Floer homology and the symplectic isotopy problem

Paul Seidel
St. Hugh's College

Thesis submitted for the degree of Doctor in Philosophy
in the University of Oxford

Trinity term, 1997

Floer homology and the symplectic isotopy problem

Paul Seidel, St. Hugh's College

Thesis submitted for the degree of Doctor of Philosophy in the University of Oxford, Trinity term 1997

Abstract

The symplectic isotopy problem is a question about automorphisms of a compact symplectic manifold. It asks whether the relation of symplectic isotopy between such automorphism is finer than the relation of diffeotopy (smooth isotopy). The principal result of this thesis is that there are symplectic manifolds for which the answer is positive; in fact, a large class of symplectic four-manifolds is shown to have this property. This result is the consequence of the study of a special class of symplectic automorphisms, called generalized Dehn twists.

The hard part of studying the symplectic isotopy problem is how to prove that two given symplectic automorphisms are *not* symplectically isotopic. Symplectic Floer homology theory assigns a 'homology group' to any symplectic automorphism. These groups are invariant under symplectic isotopy, hence an obvious candidate for the task. While there is no general procedure for computing the Floer homology groups, it turns out that this is feasible for generalized Dehn twists.

The computation involves an extension of the functorial structure of the Floer homology groups: we introduce homomorphisms induced by certain symplectic fibrations with singularities. Then we use the fact that generalized Dehn twists appear as monodromy maps of such fibrations. These induced maps on Floer homology groups may be of interest independently of their contribution to the symplectic isotopy problem.

The thesis is divided into three parts: the first part presents the symplectic isotopy problem, introduces generalized Dehn twists, and explains the consequences of the determination of their Floer homology groups. The second part is devoted to Floer homology; its main focus are the new induced maps. The final part describes the computation of the Floer homology groups of a generalized Dehn twist.

Acknowledgments. My foremost thanks are to my supervisor, Simon Donaldson, who provided encouragement, criticism, and good ideas for three years. To my former teachers, Albrecht Dold and John Roe, I am equally indebted. Peter Kronheimer pointed out simultaneous resolution, which proved to be essential. I have profited from discussions with many other mathematicians, among them Michael Callahan, Norbert A'Campo, Dusa McDuff, Leonid Polterovich, Dietmar Salamon, Andrey Todorov and Claude Viterbo. A part of this work was done during visits to Tel Aviv University and the Université de Paris-Sud; I am grateful for their hospitality. Many thanks to friends at Oxford and elsewhere who have borne with me while I was writing this thesis. Financial support was given by the German Academic Exchange Service, St. Hugh's College, Oxford, and the European Community.

Contents

Part I

1 Introduction

The symplectic isotopy problem. Let (M, ω) be a compact symplectic manifold. A symplectic automorphism of (M, ω) is a diffeomorphism $\phi : M \longrightarrow M$ such that $\phi^*\omega = \omega$. Recall that two diffeomorphisms $\phi_0, \phi_1 : M \longrightarrow M$ are called *diffeotopic* if they can be connected by a smooth family $(\phi_t)_{0 \leq t \leq 1}$ of diffeomorphisms. Similarly, two symplectic automorphisms are *symplectically isotopic* if there is a diffeotopy (ϕ_t) between them such that all the ϕ_t are symplectic automorphisms. The symplectic isotopy problem is the following

Question. *Are there symplectic automorphisms ϕ_0, ϕ_1 which are diffeotopic but not symplectically isotopic?*

The answer does not change if we consider only the case $\phi_1 = \mathrm{id}$. A symplectic automorphism which is not symplectically isotopic to the identity is called essential. Then the question can be phrased as follows: are there essential symplectic automorphisms which are diffeotopic to the identity? Let $\mathrm{Diff}^+(M)$ be the group of orientation-preserving diffeomorphisms of M (with the C^∞-topology) and $\mathrm{Aut}(M, \omega)$ the subgroup of symplectic automorphisms. Both groups are locally contractible, and any continuous path in $\mathrm{Diff}^+(M)$ or $\mathrm{Aut}(M, \omega)$ can be deformed into a smooth one while keeping its endpoints fixed. Hence our question is this: let

$$\pi_0(\mathrm{Aut}(M, \omega)) \longrightarrow \pi_0(\mathrm{Diff}^+(M))$$

be the homomorphism induced by inclusion. Is its kernel nontrivial? The motivations for studying this problem come from different parts of geometry:

Negative examples. There is a small number of cases in which the topology of $\mathrm{Aut}(M, \omega)$ is completely known. In all these cases the answer to our question turns out to be negative. Moser's theorem on volume forms shows that for two-dimensional M, $\mathrm{Aut}(M, \omega)$ is a deformation retract of $\mathrm{Diff}^+(M)$. The two other examples are due to Gromov [13]. He showed that for $M = \mathbb{C}\mathrm{P}^2$ or $\mathbb{C}\mathrm{P}^1 \times \mathbb{C}\mathrm{P}^1$, the group $\mathrm{Iso}(M)$ of Kähler isometries is a deformation retract of $\mathrm{Aut}(M, \omega)$. $\mathrm{Iso}(\mathbb{C}\mathrm{P}^2) = PU(3)$ is path-connected; hence any two symplectic automorphisms are symplectically isotopic. $\mathrm{Iso}(\mathbb{C}\mathrm{P}^1 \times \mathbb{C}\mathrm{P}^1)$ is a semi-direct product $(PU(2) \times PU(2)) \rtimes \mathbb{Z}/2$. It has two components; one contains the identity and the other one the involution which exchanges the two $\mathbb{C}\mathrm{P}^1$'s. By Gromov's theorem, the same holds for the symplectic automorphism group. In particular, a symplectic automorphism is essential iff it acts nontrivially on homology. Gromov's approach seems to be confined to the class of rational or ruled symplectic four-manifolds.

Gauge theory. Let E be the unique nontrivial principal $SO(3)$-bundle over a closed oriented surface Σ of genus $g \geq 2$. The group $\mathcal{G}(E)$ of gauge transformations acts by pullback on the space $\mathcal{A}^{\text{flat}}(E)$ of flat connections on E. Its maximal connected subgroup $\mathcal{G}_0(E)$ acts freely, and the quotient $\mathcal{N}_g = \mathcal{A}^{\text{flat}}(E)/\mathcal{G}_0(E)$ is a smooth compact manifold. $\mathcal{N}_g$ has a canonical symplectic structure ω_g.

The extended diffeomorphism group $\text{Diff}^+(\Sigma, E)$ is the group of pairs $(f, \hat{f})$ which consist of an oriented diffeomorphism f of Σ and an isomorphism $\hat{f} : E \longrightarrow f^*E$. $\widehat{\Gamma}_g = \pi_0(\text{Diff}^+(\Sigma, E))$ is called the extended mapping class group; it is an extension of $\Gamma_g = \pi_0(\text{Diff}^+(\Sigma))$ by $H^1(\Sigma; \mathbb{Z}/2)$. The operation of pulling back connections defines an action of $\widehat{\Gamma}_g$ on $(\mathcal{N}_g, \omega_g)$. Dostoglou and Salamon [8] raised the question whether the induced homomorphism

$$\widehat{\Gamma}_g \longrightarrow \pi_0(\text{Aut}(\mathcal{N}_g, \omega_g)) \tag{1.1}$$

is injective. Indications in favour of a positive answer come from unpublished work of Callahan, who gave an example of a $\tau \in \widehat{\Gamma}_2$ whose action on $(\mathcal{N}_2, \omega_2)$ is essential even though it acts trivially on $H_*(\mathcal{N}_2)$. In contrast, it seems likely that the differentiable counterpart of (1.1),

$$\widehat{\Gamma}_g \longrightarrow \pi_0(\text{Diff}^+(\mathcal{N}_g)), \tag{1.2}$$

is not injective. This likelihood comes from general results about diffeomorphism groups, e.g. [28, Theorem 12.4 or Theorem 13.3], together with the fact that $\mathcal{N}_g$ is simply connected (recall that in the case of Σ itself, the diffeomorphism group is usually studied through its action on $\pi_1(\Sigma)$, which distinguishes fully between diffeotopy classes).

The different expectations for the homomorphisms (1.1) and (1.2) leave the possibility that the symplectic isotopy question on $(\mathcal{N}_g, \omega_g)$ has a positive answer. For example, this would hold if Callahan's element τ lies in the kernel of (1.2).

Holomorphic functions. This motivation is particularly close to the point of view of the present work. Let π be a holomorphic function from a compact Kähler manifold (E, J, Ω) to a Riemann surface Σ. The set Σ^{crit} of critical values of π is finite, and the regular fibres $E_z = \pi^{-1}(z)$ form a smooth fibre bundle over $\Sigma^{\text{reg}} = \Sigma \setminus \Sigma^{\text{crit}}$. It is a classical idea to study the holomorphic function π through the monodromy homomorphism

$$\pi_1(\Sigma^{\text{reg}}, z_0) \longrightarrow \pi_0(\text{Diff}^+(E_{z_0}))$$

associated to this smooth fibre bundle. The Kähler form Ω induces a symplectic structure $\Omega_z = \Omega|E_z$ on each regular fibre. Since the cohomology class $[\Omega_z]$ is locally constant in z, (E_z, Ω_z) is a locally trivial family of symplectic manifolds, by a theorem of Moser. As a consequence, there is a symplectic monodromy homomorphism $\pi_1(\Sigma^{\text{reg}}, z_0) \longrightarrow \pi_0(\text{Aut}(E_{z_0}, \Omega_{z_0}))$

which fits into a commutative diagram

$$\begin{array}{ccc} & & \pi_0(\mathrm{Aut}(E_{z_0}, \Omega_{z_0})) \\ & \nearrow & \downarrow \\ \pi_1(\Sigma^{\mathrm{reg}}, z_0) & \longrightarrow & \pi_0(\mathrm{Diff}^+(E_{z_0})). \end{array}$$

Whether the symplectic monodromy is actually a finer invariant depends on the answer to the symplectic isotopy problem.

An important source of holomorphic functions (with $\Sigma = \mathbb{C}\mathrm{P}^1$) are the *Lefschetz fibrations* obtained from a generic pencil of hyperplane sections of a projective variety after blowing up its base. Lefschetz used the monodromy action on the homology of a regular fibre to study the topology of varieties by induction on their dimension. One might imagine a parallel attempt to analyse the symplectic geometry of algebraic varieties in which the differentiable monodromy would be replaced by its symplectic counterpart. For example, Lefschetz fibrations with the same fibres and differentiable monodromy but with different symplectic monodromy might be a potential source of inequivalent symplectic structures on the same smooth manifold. Recent work of Donaldson provides the foundation for extending such an approach even to non-Kähler symplectic manifolds.

Lefschetz fibrations have an interesting global structure, but their local structure is simple. For a holomorphic function with critical points of a more complex kind even the local aspect, that is, the symplectic monodromy along a small loop which winds around a single critical value of π, is interesting. The smooth monodromy along such small loops is a much-studied object; the relevance of symplectic geometry to questions of this kind has been advocated by Arnol'd [1].

Fragile automorphisms. The symplectic isotopy problem compares the topology of $\mathrm{Aut}(M,\omega)$ with that of $\mathrm{Diff}^+(M)$. A related question is how $\mathrm{Aut}(M,\omega)$ changes under variations of ω. The contribution of the present work to this question is the discovery of an unexpected phenomenon which we have christened *fragility*. Roughly speaking, a symplectic automorphism is fragile if after an arbitrarily small perturbation of the symplectic form, it becomes symplectically isotopic to the identity. The precise formulation uses deformations $(\omega_t, \phi_t)_{0 \leq t < \epsilon}$, where ω_t is a smooth family of symplectic forms and ϕ_t a smooth family of diffeomorphisms such that $\phi_t \in \mathrm{Aut}(M, \omega_t)$ for all t.

Definition 1.1. $\phi \in \mathrm{Aut}(M,\omega)$ is fragile if there is such a deformation with $\omega_0 = \omega$, $\phi_0 = \phi$, and such that for all $t > 0$, ϕ_t is isotopic to the identity symplectically, that is, within $\mathrm{Aut}(M, \omega_t)$.

If ϕ is symplectically isotopic to the identity, it is fragile for trivial reasons (take the constant deformation). The interesting case is when ϕ is essential: since a fragile automorphism is diffeotopic to the identity, any such example provides a positive answer to the symplectic isotopy problem. Note that by concentrating on the notion of fragility, we narrow the class of symplectic manifolds under consideration:

Lemma 1.2. *On a compact symplectic manifold with* $b_2(M) = 1$ *every fragile automorphism is symplectically isotopic to the identity.*

Proof. Let (ω_t, ϕ_t) be a deformation of (ω, ϕ) as in Definition 1.1. Since $b^1(M) = 1$ and we are free to rescale ω_t, we may assume that $[\omega_t] = [\omega] \in H^2(M, \mathbb{R})$ for all t. Moser's theorem on deformations of symplectic forms says that there is a smooth family (ρ_t) of diffeomorphisms of t with $\rho_0 = \mathrm{id}$ and such that $\rho_t^* \omega_t = \omega$ for all t. Fix some $t_0 > 0$ and an isotopy (ψ_s) between $\psi_0 = \phi_{t_0}$ and $\psi_1 = \mathrm{id}$ in $\mathrm{Aut}(M, \omega_{t_0})$. The conjugates $\phi'_t = \rho_t^{-1} \phi_t \rho_t$ and $\psi'_s = \rho_{t_0}^{-1} \psi_s \rho_{t_0}$ lie in $\mathrm{Aut}(M, \omega)$. ϕ'_{t_0} is symplectically isotopic to ϕ through (ϕ'_t) and to the identity through (ψ'_s). Therefore ϕ is symplectically isotopic to the identity. □

We will show that many compact symplectic four-manifolds (M, ω) with $b_1(M) = 0$ have automorphisms which are essential and fragile. The principal requirement of our approach is that (M, ω) must contain an embedded Lagrangian two-sphere. For such manifolds, we prove the existence of an essential and fragile automorphism under a certain algebraic condition, stated in terms of its quantum homology ring. This algebraic condition can be dealt with easily as long as (M, ω) is not rational or ruled. In this way one obtains the following result:

Theorem 1.3. *Let* (M, ω) *be a compact symplectic four-manifold, with* $b_1(M) = 0$*, which contains an embedded Lagrangian two-sphere. Assume that* (M, ω) *is minimal and irrational, and that* $\dim H_2(M; \mathbb{Z}/2) \geq 3$*. Then* (M, ω) *admits an essential and fragile symplectic automorphism.*

There is a slightly more complicated version of this theorem which does not require that (M, ω) is minimal (Theorem 5.4). One can also consider rational four-manifolds, but then the condition on the quantum homology ring needs to be checked separately in every instance.

Our approach is particularly effective for algebraic surfaces, where the existence of a Lagrangian two-sphere can often be deduced from algebro-geometric considerations. As an example, consider the class of algebraic surfaces which are complete intersections in some projective space $\mathbb{CP}^n$, with the induced symplectic structures. The first such surfaces are $\mathbb{CP}^2$ itself and $\mathbb{CP}^1 \times \mathbb{CP}^1$ which is a quadric hypersurface in $\mathbb{CP}^3$. As explained before, the answer to the symplectic isotopy problem is negative in these two cases. In all other cases, the answer is positive:

Theorem 1.4. *Any complete intersection of complex dimension two other than* $\mathbb{C}\mathrm{P}^2$ *or* $\mathbb{C}\mathrm{P}^1 \times \mathbb{C}\mathrm{P}^1$ *admits an essential and fragile symplectic automorphism.*

The remainder of this part is structured as follows: the next section introduces a special class of automorphisms of a symplectic four-manifold called generalized Dehn twists. We prove that the square of any such automorphism is fragile. In section 3 we present some of the main properties of Floer homology and state the result of its computation for generalized Dehn twists. We also explain how to apply this to the symplectic isotopy problem. The outcome is summarized in Corollary 3.6, which is the main result of this thesis. Section 4 explains how to use degenerations of algebraic surfaces to produce Lagrangian two-spheres on them. In section 5 we recall some recent results on symplectic four-manifolds which are not rational or ruled, and prove Theorem 5.4. Section 6 contains a few sample computations for rational algebraic surfaces which are necessary to prove Theorem 1.4.

Notation. Unless otherwise specified, (M, ω) always denotes a compact symplectic four-manifold with $b_1(M) = 0$.

This is the class of symplectic manifolds in which we will work throughout. The condition $b_1(M) = 0$ can be removed if one replaces symplectic isotopy by Hamiltonian isotopy everywhere. This notion of isotopy, which coincides with symplectic isotopy for $b_1(M) = 0$ but is more restrictive in general, is the natural one in Floer homology theory. The restriction to four-dimensional manifolds has two reasons: one is of a technical nature (the construction of Floer homology groups is simpler for a class of symplectic manifolds which contains all four-dimensional ones). The other reason is that our source of fragile automorphisms is an elementary feature which is specific to four dimensions. Higher-dimensional examples can be produced from this by taking products. For example, the products $(M, \omega) \times \cdots \times (M, \omega)$ of a $K3$ surface (M, ω) with itself admit an essential and fragile symplectic automorphism. However, we will not pursue this further here.

2 Generalized Dehn twists

The first step towards a positive answer to the symplectic isotopy problem is to find interesting symplectic automorphisms which are diffeotopic to the identity. Our examples are constructed from a common local model; we begin by explaining this model.

Let T^*S^2 be the cotangent bundle of S^2 and η its canonical symplectic form. The zero section $S^2 \subset T^*S^2$ is a Lagrangian submanifold. We use the representation

$$T^*S^2 = \{(u,v) \in \mathbb{R}^3 \times \mathbb{R}^3 \mid |u| = 1 \text{ and } \langle u,v\rangle = 0\}.$$

In these coordinates, $\eta = -\sum_j du_j \wedge dv_j$ and $S^2 = \{(u,v) \in T^*S^2 \mid v = 0\}$. Let $T^*_\epsilon S^2 = \{(u,v) \in T^*S^2 \mid |v| < \epsilon\}$ be the subbundle of ϵ-discs, for $\epsilon > 0$. We denote the subgroup of automorphisms $\phi \in \mathrm{Aut}(T^*S^2, \eta)$ which are supported inside $T^*_\epsilon S^2$ (that is, $\phi = \mathrm{id}$ outside some compact subset of $T^*_\epsilon S^2$) by $\mathrm{Aut}^c(T^*_\epsilon S^2, \eta)$.

Consider the Hamiltonian function $\mu(u,v) = |v|$ on $T^*S^2 \setminus S^2$. It is well-known that $\frac{1}{2}\mu^2$ induces the geodesic flow (this is true for the corresponding function on the cotangent bundle of any Riemannian manifold). Given this, it is not difficult to see what the flow of μ is: it transports every cotangent vector along the geodesic emanating from it *with unit speed*, irrespective of how long the vector is. On S^2, all geodesics are closed and of period 2π; therefore μ induces a Hamiltonian circle action on $T^*S^2 \setminus S^2$. Since it is clear what the geodesic flow is in our coordinates, we can write down this action explicitly:

$$\sigma(e^{it})(u,v) = \left(\cos(t)u + \sin(t)\frac{v}{|v|}, \cos(t)v - \sin(t)u|v|\right).$$

$\sigma(-1)(u,v) = (-u,-v)$ can be extended to an involution of T^*S^2. We call this involution the antipodal map and denote it by A.

Notation. The Hamiltonian flow induced by a (time-indepedent or time-dependent) Hamiltonian function H will be denoted by $(\phi^H_t)_{t\in\mathbb{R}}$.

Take a function $r \in C^\infty(\mathbb{R},\mathbb{R})$. The flow induced by $r(\mu)$ on $T^*S^2 \setminus S^2$ is

$$\phi^{r(\mu)}_t(x) = \sigma(e^{itr'(\mu(x))})(x); \tag{2.1}$$

this is an elementary fact which holds for any Hamiltonian circle action. If r is even, $r(\mu(u,v)) = \sqrt{r}(|v|^2)$ is a smooth function on all of T^*S^2 and every point in S^2 is a critical point of it. As a consequence (2.1) can be extended to a Hamiltonian flow on T^*S^2 which keeps S^2 pointwise fixed.

Lemma 2.1. *Let $r \in C^\infty(\mathbb{R},\mathbb{R})$ be a function which satisfies*

$$r(t) = 0 \text{ for } t \geq \tfrac{\epsilon}{2} \quad \text{and} \quad r(-t) = r(t) - t \text{ for all } t. \tag{2.2}$$

(a) The map τ of T^*S^2 into itself given by

$$\tau(x) = \begin{cases} \phi_{2\pi}^{r(\mu)}(x) & x \notin S^2, \\ A(x) & x \in S^2 \end{cases}$$

is a symplectic automorphism supported inside $T^*_\epsilon S^2$.

(b) τ commutes with the antipodal map A.

(c) The automorphisms obtained from different choices of r are isotopic in $\mathrm{Aut}^c(T^*_\epsilon S^2, \eta)$.

Proof. (a) The second condition in (2.2) shows that $R(t) = r(t) - \frac{1}{2}t$ is even. As explained above, this implies that $R(\mu)$ induces a Hamiltonian flow on T^*S^2. A simple computation using (2.1) shows that the time-2π map of this flow is $A \circ \tau$. It follows that τ itself is also a symplectic automorphism. The first condition in (2.2) says that for $(u, v) \notin T^*_{\epsilon/2}S^2$ we have $e^{2\pi i r'(|v|)} = 1$, hence $\tau(u, v) = (u, v)$ by (2.1).

(b) is clear from (2.1).

(c) Let r, r' be two functions satisfying (2.2) and τ, τ' the corresponding symplectic automorphisms. By (2.1), $\tau^{-1}\tau'$ is the time-2π map of the flow induced by $\delta(\mu)$, where $\delta = r' - r$. (2.2) implies that δ is even and $\delta(\mu)$ is supported inside $T^*_\epsilon S^2$. Therefore its Hamiltonian flow defines an isotopy from $\tau^{-1}\tau'$ to the identity in $\mathrm{Aut}^c(T^*_\epsilon S^2, \eta)$. □

Let (M, ω) be a compact symplectic four-manifold and $V \subset M$ an embedded Lagrangian two-sphere. The symplectic geometry of M near V is described by the following tubular neighbourhood theorem:

Lemma 2.2 (Weinstein [32]). *(a) There is an $\epsilon > 0$ and a symplectic embedding $i : T^*_\epsilon S^2 \longrightarrow M$ with $i(S^2) = V$.*

(b) Let i, i' be two embeddings as above and assume that $i^{-1}i'|S^2 \in \mathrm{Diff}(S^2)$ is diffeotopic to the identity. Then there is a $\delta < \epsilon$ such that $i|T^*_\delta S^2$ can be deformed into $i'|T^*_\delta S^2$ within the space of symplectic embeddings which map S^2 to V. □

Proposition 2.3. *Choose an $\epsilon > 0$, an embedding i as in Lemma 2.2, and a function r as in Lemma 2.1. Then*

$$\tau_V(x) = \begin{cases} i\tau i^{-1}(x) & x \in \mathrm{im}(i), \\ x & x \notin \mathrm{im}(i) \end{cases}$$

defines a symplectic automorphism of M. τ_V maps V to itself but reverses its orientation. It is independent of r and i up to symplectic isotopy.

We call τ_V the *generalized Dehn twist along V* because maps of this kind are natural four-dimensional analogues of the (positive) Dehn twists along simple closed curves on a surface. Generalized Dehn twists are common in singularity theory, but they are usually considered only as diffeomorphisms; the observation that their local model τ is symplectic was made by Arnol'd.

Most of Proposition 2.3 follows immediately from the previous Lemmata. We need to make one remark on the uniqueness of τ_V up to symplectic isotopy. Lemma 2.1 and Lemma 2.2(b) show that $[\tau_V] \in \pi_0(\mathrm{Aut}(M,\omega))$ depends only on the diffeotopy class of $i|S^2$. Since we have not fixed an orientation of V, there are two such classes, represented by any embedding i and its opposite $i' = i \circ A$. Let τ_V and τ_V' be the automorphisms of M constructed from the same local model using these two embeddings. By Lemma 2.1(b),

$$\tau_V'(x) = iA\tau A^{-1}i^{-1}(x) = i\tau i^{-1}(x) = \tau_V(x)$$

for $x \in \mathrm{im}(i)$ and hence $\tau_V = \tau_V'$. This completes the proof of the independence of $[\tau_V]$. Our discussion has a parallel in the two-dimensional case, where the distinction between positive and negative Dehn twists is also independent of the orientation of the curve.

For completeness' sake we include a stronger uniquess result which will not be used later; it is again an analogue of a well-known property of ordinary Dehn twists. Recall that two submanifolds $V_0, V_1 \subset M$ are called isotopic if there is a submanifold $\mathbf{V} \subset M \times [0;1]$ which intersects $M \times \{t\}$ transversely for all t and such that $V_t = \mathbf{V} \cap (M \times \{t\})$ for $t = 0, 1$. An isotopy is Lagrangian if $\mathbf{V} \cap (M \times \{t\}) \subset M$ is a Lagrangian submanifold for all t.

Lemma 2.4. *The symplectic isotopy class of τ_V depends only on the Lagrangian isotopy class of V.*

Proof. It is convenient to consider an apparently stronger notion: two Lagrangian submanifolds $V_0, V_1 \subset M$ are *ambient isotopic* if there is a symplectic isotopy $(\phi_t)_{0 \leq t \leq 1}$ with $\phi_0 = \mathrm{id}$ and $\phi_1(V_0) = V_1$. It is clear from the construction that $[\tau_V] \in \pi_0(\mathrm{Aut}(M,\omega))$ depends on V only up to ambient isotopy. However, for two-spheres, as for all compact Lagrangian submanifolds with vanishing first Betti number, any Lagrangian isotopy can be embedded into an ambient isotopy. □

Lemma 2.5 (Picard-Lefschetz formula). *The action of τ_V on homology is given by*

$$(\tau_V)_*(d) = \begin{cases} d + (d \cdot [V])[V] & d \in H_2(V) \\ d & d \in H_i(V),\ i \neq 2. \end{cases}$$

Proof. τ_V is trivial outside a tubular neighbourhood U of V. For such a map there is a variation homomorphism $\operatorname{var}(\tau_V) : H_*(M, M \setminus U) \longrightarrow H_*(U)$ such that the diagram

$$\begin{array}{ccc} H_i(M) & \xrightarrow{(\tau_V)_* - \mathrm{id}} & H_i(M) \\ \downarrow & & \uparrow \\ H_i(M, M \setminus U) & \xrightarrow{\operatorname{var}(\tau_V)} & H_i(U) \end{array}$$

commutes (the unmarked arrows are the obvious maps). Since $H_i(M, M \setminus U) \cong H^{4-i}(U) \cong H^{4-i}(S^2)$ and $H_i(U) \cong H_i(S^2)$, the variation vanishes for $i \neq 2$. For $i = 2$, both groups are infinite cyclic: the isomorphism $H_2(M, M \setminus U) \cong \mathbb{Z}$ is given by $d \longmapsto d \cdot [V]$ and $H_2(U)$ is generated by $[V]$. It follows that there is a $\lambda \in \mathbb{Z}$ such that

$$(\tau_V)_*(d) = d + \lambda(d \cdot [V])[V]$$

for all $d \in H_2(M)$. Take $d = [V]$: the local model, Lemma 2.2, shows that the self-intersection number $[V] \cdot [V]$ is the Euler number of the oriented two-plane bundle $T^*S^2 \longrightarrow S^2$, which is $e(T^*S^2) = -e(TS^2) = -\chi(S^2) = -2$. Therefore $(\tau_V)_*[V] = (1 - 2\lambda)[V]$. But since $\tau_V|V$ is orientation-reversing, $(\tau_V)_*[V] = -[V]$ and hence $\lambda = 1$. $\square$

Recall that we are searching for a class of symplectic automorphisms which are diffeotopic to the identity. This is certainly not true for τ_V since it induces a nontrivial map on homology. However, the Picard-Lefschetz formula shows that this induced map is an involution; consequently, the square τ_V^2 acts trivially on homology. We will now prove a much stronger result:

Proposition 2.6. *The square of a generalized Dehn twist is fragile. In particular, it is diffeotopic to the identity.*

The proof will be carried out in the local model $(T^*_\epsilon S^2, \eta)$. Let $\beta \in \Omega^2(S^2)$ be the standard volume form, that is, $\beta_u(\xi, \eta) = \langle u, \xi \times \eta \rangle$, and $\pi : T^*S^2 \longrightarrow S^2$ the projection. In principle, we want to deform η by adding small multiples of $\pi^*\beta$. Some adjustments have to be made to keep that deformation trivial outside $T^*_\epsilon S^2$. Since $H^2(T^*S^2 \setminus S^2, \mathbb{R}) = 0$ ($T^*S^2 \setminus S^2$ is homotopy equivalent to $\mathbb{RP}^3$), there is a one-form $\lambda \in \Omega^1(T^*S^2 \setminus S^2)$ such that $d\lambda = \pi^*\beta | T^*S^2 \setminus S^2$. Choose a function $\psi \in C^\infty(T^*S^2, \mathbb{R})$ with $\psi(u, v) = 0$ for $|v| \leq \frac{2}{3}\epsilon$ and $\psi(u, v) = 1$ for $|v| \geq \frac{3}{4}\epsilon$. $\alpha = \pi^*\beta - d(\psi\,\lambda)$ is a compactly supported closed two-form on $T^*_\epsilon S^2$. Hence there is a $\delta > 0$ such that $\eta_s = \eta + s\alpha$ is a symplectic form for $0 \leq s < \delta$.

Choose a function r as in Lemma 2.1 and such that $r(t) = \frac{1}{2}t$ for $|t| \leq \frac{\epsilon}{4}$. Then τ^2 is given by

$$\tau^2(u, v) = \begin{cases} \sigma(e^{4\pi i\, r'(|v|)})(u, v) & \frac{\epsilon}{4} \leq |v| \leq \frac{\epsilon}{2}, \\ (u, v) & \text{otherwise.} \end{cases} \tag{2.3}$$

$a(t) = 4\pi r'(t)$ goes from $a(\frac{\epsilon}{4}) = 2\pi$ to $a(\frac{\epsilon}{2}) = 0$. Let us imagine for a moment that the circle action σ could be extended smoothly to the whole of T^*S^2. Then we could define a symplectic isotopy from τ^2 to the identity within $\mathrm{Aut}^c(T^*_\epsilon S^2, \eta)$ by

$$\phi_t(u,v) = \begin{cases} \sigma(e^{it\,a(\mu(u,v))})(u,v) & |v| \leq \frac{\epsilon}{2}, \\ (u,v) & \text{otherwise.} \end{cases}$$

What hinders us from actually doing this is that σ is not continuous at S^2. However, this lack of continuity can be removed by deforming the circle action in a way which is compatible with the deformation of η. More precisely, we will construct a family $(\sigma_s)_{0<s<\delta}$ of Hamiltonian circle actions on $(T^*_{2\epsilon/3}S^2, \eta_s)$ such that σ_s converges to σ away from S^2 as $s \to 0$. The construction necessitates a short digression on Hamiltonian $SO(3)$-actions. Our convention is to write the moment map of such an action as an $\mathbb{R}^3$-valued function, using the cross-product and scalar product to identify $\mathbb{R}^3$ with $\mathfrak{so}_3$ and its dual.

Lemma 2.7. *Let ρ be a Hamiltonian $SO(3)$-action on a symplectic manifold (N, η), with moment map $m : N \longrightarrow \mathbb{R}^3$. Then the function $h(x) = |m(x)|$ induces a Hamiltonian circle action ζ on $N \setminus m^{-1}(0)$, given by*

$$\zeta(e^{it})(x) = \rho(\exp(t\frac{m(x)}{|m(x)|}))(x). \tag{2.4}$$

Proof. Let $K_\xi \in \Gamma(TN)$ be the infinitesimal action of $\xi \in \mathbb{R}^3 \cong \mathfrak{so}_3$. Recall that moment maps are equivariant with respect to the coadjoint action; in our terms, this translates into

$$(K_\xi.m)(x) = \xi \times m(x). \tag{2.5}$$

The Hamiltonian vector field of h is

$$X(x) = K_{\frac{m(x)}{|m(x)|}}(x),$$

because $-i_X\omega = d\langle m(x), m(x)/|m(x)|\rangle = d|m(x)|$. By (2.5),

$$(X.m)(x) = \frac{m(x)}{|m(x)|} \times m(x) = 0.$$

Therefore m is constant along the orbits of X, which implies that the flow of X is given by (2.4). It is a circle action because $\exp(2\pi g) = I$ for any $g \in \mathfrak{so}_3$ which has length one. □

Consider the standard $SO(3)$-action $\bar{\rho}$ on S^2 and the induced action ρ on T^*S^2, given in our coordinates by $\rho(A)(u,v) = (Au, Av)$. The moment map of ρ with respect to η is $m(u,v) = u \times v$, and the moment map of $\bar{\rho}$ with respect to β is $\bar{m}(u) = -u$. By definition, η_s agrees with $\omega + s\,\pi^*\beta$ on $U = T^*_{2\epsilon/3}S^2$. Since π is $SO(3)$-equivariant, it follows that $\rho|U$ is Hamiltonian with respect to $\eta_s|U$, with moment map

$$m_s(u,v) = m(u,v) + s\,\bar{m}(u) = u \times v - su.$$

For $s > 0$, m_s is nowhere zero; therefore $\mu_s = |m_s|$ induces a Hamiltonian circle action σ_s on (U, η_s). The expression for ζ given in Lemma 2.7 yields an explicit formula for σ_s.

There is a $\delta' \in (0;\delta]$ such that $\mu_s(u,v) = |u \times v - su|$ satisfies

$$\mu_s(u,v) \geq \frac{\epsilon}{2} \text{ for } |v| \geq \tfrac{3}{5}\epsilon \quad \text{and} \quad \mu_s(u,v) \leq \frac{\epsilon}{4} \text{ for } |v| \leq \tfrac{1}{5}\epsilon$$

for all $s < \delta'$. Let $T_s : T^*S^2 \longrightarrow T^*S^2$, $0 < s < \delta'$, be the family of maps defined by

$$T_s(u,v) = \begin{cases} \sigma_s(e^{ia(\mu_s(u,v))})(u,v) & \frac{1}{5}\epsilon \leq |v| \leq \frac{3}{5}\epsilon, \\ (u,v) & \text{otherwise.} \end{cases} \tag{2.6}$$

These maps are smooth because $e^{ia(t)} = 1$ for $t \notin [\frac{\epsilon}{4};\frac{\epsilon}{2}]$. As $s \to 0$, η_s converges to η, and μ_s converges smoothly to μ on the region $\frac{1}{5}\epsilon \leq |v| \leq \frac{3}{5}\epsilon$. Hence σ_s converges smoothly to σ on that region. By comparing (2.6) with (2.3) it follows that $(T_s)_{0<s<\delta'}$ is a smooth deformation of τ^2.

Because it is trivial outside U, T_s is the time-4π map of the Hamiltonian function $H_s = r(\mu_s)$ with respect to the symplectic structure η_s; in particular, it lies in $\mathrm{Aut}^c(T^*_\epsilon S^2, \eta_s)$. Since $r(t) = 0$ for $t \geq \epsilon/2$, H_s is supported in $T^*_\epsilon S^2$ and its flow provides an isotopy from T_s to the identity within $\mathrm{Aut}^c(T^*_\epsilon S^2, \eta_s)$ for any $s < \delta'$. Explicitly, this isotopy $(T_{s,t})_{0\leq 1}$ is given by

$$T_{s,t}(x) = \begin{cases} \sigma_s(e^{it\,a(\mu_s(x))})(x) & x \in U, \\ x & x \notin U. \end{cases}$$

It follows that $(\eta_s, T_s)_{0<s<\delta'}$ *is a deformation of* (η, τ^2) *in the sense of Definition 1.1, supported inside* $T^*_\epsilon S^2$. For sufficiently small $\epsilon > 0$, this deformation and the maps $(T_{s,t})$ can be transported from the local model to a neighbourhood of any given Lagrangian two-sphere V in a symplectic four-manifold (M, ω). This completes the proof of Proposition 2.6.

Remark 2.8. Generalized Dehn twists can be defined in any dimension: it is sufficient to replace T^*S^2 by T^*S^r as a local model. The fact that the square of a generalized Dehn twist acts trivially on homology is true for any even r. However, the fragility of τ^2 is special to four dimensions. Indeed, a nontrivial deformation of the symplectic form localized near a Lagrangian sphere is possible only in four dimensions because $H^2_c(T^*S^r) = 0$

for $r \neq 2$. A more conceptual proof of Proposition 2.6 reveals that the fragility of τ^2 is a consequence of an algebro-geometric construction which works only in complex dimension two, namely, the *simultaneous resolution* of ordinary double points [2]. We have chosen to present the argument in elementary terms because the connection between generalized Dehn twists and singularities will only be made later on, in Part III.

3 The Floer homology of a generalized Dehn twist

Proposition 2.6 – which says that the square τ_V^2 of any generalized Dehn twist is fragile – forms the easier part of our approach to the symplectic isotopy problem. The difficult part is to prove that there are (M, ω) and V such that τ_V^2 is essential, or equivalently, such that τ_V is not symplectically isotopic to τ_V^{-1}. The theory which we use for this purpose assigns to (M, ω) a ring $QH_*(M, \omega)$ and to each $\phi \in \mathrm{Aut}(M, \omega)$ a module $HF_*(\phi)$ over $QH_*(M, \omega)$ which is unchanged under symplectic isotopy. $QH_*(M, \omega)$ is called the *quantum homology ring* of (M, ω) and $HF_*(\phi)$ the *Floer homology* of ϕ. Our main result determines $HF_*(\tau_V)$ for any V. By a simple duality property, we obtain $HF_*(\tau_V^{-1})$ at the same time. A comparison of these two $QH_*(M, \omega)$-modules yields conditions under which τ_V is essential.

We will use (with minor modifications) the quantum homology ring as defined by Ruan-Tian [23] and McDuff-Salamon [20]. In the version by Ruan and Tian, this definition works for the class of 'weakly monotone' symplectic manifolds; this is sufficient for our purpose since any four-manifold satisfies this condition (we remark in passing that there are more recent approaches which work for all compact symplectic manifolds). Additively $QH_*(M, \omega)$ does not depend on ω: it is the $\mathbb{Z}/2$-graded group obtained from the ordinary homology of M with coefficients in a certain field Λ by reducing the grading, that is,

$$QH_0(M, \omega) = H_{\mathrm{even}}(M; \Lambda), \quad QH_1(M, \omega) = H_{\mathrm{odd}}(M; \Lambda).$$

The choice of coefficients is dictated in part by convenience and in part by the exigencies of the theory of pseudo-holomorphic curves.

Definition 3.1. Let Λ be the set of functions $c : \mathbb{R} \longrightarrow \mathbb{Z}/2$ which satisfy the following condition: for any $C \in \mathbb{R}$, there are only finitely many $\epsilon \leq C$ such that $c_\epsilon \neq 0$. Addition and multiplication on Λ are defined by

$$(c^{(1)} + c^{(2)})_\epsilon = c^{(1)}_\epsilon + c^{(2)}_\epsilon, \quad (c^{(1)} c^{(2)})_\epsilon = \sum_{\delta \in \mathbb{R}} c^{(1)}_\delta c^{(2)}_{\epsilon-\delta};$$

the finiteness condition ensures that $\sum_{\delta \in \mathbb{R}} c^{(1)}_\delta c^{(2)}_{\epsilon-\delta}$ has only finitely many nonzero terms. We call Λ the *universal Novikov field* over $\mathbb{Z}/2$.

The proof that Λ is indeed a field is not difficult; it can be found e.g. in [14]. It is customary to write elements of Λ as formal power series with real exponents:

$$c = \sum_{\epsilon \in \mathbb{R}} c_\epsilon t^\epsilon,$$

because addition and multiplication take on the familiar form in this notation. For elements of $QH_*(M,\omega)$ we use the same notation with coefficients $c_\epsilon \in H_*(M;\mathbb{Z}/2)$. This is justified by the Künneth isomorphism $QH_*(M,\omega) \cong H_*(M;\mathbb{Z}/2) \otimes \Lambda$.

The product on $QH_*(M,\omega)$ is defined in terms of the threefold Gromov-Witten invariants of (M,ω). These invariants are a collection of symmetric trilinear forms

$$\Phi_A : H_*(M;\mathbb{Z}/2) \otimes H_*(M;\mathbb{Z}/2) \otimes H_*(M;\mathbb{Z}/2) \longrightarrow \mathbb{Z}/2$$

indexed by $A \in H_2(M;\mathbb{Z})$. Roughly speaking, $\Phi_A(x,y,z)$ is the number mod 2 of pseudo-holomorphic spheres in the homology class A which pass through suitable cycles representing x, y and z. By using Poincaré duality and summing over all A with the same energy $\omega(A) \in \mathbb{R}$, we convert the Gromov-Witten invariants into a family $(*_\epsilon)_{\epsilon\in\mathbb{R}}$ of symmetric bilinear forms

$$H_*(M;\mathbb{Z}/2) \otimes H_*(M;\mathbb{Z}/2) \longrightarrow H_*(M;\mathbb{Z}/2)$$

with the following properties:

(1) $*_\epsilon = 0$ for all negative ϵ.

(2) any interval $[0;C] \subset \mathbb{R}$ contains only finitely many ϵ with $*_\epsilon \neq 0$.

(3) $*_\epsilon$ respects the grading of $H_*(M;\mathbb{Z}/2)$ mod 2.

(4) $*_0$ is the ordinary intersection product.

(5) $[M] *_\epsilon x = 0$ for all $x \in H_*(M;\mathbb{Z}/2)$ and $\epsilon > 0$.

The *quantum product* $*$ on $QH_*(M,\omega)$ is defined by

$$c^{(1)} * c^{(2)} = \sum_{\epsilon\in\mathbb{R}} \left(\sum_{\delta_1,\delta_2\in\mathbb{R}} c^{(1)}_{\delta_1} *_{\epsilon-\delta_1-\delta_2} c^{(2)}_{\delta_2} \right) t^\epsilon.$$

From the first two properties of $*_\epsilon$ it follows that the sum over δ_1, δ_2 contains only finitely many nonzero terms and that the expression on the r.h.s. is an element of $QH_*(M,\omega)$. By definition, $*$ is Λ-bilinear. It is commutative because the forms $*_\epsilon$ are symmetric, and it is also associative (a much deeper result). Property (3) shows that $*$ is $\mathbb{Z}/2$-graded. The last two properties imply that $u = [M]t^0$ is the unit of $*$. By property (1) and (4), the ordinary intersection product $x \circ y$ of $x, y \in H_*(M; \mathbb{Z}/2)$ is the leading term of $x\,t^0 * y\,t^0$; that is,

$$x\,t^0 * y\,t^0 = (x \circ y)t^0 + \sum_{\epsilon > 0} c_\epsilon t^\epsilon. \tag{3.1}$$

In this sense the quantum product is a deformation of the intersection product.

