{H}(Y, \sigma)} \Lambda_{\mathbb{Z}}$, and the coefficients $n_J(\nu_-, \nu_+) = \sum_{\epsilon} n_J(\nu_-, \nu_+)_{\epsilon} t^{\epsilon} \in \Lambda_{\mathbb{Z}}$ are defined using our canonical coherent orientations. These orient the moduli space $\mathcal{M}_J(\nu_-, \nu_+)_1$, but so too does the \mathbb{R} -action. By comparing these two orientations one attaches signs to the points of $\mathcal{M}_J(\nu_-, \nu_+)_1/\mathbb{R}$, and $n_J(\nu_-, \nu_+)$ counts points with these signs. The coherent orientations directly determine coefficients $n_J(\nu_-|X|\nu_+) \in \Lambda_{\mathbb{Z}}$ for use in the cobordism maps. With these definitions, the entire discussion lifts to $\Lambda_{\mathbb{Z}}$.

3.2.3 Canonical relative gradings

The $\mathbb{Z}/2$ -grading on $HF_*(Y, \sigma)_h$ lifts to a canonical *relative $\mathbb{Z}/N_h\mathbb{Z}$ grading*, where $N_h\mathbb{Z} \subset \mathbb{Z}$ is the image of the index homomorphism on $\pi_1(\mathcal{S}(Y), h)$.

Let $H_h \subset H_2(Y; \mathbb{Z})$ be the image of $\pi_1(\mathcal{S}(Y), h)$ under the homomorphism which pushes forward the fundamental homology class of the torus $[S^1 \times S^1]$. Then N_h is specified by the equation

$$\langle 2c_1(T^{\vee}Y, \sigma), H_s \rangle = N_h\mathbb{Z}. \quad (3.17)$$

For example, when h is the class of the constant section of $S^1 \times M \rightarrow S^1$, the subgroup $H_h \subset \mathbb{Z}$ is $\text{im}(\pi_2(M) \rightarrow H_2(M; \mathbb{Z})) \subset H_2(S^1 \times M; \mathbb{Z})$, and $N_h = 2c_{\min}(M, \omega)$.

The origin of these grading properties is the additivity of the index under gluing. The relative degree of $\nu_+, \nu_- \in \mathcal{H}(Y, \sigma)$ is defined to be the index $\text{ind}(u_1)$ of any $u_1 \in \mathcal{M}(\nu_-|Y|\nu_+)$. A different path u_2 may have a different index, but the difference $\text{ind}(u_1) - \text{ind}(u_2)$ is the index of the Cauchy-Riemann operator over the torus obtained by gluing u_1 to $-u_2$. By Riemann-Roch, this index is twice the Chern number of the vector bundle $(u_1 - u_2)^*T^{\vee}Y$, and so lies in $N_h\mathbb{Z}$.

Similar principles apply to the map associated to a cobordism $(X, \pi, \Omega, e_{\pm}, \hat{e}_{\pm})$ from (Y_-, π_-, σ_-) to (Y_+, π_+, σ_+) . If one chooses a homotopy class $k \in \pi_0 \mathcal{S}(X)$, restricting on the ends to $h_{\pm} \in \pi_0 \mathcal{S}(Y_{\pm})$, and chooses absolute (cyclic) gradings on $HF_*(Y_{\pm}, \sigma_{\pm})$ then the map

$$\Phi_{X,k}: HF_*(Y_-, \sigma_-)_{h_-} \rightarrow HF_*(Y_+, \sigma_+)_{h_+}$$

will have a definite degree. However, the same need not be true of the map

$$\Phi_X: HF_*(Y_-, \sigma_-)_{h_-} \rightarrow HF_*(Y_+, \sigma_+)_{h_+},$$

where h_- and h_+ are fixed but not k ; the ambiguity is again given by the possible values of $2c_1(T^v X, \Omega)$ on a certain subgroup of $H_2(X; \mathbb{Z})$.

3.2.4 Quantum cap product

Let (Y, σ) be the mapping torus of an automorphism of (Y, ω) . Then $HF_*(Y, \sigma)$ is a module over the quantum cohomology algebra of $QH^*(M, \omega)$ (which is additively just $H^*(M; \Lambda_{\mathbb{Z}})$). The ‘quantum cap product’ action of $c \in QH^*(M, \omega)$ is obtained by choosing a smooth singular cycle ζ in M , representing $\text{PD}(c)$. One then considers moduli spaces $\mathcal{M}_J(\nu_- | \zeta | \nu_+)$ —the fibre product of the trajectory space over $S^1 \times \mathbb{R}$ with ζ (thought of as a cycle in the fibre over $(1, 0)$). Counting zero-dimensional components gives a chain endomorphism of $CF_*(Y, \sigma)$ of degree $-\deg(c)$. Homologous cycles give chain-homotopic endomorphisms, so one obtains a map $QH^*(M, \omega) \rightarrow \text{End } HF_*(Y, \sigma)$. This is well known to intertwine the quantum product on $QH^*(M, \omega)$, see e.g. Salamon’s notes [37].

There is a similar story for the relative invariant of a cobordism (X, π, Ω) over S ,

$$\Phi_X: HF_*(Y, \sigma) \rightarrow HF_*(\bar{Y}, \bar{\sigma}).$$

If one chooses points $s_1, \dots, s_m \in S$ then the invariant generalises to a map

$$\bigotimes_{k=1}^m QH^*(X_{s_k}, \Omega_{s_k}) \otimes HF_*(Y, \sigma) \rightarrow HF_*(\bar{Y}, \bar{\sigma}),$$

intertwining the quantum module structures of the Floer modules. The map varies continuously as one moves the marked points. Thus a cohomology class $c \in QH^*(X_{s_k}, \Omega_{s_k})$ operates through the image of its Poincaré dual *in the homology of X* .

3.2.5 Open-closed invariants

The Floer homology theory we have described can be enlarged so as to allow Lagrangian boundary conditions.²

²Indeed, one can unite the locally Hamiltonian and Lagrangian Floer homology theories in an open-closed field theory, see Seidel [44].

That is, suppose that S is a surface with cylindrical ends as before, but also with boundary. We are given a simple quadratic fibration $X \rightarrow S$ and a Lagrangian boundary condition $Q \subset X|_{\partial S}$. Mark a number of points $s_a \in S$ (not necessarily interior) and boundary points $\{s'_b\}_{b \in \pi_0 \partial S} \in \partial S$. This results, assuming Q monotone, in a homomorphism

$$HF_*(Y, \sigma) \otimes \bigotimes_a H^*(X_{s_a}) \otimes \bigotimes_b H^*(Q_{s'_b}; \mathbb{Z}) \rightarrow HF_*(\bar{Y}, \bar{\sigma}) \quad (3.18)$$

defined from moduli spaces of index 0, finite action pseudoholomorphic sections with boundary on Q , hitting chosen cycles in the marked fibres X_{s_a} and $Q_{s'_b}$. These homomorphisms satisfy a gluing law under joining cylindrical ends.

Remark 3.2.4. *A word is in order about how these homomorphisms are rigorously constructed, since they are not dealt with explicitly in the literature. The point is that the existing analysis extends with only cosmetic changes. The linearised operators D_u are Fredholm, as is clear from the method of patching parametrices [7]; in the transversality theory, one can take the intersection of regularity conditions coming from the various ends without losing density, by the nature of Baire subsets; the proof of Gromov-Floer compactness, which is largely local in nature, goes over unchanged. In the gluing theory, the estimates required to apply the implicit function theorem depend only on the gluing region itself.*

Remark 3.2.5. (a) *The relative invariant (3.18) varies continuously as one moves the marked points. When two interior points coincide, $s_i = s_j$, the action of the tensor product $c_1 \otimes c_2$ is the same as that of the (ordinary) cup product $c_1 \cup c_2$; similarly for boundary points. These assertions are proved using standard ‘transversality of evaluation’ arguments.*

(b) *An observation which will be useful to us—the result of a typical transversality argument—is that when $s \in \partial S$, the action of $c \in H^*(X_s; \mathbb{Z})$ equals that by $i_!(c) \in H^*(Q_s; \mathbb{Z})$ ($i: Q_s \hookrightarrow X_s$).*

A reduction to homology. The following paragraph needs to be qualified by numerical conditions. These are given in the proposition that follows it.

Take a symplectic manifold (M, ω) with a monotone Lagrangian submanifold L . Let $Y = S^1 \times M$, σ the pullback of ω . Let $X_- = (-\infty, 0] \times Y$, with Lagrangian boundary condition $Q = \{0\} \times S^1 \times L$. Mark one boundary point. This results in a homomorphism

$$\alpha: HF_*(Y, \sigma) \rightarrow H_{n-*}(L; \mathbb{Z})^\vee.$$

Similarly, $X_+ := [0, \infty) \times Y$, with its boundary condition Q and one boundary marked point, gives a map

$$\beta: H_*(L) \rightarrow HF_*(Y, \sigma).$$

There is a well-known isomorphism (of $\mathbb{Z}/2$ -graded, relatively $\mathbb{Z}/2c_{\min}$ -graded \mathbb{Z} -modules)

$$\gamma: H_*(M; \mathbb{Z}) \rightarrow HF_*(Y, \sigma).$$

This is the map obtained from the trivial fibration $M \times \mathbb{C} \rightarrow \mathbb{C}$, equipped with the marked point $0 \in \mathbb{C}$ and an outgoing cylindrical end.

Proposition 3.2.6. *Let $i: L \hookrightarrow M$ be the inclusion map. Assume that $n < \mu_{\min}$. We have $\gamma^{-1} \circ \beta = i_*$ and*

$$\langle \alpha \circ \gamma(x), y \rangle = i^! x \cap y \quad (x \in H_{2n-k}(M; \mathbb{Z}), y \in H_k(L; \mathbb{Z})).$$

Proof. Consider smooth singular cycles

$$Z_x = \prod_{j=1}^p \Delta^{2n-k} \xrightarrow{\tilde{x}} M, \quad Z_y = \prod_{j=1}^q \Delta^k \xrightarrow{\tilde{y}} L$$

representing x and y respectively. By the gluing property of the invariants, $\alpha \circ \gamma$ can be computed using the trivial fibre bundle $M \times \bar{D} \rightarrow \bar{D}$ and Lagrangian boundary condition $L \times \partial \bar{D}$. Indeed, $\langle \alpha \circ \gamma(x), y \rangle$ is the signed count of points in the zero-dimensional component a regular moduli space

$$\mathcal{M}(\tilde{x}, \tilde{y}) := \mathcal{M}_J(M \times \bar{D}, L \times \partial \bar{D}) \times_{(\text{ev}, \tilde{x} \times \tilde{y})} (Z_x \times Z_y).$$

The dimension of this space, near a disk of Maslov index μ , is $\mu - k - (n - k) = \mu - n$.

We need to make a Hamiltonian perturbation of σ to achieve non-degeneracy. A convenient choice is to take a function $H: S^1 \times M \rightarrow \mathbb{R}$ which is given, in a neighbourhood of L (identified with a disk-bundle in $S^1 \times T^*L$) by $H(t, \alpha) = -|\alpha|^2$. Put $\sigma' = \sigma + d(Hdt)$, and notice that $S^1 \times L$ is then still a Lagrangian boundary condition. Take J to be the pullback of an almost complex structure on M compatible with ω . Then elements of $\mathcal{M}_J(M \times \bar{D}, L \times \partial \bar{D})$, of Maslov index μ , are pseudoholomorphic disks $v: (\bar{D}, \partial \bar{D}) \rightarrow (M, L)$, also of index μ . When the expected dimension of $\mathcal{M}(\tilde{x}, \tilde{y})$ is zero, we have $\mu = n$; but that implies that v is constant, and hence that the map $\text{ev}: \mathcal{M}(\tilde{x}, \tilde{y}) \rightarrow L$ can be identified with the cycle $\tilde{x} \cap \tilde{y}$. Since this moduli space is naturally a manifold of the expected dimension, it is reasonable to suspect that its elements are regular (the deformation operator D_v surjects). Standard computations show that this indeed so. It follows that $\alpha \circ \gamma$ is as claimed.

The argument for $\gamma^{-1} \circ \beta$ is similar.

□

Künneth isomorphism. There is a standard Künneth formula in Floer homology:

$$HF_*(Y_1 \times_{S^1} Y_2, \sigma_1 \oplus \sigma_2) \cong H_*(CF(Y_1, \sigma_1) \otimes CF(Y_2, \sigma_2)).$$

For the proof one uses almost complex structures compatible with the product structure, and observes that these are regular precisely when the factors are (see e.g. [16]). It is then a matter of checking that the obvious additive isomorphism $CF(Y_1 \times Y_2, \sigma_1 \oplus \sigma_2) \cong CF(Y_1, \sigma_1) \otimes CF(Y_2, \sigma_2)$ is a chain map. The Künneth formula for Floer homology is compatible with that in quantum cohomology.

Another general property of Floer homology is the ‘Poincaré duality’ isomorphism

$$HF_*(Y_1, -\sigma_1) \cong HF_*(Y_1, \sigma_1)^\vee.$$

Combining this with the Künneth isomorphism, and working over a Novikov field $\Lambda_{\mathbb{F}}$, we get an isomorphism of $\mathbb{Z}/2$ -graded modules

$$HF_*(Y_1 \times_{S^1} Y_2, -\sigma_1 \oplus \sigma_2) \cong \text{Hom}(HF_*(Y_1, \sigma_1), HF_*(Y_2, \sigma_2)).$$

Remark 3.2.7. Suppose that $\Lambda \subset Y_1 \times_{S^1} Y_2$ is a Lagrangian boundary condition, for the form $-\sigma_1 \oplus \sigma_2$, inducing $\Phi(\Lambda) \in \text{Hom}(HF_*(Y_1), HF_*(Y_2))$. This is related to the quantum cap product in the following way. Let M_1, M_2 , denote the fibres over a point in S^1 of Y_1, Y_2 . Let $i: L \hookrightarrow M_1 \times M_2$ be the fibre of Λ . Take $c_1 \in H^*(M_1; \mathbb{Z})$ and $c_2 \in H^*(M_2; \mathbb{Z})$, and suppose that $i^*(\text{pr}_1^* c_1 - \text{pr}_2^* c_2) = 0$. Then

$$c_2 \circ \Phi(\Lambda) = \Phi(\Lambda) \circ c_1. \quad (3.19)$$

The left-hand side comes from a moduli space of sections which hit a cycle of form $M_1 \times \zeta_2$ in some fibre. We can take this to be a boundary fibre, and insist that the sections pass through $L \cap (M_1 \times \zeta_2)$. Similarly we can compute the right-hand side via sections which hit a cycle $L \cap (\zeta_1 \times M_2)$. Since the two cycles in $L \cap (M_1 \times M_2)$ are homologous, the chain maps underlying the homomorphisms on the two sides are chain homotopic.

A small generalisation. Suppose that we have a simple quadratic fibration $\pi_1: X_1 \rightarrow S_1$ with one outgoing boundary component $C_{1,\text{out}}$ and incoming cylindrical ends $C_{1,\text{in}}^i$. A Lagrangian boundary condition $Q_1 \rightarrow C_{1,\text{out}}$ in the boundary component $(Y_{1,\text{out}}, \sigma_1)$ induces a homomorphism

$$\Phi_{Q_1}: \bigotimes_i HF(Y_{1,i}, \sigma_{1,i}) \rightarrow \Lambda_{\mathbb{Z}}.$$

where the $(Y_{1,i}, \sigma_{1,i})$ are the cylindrical ends of X_1 . Similarly, we can consider $\pi_2: X_2 \rightarrow S_2$, where S_2 has one incoming boundary $C_{2,\text{in}}$ and outgoing cylindrical ends $C_{2,\text{out}}^j$; a Lagrangian boundary condition $Q_2 \rightarrow C_{2,\text{out}}$ in $(Y_{2,\text{in}}, \sigma_2)$ gives rise to a homomorphism

$$\Phi_{Q_2}: \Lambda_{\mathbb{Z}} \rightarrow \bigotimes_j HF(Y_{2,j}, \sigma_{2,j}).$$

Now, if we parametrise $C_{1,\text{out}}$ and $C_{2,\text{in}}$ then we can consider the product $Q_1 \times_{S^1} Q_2 \subset Y_{1,\text{out}} \times_{S^1} Y_{2,\text{in}}$. This is isotropic with respect to the form $-\sigma_1 \oplus \sigma_2$. We can compute

the composite $\Phi_{Q_2} \circ \Phi_{Q_1}$ from the fibre product moduli space $\mathcal{M}(X_1, Q_1) \times_{S^1} \mathcal{M}(X_2, Q_2)$. An alternative point of view involves the *opposite* surface $-S_1$, which has a distinguished *incoming* boundary component. Over it we have $-X_1 \rightarrow -S_1$; we can think of the fibre product $\mathcal{M}(X_1, Q_1) \times_{S^1} \mathcal{M}(X_2, Q_2)$ as the space of pairs of sections $(u_1, u_2) \in \mathcal{M}(-X_1, Q_1) \times \mathcal{M}(X_2, Q_2)$ with boundary values in $Q_1 \times_{S^1} Q_2$.

We can also do something a little more general, and consider a ‘Lagrangian boundary condition’ $Q \subset Y_{1,\text{out}} \times_{S^1} Y_{2,\text{in}}$: a middle-dimensional, isotropic subbundle which need not be a fibre product. This no longer fits exactly into the framework we have been considering.³ However, exactly the same arguments as before yield a sensible moduli space $\mathcal{M}_{J_1, J_2}(Q)$ of pairs (u_1, u_2) of pseudoholomorphic sections of $(-X_1, -J_1)$ and (X_2, J_2) with boundary values in Q . These result in an invariant

$$\Phi_Q: \bigotimes_i HF(Y_{1,i}, \sigma_{1,i}) \rightarrow \bigotimes_j HF(Y_{2,j}, \sigma_{2,j})$$

with the proviso, as before, that we should work over a Novikov *field*. The gluing theory shows that this is the composite of three maps of a more standard kind,

$$\bigotimes_i HF(Y_{1,i}, \sigma_{1,i}) \xrightarrow{\Phi_{X_1}} HF(Y_{1,\text{out}}) \rightarrow HF(Y_{2,\text{in}}) \xrightarrow{\Phi_{X_2}} \bigotimes_j HF(Y_{2,j}, \sigma_{2,j}),$$

where the middle one is deduced (by the Künneth procedure above) from the invariant associated to the inclusion of Q in $Y_{1,\text{out}} \times_{S^1} Y_{2,\text{in}}$.

Example 3.2.8. Suppose that S is split along a circle C , so $S = S_1 \cup_C S_2$. We consider a fibration $(X, \pi, \omega) \rightarrow S$, restricting to (Y, π', σ) over C . Then in the product $(Y \times_C Y, -\sigma \oplus \sigma)$ one has a *diagonal* Lagrangian boundary condition Δ , and an invariant

$$\Phi(\Delta): \bigotimes_i HF(Y_{1,i}, \sigma_{1,i}) \rightarrow \bigotimes_j HF(Y_{2,j}, \sigma_{2,j}).$$

This map is the same as the relative invariant of X itself. This can be seen as an instance of the gluing law proved in [46].

3.2.6 Weakly monotone fibres

In the framework of fibred monotonicity, bubbling in low-dimensional moduli spaces (with regularity assumptions in force) is ruled out on the grounds that the principal component of a limit curve must have non-negative index. A more refined method incorporates transversality for pseudoholomorphic spheres in the fibres and disks in the boundary fibres. One can then determine numerical conditions under which moduli spaces of sections are generically

³Except where $-S_1 \cong S_2$; in that case, Q is a Lagrangian boundary condition for a fibre product.

disjoint from those of fibrewise spheres and disks. These ‘weak monotonicity’ conditions (in the form that appears in setting up Floer homology and its invariance) were determined by Hofer and Salamon [16].

Definition 3.2.9. A compact symplectic manifold (M^{2n}, ω) is **weakly monotone** if any $S \in \pi_2(M)$ satisfying $\int_S \omega > 0$ and $c_1(S) < 0$ also satisfies $c_1(S) + n - 3 < 0$.

An equivalent condition is that (M, ω) is monotone, or $c_{\min} = 0$, or $c_{\min} \geq n - 2$.

Generic almost complex structures $J \in \mathcal{J}(M, \omega)$ are regular in the sense that for every element $\beta \in \mathcal{M}_J^{\text{si}}(M)$ of the moduli space of parametrised, *simple* J -holomorphic spheres, the relevant linearised operator D_β is surjective (‘simple’ means that the set of injective points, $\{z \in S^2 : \beta^{-1}(\beta(z)) = \{z\}\}$, is non-empty; it is then open and dense in S^2). For such J , $\mathcal{M}_J^{\text{si}}(M)$ is a manifold of local dimension $2(c_1(\beta) + n)$, and its quotient by the free action of $\text{Aut}(S^2) = \text{PSL}(2, \mathbb{C})$ has dimension $2(c_1(\beta) + n - 3)$. Thus there are no simple spheres of negative Chern number, and the same holds in 1-parameter families. A lemma of McDuff asserts that every pseudoholomorphic sphere factors through a simple one, so there are no spheres of negative Chern number at all.

(a) Suppose (Y, π, σ) is non-degenerate with fibre (M, ω) . To construct $HF_*(Y, \sigma)$ one uses almost complex structures $J \in \mathcal{J}^{\text{reg}}(Y, \pi)$ with the property that no fibre contains a J -sphere β such that (i) $c_1(\beta) < 0$; or (ii) $c_1(\beta) = 0$ which meets the image of one of the trajectory spaces $\mathcal{M}_J(\nu_-, \nu_+)_1$; or (iii) $c_1(\beta) \leq 1$ which meets the image of one of the trajectory spaces $\mathcal{M}_J(\nu_-, \nu_+)_2$. We have just seen that (i) holds generically; a ‘transversality of evaluation’ argument with somewhere injective spheres shows that (ii) and (iii) do as well. For such J , the trajectory spaces of index ≤ 2 are easily seen to be compact.

(b) To obtain the map Φ_X associated to a locally Hamiltonian fibration with cylindrical ends, one uses $J \in \mathcal{J}^{\text{reg}}(X, \pi; J_\pm, j)$ with the property that no fibre contains a J -sphere β such that (i’) $c_1(\beta) < 0$, or (ii’) $c_1(\beta) = 0$ meeting the index 0 trajectory spaces.

(c) Simple quadratic fibrations can be accommodated using ideas from Seidel’s thesis [42]. One needs to make some hypothesis about the normalisation \widetilde{M}_0 of a singular fibre M_0 . The almost complex structure $J_0 = J|_{M_0}$ on M_0 lifts to one \widetilde{J}_0 on \widetilde{M}_0 , integrable near the preimage of the normal crossing divisor; it is sufficient to assume that there \widetilde{J}_0 is compatible with a weakly monotone symplectic form on \widetilde{M}_0 .

The point is that any J_0 -sphere in M_0 lifts uniquely to a \widetilde{J}_0 -sphere in \widetilde{M}_0 . One cannot expect \widetilde{J}_0 to be regular; however, for generic J , lifts of non-constant spheres in M_0 are regular, hence have non-negative Chern number: here we use a standard argument which shows that one can achieve regularity for all spheres passing through an open set U by

making a perturbation of the almost complex structure supported in U . Another generic condition on J is that the spheres in M_0 with zero Chern number should not hit the zero-index trajectories in X .

Lagrangian boundary conditions As with spheres, transversality theory for pseudo-holomorphic disks $\delta: (\bar{D}, \partial\bar{D}) \rightarrow (M, L)$ works in a straightforward way for simple disks—those for which the set of injective points is non-empty, or equivalently open and dense—but the complication is that not every non-constant δ factors through a simple one. One can, however, use Lazzarini’s lemma [25]: *if δ is a non-constant J -disk then there is a simple J -disk δ' such that $\delta'(\partial\bar{D}) \subset \delta(\partial\bar{D})$* . (A more precise version is due to Kwon and Oh.)

Assumption: M is weakly monotone. There is a constant λ such that $\mu_L(S) = \lambda \int_S \omega$ for all $S \in \pi_2(M, L)$. Either $\lambda \geq 0$, or $\mu_{\min} > n - 2$.

Now take a simple quadratic fibration with Lagrangian boundary condition (X, Q) such that the boundary fibres are symplectomorphic to (M, L) ; X may also have cylindrical ends. Assume that the normalisations of the singular fibres have weakly monotone symplectic forms. We want to see that, for generic regular almost complex structures, the *index zero* components of the moduli space $\mathcal{M}_J(X, Q)$ are compact. There are three possibilities.

- $\lambda > 0$. This is the monotone case; since non-trivial pseudoholomorphic disks have positive Maslov index, they cannot occur as bubbles.
- $\lambda < 0$. Lazzarini calls this case **strongly negative**. The index formula that the moduli space of simple disks $\mathcal{M}_J^{\text{si}}(M, L)$ has virtual dimension $\mu + n < 2$. Since $\text{Aut}(\bar{D})$ is 3-dimensional and acts freely, we conclude that, for generic $J \in \mathcal{J}(TM, \omega)$ the moduli space of simple disks is empty, and that this remains true in 1-parameter families (in particular, for the fibres of Q). Lazzarini’s lemma allows us to remove ‘simple’ from this conclusion.

Similar arguments apply when we consider 1-dimensional moduli spaces for 1-parameter families of almost complex structures (or other deformations of the data). The integers obtained by counting points in the zero-dimensional moduli spaces are therefore invariants of the Lagrangian boundary condition.

- $\lambda = 0$. *I do not know of a reference which precisely fills in the details of the following sketch, though doing so should be a matter of making minor modifications to arguments from the literature. When we consider situations of this kind we will only draw tentative conclusions and not claim theorems.*

Here $\mathcal{M}_J^{\text{si}}(M, L)$ has dimension n . The manifold $\mathcal{M}_J^{\text{si}}(M, L) \times_{\text{Aut}(\bar{D})} \partial \bar{D}$ has dimension $n - 2$, and this gives just enough leeway to arrange that the zero-dimensional spaces of sections of (X, Q) do not hit those of (simple) boundary bubbles.

We can expect to encounter isolated boundary bubbles when we consider a 1-parameter moduli space, and for this reason we establish invariance using a continuation argument instead. That is, we consider the index zero moduli spaces for a locally Hamiltonian fibration over the cylinder, with Lagrangian boundary conditions at the two ends; we can arrange that these do not hit pseudoholomorphic disks (or spheres). We then need to invoke the gluing theorem from [46] to conclude that we have a *bona fide* invariant.

Chapter 4

Lagrangian matching invariants

Our task now is to put together the methods of chapters 2 and 3 so as to obtain a field theory for near-symplectic broken fibrations. One part of this—introducing the Floer homology groups for mapping tori of symmetric products—can be done immediately. The next stage is to bring in relative Hilbert schemes so as to obtain relative invariants for Lefschetz fibrations, the relative version of the invariants defined in Smith [49]. We then state a theorem on the existence and properties of certain Lagrangian boundary conditions for the relative Hilbert schemes, and use these to construct the Lagrangian matching invariant for broken fibrations and its relative version. We show that this invariant can be organised in a way that makes it resemble the Seiberg-Witten invariant (that is, we show how its natural ‘topological sectors’ correspond to Spin^c -structures, and describe how the gradings work), and compute it in some simple cases.

Constructing the Lagrangian boundary condition involves looking carefully at the geometry of the Hilbert scheme. The crucial point, which is an observation of Smith’s [49], is that the normal crossing divisor in the Hilbert scheme $\text{Hilb}^r(C)$ of r points on the nodal curve C is isomorphic to the $(r - 1)$ th symmetric product of its normalisation.

The computations we carry out here are deduced from the formal structure of the theory, using general properties of Floer homology. A task for future research is to calculate the invariants in cases which involve more interesting moduli spaces. A specific goal is to prove a conjecture which we shall state, that a certain triangle of maps associated to certain cobordisms is exact. This triangle, and the plan of attack for proving it, are akin to Seidel’s exact triangle [46].

Remark 4.0.10. *The reader may wish to refer to the Introduction, where the main formal properties of the invariants were set out.*

4.1 Floer homology for relative symmetric products

Quantum cohomology

The small quantum cohomology of $\text{Sym}^r(\Sigma)$ was calculated by Bertram and Thaddeus [4].¹ Their results are complete except for third-order corrections in the range $r \in [\frac{3}{4}g, g-1)$, but we shall only state them in the simpler ‘weakly monotone’ range.

