Lectures on Pseudo-Holomorphic Curves and the Symplectic Isotopy Problem

Bernd Siebert[1] and Gang Tian[2]

[1] Mathematisches Institut, Albert-Ludwigs-Universität Freiburg, Eckerstr. 1, 79104 Freiburg, Germany
bernd.siebert@math.uni-freiburg.de
[2] Department of Mathematics, MIT, Cambridge, MA 02139-4307, USA
tian@math.mit.edu

Summary. The purpose of these notes is a more self-contained presentation of the results of the authors in Siebert and Tian (Ann Math 161:955–1016, 2005). Some applications are also given.

1 Introduction

This text is an expanded version of the lectures delivered by the authors at the CIME summer school "Symplectic 4-manifolds and algebraic surfaces," Cetraro (Italy), September 2–10, 2003. The aim of these lectures were mostly to introduce graduate students to pseudo-holomorphic techniques for the study of isotopy of symplectic submanifolds in dimension four. We tried to keep the style of the lectures by emphasizing the basic reasons for the correctness of a result rather than by providing full details.

Essentially none of the content claims any originality, but most of the results are scattered in the literature in sometimes hard-to-read locations. For example, we give a hands-on proof of the smooth parametrization of the space of holomorphic cycles on a complex surface under some positivity assumption. This is usually derived by the big machinery of deformation theory together with Banach-analytic methods. For an uninitiated person it is hard not only to follow the formal arguments needed to deduce the result from places in the literature, but also, and maybe more importantly, to understand *why* it is true. While our treatment here has the disadvantage to consider only a particular situation that does not even quite suffice for the proof of the Main Theorem (Theorem 9.1) we hope that it is useful for enhancing the understanding of this result outside the community of hardcore complex analysts and algebraic geometers.

One lecture was devoted to the beautiful theorem of Micallef and White on the holomorphic nature of pseudo-holomorphic curve singularities. The original paper is quite well written and this might be the reason that a proof of this theorem has not appeared anywhere else until very recently, in the excellent monograph [McSa2]. It devotes an appendix of 40 pages length to a careful derivation of this theorem and of most of the necessary analytical tools. Following the general principle of these lecture notes our purpose here is not to give a complete and necessarily technical proof, but to condense the original proof to the essentials. We tried to achieve this goal by specializing to the paradigmical case of tacnodal singularities with a "half-integrable" complex structure.

Another section treats the compactness theorem for pseudo-holomorphic curves. A special feature of the presented proof is that it works for sequences of almost complex structures only converging in the $\mathscr{C}^0$-topology. This is essential in the application to the symplectic isotopy problem.

We also give a self-contained presentation of Shevchishin's study of moduli spaces of equisingular pseudo-holomorphic maps and of second variations of pseudo-holomorphic curves. Here we provide streamlined proofs of the two results from [Sh] by only computing the second variation in the directions actually needed for our purposes.

The last section discusses the proof of the main theorem, which is also the main result from [SiTi3]. The logic of this proof is a bit difficult, involving several reduction steps and two inductions, and we are not sure if the current presentation really helps in understanding what is going on. Maybe somebody else has to take this up again and add new ideas into streamlining this piece.

Finally there is one section on the application to symplectic Lefschetz fibrations. This makes the link to the other lectures of the summer school, notably to those by Auroux and Smith.

2 Pseudo-Holomorphic Curves

2.1 Almost Complex and Symplectic Geometry

An *almost complex structure* on a manifold M is a bundle endomorphism $J : T_M \to T_M$ with square $-\operatorname{id}_{T_M}$. In other words, J makes T_M into a complex vector bundle and we have the canonical decomposition

$$T_M \otimes_{\mathbb{R}} \mathbb{C} = T_M^{1,0} \oplus T_M^{0,1} = T_M \oplus \overline{T_M}$$

into real and imaginary parts. The second equality is an isomorphism of complex vector bundles and $\overline{T_M}$ is just another copy of T_M with complex structure $-J$. For switching between complex and real notation it is important to write down the latter identifications explicitly:

$$T_M \longrightarrow T_M^{1,0}, \quad X \longmapsto \tfrac{1}{2}\big(X - iJX\big),$$

and similarly with $X+iJX$ for $T_M^{0,1}$. Standard examples are complex manifolds with $J(\partial_{x_\mu}) = \partial_{y_\mu}$, $J(\partial_{y_\mu}) = -\partial_{x_\mu}$ for holomorphic coordinates $z_\mu = x_\mu + iy_\mu$.

Then the above isomorphism sends $\partial_{x_\mu}, \partial_{y_\mu} \in T_M$ to $\partial_{z_\mu}, i\partial_{z_\mu} \in T_M^{1,0}$ and to $\partial_{\bar z_\mu}, i\partial_{\bar z_\mu} \in T_M^{0,1}$ respectively. Such *integrable* almost complex structures are characterized by the vanishing of the Nijenhuis tensor, a $(2,1)$-tensor depending only on J. In dimension 2 an almost complex structure is nothing but a conformal structure and hence it is integrable by classical theory. In higher dimensions there are manifolds having almost complex structures but no integrable ones. For example, any symplectic manifold (M,ω) possesses an almost complex structure, as we will see instantly, but there are symplectic manifolds not even homotopy equivalent to a complex manifold, see e.g. [OzSt].

The link between symplectic and almost complex geometry is by the notion of *tameness*. An almost complex structure J is *tamed* by a symplectic form ω if $\omega(X, JY) > 0$ for any $X, Y \in T_M \setminus \{0\}$. The space $\mathscr{J}^\omega$ of ω-tamed almost complex structures is contractible. In fact, one first proves this for the space of *compatible* almost complex structures, which have the additional property $\omega(JX, JY) = \omega(X, Y)$ for all X, Y. These are in one-to-one correspondence with Riemannian metrics g via $g(X,Y) = \omega(X, JY)$, and hence form a contractible space. In particular, a compatible almost complex structure J_0 in (M,ω) exists. Then the generalized Cayley transform

$$J \longmapsto (J + J_0)^{-1} \circ (J - J_0)$$

produces a diffeomorphism of $\mathscr{J}^\omega$ with the space of J_0-antilinear endomorphisms A of T_M with $\|A\| < 1$ (this is the mapping norm for $g_0 = \omega(\,.\,, J_0\,.\,)$).

A differentiable map $\varphi : N \to M$ between almost complex manifolds is *pseudo-holomorphic* if $D\varphi$ is complex linear as map between complex vector bundles. If φ is an embedding this leads to the notion of pseudo-holomorphic submanifold $\varphi(N) \subset M$. If the complex structures are integrable then pseudo-holomorphicity specializes to holomorphicity. However, there are many more cases:

Proposition 2.1 *For any symplectic submanifold $Z \subset (M,\omega)$ the space of $J \in \mathscr{J}^\omega(M)$ making Z into a pseudo-holomorphic submanifold is non-empty and contractible.*

The proof uses the same arguments as for the contractibility of $\mathscr{J}^\omega$ outlined above. Another case of interest for us is the following, which can be proved by direct computation.