Let V be a Lagrangian sphere in M. The ideal in $(QH_*(M, \omega), *)$ generated by $v = [V]t^0$ will be denoted by I_v.

Lemma 3.2. *Let $x \in H_2(M; \mathbb{Z}/2)$ be a class which has nonzero mod 2 intersection number with V (such classes exist by Poincaré duality). Set $w = x\,t^0 \in QH_*(M, \omega)$. v and $w * v$ form a basis of I_v as a vector space over Λ; in particular,* $\dim_\Lambda I_v = 2$.

Proof. First, v itself is nonzero. Indeed, $v = 0$ would imply that $[V] \in H_2(M; \mathbb{Z})$ is divisible by 2 because, as a part of the universal coefficient theorem, the canonical homomorphism $H_2(M; \mathbb{Z}) \otimes \mathbb{Z}/2 \longrightarrow H_2(M; \mathbb{Z}/2)$ is injective. But we know that $[V]$ is not divisible by 2 because $[V] \cdot [V] = -2$ is not divisible by 4.

$v * w$ is linearly independent of v over Λ because

$$v * w = [\mathrm{pt}]\,t^0 + (\text{higher order terms})$$

by (3.1). It remains to prove that $v * c$ is a linear combination of v and $v * w$ with Λ-coefficients for any $c \in QH_*(M, \omega)$. Because the Gromov-Witten invariants are invariants of the symplectic structure, the action of $\mathrm{Aut}(M, \omega)$ on $QH_*(M, \omega) = H_*(M; \Lambda)$ preserves $*$, hence

$$(\tau_V)_*((w + v) * c) = (\tau_V)_*(w + v) * (\tau_V)_*(c). \tag{3.2}$$

The Picard-Lefschetz formula (Lemma 2.5) says that

$$(\tau_V)_*(z) = z + (z \cdot_\Lambda v)v$$

where $z \cdot_\Lambda v \in \Lambda$ is the ordinary intersection number with coefficients in Λ. In particular, because $(w + v) \cdot_\Lambda v = -t^0$, $(\tau_V)_*(w + v) = w$. Inserting this into (3.2) yields

$$v * c = (c \cdot_\Lambda v)(w * v) - [((w + v) * c) \cdot_\Lambda v]\,v. \quad \square$$

The Floer homology of a symplectic automorphism $\phi \in \mathrm{Aut}(M, \omega)$ is a $\mathbb{Z}/2$-graded Λ-vector space

$$HF_*(\phi) = HF_0(\phi) \oplus HF_1(\phi)$$

equipped with a graded Λ-bilinear map $\hat{*} : QH_*(M,\omega) \otimes HF_*(\phi) \longrightarrow HF_*(\phi)$ which makes it into a unital module over $(QH_*(M,\omega), *)$. $\hat{*}$ is called the *quantum module structure.* Floer homology has the following two basic properties:

(Isotopy invariance) If ϕ_0 and ϕ_1 are symplectically isotopic, $HF_*(\phi_0)$ and $HF_*(\phi_1)$ are isomorphic as modules over $QH_*(M,\omega)$.

(Poincaré duality) For every ϕ there is a nondegenerate graded pairing

$$\langle \cdot, \cdot \rangle : HF_*(\phi^{-1}) \otimes_\Lambda HF_*(\phi) \longrightarrow \Lambda$$

which satisfies

$$\langle c \,\hat{*}\, x, y \rangle = \langle x, c \,\hat{*}\, y \rangle \tag{3.3}$$

for all $c \in QH_*(M,\omega)$, $x \in HF_*(\phi^{-1})$ and $y \in HF_*(\phi)$.

The construction of this invariant is discussed in Part II. Our application of Floer homology is based on the following observation: the classical intersection rings $(H_*(M), \circ)$ are rings of a very special type. This is a consequence of classical Poincaré duality, which says that there is a nondegenerate pairing (the intersection pairing)

$$\langle \cdot, \cdot \rangle : H_*(M) \otimes H_*(M) \longrightarrow \mathbb{Z}$$

such that $\langle c \circ x, y \rangle = \pm \langle x, c \circ y \rangle$. Such a pairing relates $H_*(M)$ to its own dual and thereby places restrictions on its structre. In contrast the 'Poincaré duality' pairing for Floer homology groups involves two different groups $HF_*(\phi)$ and $HF_*(\phi^{-1})$. Hence its existence does not have any consequences for the structure of $HF_*(\phi)$ for general ϕ. An exception to this occurs if ϕ^2 is symplectically isotopic to the identity, because then $HF_*(\phi)$ and $HF_*(\phi^{-1})$ are isomorphic. In that case 'Poincaré duality' yields the following information on $HF_*(\phi)$:

Definition 3.3. Let P be a $\mathbb{Z}/2$-graded unital module over $QH_*(M,\omega)$,. We say that *the $QH_*(M,\omega)$-action on P is self-dual* if there is a graded nondegenerate bilinear form

$$\langle \cdot, \cdot \rangle : P \otimes P \longrightarrow \Lambda$$

such that $\langle cx, y \rangle = \langle x, cy \rangle$ for all $c \in QH_*(M,\omega)$ and $x, y \in P$.

Lemma 3.4. *If ϕ^2 is symplectically isotopic to the identity, the $QH_*(M,\omega)$-action on $HF_*(\phi)$ is self-dual.* □

The main result of this work is

Theorem 3.5. *For any Lagrangian sphere $V \subset M$ there is a $\mathbb{Z}/2$-graded isomorphism*

$$HF_*(\tau_V) \cong QH_*(M,\omega)/I_v$$

which carries $\hat{*}$ to the $QH_*(M,\omega)$-module structure on $QH_*(M,\omega)/I_v$ induced by the quantum product $*$.

The proof of this Theorem is contained in Part III. Together with Lemma 2.6 and Lemma 3.4 it leads to the following consequence:

Corollary 3.6. *If M contains a Lagrangian two-sphere V such that the $QH_*(M,\omega)$-action on $QH_*(M,\omega)/I_v$ is not self-dual, it has an essential and fragile symplectic automorphism, namely τ_V^2.* □

4 Vanishing cycles

We now discuss the two conditions on (M,ω) which appear in Corollary 3.6, beginning with the existence of a Lagrangian sphere. In a sense, most symplectic four-manifolds do not contain Lagrangian spheres: by a generic perturbation of $[\omega] \in H^2(M;\mathbb{R})$, one can achieve that the hyperplane

$$\omega^\perp = \{c \mid \omega(c) = 0\} \subset H_2(M;\mathbb{R})$$

intersects the integer lattice $H_2(M;\mathbb{Z})$ only at the origin. In other examples (such as when $\pi_2(M) = 0$) there are no Lagrangian spheres even though $[\omega]$ is integral. Once the obvious topological restrictions have been exhausted, the existence and position of Lagrangian spheres remains a rather subtle invariant. We will now explain how a degeneration of a Kähler manifold (of arbitrary dimension) can be used to produce a Lagrangian sphere in it. This sphere is called a *vanishing cycle*.

Definition 4.1. Let X be a compact Kähler manifold. An ordinary degeneration of X is a Kähler manifold E with a proper holomorphic map $\pi : E \longrightarrow D$ to the open unit disc $D \subset \mathbb{C}$ with the following properties:

(1) π has at least one critical point.

(2) Any critical point x of π lies in $\pi^{-1}(0)$, and the second derivative of π at x is a nondegenerate quadratic form.

(3) There is a $z \in D \setminus \{0\}$ such that $E_z = \pi^{-1}(z)$ is isomorphic to X.

Proposition 4.2. *Any compact Kähler manifold which admits an ordinary degeneration contains a Lagrangian sphere.*

Because of its importance, we give two proofs of this result (for the reader's benefit we point out that the second proof is much shorter).

First proof. We begin with a linear analogue:

Let β be a (complex) bilinear form on $\mathbb{C}^n$ which is symmetric and nondegenerate. There is an n-dimensional real subspace $L \subset \mathbb{C}^n$ which is Lagrangian for the standard symplectic form ω_0, and such that $\beta|L \times L$ is real and positive definite.

Let $b = \mathrm{re}(\beta)$ be the real part of β, and B the $\mathbb{R}$-linear map on $\mathbb{C}^n = \mathbb{R}^{2n}$ such that $b(v,w) = \langle Bv, w\rangle_{\mathbb{R}}$. Since b is nondegenerate and symmetric, B is invertible and self-adjoint. Let $\mathbb{R}^{2n} = L^+ \oplus L^-$ be the orthogonal splitting into the $\mathbb{R}$-linear subspaces generated by the positive and negative eigenvectors of B. B is $\mathbb{C}$-antilinear: $\langle B(iv), w\rangle_{\mathbb{R}} = b(iv,w) = b(v,iw) = \langle B(v), iw\rangle_{\mathbb{R}} = \langle -i\,Bv, w\rangle_{\mathbb{R}}$. Hence $L^- = iL^+$. The fact that L^+ is orthogonal to iL^+ means that it is a Lagrangian subspace. It also implies that $b(v,iw) = \langle Bv, iw\rangle_{\mathbb{R}} = 0$ for all $v,w \in L^+$. Hence the imaginary part of β vanishes on L^+. By construction, the real part of β is positive definite on L^+. Hence $L = L^+$ has all desired properties. We note the following consequence:

Let q be the quadratic form associated to β. For all $t > 0$, $V_t = q^{-1}(t) \cap L$ is a Lagrangian $(n-1)$-sphere in the symplectic manifold $(q^{-1}(t), \omega_0|q^{-1}(t))$.

The other preliminary is a fact from local Kähler geometry:

Let ω be a Kähler form in some ball $B_\epsilon \subset \mathbb{C}^n$ around 0, which agrees with the standard form ω_0 at 0. There is a Kähler form ω' on B_ϵ such that $\omega' = \omega$ on $B_\epsilon \setminus B_{\epsilon/2}$ and $\omega' = \omega_0$ in some neighbourhood of 0.

This is proved as follows: according to [31, p. 72, Corollaire 2], there is a $\delta \in (0;\epsilon)$ such that $\omega_0 - \omega|B_\delta = i\partial\overline{\partial}f$ for some $f \in C^\infty(B_\delta, \mathbb{R})$. Since ω agrees with ω_0 at 0, we may assume that $f(0) = f'(0) = (D^2 f)_0 = 0$ (take any f and subtract the first three terms in its Taylor expansion around 0). Hence

$$|f(x)| \leq C|x|^3 \quad \text{and} \quad |df_x| \leq C|x|^2 \tag{4.1}$$

for some constant C. Choose a $\psi \in C^\infty(\mathbb{C}^n, \mathbb{R})$ with $\psi|B_1 = 1$ and $\psi|\mathbb{C}^n \setminus B_2 = 0$, and set $\psi_r(x) = \psi(x/r)$. For any $r < \frac{\delta}{2}$, $\eta_r = i\partial\overline{\partial}(\psi_r f)$ is a real $(1,1)$-form on $\mathbb{C}^n$. We can estimate

$$|(\eta_r)_x| \leq C'(|(\omega_0)_x - \omega_x| + r^{-1}|df_x| + r^{-2}|f(x)|)$$

where C' is independent of r. Using (4.1) we conclude that

$$|(\eta_r)_x| \leq C'' r$$

for all $x \in B_{2r}$. On the other hand $(\eta_r)_x = 0$ for all $x \notin B_{2r}$. It follows that the forms η_r converge to 0 uniformly as $r \to 0$. In particular, if we choose r sufficiently small, $\omega' = \omega + \eta_r$ is a Kähler form. It is easy to see that ω' has the desired properties.

Let $\pi : E \longrightarrow D$ be an ordinary degeneration with Kähler form Ω. Choose a critical point x_0 of π. The holomorphic Morse Lemma [2, Lemma 2] shows that there is a holomorphic chart

$$\mathbb{C}^n \supset B_\epsilon \xrightarrow{c} E$$

around x_0 such that $q(z) = \pi(c(z))$ is a nondegenerate quadratic form. After a linear change of c, we can assume that $c^*\Omega$ agrees with ω_0 at 0. Using the local technique explained above, replace Ω by a Kähler form Ω' such that $c^*\Omega' = \omega_0$ near 0. This brings us back to the case discussed above: it follows that for small $t > 0$, $(E_t, \Omega'|E_t)$ contains a Lagrangian sphere. Now assume that we want to prove that $(E_z, \Omega|E_z)$, for some $z \neq 0$, contains a Lagrangian sphere. Since the modification of the Kähler form is local near x_0, we can carry it out in such a way that $\Omega'|E_z = \Omega|E_z$. It follows from Moser's theorem that all regular fibres of a degeneration are symplectically isomorphic. Applying this to (E, Ω'), we see that $(E_z, \Omega|E_z)$ is isomorphic to $(E_t, \Omega'|E_t)$ for all $t > 0$. This completes the proof. $\square$

Second proof (Donaldson). Let $\pi : E \longrightarrow D$ be an ordinary degeneration with Kähler form Ω and corresponding Riemannian metric g. Let x_0 be a critical point of π. Since the real part of a nondegenerate complex quadratic form is a nondegenerate real quadratic form with signature zero, x_0 is a nondegenerate critical point of $f = \mathrm{re}(\pi)$ whose Morse index is half the (real) dimension of E. Let (ϕ_r) be the negative gradient flow of f with respect to g, and $W^s \subset E$ the stable submanifold of x_0 under this flow.

The (real) Morse Lemma shows that for small $t > 0$, $V_t = W^s \cap f^{-1}(t)$ is an embedded sphere. A straightforward computation shows that the negative gradient flow of f is equal to the Hamiltonian flow induced by $h = \mathrm{im}(\pi)$; in particular, ϕ_r preserves h. It follows that $W^s \subset h^{-1}(0)$, hence $V_t \subset h^{-1}(0) \cap f^{-1}(t) = E_t$.

We will now prove that W^s is a Lagrangian submanifold of (E, Ω). The proof is based on the fact that the flow (ϕ_r) 'compresses' W^s. More precisely, consider the derivative

$$D\phi_r(x)|T_xW^s : T_xW^s \longrightarrow T_{\phi_r(x)}W^s$$

at some point $x \in W^s$. The fact which we use is that its norm $|D\phi_r(x)|$ with respect to g goes to 0 as $r \to \infty$. Since Ω is clearly bounded with respect to g, this implies that $(\phi_r^*\Omega)(X, Y) \to 0$ for all $X, Y \in T_xW^s$. Because $\phi_r^*\Omega = \Omega$, it follows that $\Omega(X, Y) = 0$.

Since W^s is Lagrangian, V_t is a Lagrangian sphere in $(E_t, \Omega|E_t)$ for small $t > 0$. Moser's theorem shows that all regular fibres of a degeneration are symplectically isomorphic; hence any fibre contains a Lagrangian sphere. $\square$

Proposition 4.3. *Let $X \subset \mathbb{C}\mathrm{P}^n$ be a smooth projective variety and $H \subset \mathbb{C}\mathrm{P}^n$ a smooth hypersurface of degree $d \geq 2$ which intersects X transversely. Then $X_0 = X \cap H$ admits an ordinary degeneration.*

Combining this with Proposition 4.2 one obtains

Corollary 4.4. *Any smooth complete intersection in $\mathbb{C}\mathrm{P}^n$ which is non-trivial (that is, not an intersection of hyperplanes) contains a Lagrangian sphere.* □

Proof of Proposition 4.3. The idea is to construct the degeneration of X_0 as a family of hypersurface sections of X. Let $(H_\lambda)_{\lambda \in \mathbb{C}\mathrm{P}^1}$ be a pencil of hypersurfaces of degree d, with $H_0 = H$. Consider

$$X_{(H_\lambda)} = \{(x, \lambda) \in X \times \mathbb{C}\mathrm{P}^1 \mid x \in H_\lambda\}.$$

We will call the projection $X_{(H_\lambda)} \longrightarrow \mathbb{C}\mathrm{P}^1$ the fibration induced by (H_λ), and denote it by $\pi_{(H_\lambda)}$. Recall that (H_λ) is called a *Lefschetz pencil* on X if $X_{(H_\lambda)}$ is smooth and all critical points of $\pi_{(H_\lambda)}$ are nondegenerate. It is a well-known fact that a generic pencil containing H_0 is a Lefschetz pencil. From a Lefschetz pencil for which $\pi_{(H_\lambda)}$ has at least one critical point one can obtain an ordinary degeneration of X_0, simply by restricting the fibration to a suitable subset of $\mathbb{C}\mathrm{P}^1$. It remains to show that such a Lefschetz pencil exists.

Consider the set of pencils (H_λ) containing H, which have the following property:

(N) There is a point $(x, \lambda) \in X_{(H_\lambda)}$ which is smooth and a nondegenerate critical point of $\pi_{(H_\lambda)}$.

This is an open set because nondegenerate critical points persist under perturbations. Choose a point $x \in X \setminus H$; we can assume that $x = (1 : 0, \dots, 0)$, and that the rational functions $\frac{z_i}{z_0}$ ($i = 1, \dots r$, where $r = \dim X$) are coordinates on a neighbourhood of x in X. It is not difficult to see that the pencil generated by H and

$$z_0^{d-2} \sum_{i=1}^{r} z_i^2 = 0$$

has property (N). By a small perturbation, we obtain a Lefschetz pencil with the same property, hence one whose fibration has a critical point. □

5 The irrational case

The second condition in Corollary 3.6 depends only on the Gromov-Witten invariants which define the quantum product and the mod 2 homology class of the Lagrangian sphere. These Gromov-Witten invariants carry rather less information than one might expect; in fact, they vanish for a large class of symplectic four-manifolds. A simple example of this phenomenon is

Lemma 5.1. *Assume that the first Chern class of* (M, ω) *satisfies* $c_1 = \lambda[\omega]$ *for some* $\lambda \leq 0$. *Then* $*_\epsilon = 0$ *for all* $\epsilon > 0$.

Proof. Let J be an ω-tame almost complex structure. For $A \in H_2(M;\mathbb{Z})$, we denote by $\mathcal{M}^s(A,J)$ the moduli space of simple J-holomorphic maps $w : \mathbb{C}\mathrm{P}^1 \longrightarrow M$ representing A. The group $PSL(2,\mathbb{C})$ of holomorphic automorphisms of $\mathbb{C}\mathrm{P}^1$ acts freely on $\mathcal{M}^s(A,J)$. According to the transversality theorem for pseudo-holomorphic curves [20, Theorem 3.1.2] there is an ω-tame almost complex structure J_0 which is *regular*, that is, such that for any A the quotient $\mathcal{M}^s(A,J_0)/PSL(2,\mathbb{C})$ is a manifold of dimension $2c_1(A) - 2$ (here, as later on, we write $c_1(A)$ for $\langle c_1(TM,\omega), A\rangle$). By assumption, any A with $c_1(A) > 0$ satisfies $\omega(A) \leq 0$, hence $\mathcal{M}^s(A,J_0) = \emptyset$. On the other hand, $\mathcal{M}^s(A,J_0)/PSL(2,\mathbb{C}) = \emptyset$ for all A with $c_1(A) \leq 0$ because its dimension is negative. Since any nonconstant pseudo-holomorphic sphere is a multiple cover of a simple one, it follows that there are no J_0-holomorphic spheres except for the constant ones. This implies that $*_\epsilon = 0$, by definition. □

The vanishing of $*_\epsilon$ for $\epsilon > 0$ means that

$$c * c' = c \circ_\Lambda c'$$

for all $c, c' \in QH_*(M,\omega)$, where $\circ_\Lambda$ is the intersection product with Λ-coefficients. In this case we say that the quantum product is undeformed. More specifically, we say that the quantum product with a class $X \in H_*(M;\mathbb{Z}/2)$ is undeformed if

$$X\,t^0 * c = X\,t^0 \circ_\Lambda c$$

for all c. For instance, the quantum product with $[M]$ is always undeformed.

As we have seen, there are situations in which the quantum product is undeformed because of the interplay between the nonnegativity of the energy of a J-holomorphic sphere and the dimension formula for moduli spaces of such spheres. Arguments of this kind can be applied to symplectic manifolds of any dimension. For symplectic four-manifolds there is a much deeper theory of pseudo-holomorphic curves, including Taubes' results on the connection between such curves and Seiberg-Witten invariants. McDuff [17] pointed out that this theory has the following consequence:

Theorem 5.2. *Let (M, ω) be a compact symplectic four-manifold which is not rational or the blowup of a ruled symplectic manifold. Let $E_1, \dots, E_r$ be a maximal family of disjoint embedded symplectic spheres in M with self-intersection (-1). Then the quantum product with any class in*

$$\operatorname{im}(H_*(M \setminus (E_1 \cup \cdots \cup E_r); \mathbb{Z}/2) \longrightarrow H_*(M; \mathbb{Z}/2)) \tag{5.1}$$

is undeformed.

The proof of this theorem combines results of McDuff [19], Taubes [29] [30] and Liu [16]. What the proof shows is that there is an ω-tame almost complex structure J on M such that the exceptional curves E_j and their multiple covers are the only non-constant J-holomorphic spheres. This implies the result stated above.

Now we return to our discussion of Corollary 3.6. Information about when the quantum product is undeformed is relevant for the following reason:

Lemma 5.3. *Let (M, ω) be a compact symplectic four-manifold which contains a Lagrangian sphere V. Assume that there are $X, Y \in H_2(M; \mathbb{Z}/2)$ with the following properties:*

(1) The quantum product with X and Y is undeformed;

(2) $X \cdot [V] = 1$;

(3) X, Y and $[V]$ are linearly independent elements of $H_2(M; \mathbb{Z}/2)$.

Then the $QH_*(M, \omega)$-action on $QH_*(M, \omega)/I_v$ is not self-dual.

Proof. We will use the same notation for an element of $QH_*(M, \omega)$ and its image in $QH_*(M, \omega)/I_v$ throughout the proof. This should not cause any confusion.

Let $v = [V]t^0$, $x = Xt^0$ and $y = Yt^0$. By Lemma 3.2, $I_v = \Lambda v \oplus \Lambda(v * x)$. Since the quantum product with X is undeformed, $v * x = [\mathrm{pt}]t^0$ and hence

$$I_v = \Lambda v \oplus H_0(M; \Lambda). \tag{5.2}$$

Let $\langle \cdot, \cdot \rangle$ be a graded Λ-bilinear form on $QH_*(M, \omega)/I_v$ which satisfies (3.3). We have to prove that this form is degenerate. Let $u = [M]t^0$ be the unit element of $*$. There is a nontrivial linear combination c (with coefficients in Λ) of x and y such that

$$\langle u, c \rangle = 0.$$

Equation (5.2) and the fact that V, X and Y are linearly independent imply that $c \notin I_v$. To prove that $\langle \cdot, \cdot \rangle$ is degenerate it is sufficient to show that $\langle \cdot, c \rangle$ vanishes. Using the property (3.3) one sees that

$$\langle b, c \rangle = \langle b * u, c \rangle = \langle u, b * c \rangle$$

for all $b \in QH_*(M,\omega)$. Since the quantum product with c is undeformed, $b * c \in \Lambda c \oplus H_1(M;\Lambda) \oplus H_0(M;\Lambda)$ for all b. By construction, $\langle u, c\rangle = 0$. $\langle u, H_1(M;\Lambda)\rangle$ vanishes because of the grading, and $\langle u, H_0(M;\Lambda)\rangle = 0$ because $H_0(M;\Lambda)$ goes to zero in $QH_*(M,\omega)/I_v$. This proves that $\langle u, b{*}c\rangle = 0$ for all b. □

Theorem 5.4. *Let (M,ω) be a compact symplectic four-manifold, with $b_1(M) = 0$, which contains a Lagrangian sphere V. Assume that (M,ω) is irrational, with minimal model*

$$p : M \longrightarrow \overline{M}.$$

Moreover, assume that $\dim H_2(\overline{M};\mathbb{Z}/2) \geq 3$ and that $[V] \in H_2(M;\mathbb{Z}/2)$ does not lie in the kernel of p_*. Then (M,ω) admits an essential and fragile symplectic automorphism.

Proof. Let $E_1,\dots,E_r$ be the family of exceptional curves contracted by p. The two-dimensional homology of M splits into $H_2(M;\mathbb{Z}/2) = \ker(p_*) \oplus R$, where

$$\begin{aligned} R &= \mathrm{im}(H_2(M \setminus (E_1 \cup \dots \cup E_r);\mathbb{Z}/2) \longrightarrow H_2(M;\mathbb{Z}/2)) \\ &\cong H_2(\overline{M};\mathbb{Z}/2). \end{aligned}$$

The two parts are orthogonal with respect to the intersection form. Since $[V] \notin \ker(p_*)$, there is an $X \in R$ such that $[V]\cdot X = 1$. $X \neq [V]$ because V has self-intersection (-2). Because $\dim R \geq 3$ we can find a third element $Y \in R$ such that $[V]$, X and Y are linearly independent. By Theorem 5.2, the quantum product with X and Y is undeformed. Lemma 5.3 implies that the $QH_*(M,\omega)$-action on $QH_*(M,\omega)/I_v$ is not self-dual, and Corollary 3.6 completes the proof. □

The Theorem 1.3 stated in the Introduction is the special case when (M,ω) is minimal.

6 Three rational algebraic surfaces

In this section we consider the quadric and cubic hypersurfaces $M_2, M_3 \subset \mathbb{C}\mathrm{P}^3$ and the intersection of two quadrics $M_{2,2} \subset \mathbb{C}\mathrm{P}^4$. We denote the induced symplectic forms by ω_2, ω_3 and $\omega_{2,2}$ respectively, and normalize them in such a way that their cohomology class equals the first Chern class. Corollary 4.4 shows that these three manifolds contain Lagrangian spheres. In the case of M_2, we can describe such a sphere explicitly: M_2 is symplectically isomorphic to $\mathbb{C}\mathrm{P}^1 \times \mathbb{C}\mathrm{P}^1$, and the 'antidiagonal'

$$\overline{\Delta} = \{(x,y) \in \mathbb{C}\mathrm{P}^1 \times \mathbb{C}\mathrm{P}^1 \mid x_0\overline{x_1} + y_0\overline{y_1} = 0\}$$

is a Lagrangian sphere.

The quadric is interesting for the following reason: the results of the previous section show that on a large class of symplectic four-manifolds, the square of any generalized Dehn twist is essential. One could ask whether this is always true, and M_2 provides a counterexample. In fact, Gromov's theorem, which we have quoted in the introduction, shows that $\pi_0(\mathrm{Aut}(M_2, \omega_2))$ has order two; hence the square of any generalized Dehn twist is symplectically isotopic to the identity (for $\tau_{\overline{\Delta}}$, this can be proved by an explicit isotopy). Running Corollary 3.6 backwards, it follows that for any Lagrangian sphere $V \subset M_2$, the action of $QH_*(M_2, \omega_2)$ on $QH_*(M_2, \omega_2)/I_v$ is self-dual. It is instructive to verify this directly, and we will do so now.

We will identify $M_2 = \mathbb{C}\mathrm{P}^1 \times \mathbb{C}\mathrm{P}^1$. Let

$$c_4 = [M]t^0, \quad c_2 = [\mathbb{C}\mathrm{P}^1 \times \mathrm{pt}]t^0, \quad c_2' = [\mathrm{pt} \times \mathbb{C}\mathrm{P}^1]t^0, \quad c_0 = [\mathrm{pt}]t^0$$

be the standard basis of $QH_*(M_2, \omega_2)$. The quantum product is determined by the relations

$$c_2 * c_2 = c_2' * c_2' = c_4\, t^2, \quad c_2 * c_2' = c_0.$$

Up to sign, $[\mathbb{C}\mathrm{P}^1 \times \mathrm{pt}] - [\mathrm{pt} \times \mathbb{C}\mathrm{P}^1]$ is the only class in $H_2(M; \mathbb{Z})$ with self-intersection (-2). It follows that for any Lagrangian sphere V,

$$I_v = \Lambda(c_2 - c_2') \oplus \Lambda(c_4\, t^2 - c_0).$$

Define a symmetric bilinear form on $QH_*(M_2, \omega_2)$ by

$$\langle b, c \rangle = (c_2 + c_2') \cdot_\Lambda (b * c) \in \Lambda,$$

where $\cdot_\Lambda$ denotes the intersection number with Λ-coefficients. An easy computation shows that $\langle \cdot, I_v \rangle = 0$. Therefore $\langle \cdot, \cdot \rangle$ induces a bilinear form on $QH_*(M_2, \omega_2)/I_v$. By definition, this form satisfies (3.3). The classes of c_2 and c_4 form a basis of $QH_*(M_2, \omega_2)/I_v$; with respect to this basis, $\langle \cdot, \cdot \rangle$ is given by the invertible matrix $\left(\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right)$. This shows that the $QH_*(M_2, \omega_2)$-action on $QH_*(M_2, \omega_2)/I_v$ is self-dual.

The two other examples $M_3, M_{2,2}$ are del Pezzo surfaces: they are isomorphic to $\mathbb{C}\mathrm{P}^2$ blown up at six resp. five points in 'general position'. We will treat M_3 and $M_{2,2}$ simultaneously and refer to either manifold as (M, ω).

We denote the mod 2 Poincaré dual of $c_1(M)$ by $K \in H_2(M; \mathbb{Z}/2)$. Let $\mathbf{L} \subset H_2(M; \mathbb{Z}/2)$ be the subspace of classes X with $K \cdot X = 0$. Because $K \neq 0$ (for instance, $K \cdot E = 1$ for any exceptional divisor E), $\mathbf{L}$ is a subspace of codimention one. The corresponding subspace $\mathbf{L} \otimes \Lambda \subset QH_*(M, \omega)$ will be denoted by $\mathbf{L}_\Lambda$.

Let $V \subset M$ be a Lagrangian sphere and $v = [V]\, t^0 \in QH_*(M, \omega)$.

Lemma 6.1. *$I_v \cap \mathbf{L}_\Lambda$ is one-dimensional (over Λ).*

Proof. $I_v \cap \mathbf{L}_\Lambda$ is not zero since the nontrivial class v lies in it. On the other hand, $I_v \not\subset \mathbf{L}_\Lambda$ because for any $W \in H_2(M;\mathbb{Z}/2)$ with $W \cdot [V] = 1$ we have

$$W\,t^0 * v = [\mathrm{pt}]\,t^0 + (\text{higher order terms})$$

and this does not lie in $\mathbf{L}_\Lambda$. This completes the proof because by Lemma 3.2 I_v is two-dimensional. □

Proposition 6.2. $\mathbf{L}_\Lambda * \mathbf{L}_\Lambda \subset I_v$.

We postpone the proof to an Appendix at the end of this section.

Theorem 6.3. *For any Lagrangian sphere $V \subset M$, τ_V^2 is an essential and fragile symplectic automorphism.*

Proof. Assume that $\langle \cdot, \cdot \rangle$ is a nondegenerate Λ-bilinear form on $QH_*(M,\omega)/I_v$ which satisfies (3.3). As we saw in the proof of Lemma 5.3, this implies that

$$\langle b, c \rangle = \langle [M]\,t^0, b * c \rangle. \tag{6.1}$$

Consider $\overline{\mathbf{L}_\Lambda} = \mathbf{L}_\Lambda/(\mathbf{L}_\Lambda \cap I_v) \subset QH_*(M,\omega)/I_v$. Proposition 6.2 and equation (6.1) imply that $\overline{\mathbf{L}_\Lambda}$ is an isotropic subspace for $\langle \cdot, \cdot \rangle$. Because of the nondegeneracy of the bilinear form, an isotropic subspace must satisfy

$$\dim_\Lambda \overline{\mathbf{L}_\Lambda} \leq \tfrac{1}{2} \dim_\Lambda QH_*(M,\omega)/I_v. \tag{6.2}$$

By Lemma 3.2, the r.h.s. of this inequality is

$$\tfrac{1}{2} \dim_\Lambda QH_*(M,\omega)/I_v = \tfrac{1}{2}(\dim H_*(M;\mathbb{Z}/2) - 2) = \begin{cases} \frac{7}{2} & M = M_3, \\ 3 & M = M_{2,2} \end{cases} \tag{6.3}$$

As mentioned above, $\dim_\Lambda \mathbf{L}_\Lambda = \dim H_2(M;\mathbb{Z}/2) - 1$. Lemma 6.1 says that $\dim_\Lambda \overline{\mathbf{L}_\Lambda} = \dim_\Lambda \mathbf{L}_\Lambda - 1$. Hence

$$\dim_\Lambda \overline{\mathbf{L}_\Lambda} = \begin{cases} 5 & M = M_3, \\ 4 & M = M_{2,2}. \end{cases} \tag{6.4}$$

By comparing (6.3) with (6.4) one sees that the inequality (6.2) is violated for both M_3 and $M_{2,2}$. This shows that a bilinear form $\langle \cdot, \cdot \rangle$ with the properties stated above cannot exist, hence that the $QH_*(M,\omega)$-action on $QH_*(M,\omega)/I_v$ is not self-dual. Applying Corollary 3.6 completes the proof. □

We can now prove the Theorem 1.4 stated in the Introduction.

Proof of Theorem 1.4. Let $M \subset \mathbb{C}\mathrm{P}^r$ be a complete intersection of type $d = (d_1 \dots d_{r-2})$. We assume that $r \geq 3$, $d_i \geq 2$ and $d \neq (2)$. This excludes the trivial intersection $\mathbb{C}\mathrm{P}^2$ and the quadric $M_2 = \mathbb{C}\mathrm{P}^1 \times \mathbb{C}\mathrm{P}^1$. Corollary 4.4 shows that M contains a Lagrangian sphere.

By the Lefschetz hyperplane theorem, M is simply connected; hence $b_1(M) = 0$ and $\dim H_2(M; \mathbb{Z}/2) = \chi(M) - 2$. The Euler characteristic of complete intersections is well-known:

$$\chi(M) = \frac{1}{2}\left(\prod_i d_i\right)\left[\left(\sum_i d_i - (r+1)\right)^2 - (r+1) + \sum_i d_i^2\right].$$

For $r \geq 4$, we have $-(r+1) + \sum_i d_i^2 \geq -(r+1) - 4(r-2) \geq 3$ and therefore $\chi(M) \geq \frac{1}{2}2^{r-2}3 \geq 6$. For $r = 3$, $\chi(M) = d_1[(d_1-2)^2+2] \geq 9$ because $d_1 \geq 3$ by assumption. It follows that $\dim H_2(M; \mathbb{Z}/2) \geq 3$.

The first Chern class of M is $c_1 = \lambda[\omega]$ with $\lambda = (r+1) - \sum_i d_i$. The only cases with $\lambda > 0$ are $d = (3), (2,2)$; these are the two del Pezzo surfaces which we have studied before, see Theorem 6.3. In all other cases (M, ω) is a minimal irrational surface and the result follows from Theorem 1.3. Note that it is not really necessary to appeal to Theorem 5.2 to prove that the quantum product is undeformed; Lemma 5.1 is sufficient. □

Appendix: Proof of Proposition 6.2

The quantum product of M has been computed in [5] and [6]. Proposition 6.2 could be derived from this computation, but prefer a slightly less direct route which uses only partial information about $QH_*(M, \omega)$.

Since $[\omega] = c_1(M)$, the only nontrivial coefficients in the quantum product on $H_2(M; \mathbb{Z}/2)$ are the intersection form

$$*_0 : H_2(M; \mathbb{Z}/2) \otimes H_2(M; \mathbb{Z}/2) \longrightarrow H_0(M; \mathbb{Z}/2) \cong \mathbb{Z}/2$$

and its first two 'quantum corrections'

$$*_1 : H_2(M; \mathbb{Z}/2) \otimes H_2(M; \mathbb{Z}/2) \longrightarrow H_2(M; \mathbb{Z}/2),$$
$$*_2 : H_2(M; \mathbb{Z}/2) \otimes H_2(M; \mathbb{Z}/2) \longrightarrow H_4(M; \mathbb{Z}/2) \cong \mathbb{Z}/2.$$

The first quantum correction comes from the lines on M (there are 27 on M_3 and 16 on $M_{2,2}$). It is not difficult to compute that

$$X *_1 Y = \begin{cases} 0 & M = M_3, \\ K(X \cdot Y) & M = M_{2,2} \end{cases} \tag{6.5}$$

for all $X, Y \in \mathbf{L}$.

The other property of $*$ which we will use is that it is invariant under the action of the Weyl group W. Recall that the Weyl group is the (finite) group of automorphisms of $H_2(M;\mathbb{Z})$ which preserve the intersection form and the Poincaré dual of c_1. The invariance of $*$ under W can be read off from the computations in the papers quoted above. (In order to understand this geometrically one would have to prove that any element of W is induced by a symplectic automorphism of (M,ω); this can probably be derived from known results about Del Pezzo surfaces.) It is not difficult to see that $\mathbf{L}$ is an irreducible W-module; therefore there is only one invariant quadratic form on it, up to multiples. We conclude that $*_2$ is either zero or equal to the intersection form. Together with (6.5) this shows that there is an $a \in QH_*(M,\omega)$ such that

$$x * y = (x \cdot_\Lambda y)a \tag{6.6}$$

for all $x, y \in \mathbf{L}_\Lambda$.

There is a $W \in \mathbf{L}$ such that $W \cdot [V] = 1$. This can be proved as follows: since the intersection form on $H_2(M;\mathbb{Z}/2)$ is unimodular and $\mathbf{L} = K^\perp$, $W \cdot [V] = 0$ for all $W \in \mathbf{L}$ would imply that $[V] = K$. Then the self-intersection number of V would be congruent to $c_1(M)^2$ mod 4. But $V \cdot V = -2$ whereas $c_1(M_3)^2 = 3$ and $c_1(M_{2,2})^2 = 4$.

According to (6.6),

$$v \cdot (W\,t^0) = ([V] \cdot W)a = a;$$

and this shows that $a \in I_v$. From this and (6.6) it follows that $\mathbf{L}_\Lambda * \mathbf{L}_\Lambda \subset I_v$.

Part II

7 Floer homology as a functor

In this part we get to grips with the Floer homology groups $HF_*(\phi)$. Our aim is to present some known properties of Floer homology in a certain perspective and then to introduce a new extension of its structure. The emphasis throughout will be on the functorial nature of Floer homology. This aspect is not usually considered to be of much interest: homomorphisms between Floer homology groups appear in one step of the construction of Floer homology, but they all turn out to be isomorphisms. We choose to make this part of the structure more explicit. This leads to a picture of Floer homology as a functor on a category whose objects are symplectic fibre bundles over S^1 and whose morphisms are such fibre bundles over a cylinder. This is part of a larger picture in which Floer homology appears as a topological quantum field theory for symplectic fibre bundles over Riemann surfaces, but the cylinder alone is sufficient for our purpose.

Up to that point, our description to Floer homology contains only known results. The next step, however, is new: we extend the set of morphisms from symplectic fibre bundles to a larger class of fibrations which may have singular fibres of a simple kind. This introduces new induced maps between Floer homology groups; one of these maps will be used in Part III to compute the Floer homology of generalized Dehn twists.

This first section serves as an introduction. We begin by defining the 'symplectic fibre bundles' which have been mentioned above.

Definition 7.1. Let B be a smooth manifold. A symplectic fibre bundle over B is a smooth proper submersion $\pi : E \longrightarrow B$ together with a closed two-form $\Omega \in \Omega^2(E)$ whose restriction to any fibre $E_z = \pi^{-1}(z)$ is nondegenerate.

Ω determines a connection on the fibre bundle $E \longrightarrow B$, that is, a 'horizontal' subbundle $TE^h \subset TE$ which is complementary to the 'vertical' subbundle $TE^v = \ker(D\pi)$. It is defined by

$$TE^h_x = (TE^v_x)^\perp = \{X \in TE_x \mid \Omega(X,Y) = 0 \text{ for all } Y \in TE^v_x\}.$$

We denote the horizontal lift of $Z \in TB$ by $Z^\natural$.

Lemma 7.2. *The parallel transport $P_\gamma : E_{\gamma(0)} \longrightarrow E_{\gamma(1)}$ along any path $\gamma : [0;1] \longrightarrow B$ satisfies $P_\gamma^*(\Omega|E_{\gamma(1)}) = \Omega|E_{\gamma(0)}$.*

Proof. Consider the family of parallel transports $P_t = P_{\gamma|[0;t]} : E_{\gamma(0)} \longrightarrow E_{\gamma(t)}$. Since Ω is closed,

$$\frac{d}{dt}\left(P_t^*\Omega\right) = P_t^*(L_{(\partial_t\gamma)^\natural}\Omega) = d\left[P_t^*(i_{(\partial_t\gamma)^\natural}\Omega)\right].$$

But $i_{(\partial_t\gamma)^\natural}\Omega|E_{\gamma(t)}$ vanishes by definition. Hence $P_t^*\Omega = \Omega$ for all t. □

Remark 7.3. We have proved that $(E, \Omega|TE^v)$ is a locally trivial family of symplectic manifolds, that is, a fibre bundle whose structure group is the symplectic automorphism group of $(E_z, \Omega|E_z)$. However Ω contains more information than $\Omega|TE^v$ (for instance, it determines the connection TE^h). Therefore the name 'symplectic fibre bundle' is not entirely appropriate; a more accurate one might be 'symplectic fibre bundle with a Hamiltonian connection'.

From now on, it will be assumed that all fibre bundles which occur have four-dimensional fibres. This dimensional condition does not have any fundamental importance; it arises from the technical details of our definition of Floer homology, and could problably be removed by using more sophisticated techniques.