For any $r > 1$, the image of the Hurewicz homomorphism $\pi_2(\text{Sym}^r(\Sigma)) \rightarrow H_2(\Sigma; \mathbb{Z})$ is an infinite cyclic group generated by a pencil S in a fibre of the Abel-Jacobi map (cf. Appendix A). The Chern number of S is $r - g + 1$. Its area, with respect to a chosen Kähler form, is denoted by A . We regard the quantum cohomology $(QH^*(\text{Sym}^r(\Sigma)), *)$ as a graded $\mathbb{Q}[q]$ -algebra, where q has degree $2(r + 1 - g)$. Its underlying module is

$$QH^*(\text{Sym}^r(\Sigma)) = H^*(\text{Sym}^r(\Sigma); \mathbb{Q}) \otimes_{\mathbb{Q}} \mathbb{Q}[[t^A]], \quad \deg(t^A) = 2(r + 1 - g).$$

We think of $\mathbb{Q}[[t^A]]$ as a subring of $\Lambda_{\mathbb{Q}}$; q acts by scalar multiplication by t^A .

We shall *exclude* the range

$$(g + 1)/2 < r < g - 1, \tag{4.1}$$

which means that we can work over \mathbb{Z} rather than \mathbb{Q} , that is, we consider the $\mathbb{Z}[q]$ -module

$$QH^*(\text{Sym}^r(\Sigma); \mathbb{Z}) = H^*(\text{Sym}^r(\Sigma); \mathbb{Z}) \otimes_{\mathbb{Z}} \mathbb{Z}[[t^A]]. \tag{4.2}$$

From the point of view of algebraic geometry, the interval (4.1) is the most interesting (and the most complicated) but it remains to be seen whether its intricacies are needed in studying fibred 4-manifolds.

A general fact, proved in [4], is that quantum product with any class in H^1 is undeformed. Consequently, for any $c \in \Lambda^* H^1(\Sigma; \mathbb{Z}) \subset H^*(\text{Sym}^r(X); \mathbb{Z})$, the endomorphism

$$c * - \in \text{End}_{\mathbb{Z}[[q]]} QH^*(\text{Sym}^r(\Sigma); \mathbb{Z})$$

is equal to the classical one $c \cup -$, extended $\mathbb{Z}[q]$ -linearly. To give the quantum product, it therefore suffices to describe the degree 2 endomorphism $\eta * -$.

Let θ_i denote the *integral* class $\theta^i/i!$ when $i \geq 0$, and 0 when $i < 0$.

Theorem 4.1.1 (Bertram-Thaddeus [4]). *The submodule of the $\mathbb{Z}[q]$ -module*

¹For related computations in Seiberg-Witten Floer cohomology, see Muñoz-Wang [33]. They consider a version of the theory in which the vector space associated with $S^1 \times \Sigma$ is identified with $H^*(\text{Sym}^r(\Sigma); \mathbb{Q})$. The pair-of-pants product deforms the classical ring structure. Muñoz-Wang give a presentation for the Floer cohomology ring, but it is an open problem to determine how it deforms the classical product.

$QH^*(\text{Sym}^r(X); \mathbb{Z})$ spanned by the classes $\{\theta^m \eta^n\}_{m,n \geq 0}$ is invariant under $\eta * -$, and

$$\eta * (\eta^n) = \begin{cases} \eta^{n+1}, & r \leq (g+1)/2; \\ \eta^{n+1} + \frac{t^A}{1-t^A} \sum_{i=0}^1 (\theta_{i+n} \eta^{1-i} - \theta_i \eta^{1+n-i}), & r = g-1; \\ \eta^{n+1} + t^A (\theta_{g-r+n} - \theta_{g-r} \eta^n), & r \geq g. \end{cases}$$

Floer homology of symmetric products: general form

Consider the relative symmetric product $Y_r = \text{Sym}_{S^1}^r(Y)$ of a proper surface-bundle $Y \rightarrow S^1$ with fibre Σ (equipped with $j \in T^\vee Y$). After choosing a locally Hamiltonian structure Ω compatible with the complex structure $\text{Sym}^r(j)$ on $T^\vee Y_r$, Floer homology is defined, and has the following basic format. Again, we exclude the interval (4.1), though it can be included if one works over \mathbb{Q} (see [37]).

- For $r > 1$, the set of topological sectors, $\pi_0 \mathcal{S}(Y_r)$, is canonically identified with $H_1(Y; \mathbb{Z})_r$, the affine subgroup of $H_1(Y; \mathbb{Z})$ of classes h with $h \cdot [\Sigma] = r$.
- $HF_*(Y_r, \Omega, h)$ is a $\mathbb{Z}/2$ -graded, relatively $\mathbb{Z}/2(r+1-g)$ -graded $\Lambda_{\mathbb{Z}}$ -module.
- Quantum cap product makes $HF_*(Y_r, \Omega, h)$ a module over $QH_*(\text{Sym}^r(\Sigma); \mathbb{Z})$. For $x \in H^*(\text{Sym}^r(\Sigma); \mathbb{Z})$, $\nu \in HF_*(Y_r, \Omega, h)$,

$$x \cdot \nu = \sum_{n \geq 0} t^{An} (x \cdot \nu)_n,$$

where $(x \cdot \nu)_n \in HF_{\deg(\nu) - 2n(r+1-g)}(Y_r, \Omega, h)$.

Floer homology of symmetric products: special form

We now look more closely at the Floer homology of symmetric products in the case where Ω is induced by $\sigma \in Z_Y^2$ via the vortex construction. That is, we consider a form $\sigma_{L,\tau}$. Recall that its cohomology class is independent of the choice of $L \in H^2(Y; \mathbb{Z})$:

$$[\sigma_{L,\tau}] = 2\pi\tau[\sigma]^{(1)} - 2\pi^2 1^{(2)}.$$

Theorem 4.1.2. *Assume $r \geq g-1$. (a) The Floer homology module $HF_*(Y_r, \sigma_{L,\tau})$ is determined, up to unique isomorphism, by the cohomology class $[\tau\sigma] \in H^2(Y; \mathbb{R})$.*

(b) Assume that the fibres of $Y \rightarrow S^1$ have genus $g \geq 2$. One obtains an invariant of a mapping class, $HF_([\phi], r)$, by taking the class $[\tau\sigma] \in H^2(T(\phi); \mathbb{R})$ to be proportional to $c_1(T^\vee T(\phi))$.*

Proof. (a) This is more-or-less immediate from Theorem 2.3.12 and its corollary 2.3.13. We can consider $HF_*(Y_r, \Xi)$ for any closed, fibrewise Kähler 2-form Ξ representing the class

$[\sigma_{L,\tau}]$. Because Ξ is chosen from within a convex space, the method of continuation maps gives quasi-isomorphisms between the chain complexes $CF_*(Y_r, \Xi)$, $CF_*(Y_r, \Xi')$, and chain homotopies between the different possible maps between them. Note that the choice of complex structure $j \in T^*Y$ doesn't matter either, because that too varies in a contractible space.

(b) This is a simple generalisation of Seidel's argument from [44] which deals with the case $r = 1$.

Fix a symplectic form $\omega \in \Omega_{\Sigma}^2$. The space $\text{Diff}^+(\Sigma)$ deformation-retracts onto the subspace $\text{Aut}(\Sigma, \omega)$ of volume-preserving diffeomorphisms by a result of Moser. But when $g(\Sigma) \geq 2$, the components of $\text{Diff}^+(\Sigma)$ are contractible (this is a famous theorem of Earle and Eells). Thus there is a well-defined homomorphism

$$\text{Aut}(\Sigma, \omega) \xrightarrow{\text{Flux}} H^1(\Sigma; \mathbb{R}),$$

and $\text{Flux}^{-1}(a) \cap \text{Aut}_0(\Sigma, \omega)$ is a contractible subspace of $\text{Aut}_0(\Sigma, \omega)$. On the other hand, by choosing the flux suitably, one can find a locally Hamiltonian structure σ on $T(\phi)$ which realises any chosen cohomology class extending $[\omega]$. An obvious choice is to require that $[\sigma]$ be proportional to $c_1(T^*T(\phi))$, though there are other possibilities (one can make a different choice for each topological sector $h \in H_1(Y; \mathbb{Z})$, as is done by Hutchings and Sullivan [19]).

As in (a), the auxiliary choices involved in the definition of the Floer homology come from a contractible space, so it is well-defined up to unique isomorphism. \square

There is one way in which the situation is different from that where $r = 1$, and that concerns the periods of the action functional \mathcal{A} . We consider its periods on the group

$$P := \ker (c_1(T^*Y_r) : \pi_1(\mathcal{S}(Y_r), h) \rightarrow \mathbb{Z}).$$

By Lemma 2.1.1, $c_1(T^*Y_r) = (c^{(1)} + 1^{(2)})/2$, where $c = c_1(T^*Y) \in H^2(Y; \mathbb{Z})$. Hence each $p \in P$ satisfies

$$\mathcal{A}(p) = 2\pi \langle (\tau[\sigma] + \pi c)^{(1)}, p \rangle \tag{4.3}$$

$$= \pi \langle 2\tau[\sigma] + 2\pi c, p_{(1)} \rangle. \tag{4.4}$$

When $r = 1$, an additional reason for choosing $[\sigma]$ proportional to $c_1(T^*Y)$ is that in that case the periods vanish on P , and as a consequence the Floer homology is finitely generated over \mathbb{Z} . For $r \geq g - 1$ it does not seem to be possible to choose $\tau[\sigma]$ in such a way that the periods vanish.

Remark 4.1.3. *I have not established to what extent HF_* depends on the class $\tau[\sigma]$, but there are some general principles to bear in mind (compare e.g. Seidel-Smith [47]). Floer*

homology is the Morse homology of an exact 1-form defined on some covering space of $\mathcal{S}(Y)$, and one can expect this to be independent of exact perturbations. To prove this one has to replace the continuation-map technique by more sophisticated analysis of parametrised moduli spaces, studying birth-death processes for critical points. The Floer homology should therefore be independent of variations in the cohomology class which do not affect the sign of the action $\mathcal{A}(p)$, $p \in P$. One can look for wall-crossing phenomena where the action of some class $p \in P$ changes sign.

Definition 4.1.4. Write

$$HF_{\bullet}(Y, \pi, \sigma, r) = \bigoplus_{\gamma \in H_1(Y, \mathbb{Z})_r} HF_{\bullet}(Y, \pi, \sigma, \gamma)$$

where $HF_{\bullet}(Y, \pi, \sigma, \gamma)$ denotes the Floer homology, over $\Lambda_{\mathbb{Z}}$, of $(Y_r, \pi_r, \sigma_{PD(\gamma), \tau})$, for any (sufficiently large) $\tau \in \mathbb{R}_{>0}$. Note that we use the degree r line bundle dual to the class γ .

Abbreviate to $HF_{\bullet}(Y, \pi, r) = \bigoplus_{\gamma \in H_1(Y, \mathbb{Z})_r} HF_{\bullet}(Y, \pi, \gamma)$ the special case where $[\sigma]$ is proportional to $c_1(T^{\vee}Y)$.

These Floer homology groups are the ones discussed in the Introduction. The reason for the precise form of the conjecture made there, relating them to monopole Floer groups with a particular twisting, is that the symplectic and monopole groups then have identical periods. That is, $\pi_1(\mathcal{S}(Y), \gamma)$ is identified with π_1 of the Seiberg-Witten configuration space, and, up to a positive factor, the two action 1-forms represent the same cohomology class (thought of as a homomorphism $\pi_1 \rightarrow \mathbb{R}$).

The relative Hilbert scheme. Let $(X, \pi, \Omega, J_0, j_0)$ be a Lefschetz fibration (i.e. a four-dimensional simple quadratic fibration) over a compact surface S . Choose a vertical complex structure $J^v \in \mathcal{J}(T^{\vee}X^*, \Omega)$ which agrees with J_0 where both are defined. We can then form the relative Hilbert scheme of r points

$$\pi_r: \text{Hilb}_S^r(X) \rightarrow S.$$

The fundamental facts about it are these:

- $\text{Hilb}_S^r(X)$ is naturally a smooth manifold. (This is an important discovery of Donaldson and Smith [9].)
- There are induced (integrable) complex structures in neighbourhoods of $\text{crit}(\pi_r)$ and on the fibres of π_r .
- Over the regular values S^* of π there is a natural identification

$$\text{Hilb}_{S^*}^r(X^*) = \text{Sym}_{S^*}^r(X^*)$$

respecting the maps to S^* and the complex structures.

- The map π_r satisfies the conditions on the map in a simple quadratic fibration.
- There are cohomology operations

$$H^*(X) \rightarrow H^{*+2k-2}(\text{Hilb}_S^r(X)), \quad c \mapsto c^{(k)}.$$

extending the operations $H^*(X^*) \rightarrow H^{*+2k-2}(\text{Sym}_{S^*}^r(X^*))$ of (2.2).

- Suppose S is a disk \bar{D} , $E \rightarrow \bar{D}$ an elementary Lefschetz fibration (i.e. it has just one critical point, lying over $0 \in \bar{D}$), and $\beta \in H^2(E; \mathbb{Z})$ a class such that $\langle \beta, [E_s] \rangle = r > 0$ and $\langle \beta, \alpha \rangle \geq 0$ for each class $\alpha \in H_2(E; \mathbb{Z})$ represented by a component of a fibre E_s . Then $\beta^{(1)}$ is represented by a 2-form ξ which makes the Hilbert scheme a simple quadratic fibration.

Continuing with the last point, let $\bar{D}' \subset \bar{D}$ be a strictly smaller disk and $E' \subset E$ its preimage. We can construct a 2-form χ on $\text{Hilb}_{\bar{D}}^r(E)$ which, like ξ , makes it into a simple quadratic fibration, but which is also standard on the complement of $\text{Hilb}_{\bar{D}'}^r(E')$:

- If $\omega \in Z_E^2$ is positive on the non-singular fibres then we can arrange that χ is equal to $\omega_{\beta, \tau}$ off $\text{Hilb}_{\bar{D}'}^r(E')$, provided that $\omega\tau$ is large. Moreover, such a χ is determined uniquely up to deformation.

This can be achieved by a simple application of the Thurston-Gompf method. Notice first that $[\omega](E \setminus E') = \lambda\beta|(E \setminus E')$ for some $\lambda > 0$. We look for a form representing the class $2\pi(\lambda\tau\beta^{(1)} - \pi 1^{(2)})$. It is clear from its construction that $1^{(2)}$ is represented by a 2-form ξ' which is of type $(1, 1)$ near $\text{crit}(\pi_r)$ and on the fibres. Choose a 1-form α such that $\omega_{\beta, \tau} - 2\pi(\lambda\tau\xi - \pi\xi') = d\alpha$, and a smooth function $\rho: \bar{D} \rightarrow [0, 1]$ which is zero on $\bar{D} \setminus \bar{D}'$ and identically 1 on a smaller disk $\bar{D}'' \subset \bar{D}'$. Put

$$\chi = \begin{cases} \omega_{\beta, \tau} - d(\rho\alpha), & \text{over } \bar{D} \setminus \bar{D}'', \\ 2\pi(\lambda\tau\xi - \pi\xi'), & \text{over } \bar{D}''. \end{cases} \quad (4.5)$$

When $\lambda\tau$ is large enough, this form is Kähler near $\text{crit}(\pi_r)$ and on the fibres, as required.

The uniqueness up to deformation follows by linear interpolation, bearing in mind that spaces of Kähler forms are convex.

Grading. In the Introduction we claimed that the Floer homology of symmetric products has identical relative grading properties to the monopole Floer theory. We now make good this claim, as far as the modules themselves are concerned.

A generator $\nu \in \mathcal{H}(Y_r, \sigma_{L, \tau})$ of the Floer complex $CF_*(Y_r, \sigma_{L, \tau})$ gives rise to a closed subset ('multisection') $\bar{\nu}$ of Y which represents a 1-cycle. A simple transversality argument shows that a small perturbation ν_0 of ν will generically miss the diagonal, and moreover

that any other such perturbation ν_1 will be homotopic to ν_0 outside the diagonal. The 1-cycle $\bar{\nu}_0$ is a disjoint union of embedded loops in Y , positively transverse to the fibres, and homotopic to the loops in $\bar{\nu}_1$. By a procedure which we explained in the introduction, these loops determine a homotopy class of oriented 2-plane fields ξ_ν . Its underlying Spin^c -structure is $\mathfrak{t}_{[\bar{\nu}]}$, which depends only on the homotopy class $[\nu] \in \pi_0 \mathcal{S}(Y_r)$.

Let us write $\gamma = [\nu]$, and \mathfrak{t}_γ for the Spin^c -structure $\mathfrak{t}_{[\bar{\nu}]}$. Thus $\xi_\nu \in J(Y, \mathfrak{t}_\gamma)$, which is a transitive \mathbb{Z} -set. The stabiliser $\langle c_1(\mathfrak{t}_\gamma), H_2(Y; \mathbb{Z}) \rangle$ is equal to $2(r+1-g)\mathbb{Z}$. This is because $H_2(Y; \mathbb{Z})$ is generated by the fibre class $[\Sigma]$ (which satisfies $\langle c_1(\mathfrak{t}_\gamma), [\Sigma] \rangle = 2(r+1-g)$) and by tori on monodromy-invariant classes in $H_1(\Sigma; \mathbb{Z})$, which evaluate as zero on $c_1(\mathfrak{t}_\gamma)$. Thus the monopole Floer homology $HM_\bullet(Y, a, \mathfrak{t}_\gamma)$, like $HF_*(Y_r, \sigma, \gamma)$, is relatively $\mathbb{Z}/2(r+1-g)$ -graded.

We can sharpen this, however.

Let $\widehat{\mathcal{S}(Y_r)}$ denote the covering space of $\mathcal{S}(Y_r)$ whose points $\tilde{\nu}$ consist of a section $\nu \in \mathcal{S}(Y_r)$ and a homotopy class of symplectic trivialisations of $\nu^* T^\vee Y_r$. Consider a path u in $\mathcal{S}(Y_r)$, from ν_- to ν_+ . Choose a lift $\tilde{\nu}_-$ to $\widehat{\mathcal{S}(Y_r)}$, and let $\tilde{\nu}_+$ be the lift of ν_+ obtained by lifting the path u , starting at $\tilde{\nu}_-$.

Lemma 4.1.5. $\xi_{\nu_-}[\text{CZ}(\tilde{\nu}_+, \tilde{\nu}_-)] = \xi_{\nu_+}$.

Proof. By an argument involving concatenation of paths, we may assume without loss of generality that $\nu_+ = \nu_-$. Then $\text{CZ}(\tilde{\nu}_+, \tilde{\nu}_-)$ is the index of a Cauchy-Riemann operator over the torus. The index can be calculated by a Grothendieck-Riemann-Roch argument similar to (2.1.1) and virtually identical to that of Smith [49]. The result is that

$$\text{CZ}(\tilde{\nu}_+, \tilde{\nu}_-) = [\bar{u}]^2 + \langle c_1(T^\vee(S^1 \times Y)), [\bar{u}] \rangle,$$

where we consider \bar{u} as a 2-cycle in $S^1 \times Y$. On the other hand, by a standard calculation in Seiberg-Witten theory, $\xi_{\nu_+} = \xi_{\nu_-}[d(\mathfrak{s})]$, where $\mathfrak{s} = [\bar{u}] \cdot \mathfrak{s}_{\text{can}}$. But

$$\begin{aligned} d(\mathfrak{s}) &= (c_1(\mathfrak{s})^2 - 2e - 3\sigma)/4 \\ &= c_1(\mathfrak{s}_{\text{can}})^2/4 + [\bar{u}]^2 + \langle c_1(\mathfrak{s}_{\text{can}}), [\bar{u}] \rangle \\ &= 0 + [\bar{u}]^2 + c_1(T^\vee(S^1 \times Y)), [\bar{u}]. \end{aligned}$$

□

Now suppose that $X \rightarrow S$ is a symplectic cobordism, and $u \in \mathcal{M}_{J,j}(\nu_-|X|\nu_+)$. Recall that the index of u is given by $\text{CZ}(\tilde{\nu}_+) - \text{CZ}(\tilde{\nu}_-) + \chi(S)$, where $\tilde{\nu}_+$ is induced by $\text{CZ}(\tilde{\nu}_-)$ and a symplectic trivialisation of $u^* T^\vee X$. The same argument gives

Lemma 4.1.6. $\xi_{\nu_-}[\text{CZ}(\tilde{\nu}_+) - \text{CZ}(\tilde{\nu}_-) + \chi(S)] = \xi_{\nu_+}$.

The ‘injectivity’ part of the next result is due to Smith [49], though our proof is different.

Proposition 4.1.7. *Let $\pi: X \rightarrow S$ be a Lefschetz fibration over a compact surface S , $\text{Hilb}_S^r(X)$ its r th relative Hilbert scheme. Then, if $r > 1$, the natural map*

$$\pi_0 \mathcal{S}(\text{Hilb}_S^r(X)) \rightarrow H_2(X, \partial X; \mathbb{Z})_r$$

is bijective.

Here $H_2(X, \partial X; \mathbb{Z})_r \subset H_2(X, \partial X; \mathbb{Z})$ is the set of classes β such that (a) $\beta \cdot [X_s] = r$ for each fibre X_s , and (b) $\beta \cdot \alpha \geq 0$ when α is the class of an irreducible component of a fibre.

Proof. We will prove this in two cases: (i) S is a disk and π has one critical point; (ii) S is connected with non-empty boundary and π has no critical points. We claim that these special cases imply the general result. Indeed, we can decompose a general base into pieces of these types, and we can patch together the homology classes. More interestingly, we can also patch together the sections. Indeed, by Example B.0.4 in Appendix B, when $r > 1$, sections of the r th relative symmetric product of a mapping torus $Y = T(\phi) \rightarrow S^1$ which define homologous 1-cycles in Y are homotopic.

(i): A section $u \in \mathcal{S}(\text{Hilb}_S^r(X))$, evaluated at the critical value 0, gives a point $u(0) \in \text{Sym}^r(X_0 \setminus \{\text{node}\})$, and this establishes a bijection $\pi_0 \mathcal{S}(\text{Hilb}_S^r(X)) \rightarrow \pi_0 \text{Sym}^r(X_0 \setminus \{\text{node}\})$ (these sets contain just one element if the node is non-separating, $r+1$ elements otherwise). In the non-separating case, $H_2(X, \partial X; \mathbb{Z}) = \mathbb{Z}$, so the map is certainly bijective. In the separating case, $H_2(X, \partial X; \mathbb{Z}) = \mathbb{Z}^2$, generated by the classes C_1 and C_2 of the irreducible components of X_0 . Thus $H_2(X, \partial X)_r = \{mC_1 + (r-m)C_2 : 0 \leq m \leq r\}$, and the map is again bijective.

(ii): S deformation-retracts onto an embedded bouquet of circles $\Gamma = \bigcup_i \Gamma_i \subset S$ (Γ could just be a point). The relative symmetric product $\text{Sym}_S^r(X) \rightarrow S$ is a fibre bundle, so the restriction map $\pi_0 \mathcal{S}(\text{Sym}_S^r(X)) \rightarrow \pi_0 \mathcal{S}(\text{Sym}_\Gamma^r(\pi^{-1}(\Gamma)))$ is bijective. But we know that $\pi_0 \mathcal{S}(\pi^{-1}(\Gamma_i))$ maps bijectively to $H_1(\pi^{-1}(\Gamma_i); \mathbb{Z})_r$ (see Appendix B). Thus the natural map $\pi_0(\text{Sym}_\Gamma^r(\pi^{-1}(\Gamma))) \rightarrow H_1(\pi^{-1}(\Gamma); \mathbb{Z})_r$ is bijective.

But restriction to Γ corresponds to the map $H^2(X, \partial X; \mathbb{Z})_r \rightarrow H_1(\pi^{-1}(\Gamma); \mathbb{Z})_r$ obtained by intersection with $\pi^{-1}(\Gamma)$. The latter is Poincaré dual to restriction $H^2(X; \mathbb{Z}) \rightarrow H^2(\pi^{-1}(\Gamma); \mathbb{Z})$, which is an isomorphism; hence the result. \square

Relative invariants. Let $(X, \pi, \omega, J_0, j_0)$ be a Lefschetz fibration over a compact surface with boundary, with fibres of genus g . Let $\beta \in H_2(X, \partial X)_r$, where $r \geq g-1$. To its boundary we have associated a Floer homology group

$$HF_\bullet(\partial X, \sigma, \partial\beta).$$

For convenience we parametrise ∂X as an outgoing boundary. We can now obtain a relative invariant

$$HF_{\bullet}(X, \omega, \beta) \in HF_{\bullet}(\partial X, \sigma, \partial \beta).$$

This is a direct application of the results of Chapter 3. We use a form χ as in (4.5), equal to $\omega_{\beta, \tau}$ except over small disks D_i containing the critical values of π . Notice that χ is unique up to deformations supported away from ∂X .

By the general theory sketched there, these relative invariants are functorial under composition of cobordisms.

As remarked in the Introduction, monopole Floer homology has an absolute $\mathbb{Z}/2$ -grading induced by a labelling of elements of homotopy classes $\xi \in J(Y)$ as odd or even. The symplectic theory HF_{\bullet} has an absolute $\mathbb{Z}/2$ -grading induced by a labelling of the generators $\nu \in \mathcal{H}(Y_r, \sigma_{L, \tau})$ as odd or even.

Lemma 4.1.8. *The homotopy class ξ_{ν} and the integer $CZ(\nu)$ have the same parity.*

This is certainly true when $Y = S^1 \times \Sigma$. A general Y may be expressed as the boundary of a Lefschetz fibration X over the disk. The set $H_2(X; \mathbb{Z})_r$ is non-empty, and so the result follows from the naturality of the two sides under cobordisms.

We now bring in the quantum cap product. If we mark a number of interior points $s_1, \dots, s_m \in S$ then this takes the form of a map

$$\Psi: \bigotimes_{k=1}^m H^*(\text{Sym}^r(X_{s_k}); \mathbb{Z}) \rightarrow HF_{\bullet}(\partial X, \sigma, \partial \beta). \quad (4.6)$$

We can compute this map using cycles in $\text{Sym}^r(X_{s_i})$ derived from cycles in X_{s_k} : for points x_1, \dots, x_a and 1-cycles $\gamma_1, \dots, \gamma_b$ in X_{s_k} , one obtains a cycle

$$\delta_{x_1} \cap \dots \cap \delta_{x_a} \cap \delta_{\gamma_1} \cap \dots \cap \delta_{\gamma_b}, \quad (4.7)$$

where $\delta_{x_i} = x_i + \text{Sym}^{r-1}(X_{s_k})$, $\delta_{\gamma_j} = \gamma_j + \text{Sym}^{r-1}(X_{s_k})$. These cycles span the homology of the symmetric product, and we can use sums of them to compute the quantum cap product. In other words, we consider the surjective homomorphism

$$\mu_{s_k}: \mathbb{Z}[H_0(X_{s_k}; \mathbb{Z})] \otimes_{\mathbb{Z}} \Lambda^* H_1(X_{s_k}; \mathbb{Z}) \rightarrow H^*(\text{Sym}^r(X_s); \mathbb{Z}),$$

and the composite

$$\Psi \circ \bigotimes_{k=1}^m \mu_{s_k}: \bigotimes_{k=1}^m \mathbb{Z}[H_0(X_{s_k}; \mathbb{Z})] \otimes \Lambda^* H_1(X_{s_k}; \mathbb{Z}) \rightarrow HF_{\bullet}(\partial X, \sigma, \partial \beta).$$

The homomorphism Ψ varies continuously under deformation of the marked points, and we can use this to express the quantum cap product in terms of homology classes on X . Let us take just one marked point $\Sigma = X_{s_1}$, so that we have a homomorphism

$$\Psi \circ \mu: \mathbb{Z}[U] \otimes \Lambda^* H_1(\Sigma; \mathbb{Z}) \rightarrow HF_{\bullet}(\partial X, \sigma, \partial \beta)$$

Here U corresponds to the point class in $H_0(\Sigma; \mathbb{Z})$ and is considered to have degree 2. Notice that $\Psi \circ \mu$ vanishes

1. on the ideal in $\mathbb{Z}[U] \otimes \Lambda^* H_1(\Sigma; \mathbb{Z})$ generated by classes of form $\gamma - m_{l*} \gamma$, $\gamma \in H_1(\Sigma; \mathbb{Z})$ where m_l is the monodromy of a loop $l: (I, \partial I) \rightarrow (S^*, s)$; and
2. on the ideal generated by the classes $\gamma \in H_1(\Sigma; \mathbb{Z})$ of the vanishing cycles associated with paths running into critical points of the Lefschetz fibration.

Vanishing in the first of these two cases is a direct consequence of the continuity property. In the second case, it holds because we may transport the cycle $\mu(\gamma)$ (which lies in a smooth fibre of $\text{Hilb}_S^r(X)$) along a vanishing path and into a singular fibre—indeed, we can replace it by a cycle *inside the singular locus* of this singular fibre. Since smooth sections of the relative Hilbert scheme map to regular points, the moduli space of sections which meet this cycle is empty, and thus the invariant vanishes.