Proposition 2.2 *[SiTi3, Proposition 1.2] Let (M,ω) be a closed symplectic 4-manifold and $p : M \to B$ a smooth fiber bundle with all fibers symplectic. Then for any symplectic form ω_B on B and any almost complex structure J on M making the fibers of p pseudo-holomorphic, $\omega_k := \omega + k\, p^*(\omega_B)$ tames J for $k \gg 0$.*

The Cauchy–Riemann equation is over-determined in dimensions greater than two and hence the study of pseudo-holomorphic maps $\varphi : N \to M$

promises to be most interesting for $\dim N = 2$. Then N is a (not necessarily connected) Riemann surface, that we write Σ with almost complex structure j understood. The image of φ is called *pseudo-holomorphic curve*, or *J-holomorphic curve* if one wants to explicitly refer to an almost complex structure J on M. A pseudo-holomorphic curve is *irreducible* if Σ is connected, otherwise *reducible*, and its *irreducible components* are the images of the connected components of Σ. If φ does not factor non-trivially over a holomorphic map to another Riemann surface we call φ *reduced* or *non-multiply covered*, otherwise *non-reduced* or *multiply covered*.

2.2 Basic Properties of Pseudo-Holomorphic Curves

Pseudo-holomorphic curves have a lot in common with holomorphic curves:

(1) *Regularity.* If $\varphi : \Sigma \to (M, J)$ is of Sobolev class $W^{1,p}$, $p > 2$ (one weak derivative in L^p) and satisfies the Cauchy–Riemann equation $\frac{1}{2}(D\varphi + J \circ D\varphi \circ j) = 0$ weakly, then φ is smooth ($\mathscr{C}^\infty$; we assume J smooth). Note that by the Sobolev embedding theorem $W^{1,p}(\Sigma, M) \subset C^0(\Sigma, N)$, so it suffices to work in charts.

(2) *Critical points.* The set of critical points $\mathrm{crit}(\varphi) \subset \Sigma$ of a pseudo-holomorphic map $\varphi : \Sigma \to M$ is discrete.

(3) *Intersections and identity theorem.* Two different irreducible pseudo-holomorphic curves intersect discretely and, if $\dim M = 4$, with positive, finite intersection indices.

(4) *Local holomorphicity.* Let $C \subset (M, J)$ be a pseudo-holomorphic curve with finitely many irreducible components and $P \in C$. Then there exists a neighborhood $U \subset M$ of P and a C^1-diffeomorphism $\Phi : U \to \Phi(U) \subset \mathbb{C}^n$ such that $\Phi(C)$ is a holomorphic curve near $\Phi(P)$.

This is the content of the theorem of Micallef and White that we discuss in detail in Sect. 4. Note that this implies (2) and (3).

(5) *Removable singularities.* Let $\Delta^* \subset \mathbb{C}$ denote the pointed unit disk and $\varphi : \Delta^* \to (M, J)$ a pseudo-holomorphic map. Assume that φ has *bounded energy*, that is $\int_{\Delta^*} |D\varphi|^2 < \infty$ for any complete Riemannian metric on M. Then φ extends to a pseudo-holomorphic map $\Delta \to M$.

If ω tames the almost complex structure the energy can be bounded by the symplectic area: $\int_{\Delta^*} |D\varphi|^2 < c \cdot \int_\Sigma \varphi^*\omega$. Note that $\int_\Sigma \varphi^*\omega$ is a topological entity provided Σ is closed.

(6) *Local existence.* For any $X \in T_M$ of sufficiently small length there exists a pseudo-holomorphic map $\varphi : \Delta \to M$ with $D\varphi_{|0}(\partial_t) = X$. Here t is the standard holomorphic coordinate on the unit disk.

The construction is by application of the implicit function theorem to appropriate perturbations of the exponential map. Therefore it also works in families. In particular, any almost complex manifold can locally be fibered into pseudo-holomorphic disks. In dimension 4 this implies the local existence of complex (-valued) coordinates z, w such that $z = \text{const}$ is a pseudo-holomorphic disk with w restricting to a holomorphic coordinate. There exist then complex functions a, b with

$$T_M^{0,1} = \mathbb{C} \cdot (\partial_{\bar{z}} + a\partial_z + b\partial_w) + \mathbb{C} \cdot \partial_{\bar{w}}.$$

Conversely, any choices of a, b lead to an almost complex structure with z, w having the same properties. This provides a convenient way to work with almost complex structures respecting a given fibration of M by pseudo-holomorphic curves.

2.3 Moduli Spaces

The real use of pseudo-holomorphic curves in symplectic geometry comes from looking at whole spaces of them rather than single elements. There are various methods to set up the analysis to deal with such spaces. Here we follow the treatment of [Sh], to which we refer for details. Let $\mathbb{T}_g$ be the Teichmüller space of complex structures on a closed oriented surface Σ of genus g. The advantage of working with $\mathbb{T}_g$ rather than with the Riemann moduli space is that $\mathbb{T}_g$ parametrizes an actual family of complex structures on a fixed closed surface Σ. Let G be the holomorphic automorphism group of the identity component of any $j \in \mathbb{T}_g$, that is $G = \mathrm{PGL}(2, \mathbb{C})$ for $g = 0$, $G = U(1) \times U(1)$ for $g = 1$ and $G = 0$ for $g \geq 2$. Then $\mathbb{T}_g$ is an open ball in $\mathbb{C}^{3g-3+\dim_{\mathbb{C}} G}$, and it parametrizes a family of G-invariant complex structures. Let $\mathscr{J}$ be the Banach manifold of almost complex structures on M, of class C^l for some integer $l > 2$ fixed once and for all. The particular choice is insignificant. The *total moduli space* $\mathscr{M}$ of pseudo-holomorphic maps $\Sigma \to M$ is then a quotient of a subset of the Banach manifold

$$\mathscr{B} := \mathbb{T}_g \times W^{1,p}(\Sigma, M) \times \mathscr{J}.$$

The local shape of this Banach manifold is exhibited by its tangent spaces

$$T_{\mathscr{B},(j,\varphi,J)} = H^1(T_\Sigma) \times W^{1,p}(\varphi^* T_M) \times \mathscr{C}^l(\mathrm{End}\, T_M).$$