In its formulation in terms of symplectic fibre bundles, Floer homology theory assigns to every symplectic fibre bundle (T, Θ) over S^1 a group $HF_*(T, \Theta)$ which, as before, is a $\mathbb{Z}/2$-graded Λ-vector space. Moreover, every symplectic fibre bundle (E, Ω) over a cylinder $Z = [s_0; s_1] \times S^1$ determines a graded homomorphism

$$\Phi(E, \Omega) : HF_*(E_{s_0 \times S^1}, \Omega_{s_0 \times S^1}) \longrightarrow HF_*(E_{s_1 \times S^1}, \Omega_{s_1 \times S^1}),$$

where $(E_{s_i \times S^1}, \Omega_{s_i \times S^1})$ is the restriction of (E, Ω) to $\{s_i\} \times S^1 \subset Z$. These objects have the following properties:

(Duality) Let $\iota : S^1 \longrightarrow S^1$ be the orientation-reversing involution given by $\iota(t) = -t$ for $t \in S^1 = \mathbb{R}/\mathbb{Z}$ (from now on we always identify S^1 with $\mathbb{R}/\mathbb{Z}$). The pullback of a symplectic fibre bundle (T, Θ) over S^1 by ι will be denoted by $(\overline{T}, \overline{\Theta})$. For any (T, Θ) there is a canonical graded bilinear map

$$\langle \cdot, \cdot \rangle_{(T,\Theta)} : HF_*(\overline{T}, \overline{\Theta}) \otimes HF_*(T, \Theta) \longrightarrow \Lambda$$

which is non-degenerate. These pairings are symmetric in the sense that

$$\langle a, b \rangle_{(T,\Theta)} = \langle b, a \rangle_{(\overline{T},\overline{\Theta})}.$$

Let (E, Ω) be a symplectic fibre bundle over Z and $(\overline{E}, \overline{\Omega})$ its pullback by the involution of Z given by $(s, t) \longmapsto (s_1 + s_0 - s, -t)$. Then

$$\langle a, \Phi(E, \Omega)b \rangle_{(E_{s_1 \times S^1}, \Omega_{s_1 \times S^1})} = \langle \Phi(\overline{E}, \overline{\Omega})a, b \rangle_{(E_{s_0 \times S^1}, \Omega_{s_0 \times S^1})}.$$

(Gluing) We divide $Z = [s_0; s_1] \times S^1$ into two parts $Z_- = [s_0; s] \times S^1$ and $Z_+ = [s; s_1] \times S^1$ for some $s \in (s_0; s_1)$. Let (E, Ω) be a symplectic fibre bundle over Z and (E_-, Ω_-), (E_+, Ω_+) its restrictions to $Z_\pm$. Then

$$\Phi(E, \Omega) = \Phi(E_+, \Omega_+) \circ \Phi(E_-, \Omega_-).$$

Moreover, the homomorphism induced by a product bundle $[s_0; s_1] \times (T, \Theta)$ is

$$\Phi([s_0; s_1] \times (T, \Theta)) = \mathrm{id}_{HF_*(T,\Theta)}. \tag{7.1}$$

(Deformation invariance) Let $\pi : E \longrightarrow Z$ be a smooth proper submersion and $\Omega_0, \Omega_1 \in \Omega^2(E)$ two two-forms which turn it into a symplectic fibre bundle. (E, Ω_0) and (E, Ω_1) are called deformation equivalent if

(1) $\Omega_0|E_z = \Omega_1|E_z$ for all $z \in Z$,

(2) $\Omega_0|E_{s_0 \times S^1} = \Omega_1|E_{s_0 \times S^1}$ and $\Omega_0|E_{s_1 \times S^1} = \Omega_1|E_{s_1 \times S^1}$, and if

(3) the cohomology class $[\Omega_1 - \Omega_0] \in H^2(E, E_{s_0 \times S^1} \cup E_{s_1 \times S^1}; \mathbb{R})$ vanishes.

The reason for this terminology is that $(E, t\Omega_1 + (1-t)\Omega_0)$ is a symplectic fibre bundle for $t \in [0; 1]$. (E, Ω_0) and (E, Ω_1) induce homomorphisms of the same Floer homology groups; the deformation invariance property says that these homomorphisms coincide.

We will now explain how the groups $HF_*(\phi)$ used in Part I are related to this framework[1]. Let T_ϕ be the mapping torus of $\phi \in \mathrm{Aut}(M, \omega)$, that is, the manifold obtained from $\mathbb{R} \times M$ by identifying (t, x) with $(t-1, \phi(x))$. This manifold is canonically fibered over S^1, and since ϕ is symplectic, the pullback of ω to $\mathbb{R} \times M$ induces a closed two-form Θ_ϕ on T_ϕ. (T_ϕ, Θ_ϕ) is a symplectic fibre bundle over S^1 with fibre (M, ω); we define

$$HF_*(\phi) = HF_*(T_\phi, \Theta_\phi).$$

The quantum module structure on $HF_*(\phi)$ can also be defined in terms of invariants of symplectic fibre bundles. This requires a certain generalization of the homomorphisms $\Phi(E, \Omega)$; we postpone this to section 10. We will now consider the two basic properties (isotopy invariance and 'Poincaré duality') of $HF_*(\phi)$ used in section 3. These properties (at least, the part which does not concern the multiplicative structure) can be reduced to the properties of $HF_*(T, \Theta)$ and $\Phi(E, \Omega)$ listed above. One case is simple: $(T_{\phi^{-1}}, \Theta_{\phi^{-1}})$ is naturally isomorphic to $(\overline{T_\phi}, \overline{\Theta_\phi})$, and hence 'Poincaré duality' is a consequence of the 'duality' property of $HF_*(T, \Theta)$. The argument which derives the isotopy invariance of $HF_*(\phi)$ from the 'gluing' and 'deformation invariance' properties is more complicated; we preface it by some remarks on symplectic fibre bundles.

Fix some $\phi \in \mathrm{Aut}(M, \omega)$. A function $H \in C^\infty(T_\phi, \mathbb{R})$ can be used to perturb the form $\Theta_\phi \in \Omega^2(T_\phi)$ in the following way:

$$\Theta_{\phi,H} = \Theta_\phi - d(H\,dt),$$

[1] We remind the reader that (M, ω) denotes a compact symplectic four-manifold with zero first Betti number.

where dt is the pullback of the standard one-form on S^1. By definition, $\Theta_{\phi,H}$ is closed and agrees with Θ_ϕ on each fibre of $T_\phi \longrightarrow S^1$. Hence $(T_\phi, \Theta_{\phi,H})$ is again a symplectic fibre bundle. These bundles can be identified with more familiar objects by considering the Hamiltonian flow $(\phi_t^H)_{t\in\mathbb{R}}$ on M induced by the pullback of H to $\mathbb{R} \times M$. This flow determines a diffeomorphism $p_H : T_{\phi\circ\phi_1^H} \longrightarrow T_\phi$, given by $p_H(t,x) = (t, \phi_t^H(x))$. A straightforward computation shows that $p_H^*\Theta_{\phi,H} = \Theta_{\phi\circ\phi_1^H}$. Hence $(T_\phi, \Theta_{\phi,H})$ is isomorphic to $(T_{\phi\circ\phi_1^H}, \Theta_{\phi\circ\phi_1^H})$. We call $\phi \circ \phi_1^H$ a Hamiltonian perturbation of ϕ.

Let us denote the product fibre bundle $[s_0; s_1] \times (T_\phi, \Theta_\phi) \longrightarrow Z$ by (E_ϕ, Ω_ϕ). As in the case of T_ϕ, a function $K \in C^\infty(E_\phi, \mathbb{R})$ determines a perturbation

$$\Omega_{\phi,K} = \Omega_\phi - d(K\,dt) \in \Omega^2(E_\phi).$$

The boundary values $K_0 = K|\{0\} \times T_\phi$ and $K_1 = K|\{1\} \times T_\phi$ determine the symplectic fibre bundle $(E_\phi, \Omega_{\phi,K})$ up to deformation equivalence. Therefore the homomorphism induced by $(E_\phi, \Omega_{\phi,K})$ depends only on K_0, K_1. We denote it by

$$C(\phi, K_0, K_1) : HF_*(T_\phi, \Theta_{\phi,K_0}) \longrightarrow HF_*(T_\phi, \Theta_{\phi,K_1}).$$

Using the 'gluing' property of $\Phi(E, \Omega)$ and a suitable choice of K, it is not difficult to prove that

$$C(\phi, H, H) = \mathrm{id} \quad \text{and} \quad C(\phi, H, H'') = C(\phi, H', H'') \circ C(\phi, H, H')$$

for all $H, H', H'' \in C^\infty(T_\phi, \mathbb{R})$. Hence all maps $C(\phi, H, H')$ are isomorphisms. In particular $HF_*(T_\phi, \Theta_\phi) \cong HF_*(T_\phi, \Theta_{\phi,H})$ for any H.

We can now prove the isotopy invariance of $HF_*(\phi)$: let ϕ, ϕ' be two automorphisms of (M, ω) which are symplectically isotopic. Since $H^1(M, \mathbb{R}) = 0$ (it is here that this assumption becomes important) ϕ' is a Hamiltonian perturbation of ϕ, say $\phi' = \phi \circ \phi_1^H$. Therefore

$$HF_*(\phi') = HF_*(T_\phi, \Theta_{\phi,H}) \cong HF_*(T_\phi, \Theta_\phi) = HF_*(\phi).$$

Remark 7.4. If ϕ and ϕ' lie in the same component of $\mathrm{Aut}(M, \omega)$, a function $H \in C^\infty(T_\phi, \mathbb{R})$ such that $\phi' = \phi \circ \phi_1^H$ determines an isomorphism of $HF_*(\phi)$ with $HF_*(\phi')$. In general, this isomorphism depends on the choice of H; there is no *canonical* isomorphism between the two Floer homology groups.

Our point of view so far has been to view the groups $HF_*(T, \Theta)$ and the maps $\Phi(E, \Omega)$ as fundamental and to derive the Floer homology groups of symplectic automorphisms and their properties from them. One might wonder whether the groups $HF_*(T, \Theta)$ are in fact more general invariants than $HF_*(\phi)$, or whether the 'duality', 'gluing' and 'deformation invariance' of

$HF_*(T,\Theta)$ and $\Phi(E,\Omega)$ say more about $HF_*(\phi)$ than what we have already derived from them. The answer to the first question is negative because every symplectic fibre bundle over S^1 is isomorphic to some bundle (T_ϕ,Ω_ϕ). As to the second question, there is one more property of $HF_*(\phi)$ which we have not mentioned up to now, its conjugation invariance, which can be derived from the definition $HF_*(\phi) = HF_*(T_\phi,\Theta_\phi)$. Apart from this, the answer is again negative because every fibre bundle over a cylinder is isomorphic to $(E_\phi,\Omega_{\phi,K})$ for some ϕ and K. In particular, all maps $\Phi(E,\Omega)$ are isomorphisms.

Bibliographical note. Because of its application to the Arnol'd conjecture, it is customary to define Floer homology only for automorphisms which are Hamiltonian perturbations of the identity map. For an exposition of the basic construction in this case, see the surveys [18] and [24]. To the author's knowledge, Floer homology for general symplectic automorphisms appears in the literature only in the work of Dostoglou-Salamon [7] [8] on the Atiyah-Floer conjecture. Their definition follows the approach of Floer [9] and works for simply-connected monotone symplectic manifolds (of any dimension). They also suggested that the definition could be generalized to a larger class of symplectic manifolds using the ideas of Hofer-Salamon [14]. This is the method adopted here. As mentioned above, recent progress on the Arnol'd conjecture seems to indicate that Floer homology can be defined for automorphisms of any compact symplectic manifold.

Homomorphisms which are essentially equivalent to $\Phi(E,\Omega)$ were introduced by Floer [9] to prove isotopy invariance. A detailed exposition of his construction can be found in [25, section 6]. The 'topological quantum field theory' picture of Floer homology occurs in [22] and [27].

Almost holomorphic fibrations

The new structure on Floer homology which will be constructed in the next sections involves a generalization of the concept of symplectic fibre bundle in which the fibres are allowed to be singular, with singularities modelled on singular points of holomorphic hypersurfaces.

Definition 7.5. Let (Σ, j) be a Riemann surface (possibly with boundary). An almost holomorphic fibration over Σ consists of

(1) a smooth manifold E and a proper surjective map $\pi : E \longrightarrow \Sigma$ whose critical point set $\mathrm{Crit}(\pi)$ lies in $\pi^{-1}(\mathrm{int}\,\Sigma)$;

(2) a closed two-form $\Omega \in \Omega^2(E)$ whose restriction to $TE^v_x = \ker(D\pi_x)$ is nondegenerate for any $x \in E$. Since $TE^v_x = TE_x$ for $x \in \mathrm{Crit}(\pi)$, such an Ω is symplectic in a neighbourhood of $\mathrm{Crit}(\pi)$.

(3) an *integrable* almost complex structure J', defined in a neighbourhood U of $\mathrm{Crit}(\pi)$ in E, compatible with $\Omega|U$ and with respect to which $\pi|U$ is a holomorphic function (strictly speaking, only the germ of J' at $\mathrm{Crit}(\pi)$ matters).

An ordinary almost holomorphic fibration is one such that the second differential

$$(D^2\pi)_x : TE_x \otimes_{\mathbb{C}} TE_x \longrightarrow T_{\pi(x)}\Sigma$$

at any point $x \in \mathrm{Crit}(\pi)$ is a nondegenerate complex quadratic form. In particular, the critical points of such a fibration are isolated. The notion of an ordinary almost holomorphic fibration bears an obvious resemblance to the 'ordinary degenerations' used in section 4.

Let (E, Ω, J') be an ordinary almost holomorphic fibration over $Z = [s_0; s_1] \times S^1$. Its boundary components $(E_{s_0 \times S^1}, \Omega_{s_1 \times S^1})$ $(i = 0, 1)$ do not contain any singular points, that is, they are symplectic fibre bundles. What we will do is define induced maps

$$\Phi(E, \Omega, J') : HF_*(E_{s_0 \times S^1}, \Omega_{s_0 \times S^1}) \longrightarrow HF_*(E_{s_1 \times S^1}, \Omega_{s_1 \times S^1}).$$

These maps generalize those induced by symplectic fibre bundles and satisfy similar properties:

(Duality) The pullback of (E, Ω, J') by the holomorphic map $Z \longrightarrow Z$, $(s, t) \longmapsto (s_0 + s_1 - s, -t)$ is again an ordinary almost holomorphic fibration; we denote it by $(\overline{E}, \overline{\Omega}, \overline{J}')$. The last part of the 'duality' property extends to almost holomorphic fibrations in a straightforward way:

$$\langle a, \Phi(E, \Omega, J')b\rangle_{(E_{s_1 \times S^1}, \Omega_{s_1 \times S^1})} = \langle \Phi(\overline{E}, \overline{\Omega}, \overline{J}')a, b\rangle_{(E_{s_0 \times S^1}, \Omega_{s_0 \times S^1})}.$$

(Gluing) This property extends to the maps $\Phi(E, \Omega, J')$ provided that the circle $\{s\} \times S^1 \subset Z$ along which we cut Z does not contain any critical values.

(Deformation invariance) This property holds for $\Phi(E, \Omega, J')$ without any essential modifications; to be precise, the situation in which deformation invariance holds is when there are two two-forms Ω_0, Ω_1 and a single almost complex structure J' such that (E, Ω_0, J') and (E, Ω_1, J') are ordinary almost holomorphic fibrations.

The induced maps $\Phi(E, \Omega, J')$ are the main topic of this Part. We proceed in the following way: the next section sets out the basic notions used to define Floer homology. In section 9 we state the definition of Floer homology and of the maps $\Phi(E, \Omega, J')$ and their main properties, without proof, and explain how these imply that the Floer homology groups are independent of the various choices involved in their definition. Section 10 concerns the quantum module structure and its relationship to the maps $\Phi(E, \Omega, J')$. In sections 11–15 we discuss some of the details involved in defining the maps $\Phi(E, \Omega, J')$. The reader who is not interested in these details can skip these sections and proceed to Part III.

8 Preliminaries on sections

The Floer homology groups $HF_*(T, \Theta)$ and the homomorphisms $\Phi(E, \Omega, J')$ are defined in terms of two classes of sections with special properties: *horizontal* sections of a symplectic fibre bundle over S^1 and, more importantly, *J-holomorphic* sections of an almost holomorphic fibration over a Riemann surface. This section contains as much of the theory of these two classes of sections as is necessary to formulate the definition of $HF_*(T, \Theta)$ and $\Phi(E, \Omega, J')$.

Horizontal sections

Definition 8.1. Let (E, Ω) be a symplectic fibre bundle over a manifold B. A smooth section $\nu : B \longrightarrow E$ is horizontal if $D\nu(Z) \in TE^h$ for all $Z \in TB$.

We will denote the set of such sections by $\mathcal{H}(E, \Omega)$. An important property of a horizontal section ν is that the pullback vector bundle $\nu^*(TE^v, \Omega|TE^v)$ carries a canonical symplectic connection ∇^ν, defined as follows: let $X \in C^\infty(TB)$ and $Y \in C^\infty(\nu^*TE^v)$. Choose a $\widetilde{Y} \in C^\infty(TE^v)$ such that $\nu^*\widetilde{Y} = Y$; we define

$$\nabla^\nu_X(Y) = \nu^*([X^\natural, \widetilde{Y}]).$$

Consider first the case $Y = 0$. Then $[X^\natural, \widetilde{Y}]_{\nu(z)}$ is the derivative of $\widetilde{Y}$ in $X^\natural$-direction at $\nu(z)$. Since $X^\natural$ is tangent to ν and $\widetilde{Y}$ vanishes along ν it follows that $\nu^*([X^\natural, \widetilde{Y}]) = 0$ for all $\widetilde{Y}$. This proves that in general, $\nabla^\nu_X(Y)$ is independent of the choice of $\widetilde{Y}$. To prove that ∇^ν is a connection we use the standard formulae

$$\begin{aligned} [fW, Z] &= f[W, Z] - (Z.f)W \quad \text{and} \\ [W, gZ] &= g[W, Z] + (W.g)Z \end{aligned}$$

for $W, Z \in C^\infty(TE)$ and $f, g \in C^\infty(E, \mathbb{R})$. Set $W = X^\natural$, $Z = \widetilde{Y}$ and let f be the pullback of a function h on B. Then $(\widetilde{Y}.f)X^\natural = 0$ because $\widetilde{Y}$ is vertical. This proves that $\nabla^\nu_{hX}(Y) = h\,\nabla^\nu_X Y$ and $\nabla^\nu_X(hY) = h\,\nabla^\nu_X(Y) + (X.h)Y$. To see that ∇^ν is symplectic, consider

$$(\nabla^\nu_X \Omega)(Y_1, Y_2) = \nu^* \left[X^\natural.\Omega(\widetilde{Y}_1, \widetilde{Y}_2) - \Omega([X^\natural, \widetilde{Y}_1], \widetilde{Y}_2) - \Omega(\widetilde{Y}_1, [X^\natural, \widetilde{Y}_2]) \right].$$

Since $X^\natural$ is horizontal and $\widetilde{Y}_k$ is vertical, $\widetilde{Y}_1.\Omega(X^\natural, \widetilde{Y}_2) = \widetilde{Y}_2.\Omega(X^\natural, \widetilde{Y}_1) = 0$ and $\Omega(X^\natural, [\widetilde{Y}_1, \widetilde{Y}_2]) = 0$. It follows that

$$(\nabla^\nu_X \Omega)(Y_1, Y_2) = \nu^*(d\Omega(X^\natural, \widetilde{Y}_1, \widetilde{Y}_2)) = 0.$$

$Y \in C^\infty(\nu^* TE^v)$ is parallel for ∇^ν iff $[X^\natural, \widetilde{Y}] = 0$ for all X. Hence the parallel transport of ∇^ν is given by the derivative of the symplectic parallel transport on the fibre bundle E. This is another possible approach to ∇^ν.

Consider a horizontal section ν of a symplectic fibre bundle (T, Θ) over S^1. The monodromy of ∇^ν around S^1 defines a symplectic linear map of the vertical tangent space of T at $\nu(t)$ to itself for any $t \in S^1$. We denote this map by $m^\nu(t)$. The monodromy maps corresponding to different choices of z are conjugate, and when only conjugation-invariant properties are concerned, we will usually omit t from the notation. We call ν *nondegenerate* if $(\mathrm{id} - m^\nu)$ is invertible. (T, Θ) is called nondegenerate if all its horizontal sections are nondegenerate. The degree $\deg(\nu) \in \mathbb{Z}/2$ of a nondegenerate horizontal section ν is defined by

$$\deg(\nu) = \begin{cases} 0 & \det(\mathrm{id} - m^\nu) > 0, \\ 1 & \det(\mathrm{id} - m^\nu) < 0 \end{cases}.$$

The meaning of these notions becomes clear if we consider the fibre bundle (T_ϕ, Θ_ϕ). Sections of T_ϕ correspond to maps $v : \mathbb{R} \longrightarrow M$ which satisfy $v(t) = \phi(v(t+1))$. Horizontal sections correspond to constant maps $v(t) \equiv x \in M$. Because of the periodicity condition, there is one such section for every fixed point x of ϕ, and if ν denotes this section, m^ν is conjugate to $D\phi(x)$. Hence ν is nondegenerate iff x is a nondegenerate fixed point, and $(-1)^{\deg(\nu)}$ is the local Lefschetz fixed point index of x. Because any symplectic fibre bundle over S^1 is isomorphic to (T_ϕ, Θ_ϕ) for some ϕ, one can always think of horizontal sections of such a fibre bundle as fixed points: for example, this is the most convenient way to prove that a nondegenerate fibre bundle (T, Θ) has only finitely many horizontal sections.

Tubular ends

The setup which is used to actually define Floer homology differs from that described in section 7 in several respects. One difference is that the finite cylinders Z are replaced by the infinite cylinder $C = \mathbb{R} \times S^1$. We begin by considering a product bundle

$$(E, \Omega) = \mathbb{R} \times (T, \Theta) \longrightarrow C.$$

This bundle carries a natural action of $\mathbb{R}$ by translation. Let K be the vector field which generates this action. It satisfies $i_K \Omega = 0$ because Ω is the pullback of Θ under the quotient map $E \longrightarrow E/\mathbb{R} = T$. This implies that K is horizontal; it is the horizontal lift of the vector field $\partial/\partial s$ on C.

Lemma 8.2. *Any horizontal section of* (E, Ω) *is of the form*

$$\sigma(s,t) = (s, \nu(t)),$$

where ν *is a horizontal section of* T*. Conversely, any section* σ *of this form is horizontal.*

Proof. If σ is a horizontal section, $\sigma|\{s\} \times S^1$ is a horizontal section of T for all $s \in \mathbb{R}$. On the other hand, $\partial\sigma/\partial s = K$, which implies that σ is translation-invariant. The proof of the converse is similar. □

Instead of symplectic fibre bundles over Z as in section 7 we will use symplectic fibre bundles over C which are isomorphic to product bundles outside a compact subset:

Definition 8.3. Let (T^-, Θ^-) and (T^+, Θ^+) be symplectic fibre bundles over S^1. A symplectic fibre bundle over C with tubular ends modelled on $(T^\pm, \Theta^\pm)$ is a symplectic fibre bundle (E, Ω), together with isomorphisms

$$\begin{aligned} \eta^- : (-\infty; -R] \times (T^-, \Theta^-) &\longrightarrow (E,\Omega)|(-\infty; -R] \times S^1 \text{ and} \\ \eta^+ : [R; \infty) \times (T^+, \Theta^+) &\longrightarrow (E,\Omega)|[R;\infty) \times S^1 \end{aligned} \tag{8.1}$$

for some $R > 0$. Usually, we identify the image of $\eta^\pm$ directly with the corresponding parts of $\mathbb{R} \times (T^\pm, \Theta^\pm)$ and do not mention the maps $\eta^\pm$.

Choose a Riemannian metric g on E whose restriction to the tubular ends is the product of the standard metric on $(-\infty; -R]$ or $[R; \infty)$ and of metrics on $T^\pm$. We denote its Levi-Civita connection by ∇^g and its exponential map by $\exp^g$.

Let ν^+ be a horizontal section of (T^+, Θ^+) and σ^+ the corresponding horizontal section of $E|[R;\infty) \times S^1$. We say that a section σ of E converges exponentially to ν^+ if there is an $R' \geq R$, a $\delta > 0$, and a vector field $\xi^+ \in (\sigma^+)^* TE^v$ with

$$|\xi^+(s,t)| + |\nabla^g \xi^+(s,t)| \leq e^{-\delta s}, \tag{8.2}$$

such that

$$\sigma(s,t) = \exp^g_{\sigma^+(s,t)}(\xi^+(s,t))$$

for all $(s,t) \in [R';\infty) \times S^1$. ν^+ is called the positive limit of σ. Negative limits, which are horizontal sections of (T^-, Θ^-), are defined in the same way. A section with horizontal limits is one which has both a negative and a positive limit. These are our 'boundary conditions' for sections of E.

Energy and index

We will now define two numbers associated to a smooth section σ with horizontal limits: its energy $e(\sigma) \in \mathbb{R}$ and its (Maslov) index $\mathrm{ind}(\sigma) \in \mathbb{Z}$. The energy is very much the simpler one:

$$e(\sigma) = \int_C \sigma^*\Omega.$$

This integral converges for the following reason: let ν^+ be the horizontal limit of σ. $(\nu^+)^*\Theta^+ = 0$ because any two-form on S^1 vanishes. Let σ^+ be the section of $E|[R;\infty) \times S^1$ to which σ is asymptotic, that is, $\sigma^+(s,t) = (s, \nu^+(t))$. This section satisfies $[(\sigma^+)^*\Omega]_{(s,t)} = [(\nu^+)^*\Theta^+]_t = 0$. Because of the decay condition (8.2), it follows that

$$|\sigma^*\Omega|_{(s,t)} \leq \mathrm{Const.}e^{-2\delta s}$$

for $s \geq R'$. Hence the integral $\int \sigma^*\Omega$ converges for $s \to \infty$; the same holds on the other end.

The index is defined for sections of E whose horizontal limits are nondegenerate. Let σ be such a section, with limits ν^-, ν^+, whose asymptotic behaviour is

$$\sigma(s,t) = \exp^g_{\sigma^-(s,t)}(\xi^-(s,t)) \text{ for } s \leq -R' \text{ and}$$
$$\sigma(s,t) = \exp^g_{\sigma^+(s,t)}(\xi^+(s,t)) \text{ for } s \geq R'.$$

with $\sigma^\pm$ and $\xi^\pm$ as above. Choose a cutoff function $\psi \in C^\infty(\mathbb{R},\mathbb{R})$ with $\psi(s) = 1$ for $|s| \leq R'$ and $\psi(s) = 0$ for $|s| \geq R' + 1$. The partial sections $\sigma^\pm$ of E are horizontal. Therefore the section

$$\sigma'(s,t) = \begin{cases} \exp^g_{\sigma^-(s,t)}(\psi(s)\xi^-(s,t)) & s \leq -R' \\ \sigma(s,t) & -R' \leq s \leq R' \\ \exp^g_{\sigma^+(s,t)}(\psi(s)\xi^+(s,t)) & s \geq R' \end{cases}$$

is horizontal at any point $(s,t) \in C$ with $|s| \geq R' + 1$. From our discussion of horizontal sections it follows that the symplectic vector bundle $(\sigma')^*(TE^v, \Omega|TE^v)$ carries a canonical symplectic connection defined outside $[-R'-1; R'+1] \times S^1 \subset C$. Let $\nabla^{\sigma'}$ be an extension of that connection over all of C. After choosing a trivialization of $(\sigma')^*(TE^v, \Omega|TE^v)$ (any symplectic vector bundle over C is trivial), the monodromy of $\nabla^{\sigma'}$ around the circles $\{s\} \times S^1 \subset C$ defines a smooth path

$$m^{\sigma'} : \mathbb{R} \longrightarrow \mathrm{Sp}(4,\mathbb{R}).$$

which is locally constant for $|s| \geq R' + 1$. By construction $m^{\sigma'}(\pm(R+1))$ is conjugate to $m^{\nu^\pm}$. Hence $m^{\sigma'}|[-R'-1; R'+1]$ is a path whose endpoints lie in the subset $\mathrm{Sp}(4,\mathbb{R})^*$ of symplectic matrices A with $\det(\mathrm{id} - A) \neq 0$.

The definition of the index of σ is based on an invariant of such paths called the Maslov index. It can be characterized as follows: The Maslov index is the unique map

$$\mu : \pi_1(\mathrm{Sp}(2n, \mathbb{R}), \mathrm{Sp}(2n, \mathbb{R})^*) \longrightarrow \mathbb{Z}$$

such that

(1) μ is additive under composition of paths.

(2) The path $\gamma : [-1; 1] \longrightarrow \mathrm{Sp}(2n, \mathbb{R})$, $\gamma(t) = \exp(J\,\mathrm{diag}(t, 1, \dots 1))$, has Maslov index 1. The same holds if we replace $\mathrm{diag}(t, 1, \dots, 1)$ by $\mathrm{diag}(t, -1, 1, \dots, 1)$.

The homotopy invariance implies that μ is conjugation invariant in the following sense: if $\gamma : [0; 1] \longrightarrow \mathrm{Sp}(2n, \mathbb{R})$ is a path with $\gamma(0), \gamma(1) \in \mathrm{Sp}(2n, \mathbb{R})^*$ and $\theta : [0; 1] \longrightarrow \mathrm{Sp}(2n, \mathbb{R})$ is an arbitrary path, the Maslov index of the path $\gamma'(t) = \theta(t)\gamma(t)\theta(t)^{-1}$ is equal to the Maslov index of γ. We refer to [25, section 3] for an extensive discussion of the Maslov index (in a slightly different form).

We can now define the index of a section σ with nondegenerate horizontal limits:

$$\mathrm{ind}(\sigma) = \mu(m^{\sigma'} | [-R' - 1; R' + 1]).$$

$\mu(m^{\sigma'})$ is independent of the choice of trivialization of $(\sigma')^*(TE^v, \Omega | TE^v)$ because different choices lead to paths in $\mathrm{Sp}(4, \mathbb{R})$ which are conjugate. The path $m^{\sigma'}$ also depends on the choice of $\nabla^{\sigma'}$ and of σ' itself, but the homotopy invariance of μ ensures that $\mu(m^{\sigma'})$ is independent of these choices. Therefore $\mathrm{ind}(\sigma)$ is well-defined. We refer again to [25] for the proof of the equality

$$\mathrm{ind}(\sigma) \equiv \deg(\nu^+) - \deg(\nu^-) \bmod 2, \tag{8.3}$$

where ν^- and ν^+ are the limits of σ.

J-holomorphic sections

Let (E, Ω) be a symplectic fibre bundle over a Riemann surface (Σ, j), with projection $\pi : E \longrightarrow \Sigma$.

Definition 8.4. An almost complex structure J on E is *partially Ω-tame* if π is (J, j)-linear (that is, $D\pi \circ J = j \circ D\pi$) and $\Omega(X, JX) > 0$ for all nonzero $X \in TE^v$.

With respect to the splitting $TE = TE^v \oplus TE^h$, a partially Ω-tame almost complex structure has the form

$$J = \begin{pmatrix} J^{vv} & J^{vh} \\ 0 & J^{hh} \end{pmatrix}. \tag{8.4}$$

J^{hh} is the almost complex structure on TE^h which corresponds to j under the isomorphism $D\pi|TE^h : TE^h \longrightarrow \pi^*T\Sigma$; J^{vv} is an almost complex structure on TE^v tamed by $\Omega|TE^v$, and $J^{vh} : (TE^h, J^{hh}) \longrightarrow (TE^v, J^{vv})$ is a $\mathbb{C}$-antilinear homomorphism. Conversely, every pair (J^{vv}, J^{vh}) with these properties determines a partially Ω-tame almost complex structure.

Definition 8.5. Let J be a partially Ω-tame almost complex structure. A smooth section $\sigma : \Sigma \longrightarrow E$ is *J-holomorphic* if its differential $D\sigma : (T\Sigma, j) \longrightarrow (TE, J)$ is $\mathbb{C}$-linear.

Assume that $\Sigma = C$ and that (E, Ω) has tubular ends modelled on $(T^\pm, \Theta^\pm)$. We will use the following notation:

Notation. For $J^- \in \mathcal{J}(T^-, \Theta^-)$ and $J^+ \in \mathcal{J}(T^+, \Theta^+)$, $\mathcal{J}(E, \Omega; J^-, J^+)$ is the space of partially Ω-compatible almost complex structures on E which agree with J^- on $E|(-\infty; -R'] \times S^1$ and with J^+ on $E|[R'; \infty) \times S^1$, for some large R'. For $J \in \mathcal{J}(E, \Omega; J^-, J^+)$, $\mathcal{M}(E, J)$ is the set of J-holomorphic sections of E with horizontal limits. The subset of sections with limits $\nu^\pm \in \mathcal{H}(T^\pm, \Theta^\pm)$ will be denoted by $\mathcal{M}(E, J; \nu^-, \nu^+)$, the subset of sections with energy $\epsilon \in \mathbb{R}$ by $\mathcal{M}_\epsilon(E, J)$ and the subset of sections with index $k \in \mathbb{Z}$ by $\mathcal{M}_k(E, J)$ (the ambiguity of this notation will not cause any problems). We will also use various intersections of these subsets, e.g. $\mathcal{M}_k(E, J; \nu^-, \nu^+)$ and $\mathcal{M}_{k,\epsilon}(E, J; \nu^-, \nu^+)$.

Product bundles

Let us return for a moment to the case of a product $(E, \Omega) = \mathbb{R} \times (T, \Theta)$, where (T, Θ) is a symplectic fibre bundle over S^1. Such a product bundle satisfies

$$\Omega|TE^h = 0.$$

To prove this it is sufficient (because TE^h is two-dimensional) to find a nowhere vanishing $X \in C^\infty(TE^h)$ such that $i_X\Omega = 0$. As mentioned above, the vector field K generating the $\mathbb{R}$-action has this property. We will now introduce a particular class of almost complex structures on (E, Ω) which are important for the definition of Floer homology groups.

Definition 8.6. $\mathcal{J}(T, \Theta)$ is the space of almost complex structures J on $E = \mathbb{R} \times T$ with the following properties:

(1) The map $E \longrightarrow C$ is (J, j)-linear (here j is the standard complex structure on $\mathbb{R} \times S^1$);

(2) J is translation-invariant;

(3) $\Omega(\cdot, J\cdot)$ is a symmetric bilinear form on TE which is positive definite on the vertical subbundle TE^v.

Because of the last condition, any $J \in \mathcal{J}(T, \Theta)$ is an Ω-tame almost complex structure. Moreover, such a J preserves the subbundles $TE^v, TE^h \subset TE$; the first one is preserved by any Ω-tame almost complex structure, and since the second one is orthogonal to the first one with respect to Ω, its invariance under J follows from the fact that $\Omega(\cdot, J\cdot)$ is symmetric. $J(TE^h) = TE^h$ implies that a horizontal section of E is J-holomorphic for any $J \in \mathcal{J}(T, \Theta)$ because for such a section σ, $\mathrm{im}(D\sigma_z) = TE^h_{\sigma(z)}$ is a J-linear subspace for all $z \in C$.

$J(TE^h) = TE^h$ means that

$$J = \begin{pmatrix} J^{vv} & 0 \\ 0 & J^{hh} \end{pmatrix} \tag{8.5}$$

with respect to $TE = TE^v \oplus TE^h$, where J^{hh} is the horizontal lift of j and J^{vv} is an Ω-compatible almost complex structure on TE^v. Because J is $\mathbb{R}$-invariant J^{vv} is in fact given by a family of almost complex structures on the fibres of T. For $(T, \Theta) = (T_\phi, \Theta_\phi)$ it follows that J is determined by a family $(J_t)_{t\in\mathbb{R}}$ of ω-compatible almost complex structures on M such that $J_t = D\phi \circ J_{t+1} \circ D\phi^{-1}$ for all t.

Lemma 8.7. *If $J \in \mathcal{J}(T, \Theta)$, $\Omega(X, JX) \geq 0$ for all $X \in TE$, with equality iff $X \in TE^h$.*

Proof. Let $X = X^v + X^h$ be the vertical and horizontal parts of X. Because J preserves TE^v and TE^h,

$$\Omega(X, JX) = \Omega(X^v, JX^v) + \Omega(X^h, JX^h).$$

The second term vanishes because $\Omega|TE^h = 0$, and since J is Ω-tame, $\Omega(X^v, JX^v) \geq 0$, with equality iff $X^v = 0$. □

Note that if $J \in \mathcal{J}(T, \Theta)$ and σ is J-holomorphic, its translates

$$\sigma^r(s, t) = r \cdot \sigma(s - r, t)$$

($\cdot$ is the $\mathbb{R}$-action by translation) are again J-holomorphic. If σ has horizontal limits so do its translates; and the energy and index are invariant under translation.

Lemma 8.8. *Let σ be a J-holomorphic section of $E = \mathbb{R} \times T$ with horizontal limits, for some $J \in \mathcal{J}(T, \Theta)$. Then $e(\sigma) \geq 0$ and the following conditions are equivalent:*

(i) $e(\sigma) = 0$*;*
(ii) σ *is a horizontal section;*
(iii) σ *is translation-invariant;*
(iv) there is an $r \in \mathbb{R}$ *such that* $\sigma^r = \sigma$*.*

If T is nondegenerate, any horizontal section σ has index zero.

Proof. The nonnegativity of the energy is a consequence of Lemma 8.7 and so is the equivalence (i) $\Leftrightarrow$ (ii). If σ is a horizontal section, $\partial\sigma/\partial s$ is the horizontal lift of $\partial/\partial s$, and therefore $\partial\sigma/\partial s = K$. This shows that σ is translation-invariant and hence that (ii) $\Rightarrow$ (iii). (iii) $\Rightarrow$ (iv) is obvious. Finally, note that if $\sigma = \sigma^r$ for some r, the form $\sigma^*\Omega$ is r-periodic and (by Lemma 8.7) nonnegative. The integral $\int_C \sigma^*\Omega$ is finite because σ has horizontal limits; but the integral of a nonnegative periodic two-form on C can only be finite if the form vanishes identically. This shows that (iv) $\Rightarrow$ (i).

Finally, note that the definition of $\mathrm{ind}(\sigma)$ becomes much simpler for horizontal σ: one can choose $\sigma' = \sigma$, and the path $m^{\sigma'}$ is given by the monodromies of ∇^σ around $\{s\} \times S^1$. In our case, because σ has the form $\sigma(s,t) = (s, \nu(t))$ for some $\nu \in \mathcal{H}(T, \Theta)$, the path $m^{\sigma'}$ is constant; a constant path has zero Maslov index. $\square$

Sections of almost holomorphic fibrations

Definition 8.9. An almost holomorphic fibration over C with tubular ends modelled on symplectic fibre bundles $(T^\pm, \Theta^\pm)$ over S^1 consists of an almost holomorphic fibration (E, Ω, J') over C and isomorphisms

$$\eta^- : (-\infty; -R] \times (T^-, \Theta^-) \longrightarrow (E, \Omega)|(-\infty; -R] \times S^1 \text{ and}$$
$$\eta^+ : [R; \infty) \times (T^+, \Theta^+) \longrightarrow (E, \Omega)|[R; \infty) \times S^1$$

for some $R > 0$.

Everything we have said about sections of symplectic fibre bundles applies equally to sections of such fibrations, for the following simple reason: a smooth section σ can never go through a critical point of the map $\pi : E \longrightarrow C$. This is obvious: $\pi \circ \sigma = \mathrm{id}$ implies $D\pi \circ = \mathrm{id}$, which implies that $D\pi_{\sigma(z)}$ is onto for all $z \in C$.

In particular, we retain the definitions of $TE^v = \ker(D\pi)$ and of TE^h (note that these are only vector bundles away from $\mathrm{Crit}(\pi)$) and of the horizontal limits, energy and index of a section. The definition a partially Ω-tame almost complex structure remains the same; however, for technical reasons, we will use only those Ω-tame almost complex structures which agree with J' in a neighbourhood of $\mathrm{Crit}(\pi)$. The space of such J will be denoted by $\mathcal{J}(E,\Omega,J')$, and the subspace of almost complex structures which agree with $J^\pm \in \mathcal{J}(T^\pm,\Theta^\pm)$ outside a compact subset by $\mathcal{J}(E,\Omega,J';J^-,J^+) \subset \mathcal{J}(E,\Omega,J')$.

9 The Floer chain complex and the induced homomorphisms

Let (T,Θ) be a nondegenerate symplectic fibre bundle over S^1. The Floer chain group $CF_*(T,\Theta)$ is the $\mathbb{Z}/2$-graded Λ-vector space freely generated by the (finite) set of horizontal sections of (T,Θ), that is,

$$CF_i(T,\Theta) = \bigoplus_{\substack{\nu\in\mathcal{H}(T,\Theta)\\ \deg(\nu)=i}} \Lambda <\nu>$$

for $i=0,1$. Floer homology is the homology of a certain boundary operator on $CF_*(T,\Theta)$. The definition of this boundary operator is based on deep results about the spaces $\mathcal{M}(\mathbb{R}\times T,J)$ of J-holomorphic sections, for $J \in \mathcal{J}(T,\Theta)$. These results are summarized in the next Theorem. Recall that the spaces $\mathcal{M}(\mathbb{R}\times T,J)$ carry a natural action of $\mathbb{R}$ by translation; we denote the quotients by $\mathcal{M}(\mathbb{R}\times T,J)/\mathbb{R}$.

Theorem 9.1. *There is a dense subset $\mathcal{J}_{\mathrm{reg}}(T,\Theta) \subset \mathcal{J}(T,\Theta)$ such that any $J \in \mathcal{J}_{\mathrm{reg}}(T,\Theta)$ has the following properties:*

(1) For all $\epsilon \in \mathbb{R}$ and $\nu^-,\nu^+ \in \mathcal{H}(T,\Theta)$, $\mathcal{M}_{1,\epsilon}(\mathbb{R}\times T,J;\nu^-,\nu^+)/\mathbb{R}$ is a finite set.

(2) For any ν^-,ν^+ and any $C>0$, there are only finitely many $\epsilon \le C$ such that $\mathcal{M}_{1,\epsilon}(\mathbb{R}\times T,J;\nu^-,\nu^+) \neq \emptyset$.

(3) Each of the sets $\mathcal{M}_{2,\epsilon}(\mathbb{R}\times T,J;\nu^-,\nu^+)/\mathbb{R}$ can be given the structure of a smooth one-dimensional manifold. This manifold has a compactification which is a compact one-dimensional manifold; the boundary of this compactification is the disjoint union of the sets

$$\mathcal{M}_{1,\delta}(\mathbb{R}\times T,J;\nu^-,\nu)/\mathbb{R} \times \mathcal{M}_{1,\epsilon-\delta}(\mathbb{R}\times T,J;\nu,\nu^+)/\mathbb{R},$$

where (ν,δ) runs over $\mathcal{H}(T,\Theta)\times\mathbb{R}$ (the first two properties imply that this boundary is a finite set).

Remark 9.2. We do *not* define $\mathcal{J}_{\mathrm{reg}}(T,\Theta)$ as the set of all almost complex structures having property (1)–(3) (otherwise the results stated later on, which involve $\mathcal{J}_{\mathrm{reg}}(T,\Theta)$, would be false). The correct definition of $\mathcal{J}_{\mathrm{reg}}(T,\Theta)$ is given in section 11.