But the kernel of $H_1(\Sigma; \mathbb{Z}) \rightarrow H_1(X; \mathbb{Z})$ is generated by monodromies $\gamma - m_{l*} \gamma$ and vanishing cycles. Hence $\Psi \circ \mu$ descends to a homomorphism

$$\Phi_X: \mathbb{Z}[U] \otimes_{\mathbb{Z}} \Lambda^* I \rightarrow HF_{\bullet}(\partial X, \sigma, \partial \beta)$$

where $I = \text{im}(H_1(\Sigma; \mathbb{Z}) \rightarrow H_1(X; \mathbb{Z}))$. To reiterate, the way one evaluates Φ_X on a homogeneous element $U^a \otimes g_1 \wedge \cdots \wedge g_b$ is to apply the map (4.6) to a sum of cycles (4.7) in fibres, observing that the outcome is independent of the choices involved.

When X is closed and the base is S^2 , we also have

Lemma 4.1.9. $I = H_1(X; \mathbb{Z})$.

Corollary 4.1.10. *When the base is S^2 , the relative invariant can be viewed as a homomorphism $\Phi_X: \mathbb{Z}[U] \otimes \Lambda^* H_1(X; \mathbb{Z}) \rightarrow \mathbb{Z}$. One could equally say $\Phi_X \in \mathbb{A}(X)$.*

Another feature of the closed situation is that the invariant does not change under deformations of the symplectic form (provided that we stay within the class of forms for which the invariant makes sense). On the other hand, the Lefschetz fibration has a distinguished deformation class of symplectic forms, and hence so does the Hilbert scheme.

In sum, we can arrange the relative invariant of a Lefschetz fibration $X \xrightarrow{\pi} S^2$, restricted to the ‘monotone range’ $d \geq 1$, as a map

$$\mathcal{L}_{X,\pi}: \text{Spin}^c(X)_{\geq 1} \rightarrow \mathbb{A}(X).$$

4.2 Lagrangian matching invariants

Here we set out the most important geometric construction in this chapter—the Lagrangian boundary condition to which we will apply the methods of Chapter 3. The proofs will be given in section 4.4. Start with

- Data 4.2.1.** 1. a compact, connected symplectic surface (Σ, ω) , and $\phi \in \text{Aut}(\Sigma, \omega)$;
2. a circle $L \subset \Sigma$ with $\phi(L) = L$;
3. $j \in \mathcal{J}(\text{T}(\phi), \omega_\phi)$.

There is an assortment of related objects (involving a small amount of further choice) which we list here for later reference:

- $Q \subset \text{T}(\phi)$ is the ω_ϕ -isotropic subbundle swept out by L .
- $(\bar{\Sigma}, \bar{\phi}) := (\sigma_L \Sigma, \sigma_L \phi)$ is the result of surgery along L . Note that $\bar{\Sigma}$ need not be connected.
- $\pi_{\text{br}}: X_{\text{br}}^L \rightarrow A$ is an elementary broken fibration over an annulus $A = A[1, 2] \subset \mathbb{C}$, realising a cobordism from $\text{T}(\phi)$ to $\text{T}(\bar{\phi})$ and having Q as a fibred vanishing cycle (see (1.2)). For definiteness, we shall suppose that $\text{T}(\phi)$ lies over the *inner* boundary of A (Σ is identified with $\pi_{\text{br}}^{-1}(1)$). Write Z for the circle $\text{crit}(\pi_{\text{br}})$.
- ω is a compatible near-symplectic form on X_{br}^L , extending ω_ϕ .
- We also write (Y, π, σ) for the mapping torus of ϕ , and $(\bar{Y}, \bar{\pi}, \bar{\sigma})$ for that of $\bar{\phi}$ (with its form induced by ω).
- We write \mathcal{S} (resp. $\bar{\mathcal{S}}$) for the relative symmetric product $\text{Sym}_{S^1}^r(Y)$ (resp. $\text{Sym}_{S^1}^{r-1}(\bar{Y})$).
- Choose a line bundle $P \rightarrow X_{\text{br}}^L$ which is the determinant of a Spin^c -structure and which has degree r over Σ . Abbreviate to c (resp. \bar{c}) the restriction of $c_1(P)$ to Y (resp. \bar{Y}). Then, upon choosing a real parameter τ , we have locally Hamiltonian fibrations $(\mathcal{S}, \pi_r, \sigma_{c, \tau})$ and $(\bar{\mathcal{S}}, \bar{\pi}_{r-1}, \bar{\sigma}_{\bar{c}, \tau})$. (The dependence on c and \bar{c} is slight; this is just a convenient way to arrange things.)

Main theorem 4.2.2. *The fibre product $(\mathcal{S} \times_{S^1} \bar{\mathcal{S}}, \pi_r \times_{S^1} \bar{\pi}_{r-1}, -\sigma_{c, \tau} \oplus \bar{\sigma}_{\bar{c}, \tau})$ contains an isotropic subbundle \mathcal{Q} such that the following hold.*

1. *The projection $\mathcal{Q} \rightarrow \mathcal{S}$ is an embedding. It represents the homology class δ_Q (see below).*
2. *The projection $\mathcal{Q} \rightarrow \bar{\mathcal{S}}$ is an S^1 -bundle. Indeed, there is a non-canonical isomorphism of fibre bundles*

$$\begin{array}{ccc}
 \mathcal{Q} & \xleftrightarrow{\quad} & \mathcal{Q} \times_{S^1} \bar{\mathcal{S}} \\
 & \searrow \text{pr}_2 & \swarrow \text{pr}_2 \\
 & \bar{\mathcal{S}} &
 \end{array}$$

Hence Λ , the intersection of \mathcal{Q} with the fibre $\text{Sym}^r(\Sigma) \times \text{Sym}^{r-1}(\bar{\Sigma})$, is a trivial S^1 -bundle over $\text{Sym}^{r-1}(\bar{\Sigma})$.

The given data determine such a subbundle canonically up to isotopy.

Remark 4.2.3. Let us clarify what is meant by the statement about homology classes. The situation is very simple when Q is orientable (a torus), but a little more awkward when it is a Klein bottle.

There is a formal addition map $Q \times \text{Sym}_{S^1}^{r-1}(Y) \rightarrow \mathcal{S}$. This carries a fundamental chain δ_Q . Likewise, \mathcal{Q} has a fundamental chain. When Q is non-orientable, these will not be cycles. Each is defined only up to sign, but choosing a sign for one gives a natural sign for the other. This is because \mathcal{Q} (or rather, a subbundle smoothly isotopic to it) has an open set consisting of triples $(t, x + D, D')$ where $t \in S^1$, $x \in Q_t$, D is a divisor consisting of $r-1$ points on Σ_t lying outside a chosen neighbourhood of Q_t , and D' is D considered as a divisor in $\bar{\Sigma}_t$ —see Equation 4.31.

Now, the difference chain, $\mathcal{Q} - \delta_Q$ is a null-homologous cycle in \mathcal{S} .

Remark 4.2.4. The zeroth symmetric product of a space is, for our purposes, a point. Thus when $r = 1$, \mathcal{Q} lies inside Y itself. In this case we may take $\mathcal{Q} = Q$.

Addendum 4.2.5. The homomorphism

$$\pi_2(\text{Sym}^r(\Sigma) \times \text{Sym}^{r-1}(\bar{\Sigma})) \rightarrow \pi_2(\text{Sym}^r(\Sigma) \times \text{Sym}^{r-1}(\bar{\Sigma}), \Lambda)$$

is surjective, provided that $\bar{\Sigma}$ is connected.

Sections of \mathcal{Q} versus Spin^c -structures. Let $d = \chi(\Sigma)/2 + r = \chi(\bar{\Sigma})/2 + (r-1)$. Consider the affine subgroup $H_2^{(d)} \subset H_2(X, Z \cup \partial X)$ of classes β such that (i) $\partial\beta = [Z]$ modulo $H_1(\partial X; \mathbb{Z})$; (ii) $\beta \cdot [\Sigma] = r$; and (iii) $\beta \cdot [\bar{\Sigma}_i] \geq 0$ for each component $\bar{\Sigma}_i$ of $\bar{\Sigma}$.

Addendum 4.2.6. There is a commutative diagram

$$\begin{array}{ccc} \pi_0 \mathcal{S}(\mathcal{Q}) & \xrightarrow{\text{pr}_{2*}} & \pi_0 \mathcal{S}(\bar{\Sigma}) \\ \downarrow & & \downarrow \text{cycle map} \\ H_2^{(d)} & \longrightarrow & H_1(\bar{Y}; \mathbb{Z})_{r-1} \end{array}$$

The cycle map is surjective, and bijective when $r-1 \geq 2$ (this follows from the description of the set homotopy classes of sections of a mapping torus given in Appendix B). The first vertical map is a canonical surjection. It is bijective when $r-1 \geq 2$ and $\bar{\Sigma}$ is connected.

The lower line was considered in the Introduction and in Chapter 1. It was established that there is a commutative diagram

$$\begin{array}{ccc} \text{Spin}^c(X)_d & \xrightarrow{\text{restr.}} & \text{Spin}^c(\bar{Y}) \\ \tau \downarrow & & \downarrow \tau \\ H_2^{(d)} & \longrightarrow & H_1(\bar{Y}; \mathbb{Z})_{r-1} \end{array}$$

The vertical maps are the ‘Taubes maps’ (see section (1.1.1)) and are bijective. The map on the lower row is, as already stated, surjective. Moreover, by (1.2.12), its (affine) kernel may be identified with the image in $H_1(X; \mathbb{Z})_1$ of $H_1(Q; \mathbb{Z})_1$.

Remark 4.2.7. *The results extend with minor modifications to the situation where Q is a disjoint union of k isotropic circle-bundles. (We stated them in the $k = 1$ case just for clarity’s sake.) Here we briefly indicate how they generalise.*

The most significant difference is that we must take $\bar{S} = \text{Sym}_{S^1}^{r-k}(\bar{Y})$ (interpreted as empty if $r < k$); $\text{pr}_2(Q)$ will then be the intersection of the hypersurfaces $\text{pr}_2(Q_i)$ coming from the components of Q .

The commutative triangle no longer makes sense, but $Q \cong \mathcal{L} \times_{S^1} \bar{S}$, where $\mathcal{L} \rightarrow S^1$ is a circle-bundle, trivial if and only if the number of Klein bottle components of Q is even.

The homomorphism of Addendum 4.2.5 is no longer surjective; its cokernel is isomorphic to $\ker(\mathbb{Z}\{[L_1], \dots, [L_k]\} \rightarrow H_1(\Sigma; \mathbb{Z}))$, where the classes $[L_i] \in H_1(\Sigma; \mathbb{Z})$ are those of the components L_i of $L \subset \Sigma$.

In Addendum 4.2.6 one just needs to replace $r - 1$ by $r - k$.

We can combine Addendum 4.2.6 with the result on homotopy classes of sections of relative Hilbert schemes (Prop. 4.1.7). Let (X, π, ω) be a directional near-symplectic fibration over S . To simplify the statement, we assume that $\pi(Z) \subset S$ is connected. Choose an embedding of a narrow closed annulus $A[1 - \epsilon, 1 + \epsilon]$ in S , so that $\pi(Z)$ is its unit circle. $S \setminus \text{int}(A)$ is a disjoint union $S_1 \sqcup S_2$; let $X_i = \pi^{-1}(S_i)$. Suppose that the fibres over S_1 (resp. S_2) are connected of genus g (resp. $g - k$). Thus ∂X_1 contains a (disconnected) fibred vanishing cycle Q , unique up to deformation (cf. Chapter 1). For each $r \geq \pi_0(Z)$ there is a Lagrangian boundary condition $Q = Q_r \subset \text{Sym}^r(X_1) \times_{S^1} \text{Sym}^{r-k}(X_2)$.

From Addendum 4.2.6, and the remarks on the generalisation to $k \geq 1$, we have the important

Corollary 4.2.8. *For each $r \geq k + 2$, the components of the set*

$$\{(u_1, u_2) \in \mathcal{S}(\text{Hilb}_{S_1}^r(X_1)) \times \mathcal{S}(\text{Hilb}_{S_2}^{r-1}(X_2)) : (\partial u_1, \partial u_2) \in \mathcal{S}(Q_r)\}$$

are in natural bijection with the subset $H_2^{(d)} \subset H_2(X, \partial X \cup Z)$.

(Even when $\bar{\Sigma}$ is not connected, the components still map surjectively to $H_2^{(d)}$.)

The proofs of the addenda, though simple, are geometric; they make use of the construction of Λ as a vanishing cycle. We postpone them, along with the proof of Theorem 4.2.2, to section 4.4.

Relative pin structures

Lemma 4.2.9. *The Lagrangian boundary condition Q has a canonical relative pin structure.*

Proof. Let $\xi \rightarrow \mathcal{S} \times_{\mathcal{S}^1} \bar{\mathcal{S}}$ be the pullback by pr_2 of the vector bundle $T^\vee \bar{\mathcal{S}}$. We exhibit a pin structure on $\xi|_{\mathcal{Q}} \oplus T^\vee \mathcal{Q}$. Notice that $T^\vee \mathcal{Q} = \text{pr}_2^* T^\vee \bar{\mathcal{S}} \oplus \lambda$, where λ is a line bundle. Hence the structure group of $\xi|_{\mathcal{Q}} \oplus T^\vee \mathcal{Q}$ reduces to the subgroup

$$\text{O}(1) \times \text{SO}(2r-2) \hookrightarrow \text{O}(4r-3), \quad (A, B) \mapsto \text{diag}(A, B, B).$$

We claim that this homomorphism factors through $\text{Pin}(4n-3)$. The restriction to identity components, $\text{id} \times \text{SO}(2r-2) \rightarrow \text{SO}(4r-3)$, factors through the covering group $\text{Spin}(4r-3)$ because it kills $\pi_1 \text{SO}(2r-2)$. But in odd dimensions, $\text{Pin} = \mathbb{Z}/2 \times \text{Spin}$ (just as $\text{O} = \mathbb{Z}/2 \times \text{SO}$), so the claim follows. This gives our canonical relative pin structure. \square

Gradings and near-symplectic cobordisms

For $\beta \in H_2^{(d)}$, $\gamma \in \pi_0 \mathcal{S}(\mathcal{S})$, and $\bar{\gamma} \in \pi_0 \mathcal{S}(\bar{\mathcal{S}})$, let us write

$$\gamma \sim_\beta \bar{\gamma}$$

if the class $\tilde{\beta} \in \pi_0 \mathcal{S}(\mathcal{Q})$ corresponding to β projects to γ under pr_{1*} and to $\bar{\gamma}$ under pr_{2*} .

Now take $\mathfrak{s} \in \text{Spin}^c(X)$, restricting to \mathfrak{t} (resp. $\bar{\mathfrak{t}}$) on Y (resp. \bar{Y}). For $\rho \in J(Y, \mathfrak{t})$ and $\bar{\rho} \in J(Y, \bar{\mathfrak{t}})$, we follow [22] in writing

$$\rho \sim_{\mathfrak{s}} \bar{\rho}$$

when there exists $I \in J(X, \mathfrak{s})$ preserving ρ on Y and $\bar{\rho}$ on \bar{Y} .

These two relations correspond as follows.

Proposition 4.2.10. *Let $\xi \in J(Y)$ (resp. $\bar{\xi} \in J(\bar{Y})$) be the classes represented by the vertical tangent bundle of π (resp. $\bar{\pi}$). Then*

$$\gamma \sim_\beta \bar{\gamma} \quad \Rightarrow \quad \xi_\gamma \sim_{\mathfrak{s}_\beta} \bar{\xi}_{\bar{\gamma}}$$

(Recall that ξ_γ denotes ξ twisted along γ .)

We can represent β by a surface $\hat{\beta}$ which is the disjoint union of $(r-1)$ sections of $X_{\text{br}}^L \rightarrow A$ and a cylinder C with one boundary component in Y and the other component equal to Z . We then have $\gamma \sim_\beta \bar{\gamma}$, where γ and $\bar{\gamma}$ are the sections of \mathcal{S} and $\bar{\mathcal{S}}$ corresponding to the boundaries of $\hat{\beta}$. The proposition follows from the following lemma.

Take a tubular neighbourhood N_Z of Z , and an almost complex structure I on $X \setminus N_Z$ compatible with π . This, of course, will not extend over N_Z .

As explained in the Introduction, a 2-plane field (or stable almost complex structure) on a 3-manifold can be twisted along an embedded 1-manifold. In similar fashion (and compatibly with restriction to the boundary) an almost complex structure I on a 4-manifold X can be twisted along an embedded surface S . The twist is induced by an automorphism

of TX supported near S . Choose a path $l: [0, 1] \rightarrow \mathrm{U}(2)$ such that $l(0) = -\mathrm{id}$ and $l(1) = \mathrm{id}$, and such that the loop $t \mapsto \det(l(t)) \in \mathrm{U}(1)$ has degree one. This induces $m: (\bar{D}^2, \partial\bar{D}^2) \rightarrow (\mathrm{U}(2), \mathrm{id})$, $re^{i\theta} \mapsto l(r)$, and hence, via a tubular neighbourhood of S , an automorphism of TX . Note that this operation is homotopically trivial if one does not fix the almost complex structures over the boundary.

Now we have:

Lemma 4.2.11. *The almost complex structure I' obtained by twisting I along $C \cap (X \setminus N_Z)$ extends over N_Z .*

Proof. Consider the almost complex structure J_0 on N_Z induced by an embedding $N_Z \hookrightarrow \mathbb{C}^2$ as a neighbourhood of a circle. This induces a homotopy class $\theta \in J(\partial N_Z)$. We claim that its twist, $\theta' = \theta[-(C \cap \partial N_Z)]$, coincides with the class ξ' induced by I' . By twisting both sides by $C \cap \partial N_Z$ we then obtain the result.

We can choose a framing $N_Z \cong S^1 \times B^3$ which maps $C \cap N_Z$ to $S^1 \times \{\mathrm{pt.}\} \subset S^1 \times S^2$. Now, both θ' and ξ' have underlying Spin^c -structure \mathfrak{t}_{-1} (here we use the notation of chapter 1). The set $J(S^1 \times S^2, \mathfrak{t}_{-1})$ has only two elements, of which precisely one is symmetric under S^1 . Both θ' and ξ' , considered as classes in $J(S^1 \times S^2, \mathfrak{t}_{-1})$, are S^1 -invariant. Hence they are equal. □

4.2.1 Defining the invariants

Elementary broken fibrations

We first explain how to use \mathcal{Q} to derive a homomorphism

$$HF_\bullet(X_{\mathrm{br}}^L, \omega, \beta): HF_\bullet(Y, \sigma, \gamma) \rightarrow HF_\bullet(\bar{Y}, \bar{\sigma}, \bar{\gamma})$$

from our elementary near-symplectic broken fibration. We assume for the time being that $\bar{\Sigma}$ is connected.

We know from Section 3.2.5 in Chapter 3 how, in principle, to obtain from \mathcal{Q} a relative invariant

$$\Phi(\mathcal{Q}) \in HF_*(\mathcal{S} \times \bar{\mathcal{S}}, -\sigma_{c,\tau} \oplus \bar{\sigma}_{\bar{c},\bar{\tau}})$$

(more precisely, a collection of invariants indexed by $\pi_0 \mathcal{S}(\mathcal{Q})$). To make it rigorous we have to check that \mathcal{Q} satisfies the hypotheses discussed there (specifically, those of Theorem 3.1.8). From Addendum 4.2.6, we have:

- Providing $\bar{\Sigma}$ has no 2-sphere component, the minimum Maslov index $\mu_{\min}(\Lambda)$ is equal to $2c_{\min}(\mathcal{S}_t \times \bar{\mathcal{S}}_t) = 2|d|$. The homomorphism μ is a positive multiple of $d \times \text{area}$.

Here $d = r + \chi(\Sigma)/2 = (r - 1) + \chi(\bar{\Sigma})/2$. This means that the level of difficulty involved in constructing the relative invariant $\Phi(Q)$ for a particular d is similar to that in constructing the Floer homology groups of \mathcal{S} and $\bar{\mathcal{S}}$ themselves:

- If $d > 0$, Q is *monotone*, and constructing $\Phi(Q)$ is unproblematic.
- If $r \leq (g(\Sigma) + 1)/2$ then Q satisfies the strong negativity assumption discussed in section 3.2.6: Lazzarini's lemma, in tandem with Hofer-Salamon transversality, allows us to construct an invariant $\Phi(Q)$.
- If $d = 0$ (i.e. $r = g(\Sigma) - 1$) the Maslov index is zero; thus Λ behaves rather like a special Lagrangian in a Calabi-Yau manifold. In (3.2.6) we sketched how to obtain an invariant $\Phi(Q)$, using arguments very close to established ones. This is a borderline case, and as such it demands closer attention. For the present we regard the construction of $\Phi(Q)$ as conjectural.
- This leaves outstanding the range $(g(\Sigma) + 1)/2 < r < g - 1$. It seems plausible that the moduli space for the Lagrangian boundary condition Q carries a *closed* virtual fundamental chain, and therefore that it is possible to define a rational invariant $\Phi(Q)$. An optimistic guess is that the moduli spaces of vertical holomorphic disks with boundary on Q can be related to spaces of spheres in the factors, and thereby shown to have codimension ≥ 2 .

By the procedure explained in Section 3.2.5 (in the paragraph on the Künneth formula), and using our canonical relative pin structure, we interpret the element $\Phi(Q)$ as a homomorphism

$$\Phi(Q) \in \text{Hom}_{\Lambda_{\mathbb{F}}} (HF_*(\mathcal{S}, \sigma_{c,\tau}), HF_*(\bar{\mathcal{S}}, \bar{\sigma}_{\bar{c},\bar{\tau}})). \quad (4.8)$$

Recall that this necessitates working over a Novikov *field* $\Lambda_{\mathbb{F}}$ so as to avoid Tor-terms.

We have established (Addendum 4.2.6) that the set of topological sectors $\pi_0 \mathcal{S}(Q)$ maps bijectively to $H_2^{(d)} \subset H_2(X_{\text{br}}^L, \partial X_{\text{br}}^L \cup Z; \mathbb{Z})$. For each $\beta \in H_2^{(d)}$, with boundary components $\gamma \in H_1(Y; \mathbb{Z})$ and $\bar{\gamma} \in H_1(\bar{Y}; \mathbb{Z})$, we interpret $\Phi(Q)$ as a homomorphism

$$HF_{\bullet}(X_{\text{br}}^L, \omega, \beta): HF_{\bullet}(Y, \sigma, \gamma) \rightarrow HF_{\bullet}(\bar{Y}, \bar{\sigma}, \bar{\gamma}). \quad (4.9)$$

In particular, we can take $[\sigma]$ proportional to $c_1(T^{\vee}Y)$ and $[\bar{\sigma}]$ to $c_1(T^{\vee}\bar{Y})$, and obtain a homomorphism

$$HF_{\bullet}(X_{\text{br}}^L, \beta): HF_{\bullet}(Y, \gamma) \rightarrow HF_{\bullet}(\bar{Y}, \bar{\gamma}). \quad (4.10)$$

Notes. (a) The homomorphism preserves degree in the sense of the discussion in the paragraph on gradings and degrees above.

(c) Compatibility with quantum product: $HF_\bullet(Y, \sigma, \gamma)$ (resp. $HF_\bullet(\bar{Y}, \bar{\sigma}, \bar{\gamma})$) is a module over $QH^*(\text{Sym}^r(\Sigma); \mathbb{Z})$ (resp. $QH^*(\text{Sym}^{r-1}(\bar{\Sigma}); \mathbb{Z})$). These module structures are specified by a distinguished endomorphism U of degree -2 (quantum cap product with η) and homomorphisms

$$\begin{aligned} H_1(\Sigma; \mathbb{Z}) &\rightarrow \text{Hom}(HF_\bullet(Y, \sigma, \gamma), HF_\bullet(Y, \sigma, \gamma)[-1]), \\ H_1(\bar{\Sigma}; \mathbb{Z}) &\rightarrow \text{Hom}(HF_\bullet(\bar{Y}, \bar{\sigma}, \bar{\gamma}), HF_\bullet(\bar{Y}, \bar{\sigma}, \bar{\gamma})[-1]). \end{aligned}$$

Lemma 4.2.12. *We have*

$$HF_\bullet(X_{\text{br}}^L, \omega, \beta) \circ U = U \circ HF_\bullet(X_{\text{br}}^L, \omega, \beta).$$

If $\alpha \in H_1(\Sigma; \mathbb{Z})$, $\bar{\alpha} \in H_1(\bar{\Sigma}; \mathbb{Z})$ represent the same class in $H_1(X; \mathbb{Z})$ then

$$\bar{\alpha} \circ HF_\bullet(X_{\text{br}}^L, \omega, \beta) = HF_\bullet(X_{\text{br}}^L, \omega, \beta) \circ \alpha.$$

Proof. The classes $\eta_\Sigma \in H^2(\text{Sym}^r(\Sigma); \mathbb{Z})$ and $\eta_{\bar{\Sigma}} \in H^2(\text{Sym}^{r-1}(\bar{\Sigma}); \mathbb{Z})$ agree when pulled back to the product and restricted to Λ . Thus the first equation follows from the relation (3.19). The same applies to pairs $\alpha \in H^1(\text{Sym}^r(\Sigma); \mathbb{Z})$ and $\bar{\alpha} \in H^1(\text{Sym}^{r-1}(\bar{\Sigma}); \mathbb{Z})$. \square

(d) *There are maps in both directions!* Indeed, there is a very similar discussion when the base annulus has the opposite orientation, cf. 3.2.5. This is in marked contrast to the relative invariants of Lefschetz fibrations, where only one orientation of the base is allowed.

Open case. We can extend the above discussion to general near-symplectic broken fibrations (X, π, ω) , provided they are directional (i.e., that each circle of critical values has a consistent ‘connected’ or ‘high-genus’ side) and that all fibres are connected.

Assume that (X, π, ω) has incoming boundary $(Y_{\text{in}}, \pi_{\text{in}}, \sigma_{\text{in}})$ and outgoing boundary $(Y_{\text{out}}, \pi_{\text{out}}, \sigma_{\text{out}})$.

One defines the invariant by the following steps.

- Choose oriented embeddings $f_i: A \hookrightarrow S$ of closed annuli as narrow, disjoint neighbourhoods of the circles in $\pi(X^{\text{crit}})$. No isolated critical value is allowed to lie in $\text{im}(f_i)$, nor is any point of ∂S .
- Consider $S' = S \setminus \bigcup_i \text{int}(\text{im}(f_i))$. To each pair (Y_i, \bar{Y}_i) of boundary components of $\pi^{-1}(S')$ one associates a fibred vanishing cycle $Q_i \subset Y_i$, and hence a Lagrangian boundary condition

$$\mathcal{Q}_i \subset \text{Sym}_{S^1}^r(Y_i) \times_{S^1} \text{Sym}_{S^1}^{r-k}(\bar{Y}_i),$$

where $k = \pi_0(Q)$.

By the procedure explained in Section 3.2.5 (under the heading ‘a small generalisation’) one obtains a homomorphism

$$HF_{\bullet}(X, \omega, \beta): HF_{\bullet}(Y_{\text{in}}, \sigma_{\text{in}}, \beta_{\text{in}}) \rightarrow HF_{\bullet}(Y_{\text{out}}, \sigma_{\text{out}}, \beta_{\text{out}}).$$

We need to establish that this is independent of the auxiliary choices involved. These are

- The choice of almost complex structure on TX , away from the singular circles, compatible with the broken fibration.
- The embeddings f_i .
- The fibred vanishing cycles Q_i .
- The Lagrangian boundary conditions \mathcal{Q}_i .
- The closed 2-forms on the Hilbert schemes near the singular fibres.

It is clear that the auxiliary choices form a path-connected space. In particular, the construction of vanishing cycles is unique up to deformation, and varies continuously in families. Hence the invariance follows from the deformation-invariance of the relative invariants.

We should note, however, that the space of auxiliary choices may well not be contractible (it is possible that parametrised versions of the relative invariants detect essential spheres).

We can also see that one can make deformations of the map π and 2-form ω , supported away from ∂X , without affecting the invariant. The only non-trivial point is to see that it makes no difference whether one has two singular circles mapping to one circle in $C \subset S$ or to two parallel circles $C_1, C_2 \subset S$. This is not difficult to prove, using transversality of evaluation arguments, but we do not give the details here (nor do we use the result).

There is a more general form of the invariant, incorporating the quantum cap product. Rather than recording it in full, we simply note here that there is distinguished endomorphism U of degree -2 (cap product with an η -class in a regular fibre) which commutes with $HF_{\bullet}(X, \omega, \beta)$.