Here $H^1(T_\Sigma)$ is the cohomology with values in T_Σ, viewed as holomorphic line bundle over Σ. In more classical notation $H^1(T_\Sigma)$ may be replaced by the space of holomorphic quadratic differentials. In any case, the tangent space to $\mathbb{T}_g$ is also a subspace of $\mathscr{C}^\infty(\mathrm{End}(T_\Sigma))$ via variations of j and this is how we are going to represent its elements. To describe $\mathscr{M}$ consider the Banach bundle $\mathscr{E}$ over $\mathscr{B}$ with fibers

$$\mathscr{E}_{(j,\varphi,J)} = L^p(\Sigma, \varphi^*(T_M, J) \otimes_{\mathbb{C}} \Lambda^{0,1}),$$

where $\Lambda^{0,1}$ is our shorthand notation for $(T_\Sigma^{0,1})^*$ and where we wrote (T_M, J) to emphasize that T_M is viewed as a complex vector bundle via J. Consider the section $s : \mathscr{B} \to \mathscr{E}$ defined by the condition of complex linearity of $D\varphi$:

$$s(j, \varphi, J) = D\varphi + J \circ D\varphi \circ j.$$

Thus $s(j,\varphi,J)=0$ iff $\varphi:(\Sigma,j)\to(M,J)$ is pseudo-holomorphic. We call the operator defined by the right-hand side the *nonlinear $\bar\partial$-operator*. If $M=\mathbb{C}^n$ with its standard complex structure this is just twice the usual $\bar\partial$-operator applied to the components of φ. Define $\hat{\mathscr{M}}$ as the zero locus of s minus those (j,φ,J) defining a multiply covered pseudo-holomorphic map. In other words, we consider only generically injective φ. This restriction is crucial for transversality to work, see Lemma 2.4. There is an obvious action of G on $\mathscr{B}$ by composing φ with biholomorphisms of (Σ,j). The moduli space of our interest is the quotient $\mathscr{M}:=\hat{\mathscr{M}}/G$ of the induced action on $\hat{\mathscr{M}}\subset\mathscr{B}$.

Proposition 2.3 *$\hat{\mathscr{M}}\subset\mathscr{B}$ is a submanifold and G acts properly and freely.*

Sketch of proof. A torsion-free connection ∇ on M induces connections on φ^*T_M (also denoted ∇) and on $\mathscr{E}$ (denoted $\nabla^{\mathscr{E}}$). For $(j',v,J')\in T_{(j,\varphi,J)}$ and $w\in T_\Sigma$ a straightforward computation gives, in real notation

$$\big(\nabla^{\mathscr{E}}_{(j',v,J')}s\big)w=\nabla_w v+J\circ\nabla_{j(w)}v+\nabla_v J\circ D_{j(w)}\varphi+J'\circ D_{j(w)}\varphi+J\circ D_{j'(w)}\varphi.$$

Replacing w by $j(w)$ changes signs, so $\nabla^{\mathscr{E}}_{(j',v,J')}s$ lies in $L^p(\varphi^*T_M\otimes_{\mathbb{C}}\Lambda^{0,1})=\mathscr{E}_{(j,\varphi,J)}$ as it should. The last two terms treat the variations of J and j respectively. The first three terms compute the derivative of s for fixed almost complex structures. They combine to a first order differential operator on φ^*T_M denoted by

$$D_{\varphi,J}v=\nabla v+J\circ\nabla_{j(.)}v+\nabla_v J\circ D_{j(.)}\varphi. \tag{2.1}$$

This operator is not generally J-linear, but has the form

$$D_{\varphi,J}=2\bar\partial_{\varphi,J}+R,$$

with a J-linear differential operator $\bar\partial_{\varphi,J}$ of type $(0,1)$ and with the J-antilinear part R of order 0. (With our choices $R=N_J(\,.\,,D\varphi\circ j)$ for N_J the Nijenhuis tensor of J.) Then $\bar\partial_{\varphi,J}$ *defines* a holomorphic structure on φ^*T_M, and this ultimately is the source of the holomorphic nature of pseudo-holomorphic maps. It is then standard to deduce that $D_{\varphi,J}$ is Fredholm as map from $W^{1,p}(\varphi^*T_M)$ to $L^p(\varphi^*T_M\otimes\Lambda^{0,1})$. To finish the proof apply the implicit function theorem taking into account the following Lemma 2.4, whose proof is an exercise.

The statements on the action of G are elementary to verify.

Lemma 2.4 *If $\varphi:\Sigma\to M$ is injective over an open set in Σ, then* $\operatorname{coker} D_{\varphi,J}$ *can be spanned by terms of the form $J'\circ D\varphi\circ j$ for $J'\in T_J\mathscr{J}=\mathscr{C}^l(\operatorname{End}(T_M))$.*

The proof of the proposition together with the Riemann–Roch count (index theorem) for the holomorphic vector bundle $(\varphi^*T_M,\bar\partial_{\varphi,J})$ give the following.

Corollary 2.5 *The projection* $\pi : \hat{\mathscr{M}} \to \mathbb{T}_g \times \mathscr{J}$ *is Fredholm of index*

$$\begin{aligned} \mathrm{ind}(D_{\varphi,J}) = \mathrm{ind}(\bar\partial_{\varphi,J}) &= 2\,(\deg \varphi^* T_M + \dim_{\mathbb{C}} M \cdot (1-g)) \\ &= 2\,(c_1(M)\cdot \varphi_*[\Sigma] + \dim_{\mathbb{C}} M \cdot (1-g)) \end{aligned}$$

A few words of caution are in order. First, for $g = 0, 1$ the action of G on $\mathscr{B}$ is *not differentiable*, although it acts by differentiable transformations. The reason is that differentiating along a family of biholomorphisms costs one derivative and hence leads out of any space of finite differentiability. A related issue is that the differentiable structure on $\mathscr{B}$ depends on the choice of map from $\mathbb{T}_g$ to the space of complex structures on Σ. Because of elliptic regularity all of these issues disappear after restriction to $\hat{\mathscr{M}}$, so we may safely take the quotient by G there. One should also be aware that there is still the mapping class group of Σ acting on $\mathscr{M}$. Only the quotient is the set of isomorphism classes of pseudo-holomorphic curves on M. However, this quotient does not support a universal family of curves anymore, at least not in the naïve sense. As our interest in $\mathscr{M}$ is for Sard-type results it is technically simpler to work with $\mathscr{M}$ rather than with this discrete quotient.

Moreover, for simplicity we henceforth essentially ignore the action of G. This merely means that we drop some correction terms in the dimension counts for $g = 0, 1$.

Remark 2.6 (1) The derivative of a section of a vector bundle $\mathscr{E}$ over a manifold $\mathscr{B}$ does not depend on the choice of a connection *after restriction to the zero locus.* In fact, if $v \in \mathscr{E}_p$ lies on the zero locus, then $T_{\mathscr{E},v} = T_{M,p} \oplus \mathscr{E}_p$ canonically, and this decomposition is the same as induced by any connection. Thus Formula (2.1) has intrinsic meaning along $\hat{\mathscr{M}}$. In particular, Ds defines a section of $\mathrm{Hom}(T_{\mathscr{B}}, \mathscr{E})|_{\hat{\mathscr{M}}}$ that we need later on.

(2) The projection $\pi : \hat{\mathscr{M}} \to \mathbb{T}_g \times \mathscr{J}$ needs not be proper – sequences of pseudo-holomorphic maps can have reducible or lower genus limits. Here are two typical examples in the integrable situation.

(a) *A family of plane quadrics degenerating to two lines.*
Let $\Sigma = \mathbb{CP}^1$, $M = \mathbb{CP}^2$ and $\varphi_\varepsilon([t,u]) = [\varepsilon t^2, tu, \varepsilon u^2]$. Then $\mathrm{im}(\varphi_\varepsilon)$ is the plane algebraic curve $V(xz - \varepsilon^2 y^2) = \{[x,y,z] \in \mathbb{CP}^2 \mid xz - \varepsilon^2 y^2 = 0\}$. For $\varepsilon \to 0$ the image Hausdorff converges to the union of two lines $xz = 0$. Hence φ_ε can not converge in any sense to a pseudo-holomorphic map $\Sigma \to M$.