For $J \in \mathcal{J}_{\mathrm{reg}}(T,\Theta)$, $\nu^-,\nu^+ \in \mathcal{H}(T,\Theta)$, and $\epsilon \in \mathbb{R}$, let $n_\epsilon(J;\nu^-,\nu^+) \in \mathbb{Z}/2$ be the number of points mod 2 in $\mathcal{M}_{1,\epsilon}(\mathbb{R}\times T, J;\nu^-,\nu^+)/\mathbb{R}$. Property (2) of J implies that the formal sum

$$n(J;\nu^-,\nu^+) = \sum_{\epsilon\in\mathbb{R}} n_\epsilon(J;\nu^-,\nu^+)t^\epsilon \tag{9.1}$$

is an element of the Novikov field Λ. This 'number' can be thought of as the 'number of points' in the (possibly infinite) set $\mathcal{M}_1(\mathbb{R}\times T, J;\nu^-,\nu^+)/\mathbb{R}$. Lemma 8.8 implies that $\mathcal{M}_{1,\epsilon}(\mathbb{R}\times T, J;\nu^-,\nu^+) = \emptyset$ for all $\epsilon < 0$. By the same Lemma, this is true for $\epsilon = 0$ as well: any J-holomorphic section with zero energy is horizontal, and all horizontal sections have zero index. Therefore all nonzero coefficients in (9.1) have positive exponents.

The Floer differential $\partial(T,\Theta;J)$ on $CF_*(T,\Theta)$ is defined by its values on the generators $<\nu>$:

$$\partial(T,\Theta;J)(<\nu^->) \stackrel{\mathrm{def}}{=} \sum_{\nu^+\in\mathcal{H}(T,\Theta)} n(J;\nu^-,\nu^+)\ <\nu^+>\ .$$

Because the boundary of a compact one-manifold consists of an even number of points, property (3) of J implies that

$$\sum_{\substack{\delta\in\mathbb{R}\\ \nu\in\mathcal{H}(T,\Theta)}} n_\delta(J;\nu^-,\nu)n_{\epsilon-\delta}(J;\nu,\nu^+) = 0.$$

for all ν^-, ν^+ and ϵ. This is equivalent to $\partial(T,\Theta;J)\circ\partial(T,\Theta;J) = 0$. The mod 2 formula (8.3) for the index of a section implies that $\mathcal{M}_1(\mathbb{R}\times T;\nu^-,\nu^+) = \emptyset$ unless $\deg(\nu^-) \neq \deg(\nu^+)$. Hence $\partial(T,\Theta;J)$ interchanges the two groups $CF_i(T,\Theta)$. This shows that $(CF_*(T,\Theta),\partial(T,\Theta;J))$ is a $\mathbb{Z}/2$-graded chain complex of Λ-vector spaces. In general the differential $\partial(T,\Theta;J)$ will depend on the choice of J.

Definition 9.3. Let (T,Θ) be a nondegenerate symplectic fibre bundle over S^1 and $J \in \mathcal{J}_{\mathrm{reg}}(T,\Theta)$. The Floer homology groups $HF_i(T,\Theta;J)$ $(i = 0,1)$ are the homology groups of the chain complex $(CF_*(T,\Theta),\partial(T,\Theta;J))$.

We outline briefly the 'duality' property of these groups. Let (T,Θ) be a nondegenerate symplectic fibre bundle over S^1 and $(\overline{T},\overline{\Theta})$ its pullback by ι. There is a canonical bijection between the sets of horizontal sections on these two bundles. Using this bijection and the natural bases, one can define a nondegenerate pairing

$$\langle\cdot,\cdot\rangle_{CF_*(T,\Theta)} : CF_*(\overline{T},\overline{\Theta}) \otimes CF_*(T,\Theta) \longrightarrow \Lambda.$$

The bijection $\mathcal{H}(T,\Theta) \longrightarrow \mathcal{H}(\overline{T},\overline{\Theta})$ preserves the degree; this is an elementary fact which follows from the equality

$$\text{sign}\ \det(\text{id} - A) = \text{sign}\ \det(\text{id} - A^{-1})$$

for any symplectic matrix A which does not have 1 as an eigenvalue. As a consequence, the pairing $\langle\cdot,\cdot\rangle_{CF_*(T,\Theta)}$ is $\mathbb{Z}/2$-graded. Take an almost complex structure $J \in \mathcal{J}_{\text{reg}}(T,\Theta)$, and let $\overline{J}$ be its pullback to $\mathbb{R}\times\overline{T}$ by the holomorphic involution $(s,t) \longmapsto (-s,-t)$ of C. Pulling back sections defines a canonical bijection between the spaces $\mathcal{M}(\mathbb{R}\times T, J)$ and $\mathcal{M}(\mathbb{R}\times\overline{T},\overline{J})$. A look at the definition of $\mathcal{J}_{\text{reg}}(T,\Theta)$ given in section 13 reveals that $\overline{J} \in \mathcal{J}_{\text{reg}}(\overline{T},\overline{\Theta})$. Using these facts, it is not difficult to prove that $\langle\cdot,\cdot\rangle_{CF_*(T,\Theta)}$ is a pairing of chain complexes. Since the coefficient ring Λ is a field, the pairing induced by $\langle\cdot,\cdot\rangle_{CF_*(T,\Theta)}$ on the Floer homology groups is nondegenerate.

Let (E,Ω,J') be an ordinary almost holomorphic fibration over C whose tubular ends are modelled on nondegenerate fibre bundles (T^-,Θ^-) and (T^+,Θ^+).

Theorem 9.4. *For all $J^- \in \mathcal{J}_{\text{reg}}(T^-,\Theta^-)$ and $J^+ \in \mathcal{J}_{\text{reg}}(T^+,\Theta^+)$, there is a dense subset $\mathcal{J}_{\text{reg}}(E,\Omega,J';J^-,J^+) \subset \mathcal{J}(E,\Omega,J';J^-,J^+)$ such that any $J \in \mathcal{J}_{\text{reg}}(E,\Omega,J';J^-,J^+)$ has the following properties:*

(1) For all $\epsilon \in \mathbb{R}$, $\nu^- \in \mathcal{H}(T^-,\Theta^-)$ and $\nu^+ \in \mathcal{H}(T^+,\Theta^+)$, $\mathcal{M}_{0,\epsilon}(E,J;\nu^-,\nu^+)$ is a finite set;

(2) For any ν^-,ν^+ and $C \in \mathbb{R}$ there are only finitely many $\epsilon \leq C$ such that $\mathcal{M}_{0,\epsilon}(E,J;\nu^-,\nu^+) \neq \emptyset$.

(3) Each of the sets $\mathcal{M}_{1,\epsilon}(E,J;\nu^-,\nu^+)$ can be given the structure of a smooth one-dimensional manifold. It has a compactification which is a compact one-dimensional manifold with boundary, and the boundary of this compactification is the disjoint union of

$$\mathcal{M}_{0,\delta}(E,J;\nu^-,\nu) \times \big(\mathcal{M}_{1,\epsilon-\delta}(\mathbb{R}\times T^+,J^+;\nu,\nu^+)/\mathbb{R}\big) \qquad (\textit{for } \nu \in \mathcal{H}(T^+,\Theta^+),\ \delta \in \mathbb{R})$$

and

$$\big(\mathcal{M}_{1,\delta'}(\mathbb{R}\times T^-,J^-;\nu^-,\nu')/\mathbb{R}\big) \times \mathcal{M}_{0,\epsilon-\delta'}(E,J;\nu',\nu^+) \qquad (\textit{for } \nu' \in \mathcal{H}(T^-,\Theta^-),\ \delta' \in \mathbb{R}).$$

The almost complex structures $J \in J_{\text{reg}}(E, \Omega, J'; J^-, J^+)$ will be called regular. As in the case of $\mathcal{J}_{\text{reg}}(T, \Theta)$, $\mathcal{J}_{\text{reg}}(E, \Omega, J'; J^-, J^+)$ is not the set of all J which have the properties listed above.

For $J \in \mathcal{J}_{\text{reg}}(E, \Omega, J'; J^-, J^+)$ and $\nu^\pm \in \mathcal{H}(T^\pm, \Theta^\pm)$, let $m_\epsilon(J; \nu^-, \nu^+) \in \mathbb{Z}/2$ be the number of points mod 2 in $\mathcal{M}_{0,\epsilon}(E, J; \nu^-, \nu^+)$. Property (2) implies that $m(J; \nu^-, \nu^+) = \sum_\epsilon m_\epsilon(J; \nu^-, \nu^+) t^\epsilon$ is an element of Λ. Let $C\Phi(E, \Omega, J'; J) : CF_*(T^-, \Theta^-) \longrightarrow CF_*(T^+, \Theta^+)$ be the homomorphism defined by

$$C\Phi(E, \Omega, J'; J)(<\nu^->) = \sum_{\nu^+} m(J; \nu^-, \nu^+) <\nu^+> .$$

The mod 2 formula for the index shows that $\mathcal{M}_0(E, J; \nu^-, \nu^+) = \emptyset$ unless $\deg(\nu^-) = \deg(\nu^+)$; therefore $C\Phi(E, \Omega, J'; J)$ preserves the $\mathbb{Z}/2$-grading. By a straightforward computation, property (3) implies that $C\Phi(E, \Omega, J'; J)$ is a homomorphism of chain complexes. This homomorphism depends on the choice of J.

Definition 9.5. The homomorphism induced by $C\Phi(E, \Omega, J; J')$ will be denoted by

$$\Phi(E, \Omega, J'; J) : HF_*(T^-, \Theta^-; J^-) \longrightarrow HF_*(T^+, \Theta^+; J^+).$$

Theorem 9.6. $\Phi(E, \Omega, J'; J)$ *is independent of the choice of* J.

The proof of this theorem could be formulated like the definition of the Floer homology groups and of the homomorphism Φ; properties of the moduli spaces $\mathcal{M}(E, J_t)$ for a suitably chosen one-parameter family $(J_t)_{0 \le t \le 1}$ in $\mathcal{J}(E, \Omega, J'; J^-, J^+)$ with endpoints $J_0, J_1 \in \mathcal{J}_{\text{reg}}(E, \Omega, J'; J^-, J^+)$ show that $C\Phi(E, \Omega, J'; J_0$ and $C\Phi(E, \Omega, J'; J_1)$ are chain homotopic. We omit the precise statement.

From now on we will write $\Phi(E, \Omega, J'; J^-, J^+)$ instead of $\Phi(E, \Omega, J'; J)$; this change of notation is justified by Theorem 9.6. Let (E, Ω) be a product bundle $\mathbb{R} \times (T, \Theta)$, and choose a $J^0 \in \mathcal{J}_{\text{reg}}(T, \Theta)$. We will now sketch the proof of

$$\Phi(E, \Omega; J^0, J^0) = \mathrm{id}_{HF_*(T, \Theta; J^0)}. \tag{9.2}$$

The key idea is to use the same almost complex structure J^0 to define the Floer chain complex and the induced map. It is obvious that $J^0 \in \mathcal{J}(E, \Omega; J^0, J^0)$. Moreover, J^0 is regular, so that we can use it to define it the chain homomorphism $CF_*(T, \Theta) \longrightarrow CF_*(T, \Theta)$. We cannot prove this here, but it follows easily by comparing Definition 11.10 and 11.11. By Lemma 11.6, the J^0-holomorphic sections of $\mathbb{R} \times T$ with index zero are precisely the horizontal ones, and all these sections have zero energy. For every $\nu \in \mathcal{H}(T,\Theta)$, there is exactly one horizontal section of $E = \mathbb{R} \times T$ with positive limit ν, and the negative limit of that section is also ν. It follows that the chain homomorphism $C\Phi(E,\Omega; J^0, J^0)$ is the identity.

The next result corresponds to the 'gluing' property of Floer homology stated in section 7. We need to introduce some notation: let (E_-, Ω_-, J'_-) and (E_+, Ω_+, J'_+) be two ordinary almost holomorphic fibrations over C. Assume that the positive end of the first fibration is modelled on the same symplectic fibre bundle as the negative end of the second fibration. After choosing appropriate almost complex structures, we obtain induced maps

$$\Phi(E_-, \Omega_-, J'_-; J^-, J^0) : HF_*(T^-, \Theta^-; J^-) \longrightarrow HF_*(T^0, \Theta^0; J^0) \text{ and}$$
$$\Phi(E_+, \Omega_+, J'_+; J^0, J^+) : HF_*(T^0, \Theta^0; J^0) \longrightarrow HF_*(T^+, \Theta^+; J^+).$$

After changing the coordinates on C by a translation, we can assume that $(E_-, \Omega_-)|[-1;\infty) \times S^1 = [-1;\infty) \times (T^0, \Theta^0)$ and $(E_+, \Omega_+)|(-\infty; 1] \times S^1 = (-\infty; 1] \times (T^0, \Theta^0)$. We use these identifications to glue together $E_-|\mathbb{R}^- \times S^1$ and $E_+|\mathbb{R}^+ \times S^1$; the result is a new ordinary almost holomorphic fibration over C which we denote by

$$(E, \Omega, J') = (E_-, \Omega_-, J'_-)\#(E_+, \Omega_+, J'_+).$$

This fibration defines a homomorphism

$$\Phi(E, \Omega, J'; J^-, J^+) : HF_*(T^-, \Theta^-; J^-) \longrightarrow HF_*(T^+, \Theta^+; J^+),$$

and the 'gluing theorem' for the induced maps is

Theorem 9.7.

$$\Phi(E, \Omega, J'; J^-, J^+) = \Phi(E_+, \Omega_+, J'_+; J^0, J^+) \circ \Phi(E_-, \Omega_-, J'_-; J^-, J^0).$$

The definitions of Floer homology groups and induced maps which we have given seem to differ considerably from the presentation in section 7:

(a) The Floer homology groups have been defined only for nondegenerate fibre bundles over S^1.

(b) The definition involves an auxiliary choice of almost complex structure.

(c) Remarks similar to (a) and (b) apply to the induced maps Φ.

(d) We have not stated a result of the type of the 'deformation invariance' property in section 7.

All of these apparent differences can be overcome using Theorem 9.6 and Theorem 9.7. We begin with (b). Let (T, Θ) be a nondegenerate symplectic fibre bundle over S^1 and J^-, J^+ two almost complex structures in $\mathcal{J}_{\mathrm{reg}}(T, \Theta)$.

We denote the product bundle $\mathbb{R} \times (T, \Theta)$ by (E, Ω). An almost complex structure $J \in \mathcal{J}_{\text{reg}}(E, \Omega; J^-, J^+)$ determines a homomorphism

$$HF_*(T, \Theta; J^-) \longrightarrow HF_*(T, \Theta; J^+).$$

Let us denote this homomorphism (which is independent of the choice of J by Theorem 9.6) by $\Phi(J^-, J^+)$. Equation (9.2) says that $\Phi(J^0, J^0) = \text{id}$ for all $J^0 \in \mathcal{J}_{\text{reg}}(T, \Theta)$, and Theorem 9.7 says that

$$\Phi(J^0, J^+) \circ \Phi(J^-, J^0) = \Phi(J^-, J^+).$$

This shows that for different choices of J, the groups $HF_*(T, \Theta; J)$ are isomorphic, and in fact canonically isomorphic. This makes it possible to define Floer homology groups $HF_*(T, \Theta)$ (for nondegenerate (T, Θ)) which do not depend on a choice of almost complex structure.

The next issue which we will discuss is (d). Let (E, Ω, J') be an ordinary almost holomorphic fibration with nondegenerate tubular ends modelled on $(T^\pm, \Theta^\pm)$. Choose $J^\pm \in \mathcal{J}_{\text{reg}}(T^\pm, \Theta^\pm)$. Let Ω' be another closed two-form on E such that

(1) $\Omega(X, Y) = \Omega'(X, Y)$ for all $X, Y \in TE^v$,

(2) $\Omega = \Omega'$ outside some compact subset,

(3) $[\Omega - \Omega'] \in H^2_c(E; \mathbb{R})$ is trivial, and

(4) Ω' is compatible with J' on some neighbourhood of the critical point set of $E \longrightarrow C$.

Then (E, Ω', J) is an ordinary almost holomorphic fibration. Moreover

$$\mathcal{J}(E, \Omega', J'; J^-, J^+) = \mathcal{J}(E, \Omega, J'; J^-, J^+)$$

and the subsets of regular almost complex structures are also the same for Ω and Ω' (this follows from Definition 11.11). Moreover, and for this the assumption (3) is crucial, any section σ of E with horizontal limits has the same energy and index with respect to Ω and with respect to Ω'. It follows that

$$C\Phi(E, \Omega', J'; J) = C\Phi(E, \Omega, J'; J)$$

for any regular J, and this is the result corresponding to the 'deformation invariance' property of section 7.

Now consider problem (a). We have explained how to define Floer homology groups $HF_*(T, \Theta)$ which are independent of the choice of an almost complex structure for all nondegenerate (T, Θ). As in section 7, the deformation invariance provides canonical isomorphisms between different perturbations

$(T, \Theta - d(Hdt))$ and $(T, \Theta - d(H'dt))$ $(H, H' \in C^\infty(T, \mathbb{R})$ of a symplectic fibre bundle (T, Θ). Such a canonical isomorphism exists whenever both perturbed bundles are nondegenerate, regardless of whether (T, Θ) itself is nondegenerate or not. Hence, for an arbitrary (T, Θ), we can define

$$HF_*(T, \Theta) \stackrel{\text{def}}{=} HF_*(T, \Theta - d(Hdt))$$

where the r.h.s. is nondegenerate. The existence of an H such that $(T, \Theta - d(Hdt))$ is a nondegenerate symplectic fibre bundle can be proved using some elementary symplectic geometry: in fact, viewing (T, Θ) as a mapping torus reduces the problem to the assertion that any symplectic automorphism has a Hamiltonian perturbation whose fixed points are all transverse.

The remaining problem (c) is similar to (a) and we will not discuss it.

10 The quantum module structure

This section describes, in the informal style of section 7, the quantum module structure on Floer homology and its relationship to the induced maps $\Phi(E, \Omega, J')$. The basic objects which we will consider are ordinary almost holomorphic fibrations (E, Ω, J') over a cylinder $Z = [s_0; s_1] \times S^1$, together with a finite family $(z_1, \dots z_r)$ of points of Z which are regular values of the map $E \longrightarrow Z$. The quantum module product on Floer homology is a special case of a more general multiplicative structure: any $(E, \Omega, J', z_1, \dots, z_r)$ defines a homomorphism

$$\Phi_r(E, \Omega, J'; z_1, \dots, z_r) : \left(\bigotimes_{i=1}^{r} QH_*(E_{z_i}, \Omega | E_{z_i}) \right) \otimes \\ \otimes HF_*(E_{s_0 \times S^1}, \Omega_{s_0 \times S^1}) \longrightarrow HF_*(E_{s_1 \times S^1}, \Omega_{s_1 \times S^1}) \quad (10.1)$$

of $\mathbb{Z}/2$-graded Λ-modules. Here $QH_*(E_{z_i}, \Omega | E_{z_i})$ stands just for the homology of E_{z_i} with coefficients in Λ; we have made no assertion about the relationship of these maps with the ring structure of $QH_*(E_{z_i}, \Omega | E_{z_i})$. We call the maps Φ_k *relative Gromov-Witten invariants.* The simplest ones $(r = 0)$ are equal to the induced maps $\Phi(E, \Omega, J')$ which we have introduced before.

Notation. In the special case of a symplectic fibre bundle, there is no almost complex structure J' and we will denote the homomorphisms by $\Phi_r(E, \Omega; z_1, \dots, z_r)$. A similar convention (of simply omitting J' from the notation if there are no critical points) will also be followed in other occasions.

The relative Gromov-Witten invariants satisfy properties analogous to those of Φ:

(Duality) Let $(\overline{E}, \overline{\Omega}, \overline{J}')$ be the pullback of (E, Ω, J') by the involution $(s,t) \longmapsto (s_1 + s_0 - s, t)$ and $\bar{z}_i$ the preimage of z_i under this involution. Then

$$\langle a, \Phi_r(E, \Omega; z_1, \dots, z_r)(c_1, \dots, c_r, b)\rangle_{(E_{s_1 \times S^1}, \Omega_{s_1 \times S^1})} = \\ \langle \Phi_r(\overline{E}, \overline{\Omega}; \bar{z}_1, \dots, z_r)(c_1, \dots, c_r, a), b\rangle_{(E_{s_0 \times S^1}, \Omega_{s_0 \times S^1})}$$

for all $a \in HF_*(\overline{E}_{s_1 \times S^1}, \overline{\Omega}_{s_1 \times S^1})$, $b \in HF_*(E_{s_0 \times S^1}, \Omega_{s_0 \times S^1})$ and $c_i \in QH_*(E_{z_i}, \Omega | E_{z_i}) = QH_*(\overline{E}_{\bar{z}_i}, \overline{\Omega} | \overline{E}_{\bar{z}_i})$.

(Gluing) Divide $Z = [s_0; s_1] \times S^1$ into two parts $Z_- = [s_0; s] \times S^1$ and $Z_+ = [s; s_1] \times S^1$ for some $s \in (s_0; s_1)$. Let (E, Ω, J') be an ordinary almost holomorphic fibration over Z such that every point in $\{s\} \times S^1$ is a regular value of this fibration. We denote the restrictions of (E, Ω, J') to $Z_\pm$ by (E_-, Ω_-, J'_-) and (E_+, Ω_+, J'_+). Choose points $z_1, \dots, z_s \in Z_+$ and $z_{s+1}, \dots, z_r \in Z_-$. Then

$$\Phi_r(E, \Omega, J'; z_1, \dots, z_r)(c_1, \dots c_r, a) = \\ \Phi_s(E_+, \Omega_+, J'_+; z_1, \dots, z_s)(c_1, \dots, c_s, \\ \Phi_{r-s}(E_-, \Omega_-, J'_-; z_{s+1}, \dots, z_r)(c_{s+1}, \dots c_r, a)).$$

(Deformation invariance) This property has the same form as the corresponding one for Φ; we omit the exact statement.

In addition, the invariants Φ_r have two new properties which describe the effect of changing the marked points.

(Inserting a marked point)

$$\Phi_r(E, \Omega, J'; z_1, \dots, z_r) = \Phi_{r+1}(E, \Omega, J'; z_0, \dots z_r)(u, \dots),$$

where $u = [E_{z_0}]t^0 \in QH_*(E_{z_0}, \Omega | E_{z_0})$ is the 'fundamental class' of E_{z_0}.

(Moving the marked points) Let (E, Ω, J') be an ordinary almost holomorphic fibration over Z and $Z^{\text{reg}} \subset Z$ the set of its regular values. Take two sets $z_1, \dots, z_r \in Z^{\text{reg}}$ and $z'_1, \dots z'_r \in Z^{\text{reg}}$ of marked points, and choose paths $\gamma_1, \dots, \gamma_r : [0; 1] \longrightarrow Z^{\text{reg}}$ which connect z_i with z'_i (this is always possible since the complement of Z^{reg} is finite). Let $P_i : E_{z_i} \longrightarrow E_{z'_i}$ be the symplectic parallel transport along γ_i. This induces maps on the homology of the fibres with Λ-coefficients, which we denote by $(P_i)_* : QH_*(E_{z_i}, \Omega | E_{z_i}) \longrightarrow QH_*(E_{z'_i}, \Omega | E_{z'_i})$. Then

$$\Phi_r(E, \Omega, J'; z_1, \dots, z_r)(c_1, \dots, c_r, a) = \\ \Phi_r(E, \Omega, J'; z'_1, \dots, z'_r)((P_1)_*(c_1), \dots, (P_r)_*(c_r), a).$$

The map $(P_i)_*$ depends on the path γ_i; paths which are not isotopic in Z^{reg} may lead to different induced maps. The 'moving the marked points' property is valid for any choice of paths and hence imposes a non-trivial restriction on the homomorphisms Φ_r. We will return to this later on in a special case.

We will now define the quantum module structure on $HF_*(\phi)$ in terms of the relative invariants. Let ϕ be an automorphism of (M, ω), (T_ϕ, Θ_ϕ) its mapping torus and $(E_\phi, \Omega_\phi) = [0; 1] \times (T_\phi, \Theta_\phi)$. Recall that T_ϕ is a quotient of $\mathbb{R} \times M$. The inclusion $M = \{0\} \times M \subset \mathbb{R} \times M$ and the projection $\mathbb{R} \times M \longrightarrow T_\phi$ determine a symplectic isomorphism $F_\phi : M \longrightarrow (T_\phi)_0$, where $(T_\phi)_0$ denotes the fibre over $0 \in S^1$. Let $z_0 = (0, 0) \in [0; 1] \times S^1$. Clearly $(E_\phi)_{z_0} = (T_\phi)_0$. The quantum module product is defined by

$$x \,\hat{*}\, y = \Phi_1(E_\phi, \Omega_\phi; z_0)((F_\phi)_*(x), y)$$

for $x \in QH_*(M, \omega)$ and $y \in HF_*(\phi)$. It is not at all obvious that this makes $HF_*(\phi)$ into a module over $(QH_*(M, \omega), *)$; the fact that this is true, or more concretely, the equality

$$x_1 \,\hat{*}(x_2 \,\hat{*}\, y) = (x_1 * x_2) \,\hat{*}\, y$$

is a theorem of Piunikhin, Salamon, and Schwarz. We will not attempt to explain the proof of this theorem since our present framework is not suitable for that. The fact that $HF_*(\phi)$ is a unital $QH_*(M, \omega)$-module, that is, $[M]t^0 \,\hat{*}\, x = x$ for all x, is much simpler to prove; it follows from the formula for inserting a new marked point and the fact that the Φ-homomorphism of a product bundle is the identity map.

Using the properties stated above, it can be proved (roughly speaking) that all the maps $\Phi(E, \Omega, J')$ are homomorphisms of modules over $QH_*(M, \omega)$. A precise statement is

Proposition 10.1. *Let (E, Ω, J') be an ordinary almost holomorphic fibration over $Z = [s_0; s_1] \times S^1$ whose regular fibres are isomorphic to (M, ω). Let ϕ_0, ϕ_1 be automorphisms of (M, ω), and assume that we have fixed isomorphisms of their mapping tori with the boundary components of E:*

$$(T_{\phi_i}, \Theta_{\phi_i}) = (E_{s_i \times S^1}, \Omega_{s_i \times S^1}). \tag{10.2}$$

Let $P : E_{(s_0,0)} \longrightarrow E_{(s_1,0)}$ be the symplectic parallel transport along any path in Z from $(s_0, 0)$ to $(s_1, 0)$ which lies outside the set of critical values of $E \longrightarrow Z$. We denote the symplectic automorphism

$$M \xrightarrow{F_{\phi_0}} (T_{\phi_0})_0 = E_{(s_0,0)} \xrightarrow{P} E_{(s_1,0)} = (T_{\phi_1})_0 \xrightarrow{F_{\phi_1}^{-1}} M$$

by P'. Then

$$\Phi(E, \Omega, J')(x \,\hat{*}\, y) = P'_*(x) \,\hat{*}\, \Phi(E, \Omega, J')(y)$$

for all $x \in QH_*(M, \omega)$ and $y \in HF_*(\phi_0)$.

Sketch of the proof. Using the deformation equivalence of $\Phi(E, \Omega, J')$ one can reduce the statement to the case where (E, Ω, J') has the following property: there is an $\epsilon > 0$ such that the restriction of (E, Ω) to $[s_0; s_0 + \epsilon] \times S^1$ is isomorphic to $[s_0; s_0 + \epsilon] \times T_{\phi_0}$, and such that its restriction to $[s_1 - \epsilon; s_1] \times S^1$ is isomorphic to $[s_1 - \epsilon; s_1] \times T_{\phi_1}$; both isomorphisms will be extensions of those in (10.2). We divide (E, Ω, J') into the part (E_-, Ω_-) lying over $[s_0; s_0 + \epsilon]$, the part (E_+, Ω_+) lying over $[s_1 - \epsilon; s_1]$, and the part (E_0, Ω_0, J'_0) over $[s_0 + \epsilon; s_1 - \epsilon]$ which contains all the critical points of the fibration. $(E_\pm, \Omega_\pm)$ are product fibre bundles, and therefore $\Phi(E_\pm, \Omega_\pm) = \mathrm{id}$. Because of the gluing property of Φ, this implies that

$$\begin{aligned}\Phi(E, \Omega, J') &= \Phi(E_+, \Omega_+) \circ \Phi(E_0, \Omega_0, J'_0) \circ \Phi(E_-, \Omega_-) \\ &= \Phi(E_0, \Omega_0, J'_0).\end{aligned}$$

Let $z_0 = (s_0, 0) \in Z$. Using the gluing property of Φ_1, one sees that

$$\begin{aligned}\Phi(E, \Omega, J')\Phi_1(E_-, \Omega_-; z_0)(x, y) = \\ &= \Phi(E_0, \Omega_0, J'_0)\Phi_1(E_-, \Omega_-; z_0)(x, y) \\ &= \Phi(E_+, \Omega_+)\Phi(E_0, \Omega_0, J'_0)\Phi_1(E_-, \Omega_-; z_0)(x, y) \quad (10.3) \\ &= \Phi_1(E, \Omega, J'; z_0)(x, y).\end{aligned}$$

By moving the marked point to $z_1 = (s_1, 0)$ one obtains

$$\Phi_1(E, \Omega, J'; z_0)(x, y) = \Phi_1(E, \Omega, J'; z_1)(P_*(x), y).$$

Reversing the reasoning of (10.3) with z_1 instead of z_0 leads to the equation

$$\Phi_1(E, \Omega, J'; z_1)(P_*(x), y) = \Phi_1(E_+, \Omega_+; z_1)(P_*(x), \Phi(E, \Omega, J')(y));$$

therefore

$$\Phi(E, \Omega, J')\Phi_1(E_-, \Omega_-; z_0)(x, y) = \Phi_1(E_+, \Omega_+; z_1)(P_*(x), \Phi(E, \Omega, J')(y)).$$

Because (E_-, Ω_-) and (E_+, Ω_+) are product bundles, $\Phi_1(E_-, \Omega_-; z_0)(x, \cdot)$ is the quantum module product with $(F_{\phi_0}^{-1})_*(x)$, and $\Phi_1(E_+, \Omega_+; z_1)(P_*(x), \cdot)$ is the quantum module product with $(F_{\phi_1}^{-1})_*(P_*(x))$. This completes the proof. □

One application of Proposition 10.1 is to the fibre bundles $(E_\phi, \Omega_{\phi,K})$ which were defined in section 7. Because the action of the parallel transport P on homology does not depend on Ω at all, it is easy to see that the isomorphisms

$$C(\phi, K_0, K_1) : HF_*(\phi \circ \phi_1^{K_0}) \longrightarrow HF_*(\phi \circ \phi_1^{K_1})$$

defined by these fibre bundles are isomorphisms of $QH_*(M, \omega)$-modules. This has the important consequence (stated as one of the basic properties of Floer homology in Part I) that the Floer homology groups of two symplectically isotopic automorphisms of (M, ω) are isomorphic as $QH_*(M, \omega)$-modules.

Proposition 10.2. *The quantum product on $HF_*(\phi)$ satisfies*

$$x \mathbin{\hat{*}} y = \phi_*(x) \mathbin{\hat{*}} y$$

for all $x \in QH_*(M, \omega)$ and $y \in HF_*(\phi)$.

Sketch of the proof. Let (E_ϕ, Ω_ϕ) be the product bundle $[0;1] \times (T_\phi, \Theta_\phi)$ and $F_\phi : M \longrightarrow (E_\phi)_{(0,0)}$ the isomorphism defined above. Let

$$P : (E_\phi)_{(0,0)} \longrightarrow (E_\phi)_{(0,0)}$$

be the symplectic monodromy around the loop $\{0\} \times S^1$. By moving the marked point $z_0 = (0,0)$ around this loop one obtains

$$\begin{aligned} x \mathbin{\hat{*}} y = \Phi_1(E_\phi, \Omega_\phi; z_0)((F_\phi)_*(x), y) = \\ \Phi_1(E_\phi, \Omega_\phi; z_0)((P \circ F_\phi)_*(x), y) = (F_\phi^{-1} \circ P \circ F_\phi)_*(x) \mathbin{\hat{*}} y. \end{aligned}$$

It is clear from the definition of T_ϕ that $F_\phi^{-1} \circ P \circ F_\phi = \phi$. □

We will now briefly outline the definition of the invariants Φ_r in the framework of section 9. For simplicity, we consider only the case $r = 1$. Let (E, Ω, J') be an ordinary almost holomorphic fibration over $C = \mathbb{R} \times S^1$ with tubular ends modelled on nondegenerate symplectic fibre bundles $(T^\pm, \Theta^\pm)$. Choose $J^\pm \in \mathcal{J}_{\mathrm{reg}}(T^\pm, \Theta^\pm)$ and a $J \in \mathcal{J}(E, \Omega, J'; J^-, J^+)$. Let $z \in C$ be a regular value of $E \longrightarrow C$ and $X \subset E_z$ a 'cycle' representing some d-dimensional mod 2 homology class in E_z. There are several possibilities for the kinds of 'cycles' one can use; the simplest method (due to Schwarz) seems to be to use the unstable manifolds of a Morse function as cycles. Here, to simplify matters, we will assume that X is an embedded submanifold. For $\nu^\pm \in \mathcal{H}(T^\pm, \Theta^\pm)$, $\epsilon \in \mathbb{R}$ and $k \in \mathbb{Z}$, let

$$\mathcal{M}_{k,\epsilon}(E, J; \nu^-, \nu^+, X) \subset \mathcal{M}_{k,\epsilon}(E, J; \nu^-, \nu^+)$$

be the subset of sections σ such that $\sigma(z) \in X$. This subset is the preimage of X under the evaluation map

$$\mathrm{ev}_z : \mathcal{M}_{k,\epsilon}(E, J; \nu^-, \nu^+) \longrightarrow E_z.$$

Therefore it is plausible that (at least in a generic situation) it should be a submanifold of $\mathcal{M}_{k,\epsilon}(E, J; \nu^-, \nu^+)$ of codimension $4 - d$. In turns out that this is true: moreover, for generic J, the spaces $\mathcal{M}_{4-d,\epsilon}(E, J; \nu^-, \nu^+, X)$ are finite and the formal sum

$$p(\nu^-, \nu^+) = \sum_\epsilon \#\mathcal{M}_{4-d,\epsilon}(E, J; \nu^-, \nu^+, X) t^\epsilon$$

(# denotes the number of points mod 2) is an element of Λ. Let $m : CF_*(T^-, \Theta^-) \longrightarrow CF_*(T^+, \Theta^+)$ be the homomorphism given by

$$m(<\nu^->) = \sum_{\nu^+} p(\nu^-, \nu^+) <\nu^+> .$$

For generic J, m is a chain homomorphism. The induced homomorphism of Floer homology groups defines $\Phi_1(E, \Omega, J'; z)([X]t^0, \cdot)$. Of course it is necessary to show that this is independent of J and of the cycle representing $[X]$.

Bibliographical note. The 'quantum module structure' appears already in Floer's work [9]. It has been used by LeHong-Ono [15] and by Schwarz [26] in connection with the Arnol'd conjecture. The paper of Schwarz is the first one in which the product is defined for a broad class of symplectic manifolds. The fact that this product makes Floer homology into a module over the quantum homology ring was proved by Piunikhin, Salamon and Schwarz [22]. We have borrowed their idea that the quantum module product is a special case of a more general 'relative invariant'. [22] also describes the relationship between the quantum module product and various other definitions of products on Floer homology.

We have now completed our survey of Floer homology and of its functorial and product structures. To carry out the details of this construction and prove the properties we have stated is a major *tour de force* in nonlinear analysis. The principal steps of this programme are:

(1) Theorem 9.1, which summarizes the analytic results underlying the definition of the Floer chain complex.

(2) The analogous result (Theorem 9.4) used to define the induced maps.

(3) The gluing theorem and the invariance under a change of the almost complex structure (Theorems 9.6 and 9.7). These are not only basic properties of Floer homology, but (as explained in section 9) they are necessary even to make the Floer homology groups independent of the choice of almost complex structures.

(4) The construction of the relative Gromov-Witten invariants Φ_r and the proof of their basic properties stated above.

As has been mentioned before, some of these topics are alrady covered in the literature. Usually, the framework is that of Floer homology for symplectic automorphisms which are Hamiltonian perturbations of the identity, but the proofs remain essentially the same in the general framework. In particular, (1) is an adaptation of the work of Hofer and Salamon [14]. The other items are known in the case of symplectic fibre bundles; what is new is the generalization to almost holomorphic fibrations. However, a large part of the construction is not affected by this generalization. To be precise, a new problem appears only at one point, namely, when studying the compactification of the moduli spaces $\mathcal{M}_{k,\epsilon}(E, J; \nu^-, \nu^+)$. This new problem is the deformation theory for pseudo-holomorphic spheres which lie in a singular fibre of E: since this fibre is not a symplectic manifold, the usual transversality theory does not apply. This problem occurs at each of the steps (2)–(4) and it can be solved in the same way each time. For this reason we will not even attempt to discuss all of the four steps listed above. We will concentrate instead on (2), that is, the proof of Theorem 9.4, which is the most basic one. Even there, we will only outline the greatest part of the argument, and only the part which contains the new problem described above will be treated in detail.

To prove Theorem 9.4, we proceed as follows: first we define the subset $\mathcal{J}_{\mathrm{reg}}(E, \Omega, J'; J^-, J^+) \subset \mathcal{J}(E, \Omega, J'; J^-, J^+)$; this is done in the next section. Then we prove that any almost complex structure in this subset has the properties described in that Theorem; the arguments which yield these properties are outlined in section 12. Finally we prove that $\mathcal{J}_{\mathrm{reg}}(E, \Omega, J'; J^-, J^+)$ is dense. This is the part which we will discuss in detail, and it occupies sections 13–15.

11 Regular J-holomorphic sections

In this section we review the local properties of the spaces of J-holomorphic sections with horizontal limits. The most important concept is that of a regular J-holomorphic section. Having introduced that we define the sets $\mathcal{J}_{\mathrm{reg}}(T, \Theta)$ and $\mathcal{J}_{\mathrm{reg}}(E, \Omega, J'; J^-, J^+)$ which occur in Theorem 9.1 and Theorem 9.4. Throughout this section, (E, Ω, J') denotes an almost holomorphic fibration over C with tubular ends modelled on nondegenerate symplectic fibre bundles $(T^\pm, \Theta^\pm)$, with projection $\pi : E \longrightarrow C$.

Let J be a partially Ω-tame almost complex structure on E. There is a differential operator

$$D\bar{\partial}_J(\sigma) : C^\infty(\sigma^* TE^v) \longrightarrow C^\infty(\sigma^* TE^v)$$

canonically associated to any J-holomorphic section σ. We will give three equivalent definitions of this operator. The first one is based on the fact that a smooth section τ of E is J-holomorphic iff the expression

$$\bar{\partial}_J(\tau) = \frac{\partial \tau}{\partial s} + J(\tau) \frac{\partial \tau}{\partial t}$$

vanishes. Because π is (J, j)-linear, $D\pi(J \frac{\partial \tau}{\partial t}) = -D\pi(\frac{\partial \tau}{\partial s})$; therefore $\bar{\partial}_J(\tau)$ is a section of the pullback bundle $\tau^* TE^v$.

Definition 11.1. Let σ be a J-holomorphic section of E and X a compactly supported smooth section of the vector bundle $\sigma^*TE^v \longrightarrow C$. Choose a smooth family $(\sigma_r)_{0\leq r<\epsilon}$ of sections of E with $\sigma_0 = \sigma$ and such that $\partial\sigma_r/\partial r = X$ at $r = 0$. Then

$$D\bar{\partial}_J(\tau)X = \frac{\partial}{\partial r}\left[\bar{\partial}_J(\sigma_r)\right],$$

where the derivative is taken at $r = 0$.

To be precise, this definition should be worded as follows: the family (σ_r) defines a map $\Sigma : [0;\epsilon) \times \mathbb{R} \times S^1 \longrightarrow E$. $(\bar{\partial}_J(\sigma_r))_{0\leq r<\epsilon}$ is a section of Σ^*TE^v which vanishes at any point (r, s, t) with $r = 0$. Therefore we can form the derivative of this section in r-direction at such points (without any choice of connection). It is not obvious from Definition 11.1 that $D\bar{\partial}_J(\tau)X$ is independent of the choice of (σ_r). To prove this, choose a Riemannian metric on E and let ∇ be its Levi-Civita connection. Because this connection is torsion-free and $[\partial_r\sigma_r, \partial_s\sigma_r] = [\partial_r\sigma_r, \partial_t\sigma_r] = 0$, we have

$$\begin{aligned}
D\bar{\partial}_J(\sigma)X &= \frac{\partial}{\partial r}\left(\frac{\partial\sigma_r}{\partial s} + J(\sigma_r)\frac{\partial\sigma_r}{\partial t}\right)\\
&= \nabla_X\left(\frac{\partial\sigma_r}{\partial s}\right) + J(\sigma_r)\nabla_X\left(\frac{\partial\sigma_r}{\partial t}\right) + (\nabla_X J)\frac{\partial\sigma}{\partial t}\\
&= \nabla_{\frac{\partial\sigma}{\partial s}}X + J(\sigma)\nabla_{\frac{\partial\sigma}{\partial t}}X + (\nabla_X J)\frac{\partial\sigma}{\partial t}. \qquad (11.1)
\end{aligned}$$

This proves that $D\bar{\partial}_J(\sigma)X$ is independent of the choice of σ_r and also that $D\bar{\partial}_J(\sigma)$ is a differential operator. Alternatively, equation (11.1) can be used as the definition of $D\bar{\partial}_J(\sigma)$, and then the argument above shows that it is independent of the choice of ∇ and that the r.h.s. of (11.1) is a section of $\sigma^*TE^v \subset \sigma^*TE$ (this is not obvious because ∇ may not preserve TE^v).

The third definition of $D\bar{\partial}_J(\sigma)$ is this: choose $\widetilde{S} \in C^\infty(TE)$ and $\widetilde{X} \in C^\infty(TE^v)$ which extend the vector fields $\partial\sigma/\partial s$ and X defined along σ. Using the torsion-freeness of ∇ one sees that

$$\nabla_{\frac{\partial\sigma}{\partial s}}X = \nabla_{\widetilde{X}}\widetilde{S} + [\widetilde{S}, \widetilde{X}]$$

and

$$\begin{aligned}
J\,\nabla_{\frac{\partial\sigma}{\partial t}}X + (\nabla_X J)\frac{\partial\sigma}{\partial t} &= J\nabla_{\widetilde{X}}(J\widetilde{S}) + J[J\widetilde{S}, \widetilde{X}] + (\nabla_{\widetilde{X}}J)(J\widetilde{S})\\
&= J[J\widetilde{S}, \widetilde{X}] - \nabla_{\widetilde{X}}\widetilde{S}.
\end{aligned}$$

Together with (11.1) this means that

$$D\bar{\partial}_J(\sigma)X = [\widetilde{S}, \widetilde{X}] + J[J\widetilde{S}, \widetilde{X}]. \qquad (11.2)$$

The equivalence of this expression with the other two shows that the r.h.s. of (11.2) is independent of the choice of $\widetilde{S}$ and $\widetilde{X}$.