Closed case. Suppose that (X, π) is a broken fibration over S^2 , with singular circles Z_i mapping to a single equator in S^2 , with the directionality property, and admitting a compatible near-symplectic form (cf. Theorem 0.0.2).

A Lefschetz fibration over the disk has the property that the inclusion of a fibre induces a surjection on H_1 . From this it is easy to deduce

Lemma 4.2.13. *$H_1(X; \mathbb{Z})$ is spanned by the images of $H_1(X_b; \mathbb{Z})$ as b varies over S^2 .*

Hence in this case we can formulate the relative invariant as a map

$$\mathcal{L}_{X, \pi}: \text{Spin}^c(X)_{\geq 1} \rightarrow \mathbb{A}(X).$$

We call this the Lagrangian matching invariant. Its value $\mathcal{L}_{X,\pi}(\mathfrak{s})$ is specified by evaluating it on homogeneous elements of $\mathbb{Z}[U] \otimes_{\mathbb{Z}} \Lambda^* H_1(X; \mathbb{Z})$ by the same procedure as in the Lefschetz case (using fibrewise divisors). It follows from Lemma (4.2.12) that it does not matter whether we choose fibres on either side of a critical circle or the other, hence that $\mathcal{L}_{X,\pi}$ is well-defined.

Note also that it is independent of the choice of near-symplectic form ω , because these are determined by π , up to deformation.

4.2.2 Some simple calculations

Dimensional reduction

The category of near-symplectic cobordisms contains a much simpler category $\mathcal{C}_{\text{Morse}}$ in which all the structures have S^1 -symmetry.

An object of $\mathcal{C}_{\text{Morse}}$ is a triple (Σ, c, r) consisting of a closed, oriented surface with a locally constant function $c: \Sigma \rightarrow S^1$ and $r \in H_0(\Sigma; \mathbb{Z})$. A morphism (Y, f, a) from (Σ_1, c_1, r_1) to (Σ_2, c_2, r_2) is a cobordism Y , equipped with a circle-valued Morse function $f: Y \rightarrow S^1$ with no interior maximum or minimum, and a homology class $a \in H_1(Y, \partial Y \cup \text{crit}(f); \mathbb{Z})$. There are two requirements on a : its boundary must be

$$\partial a = r_2 - r_1 + \sum_{x \in \text{crit}(f)} (-1)^{\text{ind}(x)+1} [x] \in H_0(\partial Y \cup \text{crit}(f); \mathbb{Z}),$$

and the number $a \cap [F] - \dim(H_1(F))/2$ must be constant as F ranges over regular fibres F of f .

For example, the second condition rules out the class a of a path running from an index 2 critical point to a ‘cancelling’ index 1 critical point. An alternative would be to keep track of a line bundle $L \rightarrow Y$ instead of the classes a .

We can see $\mathcal{C}_{\text{Morse}}$ as a dimensionally-reduced version of \mathcal{NS} ; one obtains a near-symplectic cobordism by crossing everything with S^1 (strictly the near-symplectic form is only determined up to deformation, though that could be fixed by keeping track of metrics such that df is harmonic).

Elementary cobordisms. A basic case to consider is one where Y is an elementary cobordism C_{elem}^- from Σ to $\bar{\Sigma}$; the Morse function f is real valued, and has a single critical point, of index $+1$.

This gives rise to an invariant

$$\Phi: HF_{\bullet}(\Sigma_r \times S^1) \rightarrow HF_{\bullet}(\bar{\Sigma}_{r-1} \times S^1),$$

($\Sigma_r = \text{Sym}^r(\Sigma)$, $\bar{\Sigma}_{r-1} = \text{Sym}^{r-1}(\bar{\Sigma})$). Here we simplify the notation by summing over topological sectors and suppressing the 2-forms. On the other hand, we have

$$\begin{aligned} HF_\bullet(\Sigma_r \times S^1, r) &\cong H_*(\Sigma_r; \Lambda_{\mathbb{F}}) \cong X(\Sigma, r) \otimes_{\mathbb{Z}} \Lambda_{\mathbb{F}}, \\ HF_\bullet(\bar{\Sigma}_{r-1} \times S^1, r-1) &\cong H_*(\bar{\Sigma}_{r-1}; \Lambda_{\mathbb{F}}) \cong X(\bar{\Sigma}, r-1) \otimes_{\mathbb{Z}} \Lambda_{\mathbb{F}} \end{aligned}$$

Here, for a compact oriented surface S , we write

$$X(S, n) := \bigoplus_{j=0}^n \text{Sym}^j H_{\text{ev}}(S; \mathbb{Z}) \otimes_{\mathbb{Z}} \Lambda^{n-j} H_1(S; \mathbb{Z}).$$

By Prop. (3.2.6), Φ is the fundamental class $[\Lambda]$ of $\Lambda \subset \Sigma_r \times \bar{\Sigma}_{r-1}$, interpreted via the Künneth formula and Poincaré duality on Σ_r :

$$\begin{aligned} \Phi = [\Lambda] \in H_{2r-1}(\Sigma_r \times \bar{\Sigma}_{r-1}; \Lambda_{\mathbb{F}}) &\cong [H_*(\Sigma_r; \Lambda_{\mathbb{F}}) \otimes_{\Lambda_{\mathbb{F}}} H_*(\bar{\Sigma}_{r-1}; \Lambda_{\mathbb{F}})]_{2r-1} \\ &\cong \text{Hom}_{\Lambda_{\mathbb{F}}}(H_*(\Sigma_r; \Lambda_{\mathbb{F}}), H_{*-1}(\bar{\Sigma}_{r-1}; \Lambda_{\mathbb{F}})). \end{aligned}$$

We can also consider an elementary cobordism C_{elem}^+ running in the opposite direction (so the Morse function has one critical point, of index 2). Its invariant is a homomorphism

$$\Psi: HF_\bullet(\bar{\Sigma}_{r-1} \times S^1) \rightarrow HF_\bullet(\Sigma_r \times S^1).$$

Again, it has a simple interpretation, which this time involves Poincaré duality on $\bar{\Sigma}_{r-1}$:

$$\Psi = [\Lambda] \in H_{2r-1}(\bar{\Sigma}_{r-1} \times \Sigma_r; \Lambda_{\mathbb{F}}) \cong \text{Hom}_{\Lambda_{\mathbb{F}}}(H_*(\bar{\Sigma}_{r-1}; \Lambda_{\mathbb{F}}), H_{*+1}(\Sigma_r; \Lambda_{\mathbb{F}})).$$

It is not difficult to give completely explicit formulae for the maps Φ and Ψ , but what is more illuminating is to observe that they are the same as the maps that arise in a well-known (2+1)-dimensional TQFT, which we shall call the Segal-Donaldson model [41, 6]. This can be characterised as follows.

Definition 4.2.14. The **admissible (2+1)-dimensional cobordism category** of degree d is the category in which an object is an oriented 2-manifold, and a morphism is an ‘admissible cobordism’ Y equipped with a line bundle which has degree d over the incoming and outgoing ends. Admissible means that either

$$\begin{aligned} \partial Y = \emptyset, \quad H_1(Y; \mathbb{Z}) &\cong \mathbb{Z}, \text{ or} \\ \partial Y \neq \emptyset, \quad H_1(\partial Y; \mathbb{Z}) &\rightarrow H_1(Y; \mathbb{Z}) \text{ is onto.} \end{aligned}$$

The Segal-Donaldson model is a TQFT for the admissible (2+1)-dimensional cobordism category, i.e. a functor from this category to super $\Lambda_{\mathbb{F}}$ -vector spaces, preserving the multiplicative structures (disjoint union; super tensor product) and dualities. Strictly, the morphisms are defined only up to a sign ± 1 (Segal explains how to fix this ambiguity, but I have not yet checked how this applies in our situation).

Theorem 4.2.15 (Donaldson [6]). *For any sequence of non-negative integers n_0, n_1, \dots , there is a unique TQFT V_d on the level d admissible cobordism category such that*

$$V_d(\Sigma) = \bigoplus_{i \geq 0} n_i \Lambda_i \Sigma$$

for connected surfaces Σ , as representations of the mapping class group $\pi_0 \text{Diff}^+(\Sigma)$. Here $\Lambda_i \Sigma$ denotes $\Lambda^{g(\Sigma)-i} H_1(\Sigma; \mathbb{F})$. If Y^3 is closed, so that there is a unique line bundle L of degree d , then

$$V_d(Y, L) = \sum_i n_i a_i,$$

where $A(Y) = \pm(a_0 + \sum_i a_i(t^i + t^{-i}))$ is the normalised Alexander polynomial of Y .

(‘Unique’ means up to natural isomorphism.) We shall take V_d to be the unique TQFT corresponding to the sequence

$$n_i = i + 1 + d.$$

This is of interest (especially when $d \geq 0$) because we have

$$V_d(\Sigma) = X(\Sigma, g - 1 + d) \otimes \Lambda_{\mathbb{F}}$$

as representations of $\pi_0 \text{Diff}^+(\Sigma)$.

Note that one can sum over isomorphism classes of line bundles (and just write $V_d(Y)$) but one must be careful about how these maps compose.

Lemma 4.2.16. *Φ satisfies the following properties:*

1. *Let $(\alpha, \bar{\alpha}) \in H_1(\Sigma; \mathbb{Z}) \times H_1(\bar{\Sigma}; \mathbb{Z})$ be such that the images of α and $\bar{\alpha}$ in $H_1(C_{\text{elem}}^-; \mathbb{Z})$ are equal. Then $\bar{\alpha} \wedge \Phi(x) = \Phi(\alpha \wedge x)$.*
2. *$U \circ \Phi = \Phi \circ U$ and $\bar{\theta} \circ \Phi = \Phi \circ \theta$, where $\theta \in \Lambda^2 H_1(\Sigma; \mathbb{Z})$ and $\bar{\theta} \in \Lambda^2 H_1(\bar{\Sigma}; \mathbb{Z})$ are the cup product forms;*
3. *on the degree 1 component $H_1(\Sigma; \mathbb{Z})$ of $X(\Sigma, r)$, Φ is given, up to sign, by $x \mapsto x \cap [L]$.*

Moreover, it is uniquely characterised, up to sign, by these properties.

Proof. The first two properties are instances of (3.19), though they can be verified in an elementary way here. The third is another way of writing the equation $\text{pr}_{1*}[\Lambda] = [\delta_L]$.

For the uniqueness, we note that by (2) it suffices to show that Φ is uniquely determined on the subspace $\Lambda^r H_1(\Sigma; \mathbb{Z}) \subset X(\Sigma, r)$; this is a simple consequence of (1) and (3). \square

On the other hand, the Donaldson-Segal map $V_d(C_{\text{elem}}^-)$ satisfies these same properties, as one can easily verify from Donaldson’s model.

Corollary 4.2.17. $V_d(C_{\text{elem}}^-) = \pm\Phi$, and $V_d(C_{\text{elem}}^+) = \pm\Psi$.

(The statement about Ψ is proved by an identical argument.)

Corollary 4.2.18. *Let Y be a closed 3-manifold with $H_1(Y; \mathbb{Z}) \cong \mathbb{Z}$, and $f: Y \rightarrow S^1$ a Morse function with connected fibres and no extrema. Let*

$$\text{id} \times f: S^1 \times Y \rightarrow S^1 \times S^1$$

be the corresponding broken fibration. Then $\mathcal{L}_{S^1 \times Y, \text{id} \times f}(\mathfrak{s}_d) = \pm \sum_{i \geq 0} (i + 1 + d)a_i$, where $A(Y) = \pm(a_0 + \sum_i a_i(t^i + t^{-i}))$ is the Alexander polynomial.

(The assumption $H_1(Y; \mathbb{Z}) \cong \mathbb{Z}$ implies that such Morse functions exist.) The result presumably holds for general H_1 , provided that the Alexander polynomial is replaced by the Milnor-Turaev torsion.

Connected sum model

Disconnected fibres. In what follows, disconnected fibres $\bar{\Sigma}$ are unavoidable. We insert here a brief discussion on Lagrangian matching invariants where there are disconnected fibres. To be precise, we explain how to modify the construction of (4.9) when L is connected but separating. We suppose that $\bar{\Sigma}$ has components $\bar{\Sigma}_1$ and $\bar{\Sigma}_2$ of genera $g_1 > 0$ and $g_2 > 0$.

There are two points we need to take care of. The first is that the homotopy classes of sections of \mathcal{Q} map surjectively, but not bijectively, to $H_2^{(d)}$. Thus in (4.9) we should sum over homotopy classes inducing the same class β . Gromov-Floer compactness implies that the summation is finite. Indeed, for any $(\nu, \bar{\nu}) \in \mathcal{H}(\text{Sym}_{S^1}^r(Y)) \times \mathcal{H}(\text{Sym}_{S^1}^{r-1}(\bar{Y}))$, the action of an index zero pseudoholomorphic section of

$$(-\infty, 0] \times \text{Sym}_{S^1}^r(Y) \times_{S^1} \mathcal{H}(\text{Sym}_{S^1}^{r-1}(\bar{Y})) \rightarrow (-\infty, 0] \times S^1,$$

asymptotic to $(\nu, \bar{\nu})$ and with boundary on \mathcal{Q} , depends only on β .

The second point concerns bubbling. We have

$$\text{Sym}^{r-1}(\bar{\Sigma}) = \prod_{r_1=0}^{r-1} \text{Sym}^{r_1}(\bar{\Sigma}_1) \times \text{Sym}^{r-1-r_1}(\bar{\Sigma}_2).$$

Thus there are positive-area bubbles of Chern numbers $r' + 1 - g_1$ and $r' + 1 - g_2$ for each $2 \leq r' \leq r$. One finds that $\text{Sym}^{r-1}(\bar{\Sigma})$ is weakly monotone provided that $r \leq \min(g_1/2, g_2/2)$. As to boundary bubbles, we now have

$$\text{coker}(\pi_2(\text{Sym}^r(\Sigma) \times \text{Sym}^{r-k}(\bar{\Sigma})) \rightarrow \pi_2(\text{Sym}^r(\Sigma) \times \text{Sym}^{r-k}(\bar{\Sigma}, \Lambda)) \cong \mathbb{Z}.$$

It is not difficult to verify that there are positive-area disks, representing generators of the cokernel, having Maslov indices $2[(r-1) + (1-g_i)] - [r+1-g_i]$. When $r \leq \min(g_1/2, g_2/2)$, the ‘strong negativity’ assumption discussed at the end of Chapter 3 is valid, and the construction of the invariant goes through.

Vanishing of the invariant. A basic fact in Seiberg-Witten theory is that $\mathcal{SW}_{X_1 \# X_2}$ vanishes identically when the connected summands X_1 and X_2 have positive b_2^+ . (Similarly, the relative invariant in monopole Floer theory associated to a cobordism which is an interior connected sum of two cobordisms is zero.)

At the end of Chapter 1 we described a simple model for the connected sum of two broken fibrations, $(X_1 \# X_2, \pi_1 \# \pi_2)$. Recall that $X_1 \# X_2$ has a portion X' which is a broken fibration over the disk \bar{D} ; the central fibre $\Sigma_1 \# \Sigma_2$ is connected, the boundary fibre $\Sigma_1 \sqcup \Sigma_2$ disconnected. Deleting a small central disk we obtain a broken fibration $X'' \rightarrow A$ over an annulus.

Its Lagrangian matching invariant vanishes:

Proposition 4.2.19. $\mathcal{L}_{(X_1 \# X_2, \pi_1 \# \pi_2)}(\beta)$ is well-defined, and equal to zero, for each $\beta \in H_2^{(d)}$ such that $\beta \cdot ([\Sigma_1] + [\Sigma_2]) \leq g(\Sigma_i)/2$, $i = 1, 2$.

Proof. The relative invariants of $X'' \rightarrow A$ are zero, for a simple homological reason. They are given by components of the Segal-Donaldson homomorphism $\Phi: V_d(\Sigma_1 \sqcup \Sigma_2) \rightarrow V_d(\Sigma_1 \# \Sigma_2)$ corresponding to surgery along the separating circle L along which the connected sum takes place. But when the homology class $[L]$ is zero, Φ is zero, as is clear from the lemma about Φ above. □

4.3 Geometry of the compactified families

Remark 4.3.1. *Compactified Jacobians and Hilbert schemes, and their relative versions, have been studied by algebraic geometers for many years. The modern approach, using geometric invariant theory to construct projective compactifications over the moduli space of stable curves [5], is a systematic and powerful one. However, for us, functoriality is less important than explicitness. We give concrete descriptions of the spaces involved, for the most part avoiding group quotients, and we do not need to appeal to general existence theorems.*

The idea of using such compactifications in the present context is due to Donaldson and Smith; see their paper [9] for further information and references.

Our presentation draws both on [9] and on the explicit formulae in the recent paper of Ran [36].

In this section, we work in the category of complex analytic spaces (purists might therefore object to our use of the term ‘Hilbert scheme’ rather than ‘Douady space’).

4.3.1 The Jacobian of a compact nodal curve

The compactified Jacobian $\mathcal{P}^0(C)$ of a compact nodal curve is a projective compactification of the identity component of the Picard group $\text{Pic}(C) = H^1(C, \mathcal{O}^\times)$. It parametrises torsion-free coherent analytic sheaves of rank 1 and degree 0 on C .

We will construct the compactified Jacobian only in the case where there is just one node x_0 , so we may as well impose that condition from the beginning. The normalisation map $\nu: \tilde{C} \rightarrow C$ collapses a pair of points (x_+, x_-) to x_0 .

The Picard group $\text{Pic}(C) = H^1(C, \mathcal{O}^\times)$ parametrises invertible sheaves. The exponential sequence on C gives a homomorphism $c_1: \text{Pic}(C) \xrightarrow{c_1} H^2(C; \mathbb{Z}) \cong H^2(\tilde{C}; \mathbb{Z})$; the sequence identifies $\text{Pic}^0(C) := \ker(c_1)$ with $H^1(C, \mathcal{O})/H^1(C, \mathbb{Z})$. As a quotient of a vector space by a discrete subgroup, $\text{Pic}^0(C)$ has a natural structure of complex Lie group.

The short exact sequence of sheaves on C

$$0 \rightarrow \mathcal{O}_C^\times \rightarrow \nu_* \mathcal{O}_{\tilde{C}}^\times \rightarrow \mathcal{O}_C^\times / \nu_* \mathcal{O}_{\tilde{C}}^\times \rightarrow 0$$

gives rise to an exact sequence of cohomology groups which, when C is irreducible, reads

$$0 \rightarrow \mathbb{C}^* \rightarrow \text{Pic}^0(C) \xrightarrow{\nu^*} \text{Pic}^0(\tilde{C}) \rightarrow 0. \quad (4.11)$$

When C is reducible, $\text{Pic}^0(C) \xrightarrow{\nu^*} \text{Pic}^0(\tilde{C})$ is an isomorphism.

There is also a ‘Jacobian’ description of $\text{Pic}^0(C)$, related to the previous one through Serre duality. A description of the dualising sheaf \mathcal{K}_C which is valid on any compact nodal curve is as follows: for an open set U , $\mathcal{K}_C(U)$ is the $\mathcal{O}_C(U)$ -module of meromorphic differentials η on $\nu^{-1}(U)$, holomorphic apart from simple poles at the points x_\pm^α lying over nodes x_0^α , and with

$$\text{Res}(\eta; x_+^\alpha) + \text{Res}(\eta; x_-^\alpha) = 0.$$

This definition makes \mathcal{K}_C an invertible subsheaf of $\nu_*(\mathcal{K}_{\tilde{C}}(\sum x_+^\alpha + x_-^\alpha))$. When C is irreducible with one node, the condition on the residues is vacuous by the residue theorem.

Now, Serre duality takes the form of a perfect pairing

$$H^0(C, \mathcal{K}_C) \otimes H^1(C, \mathcal{O}_C) \rightarrow H^1(C, \mathcal{K}_C) \xrightarrow{\text{Res}} \mathbb{C},$$

and we have $\text{Pic}(\tilde{C}) \cong H^0(\tilde{C}, \mathcal{K}_{\tilde{C}})^\vee / H_1(\tilde{C}; \mathbb{Z})$. Likewise, when C is irreducible, there is an isomorphism

$$\text{Pic}^0(C) \cong \frac{H^0(C, \mathcal{K}_C)^\vee}{H_1(C \setminus \{x_0\}; \mathbb{Z})}.$$

(Note that $H^1(C; \mathbb{Z}) \cong H_1(C \setminus \{x_0\}; \mathbb{Z})$.) We can see again from this description that $\text{Pic}^0(C)$ fibres over $\text{Pic}^0(\tilde{C})$, with fibre $\mathbb{C}/\mathbb{Z} = \mathbb{C}^*$.

Constructing the compactified Jacobian

If C is reducible (with, as before, just one node) then $\text{Pic}^0(C)$ is already compact. We proceed to describe a compactification $\mathcal{P}(C)$ in the case where C is irreducible with one node. The points of $\mathcal{P}(C)$ correspond to torsion-free sheaves of rank 1 (here ‘rank’ is the unique Euler-Poincaré map on coherent sheaves such that line bundles have rank 1). It is a disjoint union of components $\mathcal{P}^r(C)$, where the sheaves have degree (Chern number) r .

If \mathcal{F} is a torsion-free sheaf of rank 1, then we can pull back by ν and double-dualise (killing any torsion) to obtain an invertible sheaf

$$\mathcal{F}^\nu := (\nu^* \mathcal{F})^{\vee\vee} \in \text{Pic}(\tilde{C}).$$

Take a Poincaré line bundle $P \rightarrow \tilde{C} \times \text{Pic}(\tilde{C})$, and let $P_x \rightarrow \text{Pic}(\tilde{C})$ denote its restriction to the slice $\{x\} \times \text{Pic}(\tilde{C})$. Let $E = P_{x_+}^\vee \oplus P_{x_-}^\vee$; this is a rank 2 holomorphic vector bundle whose fibre over the line bundle \mathcal{L} is $(\mathcal{L}_{x_+})^\vee \oplus (\mathcal{L}_{x_-})^\vee$. We consider the projectivised bundle $\mathbb{P}E$. For each point $(\mathcal{L}, (a : b)) \in \mathbb{P}E$, we define a sheaf $\mu(\mathcal{L}, (a : b))$ on C as the subsheaf of $\nu_* \mathcal{L}$ given by ‘gluing the fibres \mathcal{L}_{x_+} and \mathcal{L}_{x_-} according to $(a : b)$ ’: that is, as the sheaf associated to the presheaf

$$U \mapsto \{s \in \mathcal{L}(\nu^{-1}U) : a(s(x^+)) = b(s(x^-)) \text{ if } x_0 \in U\}.$$

Lemma 4.3.2. *$\mu(\mathcal{L}, (a : b))$ is a torsion-free coherent sheaf of rank 1 with the same Chern class as \mathcal{L} .*

We verify this by distinguishing two cases.

1. The \mathbb{C}^* -subbundle of $\mathbb{P}E$ where a and b are non-zero.

In this case $\mu(\mathcal{L}, (a : b))$ is invertible. If \mathcal{F} is an invertible sheaf on C then \mathcal{F} is isomorphic to $\mu(\mathcal{F}^\nu, (a : b))$ for some unique $(a : b)$. This establishes an isomorphism between $\text{Pic}(C)$ and the \mathbb{C}^* -subbundle of $\mathbb{P}E$.

2. The zero- and infinity-sections of $\mathbb{P}E$.

The sheaf $\mu(\mathcal{L}, (1 : 0))$ is not invertible, as the stalk at x_0 is not singly generated, but it is the push-forward of a locally free sheaf, and hence coherent:

$$\mu((\mathcal{L}, (1 : 0))) = \nu_*(\mathcal{L}(-x_+)).$$

Likewise $\mu((\mathcal{L}, (0 : 1))) = \nu_*(\mathcal{L}(-x_-))$.

Notice that $\nu^* \nu_*(\mathcal{L}(-x_+))$ *does* have torsion. If \mathcal{F} is a rank 1 torsion-free sheaf which is *not* invertible, then $\mathcal{F} \cong \mu(\mathcal{F}^\nu(-x_+), (1 : 0))$. Thus these sheaves \mathcal{F} correspond bijectively with the zero-section of $\mathbb{P}E$.

Here is the definition of $\mathcal{P}(C)$. There are two embeddings $i_{\pm}: \text{Pic}(\tilde{C}) \rightarrow \mathbb{P}E$,

$$\begin{aligned} i_+(\mathcal{L}') &= (\mathcal{L}'(x_+), (1 : 0)), \\ i_-(\mathcal{L}') &= (\mathcal{L}'(x_-), (0 : 1)), \end{aligned}$$

and these satisfy $\mu \circ i_+ = \mu \circ i_-$. Moreover, $i_+(\text{Pic}(\tilde{C}))$ and $i_-(\text{Pic}(\tilde{C}))$ have naturally isomorphic neighbourhoods, so it makes sense to define $\mathcal{P}(C)$ as the colimit

$$\text{Pic}(\tilde{C}) \begin{array}{c} \xrightarrow{i_+} \\ \xrightarrow{i_-} \end{array} \mathbb{P}(E) \longrightarrow \mathcal{P}(C). \quad (4.12)$$

Thus, as a topological space, $\mathcal{P}(C)$ is the quotient $\mathbb{P}E/(i_+\mathcal{L}' \sim i_-\mathcal{L}')$. As an analytic space, $\mathcal{P}(C)$ is irreducible and has normal crossing singularities along the gluing locus $\mathcal{P}(C)_{\text{sing}} = \text{Pic}(\tilde{C})$. The normal bundles to $\text{Pic}(\tilde{C})$ along the two branches correspond to the normal bundles to $\mathbb{P}E$ along the zero- and infinity-sections, that is, to $P_{x_-} \otimes P_{x_+}^{\vee}$ and its dual $P_{x_+} \otimes P_{x_-}^{\vee}$.

These normal bundles are differentiably trivial—they are complex line bundles of Chern class zero—but not holomorphically so, because x_+ and x_- are linearly inequivalent as divisors on \tilde{C} .

4.3.2 Compactifying the Picard family

Next we consider a proper family of curves, $\pi: X \rightarrow \bar{D}$, over the unit disk. We suppose that the family is smooth over the punctured disk, and that X_0 has one node x_0 and no other singularities. There is then a Picard family

$$\text{Pic}_{\bar{D}}(X) = R^1\pi_*\mathcal{O}_X^{\times} \rightarrow \bar{D}$$

with fibres $\text{Pic}_{\bar{D}}(X)_z = \text{Pic}(X_z)$. It breaks into connected components $\text{Pic}_{\bar{D}}^r(X)$, $r \in \mathbb{Z}$. When X_0 is reducible, $\text{Pic}_{\bar{D}}^r(X)$ is already proper, and we shall not modify it. When X_0 is irreducible, $\text{Pic}_{\bar{D}}^r(X)$ embeds in a proper family $\mathcal{P}_{\bar{D}}^r(X) \rightarrow \bar{D}$. Its fibres are the spaces $\mathcal{P}^r(X_z)$. We give a ‘hands-on’ construction of this family, taking $r = 0$, using period matrices, following [9].

Since the inclusion $j: X_0 \rightarrow X$ induces an isomorphism on cohomology, we can find classes $(\alpha_1, \dots, \alpha_{g-1}, \beta_1, \dots, \beta_{g-1})$ in $H^1(X; \mathbb{Z})$ such that $(\nu^*j^*\alpha_i, \nu^*j^*\beta_i)_i$ give a symplectic basis for $H^1(\tilde{X}_0; \mathbb{Z})$. By restriction we obtain, for every $s \in \bar{D}$, classes $\alpha_i(s), \beta_i(s) \in H^1(X_s; \mathbb{Z})$, spanning a symplectic summand.

There is, up to sign, a unique non-multiple class $\beta_g \in H^1(X; \mathbb{Z})$ such that $\nu^*j^*\beta_g = 0$. This induces $\beta_g(s) \in H^1(X_s; \mathbb{Z})$, orthogonal to $\alpha_i(s)$ and $\beta_i(s)$ for $i < g$. We can also find a family of classes $\alpha_g(s)$, $s \neq 0$, so as to obtain a complete symplectic basis

$(\alpha_1, \dots, \alpha_g, \beta_1, \dots, \beta_g)$ for $H^1(X_s; \mathbb{Z})$; however, these are only defined modulo addition of $\beta_g(s)$, because of their monodromy.