(b) *A drop of genus by the occurrence of cusps.*
For $\varepsilon \neq 0$ the cubic curve $V(x^2 z - y^3 - \varepsilon z^3) \subset \mathbb{CP}^2$ is a smooth holomorphic curve of genus 1, hence the image of a holomorphic map φ_ε from $\Sigma = S^1 \times S^1$. For $\varepsilon \to 0$ the image of φ_ε converges to the cuspidal cubic $V(x^2 z - y^3)$, which is the bijective image of

$$\mathbb{CP}^1 \to \mathbb{CP}^2, \quad [t,u] \to [t^3, t^2 u, u^3].$$

The *Gromov compactness theorem* explains the nature of this non-compactness precisely. In its modern form it states the compactness of a Hausdorff enlargement of the space of isomorphism classes of pseudo-holomorphic curves over any compact subset of the space of almost complex structures $\mathscr{J}$. The elements of the enlargement are so-called *stable maps*, which are maps with domains nodal Riemann surfaces. For a detailed discussion see Sect. 6.

2.4 Applications

I. *Ruled surfaces.* In complex geometry a ruled surface is a holomorphic $\mathbb{CP}^1$-bundle. They are all Kähler. A ruled surface is rational (birational to $\mathbb{CP}^2$) iff the base of the fibration is $\mathbb{CP}^1$. A symplectic analogue are S^2-bundles with a symplectic structure making all fibers symplectic. The fibers are then symplectic spheres with self-intersection 0. Conversely, a result of McDuff says that a symplectic manifold with a symplectic sphere $C \subset M$ with $C \cdot C \geq 0$ is either $\mathbb{CP}^2$ with the standard structure (and C is a line or a conic) or a symplectic ruled surface (and C is a fiber or, in the rational case, a positive section), up to symplectic blowing-up [MD]. The proof employs pseudo-holomorphic techniques similar to what follows.

Now let $p : (M^4, I) \to (S^2, i)$ be a holomorphic rational ruled surface. In the notation we indicated the complex structures by I and i. Let ω be a Kähler form on M and $\mathscr{J}^\omega$ the space of ω-tamed almost complex structures. The following result was used in [SiTi3] to reduce the isotopy problem to a fibered situation.

Proposition 2.7 *For any $J \in \mathscr{J}^\omega$, M arises as total space of an S^2-bundle $p' : M \to S^2$ with all fibers J-holomorphic. Moreover, p' is homotopic to p through a homotopy of S^2-bundles.*

Sketch of proof. By connectedness of $\mathscr{J}^\omega$ there exists a path $(J_t)_{t\in[0,1]}$ connecting $I = J_0$ with $J = J_1$. By a standard Sard-type argument

$$\mathscr{M}_{(J_t)} := [0,1] \times_{\mathscr{J}} \mathscr{M} = \{(t, j, \varphi, J) \in [0,1] \times \mathscr{M} \mid J = J_t\}$$

is a manifold for an appropriate ("general") choice of (J_t). Let $\mathscr{M}_{F,J_t} \subset \mathscr{M}_{(J_t)}$ be the subset of J_t-holomorphic curves homologous to a fiber $F \subset M$ of p. The exact sequence of complex vector bundles over the domain $\Sigma = S^2$ of such a J_t-holomorphic curve (see Sect. 3.2)

$$0 \longrightarrow T_\Sigma \longrightarrow T_M|_\Sigma \longrightarrow N_{\Sigma|M} \longrightarrow 0$$

gives $c_1(M) \cdot [F] = c_1(T_\Sigma) \cdot [\Sigma] + F \cdot F = 2$. Then the dimension formula from Corollary 2.5 shows

$$\dim_{\mathbb{C}} \mathscr{M}_{F,(J_t)} = c_1(M) \cdot [F] + \dim_{\mathbb{C}} M \cdot (1-g) - \dim G = 2 + 2 - 3 = 1.$$

Moreover, $\mathscr{M}_{F,J_t}$ is compact by the Gromov compactness theorem since $[F]$ is a primitive class in $\{A \in H_2(M, \mathbb{Z}) \mid \int_{[A]} \omega > 0\}$. In fact, by primitivity any pseudo-holomorphic curve C representing $[F]$ has to be irreducible. Moreover, the genus formula (Proposition 3.2 below) implies that any irreducible pseudo-holomorphic curve representing $[F]$ has to be an embedded sphere. We will see in Proposition 3.4 that then the deformation theory of any $C \in \mathscr{M}_{F,J_t}$ is unobstructed.

Next, the positivity of intersection indices of pseudo-holomorphic curves implies that any two curves in $\mathscr{M}_{F,J_t}$ are either disjoint or equal. Together with unobstructedness we find that through any point $P \in M$ passes exactly one J_t-holomorphic curve homologous to F. Define

$$p_t : M \to \mathscr{M}_{F,J_t}, \quad P \longmapsto C, \quad C \text{ the curve passing through } P.$$

Since $\mathscr{M}_{F,J_0} \simeq S^2$ via $C \leftrightarrow p^{-1}(x)$ we may identify $p_0 = p$. A computation on the map of tangent spaces shows that p_t is a submersion for any t. Finally, for homological reasons $\mathscr{M}_{F,J_t} \simeq S^2$ for any t. The proof is finished by setting $p' = p_1$.

II. *Isotopy of symplectic surfaces.* The main topic of these lectures is the isotopy classification of symplectic surfaces. We are now ready to explain the relevance of pseudo-holomorphic techniques for this question. Let (M^4, I, ω) be a Kähler surface. We wish to ask the following question.

If $B \subset M$ is a symplectic surface then is B isotopic to a holomorphic curve?

By *isotopy* we mean connected by a path inside the space of smooth symplectic submanifolds. In cases of interest the space of smooth holomorphic curves representing $[B] \in H_2(M, \mathbb{Z})$ is connected. Hence a positive answer to our question shows uniqueness of symplectic submanifolds homologous to B up to isotopy.

The use of pseudo-holomorphic techniques is straightforward. By the discussion in Sect. 2.2 there exists a tamed almost complex structure J making B a pseudo-holomorphic curve. As in 2.4,I choose a generic path $(J_t)_{t\in[0,1]}$ in $\mathscr{J}^\omega$ connecting J with the integrable complex structure I. Now try to deform B as pseudo-holomorphic curve together with J. In other words, we want to find a family $(B_t)_{t\in[0,1]}$ of submanifolds with B_t pseudo-holomorphic for J_t and with $B_0 = B$.