From now on we assume that $J \in \mathcal{J}(E, \Omega, J'; J^-, J^+)$ for some $J^\pm \in \mathcal{J}(T^\pm, \Theta^\pm)$. The next theorem summarizes the main properties of the operator $D\bar{\partial}_J(\sigma)$ for a J-holomorphic section σ with horizontal limits. Proofs can be found in [25] (with $p = 2$) and in [27] (general case).

Theorem 11.2. *Let σ be a J-holomorphic section with horizontal limits. For $p \geq 2$, let $L^p(\sigma^* TE^v)$ and $W^{1,p}(\sigma^* TE^v)$ be the spaces of sections of $\sigma^* TE^v$ of class L^p resp. $W^{1,p}$ (these spaces should be defined using a Riemannian metric on E which has product form on the tubular ends). Then $D\bar{\partial}_J(\sigma)$ defines a bounded operator $W^{1,p}(\sigma^* TE^v) \longrightarrow L^p(\sigma^* TE^v)$. This operator is Fredholm, and its index equals* $\mathrm{ind}(\sigma)$.

From now on, we will use some fixed $p > 2$.

Definition 11.3. Let σ be a J-holomorphic section with horizontal limits. σ is called *regular* if the operator

$$D\bar{\partial}_J(\sigma) : W^{1,p}(\sigma^* TE^v) \longrightarrow L^p(\sigma^* TE^v)$$

is onto.

We will now outline the framework which was introduced by Floer to study J-holomorphic sections using the tools of nonlinear analysis. Let σ be a section of E of class $W^{1,p}_{\mathrm{loc}}$, ν^+ a horizontal section of (T^+, Θ^+), and $\sigma^+(s,t) = (s, \nu^+(t))$ the corresponding partial section of E (defined for $s \gg 0$). We say that σ is $W^{1,p}$-convergent to ν^+ if there is a vector field $\xi^+ \in W^{1,p}((\sigma^+)^* TE^v)$ such that

$$\sigma(s,t) = \exp_{\sigma^+(s,t)}(\xi(s,t))$$

for large s. As before, exp and $W^{1,p}$ are defined using a Riemannian metric which has product form on the tubular ends. $W^{1,p}$-convergence to a horizontal section of (T^-, Θ^-) as $s \to -\infty$ is defined in the same way. The space of $W^{1,p}_{\mathrm{loc}}$-sections σ which converge to $\nu^\pm \in \mathcal{H}(T^\pm, \Theta^\pm)$ in this sense and such that $\mathrm{im}(\sigma) \cap \mathrm{Crit}(\pi) = \emptyset$ will be denoted by $\mathcal{S}(E; \nu^-, \nu^+)$.

Remark. In contrast to the case of a smooth section, it is not clear whether a $W^{1,p}_{\mathrm{loc}}$-section σ automatically satisfies $\mathrm{im}(\sigma) \cap \mathrm{Crit}(\pi) = \emptyset$. This is the reason why we have included this condition in our definition of $\mathcal{S}(E; \nu^-, \nu^+)$.

$\mathcal{S}(E; \nu^-, \nu^+)$ carries a natural topology and the structure of a smooth Banach manifold compatible with this topology. The tangent space at a point $\sigma \in \mathcal{S}(E; \nu^-, \nu^+)$ is canonically isomorphic to the space $W^{1,p}(\sigma^* TE^v)$ of sections of $\sigma^* TE^v$ of class $W^{1,p}$.

Remark. The space $W^{1,p}(\sigma^*TE^v)$ is well-defined for any $\sigma \in \mathcal{S}(E;\nu^-,\nu^+)$ because the pullback σ^*TE^v is a vector bundle over C of class $W^{1,p}_{\mathrm{loc}}$. More precisely, a family of smooth local trivializations of TE^v induces a family of trivializations of σ^*TE^v whose transition functions lie in $W^{1,p}_{\mathrm{loc}}$.

Like $W^{1,p}(\sigma^*TE^v)$, the space $L^p(\sigma^*TE^v)$ of L^p-sections is also well-defined for any $\sigma \in \mathcal{S}(E;\nu^-,\nu^+)$. There is a Banach vector bundle

$$\mathcal{E} \longrightarrow \mathcal{S}(E;\nu^-,\nu^+)$$

with fibres $\mathcal{E}_\sigma = L^p(\sigma^*TE^v)$. Every $J \in \mathcal{J}(E,\Omega,J';J^-,J^+)$ determines a smooth section $\bar\partial_J$ of $\mathcal{E}$ given by

$$\bar\partial_J(\sigma) = \frac{\partial\sigma}{\partial s} + J(\sigma)\frac{\partial\sigma}{\partial t}.$$

Theorem 11.4. *$\bar\partial_J^{-1}(0) = \mathcal{M}(E,J;\nu^-,\nu^+)$.*

This statement is actually a combination of two results: an elliptic regularity result which shows that any $\sigma \in \mathcal{S}(E;\nu^-,\nu^+)$ with $\bar\partial_J(\sigma) = 0$ is smooth, and a (more difficult) exponential convergence result which shows that any such σ converges to its horizontal limits in the stricter sense of section 8. This shows that $\bar\partial_J^{-1}(0) \subset \mathcal{M}(E,J;\nu^-,\nu^+)$; the converse is obvious.

The derivative of the section $\bar\partial_J$ at a point $\sigma \in \bar\partial_J^{-1}(0)$ is a homomorphism

$$T_\sigma\mathcal{S}(E;\nu^+,\nu^-) \longrightarrow \mathcal{E}_\sigma.$$

$T_\sigma\mathcal{S}(E;\nu^+,\nu^-) = W^{1,p}(\sigma^*TE^v)$ and $\mathcal{E}_\sigma = L^p(\sigma^*TE^v)$, and using our first definition it is not difficult to see that this homomorphism is given by the operator $D\bar\partial_J(\sigma)$. Using the implicit function theorem we obtain the following consequence:

Corollary 11.5. *A regular $\sigma \in \mathcal{M}(E,J;\nu^-,\nu^+)$ has a neighbourhood in $\mathcal{M}(E,J;\nu^-,\nu^+) \subset \mathcal{S}(E;\nu^-,\nu^+)$ which is a smooth submanifold of dimension $\mathrm{ind}(\sigma)$. In particular, if for some k, all $\sigma \in \mathcal{M}_k(E,J;\nu^-,\nu^+)$ are regular, then $\mathcal{M}_k(E,J;\nu^-,\nu^+)$ is a smooth k-dimensional submanifold of $\mathcal{S}(E;\nu^-,\nu^+)$.*

For a product bundle $(E,\Omega) = \mathbb{R}\times(T,\Theta)$, the spaces $\mathcal{S}(E;\nu^-,\nu^+)$ have a natural smooth $\mathbb{R}$-action by translation. This action preserves the subsets $\mathcal{M}_{k,\epsilon}(E,J;\nu^-,\nu^+)$ for any $J \in \mathcal{J}(T,\Theta)$.

Lemma 11.6. *Let $J \in \mathcal{J}(T,\Theta)$ and assume that all J-holomorphic sections of $\mathbb{R}\times T$ with horizontal limits are regular. Then the following conditions for $\sigma \in \mathcal{M}(\mathbb{R}\times T,J)$ are equivalent:*

(i) $\mathrm{ind}(\sigma) = 0$;

(ii) $e(\sigma) = 0$;
(iii) σ *is horizontal;*
(iv) σ *has a nontrivial stabilizer under the* $\mathbb{R}$*-action.*

Proof. The equivalence of the final three conditions is part of Lemma 8.8. The description of a horizontal section given in Lemma 8.2:

$$\sigma(s,t) = (s, \nu(t)),$$

implies that the index of such a section vanishes. Conversely, let σ be a J-holomorphic section of index 0. It is an isolated point of $\mathcal{M}(\mathbb{R} \times T, J)$; therefore it must be equal to its translate σ^r for sufficiently small r. □

Corollary 11.7. *If all* $\sigma \in \mathcal{M}(\mathbb{R} \times T, J)$ *are regular,* $\mathcal{M}_{k,\epsilon}(\mathbb{R} \times T, J) = \emptyset$ *for all* $\epsilon > 0$ *and* $k \leq 0$.

It is not difficult to see that outside the subset of horizontal sections in $\mathcal{M}(\mathbb{R}\times T, J)$ the $\mathbb{R}$-action is not only free but also proper. As a consequence, one has

Corollary 11.8. *Assume that for some* $\nu^-, \nu^+ \in \mathcal{H}(T, \Theta)$ *and* ϵ, k *with* $\epsilon \neq 0$ *or* $k \neq 0$*, any* $\sigma \in \mathcal{M}_{k,\epsilon}(\mathbb{R} \times T, J; \nu^-, \nu^+)$ *is regular. Then the quotient*

$$\mathcal{M}_{k,\epsilon}(\mathbb{R} \times T, J; \nu^-, \nu^+)/\mathbb{R}$$

is a smooth $(k-1)$*-dimensional manifold.*

In order to define $\mathcal{J}_{\mathrm{reg}}(E, \Omega, J'; J^-, J^+)$ we need to introduce a different class of pseudo-holomorphic curves in E.

Definition 11.9. Let J be an Ω-tame almost complex structure on E. A smooth nonconstant map $w : \mathbb{C}\mathrm{P}^1 \longrightarrow E$ is called a *J-bubble* if $Dw \circ i = J \circ Dw$ (i denotes the complex structure on $\mathbb{C}\mathrm{P}^1$) and $\mathrm{im}(w)$ lies in a single fibre of π.

The numbers

$$\Omega(w) = \int_{\mathbb{C}\mathrm{P}^1} w^*\Omega \quad \text{and} \quad c_1(w) = \langle c_1(TE, J), [w] \rangle$$

are called the energy and the Chern number of the bubble w. We denote the set of points which lie on the image of a J-bubble with Chern number $\leq k$ by $V_k(J) \subset E$. The importance of J-bubbles for the compactification of the space $\mathcal{M}(E, J)$ will be explained in the next section. We can now supply the definitions of $\mathcal{J}_{\mathrm{reg}}(T, \Theta)$ and $\mathcal{J}_{\mathrm{reg}}(E, \Omega, J'; J^-, J^+)$.

Definition 11.10. Let (T, Θ) be a nondegenerate symplectic fibre bundle over S^1. $\mathcal{J}_{\mathrm{reg}}(T, \Theta) \subset \mathcal{J}(T, \Theta)$ is the subset of almost complex structures J such that

(1) any $\sigma \in \mathcal{M}(\mathbb{R} \times T, J)$ is regular.

(2) $V_0(J) = \emptyset$ (this means that there are no J-bubbles with nonpositive Chern number).

(3) Any horizontal section σ of $\mathbb{R} \times T$ satisfies $\mathrm{im}(\sigma) \cap V_1(J) = \emptyset$.

Definition 11.11. Let (E, Ω, J') be an ordinary almost holomorphic fibration over C, with tubular ends modelled on nondegenerate fibre bundles $(T^{\pm}, \Theta^{\pm})$. For $J^- \in \mathcal{J}_{\mathrm{reg}}(T^-, \Theta^-)$ and $J^+ \in \mathcal{J}_{\mathrm{reg}}(T^+, \Theta^+)$,

$$\mathcal{J}_{\mathrm{reg}}(E, \Omega, J'; J^-, J^+) \subset \mathcal{J}(E, \Omega, J'; J^-, J^+)$$

is the subset of almost complex structures J such that

(1) any $\sigma \in \mathcal{M}(E, J)$ is regular.

(2) $V_{-1}(J) = \emptyset$ (this means that there are no J-bubbles with negative Chern number).

(3) Any $\sigma \in \mathcal{M}_0(E, J) \cup \mathcal{M}_1(E, J)$ satisfies $\mathrm{im}(\sigma) \cap V_0(J) = \emptyset$.

12 Compactness

In this section we state Gromov's compactness theorem and Floer's gluing theorem for J-holomorphic sections. Both theorems are familiar in the case of symplectic fibre bundles, and they carry over unchanged to almost holomorphic fibrations. Then we describe the application of these results to Theorem 9.4.

The central notion in this section is the *geometric convergence* of a sequence of J-holomorphic sections to a *broken J-holomorphic section*. We fix the following notation: (E, Ω, J') is an almost holomorphic fibration over C with projection $\pi : E \longrightarrow C$. It has tubular ends modelled on nondegenerate symplectic fibre bundles $(T^{\pm}, \Theta^{\pm})$. The product bundles $\mathbb{R} \times (T^{\pm}, \Theta^{\pm})$ will be denoted by $(E^{\pm}, \Omega^{\pm})$. J^- and J^+ are almost complex structures in $\mathcal{J}(T^-, \Theta^-)$ and $\mathcal{J}(T^+, \Theta^+)$, respectively, and J is an almost complex structure in $\mathcal{J}(E, \Omega, J'; J^-, J^+)$.

Definition 12.1. A broken J-holomorphic section

$$\hat{\sigma} = (\sigma_1^-, \dots, \sigma_m^-, \sigma, \sigma_1^+, \dots, \sigma_n^+)$$

consists of $\sigma_1^-, \dots, \sigma_m^- \in \mathcal{M}(E^-, J^-)$, $\sigma \in \mathcal{M}(E, J)$ and $\sigma_1^+, \dots, \sigma_n^+ \in \mathcal{M}(E^+, J^+)$, which have the following property: there are $\nu_0^-, \dots \nu_m^- \in$

$\mathcal{H}(T^-, \Theta^-)$ and $\nu_0^+, \dots \nu_n^+ \in \mathcal{H}(T^+, \Theta^+)$ such that

$$\begin{aligned}\sigma_j^- &\in \mathcal{M}(E^-, J^-; \nu_{j-1}^-, \nu_j^-) \text{ for } j = 1, \dots, m,\\ \sigma &\in \mathcal{M}(E, J; \nu_m^-, \nu_0^+), \text{ and}\\ \sigma_j^+ &\in \mathcal{M}(E^+, J^+; \nu_{j-1}^+, \nu_j^+) \text{ for } j = 1, \dots, n.\end{aligned}$$

Moreover, we assume that none of the $\sigma_j^\pm$ is a horizontal section. Two broken J-holomorphic sections will be regarded as identical if one arises from the other by translating the components $\sigma_j^\pm$ (by an amount which may be different for each component).

The basic notions of the theory of J-holomorphic sections can be extended to broken sections: ν_0^- and ν_n^+ as above are called the horizontal limits of $\hat{\sigma}$, and its energy (or index) is defined by adding the energies (or indices) of all components. We will denote the set of broken J-holomorphic sections with $m + n + 1$ components as above by $\widehat{\mathcal{M}}^{m,n}(E, J)$, and the subset of broken J-holomorphic sections with given index k, energy ϵ and limits $\nu^\pm$ by $\widehat{\mathcal{M}}_{k,\epsilon}^{m,n}(E, J; \nu^-, \nu^+)$. A broken J-holomorphic section $\hat{\sigma}$ is called regular if σ is a regular J-holomorphic section and the $\sigma_j^\pm$ are regular $J^\pm$-holomorphic sections.

Let $\sigma_1, \sigma_2, \dots$ be a sequence of J-holomorphic sections of E and $r_1, r_2, \dots$ a sequence of real numbers with $\lim_j r_j = -\infty$. By assumption there is an $R > 0$ such that $E|[R; \infty) \times S^1 = E^+|[R; \infty) \times S^1$. Therefore the restriction

$$\sigma_j^{r_j}|[R + r_j; \infty) \times S^1 \tag{12.1}$$

of the translate $\sigma_j^{r_j}$ can be viewed as a section of $E^+|[R + r_j; \infty) \times S^1$. Since $r_j \to -\infty$, (12.1) is a sequence of partial sections of E^+ defined over increasingly larger subsets. Therefore it makes sense to say that $\sigma_j^{r_j}$ converges on compact subsets to a section σ^+ of E^+ (the limit must be J^+-holomorphic). For a sequence (r_j) with limit ∞, the translates $\sigma_j^{r_j}$ can converge on compact subsets, in the same sense, to a J^--holomorphic section of E^-.

Definition 12.2. Let $\sigma_1, \sigma_2, \dots$ be a sequence in $\mathcal{M}(E, J)$. We say that this sequence converges geometrically to a broken J-holomorphic section

$$\hat{\sigma} = (\sigma_1^-, \dots, \sigma_m^-, \sigma, \sigma_1^+, \dots, \sigma_n^+)$$

if there are numbers $r_{j,k}^-$ ($j \in \mathbb{N}$, $k = 1, \dots, m$) and $r_{j,l}^+$ ($j \in \mathbb{N}$, $l = 1, \dots, n$), such that

(1) $\lim_j r_{j,k}^- = \infty$ and $\lim_j r_{j,l}^+ = -\infty$ for all k, l. Moreover,

$$\begin{aligned}&\lim_j \left(r_{j,k-1}^- - r_{j,k}^-\right) = \infty \text{ and}\\ &\lim_j \left(r_{j,l-1}^+ - r_{j,l}^+\right) = \infty.\end{aligned} \tag{12.2}$$

(2) σ_j converges to σ on compact subsets in the C^∞-sense. For every $k = 1, \dots, m$ the translates $\sigma_j^{r_{j,k}}$ converge to σ_k^- on compact subsets in C^∞ as $j \to \infty$. The translates $\sigma_j^{r_{j,l}}$ converge to σ_l^+ in the same sense.

(3) $e(\sigma_j) = e(\hat{\sigma})$ for all sufficiently large j.

The equation (12.2) serves to exclude counting the same component of the limit twice: it says that any two $\sigma_k^\pm$ express the behaviour of (σ_j) when translated at different rates. In contrast, condition (3) says that the broken J-holomorphic sections obtained from a geometric limit of (σ_k) by removing some components are no longer geometric limits of the same sequence (the reason is that the $\sigma_k^\pm$ may not be horizontal and any non-horizontal $J^\pm$-holomorphic section has positive energy, by Lemma 8.8). It is a non-obvious fact that these two conditions are sufficient to ensure that the geometric limit of a sequence (if it exists) is unique. Moreover, the geometric limit has the following important properties:

Proposition 12.3. *Assume that (σ_j) converges geometrically to a broken J-holomorphic section $\hat{\sigma}$. Then $\hat{\sigma}$ has the same limits and index as σ_j for all sufficiently large j.*

Proposition 12.4. *Let $\hat{\sigma}$ be a broken J-holomorphic section which is not 'really' broken: that is, $m = n = 0$ and $\hat{\sigma}$ is given by a single J-holomorphic section σ, with limits ν^-, ν^+. Then a sequence (σ_j) converges geometrically to $\hat{\sigma}$ iff converges to σ in the Banach manifold topology of $\mathcal{S}(E; \nu^-, \nu^+)$.*

We end our account of geometric convergence with a special case of Floer's 'gluing theorem' which describes how to attach broken J-holomorphic sections to $\mathcal{M}(E, J)$ as points at infinity.

Theorem 12.5. *Let*

$$\hat{\sigma} = (\sigma_1^-, \dots, \sigma_m^-, \sigma, \sigma_1^+, \dots, \sigma_n^+)$$

be a regular broken J-holomorphic section with with limits $\nu^\pm$, $\mathrm{ind}(\sigma) = 0$ and $\mathrm{ind}(\sigma_j^\pm) = 1$ for all j. There is a proper smooth embedding

$$\# : (0;1]^{m+n} \longrightarrow \mathcal{M}_{m+n}(E, J; \nu^-, \nu^+) \subset \mathcal{S}(E; \nu^-, \nu^+)$$

with the following property: if $p_1, p_2, \dots$ is a sequence in $(0;1]^{m+n}$ which converges to 0, $\#(p_j)$ converges geometrically to $\hat{\sigma}$. Conversely, if (σ_j) is any sequence in $\mathcal{M}(E, J)$ which converges geometrically to $\hat{\sigma}$, it is given by $\sigma_j = \#(p_j)$ for large j, with $\lim_j p_j = 0$.

Not every sequence in $\mathcal{M}(E, J)$ has a geometrically convergent subsequence. For instance, take a sequence (σ_j) with $e(\sigma_j) \to \infty$; then no subsequence can be geometrically convergent. More significantly, even sequences with bounded energy may not have a geometrically convergent subsequence. The behaviour of such sequences is described by the Gromov compactification of $\mathcal{M}(E, J)$, which is essentially the space of J-holomorphic cusp sections. These 'cusp sections' consist of a J-holomorphic section and a finite collection of J-bubbles, and their structure can be rather complicated. For our purpose, however, a much simplified version of the compactness theorem is sufficient.

Definition 12.6. J is called semi-positive if any $\sigma \in \mathcal{M}(E, J)$ has nonnegative index, any $\sigma^{\pm} \in \mathcal{M}(\mathbb{R} \times T^{\pm}, J^{\pm})$ has nonnegative index, and $V_{-1}(J) = \emptyset$ (that is, any J-bubble has nonnegative Chern number). Note that this implies that $V_{-1}(J^{\pm}) = \emptyset$ as well.

Theorem 12.7. *Let $J \in \mathcal{J}(E, \Omega, J'; J^-, J^+)$ be a semi-positive almost complex structure. Let (σ_j) be a sequence in $\mathcal{M}(E, J)$ with bounded energy. Then*

$$k = \overline{\lim}_j \operatorname{ind}(\sigma_j)$$

is finite and one of the following (not mutually exclusive) possibilities holds:

(1) (σ_j) has a geometrically convergent subsequence.

(2) There is a J-holomorphic section σ and a J-bubble w with

$$\operatorname{ind}(\sigma) + 2c_1(w) \leq k$$

and such that $\operatorname{im}(\sigma) \cap \operatorname{im}(w) \neq \emptyset$.

(3) There is a J^--holomorphic section σ^- of $\mathbb{R} \times T^-$ and a J^--bubble w with the same properties as in (2).

(4) The same as in (3) holds with T^- and J^- replaced by T^+ and J^+.

A few words about this Theorem are in order. The original compactness theorem of Gromov concerns compact pseudo-holomorphic curves in a symplectic manifold. The literature on Floer homology contains several versions of this theorem for maps from $\mathbb{R} \times S^1$ to a symplectic manifold which satisfy an 'inhomogeneous' version of the Cauchy-Riemann equation for pseudo-holomorphic curves. From our point of view these results correspond more or less to Theorem 12.7 in the case when (E, Ω) is a symplectic fibre bundle. Contrary to what one might suspect, the presence of critical points does not change the situation much. The essential idea is to work on the total space E, and to turn it into a symplectic manifold using the following familiar trick:

Lemma 12.8. *Let J be an almost complex structure in $\mathcal{J}(E, \Omega, J'; J^-, J^+)$ for some J^-, J^+, and $\beta = ds \wedge dt$ the standard volume form on C. There is a $c > 0$ such that $\Omega' = \Omega + c(\pi^*\beta)$ is a symplectic form on E which tames J.*

Proof. If the contrary were true there would be a sequence $X_1, X_2, \dots$ of nonzero tangent vectors on E and a sequence $c_1, c_2, \dots$ of positive numbers, with $\lim_i c_i = \infty$, such that

$$(\Omega + c_i(\pi^*\beta))\,(X_i, JX_i) = \Omega(X_i, JX_i) + c_i|D\pi(X_i)|^2 \leq 0$$

for all i. On the two tubular ends of E, $J = J^+$ or $J = J^-$, and any tangent vector Y in that region satisfies $\Omega(Y, JY) \geq 0$, with equality iff Y is horizontal (Lemma 8.7). Therefore none of the X_i can lie on one of the ends. After rescaling and passing to a subsequence we may now assume that the X_i converge to a nonzero $X \in TE$. In the limit, the inequality

$$|D\pi(X_i)|^2 \leq -c_i^{-1}\Omega(X_i, JX_i)$$

yields $|D\pi(X)|^2 \leq 0$, that is, $X \in TE^v$. On the other hand, $\Omega(X_i, JX_i) \leq 0$ for all i and therefore $\Omega(X, JX) \leq 0$, which is impossible because J is partially Ω-tame. □

In this way many questions about J-holomorphic sections can be reduced to the theory of pseudo-holomorphic curves in symplectic manifolds.

For the remainder of this section we assume that $J^\pm \in \mathcal{J}_{\mathrm{reg}}(T^\pm, \Theta^\pm)$ and $J \in \mathcal{J}_{\mathrm{reg}}(E, \Omega, J'; J^-, J^+)$. Using the results above and in the previous section, we will show that such a J has the properties stated in Theorem 9.4. The first step is

Lemma 12.9. *Every sequence (σ_j) in $\mathcal{M}(E, J)$ with bounded energy and such that* $\mathrm{ind}(\sigma_j) \leq 1$ *for all j has a geometrically convergent subsequence.*

Proof. We must exclude the other possibilities which occur in Theorem 12.7. The definitions of $\mathcal{J}_{\mathrm{reg}}(T^\pm, \Theta^\pm)$ and $\mathcal{J}_{\mathrm{reg}}(E, \Omega, J'; J^-, J^+)$ imply that J is semi-positive.

Assume first that case (2) in Theorem 12.7 holds, and let (σ, w) be a pair as stated there. Because σ is regular and $V_{-1}(J) = \emptyset$, $\mathrm{ind}(\sigma) \geq 0$ and $c_1(w) \geq 0$. Therefore $\mathrm{ind}(\sigma) + 2c_1(w) \leq 1$ can only be satisfied if $\mathrm{ind}(\sigma) = 0$ or 1 and $c_1(w) = 0$. However, condition (3) in Definition 11.11 excludes such pairs (σ, w).

If case (3) holds, the pair (σ^-, w) would satisfy $c_1(w) > 0$ (because $V_0(J^-) = \emptyset$) and therefore $\mathrm{ind}(\sigma^-) < 0$, which is impossible since σ^- is regular. The argument for case (4) is identical. □

Lemma 12.10. *Let $\sigma_1, \sigma_2, \dots$ be a sequence in $\mathcal{M}(E, J; \nu^-, \nu^+)$ such that $e(\sigma_j)$ is bounded and* $\mathrm{ind}(\sigma_j) = 0$ *for all j. Then (σ_j) has a subsequence which converges in the Banach manifold topology of $\mathcal{S}(E; \nu^-, \nu^+)$ to some J-holomorphic section in $\mathcal{M}_0(E, J; \nu^-, \nu^+)$.*

Proof. By Lemma 12.9, there is a subsequence $\sigma_{j_1}, \sigma_{j_2}, \dots$ which converges geometrically to a broken J-holomorphic section

$$\hat{\sigma} = (\sigma_1^-, \dots, \sigma_m^-, \sigma, \sigma_1^+, \dots, \sigma_n^+).$$

Since the $\sigma_j^\pm$ are regular and not horizontal, $\mathrm{ind}(\sigma_j^\pm) \geq 1$ by Corollary 11.7. Moreover, $\mathrm{ind}(\sigma) \geq 0$. On the other hand, Proposition 12.3 implies that

$$\sum_{j=1}^{m} \mathrm{ind}(\sigma_j^-) + \mathrm{ind}(\sigma) + \sum_{j=1}^{n} \mathrm{ind}(\sigma_j^+) = 0.$$

It follows that $m = n = 0$ and, by Proposition 12.4, that σ is the limit of (σ_{j_ν}) in the Banach manifold topology of $\mathcal{S}(E; \nu^-, \nu^+)$. □

This result implies the compactness of the subspace

$$\bigcup_{\epsilon \leq C} \mathcal{M}_{0,\epsilon}(E, J; \nu^- \nu^+) \subset \mathcal{S}(E; \nu^-, \nu^+) \tag{12.3}$$

for any $C > 0$ and $\nu^\pm \in \mathcal{H}(T^\pm, \Theta^\pm)$. Since this space is also discrete, it must be finite. Therefore J satisfies the first two parts of Theorem 9.4.

Lemma 12.11. *Every sequence (σ_j) in $\mathcal{M}_{1,\epsilon}(E, J; \nu^-, \nu^+)$ has a geometrically convergent subsequence whose limit lies in*

$$\mathcal{M}'_\epsilon(E, J; \nu^-, \nu^+) \stackrel{\text{def}}{=} \widehat{\mathcal{M}}^{1,0}_{1,\epsilon}(E, J; \nu^-, \nu^+) \cup \widehat{\mathcal{M}}^{0,1}_{1,\epsilon}(E, J; \nu^-, \nu^+).$$

This is proved by the same kind of argument as Lemma 12.10. We omit the details.

Lemma 12.12. *For all $\nu^\pm \in \mathcal{H}(T^\pm, \Theta^\pm)$ and $\epsilon \in \mathbb{R}$, $\mathcal{M}'(E, J; \nu^-, \nu^+)$ is a finite set.*

Proof. By definition, $\widehat{\mathcal{M}}^{1,0}_{1,\epsilon}(E, J; \nu^-, \nu^+)$ consists of all pairs

$$(\sigma^-, \sigma) \in \bigsqcup_{\nu} \mathcal{M}(E^-, J^-; \nu^-, \nu)/\mathbb{R} \times \mathcal{M}(E, J; \nu, \nu^+)$$

which satisfy $\mathrm{ind}(\sigma^-) + \mathrm{ind}(\sigma) = 1$ and $e(\sigma^-) + e(\sigma) = \epsilon$, and such that σ^- is not horizontal. The regularity of J^- and J implies that such a pair must have $\mathrm{ind}(\sigma^-) = 1$ and $\mathrm{ind}(\sigma) = 0$. By Theorem 9.1 the space

$$\bigcup_{\epsilon \leq C} \mathcal{M}_{1,\epsilon}(E^-, J^-)/\mathbb{R}$$

is finite for any $C > 0$. Together with the finiteness of (12.3) this implies that $\widehat{\mathcal{M}}^{1,0}_{1,\epsilon}(E, J; \nu^-, \nu^+)$ is finite. The same argument can be used for $\widehat{\mathcal{M}}^{0,1}_{1,\epsilon}(E, J; \nu^-, \nu^+)$. □

By Corollary 11.5, $\mathcal{M}_{1,\epsilon}(E, J; \nu^-, \nu^+)$ is a smooth one-dimensional manifold for all ν^-, ν^+ and ϵ. Define

$$\overline{\mathcal{M}}_{1,\epsilon}(E, J; \nu^-, \nu^+) = \mathcal{M}_{1,\epsilon}(E, J; \nu^-, \nu^+) \sqcup \mathcal{M}'_\epsilon(E, J; \nu^-, \nu^+).$$

Lemma 12.11, Lemma 12.12 and the gluing Theorem 12.5 can be used to endow $\overline{\mathcal{M}}_{1,\epsilon}(E, J; \nu^-, \nu^+)$ with the structure of a compact one-dimensional manifold. This is precisely the compactification required in Theorem 9.4.

13 Transversality for sections

The discussion at the end of the previous section completes one half of the proof of Theorem 9.4. We will now begin to explain the remaining half. We retain the conventions for (E, Ω, J'), π, and $(T^\pm, \Theta^\pm)$, and assume that J^- and J^+ are almost complex structures in $\mathcal{J}_{\mathrm{reg}}(T^-, \Theta^+)$ and $\mathcal{J}_{\mathrm{reg}}(T^+, \Theta^+)$, respectively. What we have to prove is

Theorem 13.1. *If in addition to the assumptions above, E is an* ordinary *almost holomorphic fibration, the subset*

$$\mathcal{J}_{\mathrm{reg}}(E, \Omega, J'; J^-, J^+) \subset \mathcal{J}(E, \Omega, J'; J^-, J^+)$$

is dense in the C^∞-topology.

This section contains the first part of the proof, in which the assumption that E is ordinary is not yet necessary:

Proposition 13.2. *Let $\mathcal{J}^{(1)}_{\mathrm{reg}} \subset \mathcal{J}(E, \Omega, J'; J^-, J^+)$ be the subset of almost complex structures J such that all $\sigma \in \mathcal{M}(E, J)$ are regular. Then $\mathcal{J}^{(1)}_{\mathrm{reg}}$ is dense (in the C^∞-topology).*

Following a standard pattern, the proof of this result has an abstract and a more specific part. The first one consists in setting up a suitable analytic framework and in applying the Sard-Smale and implicit function theorems. The second one is the study of a certain operator $D\bar{\partial}^{\mathrm{univ}}(\sigma, J)$ associated to $\sigma \in \mathcal{M}(E, J)$. We begin by indicating the idea of the proof in the C^∞-topology; this is not actually the appropriate topology for carrying out the argument, but it serves to motivate the introduction of the operators $D\bar{\partial}^{\mathrm{univ}}(\sigma, J)$. After that we introduce Floer's C^∞_ϵ-topology which replaces the C^∞-topology; and then we prove the required property of the operators $D\bar{\partial}^{\mathrm{univ}}(\sigma, J)$.

Fix some $J_0 \in \mathcal{J}(E, \Omega, J'; J^-, J^+)$. By assumption, there is an $R > 0$ such that J_0 agrees with J^- on $\pi^{-1}((-\infty; -R] \times S^1)$ and with J^+ on $\pi^{-1}([R; \infty) \times S^1)$. Moreover, there is a closed neighbourhood U of $\mathrm{Crit}(\pi)$ such that $J_0 = J'$ on U. We assume that R has been chosen large and U small; more precisely, it is enough if $\pi(U) \subset [-R+1; R-1] \times S^1$. This can certainly be arranged since $\mathrm{Crit}(\pi)$ is compact.

Let $\mathcal{J} \subset \mathcal{J}(E, \Omega, J'; J^-, J^+)$ be the subset of almost complex structures J which agree with J_0 on

$$E_0 = \pi^{-1}((-\infty; -R] \times S^1 \cup [R; \infty) \times S^1) \cup U.$$

$\mathcal{J}$ is a Fréchet manifold; the tangent space $T_J\mathcal{J}$ is the space of J-antilinear smooth homomorphisms $Y : TE \longrightarrow TE^v$ which vanish on E_0. A explicit collection of charts on $\mathcal{J}$ can be constructed in the following way:

Lemma 13.3. *The map*

$$q_J(Y) = J(1 - \tfrac{1}{2}JY)(1 + \tfrac{1}{2}JY)^{-1}$$

is a homeomorphism from a neighbourhood $N \subset T_J\mathcal{J}$ of 0 to a neighbourhood of J in $\mathcal{J}$.

The proof is an adaptation of the standard proof of contractibility of the space of tame almost complex structures on a symplectic manifold, see e.g. Proposition 1.1.6 in Chapter II of [3]. We omit the details.

A rough (and not entirely correct) idea of the proof of Proposition 13.2 is this: one considers the *universal moduli spaces*

$$\mathcal{M}^{\mathrm{univ}}(E; \nu^-, \nu^+) = \bigcup_{J \in \mathcal{J}} \mathcal{M}(E, J; \nu^-, \nu^+) \times \{J\} \subset \mathcal{S}(E; \nu^-, \nu^+) \times \mathcal{J}$$

for $\nu^\pm \in \mathcal{H}(T^\pm, \Theta^\pm)$ and proves that these are smooth. This is done in the following way: let $\mathcal{E}^{\mathrm{univ}} \longrightarrow \mathcal{S}(E; \nu^-, \nu^+) \times \mathcal{J}$ be the pullback of the Banach space bundle $\mathcal{E} \longrightarrow \mathcal{S}(E; \nu^-, \nu^+)$. This bundle has a canonical section $\bar\partial^{\mathrm{univ}}$, given by

$$\bar\partial^{\mathrm{univ}}(\sigma, J) = \frac{\partial\sigma}{\partial s} + J(\sigma)\frac{\partial\sigma}{\partial t},$$

and $\mathcal{M}^{\mathrm{univ}}(E; \nu^-, \nu^+)$ is the zero set of this section. The derivative of $\bar\partial^{\mathrm{univ}}$ at a point $(\sigma, J) \in \mathcal{M}^{\mathrm{univ}}(E; \nu^-, \nu^+)$ is an operator

$$D\bar\partial^{\mathrm{univ}}(\sigma, J) : W^{1,p}(\sigma^*TE^v) \times T_J\mathcal{J} \longrightarrow L^p(\sigma^*TE^v)$$

(recall that $T_\sigma\mathcal{S}(E; \nu^-, \nu^+) = W^{1,p}(\sigma^*TE^v)$ and $\mathcal{E}_\sigma = L^p(\sigma^*TE^v)$). It is not difficult to determine this operator explicitly:

$$D\bar\partial^{\mathrm{univ}}(\sigma, J)(X, Y) = D\bar\partial_J(\sigma)X + Y\frac{\partial\sigma}{\partial t}. \tag{13.1}$$

The smoothness of the universal moduli space follows (using the implicit function theorem) from the following result:

Lemma 13.4. *The operator $D\bar{\partial}^{\mathrm{univ}}(\sigma, J)$ is surjective for all $J \in \mathcal{J}$ and $\sigma \in \mathcal{M}(E, J)$.*

Having proved this, one observes that the subset of $J \in \mathcal{J}$ such that all J-holomorphic sections with limits $\nu^{\pm}$ are regular is the set of regular values of the projection

$$\mathcal{M}^{\mathrm{univ}}(E; \nu^-, \nu^+) \longrightarrow \mathcal{J}.$$

This projection is a Fredholm map, and the Sard-Smale theorem shows that the set of regular values is dense. Since this holds for any J_0, it implies the denseness of $\mathcal{J}_{\mathrm{reg}}^{(1)}$ in $\mathcal{J}(E, \Omega, J'; J^-, J^+)$.

The argument as we have just presented it contains a serious error: $\mathcal{J}$ is only a Fréchet manifold and therefore neither the implicit function theorem nor the Sard-Smale theorem can be applied to it. We will now modify the framework to remove this technical obstacle.

Let $\|\cdot\|_{C^k}$ be the C^k-norm on $T_{J_0}\mathcal{J}$ determined by some Riemannian metric on E, and let $\epsilon = (\epsilon_k)_{k \in \mathbb{N}}$ be a sequence of positive numbers. Floer's C^∞_ϵ-seminorm on $T_{J_0}\mathcal{J}$ is defined by

$$\|Y\|_\epsilon = \sum_{k=0}^{\infty} \epsilon_k \|Y\|_{C^k}.$$

Lemma 13.5. *There is an (ϵ_k) such that the subset of $Y \in T_{J_0}\mathcal{J}$ with $\|Y\|_\epsilon < \infty$ is dense in the C^∞-topology.*

We refer to [27, pp. 101–104] for a detailed exposition of the C^∞_ϵ-topology and the proof of this Lemma. Let $N \subset T_{J_0}\mathcal{J}$ be a neighbourhood of 0 as in Lemma 13.3 above and N_ϵ the subset of those $Y \in N$ such that $\|Y\|_\epsilon < \infty$. Evidently

$$\mathcal{J}_\epsilon = q_{J_0}(N_\epsilon) \subset \mathcal{J}$$

carries a natural Banach manifold structure, induced by that of N_ϵ. If (ϵ_k) is as in Lemma 13.5, the tangent space $T_J\mathcal{J}_\epsilon$ at any point $J \in \mathcal{J}_\epsilon$ is C^∞-dense in $T_J\mathcal{J}$. We will assume from now on that (ϵ_k) has been chosen in this way; $\mathcal{J}_\epsilon$ is called a C^∞_ϵ-neighbourhood of J_0.

Now we carry out the argument outlined above in the C^∞_ϵ-topology: the pullback $\mathcal{E}^{\mathrm{univ}}$ of $\mathcal{E}$ to $\mathcal{S}(E; \nu^-, \nu^+) \times \mathcal{J}_\epsilon$ has a smooth section $\bar{\partial}^{\mathrm{univ}}$ given by the same formula (13.1) as above (the smoothness is easy to prove because the C^∞_ϵ-norm is a very strong one). $(\bar{\partial}^{\mathrm{univ}})^{-1}(0)$ is the C^∞_ϵ-version of the universal moduli space,

$$\mathcal{M}^{\mathrm{univ}}(E; \nu^-, \nu^+) = \bigcup_{J \in \mathcal{J}_\epsilon} \mathcal{M}(E, J; \nu^-, \nu^+) \times \{J\}$$

and the derivative of $\bar{\partial}^{\mathrm{univ}}$ at a point of this space is given by an operator

$$D\bar{\partial}^{\mathrm{univ}}(\sigma, J) : W^{1,p}(\sigma^* TE^v) \times T_J \mathcal{J}_\epsilon \longrightarrow L^p(\sigma^* TE^v) \tag{13.2}$$

which is the restriction of the corresponding operator above to $T_J \mathcal{J}_\epsilon \subset T_J \mathcal{J}$.

Lemma 13.6. *The operator* (13.2) *is onto for all* (σ, J)*.*

We postpone the proof of this and continue with our argument. Since $\mathcal{S}(E; \nu^-, \nu^+) \times \mathcal{J}_\epsilon$ is a Banach manifold and $\mathcal{E}^{\mathrm{univ}}$ a bundle of Banach spaces, we can use the implicit function theorem; it shows that the universal moduli space is a smooth Banach submanifold.

Lemma 13.7. *Let* $J \in \mathcal{J}_\epsilon$ *be a regular value of the projection*

$$\Pi : \mathcal{M}^{\mathrm{univ}}(E; \nu^-, \nu^+) \longrightarrow \mathcal{J}_\epsilon.$$

Then every $\sigma \in \mathcal{M}(E, J; \nu^-, \nu^+)$ *is regular.*

This is straightforward given the definition of regularity and Lemma 13.6.