Formally, the $\alpha_i(s)$ and $\beta_i(s)$ are sections of the sheaf $R^1\pi_*\mathbb{Z}_X$ (though α_g is not a *global* section). The exponential sequence of sheaves on X induces an injective map

$$\iota: R^1\pi_*\mathbb{Z}_X \rightarrow R^1\pi_*\mathcal{O}_X.$$

The sheaf $R^1\pi_*\mathcal{O}_X$ is actually a rank g vector bundle, the dual of the Hodge bundle. We will not use its local triviality near $0 \in \bar{D}$, only on the complement \bar{D}^* , where it is clear; in fact $(\iota\alpha_1(s), \dots, \iota\alpha_g(s))$ gives a local trivialisation. So too does $(\iota\beta_1(s), \dots, \iota\beta_g(s))$, and they are related by the period matrix Z :

$$(\iota\alpha_1(s), \dots, \iota\alpha_g(s)) = -(\iota\beta_1(s), \dots, \iota\beta_g(s))Z(s). \quad (4.13)$$

The matrix $Z(s)$ has holomorphic entries, and, according to the Riemann period relations, it is symmetric with positive-definite imaginary part. The entry $Z_{gg}(s)$ is multivalued (since α_g is); its monodromy is $Z_{gg}(s) \rightarrow Z_{gg}(s) - 1$. The function $e^{2\pi i Z_{gg}}$ is bounded near 0, and hence extends to a function on \bar{D} which must have a simple zero at the origin (because of the monodromy); hence $Z_{gg}(s) - (2\pi i)^{-1} \log s$ is a single-valued holomorphic function on \bar{D} , non-vanishing at 0.

Let (e_1, \dots, e_g) be the standard basis for \mathbb{C}^g , and let $e'_i(s) = -\sum_{k=1}^g Z_{ki}(s)e_k$ (so e'_g is defined up to addition by e_g , and only when $s \neq 0$). Using the classes $\beta_i(s)$, we can identify $\text{Pic}_{\bar{D}}^0(X) \subset R^1\pi_*\mathcal{O}_X^\times$ with

$$(\bar{D} \times \mathbb{C}^g)/\Lambda, \quad \Lambda = \{(s, \lambda) : \lambda \in \Lambda_s\}$$

where

$$\begin{aligned} \Lambda_s &= \mathbb{Z}\{e_1, \dots, e_g, e'_1(s), \dots, e'_g(s)\}, \quad s \neq 0 \\ \Lambda_0 &= \mathbb{Z}\{e_1, \dots, e_g, e'_1(0), \dots, e'_{g-1}(0)\}. \end{aligned}$$

(The notation $\bar{D} \times \mathbb{C}^g/\Lambda$ is shorthand for the quotient of $\bar{D} \times \mathbb{C}^g$ by the relation $(s, x) \sim (s, x + \lambda)$, $\lambda \in \Lambda_s$.)

As previously noted, there is a (bounded) holomorphic function $h: \bar{D} \rightarrow \mathbb{C}$ such that $e'_g(s) = h(s) - \frac{\log s}{2\pi i} e_g$. Define a map

$$f: \mathbb{C}^{g-1} \times ((\bar{D} \times \bar{D}) \setminus (0, 0)) \rightarrow (\bar{D} \times \mathbb{C}^g)/\Lambda$$

by the formulae

$$(x, a, b) \mapsto \begin{cases} [ab, x + \frac{\log a}{2\pi i} e_g + \frac{h(ab)}{2}], & \text{if } |b| \leq |a|, \\ [ab, x - \frac{\log b}{2\pi i} e_g - \frac{h(ab)}{2}], & \text{if } |a| \leq |b|. \end{cases}$$

The two lines agree when both are operative.

Let $\Lambda'_s \subset \mathbb{C}^{g-1}$ be the lattice obtained by projecting Λ_s to $\mathbb{C}^{g-1} = \mathbb{C}^g / \langle e_g \rangle$. Take a set of balls $B_1, \dots, B_N \subset \mathbb{C}^{g-1}$ such that the following hold for any $s \in \bar{D}$: each $B_i \cap \Lambda'_s$ contains at most one point; and $\bigcup B_i$ is a connected set covering a fundamental domain for Λ'_s . Let f_i be the restriction of f to $B_i \times (\bar{D}^2 \setminus 0)$. The following lemma then follows from the definitions.

Lemma 4.3.3. *Each of the maps f_i is injective, and thus an open embedding. The complement of $\bigcup_{i=1}^n \text{im}(f_i)$ is relatively compact in $(\bar{D} \times \mathbb{C}^g) / \Lambda$.*

We can now define our compactified space,

$$\mathcal{P}_D^0(X) = \left((\bar{D} \times \mathbb{C}^g / \Lambda) \cup \bigcup (B_i \times \bar{D} \times \bar{D}) \right) / \sim,$$

where the \sim is the equivalence relation generated by the following:

$$\begin{aligned} (x, a, b)_i &\sim f_i(x, a, b)_i, & (a, b) &\in \bar{D} \times \bar{D} \setminus \{(0, 0)\}; \\ (x, a, b)_i &\sim (x, a, b)_j, & x &\in B_i \cap B_j; \\ (x, a, 0)_i &\sim (x + \lambda, a, 0)_j, & \lambda &\in \Lambda'_0; \\ (x, 0, b)_i &\sim (x + \lambda, 0, b)_j, & \lambda &\in \Lambda'_0. \end{aligned}$$

Here the subscripts keep track of the sets $B_i \times \bar{D} \times \bar{D}$.

Proposition 4.3.4. *The analytic space $\mathcal{P}_D^0(X)$ is a complex manifold, and the natural map $\pi: \mathcal{P}_D^0(X) \rightarrow \bar{D}$ is proper. It has 0 as only critical value, and its critical set $\text{crit}(\pi)$ is a $(g-1)$ -dimensional complex submanifold of $\mathcal{P}_D^0(X)$. At a critical point, π is modelled on $(z_0, z_1, \dots, z_g) \mapsto z_0 z_1$. The open set of regular values is naturally identified with $\text{Pic}_D^0(X)$, and the zero-fibre with $\mathcal{P}^0(X_0)$; in particular, $\text{crit}(\pi) \cong \text{Pic}^0(\widetilde{X}_0)$.*

Proof. Smoothness is immediate from the construction; compactness holds because the f_i cover the complement of a compact set in $(\bar{D} \times \mathbb{C}^g) / \Lambda$. The obvious map from $(\bar{D} \times \mathbb{C}^g) / \Lambda$ to the open set

$$\left((\bar{D} \times \mathbb{C}^g / \Lambda) \cup \bigcup B_i \times (\bar{D} \times \bar{D} \setminus (0, 0)) \right) / \sim$$

is invertible, and so identifies $\text{Pic}_D^0(X)$ with $\mathcal{P}_D^0(X)$ minus its compactification divisor. The statements about the critical points of π are trivial. \square

One can see directly that the natural ‘theta’ line bundle $\Theta \rightarrow \text{Pic}_D^0(X)$ extends over the compactification. This is the bundle

$$\Theta = \mathbb{C} \times \mathbb{C}^g \times \bar{D} / \sim,$$

where the equivalence relation \sim is generated by

$$(\xi, x + e_i, s) \sim (\xi, x, i), \quad (\xi, x + e_i, s) \sim (e^{-2\pi i x_i} \xi, x + e'_i, s)$$

for $i = 1, \dots, g$ (when $i = g$, the second relation only applies when $s \neq 0$).

The map f lifts to a map of line bundles \tilde{f} from the trivial line bundle over $\mathbb{C}^{g-1} \times (\bar{D} \times \bar{D} \setminus (0, 0))$ to Θ ,

$$\tilde{f}(\xi, x, a, b) = [\xi, f(x, a, b)].$$

One can use this to patch together trivial line bundles over the spaces $B_i \times \bar{D} \times \bar{D}$ with Θ , and so obtain a line bundle $\widehat{\Theta} \rightarrow \mathcal{P}_D^0(X)$.

4.3.3 The Hilbert scheme of r points on a nodal curve

For each nodal curve C , not necessarily compact, the r -point Hilbert scheme $\text{Hilb}^r(C)$ is a reduced, r -dimensional complex space with normal crossing singularities. Its points are in canonical bijection with the set

$$\{\text{coherent ideal sheaves } \mathcal{I} \subset \mathcal{O}_C : \sum_{z \in C} \dim_{\mathbb{C}}(\mathcal{O}_C/\mathcal{I})_z \text{ is finite and equal to } r\}.$$

There is an ideal sheaf $\mathcal{I}^{\text{univ}} \subset \mathcal{O}_{\text{Hilb}^r(C) \times C}$, flat over C , and the bijection sends $x \in \text{Hilb}^r(C)$ to the ideal obtained by specialising $\mathcal{I}^{\text{univ}}$ to $\{x\} \times C$. There is a cycle map from $\text{Hilb}^r(C)$ to the r th symmetric product of C , given by

$$\text{Hilb}^r(C) \rightarrow \text{Sym}^r(C), \quad \mathcal{I} \mapsto \sum_{z \in C} \dim_{\mathbb{C}}(\mathcal{O}_C/\mathcal{I})_z [z].$$

If C is non-singular this is an isomorphism, because any ideal in the local ring $\mathcal{O}_{\mathbb{C},0}$ is of the form (z^n) . Moreover, in this case the ideal $\mathcal{I}^{\text{univ}} \subset \mathcal{O}_{\text{Sym}^r(C) \times C}$ defines the universal divisor $\Delta^{\text{univ}} \subset \text{Sym}^r(C) \times C$.

Ideals in the local ring. We begin by describing the ideals in the local ring at a node. Let $\mathcal{O}_{\mathbb{C}^2,0}$ be the \mathbb{C} -algebra of germs of holomorphic functions at $0 \in \mathbb{C}^2$, and write $R = \mathcal{O}_{\mathbb{C}^2,0}/\langle xy \rangle$.

Lemma 4.3.5. *A proper ideal $I \subset R$ has finite length $r = \dim_{\mathbb{C}}(R/I)$ if and only if it is of the form*

$$I = \langle x^{m+1}, y^{r-m+1}, ax^m + by^{r-m} \rangle, \quad 1 \leq m \leq r-1,$$

for some $(a : b) \in \mathbb{P}^1$.

When $ab \neq 0$, this may be rewritten as $I = \langle ax^m + by^{r-m} \rangle$.

Proof. The formula evidently defines an ideal of length r . Conversely, given an ideal $I \neq R$ of finite length, let m and n be minimal such that $x^{m+1} \in I$ and $y^{n+1} \in I$. If I is not equal to $\langle x^{m+1}, y^{n+1} \rangle$ then it must contain an element of form $ax^k + by^l$, where $ab \neq 0$, $k < m$, and $l < n$. But then $x^{k+1} \in I$, so $k+1 = m$ by minimality of m . Similarly $n = l+1$. Hence $\langle x^{m+1}, y^{n+1}, ax^m + by^n \rangle \subset I$, and running the same argument again we see that equality holds. \square

Let $C_0 = \{(x, y) : xy = 0\}$. By the lemma, the length r ideal sheaves on C_0 , supported at 0, are naturally parametrised by a chain of $r-1$ copies of \mathbb{P}^1 . Write \mathbb{P}_m^1 for the m th copy, which has homogeneous coordinates $(a_m : b_m)$. A model for the chain E is given by the subspace of $\mathbb{P}_1^1 \times \cdots \times \mathbb{P}_{r-1}^1$ defined by the equations

$$b_1 a_2 = 0, \dots, b_{r-2} a_{r-1} = 0, \quad (4.14)$$

Abstractly,

$$E = E_1 \cup_{q_1 \sim p_2} E_2 \cup_{q_2 \sim p_3} \cdots \cup_{q_{r-2} \sim p_{r-1}} E_{r-1}, \quad (4.15)$$

where E_m is a copy of \mathbb{P}^1 , and $p_m = (0 : 1) \in E_m$, $q_m = (1 : 0) \in E_m$.

Defining the Hilbert scheme. Suppose that C has nodes x_0^1, \dots, x_0^k , with preimages $\{x_+^\alpha, x_-^\alpha\} = \nu^{-1}(x_0^\alpha)$ in the normalisation $\nu : \tilde{C} \rightarrow C$. The complex manifold $\text{Sym}^r(\tilde{C})$ contains codimension-2 submanifolds $S_\alpha := x_+^\alpha + x_-^\alpha + \text{Sym}^{r-2}(\tilde{C})$. Let $S = \bigcup_{\alpha=1}^k S_\alpha$, and form the blow-up

$$H := \text{Bl}_S(\text{Sym}^r(\tilde{C})).$$

Define $\text{Hilb}^r(C)$ as the colimit

$$\coprod_{\{1, \dots, k\}} \text{Sym}^{r-1}(\tilde{C}) \xrightarrow[i_-]{i_+} H \longrightarrow \text{Hilb}^r(C), \quad (4.16)$$

where the maps i_\pm are the proper transforms of the maps $(\alpha, \mathbf{z}) \mapsto x_\pm^\alpha + \mathbf{z} \in \text{Sym}^r(\tilde{C})$. That is, $\text{Hilb}^r(C)$ is the topological space obtained from H by gluing the two copies of $\text{Sym}^{r-1}(\tilde{C})$ associated to each node to one another. Note that in the blow up these two copies are *disjoint*, since before blowing up they intersect transversely. The quotient space has a natural analytic structure, with normal crossings along the gluing locus. When $k = 1$, the singular locus is smooth, isomorphic to $\text{Sym}^{r-1}(\tilde{C})$.

To examine the structure of this space we consider the model curve C_0 . Its normalisation $\tilde{C}_0 = C_+ \cup C_-$ is the disjoint union of two copies of \mathbb{C} , and

$$\text{Sym}^r(\tilde{C}) = \bigcup_{i=0}^r \text{Sym}^i(C_+) \times \text{Sym}^{r-i}(C_-) \cong \bigcup_{i=0}^r \mathbb{C}^i \times \mathbb{C}^{r-i}.$$

The isomorphism here is induced by elementary symmetric functions in the usual way. Writing the coordinates on $\mathbb{C}^i \times \mathbb{C}^{r-i}$ as $(\sigma_1, \dots, \sigma_i; \tau_1, \dots, \tau_{r-i})$, the submanifold $x_+ + x_- + \text{Sym}^{r-2}(\widetilde{C}_0)$ is defined by $\sigma_i = \tau_{r-i} = 0$. The blow-up H is also a disjoint union,

$$H = \bigcup_{i=0}^r H_i,$$

where $H_0 = \{0\} \times \mathbb{C}^r$, $H_r = \mathbb{C}^r \times \{0\}$, and for $i = 1, \dots, r-1$,

$$H_i = \{(\sigma_1, \dots, \sigma_i; \tau_1, \dots, \tau_{r-i}; (a : b)) \in \mathbb{C}^i \times \mathbb{C}^{r-i} \times \mathbb{P}^1 : \sigma_i b = \tau_{r-i} a\}.$$

We define a map $p: H \rightarrow \mathbb{C}^r \times \mathbb{C}^r \times E$, as follows:

$$\begin{aligned} H_0 \ni (\tau_1, \dots, \tau_r) &\mapsto (0_r; (\tau_1, \dots, \tau_r); (0 : 1)_1), \\ H_i \ni (\sigma_1, \dots, \sigma_i; \tau_1, \dots, \tau_{r-i}; (a : b)) &\mapsto (\sigma_1, \dots, \sigma_i, 0_{r-i}; \tau_1, \dots, \tau_{r-i}, 0_i; (a : b)_i), \\ H_r \ni (\sigma_1, \dots, \sigma_r) &\mapsto (\sigma_1, \dots, \sigma_r; 0_r; (1 : 0)_{r-1}), \end{aligned} \quad (4.17)$$

The restriction $p|_{H_i}$ is an embedding, but p is not injective globally. One can easily check that it descends to $\text{Hilb}^r(C)$, and that the resulting map is an embedding. Its image is the subspace of $\mathbb{C}^r \times \mathbb{C}^r \times E$ cut out by the equations

$$x_1 y_r = 0, \quad x_r y_1 = 0; \quad x_1 b_1 = y_{r-1} a_1, \dots, x_{r-1} b_{r-1} = y_1 a_{r-1} \quad (4.18)$$

(here $(x_1, \dots, x_r; y_1, \dots, y_r)$ are coordinates on $\mathbb{C}^r \times \mathbb{C}^r$. These are explicit equations for $\text{Hilb}^r(C_0)$).

Cycle map. The image of the projection $\text{Hilb}^r(C_0) \rightarrow \mathbb{C}^{2r}$ is cut out by the equations

$$x_1 y_r = x_2 y_{r-1} = \dots = x_r y_1 = 0. \quad (4.19)$$

The cycle map $\text{Hilb}^r(C_0) \rightarrow \text{Sym}^{r-1}(C_0)$ corresponds, via elementary symmetric functions, to this projection map.

Universal ideal sheaf. We now define the universal divisor $\Delta^{\text{univ}} \subset \text{Hilb}^r(C) \times C$ which corresponds to the universal ideal sheaf $\mathcal{I}^{\text{univ}}$. We do this first in the case where $C = C_0 = \{zw = 0\}$ with components $C_+ = \{w = 0\}$, $C_- = \{z = 0\}$. Introduce the polynomials

$$f_i(x; z) = z^i + \sum_{j=1}^i (-1)^j x_j z^{i-j} = 0, \quad i = 1, \dots, r.$$

The defining equations for Δ^{univ} are then

$$\begin{aligned}
f_r(x; z) &= 0; \\
b_{r-1}f_{r-1}(x; z) + a_{r-1}f_1(y; w) &= 0; \\
&\vdots \\
b_1f_1(x; z) + a_1f_{r-1}(y; w) &= 0; \\
f_r(y; w) &= 0.
\end{aligned} \tag{4.20}$$

On the component $p(H_i)$, Δ^{univ} is cut out by

$$z^i f_i(x; z) = 0, \quad w^i f_{r-i}(y; w) = 0, \quad a_i f_{r-i}(y; w) + b_i f_i(x; z) = 0. \tag{4.21}$$

From this one can see that (i) Δ^{univ} does give a bijection with the ideal sheaves of length r (by comparing with our lemma describing these sheaves at the node); (ii) over the non-singular region $\text{Hilb}^r(C_0 \setminus \{x_0\}) \times (C \setminus \{x_0\}) = \text{Sym}^r(C_0 \setminus \{x_0\}) \times (C \setminus \{x_0\})$, Δ^{univ} coincides with the universal divisor for the symmetric product. It follows that one can define Δ^{univ} for arbitrary nodal curves by patching together locally-defined divisors.

Abel-Jacobi map. Suppose that C is irreducible with one node. Any global section $s \neq 0$ of a sheaf $\mathcal{F} \in \mathcal{P}^r(C)$ defines an ideal sheaf $\mathcal{I}_s \in \text{Hilb}^r(C)$. Moreover, there is a holomorphic map

$$\mathbf{aj}: \text{Hilb}^r(C) \rightarrow \mathcal{P}^r(C)$$

such that $\mathbf{aj}^{-1}(\mathcal{F}) = \mathbb{P}H^0(C, \mathcal{F})$. An ideal sheaf \mathcal{I} is mapped by \mathbf{aj} to the isomorphism class of the torsion-free sheaf $\mathcal{O}_C/\mathcal{I}$.

Remark 4.3.6. *The Abel-Jacobi map can be seen in terms of our explicit constructions. Here is a sketch. We begin with the Abel-Jacobi map $\mathbf{aj}_{\tilde{C}}: \text{Sym}^r(\tilde{C}) \rightarrow \text{Pic}^r(\tilde{C})$. This lifts to a map $H \rightarrow \mathbb{P}E$, that is, a map between the normalisations $\widehat{\text{Hilb}^r(C)} \rightarrow \widehat{\mathcal{P}^r(C)}$: the reason is that the normal bundle $N := N_{S/\text{Sym}^r(\tilde{C})}$ is isomorphic to $(\mathbf{aj}_{\tilde{C}}^* E)|_S$. The proper transform of S is a copy of $\mathbb{P}N$, and it has a neighbourhood $U \subset H$ which embeds into $\mathcal{O}_{\mathbb{P}N}(-1)$; this maps to $\mathcal{O}_{\mathbb{P}E}(-1)$, and hence to $\mathbb{P}E$. The resulting map $U \rightarrow \mathbb{P}E$ may be identified, away from the zero sections, with $\mathbf{aj}_{\tilde{C}}$. It therefore globalises to a map $H \rightarrow \mathbb{P}E$, as desired. This map descends, in a unique way, to a map $\text{Hilb}^r(C) \rightarrow \mathcal{P}^r(C)$; this is \mathbf{aj} .*

For clarity, we state the following proposition for the case where there is just one node.

Proposition 4.3.7. (a) *The set of singular points in $\text{Hilb}^r(C)$ is the preimage, under the cycle map, of $x_0 + \text{Sym}^{r-1}(C) \subset \text{Sym}^r(C)$. It is a normal crossing divisor, itself smooth. Indeed, $\text{Hilb}^r(C)_{\text{sing}}$ is an embedded copy of $\text{Sym}^{r-1}(\tilde{C})$, and the universal sheaf*

over $\mathrm{Hilb}^r(C) \times C$ pulls back to the universal sheaf over $\mathrm{Sym}^{r-1}(\tilde{C}) \times \tilde{C}$ under the natural map $\mathrm{Sym}^{r-1}(\tilde{C}) \times \tilde{C} \rightarrow \mathrm{Hilb}^r(C) \times C$.

(b) Along the normal crossing divisor $\mathrm{Sym}^{r-1}(\tilde{C})$ there are two normal line bundles \mathcal{N} and \mathcal{N}^\vee , corresponding to the two branches of $\mathrm{Hilb}^r(C)$. These are topologically trivial.

(c) There is a well-defined and holomorphic Abel-Jacobi map $\mathbf{aj}: \mathrm{Hilb}^r(C) \rightarrow \mathcal{P}^r(C)$, whose fibres are the projective spaces of sections of sheaves representing points in $\mathcal{P}^r(C)$. It sends the normal crossing divisor in $\mathrm{Hilb}^r(C)$ to that in $\mathcal{P}^r(C)$, and this restricted map corresponds to the Abel-Jacobi map $\mathrm{Sym}^{r-1}(\tilde{C}) \rightarrow \mathrm{Pic}^{r-1}(\tilde{C})$.

Proof. We have seen most of this already, but we should still prove that the normal line bundles to the normal crossing divisor are dual to one another and have Chern class zero.

A general model here consists of a pair Y_1, Y_2 of codimension 1 complex submanifolds in a complex manifold X , with transverse intersection $Z = Y_1 \cap Y_2$. In the blow-up $\tilde{X} = \mathrm{Bl}_Z(X)$, the proper transforms \tilde{Y}_i of Y_i have normal bundles $N_{\tilde{Y}_i/\tilde{X}} \rightarrow \tilde{Y}_i \cong Y_i$ isomorphic to $N_{Y_i/X} \otimes \mathcal{O}_{Y_i}(-Z)$.

Let δ_-^n (resp. δ_+^n) be the image in $\mathrm{Sym}^n(\tilde{C})$ of $x_- + \mathrm{Sym}^{n-1}(\tilde{C})$, (resp. $x_+ + \mathrm{Sym}^{n-1}(\tilde{C})$). Applying the model to $Y_1 = \delta_-^r, Y_2 = \delta_+^r$ in $X = \mathrm{Sym}^r(\tilde{C})$, we obtain

$$\begin{aligned} N_{\tilde{Y}_1/\tilde{X}} &\cong N_{\delta_-^r/\mathrm{Sym}^r(\tilde{C})} \otimes \mathcal{O}_{\delta_-^r}(-\delta_+^{r-1}) \\ &\cong N_{\delta_-^r/\mathrm{Sym}^r(\tilde{C})} \otimes \mathcal{O}_{\delta_-^r}(-\delta_-^{r-1}) \otimes \mathcal{O}_{\delta_-^r}(\delta_-^{r-1} - \delta_+^{r-1}) \\ &\cong N_{\delta_-^r/\mathrm{Sym}^r(\tilde{C})} \otimes (\mathcal{O}_X(-\delta_-^r))|_{\delta_-^r} \otimes \mathcal{O}_{\delta_-^r}(\delta_-^{r-1} - \delta_+^{r-1}) \\ &\cong \mathcal{O}_{\delta_-^r}(\delta_-^{r-1} - \delta_+^{r-1}), \end{aligned}$$

where the last isomorphism uses the adjunction formula. Similarly,

$$N_{\tilde{Y}_2/\tilde{X}} \cong \mathcal{O}_{\delta_+^r}(\delta_+^{r-1} - \delta_-^{r-1}),$$

which is dual to $\mathcal{O}_{\delta_-^r}(\delta_-^{r-1} - \delta_+^{r-1})$. The two line bundles $N_{\tilde{Y}_i/\tilde{X}}$ can naturally be identified with the normal bundles along the normal crossing divisor in $\mathrm{Hilb}^r(C)$. Since the divisors δ_+^{r-1} and δ_-^{r-1} in $\mathrm{Sym}^{r-1}(\tilde{C})$ are homologous, the line bundles are topologically trivial (the divisors will not usually be linearly equivalent, so the bundles need not be holomorphically trivial). □

4.3.4 The relative Hilbert scheme

We describe the relative Hilbert scheme $\mathrm{Hilb}^r(\pi) \rightarrow \mathbb{C}$ of the map $\pi: \mathbb{C}^2 \rightarrow \mathbb{C}, (z, w) \mapsto zw$, using Ran's model. The most striking feature of this space, uncovered by Donaldson and Smith [9, 49], is that, like the compactified relative Jacobian, it is *non-singular*.

Assume $r > 1$. The construction begins with the definition of complex surfaces X_r , starting from

$$X_2 := \mathbb{P}^1 \times \mathbb{C}. \quad (4.22)$$

For higher r , X_r is the subspace of $\mathbb{P}_1^1 \times \cdots \times \mathbb{P}_{r-1}^1 \times \mathbb{C}$ defined by the equations

$$b_i a_{i+1} = t a_i b_{i+1}, \quad i = 1, \dots, r-2. \quad (4.23)$$

This is non-singular of dimension 2, with non-singular rational fibres over points in $\mathbb{C}^* \subset \mathbb{C}$. The zero-fibre is the rational chain E constructed earlier, and one can obtain X_r from X_{r-1} by blowing up a point in the last curve in the chain. Now,

$$\text{Hilb}^r(\pi) \subset \mathbb{C}^r \times \mathbb{C}^r \times X_r$$

is the complex subspace defined by the equations

$$\begin{aligned} a_1 y_r &= t b_1; \\ x_i b_i &= y_{r-i} a_i, \quad i = 1, \dots, r-1; \\ b_{r-1} x_r &= t a_{r-1}. \end{aligned} \quad (4.24)$$

This maps to \mathbb{C} (via X_r), and the zero-fibre is $\text{Hilb}^r(C_0)$ (the defining equations reduce to those we wrote down earlier). The image of the projection $\text{Hilb}^r(\pi) \rightarrow \mathbb{C}^r \times \mathbb{C}^r \times \mathbb{C}$ is cut out by the equations

$$\begin{aligned} x_1 y_r &= t y_{r-1}; \\ x_{i+1} y_{r-i} &= t y_{r-i-1} x_i, \quad i = 1, \dots, r-1 \\ x_r y_1 &= t x_{r-1}, \end{aligned} \quad (4.25)$$

and the projection is an embedding except where $t = 0$. On the other hand, the relative symmetric product $\text{Sym}^r(\pi)$ also embeds in $\mathbb{C}^r \times \mathbb{C}^r \times \mathbb{C}$, using the elementary symmetric functions in z and in w , and its image is precisely the image of $\text{Hilb}^r(\pi)$. Thus, the projection to $\mathbb{C}^r \times \mathbb{C}^r \times \mathbb{C}$ is an explicit form of the cycle map.

Proposition 4.3.8. *$\text{Hilb}^r(\pi)$ is non-singular, hence a complex manifold of dimension $r+1$; it is equipped with a proper projection map $\pi_r: \text{Hilb}^r(\pi) \rightarrow \mathbb{C}$. The set of critical points, $Z := \text{crit}(\pi)$, is the normal crossing divisor in $\text{Hilb}^r(C_0)$. There is a commutative diagram*

$$\begin{array}{ccc} \mathbb{C}^2 \times Z & \xrightarrow{e} & \text{Hilb}^r(\pi) \\ & \searrow \pi \circ \text{pr}_1 & \swarrow \pi_r \\ & \mathbb{C} & \end{array}$$

where e is a holomorphic immersion which restricts to the identity on $\{0\} \times Z = Z$.