There are two obstructions to carrying this through. Let $\mathscr{M}_{B,(J_t)}$ be the moduli space of pseudo-holomorphic submanifolds $C \subset M$ homologous to B and pseudo-holomorphic for some J_t. The first problem arises from the fact that the projection $\mathscr{M}_{B,(J_t)} \to [0,1]$ may have critical points. Thus if $(B_t)_{t\in[0,t_0]}$ is a deformation of B with the requested properties over the interval $[0, t_0]$ with t_0 a critical point, it might occur that B_{t_0} does not deform to a J_t-holomorphic curve for any $t > t_0$. We will see in Sect. 3 that this phenomenon does indeed not occur under certain positivity conditions on M. The second reason is non-properness of the projection $\mathscr{M}_{B,(J_t)} \to [0,1]$. It might happen that a family (B_t) exists on $[0, t_0)$, but does not extend to t_0. In view of the

Gromov compactness theorem a different way to say this is to view $\mathscr{M}_{B,(J_t)}$ as an open subset of a larger moduli space $\tilde{\mathscr{M}}_{B,(J_t)}$ of pseudo-holomorphic cycles. Then the question is if the closed subset of singular cycles does locally disconnect $\tilde{\mathscr{M}}_{B,(J_t)}$ or not. Thus for this second question one has to study deformations of singular pseudo-holomorphic cycles.

III. *Pseudo-holomorphic spheres with prescribed singularities.* Another variant of the above technique allows the construction of pseudo-holomorphic spheres with prescribed singularities. The following is from [SiTi3], Proposition 7.1.

Proposition 2.8 *Let $p : (M, I) \to \mathbb{CP}^1$ be a rational ruled surface and*

$$\varphi : \Delta \longrightarrow M$$

an injective holomorphic map, $\Delta \subset \mathbb{C}$ the unit disk. Let J be an almost complex structure on M making p pseudo-holomorphic and agreeing with I in a neighborhood of $\varphi(0)$.

Then for any $k > 0$ there exists a J-holomorphic sphere

$$\psi_k : \mathbb{CP}^1 \longrightarrow M$$

approximating φ to kth order at 0:

$$d_M(\varphi(\tau), \psi_k(\tau)) = o(|\tau|^k).$$

Here d_M is the distance function for some Riemannian metric on M. This result says that any plane holomorphic curve singularity arises as the singularity of a J-holomorphic sphere, for J with the stated properties. The proof relies heavily on the fact that M is a rational ruled surface. Note that the existence of a J-holomorphic sphere not homologous to a fiber excludes ruled surfaces over a base of higher genus.

Sketch of proof. It is not hard to see that J can be connected to I by a path of almost complex structures with the same properties as J. Therefore the idea is again to start with a holomorphic solution to the problem and then to deform the almost complex structure. Excluding the trivial case $p \circ \varphi = \mathrm{const}$ write

$$\varphi(\tau) = (\tau^m, h(\tau))$$

in holomorphic coordinates on $M \setminus (F \cup H) \simeq \mathbb{C}^2$, F a fiber and H a positive holomorphic section of p $(H \cdot H \geq 0)$. Then h is a holomorphic function. Now consider the space of pseudo-holomorphic maps $\mathbb{CP}^1 \to M$ of the form

$$\tau \longmapsto \big(\tau^m, h(\tau) + \mathrm{o}(|\tau|^l)\big).$$

For appropriate l the moduli space of such maps has expected dimension 0. Then for a generic path $(J_t)_{t\in[0,1]}$ of almost complex structures the union of such moduli spaces over this path is a differentiable one-dimensional manifold $q : \mathscr{M}_{\varphi,(J_t)} \to [0,1]$ without critical points over $t = 0, 1$. By a straightforward dimension estimate the corresponding moduli spaces of reducible pseudo-holomorphic curves occurring in the Gromov compactness theorem are empty. Hence the projection $q : \mathscr{M}_{\varphi,(J_t)} \to [0,1]$ is proper. Thus $\mathscr{M}_{\varphi,(J_t)}$ is a compact one-dimensional manifold with boundary and all boundary points lie over $\{0,1\} \subset [0,1]$. The closed components of $\mathscr{M}_{\varphi,(J_t)}$ do not contribute to the moduli spaces for $t = 0, 1$. The other components have two ends each, and they either map to the same boundary point of $[0,1]$ or to different ones. In any case the parity of the cardinality of $q^{-1}(0)$ and of $q^{-1}(1)$ are the same, as illustrated in the following Fig. 2.1. Finally, an explicit computation shows that in the integrable situation the moduli space has exactly one element. Therefore $q^{-1}(1)$ can not be empty either. An element of this moduli space provides the desired J-holomorphic approximation of φ.

2.5 Pseudo-Analytic Inequalities

In this section we lay the foundations for the study of critical points of pseudo-holomorphic maps. As this is a local question we take as domain the unit disk, $\varphi : \Delta \to (M, J)$. The main point of this study is that any singularity bears a certain kind of holomorphicity in itself, and the amount of holomorphicity indeed increases with the complexity of the singularity. The reason for this to happen comes from the following series of results on differential inequalities for the $\bar{\partial}$-operator. For simplicity of proof we formulate these only for functions with values in $\mathbb{C}$ rather than $\mathbb{C}^n$ as needed, but comment on how to generalize to $n > 1$.

Lemma 2.9 *Let $f \in W^{1,2}(\Delta)$ fulfill $|\partial_{\bar{z}} f| \le \phi \cdot |f|$ almost everywhere for $\phi \in L^p(\Delta)$, $p > 2$. Then either $f = 0$ or there exist a uniquely determined integer μ and $g \in W^{1,p}(\Delta)$, $g(0) \neq 0$ with*

$$f(z) = z^\mu \cdot g(z) \quad \textit{almost everywhere.}$$

Proof. A standard elliptic bootstrapping argument shows $f \in W^{1,p}(\Delta)$, see for example [IvSh1], Lemma 3.1.1,(i). This step requires $p > 2$. Next

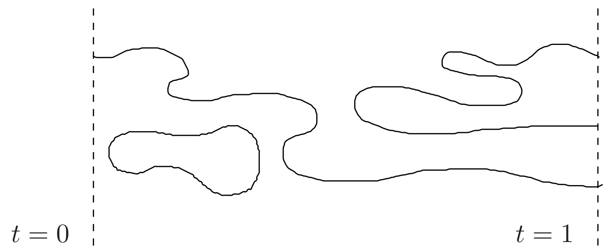

Fig. 2.1. A one-dimensional cobordism

comes a trick attributed to Carleman to reduce to the holomorphic situation: By hypothesis $\left|\frac{\partial_{\bar z} f}{f}\right| \le \phi$. We will recall in Proposition 3.1 below that $\partial_{\bar z} : W^{1,p}(\Delta) \to L^p(\Delta)$ is surjective. Hence there exists $\psi \in W^{1,p}(\Delta)$ solving $\partial_{\bar z}\psi = \frac{\partial_{\bar z} f}{f}$. Then

$$\partial_{\bar z}(e^{-\psi} f) = e^{-\psi}(-\partial_{\bar z}\psi) f + e^{-\psi}\partial_{\bar z} f = 0$$

shows that $e^{-\psi} f$ is a holomorphic function (*Carleman similarity principle*). Now complex function theory tells us that $e^{-\psi} f = z^\mu \cdot h$ for $h \in \mathcal{O}(\Delta) \cap W^{1,p}(\Delta)$, $h(0) \neq 0$. Putting $g = e^{\psi} \cdot h$ gives the required representation of f. □

Remark 2.10 (1) As for an intrinsic interpretation of μ note that it is the intersection multiplicity of the graph of h with the graph of the zero function inside $\Delta \times \mathbb{C}$ at $(0,0)$. Multiplication by e^{ψ} induces a homeomorphism of $\Delta \times \mathbb{C}$ and transforms the graph of h into the graph of f. Hence μ is a topologically defined entity depending only on f.