Lemma 13.8. *The derivatives of* Π *are Fredholm operators.*

This is a simple consequence of the Fredholm property of $D\bar{\partial}_J(\sigma)$ (Theorem 11.2). We can now apply the Sard-Smale theorem which says that the regular values of Π form a subset of second category in $\mathcal{J}_\epsilon$. This proves Proposition 13.2. In fact the argument yields the following stronger result:

Proposition 13.9. *For every* C^∞_ϵ*-neighbourhood* $\mathcal{J}_\epsilon \subset \mathcal{J}$*, the intersection* $\mathcal{J}^{(1)}_{\mathrm{reg}} \cap \mathcal{J}_\epsilon$ *is a subset of second category in* $\mathcal{J}_\epsilon$ *with respect to the* C^∞_ϵ*-topology.*

The importance of this stronger version is that the intersection of any (countable) number of subsets of $\mathcal{J}$ with this property remains dense. We will now fill the remaining gap in the argument, that is, prove Lemma 13.6. We begin with a proof of its C^∞-analogue:

Proof of Lemma 13.4. Let $\mathrm{Hom}^{0,1}_J(TE, TE^v)$ be the vector bundle over $E \setminus \mathrm{Crit}(\pi)$ which consists of J-antilinear homomorphisms $TE \longrightarrow TE^v$. The map

$$\sigma^*(\mathrm{Hom}^{0,1}_J(TE, TE^v)) \longrightarrow \sigma^* TE^v, \quad Y' \longmapsto Y' \frac{\partial \sigma}{\partial t} \tag{13.3}$$

is a homomorphism of vector bundles over C. Since σ is a section, $\frac{\partial \sigma}{\partial t}$ is nowhere zero. Using some simple linear algebra, it follows that (13.3) is surjective. Therefore, given a $Z \in C^\infty(\sigma^* TE^v)$ which is supported in $(R-1; R) \times S^1$, we can find a $Y' \in C^\infty(\sigma^*(\mathrm{Hom}^{0,1}_J(TE, TE^v))$ which is supported in the same subset and such that $Y' \frac{\partial \sigma}{\partial t} = Z$. Because $\sigma : C \longrightarrow E$

is an embedding, Y' can be extended to a J-antilinear homomorphism $Y : TE \longrightarrow TE^v$ supported inside $\pi^{-1}((R-1;R)\times S^1)$. Such a homomorphism is an element of $T_J\mathcal{J}$ since $E_0 \cap \pi^{-1}((R-1;R)\times S^1) = \emptyset$. In view of the formula (13.1), this proves that every Z as above lies in the image of

$$D\bar{\partial}^{\text{univ}}(\sigma, J) : W^{1,p}(\sigma^* TE^v) \times T_J\mathcal{J} \longrightarrow L^p(\sigma^* TE^v).$$

Because $D\bar{\partial}_J(\sigma)$ is Fredholm the image of this operator is closed. Let us assume that $D\bar{\partial}^{\text{univ}}(\sigma, J)$ is not onto. Then there is a nontrivial L^q-section ($1/p + 1/q = 1$) of the dual bundle $\sigma^*(TE^v)^*$ which is orthogonal to the image of $D\bar{\partial}^{\text{univ}}(\sigma, J)$. If we denote this section by W, this means that

$$\int_C \langle D\bar{\partial}_J(\sigma)X, W\rangle + \langle Y\frac{\partial\sigma}{\partial t}, W\rangle = 0 \tag{13.4}$$

for all (X, Y). Restricting to $X = 0$ yields

$$\int_C \langle Y\frac{\partial\sigma}{\partial t}, W\rangle = 0.$$

Since $Y\frac{\partial\sigma}{\partial t}$ can take on any smooth value which is supported in $(R-1;R)\times S^1$ it follows that $W|(R-1;R)\times S^1 = 0$. On the other hand, restricting (13.4) to $Y = 0$ shows that

$$D\bar{\partial}_J(\sigma)^* W = 0, \tag{13.5}$$

where $D\bar{\partial}_J(\sigma)^*$ is the differential operator dual to $D\bar{\partial}_J(\sigma)$. This operator has the form

$$D\bar{\partial}_J(\sigma)^* = -\frac{\partial}{\partial s} + J(\sigma)\frac{\partial}{\partial t} + \text{a term of order zero.}$$

Solutions of (13.5) satisfy a unique continuation property (see [12]). Therefore $W|(R-1;R)\times S^1 = 0$ implies that $W = 0$, contradicting the original assumption. This shows that $D\bar{\partial}^{\text{univ}}(\sigma, J)$ is onto. □

When trying to adapt this proof to the C^∞_ϵ-case one runs into a problem: it is certainly not true that any $Z \in C^\infty(\sigma^* TE^v)$ supported in $(R;R+1)\times S^1$ is of the form $Z = Y\frac{\partial\sigma}{\partial t}$ with $Y \in T_J\mathcal{J}_\epsilon$. This problem can be overcome by modifying the argument slightly. Alternatively, the C^∞_ϵ-version of the surjectivity result can be reduced to its C^∞-version, and this is what we will do.

Proof of Lemma 13.6. We have to show that

$$D\bar{\partial}^{\text{univ}}(\sigma, J)(W^{1,p}(\sigma^* TE^v) \times T_J\mathcal{J}_\epsilon) = L^p(\sigma^* TE^v).$$

By Lemma 13.4,

$$D\bar{\partial}^{\text{univ}}(\sigma, J)(W^{1,p}(\sigma^* TE^v) \times T_J\mathcal{J}) = L^p(\sigma^* TE^v).$$

Since $D\bar{\partial}^{\mathrm{univ}}(\sigma, J)$ is continuous and $T_J\mathcal{J}_\epsilon \subset T_J\mathcal{J}$ is dense, it follows that $D\bar{\partial}^{\mathrm{univ}}(\sigma, J)(W^{1,p}(\sigma^*TE^v) \times T_J\mathcal{J}_\epsilon)$ is dense in $L^p(\sigma^*TE^v)$. On the other hand, this subset is closed because it contais the image if $D\bar{\partial}_J(\sigma)$, which is closed and of finite codimension. □

14 Transversality for bubbles

We retain the same notation as in the previous section, assuming, however, that E is ordinary. This section contains the proofs of the following two results:

Proposition 14.1. *Let $\mathcal{J}^{(2)}_{\mathrm{reg}} \subset \mathcal{J}(E, \Omega, J'; J^-, J^+)$ be the subset of almost complex structures J such that any J-bubble which does not lie in a singular fibre of π has nonnegative Chern number. For every C^∞_ϵ-neighbourhood $\mathcal{J}_\epsilon \subset \mathcal{J}$, the intersection $\mathcal{J}^{(2)}_{\mathrm{reg}} \cap \mathcal{J}_\epsilon$ is a subset of second category in $\mathcal{J}_\epsilon$ with respect to the C^∞_ϵ-topology.*

Proposition 14.2. *Let $\mathcal{J}^{(3)}_{\mathrm{reg}} \subset \mathcal{J}(E, \Omega, J'; J^-, J^+)$ be the subset of almost complex structures J such that any pair (σ, w), consisting of a $\sigma \in \mathcal{M}(E, J)$ with index ≤ 1 and a J-bubble w with Chern number 0 which does not lie in a singular fibre, satisfies $\mathrm{im}(\sigma) \cap \mathrm{im}(w) = \emptyset$. $\mathcal{J}^{(3)}_{\mathrm{reg}}$ has the same property as the subset $\mathcal{J}^{(2)}_{\mathrm{reg}}$ above.*

The subset $\mathcal{J}^{(1)}_{\mathrm{reg}}$ is defined by a certain transversality condition (that $\bar{\partial}_J$ should be transverse to the zero section) and that was the starting point for the proof of Proposition 13.9. The first step in the proof of the two Propositions above is to write the conditions defining $\mathcal{J}^{(k)}_{\mathrm{reg}}$, $k = 2, 3$ (or some smaller subset) in a similar way. In order to do this we need the transversality theory for J-bubbles (which do not lie in a singular fibre of π). This theory is well-known: it is the parametrized version of the transversality theory of J-holomorphic curves in a symplectic manifold, and appears e.g. in [20, Theorem 3.1.3]. We will only give a quick overview. Note that this theory does not apply to J-bubbles in a singular fibre of π; such bubbles will be dealt with in a slightly different way in the next section.

Let J be an Ω-tame almost complex structure on E and w a J-bubble. A point $z \in \mathbb{CP}^1$ is an *injective point* of w if $Dw(z) \neq 0$ and $w^{-1}(w(z)) = \{z\}$. We call a J-bubble *simple* if it has an injective point. In that case, the set of injective points is open and dense in $\mathbb{CP}^1$ (this is part of Proposition 2.3.1 in [20]).

Lemma 14.3. *Any J-bubble w which is not simple is multiply-covered: that is, it can be represented as*

$$w = \bar{w} \circ c,$$

where $\bar{w}$ is a simple J-bubble and $c : \mathbb{C}\mathrm{P}^1 \longrightarrow \mathbb{C}\mathrm{P}^1$ is a holomorphic map of degree ≥ 2.

We refer to [20, Chapter 2] for the proof. $\bar{w}$ is unique up to complex automorphisms of $\mathbb{C}\mathrm{P}^1$. By a slight abuse of notation, we will call it the simple J-bubble underlying w. Clearly $\Omega(w) = d\,\Omega(\bar{w})$ and $c_1(w) = d\,c_1(\bar{w})$ with $d = \deg(c) \geq 2$. The set of simple J-bubbles with Chern number k which do not lie in a singular fibre of π will be denoted by $\mathcal{B}_k(E, J)$. The group $PSL(2, \mathbb{C})$ of automorphisms of $\mathbb{C}\mathrm{P}^1$ acts (by composition) freely and properly on $\mathcal{B}_k(E, J)$ for any k.

Let $\mathcal{P}(E)$ be the space of non-constant $W^{1,p}$-maps $w : \mathbb{C}\mathrm{P}^1 \longrightarrow E$ such that $\pi \circ w$ is constant and not a singular value of π. $\mathcal{P}(E)$ is a smooth Banach manifold. Its tangent space at a point w consists of all sections X of the vector bundle $w^*TE \longrightarrow \mathbb{C}\mathrm{P}^1$ of class $W^{1,p}$ which have the following property: $D\pi(X(z_1)) = D\pi(X(z_2))$ for any $z_1, z_2 \in \mathbb{C}\mathrm{P}^1$. This tangent space $T_w\mathcal{P}(E)$ contains $W^{1,p}(w^*TE^v)$, which has codimension two in it.

Let J be an Ω-tame almost complex structure on E. For $w \in \mathcal{P}(E)$, let $\Lambda^{0,1}_J(w^*TE^v)$ be the vector bundle over $\mathbb{C}\mathrm{P}^1$ which consists of the J-antilinear homomorphisms $T\mathbb{C}\mathrm{P}^1 \longrightarrow w^*TE^v$. The spaces $L^p(\Lambda^{0,1}_J(\sigma^*TE^v))$ of L^p-sections of this vector bundle are fibres of a Banach space bundle over $\mathcal{P}(E)$ which we denote by Λ_J. Λ_J has a canonical smooth section, given by

$$\bar{\partial}'_J(w) = Dw + J \circ Dw \circ i$$

(i denotes the complex structure on $\mathbb{C}\mathrm{P}^1$). The standard regularity result for J-holomorphic curves ([20, Theorem B.4.1]) shows that $(\bar{\partial}'_J)^{-1}(0)$ is the set of (smooth) J-bubbles which do not lie in a singular fibre of π. The derivative of $\bar{\partial}'_J$ at a point $w \in (\bar{\partial}'_J)^{-1}(0)$ is an operator

$$D\bar{\partial}'_J(w) : T_w\mathcal{P}(E) \longrightarrow L^p(\Lambda^{0,1}_J(w^*TE^v)).$$

Let D_w be the restriction of $D\bar{\partial}'_J(w)$ to the codimension two subspace $W^{1,p}(w^*TE^v)$. D_w is a differential operator of the form

$$D_w = \bar{\partial}_\nabla + \text{a term of order zero}, \tag{14.1}$$

where $\bar{\partial}_\nabla$ is the $\bar{\partial}$-operator on w^*TE^v defined using some connection ∇; of course, the choice of ∇ is irrelevant for (14.1). Using the Riemann-Roch theorem, one obtains

Proposition 14.4. *$D\bar{\partial}'_J(w)$ is a Fredholm operator of index $6 + 2c_1(w)$.*

We call $w \in (\bar{\partial}'_J)^{-1}(0)$ regular if $D\bar{\partial}'_J(w)$ is onto.

Lemma 14.5. *Let $J \in \mathcal{J}(E, \Omega, J'; J^-, J^+)$ be an almost complex structure such that any $w \in \mathcal{B}_k(E, J)$ with $k < 0$ is regular. Then $J \in \mathcal{J}^{(2)}_{\mathrm{reg}}$.*

Proof. From the implicit function theorem it follows that for all $k < 0$, $\mathcal{B}_k(E, J)$ is a smooth manifold of dimension $6 + 2k$. Since $PSL(2, \mathbb{C})$ acts freely and properly on this manifold, the quotient $\mathcal{B}_k(E, J)/PSL(2, \mathbb{C})$ is smooth of dimension $2k$, hence empty for negative k. By reduction to the underlying simple J-bubble it follows that any J-bubble which does not lie in a singular fibre of π has nonnegative Chern number. □

The rest of the proof of Proposition 14.1 follows the same strategy as in the previous section. We start with a fixed $J_0 \in \mathcal{J}(E, \Omega, J'; J^-, J^+)$ and consider a subspace $\mathcal{J} \subset \mathcal{J}(E, \Omega, J'; J^-, J^+)$ as we did then. However, we will make an additional assumption on the neighbourhood U of $\mathrm{Crit}(\pi)$. Namely, we assume that $[\Omega|U] \in H^2(U; \mathbb{R})$ is zero. This is certainly possible because the critical points of π are isolated (this follows from the assumption that E is an ordinary holomorphic fibration). We will not explain how the universal moduli problem is set up since this parallels closely the case of J-holomorphic sections. Eventually, Proposition 14.1 is reduced to the following technical statement:

Lemma 14.6. *Let w be a simple J-bubble with negative Chern number for some $J \in \mathcal{J}$. Then the operator*

$$(D\bar{\partial}')^{\mathrm{univ}}(w, J) : T_w\mathcal{P}(E) \times T_J\mathcal{J} \longrightarrow L^p(\Lambda^{0,1}_J(w^*TE^v)),$$
$$(D\bar{\partial}')^{\mathrm{univ}}(w, J)(X, Y) = D\bar{\partial}'_J(w)X + Y \circ Dw \circ i$$

is onto.

Note that the transversality argument uses C^∞_ϵ-norms and therefore leads to a C^∞_ϵ-version of Lemma 14.6. However, this version can be reduced to the one which we have just stated in the same way as Lemma 13.6 has been reduced to Lemma 13.4.

Proof of Lemma 14.6. The first step is to show that $\mathrm{im}(w)$ is not contained in

$$E_0 = \pi^{-1}((-\infty; -R] \times S^1 \cup [R; \infty) \times S^1) \cup U \subset E.$$

$\mathrm{im}(w)$ cannot be lie in one of the fibres over $(-\infty; -R] \times S^1 \subset C$, for the following reason: on these fibres J agrees with J^-, but since (by assumption) $J^- \in \mathcal{J}_{\mathrm{reg}}(T^-, \Theta^-)$, there are no J^--holomorphic bubbles with negative Chern number (compare Definition 11.10). The same holds for $\pi^{-1}([R; \infty) \times S^1)$. Finally, w can not lie completely within U since $[\Omega|U] = 0$, whereas $\Omega(w) > 0$ because w is not constant. Since E_0 is closed and the set of injective points is dense in $\mathbb{CP}^1$, w has an injective point z_0 such that $w(z_0) \notin E_0$. Because the of injective points of w is open, it follows that there is an open subset $Q \subset \mathbb{CP}^1$ such that any $z \in Q$ is an injective point and satisfies $w(z) \notin E_0$.

Consider the map of vector bundles over $\mathbb{C}\mathrm{P}^1$

$$\begin{aligned} w^*(\mathrm{Hom}_J^{0,1}(TE, TE^v)) &\longrightarrow \Lambda_J^{0,1}(w^*TE^v), \\ Y &\longmapsto Y \circ Dw \circ i. \end{aligned} \tag{14.2}$$

Since Dw does not vanish anywhere in Q, a simple argument from linear algebra shows that for any smooth section Z of $\Lambda_J^{0,1}(w^*TE^v)$ which is supported in Q there is a section Y' of $w^*(\mathrm{Hom}_J^{0,1}(TE, TE^v))$, equally supported in Q, which is mapped to Z under (14.2). Since $w|Q$ is an embedding, Y' can be extended to a section Y of $\mathrm{Hom}_J^{0,1}(TE, TE^v)$ which is supported outside E_0, that is, $Y \in T_J\mathcal{J}$. This shows that any such Z lies in the image of $(D\bar{\partial}')^{\mathrm{univ}}(w, J)$.

Assume that $(D\bar{\partial}')^{\mathrm{univ}}(w, J)$ is not onto. By the same argument as in the proof of Lemma 13.4, one obtains a nonzero L^q-section Z of the bundle dual to $\Lambda_J^{0,1}(w^*TE^v)$ which vanishes on Q and satisfies

$$\int_{\mathbb{C}\mathrm{P}^1} \langle D\bar{\partial}'_J(\sigma)X, Z\rangle = 0$$

for all $X \in T_w\mathcal{P}(E)$. By restricting to $W^{1,p}(\sigma^*TE^v) \subset T_w\mathcal{P}(E)$ one sees that

$$D_w^* Z = 0.$$

Again, a unique continuation principle holds for the solutions of this equation: therefore $Z = 0$, which completes the proof of Lemma 14.6. □

Lemma 14.6 is familiar (compare [20, p. 35]); we have repeated the usual proof to show that the condition that J should be equal to J' near $\mathrm{Crit}(\pi)$ does not prevent transversality.

It remains to prove Proposition 14.2. We will be even more brief in this case, and write down only the proof of the basic technical result:

Lemma 14.7. *Choose a $J \in \mathcal{J}(E, \Omega, J'; J^-, J^+)$. Let w be a simple J-bubble of Chern number 0 which lies in a regular fibre of π, $\sigma \in \mathcal{M}(E, J)$ a J-holomorphic section, and $z_1 \in C$, $z_2 \in \mathbb{C}\mathrm{P}^1$ points such that $\sigma(z_1) = w(z_2)$. Then the operator*

$$\begin{aligned} D : W^{1,p}(\sigma^*TE^v) \times T_w\mathcal{P}(E) \times T_J\mathcal{J} \times T_{z_1}C \times T_{z_2}\mathbb{C}\mathrm{P}^1 \longrightarrow \\ \longrightarrow L^p(\sigma^*TE^v) \times L^p(\Lambda_J^{0,1}(w^*TE^v)) \times TE_{\sigma(z_1)} \end{aligned}$$

given by

$$\begin{aligned} D(X_1, X_2, Y, V_1, V_2) = \big(D\bar{\partial}_J(\sigma)X_1 + Y\frac{\partial\sigma}{\partial t}, \\ D\bar{\partial}'_J(w)X_2 + Y \circ Dw \circ i, X_1(z_1) + D\sigma(V_1) - X_2(z_2) - Dw(V_2)\big) \end{aligned}$$

is surjective.

Proof. As usual, we assume that D is not onto and consider a nonzero $W = (W_1, W_2, W_3)$ which is orthogonal to its image. This means that

$$\left(\int_C \langle D\bar{\partial}_J(\sigma)X_1 + Y\frac{\partial\sigma}{\partial t}, W_1\rangle\right) + \left(\int_{\mathbb{C}\mathrm{P}^1} \langle D\bar{\partial}'_J(w)X_2 + Y \circ Dw \circ i, W_2\rangle\right) + \\ + \langle X_1(z_1) + D\sigma(V_1) - X_2(z_2) - Dw(V_2), W_3\rangle = 0 \quad (14.3)$$

for all X_1, X_2, Y, V_1, V_2. By using only the component X_1, one obtains

$$\left(\int_C \langle D\bar{\partial}_J(\sigma)X_1, W_2\rangle\right) + \langle X_1(z_1), W_3\rangle = 0.$$

Therefore

$$D\bar{\partial}_J(\sigma)^* W_1 = -\delta_{z_1} W_3^v \tag{14.4}$$

where δ_{z_1} is the δ-function at $z_1 \in C$ and W_3^v is image of W_3 under the projection $(TE_{\sigma(z_1)})^* \longrightarrow (TE^v_{\sigma(z_1)})^*$. In particular $D\bar{\partial}_J(\sigma)^* W_1 = 0$ *away from the point* z_1.

Because σ is a section and w lies in one fibre of π they can intersect at most in one point. It follows that there are $r, r' \in \mathbb{R}$ with $R - 1 \leq r < r' \leq R$ and $z_1 \notin (r; r') \times S^1$, such that $\sigma([r; r'] \times S^1) \cap \mathrm{im}(w) = \emptyset$. From (14.3) it follows that

$$\int_C \langle Y\frac{\partial\sigma}{\partial t}, W_1\rangle = 0$$

for every $Y \in T_J\mathcal{J}$ supported in $\pi^{-1}((r; r') \times S^1)$. The same argument as in the proof of Lemma 13.4 shows that $W_1 | (r; r') \times S^1 = 0$. By unique continuation, W_1 vanishes everywhere except possibly at z_1; and since it is an L^q-function, this means that $W_1 = 0$, and, by (14.4), $W_3^v = 0$.

Using only the component V_1, one obtains from (14.3) that

$$\langle D\sigma(V_1), W_3\rangle = 0.$$

We have already seen that $\langle Z, W_3\rangle = 0$ for all $Z \in TE^v_{\sigma(z_1)}$. Since $\mathrm{im}(D\sigma_{z_1}) \oplus TE^v_{\sigma(z_1)} = TE_{\sigma(z_1)}$, it follows that $W_3 = 0$.

The remaining component W_2 satisfies

$$\int_{\mathbb{C}\mathrm{P}^1} \langle D\bar{\partial}'_J(w)X_2 + Y \circ Dw \circ i, W_2\rangle = 0$$

for all (X_2, Y). This is the situation of Lemma 14.6, and we have proved that in that case, $W_2 = 0$. □

Problems similar to Lemma 14.7 appear at several points in the theory of pseudo-holomorphic curves. The use of δ-functions, which seems to be new, simplifies the usual argument (compare [20, Lemma 6.1.2]).

15 Resolution of the singular fibres

We retain the same notation as in the previous section. The aim of this section is to prove

Proposition 15.1. *Let* $\mathcal{J}^{(4)}_{\mathrm{reg}} \subset \mathcal{J}(E,\Omega,J';J^-,J^+)$ *be the subset of almost complex structures* J *such that any* J*-bubble which lies in a singular fibre of* π *has positive Chern number. For every* C^∞_ϵ*-neighbourhood* $\mathcal{J}_\epsilon \subset \mathcal{J}$*, the intersection* $\mathcal{J}^{(4)}_{\mathrm{reg}} \cap \mathcal{J}_\epsilon$ *is a subset of second category in* $\mathcal{J}_\epsilon$ *with respect to the* C^∞_ϵ*-topology.*

Note that together with Propositions 13.9, 14.1 and 14.2 this completes the proof of Theorem 13.1. This is because these four results imply that the intersection $\mathcal{J}^{(1)}_{\mathrm{reg}} \cap \mathcal{J}^{(2)}_{\mathrm{reg}} \cap \mathcal{J}^{(3)}_{\mathrm{reg}} \cap \mathcal{J}^{(4)}_{\mathrm{reg}}$ is dense in any C^∞_ϵ-neigbourhood, and hence dense in $\mathcal{J}(E,\Omega,J';J^-,J^+)$. By definition, this intersection is a subset of $\mathcal{J}_{\mathrm{reg}}(E,\Omega,J';J^-,J^+)$, and therefore $\mathcal{J}_{\mathrm{reg}}(E,\Omega,J';J^-,J^+)$ is also dense.

The difference between Proposition 15.1 and the previous results of a similar nature is that the familiar setup for transversality theory does not work for J-bubbles in a singular fibre; for instance, if E_z is such a fibre, the space of $W^{1,p}$-maps $\mathbb{C}\mathrm{P}^1 \longrightarrow E_z$ is not a Banach manifold. We will avoid this problem by resolving the singular fibre and lifting the J-bubbles to the resolution. This uses the fact that (E,Ω,J') is ordinary in an essential way (in the previous section we have only used that the critical points are isolated). For simplicity, we assume that $\pi : E \longrightarrow C$ has only one critical point $x_0 \in E_{z_0}$ (if there are no critical points, Proposition 15.1 is vacuous). We begin by considering the local model for the resolution.

Lemma 15.2. *Let* $Q \subset \mathbb{C}^3$ *be the singular quadric defined by*

$$x_1^2 + x_2^2 + x_3^2 = 0,$$

$\widehat{Q}$ *its proper transform with respect to the blowup of* $0 \in \mathbb{C}^3$*, and* $r : \widehat{Q} \longrightarrow Q$ *the canonical projection. Then*

(1) $\widehat{Q}$ *is smooth, and* $D = r^{-1}(0)$ *is a rational curve with self-intersection* (-2)*.*

(2) $r|\widehat{Q} \setminus D : \widehat{Q} \setminus D \longrightarrow Q \setminus \{0\}$ *is an isomorphism.*

(3) Let $w : \mathbb{C} \longrightarrow Q$ *be a non-constant holomorphic map with* $w(0) = 0$*. There is a unique holomorphic map* $\hat{w} : \mathbb{C} \longrightarrow \widehat{Q}$ *such that* $w = r \circ \hat{w}$*.*

$\widehat{Q}$ is the subvariety of $\mathbb{C}^3 \times \mathbb{C}\mathrm{P}^2$ defined by the equations

$$\xi_1^2 + \xi_2^2 + \xi_3^2 = 0 \quad \text{and} \quad \xi_i x_j = \xi_j x_i$$

for $(x,\xi) \in \mathbb{C}^3 \times \mathbb{CP}^2$. Starting from this description, all properties stated above can be proved in an elementary way. We omit the details. Note that (3) is a local property, that is, it holds for germs of holomorphic maps w.

Because we have assumed that the second derivative $D^2\pi_{x_0}$ is nondegenerate, the complex Morse Lemma [2, Lemma 2] says that the singularity of the fibre E_{z_0} is locally isomorphic to the singularity of Q at 0. A precise statement is this:

Lemma 15.3. *There is a neighbourhood $U \subset E$ of x_0 and a holomorphic embedding*

$$f : (U, J') \longrightarrow \mathbb{C}^3$$

with $f(x_0) = 0$ and $f(U \cap E_{z_0}) \subset Q$. □

Note that the fact that J' is integrable is essential here. We can use this local description of E_{z_0} to glue in the resolution $\widehat{Q}$. More precisely, let $U' = f(U \cap E_{z_0}) \subset Q$; we glue together $E_{z_0} \setminus \{x_0\}$ and $r^{-1}(U') \subset \widehat{Q}$ using the diffeomorphism

$$r^{-1}(U') \setminus D \xrightarrow{r} U' \setminus \{0\} \xrightarrow{f^{-1}} (U \cap E_{z_0}) \setminus \{x_0\}.$$

This yields a smooth compact four-manifold $\widehat{E}_{z_0}$ with a map $r_{x_0} : \widehat{E}_{z_0} \longrightarrow E_{z_0}$. We call $\widehat{E}_{z_0}$ the resolution of E_{z_0}.

Let J be an almost complex structure in $\mathcal{J}(E, \Omega, J'; J^-, J^+)$. By definition, J must agree with J' on a neighbourhood of x_0. Using the complex structure on $\widehat{Q}$ we can lift J to an almost complex structure $\widehat{J}$ on $\widehat{E}_{z_0}$ such that the derivatives of r_{x_0} are $(\widehat{J}, J)$-linear. The lift $\widehat{J}$ is unique and depends smoothly on J. By Lemma 15.2(3) any J-bubble $w : \mathbb{CP}^1 \longrightarrow E_{z_0}$ has a unique $\widehat{J}$-holomorphic lift $\hat{w} : \mathbb{CP}^1 \longrightarrow \widehat{E}_{z_0}$. We will not try to equip $\widehat{E}_{z_0}$ with a symplectic structure, since that is not necessary for our argument.

Now we can return to the usual strategy: let $\mathcal{P}(\widehat{E}_{z_0})$ be the space non-constant $W^{1,p}$-maps from $\mathbb{CP}^1$ to $\widehat{E}_{z_0}$. For $\hat{w} \in \mathcal{P}(\widehat{E}_{z_0})$ and a $\widehat{J}$ as above, let $\Lambda^{0,1}_{\widehat{J}}(\hat{w}^*T\widehat{E}_{z_0}) \subset \mathrm{Hom}(T\mathbb{CP}^1, \hat{w}^*T\widehat{E}_{z_0})$ be the vector bundle of $\widehat{J}$-antilinear homomorphisms $T\mathbb{CP}^1 \longrightarrow \hat{w}^*T\widehat{E}_{z_0}$. The spaces of L^p-sections of these vector bundles form a bundle of Banach spaces

$$\Lambda_{\widehat{J}} \longrightarrow \mathcal{P}(\widehat{E}_{z_0})$$

which has a canonical section $\bar{\partial}_{\widehat{J}}$, given by $\bar{\partial}_{\widehat{J}}(\hat{w}) = D\hat{w} + \widehat{J} \circ D\hat{w} \circ i$. $\bar{\partial}_{\widehat{J}}^{-1}(0)$ is the space of non-constant $\widehat{J}$-holomorphic maps. We call such a map regular if it is a regular zero of $\bar{\partial}_{\widehat{J}}$, that is, if the differential

$$D\bar{\partial}_{\widehat{J}}(\hat{w}) : T_w\mathcal{P}(\widehat{E}_{z_0}) \longrightarrow L^p(\Lambda^{0,1}_{\widehat{J}}(\hat{w}^*T\widehat{E}_{z_0}))$$

is surjective. This differential is a Fredholm operator of index

$$\mathrm{ind} D\bar{\partial}_{\widehat{J}}(\hat{w}) = 4 + 2\langle c_1(T\widehat{E}, \widehat{J}), [\hat{w}]\rangle. \tag{15.1}$$

This is just the ordinary transversality theory of pseudo-holomorphic curves on an almost complex four-manifold. The only difference is that we consider only almost complex structures which come from an almost complex structure J on E. These almost complex structures are not generic: since they agree with the lift of J' on a neighbourhood of $D_{x_0} = r_{x_0}^{-1}(x_0)$, D_{x_0} is a $\widehat{J}$-holomorphic sphere with self-intersection (-2). Such a sphere is not a regular $\widehat{J}$-holomorphic map. However, the space of almost complex structures $\widehat{J}$ is sufficiently large to prove the following result:

Lemma 15.4. *Let $\mathcal{J}_\epsilon \subset \mathcal{J}(E, \Omega, J'; J^-, J^+)$ be a C^∞_ϵ-neighbourhood. There is a subset $\mathcal{J}_{\epsilon,\mathrm{reg}} \subset \mathcal{J}_\epsilon$ of second category such that for $J \in \mathcal{J}_{\epsilon,\mathrm{reg}}$, any simple $\widehat{J}$-holomorphic map $\hat{w} : \mathbb{C}\mathrm{P}^1 \longrightarrow \widehat{E}_{z_0}$ which does not lie in the exceptional curve D_{x_0} is regular.*

Because of the usual $PSL(2, \mathbb{C})$-action, there can be no simple regular $\widehat{J}$-holomorphic spheres of index < 6. (15.1) implies that if $J \in \mathcal{J}_{\epsilon,\mathrm{reg}}$, any simple $\widehat{J}$-holomorphic sphere has positive Chern number, with the exception of those spheres which lie in D_{x_0}. By passing to the underlying simple holomorphic map, the same can be proved for all non-constant J-holomorphic maps $\hat{w} : \mathbb{C}\mathrm{P}^1 \longrightarrow \widehat{E}_{z_0}$, again with the same exception. In particular, if w is a J-bubble in E_{z_0}, its lift $\hat{w}$ to $\widehat{E}_{z_0}$ has positive Chern number ($\mathrm{im}(\hat{w}) \neq D_{x_0}$ because r_{x_0} contracts D_{x_0} to the point x_0). To prove Proposition 15.1, it remains to show that the Chern numbers of w and $\hat{w}$ coincide, and that is a consequence of the following computation:

Lemma 15.5. *$c_1(T\widehat{E}_{z_0}, \widehat{J}) = r_{x_0}^* c_1(TE, J)$.*

Proof. Let $N \subset \widehat{E}_{z_0}$ be a tubular neighbourhood of D_{x_0}. Consider the Mayer-Vietoris sequence

$$H^1(N \setminus D_{x_0}) \longrightarrow H^2(\widehat{E}_{z_0}) \longrightarrow H^2(\widehat{E}_{z_0} \setminus D_{x_0}) \oplus H^2(N).$$

Because D_{x_0} has self-intersection (-2), $N \setminus D_{x_0}$ is homotopy equivalent to $\mathbb{R}P^3$; in particular $H^1(N \setminus D_{x_0}) = 0$. Let $\delta = c_1(T\widehat{E}_{z_0}, \widehat{J}) - r_{x_0}^* c_1(TE, J) \in H^2(\widehat{E}_{z_0})$. The image of $r_{x_0}^* c_1(TE, J)$ in $H^2(N)$ vanishes because r_{x_0} collapses D_{x_0} to a point. On the other hand, $c_1(T\widehat{E}, \widehat{J})|N$ vanishes because D_{x_0} is a rational curve of self-intersection (-2) and therefore has Chern number zero. It follows that δ maps to zero in $H^2(N)$.

$r_{z_0} : \widehat{E}_{z_0} \setminus D_{x_0} \longrightarrow E_{z_0} \setminus \{x_0\}$ is a diffeomorphism and carries $\widehat{J}$ to J. Therefore

$$T\widehat{E}|(\widehat{E}_{z_0} \setminus D_{x_0}) \cong r_{z_0}^* T(E_{z_0} \setminus \{x_0\}). \tag{15.2}$$

$T(E_{z_0} \setminus \{x_0\})$ is the restriction of TE^v to $E_{z_0} \setminus \{x_0\}$ and hence

$$\begin{aligned} TE|E_{z_0} \setminus \{x_0\} &= (TE^h \oplus TE^v)|E_{z_0} \setminus \{x_0\} \\ &\cong \mathbb{C} \oplus T(E_{z_0} \setminus \{x_0\}). \end{aligned}$$

Together with (15.2) this shows that δ maps to zero in $H^2(\widehat{E}_{z_0} \setminus D_{x_0})$. $\square$

Part III

16 An outline of the argument

This final part contains the computation of the Floer homology of a generalized Dehn twist. The result was stated in Theorem 3.5 and its proof uses the tools introduced in Part II. The proof is based on three results. Two of these are general properties of Floer homology:

(Floer homology of the identity) $HF_*(\mathrm{id}_M)$ as a $QH_*(M,\omega)$-module is canonically isomorphic to $QH_*(M,\omega)$ as a module over itself.

(Mapping invariance) The quantum module product satisfies

$$c \,\hat{*}\, x = \phi_*(c) \,\hat{*}\, x$$

for all $c \in QH_*(M,\omega)$ and $x \in HF_*(\phi)$. Here ϕ_* is the action of ϕ on $QH_*(M,\omega) = H_*(M;\Lambda)$.

The isomorphism of $HF_*(M)$ and $H_*(M;\Lambda)$ as graded groups is fundamental for the original application of Floer homology to the Arnol'd conjecture; it was proved in successively more general versions by Floer [10], Hofer-Salamon [14] and Piunikhin-Salamon-Schwarz [22]. The last-mentioned paper contains the proof of the statement about the product structures. The 'mapping invariance' property is Proposition 10.2.

The third result which enters into the proof of Theorem 3.5 is the following exact sequence:

Theorem 16.1. *Let τ_V be the generalized Dehn twist along $V \subset M$. There is a homomorphism of $\mathbb{Z}/2$-graded $QH_*(M,\omega)$-modules $\Phi : HF_*(\mathrm{id}_M) \longrightarrow HF_*(\tau_V)$ which fits into a long exact sequence*

$$0 \longrightarrow HF_1(\mathrm{id}_M) \xrightarrow{\Phi} HF_1(\tau_V) \longrightarrow$$
$$\longrightarrow \Lambda \oplus \Lambda \longrightarrow HF_0(\mathrm{id}_M) \xrightarrow{\Phi} HF_0(\tau_V) \longrightarrow 0.$$

The homomorphism Φ is induced by a suitable almost holomorphic fibration. The construction of this fibration is based on the special role of generalized Dehn twists as 'symplectic monodromy maps', and the exact sequence is a consequence of a partial determination of Φ at chain level. Theorem 16.1 is, at least in principle, a special case of a more general long exact sequence for Floer homology groups; this will be the topic of a future publication.

Proof of Theorem 3.5. The exact sequence of Theorem 16.1 has the following consequence:

the kernel of Φ satisfies $\dim_\Lambda \ker(\Phi) \leq 2$. Moreover, if equality holds, $HF_*(\tau_V) \cong HF_*(\mathrm{id}_M)/\ker(\Phi)$ as modules over $QH_*(M,\omega)$.

Let $\Psi : QH_*(M,\omega) \longrightarrow HF_*(\mathrm{id}_M)$ be the canonical isomorphism, and $v = [V]t^0 \in QH_*(M,\omega)$. The next step is to prove

$$\Psi(v) \in \ker(\Phi). \tag{16.1}$$

To see this, take the unit element $u = [M]t^0 \in QH_*(M,\omega)$ and another element $w = W\,t^0$, where $W \in H_2(M;\mathbb{Z}/2)$ satisfies $[V]\cdot W = 1$. Because Φ is a homomorphism of $QH_*(M,\omega)$-modules,

$$\Phi(\Psi(w)) = \Phi(\Psi(w * u)) = \Phi(w \mathbin{\hat{*}} \Psi(u)) = w \mathbin{\hat{*}} \Phi(\Psi(u)).$$

Similarly $\Phi(\Psi(v{+}w)) = (v{+}w) \mathbin{\hat{*}} \Phi(\Psi(u))$. By the Picard-Lefschetz formula, $(\tau_V)_*(w) = w + (w \cdot_\Lambda v)v = w + v$. From the 'mapping invariance' property of the quantum module structure it follows that

$$w \mathbin{\hat{*}} \Phi(\Psi(u)) = (v + w) \mathbin{\hat{*}} \Phi(\Psi(u));$$

hence $\Phi(\Psi(v)) = \Phi(\Psi(v + w)) - \Phi(\Psi(w)) = 0$. Having proved (16.1) we can immediately strengthen it:

$$\Psi(I_v) \subset \ker(\Phi),$$

because $\Phi(\Psi(c * v)) = c \mathbin{\hat{*}} \Phi(\Psi(v)) = c \mathbin{\hat{*}} 0 = 0$ for any $c \in QH_*(M,\omega)$. Now $\dim_\Lambda I_v = 2$ by Lemma 3.2. Since Ψ is an isomorphism, this means that equality holds in the condition $\dim_\Lambda \ker(\Phi) \leq 2$ derived above and hence that $\ker(\Phi) = \Psi(I_v)$. In this case, the exact sequence yields an isomorphism

$$HF_*(\tau_V) \cong HF_*(\mathrm{id}_M)/\Psi(I_v).$$

Clearly this implies that $HF_*(\tau_V) \cong QH_*(M,\omega)/I_v$ as $QH_*(M,\omega)$-modules, with the quantum module structure on $HF_*(\tau_V)$ and the one induced by the quantum product on $QH_*(M,\omega)/I_v$. □

It remains to prove the exact sequence in Theorem 16.1. The proof takes up the rest of this part. It is based on a simple observation about the energy of J-holomorphic sections of bundles which satisfy a certain 'curvature property'; this idea is presented in the next section. The two following sections describe the almost holomorphic fibration which induces Φ; its construction is elementary but lengthy. The final section contains a technical transversality result for J-holomorphic sections in the spirit of section 13.

17 Nonnegative fibrations

Let (E, Ω, J') be an almost holomorphic fibration over a Riemann surface Σ, with projection $\pi : E \longrightarrow \Sigma$.

Definition 17.1. Let J be a partially Ω-tame almost complex structure on E. J is called *horizontal* if $J(TE^h_x) = TE^h_x$ for all $x \in E$ (note that for $x \in \mathrm{Crit}(\pi)$ this condition is vacuous because $TE^h_x = 0$).

Let us assume for a moment that π has no critical points, that is, that (E, Ω) is a symplectic fibre bundle. We know that $(E_z, \Omega|E_z)$ is a locally trivial family of symplectic manifolds. Taken together with Darboux's theorem, this says that the vertical component $\Omega|TE^v$ does not have any local differential-geometric invariants. In contrast, the local geometry of the whole form Ω is nontrivial. The importance of the class of horizontal almost complex structures is that it is linked with TE^h and hence with the differential geometry of Ω, whereas the class of partially Ω-tame almost complex structures depends only on $\Omega|TE^v$.

Definition 17.2. (E, Ω, J') is *nonnegative* if it satisfies one of the following two equivalent conditions:

(i) Let $\beta \in \Omega^2(\Sigma)$ be a positively oriented volume form. For every $x \in E \setminus \mathrm{Crit}(\pi)$ there is a $\rho(x) \geq 0$ such that

$$\Omega|TE^h_x = \rho(x)(\pi^*\beta|TE^h_x)$$

(this is obviously independent of the choice of β).

(ii) Any horizontal almost complex structure J on E satisfies

$$\Omega(X, JX) \geq 0 \text{ for all } X \in TE.$$

Moreover, $\Omega(X, JX) = 0$ implies that $X \in TE^h$.

The equivalence of the two conditions is proved as follows:

(i) $\Rightarrow$ (ii) Let J be a horizontal almost complex structure. For $x \in \mathrm{Crit}(\pi)$ we have $\Omega(X, JX) > 0$ for all nonzero $X \in TE_x$ because J is partially Ω-tame and $TE_x = TE^v_x$.

Take a point $x \notin \mathrm{Crit}(\pi)$ and an $X \in TE_x$, and let $X = X^v + X^h$ be its vertical and horizontal parts.

$$\begin{aligned}\Omega(X, JX) &= \Omega(X^v + X^h, J(X^v + X^h)) = \Omega(X^v, JX^v) + \Omega(X^h, JX^h) \\ &= \Omega(X^v, JX^v) + \rho(x)\beta(D\pi(X^h), j\, D\pi(X^h)).\end{aligned}$$

The second term is nonnegative because β is a positive volume form and $\rho(x) \geq 0$. The first term is nonnegative (because J is partially Ω-tame) and vanishes iff $X^v = 0$.

(ii) $\Rightarrow$ (i) Fix a point $x \notin \mathrm{Crit}(\pi)$. TE^h_x is two-dimensional and $D\pi_x|TE^h_x : TE^h_x \longrightarrow T\Sigma_{\pi(x)}$ is an isomorphism. Hence there is always a $\rho(x) \in \mathbb{R}$ such that $\Omega|TE^h_x = \rho(x)(\pi^*\beta|TE^h_x)$. Choose a nonzero $X \in TE^h_x$ and a horizontal almost complex structure J. By assumption

$$0 \leq \Omega(X, JX) = \rho(x)\beta(D\pi(X), j\, D\pi(X)).$$

Because β is a positive volume form and $D\pi(X) \neq 0$, $\beta(D\pi(X), j\, D\pi(X))$ is positive. Therefore $\rho(x) \geq 0$. □

Let us return briefly to the situation where $\mathrm{Crit}(\pi) = \emptyset$. We have pointed out that the local geometry of Ω is nontrivial. The principal local invariant of Ω is its Hamiltonian curvature r. This is a family of homomorphisms

$$r_z : \Lambda^2(T\Sigma_z) \longrightarrow C^\infty(E_z, \mathbb{R})$$

parametrized by $z \in \Sigma$, which is defined simply by $r_z(Z_1, Z_2) = -\Omega(Z_1^\natural, Z_2^\natural)$. The name 'Hamiltonian curvature' can be explained as follows: the curvature of the symplectic connection TE^h on the fibre bundle E is a two-form on Σ with values in the symplectic vector fields on the fibres E_z; that is, for $Z_1, Z_2 \in (T\Sigma)_z$ the curvature $R(Z_1, Z_2)$ is a symplectic vector field on $(E_z, \Omega|E_z)$. With our sign conventions, $R(Z_1, Z_2)$ turns out to be the Hamiltonian vector field associated to $r(Z_1, Z_2) \in C^\infty(E_z, \mathbb{R})$.