Proof. To prove non-singularity of $\text{Hilb}^r(\pi)$ near a point P , and to analyse π_r there, we distinguish three cases:

- $\pi_r(P) = 0$ and P lies on a non-singular point of $\text{Hilb}^r(C_0)$. Say $P' \in p(H_i)$. Near P' , the projection to $\mathbb{C}^i \times \mathbb{C}^{r-i} \times \mathbb{C}$ given by the variables $(x_1, \dots, x_i; y_1, \dots, y_{r-i}; t)$ is an open embedding. The proof that it is injective is a simple matter of eliminating the remaining variables, one by one; to see that it is locally surjective one reverses this procedure. To do so, first determine $(a_i : b_i)$, using the fact that x_i and y_{r-i} are not both zero, which follows from the hypothesis on P . The remaining variables are easy to reconstruct. Hence $\text{Hilb}^r(\pi)$ is non-singular at P , and π_r is a submersion.
- $\pi_r(P) \neq 0$. The projection to $\mathbb{C}^r \times \mathbb{C}$ given by $(x_1, \dots, x_r; t)$ is locally an embedding.
- $\pi_r(P) = 0$ and P lies on a singular point of $\text{Hilb}^r(C_0)$. This is a more interesting case, and we deal with it below by constructing the embedding e .

For integers i with $1 \leq i \leq r-2$, we consider the open subset $U_i \subset \text{Hilb}^r(\pi)$ on which $a_i \neq 0$ and $b_{i+1} \neq 0$. This contains $p(H_i) \cap p(H_{i+1})$. Define a map

$$f_i = (b_i/a_i, a_{i+1}/b_{i+1}; x_1, \dots, x_i; y_1, \dots, y_{r-i-1}): U_i \rightarrow \mathbb{C}^2 \times \mathbb{C}^i \times \mathbb{C}^{r-i-1}.$$

One can check that f_i is an isomorphism. We have $\pi_r = \pi \circ \text{pr}_{\mathbb{C}^2} \circ f_i$. Hence $\text{Hilb}^r(\pi)$ is smooth along $p(H_i) \cap p(H_{i+1})$, and π_r has a simple quadratic singularity there.

Similar remarks apply if we define $U_0 \subset \text{Hilb}^r(\pi)$ by $\{b_1 \neq 0\}$, and U_{r-1} by $\{a_{r-1} \neq 0\}$, so that U_0 (resp. U_{r-1}) contains $p(H_0) \cap p(H_1)$ (resp. $p(H_{r-1}) \cap p(H_r)$), and define maps

$$\begin{aligned} f_0 &= (a_1/b_1, y_r; y_1, \dots, y_{r-1}): U_0 \rightarrow \mathbb{C}^2 \times \mathbb{C}^{r-1}, \\ f_{r-1} &= (b_{r-1}/a_{r-1}; x_1, \dots, x_{r-1}): U_{r-1} \rightarrow \mathbb{C}^2 \times \mathbb{C}^{r-1}. \end{aligned}$$

Now set

$$j_i: \mathbb{C}^i \times \mathbb{C}^{r-i-1} \rightarrow Z, \quad (x; y) \mapsto (x_1, \dots, x_i, 0_{r-i}, y_1, \dots, y_{r-i-1}, 0_{i+1}; p_i);$$

this is an isomorphism onto the i th connected component Z_i of Z (recall $p_i \in E \subset X_r$ from our earlier discussion). Now set

$$e_i = f_i^{-1} \circ (\text{id}_{\mathbb{C}^2} \times j_i)^{-1}: \mathbb{C}^2 \times Z_i \rightarrow U_i.$$

Notice that e_i restricts to the identity map on $p(H_i) \cap p(H_{i+1})$. Thus the map

$$e = \prod_{i=1}^{r+1} e_i: \mathbb{C}^2 \times Z \rightarrow \bigcup_i U_i \subset \text{Hilb}^r(\pi)$$

fulfils the requirements of the proposition. □

Further remarks on the relative Hilbert scheme.

- Having dealt with the local model, one can define $\text{Hilb}_S^r(X)$ for a general Lefschetz fibration $\pi: X \rightarrow S$ —with integrable complex structure near the critical points, as usual—by patching the local model with the non-singular part of the relative symmetric product.
- There is a universal divisor $\Delta^{\text{univ}} \subset \text{Hilb}_S^r(X) \times_S X$. In the local model, $\text{Hilb}_{\mathbb{C}}^r(\mathbb{C}^2) \times_{\mathbb{C}} \mathbb{C}^2$, this is defined by exactly the same equations as those that gave the universal divisor for $\text{Hilb}^r(C_0)$. The universal divisor induces a bijection between ideal sheaves of colength r in the fibres of π and points in $\text{Hilb}^r(\pi)$.
- The global Abel-Jacobi map $\text{Hilb}_S^r(X) \rightarrow \mathcal{P}_S^r(X)$ over S is pseudoholomorphic. I do not give the proof here, but nor do I use this fact.

Cohomology. Because of the universal divisor Δ^{univ} , and its dual class $\ell \in H^2(X \times_S \text{Hilb}_S^r(X); \mathbb{Z})$, there are operations

$$H^*(X; \mathbb{Z}) \rightarrow H^{*+2k-2}(\text{Hilb}_S^r(X); \mathbb{Z}), \quad c \mapsto c^{(k)},$$

defined in just the same way as for relative symmetric products (see (2.2)). Consider in particular the Hilbert scheme $\text{Hilb}_D^r(E)$ of an elementary Lefschetz fibration $E \rightarrow \bar{D}$. We look at the relation between the cohomologies of the smooth fibre $\text{Sym}^r(E_1)$ and of $\text{Hilb}_D^r(E)^{\text{crit}} = \text{Sym}^{r-1}(\tilde{E}_0)$ brought about by the inclusion maps

$$\text{Sym}^r(E_1) \xrightarrow{i_1} \text{Hilb}_D^r(E) \xleftarrow{i_{\text{crit}}} \text{Sym}^{r-1}(\tilde{E}_0).$$

Write $j_s: E_s \rightarrow E$ for the inclusion of the fibre E_s of E .

Lemma 4.3.9. *For any $c \in H^*(E; \mathbb{Z})$, we have $i_1^* c^{(1)} = (j_1^* c)^{(1)}$ and $i_{\text{crit}}^* c^{(1)} = (\nu^* j_0^* E_0)^{(1)}$, where $\nu: \tilde{E}_0 \rightarrow E_0$ is the normalisation map.*

Proof. The universal sheaf over $E \times_D \text{Hilb}_D^r(E)$ pulls back to the universal sheaves over $E_1 \times \text{Sym}^r(E_1)$ (clear) and over $\tilde{E}_0 \times \text{Sym}^{r-1}(\tilde{E}_0)$ (Prop. 4.3.7). □

There is an internal direct sum decomposition $H^*(E_1; \mathbb{Z}) = K \oplus \langle \lambda \rangle$, where $K = \text{im}(H^*(E; \mathbb{Z}) \rightarrow H^*(E_1; \mathbb{Z}))$, and $\lambda = \text{PD}[L]$ (λ may be zero). Lemma 4.3.9 implies that

$$i_{\text{crit}}^* \circ (i_1^*)^{-1}(x^{(1)}) = (\nu^* \circ j_0^* \circ (j_1^*)^{-1} x)^{(1)}, \quad x \in K. \quad (4.26)$$

(The map $\nu^* \circ j_0^* \circ (j_1^*)^{-1}$ is well-defined on K , and similarly $i_{\text{crit}}^* \circ (i_1^*)^{-1}$ is well-defined on $K^{(1)}$.) In particular,

$$i_{\text{crit}}^* \circ (i_1^*)^{-1}(\eta_{E_1}) = \eta_{\tilde{E}_0}, \quad (4.27)$$

where the η -classes come from the orientation classes: $\eta_{E_1} = (o_{E_1})^{(1)}$, $\eta_{\tilde{E}_0} = (o_{\tilde{E}_0})^{(1)}$.

If L is separating, (4.26) tells us all we need to know: there is a commutative square

$$\begin{array}{ccc} \Lambda^* H^*(E_1; \mathbb{Z}) & \xrightarrow{\Lambda^*(\nu^* j_0^* j_1^{*-1})} & \Lambda^* H^*(\tilde{E}_0; \mathbb{Z}) \\ \downarrow & & \downarrow \\ H^*(\text{Sym}^r(E_1); \mathbb{Z}) & \xrightarrow{i_{\text{crit}}^* \circ (i_1^*)^{-1}} & H^*(\text{Sym}^{r-1}(\tilde{E}_0); \mathbb{Z}) \end{array}$$

in which the vertical maps are onto.

If L is non-separating then we can write $H^*(E; \mathbb{Z}) = I \oplus \langle \lambda, \lambda' \rangle$, where $\lambda' \in H^1(E_1; \mathbb{Z})$ is dual to a loop L' with intersection number $L \cdot L' = 1$, and where I is the group of monodromy-invariant classes. Notice that the class $\lambda^{(1)} \cup \lambda'^{(1)} \in H^2(\text{Sym}^r(E_1); \mathbb{Z})$ is also monodromy-invariant, as a consequence of the Picard-Lefschetz formula for the monodromy of E . In this case there is a commutative square

$$\begin{array}{ccc} \langle 1, \lambda \wedge \lambda' \rangle \wedge \Lambda^* I & \xrightarrow{\Lambda^*(\nu^* j_0^* j_1^{*-1})} & \Lambda^* H^*(\tilde{E}_0; \mathbb{Z}) \\ \downarrow & & \downarrow \\ H^*(\text{Sym}^r(E_1); \mathbb{Z}) & \xrightarrow{i_{\text{crit}}^* \circ (i_1^*)^{-1}} & H^*(\text{Sym}^{r-1}(\tilde{E}_0); \mathbb{Z}) \end{array}$$

where the upper-left group is a subgroup of $\Lambda^* H^1(E_1; \mathbb{Z})$. For example, we have

$$i_{\text{crit}}^* \circ (i_1^*)^{-1}(\theta_{X_1}) = \theta_{\tilde{E}_0}. \quad (4.28)$$

4.4 Constructing the Lagrangian boundary conditions

4.4.1 Vanishing cycles as Lagrangian correspondences

We now consider the symplectic geometry of simple quadratic fibrations, as defined in Chapter 3. This is a fibred version of Picard-Lefschetz theory—it is the complex analogue of Morse-Bott theory.² A complete account of symplectic Picard-Lefschetz theory over the disk appears in [46]. The fibred generalisation has not yet been fully worked out, and there are technical difficulties to overcome in seeking to prove, for example, that the monodromy of a simple quadratic fibration is Hamiltonian-isotopic to a ‘model fibred Dehn twist’ about its vanishing cycle. Whilst the computations of Lagrangian matching invariants in this thesis only use the simplest properties of vanishing cycles, I expect that to make progress one will need a more precise understanding. In particular, I expect that precision will be vital in any proof of the conjectural exact triangle discussed at the end of this chapter. This is one reason for giving a rather lengthy account of the geometry of the Hilbert scheme.

²In a more systematic treatment one would also consider complex critical manifolds of arbitrary codimension, with non-degenerate normal Hessian forms.

Recall that a simple quadratic fibration is defined by data $(X, \pi, \Omega, J_0, j_0)$, where π is a map to a surface S . Take a smoothly embedded path $\gamma: [a, b] \hookrightarrow S$ such that $\gamma^{-1}(\pi(X^{\text{crit}})) = \{b\}$. Parallel transport along $\gamma|_{[a, b-\epsilon]}$ is a symplectomorphism

$$l_\epsilon: X_{\gamma(a)} \rightarrow X_{\gamma(a-\epsilon)}, \quad 0 < \epsilon \leq b - a.$$

Since π is proper, the points $l_\epsilon(x)$ ($x \in X_{\gamma(a)}$ fixed) have a limit $l(x) \in X_{\gamma(b)}$ as $\epsilon \rightarrow 0$. The **vanishing cycle** $V_\gamma \subset X_{\gamma(a)}$ associated to the path is the set $l^{-1}(Z)$ of points which flow into the critical set $X_{\gamma(b)}^{\text{crit}}$. The **thimble** is $W_\gamma := X_{\gamma(b)}^{\text{crit}} \cup \bigcup_\epsilon l_\epsilon(V)$.

Lemma 4.4.1. *W_γ is a smooth submanifold of X with boundary $V_\gamma \subset X_{\gamma(a)}$. The limit map $W_\gamma \rightarrow X_{\gamma(b)}^{\text{crit}}$ has a structure of disk bundle in a real two-plane vector bundle, with $l: V_\gamma \rightarrow X_{\gamma(b)}^{\text{crit}}$ as unit circle bundle.*

A proof is provided by Seidel and Smith [47], and we briefly recall the argument. One considers a symplectic form $\Omega + \pi^*\beta$ on X , and a function h on S which is equal to γ^{-1} on $\text{im}(\gamma)$. Then the horizontal lift of $\gamma'(t)$ coincides, up to positive reparametrisations, with the Hamiltonian vector field v of $\pi \circ h$. On the other hand, there exist local ‘holomorphic Morse-Bott’ charts in which π takes the form $(z_1, \dots, z_{n+1}) \mapsto z_1 z_2$; in these charts, v is given by $J_0 \nabla h$, where the gradient is defined using the Kähler metric associated with J_0 and $\Omega + \pi^*\beta$. But near $\text{im}(\gamma)$, h is a Morse-Bott function, with critical manifold of dimension $2n - 2$. The Hessian on the normal bundle is hyperbolic. Hence v is also a Morse-Bott vector field. A general result in Morse-Bott theory then says that its stable manifold W_γ is smooth, and identified with a disk bundle of a 2-plane bundle over $X_{\gamma(b)}^{\text{crit}}$.

Write (M, ω) for $(X_{\gamma(a)}, \Omega|_{X_{\gamma(a)}})$ and $(\bar{M}, \bar{\omega})$ for $(X_{\gamma(b)}^{\text{crit}}, \Omega|_{X_{\gamma(b)}^{\text{crit}}})$. We consider the product $(M \times \bar{M}, -\omega \oplus \bar{\omega})$, and the correspondence

$$\Lambda_\gamma := \{(x, l(x)) : x \in V_\gamma\} \subset M \times \bar{M}. \quad (4.29)$$

Lemma 4.4.2. *Λ_γ is a Lagrangian submanifold of $M \times \bar{M}$.*

Proof. The previous lemma shows that Λ_γ is a middle-dimensional submanifold of $M \times \bar{M}$; we have to show that it is isotropic. A tangent vector to Λ_γ is a pair $(v, Dl(v))$ with $v \in TV_\gamma$, and the symplectic form evaluates on a pair of tangent vectors $(v_1, Dl(v_2)), (v_2, Dl(v_1))$ as

$$\begin{aligned} -\omega(v_1, v_2) + l^* \bar{\omega}(v_1, v_2) &= -\omega(v_1, v_2) + \Omega(Dl(v_1), Dl(v_2)) \\ &= -\omega(v_1, v_2) + \lim_{\epsilon \rightarrow 0} \Omega(Dl(v_1), Dl(v_2)) \\ &= -\omega(v_1, v_2) + \lim_{\epsilon \rightarrow 0} l_\epsilon^* \omega(v_1, v_2) \\ &= 0, \end{aligned}$$

the last equality being due to the fact that l_ϵ is a symplectomorphism. □

The proof also makes clear that when the normal crossing divisor in the singular fibre has two distinct normal line bundles which are topologically trivial, the vanishing cycle is a *trivial* circle-bundle over \bar{M} .

We now apply these ideas to the relative Hilbert scheme $\text{Hilb}_{\bar{D}}^r(E)$ of an elementary Lefschetz fibration $E \rightarrow \bar{D}$. For γ we take the ray $[0, 1] \subset \bar{D}$. Since the critical set $\text{Hilb}_{\bar{D}}^r(E)^{\text{crit}}$ is identified with $\text{Sym}^{r-1}(\tilde{E}_0)$, the vanishing cycle is a Lagrangian subspace

$$\Lambda \subset \text{Sym}^r(E_1) \times \text{Sym}^{r-1}(\tilde{E}_0);$$

the projection on the first factor is injective, and on the second factor a trivial S^1 -bundle.

Likewise, the Picard family of E has a vanishing cycle contained in $\text{Pic}^r(E_1) \times \text{Pic}^{r-1}(\tilde{E}_0)$ which is a trivial circle-bundle over $\text{Pic}^{r-1}(\tilde{E}_0)$.

Remark 4.4.3. *In the present work we will not need to give an explicit representative of the isotopy class of Λ , but we do make two observations about Λ' , its image in $\text{Sym}^r(E_1)$. Consider E_1 as the union $\Sigma_0 \cup A$ of two surfaces along their common boundary, with A an annulus $A[\epsilon, \epsilon^{-1}]$ embedded in E_1 so that L is its unit circle. Then $\text{Sym}^r(E_1) = \bigcup_{i=0}^r S_i$, where $S_i = \text{Sym}^i(A) \times \text{Sym}^{r-i}(\Sigma_0)$. Now Λ' is isotopic to a hypersurface Λ'' such that*

$$\Lambda'' \cap S_0 = \emptyset, \tag{4.30}$$

$$\Lambda'' \cap S_1 = L \times \text{Sym}^{r-1}(\Sigma_0), \tag{4.31}$$

$$\Lambda'' \cap S_r = \{[z_1, \dots, z_r] \in \text{Sym}^r(A) : |z_1 \dots z_r| = 1\}. \tag{4.32}$$

Consider first the ‘local’ relative Hilbert scheme $\text{Hilb}_{\mathbb{C}}^r(zw = 0)$. This has singular set $\text{Sym}^{r-1}(\mathbb{C} \sqcup \mathbb{C})$, and this has a neighbourhood $\mathbb{C}^2 \times \text{Sym}^{r-1}(\mathbb{C} \sqcup \mathbb{C})$ with the map to \mathbb{C} given by $(a, b, \mathbf{z}) \mapsto ab$. If we gave this a product symplectic form, $\omega_{\mathbb{C}^2} \oplus \omega'$, standard on the first factor, then the vanishing cycle would be as in the right-hand side (4.32), as one sees by examining the explicit formulae (see Section 4.3.4) defining the Hilbert scheme. The other two lines likewise arise from model situations. Now, realistically the symplectic forms will not be standard ones. However, one can isotope the connection defining the symplectic parallel transport, not through symplectic connections but through connections for which the vanishing cycle remains well-defined, so that after the isotopy it coincides (on suitable open sets) with the parallel transport one obtains from the models.

4.4.2 Some auxiliary spaces

Definition 4.4.4. A family of simple quadratic fibrations parametrised by a manifold B is given by data $(\mathcal{X}, \pi_B, \Omega, \mathbf{J}_0, \mathbf{j}_0)$:

- $\pi_B: \mathcal{X} \rightarrow B \times S$ is a smooth, proper map such that $\text{pr}_1 \circ \pi_B: \mathcal{X} \rightarrow B$ is a fibre bundle. Write $T^\vee \mathcal{X} = \ker(D(\text{pr}_1 \circ \pi_B))$ and $\mathcal{X}^{\text{crit}} := \text{crit}(\pi_B)$.

- $\Omega \in Z_{\mathcal{X}}^2$.
- \mathbf{J}_0 (resp. \mathbf{j}_0) is the germ of a complex structure on $T^v\mathcal{X}$ (resp. on $\ker D(\text{pr}_1) \subset T(B \times S)$) near $\mathcal{X}^{\text{crit}}$ (resp. near $\pi_B(\mathcal{X}^{\text{crit}})$). Moreover, the restriction of \mathbf{J}_0 to any fibre of $\text{pr}_1 \circ B$ is (the germ near the critical set of) an integrable complex structure on this fibre.

As the name suggests, the defining condition is that, for each $b \in B$, the restricted data $(\mathcal{X}_b, \pi_b, \Omega_b, \mathbf{J}_{0,b}, \mathbf{j}_{0,b})$ constitute a simple quadratic fibration over S .

The case we shall consider has $B = S^1$ and $S = \bar{D}$. In this case we can consider the family of rays

$$\Gamma: S^1 \times [0, 1] \rightarrow S, \quad (t, s) \mapsto (t, 1 - s).$$

Then parallel transport along the family $\Gamma|_{[0, 1 - \epsilon]}$ gives a map

$$m_{\Gamma, \epsilon}: \mathcal{X}_{S^1 \times \{1\}} \rightarrow \mathcal{X}_{S^1 \times \{\epsilon\}}, \quad \epsilon \in [0, 1].$$

As observed in Chapter 2, $m_{\Gamma, \epsilon}^*(\Omega|_{\mathcal{X}_{S^1 \times \{\epsilon\}}}) = \Omega|_{\mathcal{X}_{S^1 \times \{1\}}}$. Set

$$V_{\Gamma} = \bigcup_{t \in S^1} V_{\Gamma|_{\{t\} \times [0, 1]}},$$

the union of the vanishing cycles of the rays making up Γ .

Lemma 4.4.5. *V_{Γ} is a submanifold of $\mathcal{X}_{S^1 \times \{1\}}$.*

Indeed, a straightforward generalisation of the earlier arguments identifies V_{Γ} with the unit circle bundle in the stable normal bundle over $\mathcal{X}^{\text{crit}}$ of a Morse-Bott vector field with critical manifold $\mathcal{X}^{\text{crit}}$.

The limits of the flow $m_{\Gamma, \epsilon}$ give a map $l: V_{\Gamma} \rightarrow \mathcal{X}^{\text{crit}}$. The **fibred vanishing cycle** of Γ is the correspondence

$$\Lambda_{\Gamma} = \{(x, lx) : x \in V_{\Gamma}\} \subset V_{\Gamma} \times_{S^1} \mathcal{X}^{\text{crit}}.$$

It is a smooth fibre bundle over S^1 . The same argument which shows that vanishing cycles are Lagrangian now shows:

Lemma 4.4.6. *Λ_{Γ} is isotropic with respect to the form $-\Omega \oplus \Omega$.*

We shall apply this lemma to a family of relative Hilbert schemes associated with the following family E .

Construction 4.4.7. *Suppose given a symplectic surface (Σ, ω) and a compact embedded 1-manifold L , invariant under $\phi \in \text{Aut}(\Sigma, \omega)$. Thus each fibre Σ_t of the mapping torus $T(\phi)$ contains a 1-manifold L_t . Then we can construct a family of simple quadratic fibrations parametrised by S^1 , $(E^5, \pi_{S^1}, \chi, \mathbf{J}_0, \mathbf{j}_0)$, where π_{S^1} maps to $S^1 \times \bar{D}$, such that:*

1. $\pi_{S^1}(E^{\text{crit}}) = S^1 \times \{0\} \subset S^1 \times \bar{D}$;

2. *there is an isomorphism of locally Hamiltonian fibrations over S^1*

$$\rho: T(\phi) \rightarrow (E, \pi, \chi)|_{S^1 \times \{1\}};$$

3. *for each $t \in S^1$, the elementary Lefschetz fibration $(E_t, \pi_t, \chi_t, \mathbf{J}_{0,t}, \mathbf{j}_{0,t})$ has the property that $\rho(L_t) \subset E_{(t,1)}$ is the vanishing cycle associated to the ray $\{t\} \times [0, 1]$.*

This is an S^1 -parametric version of Seidel's construction [46] of an elementary Lefschetz fibration starting from a surface with embedded circle L .

1. Define two maps $r_{\pm}: S^1 \rightarrow S^1$ by $r_+(z) = z$, $r_-(z) = \bar{z}$. The induced maps $r_{\pm}^*: T^*S^1 \rightarrow T^*S^1$ preserve the canonical symplectic form κ on T^*S^1 , and hence induce closed 2-forms κ_{\pm} on the mapping tori $T_{T^*S^1}(r_{\pm}^*)$. We abbreviate $T_{T^*S^1}(r_{\pm}^*)$ to T_{\pm} , and write $T_{\pm}^{\leq \lambda}$ for its subspace where the 1-form component has length $\leq \lambda$ with respect to a standard metric on T^*S^1 .

The circles $L_t \subset \Sigma_t$ sweep out an isotropic surface $Q = \bigcup_{t \in S^1} L_t$ in $(T(\phi), \omega_{\phi})$ which fibres over S^1 . A version of the Lagrangian neighbourhood theorem says that there is an $\lambda > 0$ and a fibre-preserving embedding

$$e_{\lambda}: T_{\pm}^{\leq \lambda} \rightarrow T(\phi),$$

which maps the zero-section to L , such that $e^*\omega_{\phi} = \kappa_{\pm}$.

2. We will build E by gluing one of two standard pieces E_0^{\pm} to another piece E_1 . The latter is defined as

$$E_1 = (T(\phi) \setminus \text{int}(\text{im } e_{\lambda/2})) \times \bar{D};$$

it maps to $S^1 \times \bar{D}$ by combining $T(\phi) \rightarrow S^1$ with the trivial projection to \bar{D} . The standard pieces, E_0^{\pm} , are regions in the mapping tori of linear automorphisms R_{\pm} of \mathbb{C}^2 : $R_+ = \text{id}_{\mathbb{C}^2}$, and $R_-(a, b) = (b, -\bar{a})$. Notice that the maps R_{\pm} preserve both the symplectic form $\omega_{\mathbb{C}^2}$ and the fibres of the map $q: \mathbb{C}^2 \rightarrow \mathbb{C}$, $(a, b) \mapsto a^2 + b^2$ —thus they lie in $\text{Sp}(4, \mathbb{R}) \cap \text{O}(2, \mathbb{C}) = \text{O}(2, \mathbb{R})$. In particular, $h(R_{\pm}x) = h(x)$, where $h(x) = \|x\|^4 - |q(x)|^2$. We set

$$E_0^{\pm} = \{[t, x] \in T(R_{\pm}) : |q(x)| \leq 1, h(x) \leq \sqrt{\lambda}/2\},$$

(the choice of $\lambda > 0$ is immaterial). This maps to $S^1 \times \bar{D}$ by $[t, x] \mapsto (t, q(x))$.

The following observations are extracted from [46, p.13-14]. First, the significance of h is that its level sets are symplectically perpendicular to the fibres of q . Radial parallel transport in \mathbb{C} gives a symplectic trivialisation

$$\Phi: \mathbb{C}^2 \setminus h^{-1}(0) \rightarrow \mathbb{C} \times (T^*S^1 \setminus S^1)$$

of $q: \mathbb{C}^2 \setminus h^{-1}(0) \rightarrow \mathbb{C}$ considered as a locally Hamiltonian fibration. It turns out that $|\Phi(x)|^2 = 4h(x)$, so the map

$$[\text{id}_{\mathbb{R}} \times \Phi]: E_0^{\pm} \setminus \{[t, x] : h(x) = 0\} \rightarrow \bar{D} \times T_{\pm}$$

has image $\bar{D} \times T_{\pm}^{\leq \lambda}$.

We can now glue E_0^{\pm} to E_1 : the gluing region in E_0^{\pm} is $h^{-1}(\sqrt{\lambda/8}, \sqrt{\lambda/2}]$, and that in E_1 is $\bar{D} \times (\text{im}(e_{\lambda}) \setminus \text{im}(e_{\lambda/2}))$. These are joined together using $[\text{id}_{\mathbb{R}} \times e_{\lambda} \circ \Phi]$ to form E . Because Φ preserves q , there is an induced map $\pi_{S^1}: E \rightarrow S^1 \times \bar{D}$.

3. We now verify that E has the right properties. By construction, $\{0\} \times S^1$ is the set of critical values of π_{S^1} ; it is clear from our gluing procedure how to construct the isomorphism ρ . For the (germs of) complex structures we use the standard ones on \mathbb{C}^2 and \mathbb{C} . The thimble for the path $\{t\} \times [-1, 0]$ is $\pi_{S^1}^{-1}(\{t\} \times [-1, 0]) \cap h^{-1}(0)$; this implies that the vanishing cycle is as claimed.

Remark 4.4.8. Each singular fibre $E_{(t,0)}$ has a normalisation $\tilde{E}_{(t,0)}$, defined using the complex structure near the singular points of $E_{(t,0)}$. Letting t vary, these form a manifold $\tilde{E}_{S^1 \times \{0\}}$ which fibres over S^1 . Moreover, essentially by definition of the two spaces involved, there is a fibred diffeomorphism $f: T(s_L \phi) \rightarrow \tilde{E}_{S^1 \times \{0\}}$.

We now form the set

$$\mathcal{E} = \bigcup_{t \in S^1} \text{Hilb}_{\bar{D}}^r(E_{\{t\} \times \bar{D}}).$$

This carries a natural topology and differentiable structure: one can consider it as a mapping torus, for example.

Lemma 4.4.9. The family $\pi_S^1: \mathcal{E} \rightarrow S^1 \times \bar{D}$ has a structure of family of simple quadratic fibrations $(\mathcal{X}, \pi_{S^1}, \Omega, \mathbf{J}_0, \mathbf{j}_0)$. One can choose Ω so that the map

$$\text{Sym}^r(\rho): \text{Sym}_{S^1}^r T(\phi) \rightarrow \mathcal{E}_{S^1 \times \{-1\}}$$

is an isomorphism of locally Hamiltonian fibrations, where one gives $\text{Sym}_{S^1}^r T(\phi)$ the form $\omega_{c, \tau}$; and so that there is a diffeomorphism

$$F: \text{Sym}_{S^1}^{r-1} T(s_L \phi) \rightarrow \mathcal{E}^{\text{crit}},$$

fibred over S^1 , such that $F^*(\Omega|_{\mathcal{E}^{\text{crit}}}) = \bar{\omega}_{\bar{e}, \tau}$.