(2) The Carleman trick replaces the use of a general removable singularities theorem for solutions of differential inequalities due to Harvey and Polking that was employed in [IvSh1], Lemma 3.1.1. Unlike the Carleman trick this method generalizes to maps $f : \Delta \to \mathbb{C}^n$ with $n > 1$. Another possibility that works also for $n > 1$ is to use the Hartman-Wintner theorem on the polynomial behavior of solutions of certain partial differential equations in two variables, see e.g. [McSa2]. A third approach appeared in the printed version [IvSh2] of [IvSh1]; here the authors noticed that one can deduce a Carleman similarity principle also for maps $f : \Delta \to \mathbb{C}^n$ by viewing f as a holomorphic section of $\Delta \times \mathbb{C}^n$, viewed as holomorphic vector bundle with non-standard $\bar\partial$-operator. This is arguably the easiest method to deduce the result for all n.

A similar looking lemma of quite different flavor deduces holomorphicity up to some order from a polynomial estimate on $|\partial_{\bar z} f|$. Again we took this from [IvSh1], but the proof given there makes unnecessary use of Lemma 2.9.

Lemma 2.11 *Let $f \in L^2(\Delta)$ fulfill $|\partial_{\bar z} f| \le \phi \cdot |z|^\nu$ almost everywhere for some $\phi \in L^p(\Delta)$, $p > 2$ and $\nu \in \mathbb{N}$. Then either $f = 0$ or there exists $P \in \mathbb{C}[z]$, $\deg P \le \nu$ and $g \in W^{1,p}(\Delta)$, $g(0) = 0$ with*

$$f(z) = P(z) + z^\nu \cdot g(z) \quad \textit{almost everywhere.}$$

Proof. By induction over ν, the case $\nu = 0$ being trivial. Assume the case $\nu - 1$ is true. Elliptic regularity gives $f \in W^{1,p}(\Delta)$. By the Sobolev embedding theorem f is Hölder continuous of exponent $\alpha = 1 - \frac{2}{p} \in (0,1)$. Hence $f_1 = \dfrac{f - f(0)}{z}$ is L^2 and

$$|\partial_{\bar z} f_1| = \left|\frac{\partial_{\bar z} f}{z}\right| \le \phi \cdot |z|^{\nu - 1}.$$

Therefore induction applies to f_1 and we see $f_1 = P_1 + z^{\nu-1} \cdot g$ with g of the required form. Plugging in the definition of f_1 gives $f = (f(0) + zP_1) + z^\nu \cdot g$, so $P = f(0) + zP_1$ is the correct definition of P. $\square$

Remark 2.12 The lemma generalizes in a straightforward manner to maps $f : \Delta \to \mathbb{C}^n$. In this situation the line $\mathbb{C} \cdot P(0)$ has a geometric interpretation as complex tangent line of $\mathrm{im}(f)$ in $f(0)$, which by definition is the limit of lines $\mathbb{C} \cdot (f(z) - f(0)) \in \mathbb{CP}^{n-1}$ for $z \neq 0$, $z \to 0$.

Combining the two lemmas gives the following useful result, which again generalizes to maps $f : \Delta \to \mathbb{C}^n$.

Proposition 2.13 *Let $f \in L^2(\Delta)$ fulfill $|\partial_{\bar z} f| \le \phi |z|^\nu |f|$ almost everywhere for $\phi \in L^p(\Delta)$, $p > 2$ and $\nu \in \mathbb{N}$. Then either $f = 0$ or there exist uniquely determined $\mu, \nu \in \mathbb{N}$ and $P \in \mathbb{C}[z]$, $\deg P \le \nu$, $P(0) \neq 0$, $g \in W^{1,p}(\Delta)$, $g(0) = 0$ with*

$$f(z) = z^\mu \big(P(z) + z^\nu \cdot g(z) \big) \quad \textit{almost everywhere.}$$

Proof. Lemma 2.9 gives $f = z^\mu g$. Now by hypothesis g also fulfills the stated estimate:

$$|\partial_{\bar z} g| \;=\; \left| \frac{\partial_{\bar z} f}{z^\mu} \right| \;\le\; \phi |z|^\nu \left| \frac{f}{z^\mu} \right| \;=\; \phi |z^\nu| \cdot |g|.$$

Thus replacing f by g reduces to the case $f(0) \neq 0$ and $\mu = 0$. The result then follows by Lemma 2.11 applied to f with ϕ replaced by $\phi \cdot |f|$. $\square$

3 Unobstructedness I: Smooth and Nodal Curves

3.1 Preliminaries on the $\bar\partial$-Equation

The crucial tool to study the $\bar\partial$-equation analytically is the inhomogeneous Cauchy integral formula. It says

$$f = Hf + T(\partial_{\bar z} f) \tag{3.1}$$

for all $f \in \mathscr{C}^1(\overline{\Delta})$ with integral operators

$$Hf(z) = \frac{1}{2\pi i} \int_{\partial\Delta} \frac{f(w)}{w - z}\, dw\,, \quad Tg(z) = \frac{1}{2\pi i} \int_{\Delta} \frac{g(w)}{w - z}\, dw \wedge d\bar w\,.$$

(All functions in this section are $\mathbb{C}$-valued.) The first operator H maps continuous functions defined on $S^1 = \partial\Delta$ to holomorphic functions on Δ. Continuity of Hf along the boundary is not generally true if f is just continuous. To understand this note that any $f \in \mathscr{C}^0(S^1)$ can be written as a Fourier series $\sum_{n\in\mathbb{Z}} a_n z^n$ and then $Hf = \sum_{n\in\mathbb{N}} a_n z^n$ is the projection to the space of functions spanned by non-negative Fourier modes. This function needs not be continuous.

The integrand of the second integral operator T looks discontinuous, but in fact it is not as one sees in polar coordinates with center $w = z$. For differentiability properties one computes $\partial_{\bar z} T = \mathrm{id}$ from (3.1), while $\partial_z T = S$ with S the singular integral operator

$$Sg(z) = \frac{1}{2\pi i} \lim_{\varepsilon \to 0} \int_{\Delta \setminus B_\varepsilon(z)} \frac{g(w)}{(w-z)^2}\, dw \wedge d\bar w\,.$$

The Calderon–Zygmund theorem says that S is a continuous map from $L^p(\Delta)$ to itself for $1 < p < \infty$. Recall also that the Sobolev space $W^{1,p}(\Delta)$ consists of L^p-functions with weak partial derivatives in L^p too. For $p > 2$ it holds $1 - \frac{2}{p} > 0$, so the Sobolev embedding theorem implies that any $f \in W^{1,p}(\Delta)$ has a continuous representative. Moreover, the map $W^{1,p}(\Delta) \to C^0(\overline{\Delta})$ thus defined is continuous. Thus (3.1) holds for $f \in W^{1,p}(\Delta)$ for any $p > 2$. Summarizing the discussion, the Cauchy integral formula induces the following remarkable direct sum decomposition of $W^{1,p}(\Delta)$.