The nonnegativity condition says that $r(Z_1, Z_2)$ is a *nonpositive* function on E_z for a positively oriented basis Z_1, Z_2 of $T\Sigma_z$. More succinctly, one could say that (E, Ω) is nonnegative iff it has nonpositive Hamiltonian curvature. This clash of signs is unfortunate; we have chosen to retain the name 'nonnegative fibration' because it is more intuitive.

Lemma 17.3. *Let J be a horizontal almost complex structure on E.*

(a) A horizontal section $\sigma : \Sigma \longrightarrow E$ is J-holomorphic for any horizontal almost complex structure J.

(b) If (E, Ω, J') is nonnegative, any J-holomorphic section has nonnegative energy, and any J-holomorphic section with zero energy is horizontal.

Proof. (a) $D\pi : \sigma^*(TE^h, J|TE^h) \longrightarrow (T\Sigma, j)$ is an isomorphism and $D\sigma$ is its inverse. Since $D\pi$ is (J, j)-linear, $D\sigma$ is (j, J)-linear.

(b) Let σ be a J-holomorphic section. The definition of nonnegativity in terms of J shows that $\sigma^*\Omega(Z, jZ) = \Omega(D\sigma(Z), J\, D\sigma(Z)) \geq 0$ for any $Z \in T\Sigma$, and that equality can hold only if $D\sigma(Z) \in TE^h$. Hence $\sigma^*\Omega$ is a nonnegative two-form, and it vanishes only if $D\sigma(Z) \in TE^h$ for all Z. □

Example 17.4. It is instructive to review some of the results of section 8 from the present point of view. Let (E, Ω) be a product fibre bundle $\mathbb{R} \times (T, \Theta)$. The almost complex structures in $\mathcal{J}(T, \Theta)$ are all horizontal, and Lemma 8.7 proves that (E, Ω) is nonnegative. In fact, it is *flat*, that is, $\Omega|TE^h \equiv 0$.

Example 17.5. Now consider the same smooth fibre bundle E, but with a perturbed two-form $\Omega' = \Omega - d(H\,dt)$ as in section 7. Let K be the vector field which generates the group of translations of E. Take the standard basis $Z_1 = (1,0), Z_2 = (0,1)$ of vector fields on C and let $Z_1^\natural, Z_2^\natural$ be their horizontal lifts with respect to Ω'. Ω' is nonnegative iff $\Omega'(Z_1^\natural, Z_2^\natural) \geq 0$. Because Z_2 is horizontal and $Z_1^\natural - K$ is vertical,

$$\Omega'(Z_1^\natural, Z_2^\natural) = \Omega'(K, Z_2^\natural) = \Omega(K, Z_2^\natural) - (dH \wedge dt)(K, Z_2^\natural).$$

Since $i_K\Omega = 0$, $dt(K) = 0$, and $dt(Z_2^\natural) = 1$, it follows that

$$\Omega'(Z_1^\natural, Z_2^\natural) = -(dH \wedge dt)(K, Z_2^\natural) = -dH(K).$$

If we see H as a family $(H_s)_{s\in\mathbb{R}}$ of functions on T, $dH(K) = \partial_s H_s$. Hence (E, Ω') is nonnegative iff $\partial_s H_s \leq 0$.

For example, let H be the pullback of function $\overline{H}$ on C which is constant outside a compact subset. (E, Ω') is nonnegative if $\partial\overline{H}/\partial s \leq 0$. It is not difficult to see that Ω and Ω' induce the same symplectic connection on E. However, the horizontal sections of E, which have zero energy with respect to Ω, have energy $-\int_C \partial\overline{H}/\partial s$ with respect to Ω'. This example illustrates the fact that the last sentence of Lemma 17.3 does not have a converse: horizontal sections of a nonnegative fibration may have positive energy.

From now on we assume that $\Sigma = C$ and that (E, Ω, J') has tubular ends modelled on $(T^\pm, \Theta^\pm)$. We will use the following simple properties of horizontal sections of E:

Lemma 17.6. *(a) Any horizontal section of E is a section with horizontal limits.*

(b) Two horizontal sections of E with the same positive (or negative) limit coincide.

(c) Let σ be a horizontal section of E such that $\Omega|TE^h$ vanishes in a neighbourhood of $\mathrm{im}(\sigma)$. *Then the canonical connection ∇^σ on σ^*TE^v is flat.*

(d) Let σ be as in (c) and assume that $(T^\pm, \Theta^\pm)$ are nondegenerate. Then $\mathrm{ind}(\sigma) = 0$.

Proof. (a) Because (E, Ω) has tubular ends, there is an $R > 0$ such that the restriction of a horizontal section of E to $[R; \infty) \times S^1$ is a horizontal section of $[R; \infty) \times (T^+, \Theta^+)$. As shown in Lemma 8.2, any such section is given by

$$\sigma(s,t) = (s, \nu^+(t)), \tag{17.1}$$

where ν^+ is a horizontal section of (T^+, Θ^+). The same holds on the other end.

(b) Equation (17.1) shows that the behaviour of a horizontal section for large s is determined by its positive limit. In particular, two horizontal sections with the same positive limit coincide on an open subset of C. Horizontal sections satisfy a strong unique continuation condition (two horizontal sections which have equal values at a single point coincide everywhere). Hence two distinct horizontal sections with the same positive (or negative) limit cannot exist.

(c) We have defined ∇^σ by

$$\nabla^\sigma_X W = [X^\natural, \widetilde{W}] \tag{17.2}$$

where $\widetilde{W}$ is some extension of W from $\sigma^* TE^v$ to all of TE^v. By using the Jacobi identity it follows that

$$\nabla^\sigma_X \nabla^\sigma_Y W - \nabla^\sigma_Y \nabla^\sigma_X W = [[X^\natural, Y^\natural], \widetilde{W}].$$

Therefore ∇^σ is flat if

$$[X^\natural, Y^\natural] = [X, Y]^\natural$$

in some neighbourhood of $\mathrm{im}(\sigma)$. By a general fact about the Lie bracket on smooth fibre bundles, $[X^\natural, Y^\natural]$ is a lift of $[X, Y]$; it remain to show that $[X^\natural, Y^\natural]$ is horizontal in a neigbourhood of $\mathrm{im}(\sigma)$. Because Ω is closed,

$$0 = d\Omega(X^\natural, Y^\natural, Z) = Z.\Omega(X^\natural, Y^\natural) - \Omega([X^\natural, Y^\natural], Z)$$

for any $Z \in C^\infty(TE^v)$ (all other terms in the standard formula for $d\Omega$ vanish because $X^\natural$ and $Y^\natural$ are horizontal and Z is vertical). By assumption, $Z.\Omega(X^\natural, Y^\natural) = 0$ in a neighbourhood of $\mathrm{im}(\sigma)$; therefore $\Omega([X^\natural, Y^\natural], Z) = 0$ in that neighbourhood, that is, $[X^\natural, Y^\natural]$ is horizontal there.

(d) Recall that the definition of the index goes as follows: one has to replace σ by a section σ' which is horizontal outside a compact subset. $(\sigma')^* TE^v$ has a canonical symplectic connection $\nabla^{\sigma'}$ defined again outside some compact subset in C. Extend this connection on all of C and consider its monodromy around the circles $\{s\} \times S^1$. This yields a path in the symplectic group whose Maslov index is the index of σ.

This procedure can be simplified considerably if σ is horizontal: one can take $\sigma' = \sigma$, and the canonical connection ∇^σ is defined on all of C. In our case ∇^σ is flat and hence its monodromy around $\{s\} \times S^1$ is independent of s in a suitable trivialization of $\sigma^* TE^v$. A constant path in $\mathrm{Sp}(4, \mathbb{R})$ has zero Maslov index; therefore $\mathrm{ind}(\sigma) = 0$. □

Theorem 17.7. *Let (E, Ω, J') be an ordinary almost holomorphic fibration over C, with tubular ends modelled on nondegenerate fibre bundles $(T^\pm, \Theta^\pm)$. Choose $J^- \in \mathcal{J}_{\mathrm{reg}}(T^-, \Theta^-)$ and $J^+ \in \mathcal{J}_{\mathrm{reg}}(T^+, \Theta^+)$. We assume that (E, Ω, J') is nonnegative and that any $\nu^+ \in \mathcal{H}(T^+, \Theta^+)$ has the following* *property: there is a horizontal section σ which has ν^+ as its positive limit and such that $\Omega|TE^h_x$ vanishes for all x in a neighbourhood of $\mathrm{im}(\sigma)$. Then the homomorphism*

$$C\Phi(E,\Omega,J';J) : CF_*(T^-,\Theta^-) \longrightarrow CF_*(T^+,\Theta^+)$$

is surjective for any $J \in \mathcal{J}_{\mathrm{reg}}(E,\Omega,J';J^-,J^+)$ which is horizontal.

The principle underlying this Theorem was first used by Floer [11] to prove his exact sequence for instanton Floer homology; in the expository paper [4] (where it is called the 'monotonicity property') it is formulated as an observation about filtered chain complexes. In our case, thanks to the use of the Novikov field Λ, this observation takes on a simpler form.

Let $A : \Lambda^m \longrightarrow \Lambda^n$ be a Λ-linear map. The most obvious description of such a map is as an $(n \times m)$-matrix with entries in Λ, but we prefer to write it as a formal power series

$$A = \sum_{\epsilon\in\mathbb{R}} A_\epsilon\, t^\epsilon$$

whose coefficients A_ϵ are homomorphisms $(\mathbb{Z}/2)^m \longrightarrow (\mathbb{Z}/2)^n$ and satisfy the usual finiteness condition

$$\#\{\epsilon \in \mathbb{R} \mid \epsilon \leq C \text{ and } A_\epsilon \neq 0\} < \infty \text{ for all } C \in \mathbb{R}. \tag{17.3}$$

A is called nonnegative if $A_\epsilon = 0$ for all $\epsilon < 0$. It is called positive if $A_\epsilon = 0$ for all $\epsilon \leq 0$. Note that these are properties of homomorphisms of *based* Λ-vector spaces; they are not invariant under base change.

Lemma 17.8. *If $P : \Lambda^n \longrightarrow \Lambda^n$ is positive, $(\mathrm{id} - P)$ is invertible.*

Proof. Since the formal series

$$P = \sum_{\epsilon>0} P_\epsilon\, t^\epsilon$$

satisfies (17.3) there is a $\delta > 0$ such that $P_\epsilon = 0$ for $\epsilon < \delta$. It is easy to see that the formal series

$$\mathrm{id} + P + P^2 + \cdots = \sum_{k=0}^{\infty} \Big(\sum_{\epsilon\geq\delta} P_\epsilon\, t^\epsilon \Big)^k$$

also satisfies (17.3). The endomorphism of Λ^n defined by this series is the inverse of $(\mathrm{id} - P)$. □

Lemma 17.9. *Let $A : \Lambda^m \longrightarrow \Lambda^n$ be a nonnegative homomorphism. If its leading coefficient A_0 is surjective (or injective), so is A itself.*

Proof. If A_0 is onto, there is a $D_0 : (\mathbb{Z}/2)^n \longrightarrow (\mathbb{Z}/2)^m$ with $A_0 \circ D_0 = \mathrm{id}$. Let $D = D_0\, t^0 : \Lambda^n \longrightarrow \Lambda^m$.

$$A \circ D = \mathrm{id}_{(\mathbb{Z}/2)^n}\, t^0 + \sum_{\epsilon>0} (A_\epsilon \circ D_0)\, t^\epsilon;$$

Lemma 17.8 shows that $A \circ D$ is invertible, and this implies that A is onto. The parallel statement about injectivity follows by taking the duals. □

Proof of Theorem 17.7. Let V^- and V^+ be the $\mathbb{Z}/2$-vector spaces whose bases are the sets $\mathcal{H}(T^-, \Theta^-)$ resp. $\mathcal{H}(T^+, \Theta^+)$. $C\Phi(E, \Omega, J'; J)$ can be written as a power series

$$C\Phi(E, \Omega, J'; J) = \sum_\epsilon C\Phi_\epsilon\, t^\epsilon$$

with coefficients $C\Phi_\epsilon \in \mathrm{Hom}(V^-, V^+)$. By definition, $C\Phi_\epsilon$ is given by the matrix

$$(m_\epsilon(J; \nu^-, \nu^+))_{\nu^+, \nu^-}$$

which encodes the asymptotic behaviour of J-holomorphic sections with index 0 and energy ϵ. By Lemma 17.3 and Lemma 17.6(d),

(a) $C\Phi_\epsilon = 0$ for $\epsilon < 0$ and

(b) $C\Phi_0$ counts precisely the horizontal sections with zero energy.

Part (a) says that $C\Phi(E, \Omega, J'; J)$ is nonnegative with respect to the standard bases of $CF_*(T^-, \Theta^-)$ and $CF_*(T^+, \Theta^+)$. We will use (b) and the assumption about horizontal sections of E to prove that $C\Phi_0$ is onto. In view of Lemma 17.9 this completes the proof of Theorem 17.7.

Let ν^+ be a horizontal section of (T^+, Θ^+). By assumption there is a horizontal section $\sigma(\nu^+)$ of E which has ν^+ as its positive limit. Lemma 17.6 shows that $\sigma(\nu^+)$ is unique and that any horizontal section equals $\sigma(\nu^+)$ for some ν^+. Let $i(\nu^+) \in \mathcal{H}(T^-, \Theta^-)$ be the negative limit of $\sigma(\nu^-)$. Applying Lemma 17.6 again shows that $i : \mathcal{H}(T^+, \Theta^+) \longrightarrow \mathcal{H}(T^-, \Theta^-)$ is injective.

The assumption that $\Omega|TE^h$ vanishes in a neighbourhood of $\sigma(\nu^+)$ implies that $\sigma(\nu^+)^*\Omega = 0$. It follows that

$$m_0(J; \nu^-, \nu^+) = \begin{cases} 1 & \nu^- = i(\nu^+), \\ 0 & \text{otherwise.} \end{cases}$$

In other words, $C\Phi_0$ maps the element of the natural basis of V^- corresponding to $i(\nu^+)$ to the basis element of V^+ corresponding to ν^+, for all $\nu^+ \in \mathcal{H}(T^+, \Theta^+)$. Clearly, this means that $C\Phi_0$ is onto. □

Theorem 17.7 is complemented by a technical result whose proof we postpone to section 21:

Theorem 17.10. *Let (E, Ω, J') be as in Theorem 17.7. For all $J^\pm \in \mathcal{J}_{\mathrm{reg}}(T^\pm, \Theta^\pm)$ there is a $J \in \mathcal{J}_{\mathrm{reg}}(E, \Omega, J'; J^-, J^+)$ which is horizontal.*

Corollary 17.11. *Let (E, Ω, J'), $J^\pm$ and J be as in Theorem 17.7 and assume that $\dim_\Lambda CF_1(T^-, \Theta^-) = \dim_\Lambda CF_1(T^+, \Theta^+)$. Then $\Phi(E, \Omega, J')$ fits into a long exact sequence*

$$0 \longrightarrow HF_1(T^-, \Theta^-; J^-) \xrightarrow{C\Phi_1(E,\Omega,J';J)} HF_1(T^+, \Theta^+; J^+) \longrightarrow$$
$$\longrightarrow \Lambda^d \longrightarrow HF_0(T^-, \Theta^-; J^-) \xrightarrow{C\Phi_0(E,\Omega,J';J)} HF_0(T^+, \Theta^+; J^+) \longrightarrow 0,$$

where $d = \dim_\Lambda CF_0(T^-, \Theta^-) - \dim_\Lambda CF_0(T^+, \Theta^+)$.

Proof. Because $\dim_\Lambda CF_1(T^-, \Theta^-) = \dim_\Lambda CF_1(T^+, \Theta^+)$, the subcomplex $\ker(C\Phi(E, \Omega, J'; J)) \subset CF_*(T^-, \Theta^-)$ is zero in degree 1. The long exact sequence is induced by the sequence of chain complexes

$$0 \longrightarrow \ker(C\Phi(E, \Omega, J; J')) \hookrightarrow CF_*(T^-, \Theta^-) \longrightarrow$$
$$\xrightarrow{C\Phi(E,\Omega,J;J')} CF_*(T^+, \Theta^+) \longrightarrow 0.$$

□

This is the mechanism which we will use to produce the exact sequence of Theorem 16.1.

18 The quadratic fibration

Let (E, Ω, J') be an almost holomorphic fibration over the closed disc $D_\delta \subset \mathbb{C}$ of radius $\delta > 0$, such that $0 \in D_\delta$ is the only critical value of $E \longrightarrow D_\delta$. The symplectic parallel transport $P_{\partial D_\delta} \in \mathrm{Aut}(E_\delta, \Omega|E_\delta)$ along ∂D_δ is called the symplectic monodromy of (E, Ω, J').

In this section we will study the monodromy of a particularly simple fibration. This fibration will not be used later. However, a modified version of it, which will be introduced in the next section, is essential for the proof of Theorem 16.1. Many arguments can be carried out in parallel for both versions, and we prefer to present them in the simpler case.

Our example is $E = \pi^{-1}(D_\delta) \subset \mathbb{C}^3$, where $\pi : \mathbb{C}^3 \longrightarrow \mathbb{C}$ is the holomorphic function $\pi(x) = x_1^2 + x_2^2 + x_2^3$ and δ is some positive number. Ω and J' are the standard symplectic and complex structures on E. $\pi : E \longrightarrow D_\delta$ has a single critical point $0 \in \mathbb{C}^3$. Strictly speaking, (E, Ω, J') is not an almost holomorphic fibration over D_δ because its fibres are noncompact. In general, such a lack of compactness causes a problem because the symplectic monodromy may not exist. We will ignore this problem, that is, we will proceed as if we knew that (E, Ω, J') has a symplectic monodromy and then compute this monodromy explicitly. To put the argument on a strictly sound basis one would have to reverse it by starting with the explicit formula and working backwards to show that this formula describes the monodromy of (E, Ω, J'). This is perfectly possible; however, it would obscure the argument.

We begin by identifying the regular fibres of E.

Lemma 18.1. *The restriction of* (E, Ω) *to* $(0; \delta] \subset D_\delta$ *is isomorphic to the trivial symplectic fibre bundle* $(0; \delta] \times (T^*S^2, \eta)$. *An explicit isomorphism* $f : E|(0; \delta] \longrightarrow (0; \delta] \times T^*S^2$ *is given by*

$$f(x) = \Big(\pi(x), \frac{\mathrm{re}(x)}{|\mathrm{re}(x)|}, -\mathrm{im}(x)|\mathrm{re}(x)|\Big). \tag{18.1}$$

In (18.1) we view $(0; \delta] \times T^*S^2$ as a subset of $\mathbb{R} \times \mathbb{R}^3 \times \mathbb{R}^3$ by using the same coordinates on T^*S^2 as in section 2.

Proof. It is convenient to separate the real and imaginary parts $x = p + iq$ of $x \in \mathbb{C}^3$. The equation $\pi(x) = s \in (0; \delta]$ translates into

$$|p|^2 - |q|^2 = s, \quad \langle p, q\rangle = 0. \tag{18.2}$$

The first equation implies that $p \neq 0$ and this shows that (18.1) defines a smooth map $E|(0; \delta] \longrightarrow (0; \delta] \times \mathbb{R}^3 \times \mathbb{R}^3$. The second equation implies that $\mathrm{im}(f)$ lies in $(0; \delta] \times T^*S^2 \subset (0; \delta] \times \mathbb{R}^3 \times \mathbb{R}^3$. We will write η for the standard symplectic structure on T^*S^2 and for its pullback to $(0; \delta] \times T^*S^2$; this should not cause any confusion.

$$\begin{aligned} f^*\eta &= d\Big(f^*\Big(-\sum_{j=1}^{3} u_j dv_j\Big)\Big) = d\Big(\sum_{j=1}^{3} \frac{p_j}{|p|} d(q_j|p|)\Big) \\ &= \sum_{j=1}^{3} dp_j \wedge dq_j + d\Big(\frac{d|p|}{|p|}\langle p, q\rangle\Big); \end{aligned}$$

the second term vanishes by (18.2). Therefore $f^*\eta$ equals the restriction of Ω to $E|(0; \delta]$. This is nearly sufficient to prove that f is an isomorphism of symplectic fibre bundles; an additional consideration is necessary because of the non-compactness of the fibres.

Let $f_s : E_s \longrightarrow T^*S^2$ be the restriction of f to the fibre over $s \in (0; \delta]$. We have proved that $f_s^*\eta = \Omega|E_s$; therefore f_s is a local diffeomorphism. For all $R > 0$ we have

$$f_s^{-1}(T_R^*S^2) = \{p + iq \in E_s \mid |p|\,|q| \leq R\}.$$

Using the first part of (18.2) it is easy to see that a bound on $|p|\,|q|$ implies a bound on $|p|^2 + |q|^2$; therefore $f^{-1}(T_R^*S^2)$ is compact. This proves that f_s

is proper. A proper local diffeomorphism is a covering map; since any point in $S^2 \subset T^*S^2$ has a unique preimage (this can be easily proved by looking at the formula for f_s) this covering map is a diffeomorphism.

We have proved that f is a fibrewise diffeomorphism and that $f^*\eta = \Omega$; therefore f is an isomorphism of symplectic fibre bundles over $(0;\delta]$. $\square$

Lemma 18.2. *Let $H \in C^\infty(E,\mathbb{R})$ be the function given by $H(x) = \frac{1}{4}|x|^2$. $h = H \circ f^{-1} \in C^\infty((0;\delta] \times T^*S^2,\mathbb{R})$ is given by*

$$h(s,u,v) = \tfrac{1}{4}\sqrt{s^2 + 4\,|v|^2}.$$

Proof. Using equation (18.2) one checks that

$$\begin{aligned} h(f(p+iq)) &= \tfrac{1}{4}\sqrt{s^2+4|p|^2|q|^2} = \\ &= \tfrac{1}{4}\sqrt{(|p|^2-|q|^2)^2+4|p|^2|q|^2} = \tfrac{1}{4}(|p|^2+|q|^2) = H(p+iq). \quad \square \end{aligned}$$

Lemma 18.3. *(E,Ω,J') is a nonnegative almost holomorphic fibration.*

The simplest proof of this is based on the fact that (E,Ω,J') is a Kähler manifold. We choose a different proof which avoids using J'.

Proof. We must prove that for any $x \in E \setminus \{0\}$, $\Omega|TE^h_x$ is a nonnegative multiple of the pullback of the standard volume form on D_δ to TE^h_x. Since the set of points with this property is closed, it is sufficient to prove this for $x \in \pi^{-1}(D_\delta \setminus 0)$.

Let $X^\natural, Y^\natural$ be the horizontal lifts of the vector fields $X(z) = z$ and $Y(z) = iz$ on D_δ. They are defined on $E \setminus \{0\}$. Since X and Y form an oriented basis of the tangent space at any point in $D_\delta \setminus 0$, what we have to prove is that

$$\Omega(X^\natural, Y^\natural) \geq 0 \tag{18.3}$$

at any point in $\pi^{-1}(D_\delta \setminus \{0\}) \subset E$.

Let σ_E be the standard circle action on $\mathbb{C}^3$ and K_E the vector field which generates it. Ω is invariant under σ_E and $\pi(\sigma(e^{it})x) = e^{2it}\pi(x)$. Because of this symmetry, it is sufficient to show that (18.3) holds on $\pi^{-1}((0;\delta]) \subset \pi^{-1}(D_\delta \setminus \{0\})$.

Since $X^\natural$ is horizontal, $\Omega(X^\natural, Y^\natural) = \Omega(X^\natural, Y')$ for any vector field Y' which is a lift of Y. In particular, one can take $Y' = \frac{1}{2}K_E$. Since $\Omega(\cdot, \frac{1}{2}K_E) = dH$ with H as in Lemma 18.2, we have

$$\Omega(X^\natural, Y^\natural) = X^\natural.H.$$

Now we transfer the whole situation to $(0;\delta] \times T^*S^2$ using f. $f_*X^\natural$ is a horizontal vector field on the trivial symplectic fibre bundle $(0;\delta] \times (T^*S^2,\eta)$

and has the form $(f_*X^\natural)_{(s,u,v)} = (s,\cdots)$; hence $(f_*X^\natural)_{(s,u,v)} = (s,0,0)$. It follows that

$$(f_*X^\natural).(H\circ f^{-1}) = s\frac{\partial}{\partial s}(H\circ f^{-1}).$$

In Lemma 18.2 we gave an explicit formula for $h = H\circ f^{-1}$; this formula shows that $\partial h/\partial s > 0$ everywhere. □

Lemma 18.4. *The symplectic monodromy of (E,Ω) is equal to*

$$a\circ\phi_{-2\pi}^{H|E_\delta} \in \mathrm{Aut}(E_\delta,\Omega|E_\delta),$$

where a is the involution of E_δ given by $a(x) = -x$ and $(\phi_t^{H|E_\delta})_{t\in\mathbb{R}}$ is the Hamiltonian flow on $(E_\delta,\Omega|E_\delta)$ induced by $H|E_\delta$.

Proof. Let $Y^\natural$ be the horizontal lift of the same vector field $Y(z) = iz$ as in the proof of the previous Lemma. The map which we seek to determine is the time-(2π) map of the flow $(\psi_t)_{t\in\mathbb{R}}$ induced by $Y^\natural$.

Since Y is invariant under rotations of D_δ, $Y^\natural$ is invariant under σ_E; hence

$$\widehat{\psi}_t = \sigma_E(e^{-\frac{it}{2}})\circ\psi_t$$

is a flow on $E\setminus E_0$. This flow is generated by the vector field $V = Y^\natural - \frac{1}{2}K_E$ and since $D\pi(V) = 0$, it maps each fibre to itself. $\Omega(\cdot,Y^\natural)|E_\delta = 0$ because $Y^\natural$ is horizontal, and $\Omega(\cdot,\frac{1}{2}K_E) = dH$. Therefore

$$i_V\Omega|E_\delta = d(H|E_\delta).$$

This means that $\widehat{\psi}_{2\pi}|E_\delta$ is the time-(-2π) map of the Hamiltonian flow generated by $H|E_\delta$. By definition we have $\psi_{2\pi} = \sigma_E(-1)\widehat{\psi}_{2\pi}$, and $\sigma_E(-1)$ is precisely the map a. □

Now we transport the monodromy map from E_δ to T^*S^2: define

$$\tau = f_\delta\circ(a\circ\phi_{-2\pi}^{H|E_\delta})\circ f_\delta^{-1}.$$

Since $f_\delta\circ a\circ f_\delta^{-1}$ is the antipodal involution A on T^*S^2,

$$\tau = A\circ\phi_{-2\pi}^{h|\{\delta\}\times T^*S^2}.$$

Recall from section 2 that $\mu(u,v) = |v|$ generates a Hamiltonian circle action σ on $T^*S^2\setminus S^2$ with $\sigma(-1) = A|T^*S^2\setminus S^2$. h is invariant under this circle action; therefore

$$\tau\,|\,T^*S^2\setminus S^2 = \phi_\pi^\mu\circ\phi_{-2\pi}^{h|\{\delta\}\times T^*S^2} = \phi_{2\pi}^k,$$

where $k = \frac{1}{2}\mu - (h|\{\delta\}\times T^*S^2)$. Since every point in S^2 is a critical point of $h|\{\delta\}\times T^*S^2$, $\tau|S^2 = A|S^2$. Using Lemma 18.2 we can write $k = r(\mu)$ with

$$r(t) = \tfrac{1}{2}t - \tfrac{1}{4}\sqrt{\delta^2+4t^2}.$$

We have now shown that

$$\tau(x) = \begin{cases} \phi_{2\pi}^{r(\mu)}(x) & x \notin S^2, \\ A(x) & x \in S^2. \end{cases}$$

This expression is the same as that which defines the local model of a generalized Dehn twist (see Lemma 2.1). The function r which occurs in our situation satisfies $r(-t) = r(t) - t$ and $\lim_{t\to\infty} r(t) = \lim_{t\to\infty} r'(t) = 0$, but it does not vanish for $t \gg 0$. Therefore τ is not compactly supported but only asymptotic to the identity at infinity. For this reason it narrowly misses being a generalized Dehn twist.

We end our discussion of (E, Ω, J') with the following observation:

Lemma 18.5. *τ is fixed point free.*

Proof. Using the formula (2.1) one sees that

$$\tau(x) = \begin{cases} \sigma(e^{2\pi i\, r'(\mu(x))})(x) & x \notin S^2, \\ A(x) & x \in S^2. \end{cases}$$

A is clearly fixed point free. σ is a free circle action, and

$$r'(t) = \frac{1}{2} - \frac{t}{\sqrt{\delta^2 + 4t^2}} \in (0; \tfrac{1}{2})$$

for all $t > 0$. □

19 Generalized Dehn twists as monodromy maps

Proposition 19.1. *Let (M, ω) be a compact symplectic four-manifold and $V \subset M$ a Lagrangian two-sphere. There is an ordinary almost holomorphic fibration (E_V, Ω_V, J'_V) over some disc D_δ whose only critical value is $0 \in D_\delta$, whose fibre over $\delta \in D_\delta$ is isomorphic to (M, ω) and whose symplectic monodromy is (for a suitable choice of isomorphism) a generalized Dehn twist along V.*

This is an analogue of the well-known fact that Dehn twists on surfaces occur as monodromy maps. The proof takes up the whole of this section. A more detailed statement of the result can be found in Proposition 19.10 below.

Like the definition of generalized Dehn twists, our proof of Proposition 19.1 is based on a local model. The quadratic fibration considered in the previous section already comes very near to being this local model; however, as we have seen, its monodromy is not quite a generalized Dehn twist. We will now introduce a modified version of it. The modification achieves several goals: first, the monodromy becomes a genuine generalized Dehn twist; secondly, the fibration becomes 'trivial at infinity'; and finally, the regular fibre shrinks to a small tubular neighbourhood of S^2 in T^*S^2. The same remark about the non-compactness of the fibres as in the previous section applies.

Fix an $\epsilon > 0$ and a function $\xi \in C^\infty(\mathbb{R}, \mathbb{R})$ such that $\xi|[0; \epsilon/2] = 1$, $\xi|[\epsilon; \infty) = 0$, and $\xi'(t) < 0$ for $t \in (\epsilon/2; \epsilon)$. For $\delta > 0$, we define

$$\widetilde{E} = \{(x, z) \in \mathbb{C}^3 \times D_\delta \mid |x|^2 < 2\epsilon \text{ and } x_1^2 + x_2^2 + x_3^2 = \xi(|x|^2)z\}.$$

Let $\tilde{\pi} : \widetilde{E} \longrightarrow D_\delta$ be the projection, $\widetilde{\Omega} \in \Omega^2(\widetilde{E})$ the pullback of the standard symplectic form on $\mathbb{C}^3$, and $\widetilde{J}'$ the standard complex structure on the subset

$$\begin{aligned} U &= \{(x, z) \in \widetilde{E} \mid |x|^2 < \epsilon/2\} \\ &= \{(x, z) \in \mathbb{C}^3 \times D_\delta \mid x_1^2 + x_2^2 + x_3^2 = z \text{ and } |x|^2 < \epsilon/2\}. \end{aligned}$$

Lemma 19.2. *If δ is sufficiently small, $(\widetilde{E}, \widetilde{\Omega}, \widetilde{J}')$ is an ordinary almost holomorphic fibration, and $(0, 0) \in \widetilde{E}$ is the only critical point of $\tilde{\pi}$.*

Proof. The 'Zariski tangent space' $T\widetilde{E}_{(x,z)}$ of $\widetilde{E}$ at a point (x, z) consists of all $(X, Z) \in \mathbb{C}^3 \times \mathbb{C}$ such that

$$2(x_1X_1 + x_2X_2 + x_3X_3) = \xi(|x|^2)Z + 2z\,\xi'(|x|^2)\,\mathrm{re}(\langle x, X\rangle). \tag{19.1}$$

$\widetilde{E}$ is smooth at (x, z) if $T\widetilde{E}_{(x,z)}$ has (real) dimension 6. A smooth point (x, z) is a regular point of $\tilde{\pi}$ if $T\widetilde{E}^v_{(x,z)} = T\widetilde{E}_{(x,z)} \cap (\mathbb{C}^3 \times 0)$ is four-dimensional.

Let $R \subset \widetilde{E}$ be the set of points (x, z) such that $z = 0$ or $|x|^2 \geq \epsilon$. At any point of R the equation (19.1) is $\mathbb{C}$-linear and (because $x \neq 0$ or $\xi(|x|)^2 \neq 0$ holds) nontrivial. Hence all points in R are smooth points of $\widetilde{E}$. Since the smoothness condition is open, the set of smooth points is a neighbourhood of R. It is not difficult to see that any neighbourhood of R contains $\tilde{\pi}^{-1}(D_{\delta'})$ for sufficiently small $\delta' > 0$. Therefore we can ensure that $\widetilde{E}$ is smooth by making δ smaller.

The proof that $(0, 0)$ is the only critical point of $\tilde{\pi}$ is along the same lines: any point in $R \setminus (0, 0)$ is a regular of point of $\tilde{\pi}$. The regularity condition is open, and a neighbourhood of $R \setminus (0, 0)$ contains $\tilde{\pi}^{-1}(D_{\delta'}) \setminus (0, 0)$ for sufficiently small δ'.

$\widetilde{\Omega}|T\widetilde{E}^v_x$ is nondegenerate for any $x \in R$ because $T\widetilde{E}^v_x$ is complex-linear. The nondegeneracy condition is open; by shrinking δ again we can achieve that $\widetilde{\Omega}|T\widetilde{E}^v_x$ is nondegenerate everywhere. The remaining condition (3) in Definition 7.5 is satisfied because $(U, \widetilde{\Omega}|U, \widetilde{J}')$ is a Kähler manifold. The fact that $(0, 0)$ is an ordinary critical point is obvious. □

From now on we assume that δ has been chosen such that the conclusions of Lemma 19.2 apply. $(\widetilde{E}, \widetilde{\Omega}, \widetilde{J}')$ has many points in common with the simpler quadratic fibration (E, Ω, J'). For instance, $\sigma_{\widetilde{E}}(e^{it})(x, z) = (e^{it}x, e^{2it}z)$ defines a circle action on $\widetilde{E}$ which preserves $\widetilde{\Omega}$. This action covers the double of the standard circle action on the base D_δ, and the vector field $K_{\widetilde{E}}$ which generates $\sigma_{\widetilde{E}}$ satisfies

$$\widetilde{\Omega}(\cdot, \tfrac{1}{2}K_{\widetilde{E}}) = d\widetilde{H}$$

where $\widetilde{H}(x, z) = \frac{1}{4}|x|^2$. Some of the arguments in the previous section use only these properties and carry over to $(\widetilde{E}, \widetilde{\Omega}, \widetilde{J}')$ unchanged. In this way one obtains

Lemma 19.3. *(a) Let $X^\natural$ be the horizontal lift of $X(z) = z$. $(\widetilde{E}, \widetilde{\Omega}, \widetilde{H})$ is nonnegative iff $X^\natural.\widetilde{H} \geq 0$ at any point in $\tilde{\pi}^{-1}(0; \delta]$.*

(b) The monodromy of $(\widetilde{E}, \widetilde{\Omega}, \widetilde{H})$ is

$$(\sigma_{\widetilde{E}}(-1)|\widetilde{E}_\delta) \circ \phi_{-2\pi}^{\widetilde{H}|\widetilde{E}_\delta}. \quad \square$$

The analogue of Lemma 18.1 for $\widetilde{E}$ is

Lemma 19.4. *The restriction of $(\widetilde{E}, \widetilde{\Omega})$ to $(0; \delta] \subset D_\delta$ is isomorphic to the trivial symplectic fibre bundle $(0; \delta] \times (T^*_\epsilon S^2, \eta)$. An explicit isomorphism $\tilde{f} : \widetilde{E}|(0; \delta] \longrightarrow (0; \delta] \times T^*_\epsilon S^2$ is given by*

$$\tilde{f}(x, z) = \left(z, \frac{\mathrm{re}(x)}{|\mathrm{re}(x)|}, -\mathrm{im}(x)|\mathrm{re}(x)|\right). \tag{19.2}$$

Proof. We write again $x = p + iq$. The real and imaginary parts of the equation

$$x_1^2 + x_2^2 + x_3^2 = s\,\xi(|x|^2)$$

with $s \in (0; \delta]$ are

$$|p|^2 - |q|^2 = s\,\xi(|p|^2 + |q|^2), \quad \langle p, q\rangle = 0. \tag{19.3}$$

The first equation implies that $p \neq 0$ because $p = 0 \Rightarrow$ (since $s\,\xi(|p|^2 + |q|^2)$ is nonnegative) $q = 0 \Rightarrow \xi(|p|^2 + |q|^2) = 1$, which leads to a contradiction. Together with the second equation, this shows that $\tilde{f}$ is a well-defined map from $\widetilde{E}|(0; \delta]$ to $(0; \delta] \times T^*S^2$.

If $\tilde{f}(p + iq) = (s, u, v)$,

$$\begin{aligned} |v|^2 = |p|^2|q|^2 &= \tfrac{1}{4}\left((|p|^2 + |q|^2)^2 - (|p|^2 - |q|^2)^2\right) \\ &= \tfrac{1}{4}\left(|p + iq|^2 - s^2\xi(|p + iq|^2)^2\right) \leq \tfrac{1}{4}|p + iq|^2. \end{aligned} \tag{19.4}$$

Since $|p + iq|^2 \leq 4\epsilon^2$, $|v| \leq \epsilon$. This shows that the image of $\tilde{f}$ lies in $(0; \delta] \times T^*_\epsilon S^2$.

Since $\tilde{f}$ is given by the same formula as the map f in the previous section, we have again

$$\tilde{f}^*\eta = \sum_{j=1}^{3} dp_j \wedge dq_j + d\Big(\frac{d|p|}{|p|}\langle p, q\rangle\Big).$$

The second term vanishes by (19.3) and therefore $\tilde{f}^*\eta = \widetilde{\Omega}|\tilde{\pi}^{-1}(0;\delta]$.

Let $\tilde{f}_s : \widetilde{E}_s \longrightarrow T^*_\epsilon S^2$ be the restriction of $\tilde{f}$ to the fibre over $s \in (0;\delta]$. The computation (19.4) shows that the equality $|v|^2 = \frac{1}{4}|p+iq|^2$ holds whenever $|p|^2 + |q|^2 \geq \epsilon$. Therefore

$$\tilde{f}_s^{-1}(T^*_{\epsilon'}S^2) \subset \{(x,s) \in \widetilde{E}_s \mid |x| \leq 2\epsilon'\}$$

for all $\epsilon' \in [\epsilon/2;\epsilon)$. Since the r.h.s. is a compact subset of $\widetilde{E}_s$, it follows that $\tilde{f}_s$ is proper. The rest of the argument is as in Lemma 18.1. □

Let

$$\begin{aligned}\widetilde{E}^{\mathrm{triv}} &= \{(x,z) \in \widetilde{E} \mid |x| \geq \epsilon\}\\ &= \{x \in \mathbb{C}^3 \mid \epsilon \leq |x|^2 < 2\epsilon \text{ and } x_1^2 + x_2^2 + x_3^2 = 0\} \times D_\delta.\end{aligned}$$

$(\widetilde{E}^{\mathrm{triv}}, \widetilde{\Omega}|\widetilde{E}^{\mathrm{triv}})$ is clearly a trivial symplectic fibre bundle over D_δ, and the complement of $\widetilde{E}^{\mathrm{triv}}$ in $\widetilde{E}$ is relatively compact. This means that $(\widetilde{E}, \widetilde{\Omega}, \widetilde{J}')$ is 'trivial at infinity'. More precisely, we have

Lemma 19.5. *The same expression as in* (19.2) *defines a diffeomorphism*

$$F : \widetilde{E}^{\mathrm{triv}} \longrightarrow D_\delta \times (T^*_\epsilon S^2 \setminus T^*_{\epsilon/2}S^2)$$

such that $F^*\eta = \widetilde{\Omega}|\widetilde{E}^{\mathrm{triv}}$.

This follows from the same arguments as the preceding Lemma; we omit the proof. The result can be phrased as follows: $(\widetilde{E}, \widetilde{\Omega}, \widetilde{J}')$ is an almost holomorphic fibration whose regular fibre is $T^*_\epsilon S^2$ and which contains a trivial subbundle $D_\delta \times (T^*_\epsilon S^2 \setminus T^*_{\epsilon/2}S^2)$. This will become important later when we glue together $\widetilde{E}$ and another symplectic fibre bundle along this trivial subbundle.

We can use $\tilde{f}$ to transfer questions about $\widetilde{E}$ to $(0;\delta] \times T^*_\epsilon S^2$. For example, from Lemma 19.3(a) one obtains that $(\widetilde{E}, \widetilde{\Omega}, \widetilde{J}')$ is nonnegative iff

$$\tilde{f}^{-1}(X^\natural.\widetilde{H}) = \partial/\partial s(\widetilde{H} \circ \tilde{f}^{-1}) \geq 0$$

for all $(s,u,v) \in (0;\delta] \times T^*_\epsilon S^2$. This shows that the function $\widetilde{h} = \widetilde{H} \circ \tilde{f}^{-1}$ has the same important role as its counterpart h from the last section. There is no simple explicit formula for $\widetilde{h}$, but we can approach it in following way:

Lemma 19.6. *$\widetilde{h}$ satisfies*

$$\beta(s, \widetilde{h}(s,u,v)) = |v|^2,$$

where $\beta(s,t) = \frac{1}{4}\left(16t^2 - s^2\xi(4t)^2\right)$.