Proof. There are identifications

$$\text{Sym}_{S^1}^{r-1} T(s_L \phi) \xrightarrow{\text{Sym}^r(f)} \text{Sym}_{S^1}^{r-1}(\tilde{E}_{S^1 \times \{0\}}) \cong \mathcal{E}^{\text{crit}}.$$

We write F' for the composite.

The main point of the proof is to understand how to extend cohomology classes from $\mathcal{E}_{S^1 \times \{1\}}$ to \mathcal{E} . There are operations

$$H^*(E; \mathbb{Z}) \rightarrow H^{*+2k-2}(\mathcal{E}; \mathbb{Z}), \quad c \mapsto c^{(k)}$$

defined using the universal divisor in a by-now-familiar way. These, moreover, are compatible with restriction to the critical set in the sense that one has

$$F'^*(c^{(k)}|_{\mathcal{E}^{\text{crit}}}) = (\nu^*(c|_{E^{\text{crit}}}))^{(k)}, \quad (4.33)$$

as we now prove. The formula is true in two special cases:

(a) (4.33) holds when c is induced by an invariant class in $H^k(E_{\{t\} \times \bar{D}}; \mathbb{Z})$.

(b) (4.33) holds when $c|_{E_{\{t\} \times \bar{D}}} = 0$. For then there exists $b \in H^{*-1}(E_{\{t\} \times \bar{D}}; \mathbb{Z})$ such that $\delta(b) = c$, where δ is the connecting map in the exact triangle for cohomology of mapping tori (cf. Appendix B). We already know that

$$F'_t(b^{(k)}|_{\mathcal{E}_{\{t\} \times \bar{D}}^{\text{crit}}}) = (\nu^*(b|_{E_{\{t\} \times S^1}^{\text{crit}}}))^{(k)},$$

on $\text{Sym}^{r-1}(\tilde{E}_{(t,0)})$. But it is clear that $(\delta b)^{(k)} = \delta(b^{(k)})$, so (4.33) holds for $c = \delta b$.

But $H^k(E; \mathbb{Z})$ is generated by classes arising through the mechanisms of (a) and (b), so the formula holds generally.

We can now construct Ω using the Thurston-Gompf method. Its cohomology class is to be $2\pi^2[\chi]^{(1)} - \pi 1^{(2)}$, where χ is as in the previous paragraph. We shall be brief here, since we have already seen one instance of this. We can find a cover $(U_\alpha)_\alpha$ of $S^1 \times \bar{D}$, and forms $\theta_\alpha \in Z^2(\pi_{S^1}^{-1}(U_\alpha))$, all representing restrictions of this same class, and all tamed by the same complex structure on $\ker(D\pi_{S^1})$. We take one of the members of the open cover to be an open neighbourhood of $S^1 \times \{1\}$; on this we have the form $\chi_{c,\tau}$. The standard patching procedure furnishes us with a 2-form Ω (unique up to deformation) which represents the class $2\pi^2[\chi]^{(1)} - \pi 1^{(2)}$, which restricts to $\omega_{c,\tau}$ over $S^1 \times \{1\}$, and which is tamed by the chosen almost structure on $\ker(D\pi_{S^1})$.

Now, $F'^*(\Omega|_{\mathcal{E}^{\text{crit}}})$ represents the same class as $\bar{\omega}_{\bar{c},\tau}$, and is tamed by the same vertical almost complex structure. Hence the two are isotopic, and therefore isomorphic. \square

Proof of Theorem 4.2.2. We form the 5-dimensional space E , and hence \mathcal{E} . We endow it with a structure of family of simple quadratic fibrations as in the last lemma, with 2-form Ω , and form the fibred vanishing cycle

$$\Lambda_\Gamma \subset \mathcal{E}_{S^1 \times \{1\}} \times_{S^1} \mathcal{E}^{\text{crit}}.$$

But $(\mathcal{E}_{S^1 \times \{1\}}, \Omega) \cong (\mathcal{S}, \sigma_{c,\tau})$, and $(\mathcal{E}^{\text{crit}}, \Omega) \cong (\bar{\mathcal{S}}, \bar{\sigma}_{\bar{c},\tau})$. Via these isomorphisms we obtain

$$\mathcal{Q} \subset \mathcal{S} \times_{S^1} \bar{\mathcal{S}}.$$

We have already seen why this is an isotropic subbundle.

The space E , and its 2-form χ , are not determined uniquely, but they are unique up to deformation. Likewise, the choice of Ω is unique up to deformation. The fibred vanishing-cycle construction varies continuously with parameters, and consequently \mathcal{Q} is unique up to deformations through submanifolds satisfying the conclusions of the theorem.

It is clear that pr_1 embeds \mathcal{Q} as a subbundle of \mathcal{S} , and $\text{pr}_2: \mathcal{Q} \rightarrow \bar{\mathcal{S}}$ an S^1 -bundle. The bundle $\mathcal{Q} \rightarrow S^1$ has fibres \mathcal{Q}_t , $t \in S^1$. We look at one of these, \mathcal{Q}_1 . This is a trivial S^1 -bundle over $\text{Sym}^{r-1}(\bar{\Sigma})$ because the normal bundle to $\text{Hilb}_D^r(X)^{\text{crit}}$ is differentiably trivial.

Now, \mathcal{Q} is the graph of an S^1 -bundle over $\mathcal{E}^{\text{crit}}$ which is the unit circle bundle of a two-plane vector bundle. This vector bundle is a subbundle of the normal bundle to $\mathcal{E}^{\text{crit}}$ (it is the stable bundle for a suitable Morse-Bott function) and it will be globally trivial if and only if \mathcal{Q} is orientable. This accounts for the commutative diagram.

We next show that $\text{pr}_{2*}[\Lambda] = \delta_L$. Indeed, the class $\text{pr}_{2*}[\Lambda] \in H_{2r-1}(\text{Sym}^r(\Sigma); \mathbb{Z})$ is invariant under symmetric products of diffeomorphisms $\psi \in \text{Diff}^+(\Sigma)$ which act as the identity near L . It follows that $\text{pr}_{2*}[\Lambda]$ is a multiple of δ_L . Let L' be a loop with $L \cdot L' = 1$. Let rL' denote the 1-cycle $\{rx : x \in L'\}$ on $\text{Sym}^r(\Sigma)$. Then $\text{pr}_{2*}[\Lambda] \cdot rL' = 1$, by Remark (4.4.3).

Now, $\text{pr}_{2*}[\mathcal{Q}]$ is the class swept out by the invariant class $\text{pr}_{2*}[\Lambda]$. But this is δ_Q . □

Proof of Addendum 4.2.6. The lower part of this diagram was established in Chapter 1; what remains is to set up the map $\pi_0 \mathcal{S}(\mathcal{Q}) \rightarrow H_2^{(d)}$. Decompose Y as $Y_0 \cup_{\partial Y_0} N_Q$, where N_Q is a closed tubular neighbourhood of Q . Thus there is a corresponding decomposition $\bar{Y} = Y_0 \cup_{\partial Y} N'$. By Remark 4.4.3, \mathcal{Q} , or rather a ‘topological vanishing cycle’ isotopic to it, contains an open subbundle $\mathcal{Q}_0 = \{(t, q + \mathbf{y}, \mathbf{y}) : t \in S^1, q \in Q_t, \mathbf{y} \in \text{Sym}^{r-1}(Y_{0,t})\}$. There is a ‘cycle map’ $\mathcal{S}(\mathcal{Q}_0) \rightarrow \pi_0 \mathcal{S}(Q) \times H_1(\bar{Y}; \mathbb{Z})_{r-1}$ which is bijective if $r > 1$. But we already know that $\pi_0 \mathcal{S}(\mathcal{Q}) \cong \pi_0(Q \times_{S^1} \text{Sym}^{r-1}(\bar{Y})) \cong \pi_0(Q) \times H_1(\bar{Y}; \mathbb{Z})_{r-1}$, and from this we see that $\pi_0 \mathcal{S}(\mathcal{Q}) = \pi_0 \mathcal{S}(\mathcal{Q}_0)$. Thus we have a bijective map

$$\pi_0 \mathcal{S}(\mathcal{Q}) \rightarrow \pi_0 \mathcal{S}(Q) \times H_1(\bar{Y}; \mathbb{Z})_{r-1}.$$

But the right-hand side is naturally identified with $H_2^{(d)}$. □

Proof of Addendum 4.2.5. We write $\Sigma_r = \text{Sym}^r(\Sigma)$, $\bar{\Sigma}_{r-1} = \text{Sym}^{r-1} \bar{\Sigma}$. It suffices, by the exact sequence of homotopy groups for the pair $(\Sigma_r \times \bar{\Sigma}_{r-1}, \Lambda)$, to show that the homomorphism $\pi_1(\Lambda) \rightarrow \pi_1(\Sigma_r \times \bar{\Sigma}_{r-1})$ is injective; we prove this in three steps.

1. *On the relative Hilbert scheme $\pi_r: \text{Hilb}_{\bar{D}}^r(E) \rightarrow \bar{D}$ of an elementary Lefschetz fibration $E \rightarrow \bar{D}$, the inclusion of the critical set, $i: \text{Hilb}_{\bar{D}}^r(E)^{\text{crit}} \rightarrow \text{Hilb}_{\bar{D}}^r(X)$, induces an injection on H_1 .*

For if $h \in \ker(i_*)$ then $\mathbf{aj}_*(i_*h) = 0 \in H_1(\mathcal{P}_D^r(X); \mathbb{Z})$. This implies, by our explicit picture of the Picard family, that $\mathbf{aj}_*(h)$ is zero within H_1 of the critical set $\mathcal{P}^r(X_0)^{\text{crit}} \subset \mathcal{P}^r(X_0)$. But $\mathbf{aj}: \text{Hilb}_{\bar{D}}^r(E)^{\text{crit}} \rightarrow \mathcal{P}^r(X_0)^{\text{crit}}$ can be identified with the usual Abel-Jacobi map $\text{Sym}^{r-1}(\tilde{X}_0) \rightarrow \text{Pic}^{r-1}(\tilde{X}_0)$, which is injective on H_1 .

2. *The homomorphism $\text{pr}_{1*}: \pi_1(\Lambda) \rightarrow \pi_1(\Sigma_r)$ is injective.*

Since $\text{pr}_1: \Lambda \rightarrow \Sigma_r$ is a diffeomorphism onto its image Λ' , we may consider the map $\pi_1(\Lambda') \rightarrow \pi_1(\Sigma_r)$. Suppose γ lies in its kernel. Then its homology class $\bar{\gamma} \in H_1(\Lambda; \mathbb{Z})$ becomes trivial in $H_1(\text{Hilb}_{\bar{D}}^r(E); \mathbb{Z})$; that is, $k_*\bar{\gamma} = 0$, where $k: \Lambda' \rightarrow \text{Hilb}_{\bar{D}}^r(X)$ is the inclusion in the fibre over 1. Parallel transport into the critical set gives an S^1 -bundle $l: \Lambda' \rightarrow \bar{\Sigma}_{r-1}$. We have $i_*l_*\bar{\gamma} = k_*\bar{\gamma} = 0$. By (1), this implies that $l_*\bar{\gamma} = 0$, so γ is represented by a loop on one of the S^1 -fibres of Λ . The fibres are homotopic to the circles $\mathbf{x}_0 + L$, where $\mathbf{x}_0 \in \text{Sym}^{r-1}(\Sigma)$ is some fixed divisor, supported away from the circle L , and these circles π_1 -inject. Hence γ is trivial.

3. *The homomorphism $\pi_1(\Lambda) \rightarrow \pi_1(\Sigma_r \times \bar{\Sigma}_{r-1})$ is injective.*

This follows immediately from (2). □

4.5 Floer homology of the Jacobian families

All the invariants set up in this chapter have simpler variants, derived from the compactified Picard families rather than the Hilbert schemes. These constitute a TQFT which is a simpler counterpart of our symmetric product theory. It may well be possible to give an algebro-topological model for this TQFT in terms of Spin^c -structures and homology groups of the 3- and 4-manifolds, and it would be interesting to carry this out. The Floer homology groups themselves are known, as we explain here.

As a matter of convenience of language, we will phrase the following discussion in terms of symplectic automorphisms rather than their mapping tori.

The Floer homology of linear automorphisms of linear symplectic tori was determined by Poźniak [35]³; it is simply the homology of the fixed point set, and as such can be computed explicitly.

Let $T = V/\Lambda$, where V is a real vector space of dimension $2k$ and $\Lambda \subset V$ a lattice. Let ω_0 be a linear symplectic form on the manifold V ; ω_0 also denotes the induced form on T . A

³Poźniak works in the setting of Floer homology of Lagrangian intersections. The results that we state are deduced from his using the standard isomorphism $HF_*^M(\phi) \cong HF_*^{M \times M}(\Gamma_{\text{id}}, \Gamma_\phi)$; alternatively, his proof could easily be adapted.

linear symplectomorphism $\tilde{\phi} \in \text{Aut}(V, \omega_0)$ which preserves Λ induces a symplectomorphism $\phi \in \text{Aut}(T, \omega_0)$. Such linear symplectomorphisms form a group $\text{Aut}^l(T, \omega_0)$. Write

$$V^{\tilde{\phi}} := \ker(\text{id}_V - \tilde{\phi}), \quad V_{\tilde{\phi}} := \text{coker}(\text{id}_V - \tilde{\phi})$$

for the invariants and coinvariants of $\tilde{\phi}$. Write $\Lambda^{\tilde{\phi}}, \Lambda_{\tilde{\phi}}$ for the kernel and cokernel of $\text{id} - \tilde{\phi}$ operating on Λ .

(a) The fixed-point set of $\phi \in \text{Aut}^l(T, \omega_0)$ is the closed subgroup

$$\{x \in V : (\text{id}_V - \tilde{\phi})x \in \Lambda\} / \Lambda.$$

The exact sequence of topological groups $0 \rightarrow \text{Fix}(\phi)_0 \rightarrow \text{Fix}(\phi) \rightarrow \pi_0(\text{Fix}(\phi)) \rightarrow 0$ takes the concrete form

$$0 \longrightarrow \frac{\ker(\text{id}_V - \tilde{\phi})}{\Lambda \cap \ker(\text{id}_V - \tilde{\phi})} \longrightarrow \text{Fix}(\phi) \xrightarrow{\text{id} - \phi} (\Lambda_{\tilde{\phi}})^{\text{tors}} \longrightarrow 0.$$

Hence $\text{Fix}(\phi)$ is a disjoint union of affine subtori, each of dimension $\dim(V^{\tilde{\phi}})$, indexed by the torsion subgroup of the coinvariants, $(\Lambda_{\tilde{\phi}})^{\text{tors}}$, and

$$H_*(\text{Fix}(\phi); \mathbb{Z}) \cong \bigoplus_{(\Lambda_{\tilde{\phi}})^{\text{tors}}} \bigwedge^* \Lambda^{\tilde{\phi}}.$$

(b) By Appendix B, Lemma B.0.5, there is a natural exact sequence

$$0 \longrightarrow \pi_1(L_\phi T, l) \longrightarrow \Lambda \xrightarrow{\text{id} - \tilde{\phi}} \Lambda \longrightarrow \pi_0(L_\phi T) \longrightarrow 0$$

where $L_\phi T$ is the twisted loop space (i.e. the space of sections of the mapping torus of ϕ). Comparing this sequence with the previous one, one deduces that there is at most one component of $\text{Fix}(\phi)$ in any given component of $L_\phi(T)$, and that the natural map $\pi_1(\text{Fix}(\phi)_0, [\Lambda]) \rightarrow \pi_1(L_\phi T, \text{const.})$ is an isomorphism.

Theorem 4.5.1 (Poźniak). *Let $\phi \in \text{Aut}(M, \omega)$ be an automorphism of a symplectic manifold, $\gamma \in \pi_0(L_\phi M)$, and suppose that the set $\text{Fix}(\phi)_\gamma \subset \text{Fix}(\phi)$ of fixed points representing γ is a compact smooth submanifold of M .*

(a) *The local Floer homology $HF_*(\phi)_{\gamma, \text{loc}}$ is defined over $\mathbb{Z}/2$, and is canonically isomorphic to $H_*(\text{Fix}(\phi)_\gamma; \mathbb{Z}/2)$.*

(b) *If in addition, the action functional is a single-valued function on γ and constant on $\text{Fix}(\phi)_\gamma$, then $HF_*(\phi)_{\gamma, \text{loc}} = HF_*(\phi)_\gamma$.*

Corollary 4.5.2. *If $\phi \in \text{Aut}^l(T, \omega_0)$ then (using $\mathbb{Z}/2$ coefficients) $HF_*(\phi) \cong H_*(\text{Fix}(\phi); \mathbb{Z}/2)$.*

We apply this to a surface-diffeomorphism $\psi \in \text{Diff}^+(\Sigma)$. Let $V = H^1(\Sigma; \mathbb{R})$, $\Lambda = H^1(\Sigma; \mathbb{Z})$, and $\tilde{\phi} = \psi^*$. Let $Y = T(\psi)$ be the mapping torus. Then the Floer homology of the Picard family of Y is just $HF_*(V/\Lambda, \phi)$. The set of topological sectors $(H^1(\Sigma; \mathbb{Z})_\psi)^{\text{tors}}$ may be identified with $H_1(Y; \mathbb{Z})_0^{\text{tors}}$, where $H_1(Y; \mathbb{Z})_0 \subset H_1(Y; \mathbb{Z})$ is the kernel of the map $H_1(Y; \mathbb{Z}) \rightarrow H_0(Y; \mathbb{Z})$ given by intersection with the fibre. The invariants $H^1(\Sigma; \mathbb{Z}/2)^\phi$ may be identified with

$$H^1(Y; \mathbb{Z}/2)_0 := \text{coker}(H^0(\Sigma; \mathbb{Z}/2) \rightarrow H^1(Y; \mathbb{Z}/2)).$$

Corollary 4.5.3. *We have*

$$HF_*(T, \phi) = \bigoplus_{h \in H_1(Y; \mathbb{Z})_0^{\text{tors}}} HF_*(T, \phi)_h = \bigoplus_{h \in H_1(Y; \mathbb{Z})_0^{\text{tors}}} \Lambda^* H^1(Y; \mathbb{Z}/2)_0.$$

Remark 4.5.4. *There are two refinements of this corollary which I expect to be true but have not yet verified. The first is that it lifts to \mathbb{Z} -coefficients in the obvious way. The second is that the absolute grading on the Floer homology implied by the isomorphism corresponds to that which one uses the Lie group structure of T to consider it as a ‘graded symplectic manifold’, and ϕ as a ‘graded symplectic automorphism’ of it (see Seidel [43]).*

4.6 An exact triangle?

The purpose of this section is to state a conjecture relating the Floer homology maps associated with near-symplectic broken fibrations to those associated with Lefschetz fibrations.

Suppose that $\phi \in \text{Aut}(\Sigma, \omega)$ preserves a circle $L \subset \Sigma$. Then one can consider $\tau_L \circ \phi$, that is, ϕ composed with a Dehn twist along L , but also $\sigma_L \phi$, the result of surgery along L .

There is a Lefschetz fibration $X_{\text{symp}}^L \rightarrow A[r_2, r_3]$ over an annulus, with just one critical point, which is isomorphic to $T(\phi)$ over its inner boundary $C(r_2)$ and to $T(\tau_L \circ \phi)$ over its outer boundary $C(r_3)$. There is an elementary broken fibration $X_{\text{br}}^L \rightarrow A[r_1, r_2]$ which restricts to $T(\sigma_L \phi)$ over $C(r_1)$ and to $T(\phi)$ over $C(r_2)$. There is also an elementary broken fibration $W \rightarrow A[r_1, r_3]$ (no isolated critical points!) which restricts to $T(\sigma_L \phi)$ over $C(r_1)$ and to $T(\tau_L \circ \phi)$ over $C(r_3)$. Let $\bar{W} \rightarrow A[r_1, r_3]$, the pullback of W by the orientation-reversing self-diffeomorphism $re^{i\theta} \mapsto (-r + r_1 + r_3)e^{i\theta}$ of $A[r_1, r_3]$. Note that \bar{W} restricts to $T(\phi)$ over $C(r_1)$ and to $T(\sigma_L \phi)$ over $C(r_3)$.

Conjecture 4.6.1. *The sequence*

$$\begin{array}{ccc} H_*(L; \Lambda_{\mathbb{F}}) \otimes HF_{\bullet}(T(\sigma_L \phi), \gamma_1) & \xrightarrow{\lambda \otimes x \mapsto \delta_{\lambda} \circ HF_{\bullet}(X_{\text{br}}^L)} & HF_{\bullet}(T(\phi), \gamma_2) \\ & \nwarrow HF_{\bullet}(\bar{W}) \quad \nearrow HF_{\bullet}(X_{\text{symp}}^L) & \\ & HF_{\bullet}(T(\tau_L \circ \phi), \gamma_3) & \end{array}$$

is exact, for homology classes γ_i which arise from a class $\beta \in H_2^{(d)}(X_{\text{br}}^L \circ X_{\text{symp}}^L)$.

Here δ_λ denotes quantum cap product with the codimension-one cycle δ_λ on $\text{Sym}^r(\Sigma)$. We use Seidel's homomorphism HF_\bullet^{symp} , induced by the Lefschetz fibration, and two of our homomorphisms induced by broken fibrations, $HF_\bullet(X_{\text{br}}^L)$ and $HF_\bullet(\bar{W})$. Note that we have used the special feature of elementary broken fibrations that one construct maps running in either direction (we expressed this by reversing the fibration, replacing W by \bar{W} ; alternatively one can construct the reversed homomorphism directly, but the result will be the same).

We conjecture that the Floer complexes form a short exact sequence, inducing the triangle above, so that $HF_\bullet(\bar{W})$ coincides with the algebraic connecting map. It is important that we sum over the topological sectors of the map $HF_\bullet(X_{\text{br}}^L)$ (these are parametrised by \mathbb{Z} once one fixes γ_2 ; part of the conjecture is that the sum is finite).

The conjecture is motivated by the exact triangles of Seidel [42, 46]. When $r = 1$, we recover one of these,

$$H_*(S^1; \Lambda_{\mathbb{F}}) \rightarrow HF_*(\phi) \rightarrow HF_*(\tau_L \circ \phi).$$

For higher r , taking $\phi = \text{id}$ and L non-separating, the sequence permits one to compute $HF_\bullet(\tau_L)$; the result is isomorphic to the periodic Floer homology $PFH_*(\tau_L)$ as computed by Hutchings and Sullivan [19]. There should be a similar sequence in the Floer homology for the Picard families arising from the same cobordisms. This may be useful as a simpler model, particularly since one knows the Floer modules themselves.

The composite cobordism $X_{\text{br}}^L \circ X_{\text{symp}}^L$ is known to be diffeomorphic to the blow-up $W \# \overline{\mathbb{C}P^2}$ of the elementary broken fibration W (see [3]). The exceptional sphere S_{exc} appears as a 'matching cycle', lying over a ray between the isolated critical value and the circle of critical values. The conjecture relating HF_\bullet to monopole Floer homology predicts that $HF_\bullet(X_{\text{br}}^L \circ X_{\text{symp}}^L, \hat{\beta})$ should be equal, up to sign, to the composite of $HF_\bullet(X'', \beta)$ with the quantum product with U^k , where k is determined by the intersection number of $\hat{\beta}$ with S_{exc} . The contributions of the sectors $\hat{\beta}$ to $HF_\bullet(X_{\text{br}}^L \circ X_{\text{symp}}^L)$ should cancel in pairs.

(-1) -spheres also occur in the composite cobordisms involved in the surgery exact sequence for HM_\bullet [22], and it would be interesting to work out how the situations are related.

In seeking to prove the conjecture, we should look for a null-homotopy of the map $CF_\bullet(X_{\text{br}}^L \circ X_{\text{symp}}^L)$. It may be possible, as with Seidel's exact triangle [46] (and in analogy to [22]), to do this by breaking up $X_{\text{br}}^L \circ X_{\text{symp}}^L$ into two pieces in an alternative way. One considers the part lying over a disk containing the unique critical value. This is an elementary Lefschetz fibration $E \rightarrow \bar{D}$. It has a Lagrangian boundary condition $\mathcal{Q} \subset \text{Sym}_{S^1}^r(\partial E) \times \text{Sym}^{r-1}(\sigma_L \Sigma)$, with a marked fibre Λ , and thus an open Gromov invariant with values in the bordism group $MO_*(\Lambda; \mathbb{Z})$. I expect this invariant to be zero. The moduli space $\mathcal{M}(\mathcal{Q})$ has expected dimension $2r - 1$, and one can arrange that $\text{ev}: \mathcal{M}(\mathcal{Q}) \rightarrow \Lambda$ is

a submersion, hence a covering map. It should be possible to construct a fixed-point-free involution on $\mathcal{M}(Q)$ preserving the fibres of ev , by using a symmetry of E . Then ev will be bordant to zero, and (as in [46]) one can try to use a null-bordism to construct a null-homotopy of $CF_{\bullet}(X_{\text{br}}^L \circ X_{\text{symp}}^L)$.

Appendix A

Symmetric products of non-singular complex curves

There are many accounts of the algebraic topology and algebraic geometry of the symmetric products $\text{Sym}^r(\Sigma) = (\Sigma^{\times r})/S_r$ of a compact, connected, non-singular complex curve Σ . The original reference for the topology is Macdonald [29], but other useful sources are Arbarello-Cornalba-Griffiths-Harris [2], Ozsváth-Szabó [34], Bertram-Thaddeus [4]. We give a summary, with brief indications of proofs.

The space $\text{Sym}^r(\Sigma)$ is a complex manifold, and often the most efficient way to study its topology is via complex geometry.

For a basepoint $x \in \Sigma$, we let $\delta_x: \text{Sym}^{r-1}(\Sigma) \rightarrow \text{Sym}^r(\Sigma)$ denote the map $D \mapsto x + D$.

(a) The Abel-Jacobi map

$$\mathbf{aj}: \text{Sym}^r(\Sigma) \rightarrow \text{Pic}^r(\Sigma),$$

which assigns to the divisor D the holomorphic line bundle $\mathcal{O}(D)$, is crucial to understanding the geometry and topology of $\text{Sym}^r(\Sigma)$. The fibre $\mathbf{aj}^{-1}(\mathcal{L})$ is $\mathbb{P}H^0(\Sigma, \mathcal{L})$, so its dimension can be estimated using the Riemann-Roch formula. The map \mathbf{aj} is generically 1-1 when $r \leq g$, and surjective when $r \geq g$. When $r > 2g - 2$ the fibres are all projective spaces of dimension $(r - g)$, and in fact there is a holomorphic vector bundle $\mathcal{V}_r \rightarrow \text{Pic}^r(\Sigma)$ such that $\text{Sym}^r(\Sigma) = \mathbb{P}(\mathcal{V}_r)$: start with a Poincaré line bundle $\mathcal{L} \rightarrow \Sigma \times \text{Pic}^r(\Sigma)$ normalised at x (so $\mathcal{L}|_{\Sigma \times \{L\}} = L$ for any L , and $\mathcal{L}|_{\{x\} \times \text{Pic}^r(\Sigma)} = \mathcal{O}_{\text{Pic}^r(\Sigma)}$), and put $\mathcal{V}_r = (\text{pr}_2)_* \mathcal{L}$.

(b) The line bundle $\mathcal{O}(1) = \mathcal{O}_{\text{Sym}^r(\Sigma)}(\text{Sym}^{r-1}(\Sigma))$ defined by the divisor $\text{im}(\delta_x)$ is ample. Indeed, when $r > 2g - 2$ it is isomorphic to the natural ample line bundle on $\mathbb{P}(\mathcal{V}_r)$, and since $\delta_x^* \mathcal{O}(1) = \mathcal{O}(1)$, ampleness follows by descending induction.

(c) One can use the fact that it is a projective vector bundle over a g -dimensional complex torus to establish a number of topological properties of $\text{Sym}^r(\Sigma)$ when $r > 2g - 2$.