Proposition 3.1 *Let $2 < p < \infty$. Then $(H, \bar\partial) : W^{1,p}(\Delta) \longrightarrow \big(\mathcal{O}(\Delta) \cap W^{1,p}(\Delta)\big) \times L^p(\Delta)$ is an isomorphism. In particular, for any $g \in L^p(\Delta)$ there exists $f \in W^{1,p}(\Delta)$ with $\partial_{\bar z} f = g$.*

Thus any $f \in W^{1,p}(\Delta)$ can be written in the form $h + T(\partial_{\bar z} f)$ with h holomorphic in Δ and continuously extending to $\overline{\Delta}$ and $T(\partial_{\bar z} f)|_{\partial\Delta}$ gathering all negative Fourier coefficients of $f|_{\overline{\Delta}}$.

3.2 The Normal $\bar\partial$-Operator

We have already described pseudo-holomorphic maps by a non-linear PDE. One trivial variation of a pseudo-holomorphic map is by reparametrization. It is sometimes useful to get rid of this part, especially if one is interested in pseudo-holomorphic curves rather than pseudo-holomorphic maps. This is achieved by the normal $\bar\partial$-operator that we now introduce.

Recall that for a pseudo-holomorphic map $\varphi : \Sigma \to M$ the operator $\bar\partial_{\varphi,J}$ from (2.2) in Sect. 2.3 defines a natural holomorphic structure on $\varphi^* T_M$ compatible with the holomorphic structure on T_Σ. In fact, a straightforward computation shows

$$D\varphi \circ \bar\partial_{T_\Sigma} = \bar\partial_{\varphi,J} \circ D\varphi.$$

If φ is an immersion we thus obtain a short exact sequence of holomorphic vector bundles over Σ

$$0 \longrightarrow T_\Sigma \longrightarrow \varphi^*(T_M) \longrightarrow N \longrightarrow 0.$$

This sequence *defines* the normal bundle N along φ. If φ has critical points it is still possible to define a normal bundle as follows. For a complex vector bundle V denote by $\mathcal{O}(V)$ the sheaf of holomorphic sections of V. (See Sect. 5.4 for a short reminder of sheaf theory.) While at critical points $D\varphi : T_\Sigma \to \varphi^* T_M$ is not injective the map of sheaves $\mathcal{O}(T_\Sigma) \to \mathcal{O}(\varphi^* T_M)$ still is. As an example consider $\varphi(t) = (t^2, t^3)$ as map from Δ to $\mathbb{C}^2$ with standard complex structure. Then

$$D\varphi(\partial_t) = 2t\,\partial_z + 3t^2\,\partial_w, \tag{3.2}$$

and as germ of holomorphic function the right-hand side is non-zero. Thus in any case we obtain a short exact sequence of sheaves of $\mathcal{O}_\Sigma$-modules

$$0 \longrightarrow \mathcal{O}(T_\Sigma) \xrightarrow{D\varphi} \mathcal{O}(\varphi^* T_M) \longrightarrow \mathcal{N} \longrightarrow 0. \tag{3.3}$$

From the definition, $\mathcal{N}$ is just some coherent sheaf on Σ. But on a Riemann surface any coherent sheaf splits uniquely into a skyscraper sheaf (discrete support) and the sheaf of sections of a holomorphic vector bundle. Thus we may write

$$\mathcal{N} = \mathcal{N}^{\mathrm{tor}} \oplus \mathcal{O}(N)$$

for some holomorphic vector bundle N. We call $\mathcal{N}$ the *normal sheaf* along φ and N the *normal bundle*. The skyscraper sheaf $\mathcal{N}^{\mathrm{tor}}$ is the subsheaf of $\mathcal{N}$ generated by sections that are annihilated by multiplication by some non-zero holomorphic function ("torsion sections"). In our example $\varphi(t) = (t^2, t^3)$ the section $v = 2\,\partial_z + 3t\,\partial_w$ of $\varphi^* T_M$ is contained in the image of $\mathcal{O}(T_\Sigma)$ for $t \neq 0$, but not at $t = 0$, while $tv = D\varphi(\partial_t)$. In fact, $\mathcal{N}^{\mathrm{tor}}$ is isomorphic to a copy of $\mathbb{C}$ over $t = 0$ generated by the germ of v at 0. Then N is the holomorphic line bundle generated by any $a(t)\partial_z + b(t)\partial_w$ with $b(0) \neq 0$.

As a simple but very powerful application of the normal sequence (3.3) we record the genus formula for pseudo-holomorphic curves in dimension four.

Proposition 3.2 *Let (M, J) be an almost complex 4-manifold and $C \subset M$ an irreducible pseudo-holomorphic curve. Then*

$$2g(C) - 2 \leq c_1(M) \cdot C + C \cdot C,$$

with equality if and only if C is smooth.

Proof. Let $\varphi : \Sigma \to M$ be the pseudo-holomorphic map with image C. Since $\deg T_\Sigma = 2g(C) - 2$ the normal sequence (3.3) shows

$$2g(C) - 2 \;=\; \deg \varphi^* T_M + \deg N + \lg \mathcal{N}^{\mathrm{tor}} \;=\; c_1(M) \cdot C + \deg N + \lg \mathcal{N}^{\mathrm{tor}}.$$

Here $\lg \mathcal{N}^{\mathrm{tor}}$ is the sum of the $\mathbb{C}$-vector space dimensions of the stalks of $\mathcal{N}$. This term vanishes iff $\mathcal{N}^{\mathrm{tor}} = 0$, that is, iff φ is an immersion. The degree of N equals $C \cdot C$ if C is smooth and drops by the self-intersection number of φ in the immersed case. By the PDE description of the space of pseudo-holomorphic maps from a unit disk to M, it is not hard to show that locally in Σ any pseudo-holomorphic map $\Sigma \to M$ can be perturbed to a pseudo-holomorphic immersion. At the expense of changing J away from the singularities of C slightly this statement globalizes. This process does not change any of $g(C)$, $c_1(M) \cdot C$ and $C \cdot C$. Hence the result for general C follows from the immersed case. □

To get rid of $\mathcal{N}^{\mathrm{tor}}$ it is convenient to go over to meromorphic sections of T_Σ with poles of order at most $\operatorname{ord}_P D\varphi$ in a critical point P of $D\varphi$. In fact, the sheaf of such meromorphic sections is the sheaf of holomorphic sections of a line bundle that we conveniently denote $T_\Sigma[A]$, where A is the divisor $\sum_{P \in \operatorname{crit}(\varphi)} (\operatorname{ord}_P D\varphi) \cdot P$. Then $T_\Sigma[A] = \operatorname{kern}(\varphi^* T_M \to N)$ and hence we obtain the short exact sequence of holomorphic vector bundles

$$0 \longrightarrow T_\Sigma[A] \xrightarrow{D\varphi} \varphi^* T_M \longrightarrow N \longrightarrow 0.$$