Proof. If $(s,u,v) = f(p+iq,s)$ then $|v|^2 = |p|^2|q|^2$. Therefore it is sufficient to prove that

$$\beta(s, \tfrac{1}{4}(|p|^2 + |q|^2)) = |p|^2|q|^2$$

for any $(p+iq, s) \in \widetilde{E}|(0;\delta]$. By (19.3),

$$\begin{aligned}\beta(s, \tfrac{1}{4}(|p|^2 + |q|^2) &= \tfrac{1}{4}\left((|p|^2 + |q|^2)^2 - s^2\xi(|p|^2 + |q|^2)^2\right)\\ &= \tfrac{1}{4}\left((|p|^2 + |q|^2)^2 - (|p|^2 - |q|^2)^2\right) = |p|^2|q|^2. \quad \square\end{aligned}$$

Lemma 19.7. *$(\widetilde{E}, \widetilde{\Omega}, \widetilde{J}')$ is a nonnegative almost holomorphic fibration.*

Proof. As we have seen, this reduces to proving that $\partial_s\widetilde{h} \geq 0$. By Lemma 19.6

$$\frac{\partial\beta}{\partial s}(s, \widetilde{h}(s,u,v)) + \frac{\partial\beta}{\partial t}(s, \widetilde{h}(s,u,v))\frac{\partial\widetilde{h}}{\partial s}(s,u,v) = 0. \tag{19.5}$$

Because $\xi \geq 0$ and $\xi' \leq 0$, we have $\frac{\partial\beta}{\partial s}(s,t) = -\frac{1}{2}s\,\xi(4t)^2 \leq 0$ and $\frac{\partial\beta}{\partial t}(s,t) = 8t - 2s^2\xi'(4t)\xi(4t) > 0$ for all $s, t > 0$. In view of (19.5) these two inequalities complete the proof (only positive $t = \widetilde{h}(u,v)$ occur since $\widetilde{H}$ is a positive function on $\widetilde{E}|(0;\delta]$). $\square$

We now concentrate on $\widetilde{h}_\delta = \widetilde{h}|\{\delta\} \times T^*_\epsilon S^2$.

Lemma 19.8. *$\widetilde{h}_\delta(u,v) = \gamma(|v|)$ for a function $\gamma \in C^\infty([0;\epsilon], \mathbb{R})$ which satisfies*

(1) $\gamma(t) = \frac{1}{4}\sqrt{\delta^2 + 4t^2}$ *for small* t,

(2) $\gamma(t) = \frac{1}{2}t$ *for* $t \geq \epsilon/2$, *and*

(3) $0 \leq \gamma'(t) < \frac{1}{2}$ *for* $t < \epsilon/2$.

Proof. The function

$$\beta(\delta, \cdot) : (0;\infty) \longrightarrow \mathbb{R}$$

satisfies $\beta(\delta, 0) = -\frac{1}{4}\delta^2$, $\partial\beta/\partial t > 0$ (see the proof of Lemma 19.7) and is unbounded. Therefore it has a smooth monotone inverse

$$\bar{\beta} : (-\tfrac{1}{4}\delta^2; \infty) \longrightarrow (0;\infty).$$

We emphasize that this is the inverse map and not the function $1/\beta$. By Lemma 19.6,

$$\widetilde{h}_\delta(u,v) = \bar\beta(\beta(\delta,\widetilde{h}_\delta(u,v))) = \bar\beta(|v|^2).$$

Therefore $\gamma(t) = \bar\beta(t^2)$ satisfies $\gamma(|v|) = \tilde{h}_\delta(u,v)$. $\beta(\delta,\cdot)$ has the following properties:

(1') $\beta(\delta,t) = \frac{1}{4}(16t^2 - \delta^2)$ for $t \le \epsilon/8$,

(2') $\beta(\delta,t) = 4t^2$ for $t \ge \epsilon/4$, and

(3') $\frac{\partial\beta}{\partial t}(\delta,t) > 8t$ for $t \in (\epsilon/8;\epsilon/4)$.

Property (3') holds because $\frac{\partial\beta}{\partial t}(\delta,t) = 8t - 2\delta^2\xi'(4t)\xi(4t)$ and $\xi'(t) < 0$, $\xi(t) > 0$ for $t \in (\epsilon/2;\epsilon)$. From the properties of β one obtains corresponding properties for its inverse:

(1") $\bar\beta(t) = \frac{1}{4}\sqrt{\delta^2 + 4t}$ for $0 \le t \le \epsilon^2/16 - \delta^2/4 = \beta(\delta,\epsilon/8)$,

(2") $\bar\beta(t) = \frac{1}{2}\sqrt{t}$ for $t \ge \epsilon^2/4 = \beta(\epsilon/4)$, and

(3") $0 < \bar\beta'(t) < (4\sqrt{t})^{-1}$ for all $t \in [0;\epsilon^2/4]$.

Again the last item requires some explanation: $\bar\beta'$ is positive because $\bar\beta$ is the inverse of the monotone function $\beta(\delta,\cdot)$. The second inequality follows from property (1") as long as $t \le \epsilon^2/16 - \delta^2/4$. In the other region ($t \in (\epsilon^2/16 - \delta^2/4;\epsilon^2/4)$) we have

$$\bar\beta'(t) = \left(\frac{\partial\beta}{\partial t}(\delta,\bar\beta(t))\right)^{-1} < \frac{1}{8\bar\beta(t)} \tag{19.6}$$

by (3'). Now $\beta(\delta,t) \le 4t^2$ and therefore $\bar\beta(t) \ge \frac{1}{2}\sqrt{t}$. Together with (19.6) this yields the desired inequality $\bar\beta'(t) < (4\sqrt{t})^{-1}$.

The properties of γ stated in the Lemma are immediate consequences of (1")–(3"). □

Proposition 19.9. *Let $P \in \mathrm{Aut}(\widetilde{E}_\delta,\widetilde{\Omega}|\widetilde{E}_\delta)$ be the symplectic monodromy of $\widetilde{E}$. The map $\tilde\tau = \tilde{f}_\delta \circ P \circ \tilde{f}_\delta^{-1} \in \mathrm{Aut}(T^*_\epsilon S^2,\eta)$ is given by*

$$\tilde\tau(x) = \begin{cases} \phi_{2\pi}^{r(\mu)}(x) & x \notin S^2, \\ A(x) & x \in S^2 \end{cases} \tag{19.7}$$

with a function $r \in C^\infty(\mathbb{R},\mathbb{R})$ such that $r(-t) = r(t) - t$ for all t, $r(t) = 0$ for $t \ge \epsilon/2$ and $0 < r'(t) \le \frac{1}{2}$ for $t \in [0;\epsilon/2)$. Therefore $\tilde\tau$ is a local model for generalized Dehn twists; moreover, $\mathrm{Fix}(\tilde\tau) = T^*_\epsilon S^2 \setminus T^*_{\epsilon/2}S^2$.

Proof. From Lemma 19.3(b) we obtain

$$\tilde{\tau} = (\tilde{f}_\delta \circ \sigma_{\widetilde{E}}(-1) \circ \tilde{f}_\delta^{-1}) \circ \phi_{-2\pi}^{\widetilde{h}_\delta}.$$

$\tilde{f}_\delta \circ \sigma_{\widetilde{E}}(-1) \circ \tilde{f}_\delta^{-1}$ is the restriction of the antipodal involution A to $T_\epsilon^* S^2$. From (1) in Lemma 19.8 it follows that any point in $S^2 \subset T_\epsilon^* S^2$ is a critical point of $\widetilde{h}_\delta$, and therefore $\tilde{\tau}|S^2 = A|S^2$.

Recall that $A|T^*S^2 \setminus S^2 = \sigma(-1)$, where σ is the circle action with moment map $\mu(u,v) = |v|$. Therefore

$$\tilde{\tau}|T_\epsilon^* S^2 \setminus S^2 = \phi_\pi^\mu \circ \phi_{-2\pi}^{\widetilde{h}_\delta}$$

Lemma 19.8 implies that $\widetilde{h}_\delta$ is invariant under σ; therefore

$$\tilde{\tau}|T_\epsilon^* S^2 \setminus S^2 = \phi_{2\pi}^{\frac{1}{2}\mu - \widetilde{h}_\delta}.$$

By Lemma 19.8, $\frac{1}{2}\mu - \widetilde{h}_\delta = r(\mu)$ where $r \in C^\infty([0;\epsilon), \mathbb{R})$ is given by $r(t) = \frac{1}{2}t - \gamma(t)$. Because γ is even for small t and $\gamma(t) = \frac{1}{2}t$ for $t \geq \epsilon/2$, r can be extended (in a unique way) to a smooth function on all of $\mathbb{R}$ such that $r(-t) = r(t) - t$ for all t and $r(t) = 0$ for $t \geq \epsilon/2$. Moreover, as a consequence of the corresponding property of γ, we have $0 < r'(t) \leq \frac{1}{2}$ for $t < \epsilon/2$. The fact that τ' is a model generalized Dehn twist follows from the definition of these models (Lemma 2.1). To see that $\mathrm{Fix}(\tau') = T_\epsilon^* S^2 \setminus T_{\epsilon/2}^* S^2$, write (19.7) as

$$\tilde{\tau}(x) = \begin{cases} \sigma(e^{2\pi i\, r'(\mu(x))})(x) & x \notin S^2, \\ A(x) & x \in S^2, \end{cases}$$

and use the fact that $0 < r'(t) < \frac{1}{2}$ for $t \in [0;\epsilon/2)$. □

Using $(\widetilde{\Omega}, \widetilde{E}, \widetilde{J}')$ as a local model we can now define (E_V, Ω_V, J_V').

Proposition 19.10. *Let (M,ω) be a compact symplectic four-manifold, V a Lagrangian two-sphere in M, and $i : T_\epsilon^* S^2 \longrightarrow M$ a symplectic embedding (for some $\epsilon > 0$) with $i(S^2) = V$. There is an ordinary almost holomorphic fibration (E_V, Ω_V, J_V') over some disc D_δ whose only critical value is $0 \in D_\delta$ and an isomorphism $f_V : ((E_V)_\delta, \Omega|(E_V)_\delta) \longrightarrow (M,\omega)$, with the following properties:*

(1) Let $P \in \mathrm{Aut}((E_V)_\delta, (\Omega_V)_\delta)$ be the monodromy of E_V. Then $\tau = f_V \circ P \circ f_V^{-1} \in \mathrm{Aut}(M,\omega)$ is the generalized Dehn twist along V constructed using the embedding i and a function r which satisfies the same conditions as in Proposition 19.9. In particular $\mathrm{Fix}(\tau) = M \setminus i(T_{\epsilon/2}^* S^2)$.

(2) (E_V, Ω_V, J_V') is nonnegative.

(3) For every point $x \in M \setminus i(T^*_{\epsilon/2}S^2)$ there is a horizontal section σ_x of (E_V, Ω_V, J'_V) such that $\sigma_x(\delta) = f_V^{-1}(x)$. If $x \in M \setminus \overline{i(T^*_{\epsilon/2}S^2)}$, $\mathrm{im}(\sigma_x)$ has a neighbourhood in E_V on which $\Omega_V|TE^h_V$ vanishes.

Proof. Recall that $(\widetilde{\Omega}, \widetilde{E}, \widetilde{J}')$ contains a trivial subbundle $T^*_\epsilon S^2 \setminus T^*_{\epsilon/2}S^2 \times D_\delta$. E_V is defined by gluing together $\widetilde{E}$ and $(M \setminus i(T^*_{\epsilon/2}S^2)) \times D_\delta$ along this trivial subbundle. It comes with a natural map $E_V \longrightarrow D_\delta$ whose only critical point is $(0,0) \in \widetilde{E}$. It inherits from $\widetilde{J}'$ a complex structure J'_V defined in a neighbourhood of this critical point. $\widetilde{\Omega}$ and the pullback of ω to $M \times D_\delta$ induce a closed two-form Ω_V on E_V.

It is not difficult to see that (E_V, Ω_V, J'_V) is an ordinary almost holomorphic fibration, and because both $(\widetilde{E}, \widetilde{\Omega}, \widetilde{J}')$ and $(M, \omega) \times D_\delta$ are nonnegative, so is (E_V, Ω_V, J'_V). The isomorphism f_V is defined by joining together $i \circ \tilde{f}_\delta : \widetilde{E}_\delta \longrightarrow i(T^*_\epsilon S^2) \subset M$ and the identity map $M \setminus T^*_{\epsilon/2}S^2 \to M$. The statement about the symplectic monodromy follows from Proposition 19.9, the fact that the symplectic monodromy of $(M, \omega) \times D^2$ is trivial, and the definition of generalized Dehn twists along V. The horizontal sections σ_x lie completely in the trivial part $(M \setminus T^*_{\epsilon/2}S^2) \times D_\delta$; they are given by $\sigma_x(z) = (x, z)$. □

20 Proof of Theorem 16.1

Throughout this section V is a Lagrangian two-sphere in (M, ω), $i : T^*_\epsilon S^2 \longrightarrow M$ is a symplectic embedding (for some $\epsilon > 0$) with $i(S^2) = V$, and $\tau_V \in \mathrm{Aut}(M, \omega)$ is a generalized Dehn twist along V formed using the embedding i and a function r as in Lemma 19.9.

Our first step is to modify the fibration (E_V, Ω_V, J'_V) constructed in the previous section slightly. Let $c : [0; 2\delta] \longrightarrow [0; \delta]$ be a smooth function such that $c(t) = t$ for $t < \delta/2$, $c(t) = \delta$ for $t \geq \delta$ and $c'(t) \geq 0$ everywhere. Consider the map $\kappa : D_{2\delta} \longrightarrow D_\delta$ given by

$$\kappa(z) = \begin{cases} z\frac{c(|z|)}{|z|} & z \neq 0, \\ 0 & z = 0. \end{cases}$$

Because κ is the identity near $0 \in D_\delta$, we can use it to pull back (E_V, Ω_V, J'_V) to a new ordinary almost holomorphic fibration over $D_{2\delta}$ (the point is that J'_V can be pulled back), which we call (E_1, Ω_1, J'_1). The pullback is again nonnegative; this follows from the fact that $\det(D\kappa_z) \geq 0$ for any $z \in D_{2\delta}$ together with Definition 17.2(i).

If we identify ∂D_δ with S^1 in the obvious way, $(E_V, \Omega_V)|\partial D_\delta$ is isomorphic to the mapping torus $(T_{\tau_V}, \Theta_{\tau_V})$; this is equivalent to the statement that

τ_V is the symplectic monodromy of (E_V, Ω_V). Since κ collapses $D_{2\delta} \setminus D_\delta$ radially to $\partial D\delta$, it follows that

$$(E_1, \Omega_1)|D_{2\delta} \setminus D_\delta \cong (\delta; 2\delta] \times (T_{\tau_V}, \Theta_{\tau_V}).$$

Proposition 19.10(3) implies that any horizontal section ν of $(T_{\tau_V}, \Theta_{\tau_V})$ can be extended to a horizontal section σ_1 of (E_1, Ω_1). Moreover, if ν corresponds to a point in $M \setminus \overline{i(T^*_{\epsilon/2}S^2)} \subset \mathrm{Fix}(\tau_V)$, there is a neighbourhood of $\mathrm{im}(\sigma_1)$ on which $\Omega_1|TE_1^h$ vanishes.

Lemma 20.1. *There is a $\phi \in \mathrm{Aut}(M, \omega)$ which is symplectically isotopic to the identity and has the following properties:*

(1) ϕ and $\tau_V \circ \phi$ have nondegenerate fixed points;

(2) $\mathrm{Fix}(\phi) \cap i(T^*_{\frac{3}{4}\epsilon}S^2)$ consists of two points. Both of them lie in V and their local fixed point index is $+1$;

(3) $\mathrm{Fix}(\tau_V \circ \phi) = \mathrm{Fix}(\phi) \cap (M \setminus V)$.

Proof. It is not difficult to see that there is a Morse function $h \in C^\infty(M, \mathbb{R})$ with $h(i(u, v)) = |v|$ for $\frac{1}{4}\epsilon \leq |v| \leq \frac{3}{4}\epsilon$ and such that $h|i(T^*_{\frac{3}{4}\epsilon}S^2)$ has two critical points, both of which have even Morse index and lie on V. We will prove that $\phi = \phi_t^H$, for sufficiently small $t > 0$, has the desired properties. It is clear that ϕ_t^H has only nondegenerate fixed points and that $\mathrm{Fix}(\phi_t^H) \cap i(T^*_{\frac{3}{4}\epsilon}S^2)$ is as required.

$$\begin{aligned}\mathrm{Fix}(\tau_V \circ \phi_t^H) \cap (M \setminus i(T^*_{\epsilon/2}S^2)) &= \mathrm{Fix}(\phi_t^H) \cap (M \setminus i(T^*_{\epsilon/2}S^2) \\ &= \mathrm{Fix}(\phi_t^H) \cap (M \setminus V)\end{aligned}$$

because ϕ_t^H preserves $(M \setminus i(T^*_{\epsilon/2}S^2))$ and $\tau_V|(M \setminus i(T^*_{\epsilon/2}S^2)) = \mathrm{id}$. For the same reasons any fixed point of $\tau_V \circ \phi_t^H$ which lies outside $i(T^*_{\epsilon/2}S^2)$ is nondegenerate. It remains to show that $\tau_V \circ \phi_t^H$ does not have any fixed points in $i(T^*_{\epsilon/2}S^2)$. Since $\mu(u, v) = |v|$ is the moment map of the familiar S^1-action σ on $T^*S^2 \setminus S^2$, we have

$$(i^{-1} \circ \phi_t^H \circ i)(u, v) = \sigma(e^{it})(u, v)$$

for $\frac{1}{4}\epsilon \leq |v| \leq \frac{3}{4}\epsilon$. Therefore

$$(i^{-1} \circ \tau_V \circ \phi_t^H \circ i)(u, v) = \sigma(e^{i(t + 2\pi\, r'(|v|))})(u, v).$$

by Proposition 19.10. Since $0 \leq r'(|v|) \leq \frac{1}{2}$ and σ is a free circle action, it follows that $\tau_V \circ \phi_t^H$ does not have any fixed points in $i(T^*_{\frac{1}{2}\epsilon}S^2 \setminus T^*_{\frac{1}{4}\epsilon}S^2)$ as long as $t < \pi$. The remaining region is simpler to deal with: since $\tau_V|i(T^*_{\frac{1}{3}\epsilon}S^2)$ is fixed point free, the same holds for $\tau_V \circ \phi_t^H|i(T^*_{\frac{1}{4}\epsilon}S^2)$ as long as t is small enough. □

From now on we assume that $\delta < \frac{1}{4}$ (this is possible since the fibration (E_1, Ω_1, J_1') can be transferred from a larger disc to a smaller one by a radial expansion). Let Σ be the surface obtained from $C = \mathbb{R} \times S^1$ by removing the closed δ-disc around the point $(0,0)$. We will think of Σ as the surface obtained from an infinite strip $[0;1] \times S^1$ by removing two half-discs and gluing part of the boundary together; that is,

$$\Sigma = \widehat{\Sigma}/\sim$$

where $\widehat{\Sigma} = \mathbb{R} \times [0;1] \setminus (B_\delta(0,0) \cup B_\delta(0,1))$ and $(s,1) \sim (s,0)$ for all $|s| > \delta$. We define a symplectic fibre bundle (E_2, Ω_2) over Σ in the following way: E_2 is obtained from the product bundle $\widehat{\Sigma} \times M$ by identifying $(s,1,x)$ with $(s,0,\phi(x))$ for $s < -\delta$ and with $(s,0,\tau_V(\phi(x)))$ for $s > \delta$. Ω_2 is induced from the pullback of ω to $\widehat{\Sigma} \times M$. If we identify the annulus $A = D_{2\delta}(0,0) \setminus D_\delta(0,0)$ with $(\delta; 2\delta] \times S^1$ in the obvious way, the restriction of (E_2, Ω_2) to A becomes isomorphic to $(\delta; 2\delta] \times (T_{\tau_V}, \Theta_{\tau_V})$. Therefore we can construct an ordinary almost holomorphic fibration (E, Ω, J') over C by gluing together (E_1, Ω_1, J_1') and (E_2, Ω_2) over A. By definition of (E_2, Ω_2), (E, Ω, J') has tubular ends modelled on (T_ϕ, Θ_ϕ) and $(T_{\tau_V \circ \phi}, \Theta_{\tau_V \circ \phi})$. Moreover, since (E_1, Ω_1, J_1') is nonnegative and (E_2, Ω_2) is flat (that is, $\Omega_2 | TE_2^h \equiv 0$), (E, Ω, J') is a nonnegative almost holomorphic fibration.

Let ν be the horizontal section of $(T_{\tau_V \circ \phi}, \Theta_{\tau_V \circ \phi})$ corresponding to a point $x \in \mathrm{Fix}(\tau_V \circ \phi)$. Consider the constant section $\hat{\sigma}_2 : \widehat{\Sigma} \longrightarrow \widehat{\Sigma} \times M$ given by $\hat{\sigma}_2(s,t) = (s,t,x)$. Lemma 20.1 says that $\mathrm{Fix}(\tau_V \circ \phi) \subset \mathrm{Fix}(\phi)$. Therefore

$\hat{\sigma}_2$ descends to a horizontal section σ_2 of (E_2, Ω_2). Because any horizontal section of $(T_{\tau_V}, \Theta_{\tau_V})$ can be extended to a horizontal section of (E_1, Ω_1, J_1'), there is a horizontal section σ_1 of (E_1, Ω_1, J_1') which agrees with σ_2 on the subset along which E_1 and E_2 are glued together. By piecing together σ_1 and σ_2 we obtain a horizontal section σ of (E, Ω, J') which has ν as its positive limit. Because $\mathrm{Fix}(\tau_V \circ \phi) \subset M \setminus i(T^*_{\frac{3}{4}\epsilon} S^2)$ by Lemma 20.1, $\mathrm{im}(\sigma_1) \subset E_1$ has a neighbourhood on which $\Omega_1 | TE_1^h$ vanishes. The corresponding fact for $\mathrm{im}(\sigma_2) \subset E_2$ is trivial because (E_2, Ω_2) is flat. It follows that σ has the same property.

Proof of Theorem 16.1. We have just shown that (E, Ω, J') satisfies the conditions of Theorem 17.7. Moreover, Lemma 20.1 implies that

$$\begin{aligned}\dim_\Lambda CF_1(T_\phi, \Theta_\phi) &= \dim_\Lambda CF_1(T_{\tau_V \circ \phi}, \Theta_{\tau_V \circ \phi}) \text{ and}\\ \dim_\Lambda CF_0(T_\phi, \Theta_\phi) &= \dim_\Lambda CF_0(T_{\tau_V \circ \phi}, \Theta_{\tau_V \circ \phi}) + 2.\end{aligned}$$

Therefore we can apply Corollary 17.11, which shows that $\Phi(E, \Omega, J')$ fits into an exact sequence

$$\begin{aligned}0 \longrightarrow HF_1(\phi) \longrightarrow HF_1(\tau_V \circ \phi) \longrightarrow &\\ \longrightarrow \Lambda^2 \longrightarrow HF_0(\phi) \longrightarrow HF_0(\tau_V \circ \phi) \longrightarrow 0.&\end{aligned}$$

Because ϕ is symplectically isotopic to the identity, $HF_*(\phi) \cong HF_*(\mathrm{id})$ and $HF_*(\tau_V \circ \phi) \cong HF_*(\tau_V)$. It remains to prove that $\Phi(E, \Omega, J')$ is a homomorphism of $QH_*(M, \omega)$-modules. Choose some $t \in S^1 \setminus 0$ and let $P : E_{(-1,t)} \longrightarrow E_{(1,t)}$ be the symplectic monodromy along the curve $[-1; 1] \times \{t\} \subset C$. Proposition 10.1 says that $\Phi(E, \Omega, J')$ is a homomorphism of $QH_*(M, \omega)$-modules if the composition

$$M \cong (T_\phi)_t \cong E_{(-1,t)} \xrightarrow{P} E_{(1,t)} \cong (T_{\tau_V \circ \phi})_t \cong M$$

induces the identity on homology. If we choose t such that $[-1; 1] \times \{t\} \subset \Sigma$, this is obvious because then $E | [-1; 1] \times \{t\} = E_2 | [-1; 1] \times \{t\}$ has a canonical trivialization which is compatible with the identifications of $E_{(\pm 1,t)}$ with M. □

21 Transversality for horizontal J

This section contains the proof of the technical Theorem 17.10. (E, Ω, J') denotes an almost holomorphic fibration which satisfies the assumptions of that Theorem, and we assume that almost complex structures $J^\pm \in \mathcal{J}_{\mathrm{reg}}(T^\pm, \Theta^\pm)$ have been chosen. As usual, π denotes the map $E \longrightarrow C$.

Theorem 17.10 amounts to the fact that the transversality arguments which lead to each of the four parts of Theorem 13.1 (namely, Propositions 13.9, 14.1, 14.2 and 15.1) can be carried out in the subspace

$$\mathcal{J}^h(E,\Omega,J';J^-,J^+) \subset \mathcal{J}(E,\Omega,J';J^-,J^+)$$

of almost complex structures which are horizontal. The restriction to the subspace $\mathcal{J}^h(E,\Omega,J';J^-,J^+)$ does not change the kind of general framework used in the proof. One chooses a $J_0 \in \mathcal{J}^h(E,\Omega,J';J^-,J^+)$, a large $R > 0$ such that J_0 agrees with J^- on $\pi^{-1}((-\infty;-R] \times S^1)$ and with J^+ on $\pi^{-1}([R;\infty) \times S^1)$, and a small closed neighbourhood U of $\mathrm{Crit}(\pi)$. The subspace of almost complex structures in $\mathcal{J}^h(E,\Omega,J';J^-,J^+)$ which agree with J_0 on

$$E_0 = \pi^{-1}((-\infty;-R] \times S^1 \cup [R;\infty) \times S^1) \cup U$$

will be denoted by $\mathcal{J}^h$.

$\mathcal{J}^h$ is a Fréchet manifold; to see this, consider the splitting $TE_x = TE_x^v \oplus TE_x^h$ at a point $x \notin \mathrm{Crit}(\pi)$. As an easy consequence of the definition, a horizontal almost complex structure has the form

$$J_x = \begin{pmatrix} J_x^{vv} & 0 \\ 0 & J_x^{hh} \end{pmatrix} \tag{21.1}$$

with respect to this splitting. Here J_x^{hh} is the unique horizontal lift of the complex structure on the base C; J_x^{vv} is an almost complex structure on TE_x^v which tames $\Omega|TE_x^v$. Conversely, every almost complex structure of the form (21.1) is horizontal at the point x. It follows that $\mathcal{J}^h$ is a Fréchet manifold and that its tangent space at any point J is the space of endomorphisms of TE^v which are $(J|TE^v)$-antilinear and vanish on E_0.

Recall that each part of the transversality Theorem 13.1 was eventually reduced to proving that a certain linear operator is onto. The tangent space $T_J\mathcal{J}$ was one factor in the domain of these operators. To prove Theorem 17.10 it is sufficient to check that these operators remain surjective if $T_J\mathcal{J}$ is replaced by the subspace $T_J\mathcal{J}^h$. We will carry out this check for the transversality theory of J-holomorphic sections (Proposition 13.9) which is by far the most important case. Indeed, in the other cases no problems arise (or, in the case of Proposition 14.2, they can be solved by a minor change in the argument) because the transversality theory for J-bubbles uses only the vertical component of J. Now, this component can be varied in the same way within $\mathcal{J}^h$ as within $\mathcal{J}$.

Let σ be a section in $\mathcal{M}(E,J;\nu^-,\nu^+)$ for some $J \in \mathcal{J}^h$ and $\nu^\pm \in \mathcal{H}(T^\pm,\Theta^\pm)$. What we need to prove is that the operator

$$D\bar\partial^{\mathrm{univ},h}(\sigma,J) : W^{1,p}(\sigma^*TE^v) \times T_J\mathcal{J}^h \longrightarrow L^p(\sigma^*TE^v),$$
$$D\bar\partial^{\mathrm{univ},h}(\sigma,J)(X,Y) = D\bar\partial_J(\sigma)X + Y(\partial\sigma/\partial t)^v$$

is onto. Here $(\partial\sigma/\partial t)^v$ denotes the vertical component of $\partial\sigma/\partial t$. We need to distinguish between horizontal and non-horizontal sections: indeed, if σ is horizontal, the second term in $D\bar{\partial}^{\mathrm{univ},h}(\sigma, J)$ vanishes and hence the argument which was used to prove Lemma 13.4 breaks down. We postpone discussing the horizontal sections and deal with the non-horizontal ones first.

Proposition 21.1. *If σ is not horizontal, the operator $D\bar{\partial}^{\mathrm{univ},h}(\sigma, J)$ is onto.*

Proof. The precise assumptions on R and U which are needed for this argument are: R should be so large that $E|[R-1;\infty) \times S^1 \cong [R-1;\infty) \times (T^+, \Theta^+)$, and $\pi(U) \subset [-R+1; R-1] \times S^1$.

As a first step, we prove that there is a point in $(R-1;R) \times S^1$ such that $\partial\sigma/\partial t$ is not horizontal at this point. Assume that the contrary is true: then

$$\sigma(s,t) = (s, \nu^+(t)) \quad \text{for } (s,t) \in (R-1;R) \times S^1,$$

where ν^+ is a horizontal section of (T^+, Θ^+). One of the assumptions on (E, Ω, J') in Theorem 17.10 is that there is a horizontal section σ^+ with positive limit ν^+. This section necessarily satisfies

$$\sigma^+(s,t) = (s, \nu^+(t)) \quad \text{for } s \geq R-1.$$

Hence $\sigma = \sigma^+$ on $(R-1;R) \times S^1$. The unique continuation theorem for J-holomorphic curves [20, Lemma 2.1.1] implies that $\sigma = \sigma^+$, contradicting our assumption that σ is not horizontal.

Let $(s_0, t_0) \in (R-1;R) \times S^1$ be a point such that $(\partial\sigma/\partial t)(s_0, t_0)$ is not horizontal. It is clear that the same holds for all (s,t) in some neighbourhood D of (s_0, t_0). This implies that for every $W \in C^\infty(\sigma^* TE^v)$ which is supported in D there is a $Y \in T_J \mathcal{J}^h$ such that

$$D\bar{\partial}^{\mathrm{univ},j}(\sigma, J)(0, Y) = Y(\partial\sigma/\partial t)^v = W.$$

The rest of the proof is as in Lemma 13.4. □

It remains to deal with the horizontal sections. From what we have said above, it is clear that $D\bar{\partial}^{\mathrm{univ},h}(\sigma, J)$ is surjective at a horizontal σ iff $D\bar{\partial}_J(\sigma)$ is onto, that is, iff σ is regular. This is ensured by the following result:

Proposition 21.2. *A horizontal section of (E, Ω, J') is regular with respect to any $J \in \mathcal{J}^h(E, \Omega, J'; J^-, J^+)$.*

To prove this we translate an idea of Braam and Donaldson [4] from the theory of instantons to that of holomorphic curves. Let σ be a horizontal section of E. As we saw in the proof of Theorem 17.7, σ must necessarily be one of the sections whose existence is part of our assumptions on (E, Ω, J').

In particular, $\Omega|TE^h$ vanishes in a neighbourhood on $\mathrm{im}(\sigma)$. Hence $\mathrm{ind}(\sigma) = 0$ by Lemma 17.6(d). Because of the vanishing of the index, it is sufficient to show that the kernel of

$$D\bar{\partial}_J(\sigma) : W^{1,p}(\sigma^* TE^v) \longrightarrow L^p(\sigma^* TE^v) \tag{21.2}$$

is trivial. The operator $D\bar{\partial}_J(\sigma)$ can be written in a particularly simple form using the canonical connection ∇^σ on $\sigma^* TE^v$:

Lemma 21.3. *Let σ be a horizontal section of (E, Ω, J'). Then*

$$D\bar{\partial}_J(\sigma) = \frac{\nabla^\sigma}{\partial s} + J(\sigma)\frac{\nabla^\sigma}{\partial t}$$

for all $J \in \mathcal{J}^h(E, \Omega, J'; J^-, J^+)$,

Proof. In section 11 we have obtained the formula

$$D\bar{\partial}_J(\sigma)X = [\widetilde{S}, \widetilde{X}] + J[J\widetilde{S}, \widetilde{X}],$$

where $\widetilde{S} \in C^\infty(TE)$ and $\widetilde{X} \in C^\infty(TE^v)$ are extensions of $\partial\sigma/\partial s$ and X. In the present case, since σ is horizontal, we can take $\widetilde{S}$ to be the horizontal lift of the unit vector field in s-direction on C. Then by definition

$$[\widetilde{S}, \widetilde{X}] = \frac{\nabla^\sigma}{\partial s} X.$$

Since J is horizontal, $J\widetilde{S}$ is also horizontal. In fact it is the horizontal lift of the unit vector field in t-direction on C and therefore

$$[J\widetilde{S}, \widetilde{X}] = \frac{\nabla^\sigma}{\partial t} X. \quad \square$$

Lemma 17.6(c) shows that the connection ∇^σ is flat. For such connections we have the following result:

Lemma 21.4. *Let (F, ω_F) be a symplectic vector bundle over C with a flat symplectic connection ∇_F. Let J_F be a complex structure on F (that is, (F, J_F) is a complex vector bundle) which tames ω_F, and*

$$\bar{\partial}_F = \frac{\nabla_F}{\partial s} + J_F \frac{\nabla_F}{\partial t}$$

the Dolbeault operator on (F, J) determined by ∇_F. Then every section $S \in C^\infty(F)$ such that

$$\bar{\partial}_F(S) = 0, \quad \int_C |\nabla_F S|^2 < \infty \ \textit{and} \ \int_C |S|\,|\nabla_F S| < \infty$$

($|\cdot|^2$ is the metric on F obtained from J_F and ω_F) is parallel: it satisfies $\nabla_F S = 0$.

Proof. There is a unique two-form Ω_F on the total space of F (considered as a manifold) with the following property: if $i : U \times \mathbb{R}^{2n} \longrightarrow F|U$ is a symplectic trivialization of (F, ω) over some $U \subset \Sigma$ such that $i^*(\nabla_F)$ is trivial, then $i^*\Omega_F$ is the pullback of the standard symplectic structure on $\mathbb{R}^{2n}$ to $U \times \mathbb{R}^{2n}$. This characterization shows that Ω_F is unique, and it also allows to construct it by patching together local charts.

Let S be a smooth section of F with $\bar{\partial}_F S = 0$. S is a smooth map from Σ to F, and we can use it to pull back Ω_F to a two-form $S^*\Omega_F$ on Σ. By looking at a local chart, one sees that

$$S^*\Omega_F = \tfrac{1}{2}|\nabla_F S|^2 ds \wedge dt; \tag{21.3}$$

our assumption implies that $\int_C S^*\Omega_F$ converges.

Let us consider the multiples $S_r = r\,S$. Clearly $S_r^*\Omega_F = r^2(S^*\Omega_F)$. On the other hand (again by considering local coordinates) we find that

$$\frac{\partial}{\partial r} S_r^*\Omega_F = d\theta_r$$

with $\theta_r = 2r\,S_r^*(i_{\partial S_r/\partial r}\Omega_F)$. θ_r satisfies $|\theta_r| \leq r|S|\,|\nabla_F S|$. Therefore it is integrable, and by Stokes' theorem, $(\partial/\partial r)S_r^*\Omega_F = 0$. Of course, one must be careful when applying Stokes' theorem because of the non-compactness of C, but a little reflection shows that our decay conditions are sufficient to take care of that. $S_r^*\Omega_F = r^2 S^*\Omega_F$ and $(\partial/\partial r)S_r^*\Omega_F = 0$ imply that $S^*\Omega_F = 0$, and because of (21.3) it follows that $\nabla_F S = 0$. □

Any $S \in W^{1,p}(\sigma^*TE^v)$ with $D\bar{\partial}_J(\sigma)S = 0$ satisfies the conditions of Lemma 21.4; therefore it must be ∇^σ-parallel. On the other hand, a nontrivial parallel section can never lie in $W^{1,p}$ (it is constant in $\mathbb{R}$-direction outside a compact subset and therefore does not decay); this proves that (21.2) is injective and completes the proof of Proposition 21.2.

Remark 21.5. Lemma 21.4 and its proof generalize to Riemann surfaces with tubular ends; the statement is

> *Let (Σ, j) be a Riemann surface with tubular ends and (F, ω_F) a symplectic vector bundle over it with a flat symplectic connection ∇_F. Let J_F be a complex structure on F which tames ω_F, and*
>
> $$\bar{\partial}_F = \tfrac{1}{2}(\nabla + J_F \circ \nabla \circ j)$$
>
> *the Dolbeault operator on (F, J_F) determined by ∇_F. Then every section $S \in C^\infty(F)$ with*
>
> $$\bar{\partial}_F(S) = 0, \quad \int_\Sigma |\nabla_F S|^2 < \infty \text{ and } \int_\Sigma |S|\,|\nabla_F S| < \infty$$
>
> *is parallel.*

It seems unlikely that this is new (but I have not found a reference). The result has interesting implications even in the closed case: assume that Σ is a closed surface of genus $g > 0$. The space of ∇-parallel sections has dimension $\leq 2n$, where $n = \mathrm{rank}_{\mathbb{C}}(F)$. Therefore we obtain the inequality

$$\mathrm{ind}_{\mathbb{R}} \bar{\partial}_F \leq 2n \tag{21.4}$$

for the (real) index of $\bar{\partial}_F$. Using the Riemann-Roch theorem we conclude that $c_1(F, \omega_F) \leq ng$. If equality holds in (21.4), the symplectic vector bundle (F, ω_F) must be trivial and hence $c_1(F, \omega_F) = 0$. Using this we can improve the estimate to

$$c_1(F, \omega_F) < ng.$$

The case $n = 1$ is an old result of Milnor [21] (Milnor's proof is completely different).

Bibliography

[1] V. I. Arnol'd, *Some remarks on symplectic monodromy of Milnor fibrations*, The Floer Memorial Volume (H. Hofer, C. Taubes, A. Weinstein, and E. Zehnder, eds.), Progress in Mathematics, vol. 133, Birkhäuser, 1995, pp. 99–104.

[2] M. F. Atiyah, *On analytic surfaces with double points*, Proc. Royal Soc. Ser. A **247** (1958), 237–244.

[3] M. Audin and J. Lafontaine (eds.), *Holomorphic curves in symplectic geometry*, Progress in Mathematics, vol. 117, Birkhäuser, 1994.

[4] P. Braam and S. K. Donaldson, *Floer's work on instanton homology, knots and surgery*, The Floer Memorial Volume (H. Hofer, C. Taubes, A. Weinstein, and E. Zehnder, eds.), Progress in Mathematics, vol. 133, Birkhäuser, 1995, pp. 195–256.

[5] B. Crauder and R. Miranda, *Quantum cohomology of rational surfaces*, The moduli space of curves (R. Dijkgraaf, C. Faber, and G. van der Geer, eds.), Progress in Mathematics, vol. 129, Birkhäuser, 1995, pp. 33–80.

[6] P. DiFrancesco and C. Itzykson, *Quantum intersection rings*, The moduli space of curves (R. Dijkgraaf, C. Faber, and G. van der Geer, eds.), Progress in Mathematics, vol. 129, Birkhäuser, 1995, pp. 81–148.

[7] S. Dostoglou and D. Salamon, *Instanton homology and symplectic fixed points*, Symplectic Geometry (D. Salamon, ed.), LMS Lecture Note Series, vol. 192, Cambridge Univ. Press, 1993, pp. 57–94.

[8] ———, *Self dual instantons and holomorphic curves*, Annals of Math. **139** (1994), 581–640.

[9] A. Floer, *Symplectic fixed points and holomorphic spheres*, Commun. Math. Phys. **120** (1989), 575–611.

[10] ———, *Witten's complex and infinite dimensional Morse theory*, J. Differential Geom. **30** (1989), 207–221.

[11] ———, *Instanton homology and Dehn surgery*, The Floer Memorial Volume (H. Hofer, C. Taubes, A. Weinstein, and E. Zehnder, eds.), Progress in Mathematics, vol. 133, Birkhäuser, 1995, pp. 77–98.

[12] A. Floer, H. Hofer, and D. Salamon, *Transversality in elliptic Morse theory for the symplectic action*, Duke Math. J. **80** (1995), 251–292.

[13] M. Gromov, *Pseudoholomorphic curves in symplectic manifolds*, Invent. Math. **82** (1985), 307–347.

[14] H. Hofer and D. Salamon, *Floer homology and Novikov rings*, The Floer memorial volume (H. Hofer, C. Taubes, A. Weinstein, and E. Zehnder, eds.), Progress in Mathematics, vol. 133, Birkhäuser, 1995, pp. 483–524.

[15] Le Hong Van and K. Ono, *Cup-length estimate for symplectic fixed points*, Contact and symplectic geometry (C. B. Thomas, ed.), Cambridge Univ. Press, 1996, pp. 268–295.

[16] A. Liu, *Some new applications of general wall crossing formula, Gompf's conjecture and its applications*, Math. Res. Lett. **3** (1996), 569–585.

[17] D. McDuff, *Lectures on Gromov invariants for symplectic 4-manifolds*, Preprint.

[18] ———, *Elliptic methods in symplectic geometry*, Bull. Amer. Math. Soc. **23** (1990), 311–358.

[19] ———, *Immersed spheres in symplectic 4-manifolds*, Ann. Inst. Fourier **42** (1992), 369–392.

[20] D. McDuff and D. Salamon, *J-holomorphic curves and quantum cohomology*, University Lecture Notes Series, vol. 6, Amer. Math. Soc., 1994.

[21] J. Milnor, *On the existence of a connection with curvature zero*, Comment. Math. Helv. **32** (1958), 215–223.

[22] S. Piunikhin, D. Salamon, and M. Schwarz, *Symplectic Floer-Donaldson theory and quantum cohomology*, Contact and symplectic geometry (C. B. Thomas, ed.), Cambridge Univ. Press, 1996, pp. 171–200.

[23] Y. Ruan and G. Tian, *A mathematical theory of Quantum Cohomology*, J. Differential Geom. **42** (1995), 259–367.

[24] D. Salamon, *Morse theory, the Conley index, and Floer homology*, Bull. London Math. Soc. **22** (1990), 113–140.

[25] D. Salamon and E. Zehnder, *Morse theory for periodic solutions of Hamiltonian systems and the Maslov index*, Comm. Pure Appl. Math. **45** (1992), 1303–1360.

[26] M. Schwarz, *A quantum cup-length estimate for symplectic fixed points*, Preprint.

[27] ———, *Cohomology operations from S^1-cobordisms in Floer homology*, Ph.D. thesis, ETH Zürich, 1995.

[28] D. Sullivan, *Infinitesimal computations in topology*, Publ. Math. IHES **47** (1977), 269–332.

[29] C. Taubes, *The Seiberg-Witten and the Gromov invariants*, Math. Res. Lett. **2** (1995), 221–238.

[30] ———, $SW \Rightarrow Gr$, Journal of the Amer. Math. Soc. **9** (1996), 845–918.

[31] A. Weil, *Introduction à l'étude des variétés kähleriennes*, Hermann, 1958.

[32] A. Weinstein, *Symplectic manifolds and their Lagrangian submanifolds*, Adv. Math. **6** (1971), 329–346.