- (1) $\pi_1(\text{Sym}^r(\Sigma))$ is abelian, isomorphic to $H_1(\Sigma; \mathbb{Z})$. For the homotopy exact sequence shows that \mathbf{aj} induces an isomorphism on π_1 .
- (2) $\pi_2(\text{Sym}^r(\Sigma)) \cong \mathbb{Z}$, generated by a pencil $\mathbb{CP}^1 \subset |D|$. This sphere represents a non-torsion class in $H_2(\text{Sym}^r(\Sigma); \mathbb{Z})$. The homotopy sequence shows that $\pi_i(\text{Sym}^r(\Sigma))$ is carried by the fibre \mathbb{CP}^{r-g} for $i > 1$. If S is the spherical homology class corresponding to a pencil then $S \cap \delta_x(\text{Sym}^{r-1}(\Sigma)) = 1$, so S is not torsion.¹
- (3) The cohomology of $\text{Sym}^r(\Sigma)$ is torsion-free. For the ‘projective bundle formula’ gives an isomorphism of graded abelian groups $H^*(\text{Sym}^r(\Sigma); \mathbb{Z}) \cong \Lambda^* H^1(\Sigma; \mathbb{Z}) \otimes_{\mathbb{Z}} H^*(\mathbb{CP}^{r-g}; \mathbb{Z})$.

(d) The italicised results from (c) are valid over the range $r \geq 2$, as one can prove by descending induction.² The Lefschetz hyperplane theorem, applied to δ_x , tells us that the pair $(\text{Sym}^r(\Sigma), \text{Sym}^{r-1}(\Sigma))$ is $(r-1)$ -connected. This gives (1) and (with Poincaré duality) (3) for $r \geq 2$, and (2) for $r \geq 3$.

(e) For any Hausdorff space X , the quotient map $p: X^{\times r} \rightarrow \text{Sym}^r(X)$ induces an isomorphism

$$H^*(\text{Sym}^r(X); \mathbb{Q}) \xrightarrow{p^*} H^*(X^{\times r}; \mathbb{Q})^{S_r} \subset H^*(X^{\times r}; \mathbb{Q});$$

its inverse is $p_*/r!$. By Künneth, $H^*(X^{\times r}; \mathbb{Q})^{S_r}$ is the algebra of graded-symmetric tensors $(H^*(X; \mathbb{Q})^{\otimes r})^{S_r}$. So, as graded \mathbb{Q} -algebras,

$$H^*(\text{Sym}^r(\Sigma); \mathbb{Q}) \cong \bigoplus_{i=0}^r \Lambda^{r-i} H^1(\Sigma; \mathbb{Q}) \otimes_{\mathbb{Q}} S^i \langle 1, U \rangle$$

where $U \in H^2(\Sigma; \mathbb{Q})$ is the oriented, integral generator. (One can match up this isomorphism with the one that arose from the projective bundle formula; to do so, it is enough to consider the generating classes in degree ≤ 2 .) Note:

- (1) The description is completed by specifying a linear map $\sigma: H^*(\Sigma; \mathbb{Q}) \rightarrow (H^*(\Sigma; \mathbb{Q})^{\otimes r})^{S_r}$, indicating which symmetric tensors correspond to U and to elements of $H^1(\Sigma; \mathbb{Q})$. The map is

$$\sigma(c) = \sum_{i=1}^r 1 \otimes \cdots \otimes 1 \otimes \overset{i}{c} \otimes 1 \otimes \cdots \otimes 1.$$

- (2) $c_1(\mathcal{O}(1)) = \sigma(U)$.

¹One can also calculate homotopy groups of symmetric products through the Dold-Thom theorem, which gives a functorial isomorphism $\pi_n(\text{Sym}^r(X)) \cong H_n(X; \mathbb{Z})$ for $r \gg 0$.

²There is one exception, concerning $\text{Sym}^2(\Sigma)$: the pencil generates π_2 over $\mathbb{Z}[\pi_1]$, but not always over \mathbb{Z} . In genus 2, $\text{Sym}^2(\Sigma)$ is the blow-up at a point of T^4 , and its universal cover is the blow-up along \mathbb{Z}^4 of \mathbb{C}^2 . This has $\pi_2 = H_2 \cong \bigoplus_{\mathbb{Z}^4} \mathbb{Z}$.

- (3) The isomorphism is compatible, in obvious senses, with several other structures on the cohomology: the Hodge decomposition; the module structures over $H^*(\text{Pic}^r(\Sigma); \mathbb{Q}) = \Lambda^* H^1(\Sigma; \mathbb{Q})$, induced by \mathbf{aj} ; and the action of the mapping class group.

(f) Since $H^*(\text{Sym}^r(\Sigma); \mathbb{Z})$ is torsion-free, $H^*(\text{Sym}^r(\Sigma); \mathbb{Z}) \rightarrow H^*(\text{Sym}^r(\Sigma); \mathbb{Q})$ is injective. There is a ring isomorphism $H^*(\text{Sym}^r(\Sigma); \mathbb{Z}) \cong X(g, r)$, where

$$X(g, r) = \bigoplus_{i=0}^r \Lambda^{r-i} H^1(\Sigma; \mathbb{Z}) \otimes_{\mathbb{Z}} \frac{\mathbb{Z}[U]}{(U^{i+1})}, \quad \deg(U) = 2.$$

To establish this, one needs to see that $\sigma(U)^i$ is a non-multiple class over \mathbb{Z} (hence also over $\Lambda^* H^1(\Sigma; \mathbb{Z})$). But when $r > 2g - 2$, the powers $c_1(\mathcal{O}(1))^i = \sigma(U)^i$ generate $H^*(\text{Sym}^r(\Sigma); \mathbb{Z})$ as a $H^*(\text{Pic}^*(\Sigma); \mathbb{Z})$ -module, by the projective bundle formula; then descending induction applies as usual.

(g) The mapping class group $\Gamma = \pi_0(\text{Diff}^+(\Sigma))$ acts on $H^*(\text{Sym}^r(\Sigma))$. The invariant subalgebra $H^*(\text{Sym}^r(\Sigma); \mathbb{Q})^\Gamma$ is generated by two distinguished classes η, θ , both integral of type $(1, 1)$: $\eta = \sigma(U) = c_1(\mathcal{O}(1))$; and $\theta \in \Lambda^2 H^1(\Sigma; \mathbb{Q}) = \Lambda^2 \text{Hom}(H_1(\Sigma; \mathbb{Z}), \mathbb{Q})$, characterised by $\langle \theta, \alpha \wedge \beta \rangle = \alpha \cap \beta$ for $\alpha, \beta \in H_1(\Sigma; \mathbb{Z})$. That these classes generate the invariant part amounts to the statement that θ generates the $\text{Sp}(2g, \mathbb{Q})$ -invariant part of $\Lambda^* H^1(\Sigma; \mathbb{Q})$. But that can be seen from the irreducible decomposition of the exterior algebra.

Note that $\theta^i/i!$ is an integral class (in fact, $\mathbf{aj}^* c(\mathcal{V}) = e^{-\theta}$; cf. also Poincaré's formula).

(h) The first Chern class of the tangent bundle is

$$c_1(\text{Sym}^r(\Sigma)) = (r - g + 1)\eta - \theta.$$

A more general formula is proved in Chapter 2. The formula specialises to two familiar ones: taking $g = 0$, we find $c_1(\mathbb{CP}^r) = (r + 1)\eta$; with $r = 1$ we get $c_1(\Sigma) = 2(1 - g)\eta$, since $\theta = g\eta$.

Appendix B

Topology of mapping tori

Exact triangles

The **mapping torus** $T(\phi)$ of a diffeomorphism $\phi \in \text{Diff}(M)$ is the quotient of the manifold $\mathbb{R} \times M$ by the free \mathbb{Z} -action $n \cdot (t, x) = (t - n, \phi^n(x))$. There is a natural fibre bundle $M \xrightarrow{i} T(\phi) \rightarrow S^1$.

The (co)homology of $T(\phi)$ fits into exact triangles

$$\begin{array}{ccc} H_*(M) & \xrightarrow{\text{id}-\phi_*} & H_*(M) \\ & \searrow [-1] & \downarrow i_* \\ & & H_*(T(\phi)) \end{array} \quad \begin{array}{ccc} H^*(M) & \xrightarrow{\text{id}-\phi^*} & H^*(M) \\ & \nwarrow [+1] & \uparrow i^* \\ & & H^*(T(\phi)). \end{array}$$

One way to set up these triangles is to use the mechanism of the Wang sequence. To derive the homology sequence, consider the Leray-Serre spectral sequence of the fibration $T(\phi) \rightarrow S^1$. It is convenient to use the cellular version, choosing the cell decomposition of S^1 with just two cells, so that the $E_{p,q}^1$ -term is zero unless p is 0 or 1, when it is $H_q(M)$, and $d_{1,q}^1$ is the map $(\text{id} - \phi_*) : H_q(M) \rightarrow H_q(M)$. The spectral sequence degenerates at E^2 , so there are exact sequences

$$0 \rightarrow E_{1,q}^\infty \rightarrow H_q(M) \xrightarrow{\text{id}-\phi_*} H_q(M) \rightarrow E_{0,q}^\infty \rightarrow 0.$$

The exact triangle results from splicing these with the short exact sequences $0 \rightarrow H_{p+q}(M) \rightarrow H_{p+q}(T(\phi)) \rightarrow E_{pq}^\infty \rightarrow 0$.

The connecting maps are the transfer operations $i^! : H_q(T(\phi)) \rightarrow H_{q-1}(M)$ (intersection with the fibre) and $i_! : H^q(M) \rightarrow H^{q+1}(T(\phi))$.

Algebraic model. The homology of $\text{cone}(\text{id} - \phi_*)$, the algebraic mapping cone of

$$\text{id} - \phi_* : C_*(M) \rightarrow C_*(M),$$

where $C_*(M)$ is the singular chain complex, fits into an exact triangle

$$\begin{array}{ccc} H_*(M) & \xrightarrow{\text{id}-\phi_*} & H_*(M) \\ & \nwarrow [-1] & \downarrow \\ & & H_*(\text{cone}(\text{id}-\phi_*)) \end{array}$$

just like that of $H_*(T(\phi))$. In fact, there is a quasi-isomorphism

$$Q : \text{cone}(\text{id}-\phi_*) \rightarrow C_*(M_\phi),$$

and the isomorphism Q_* identifies the two exact triangles.

Recall that $\text{cone}(\text{id}-\phi_*)$ is the complex $C_*(M)[-1] \oplus C_*(M)$ with differential $(a, b) \mapsto (-\partial a, \partial b + (\text{id}-\phi_*)a)$. The exact triangle arises from the obvious short exact sequence $0 \rightarrow C_*(M) \rightarrow \text{cone}(\text{id}-\phi_*) \rightarrow C_*(M)[-1] \rightarrow 0$, and the maps $(\text{id}-\phi_*)$ in the triangle are the connecting maps.

Using *cubical* singular chains, define

$$Q : \text{cone}(\text{id}-\phi_*)_q \rightarrow C_q(T(\phi)), \quad (a, b) \mapsto \tilde{a} + i^*b.$$

Here $a \mapsto \tilde{a}$ is the linear extension of the map which sends a singular $(q-1)$ -cube σ to the q -cube $\sigma \times \text{id} : [0, 1]^{q-1} \times [0, 1] \rightarrow M \times [0, 1]$, projected to $T(\phi)$. The reason for choosing cubical chains is that one can see that

$$\partial \tilde{a} = -\widetilde{\partial a} + i_*(a - \phi_*a),$$

which shows that Q is a chain map. It is then easy to check that Q_* , in conjunction with identity maps on $H_*(M)$, gives a morphism between the exact triangles. The five-lemma shows that the induced map $Q_* : H_*(\text{cone}(\text{id}-\phi_*)) \rightarrow H_*(T(\phi))$ is an isomorphism.

It is straightforward to dualise this procedure so as to set up a quasi-isomorphism $C^*(M_\phi) \rightarrow \text{cone}(\text{id}-\phi^*)$.

Example B.0.2. One can compute the homology of $\text{cone}(\text{id}-\phi_*)$ using a cellular chain complex for M and a cellular approximation for ϕ , because of the naturality properties of mapping cones (compatibility with quasi-isomorphisms, and the fact that chain homotopic maps have chain homotopic cones.) Suppose, in particular, that M admits a cell decomposition with zero differentials. Then we have an isomorphism

$$H_*(\text{cone}(\text{id}-\phi_*)) \cong H_*(M)^\phi[1] \oplus H_*(M)_\phi \tag{B.1}$$

splitting the exact sequence. (Here $H_*(M)^\phi = \ker(\text{id}-\phi_*)$, $H_*(M)_\phi = \text{coker}(\text{id}-\phi_*)$.) The isomorphism may depend, however, on the choice of cell complex and of cellular approximation to ϕ .

This may be applied to closed orientable surfaces and to their Jacobian tori. Less obviously, it applies to their symmetric products: a cell decomposition of a surface induces symmetric cell decompositions of its cartesian products, hence cell decompositions of its symmetric products; these inherit the property of having zero differentials.

De Rham model. The constructions of the previous paragraphs have rather explicit representations using differential forms.

Fix a Thom form τ on S^1 —a 1-form, supported in a small interval, with integral 1. For a form $\alpha \in \Omega_M^q$, define $i_!(\alpha) \in \Omega_{T(\phi)}^{q+1}$ to be the form whose pullback to $\mathbb{R} \times M$ is $\tau \wedge \alpha$. The induced map in cohomology, $i_! : H^q(M; \mathbb{R}) \rightarrow H^{q+1}(M; \mathbb{R})$, is the transfer operation—the connecting map in the exact triangle.

A typical element of $\Omega_{T(\phi)}^q$ has shape $\omega_t + dt \wedge \eta_t$. Here, for any $t \in \mathbb{R}$, $\omega_t \in \Omega_M^q$, $\eta_t \in \Omega_M^{q-1}$, and $\phi^* \omega_t = \omega_{t+1}$, $\phi^* \eta_t = \eta_{t+1}$.

The mapping cone $\text{cone}(\text{id} - \phi^*)$ of $\text{id} - \phi^* : \Omega_M^* \rightarrow \Omega_M^*$ is the complex

$$\Omega_M^*[-1] \oplus \Omega_M^*, \quad d(\alpha, \beta) = (-d\alpha + \beta - \phi^* \beta, d\beta),$$

and one checks that the map

$$\Omega_{T(\phi)}^* \rightarrow \text{cone}(\text{id} - \phi^*), \quad \omega_t + dt \wedge \eta_t \mapsto \left(\int_0^1 \eta_t dt, \omega_0 \right)$$

is a chain map. By the five-lemma it is a quasi-isomorphism.

Sections

The space of continuous sections of $T(\phi) \rightarrow S^1$ is the twisted loop space

$$L_\phi M = \{\ell : \mathbb{R} \rightarrow M : \ell(t+1) = \phi \circ \ell(t)\}.$$

Fix a basepoint $x \in M$; it is convenient to assume (isotoping ϕ if necessary) that $\phi(x) = x$. We then use the constant path \mathbf{x} as a basepoint for $L_\phi M$.

Lemma B.0.3. *There are canonical bijections of pointed sets*

$$\pi_0(L_\phi M, \mathbf{x}) \cong \{\phi\text{-twisted conjugacy classes in } \pi_1(M, x)\}.$$

The ϕ -twisted conjugates of $\alpha \in \pi_1(M, x)$ are the elements $\beta \cdot \alpha \cdot (\phi \circ \beta)^{-1}$, where $\beta \in \pi_1(M, x)$ and concatenation of paths is written from left to right.

Proof. To see how $\ell \in L_\phi M$ determines a twisted conjugacy class, choose a path γ from x to $\ell(0)$. Then the concatenation $\gamma \cdot \ell|_{[0,1]} \cdot (\phi \circ \gamma)^{-1}$ is a loop at x , and since any other choice γ' is homotopic to $\delta \cdot \gamma$, for some loop δ based at x , its twisted conjugacy class is

determined by ℓ . It is straightforward to check that this map descends to π_0 and then to construct an inverse. □

Example B.0.4. Take a surface-homeomorphism $\phi \in \text{Aut}(\Sigma)$. We consider its symmetric products $\phi_r \in \text{Aut}(\text{Sym}^r(\Sigma))$. Take as reference a fixed point x of ϕ , isotoping ϕ if none exists. A section $\ell \in L_{\phi_r}(\text{Sym}^r(\Sigma))$ represents a 1-cycle in $T(\phi)$, and passing to homology one gets a map

$$\pi_0(L_{\phi_r}(\text{Sym}^r(\Sigma))) \rightarrow H_1(\Sigma_\phi; \mathbb{Z}), \quad \ell \mapsto [\ell].$$

There is, moreover, a commutative diagram

$$\begin{array}{ccc} \pi_0(L_{\phi_r}(\text{Sym}^r(\Sigma)), \mathbf{x}) & \xrightarrow{\cong} & \pi_1(\text{Sym}^r(\Sigma), rx) / \sim \\ \downarrow \ell \mapsto [\ell] - [\mathbf{x}] & & \downarrow h \\ \ker[H_1(T(\phi); \mathbb{Z}) \rightarrow H_0(\Sigma; \mathbb{Z})] & \xrightarrow{\cong} & H_1(\Sigma; \mathbb{Z})_\phi. \end{array}$$

The top row is the bijection of the last lemma (\sim stands for ϕ_r -twisted conjugacy), while the bottom row comes from the homology exact sequence. The surjective map h is induced by the sequence

$$\pi_1(\text{Sym}^r(\Sigma)) \xrightarrow{\text{Hurewicz}} H_1(\text{Sym}^r(\Sigma); \mathbb{Z}) \xrightarrow{\cong} H_1(\Sigma; \mathbb{Z}).$$

When $r = 1$, h will typically not be injective: the set of components $\pi_1(\Sigma) / \sim$ is larger than the set of homology classes $H_1(\Sigma)_\phi$. But for $r > 1$, $\pi_1(\text{Sym}^r(\Sigma))$ is abelian, and consequently h is an isomorphism. Thus

$$\pi_0(L_{\phi_r}(\text{Sym}^r(\Sigma)), \mathbf{x}) \cong H_1(\Sigma; \mathbb{Z})_\phi, \quad r > 1.$$

(We stress that this bijection involves the basepoint x .)

Fundamental group. We now look at the group $\pi_1(L_\phi(M), \ell)$.

Lemma B.0.5. *There is a natural exact sequence*

$$1 \rightarrow \pi_2(M)_\phi \rightarrow \pi_1(L_\phi(M), \ell) \rightarrow \pi_1(M)^\phi \rightarrow 1.$$

Proof. There is no loss of generality if we assume that $x = \ell(0)$ is a fixed point of ϕ , because everything is isotopy-invariant. The Serre fibration

$$\Omega_x(M) \rightarrow L_\phi(M) \xrightarrow{\text{ev}_0} M,$$

gives an exact sequence

$$\pi_2(M) \rightarrow \pi_1(\Omega_x M) \rightarrow \pi_1(L_\phi(M), \ell) \rightarrow \pi_1(M) \rightarrow \pi_0(\Omega_x M).$$

Identifying $\pi_i(\Omega_x M)$ with $\pi_{i+1}(M, x)$, the sequence becomes

$$\pi_2(M) \xrightarrow{1-\phi_*} \pi_2(M) \rightarrow \pi_1(L_\phi(M), \ell) \rightarrow \pi_1(M) \xrightarrow{\alpha \mapsto \alpha(\phi\alpha)^{-1}} \pi_1(M).$$

The last map isn't a homomorphism, but its kernel is the group of invariants $\pi_1(M)^\phi$. This gives the result. □

Example B.0.6. For a surface this gives $\pi_1(L_\phi(\Sigma), \ell) \cong \pi_1(\Sigma)^\phi$.

Example B.0.7. For the higher symmetric products $\text{Sym}^r(\Sigma)$, $r \geq 3$, it gives

$$0 \rightarrow \mathbb{Z} \rightarrow \pi_1(L_{\phi_r}\Sigma, \ell) \rightarrow H_1(\Sigma)^\phi \rightarrow 0.$$

But this short exact sequence maps naturally to the sequence

$$0 \rightarrow \mathbb{Z} \rightarrow H_2(T(\phi); \mathbb{Z}) \rightarrow H_1(\Sigma; \mathbb{Z})^\phi \rightarrow 0,$$

which we know to be split, hence

$$\pi_1(L_{\phi_r}\Sigma, \ell) \cong H_1(T(\phi); \mathbb{Z}) \cong H_1(\Sigma; \mathbb{Z})^\phi \oplus \mathbb{Z}.$$

Bibliography

- [1] Atiyah, M., and Bott. R., *The Yang-Mills equations over Riemann surfaces*, Philos. Trans. Roy. Soc. London Ser. A 308 (1983), no. 1505, 523–615.
- [2] Arbarello, E., Cornalba, M., Griffiths, P., and Harris, J., *Geometry of algebraic curves*, Vol. 1, Springer-Verlag, 1985.
- [3] Auroux, D., Donaldson, S., and Katzarkov, L., *Singular Lefschetz pencils*, preprint, [math.DG/0410332](#), 2004.
- [4] Bertram, A., and Thaddeus, M., *On the quantum cohomology of a symmetric product of an algebraic curve*, Duke Math. J., 108 (2001), no. 2, 329–362.
- [5] Caporaso, L., *A compactification of the universal Picard variety over the moduli space of stable curves*, J. Amer. Math. Soc. 7 (1994), no. 3, 589–660.
- [6] Donaldson, S., *Topological field theories and formulae of Casson and Meng-Taubes*, Geom. Topol. Monographs (2) 1999, 87–102.
- [7] Donaldson, S., *Floer homology groups in Yang-Mills theory*, Cambridge Tracts in Mathematics 167, Cambridge University Press, 2002.
- [8] Donaldson, S., and Kronheimer, P., *The geometry of four-manifolds*, Oxford Mathematical Monographs, Oxford University Press, 1990.
- [9] Donaldson, S., and Smith, I., *Lefschetz pencils and the canonical class for symplectic four-manifolds*, Topology 42 (2003), no. 4, 743–785.
- [10] Duistermaat, J., and Heckman, G., *On the variation in the cohomology of the symplectic form of the reduced phase space*, Invent. Math. 69 (1982), no. 2, 259–268.
- [11] Eliashberg, Y., Givental, A., and Hofer, H., *Introduction to symplectic field theory*, Geom. Funct. Anal. 2000, Special Volume, Part II, 560–673.

- [12] Fukaya, K., Oh, Y.-G., Ohta, H., and Ono, K., *Lagrangian intersection Floer theory: anomaly and obstruction*, preprint, 2000.
- [13] García-Prada, O., *A direct existence proof for the vortex equations over a compact Riemann surface*, Bull. London Math. Soc. 26 (1994), no. 1, 88–96.
- [14] Gay, D., and Kirby, R., *Constructing symplectic forms on 4-manifolds which vanish on circles*, Geom. Topol. 8 (2004), 743–777.
- [15] Hirzebruch, F., and Hopf, H., *Felder von Flächenelementen in 4-dimensionalen Mannigfaltigkeiten*, Math. Annalen 136 (1958), 156–172.
- [16] Hofer, H., and Salamon, D., *Floer homology and Novikov rings*, The Floer memorial volume, 483–524, Progr. Math., 133, Birkhäuser, Basel, 1995.
- [17] Honda, K., *Local properties of self-dual harmonic 2-forms on a 4-manifold*, J. Reine Angew. Math. 577 (2004), 105–116.
- [18] Hutchings, M., *An index inequality for embedded pseudoholomorphic curves in symplectizations*, J. Eur. Math. Soc. 4 (2002), 313–361.
- [19] Hutchings, M., and Sullivan, M., *Periodic Floer homology of a Dehn twist*, preprint, SG/0410059, 2004.
- [20] Jaffe, A., and Taubes, C., *Vortices and monopoles*, Progress in Physics, 2, Birkhäuser, Boston, Mass., 1980.
- [21] Kronheimer, P., and Mrowka, T., *Monopoles and contact structures*, Invent. Math. 130, 209–255 (1997).
- [22] Kronheimer, P., Mrowka, T., Ozsváth, P., and Szabó, Z., *Monopoles and lens space surgeries*, preprint, GT/0310164, 2003.
- [23] Lalonde, F., *A field theory for symplectic fibrations over surfaces*, Geom. Topol. 8 (2004), 1189–1126.
- [24] Lalonde, F., and McDuff, D., *Symplectic structures on fiber bundles*, Topology (43) 2003, no. 2, 309–347.
- [25] Lazzarini, L., *Existence of a somewhere injective pseudo-holomorphic disc*, Geom. Funct. Anal. 10 (2000), no. 4, 829–862.
- [26] LeBrun, C., *Yamabe constants and the perturbed Seiberg-Witten equations*, Comm. Anal. Geom. 5 (1997), no. 3, 535–553.

- [27] Lee, Y.-J., *Heegaard splittings and Seiberg-Witten monopoles*, preprint, [math.GT/0409536](#), 2004.
- [28] Luttinger, K., and Simpson, C., *A normal form for the birth/flight of closed selfdual 2-form degeneracies*, ETH preprint, unpublished, 1996.
- [29] MacDonald, I., *Symmetric products of an algebraic curve*, *Topology*, 1 (1962), 319–343.
- [30] McDuff, D., *Extending the Hamiltonian group*, preprint, [math.SG/0503268](#), 2005.
- [31] McDuff, D., and Salamon, D., *Introduction to symplectic topology*, second edition, Oxford University Press, New York, 1998.
- [32] Manton, N., and Nasir, S., *Volume of vortex moduli spaces*, *Comm. Math. Phys.* 199 (1999), no. 3, 591–604.
- [33] Muñoz, V., and Wang, B.-L., *Seiberg-Witten-Floer homology of a surface times a circle*, preprint, [math.DG/9905050](#), 1999.
- [34] Ozsváth, P., and Szabó, Z., *Holomorphic disks and topological invariants for closed three-manifolds*, *Ann. of Math. (2)* 159 (2004), no. 3, 1027–1158.
- [35] Póznia, M., *Floer homology and clean intersections*, Northern California Symplectic Geometry Seminar, 119–181, *Amer. Math. Soc. Transl. Ser. 2*, 196, Amer. Math. Soc., Providence, RI, 1999.
- [36] Ran, Z., *A note on Hilbert schemes of nodal curves*, preprint, [AG/0410037](#), 2004.
- [37] Salamon, D., *Lectures on Floer homology*, *Symplectic geometry and topology* (Park City, UT, 1997), 143–229, *IAS/Park City Math. Ser.*, 7, Amer. Math. Soc., Providence, RI, 1999.
- [38] Salamon, D., *Seiberg-Witten invariants of mapping tori, symplectic fixed points, and Lefschetz numbers*, *Turkish J. Math.* 23 (1999), no. 1, 17–143.
- [39] Samols, T., *Vortex scattering*, *Comm. Math. Phys.* 145 (1992), no. 1, 149–179.
- [40] Schwarz, M., *Cohomology operations from S^1 -cobordisms in Floer homology*, PhD thesis, ETH-Zürich, 1995.
- [41] Segal, G., *The definition of conformal field theory*, preprint circa 1990, published in *Topology, geometry and quantum field theory*, ed. U. Tillmann, London Mathematical Society Lecture Note Series, 308, Cambridge University Press, Cambridge, 2004.
- [42] Seidel, P., *Floer homology and the symplectic isotopy problem*, D. Phil thesis, University of Oxford, 1997.

- [43] Seidel, P., *Graded Lagrangian submanifolds*, Bull. Soc. Math. France 128 (2000), 103–146.
- [44] Seidel, P., *More about vanishing cycles and mutation*, Symplectic geometry and mirror symmetry (Seoul, 2000), 429–465, World Sci. Publishing, River Edge, NJ, 2001.
- [45] Seidel, P., *Symplectic Floer homology and the mapping class group*, Pacific J. Math. 206 (2002), no. 1, 219–229.
- [46] Seidel, P., *A long exact sequence for symplectic Floer cohomology*, Topology 42 (2003), no. 5, 1003–1063.
- [47] Seidel, P., and Smith, I., *A link invariant from the symplectic geometry of nilpotent slices*, preprint, [math.SG/0405089](#).
- [48] De Silva, V., *Products in the symplectic Floer homology of Lagrangian intersections*, D. Phil. thesis, Oxford, 1999.
- [49] Smith, I., *Serre-Taubes duality for pseudoholomorphic curves*, Topology 42 (2003), no. 5, 931–979.
- [50] Taubes, C., *The geometry of the Seiberg-Witten invariants*, Surveys in differential geometry, Vol. III (Cambridge, MA, 1996), 299–339, Int. Press, Boston, MA, 1998.
- [51] Taubes, C., *Seiberg-Witten invariants and pseudo-holomorphic subvarieties for self-dual, harmonic 2-forms*, Geom. Topol. 3 (1999), 167–210.
- [52] Usher, M., *The Gromov invariant and the Donaldson-Smith standard surface count*, Geom. Topol. 8 (2004), 565–610.
- [53] Ye, R., *Gromov’s compactness theorem for pseudo holomorphic curves*, Trans. Amer. Math. Soc. 342 (1994), no. 2, 671–694.