Thus by (3.2) together with $R \circ D\varphi = 0$ the operator $D_{\varphi,J} = \bar\partial_{\varphi,J} + R : W^{1,p}(\varphi^* T_M) \to L^p(\varphi^* T_M \otimes \Lambda^{0,1})$ fits into the following commutative diagram with exact rows.

$$\begin{array}{ccccccccc}
0 & \to & W^{1,p}(T_\Sigma[A]) & \to & W^{1,p}(\varphi^* T_M) & \to & W^{1,p}(N) & \to & 0 \\
 & & \bar\partial_{T_\Sigma}\big\downarrow & & \big\downarrow D_{\varphi,J} & & \big\downarrow D^N_{\varphi,J} & & \\
0 & \to & L^p(T_\Sigma[A] \otimes \Lambda^{0,1}) & \to & L^p(\varphi^* T_M \otimes \Lambda^{0,1}) & \to & L^p(N \otimes \Lambda^{0,1}) & \to & 0
\end{array}$$

This defines the *normal* $\bar\partial_J$*-operator* $D^N_{\varphi,J}$. As with $D_{\varphi,J}$ we have the decomposition $D^N_{\varphi,J} = \bar\partial_N + R_N$ into complex linear and a zero order complex anti-linear part. By the snake lemma the diagram readily induces the long exact sequence

$$\begin{aligned}
0 \ &\to H^0(T_\Sigma[A]) \to \operatorname{kern} D_{\varphi,J} \to \operatorname{kern} D^N_{\varphi,J} \\
&\to H^1(T_\Sigma[A]) \to \operatorname{coker} D_{\varphi,J} \to \operatorname{coker} D^N_{\varphi,J} \to 0
\end{aligned}$$

The cohomology groups on the left are Dolbeault cohomology groups for the holomorphic vector bundle $T_\Sigma[A]$ or sheaf cohomology groups of the corresponding coherent sheaves. Forgetting the twist by A then gives the following exact sequence.

$$\begin{aligned}
0 \ &\to H^0(T_\Sigma) \to \operatorname{kern} D_{\varphi,J} \to \operatorname{kern} D^N_{\varphi,J} \oplus H^0(\mathcal{N}^{\mathrm{tor}}) \\
&\to H^1(T_\Sigma) \to \operatorname{coker} D_{\varphi,J} \to \operatorname{coker} D^N_{\varphi,J} \to 0
\end{aligned} \tag{3.4}$$

The terms in this sequence have a geometric interpretation. Each column is associated to a deformation problem. The left-most column deals with deformations of Riemann surfaces: $H^0(T_\Sigma)$ is the space of holomorphic vector fields on Σ. It is trivial except in genera 0 and 1, where it gives infinitesimal holomorphic reparametrizations of φ. As already mentioned in Sect. 2.3, the space $H^1(T_\Sigma)$ is isomorphic to the space of holomorphic quadratic differentials via Serre-duality and hence describes the tangent space to the Riemann or Teichmüller space of complex structures on Σ. Every element of this space is the tangent vector of an actual one-parameter family of complex structures on Σ – this deformation problem is unobstructed.

The middle column covers the deformation problem of φ as pseudo-holomorphic curve *with almost complex structures both on M and on Σ fixed.* In fact, $D_{\varphi,J}$ is the linearization of the Fredholm map describing the space of pseudo-holomorphic maps $(\Sigma, j) \to (M, J)$ with fixed almost complex structures (see proof of Proposition 2.3). If $\operatorname{coker} D_{\varphi,J} \neq 0$ this moduli space might not be smooth of the expected dimension $\operatorname{ind}(D_{\varphi,J})$, and this is then an obstructed deformation problem.

The maps from the left column to the middle column also have an interesting meaning. On the upper part, $H^0(T_\Sigma) \to \operatorname{kern} D_{\varphi,J}$ describes infinitesimal holomorphic reparametrizations as infinitesimal deformations of φ. On the lower part, the image of $H^1(T_\Sigma)$ in $\operatorname{coker} D_{\varphi,J}$ exhibits those obstructions of the deformation problem of φ with fixed almost complex structures that can be killed by variations of the complex structure of Σ.

The right column is maybe the most interesting. First, there are no obstructions to the deformations of φ as J-holomorphic map iff $\operatorname{coker} D^N_{\varphi,J} = 0$, provided we allow variations of the complex structure of Σ. Thus when it comes to the smoothness of the moduli space relative $\mathscr{J}$ then $\operatorname{coker} D^N_{\varphi,J}$ is much more relevant than the more traditional $\operatorname{coker} D_{\varphi,J}$. Finally, the term on the upper right corner consists of two terms, with $H^0(\mathcal{N}^{\mathrm{tor}})$ reflecting deformations of the singularities. In fact, infinitesimal deformations with vanishing component on this part can be realized by sections of $\varphi^* T_M$ with zeros of the same orders as those of $D\varphi$. This follows directly from the definition of $\mathcal{N}^{\mathrm{tor}}$. Such deformations are exactly those keeping the number of critical points of φ. Note that while $\mathcal{N}^{\mathrm{tor}}$ does not explicitly show up in the obstruction space $\operatorname{coker} D^N_{\varphi,J}$, it does influence this space by lowering the degree of N. The exact sequence also gives the (non-canonical) direct sum decomposition

$$\operatorname{kern} D^N_{\varphi,J} \oplus H^0(\mathcal{N}^{\mathrm{tor}}) = \big(\operatorname{kern} D_{\varphi,J}/H^0(T_\Sigma)\big) \oplus \operatorname{kern}\big(H^1(T_\Sigma) \to \operatorname{coker} D_{\varphi,J}\big).$$

The decomposition on the right-hand side mixes local and global contributions. The previous discussion gives the following interpretation: $\operatorname{kern} D_{\varphi,J}/H^0(T_\Sigma)$ is the space of infinitesimal deformations of φ as J-holomorphic map modulo biholomorphisms; $\operatorname{kern}\big(H^1(T_\Sigma) \to \operatorname{coker} D_{\varphi,J}\big)$ is the tangent space to the space of complex structures on Σ that can be realized by variations of φ as J-holomorphic map.

Summarizing this discussion, it is the right column that describes the moduli space of pseudo-holomorphic maps for fixed J. In particular, if $\operatorname{coker} D^N_{\varphi,J} = 0$ then the moduli space $\mathscr{M}_J \subset \mathscr{M}$ of J-holomorphic maps $\Sigma \to M$ for *arbitrary* complex structures on Σ is smooth at (φ, J, j) with tangent space

$$T_{\mathscr{M}_J,(\varphi,J,j)} = \operatorname{kern} D^N_{\varphi,J} \oplus H^0(\mathcal{N}^{\mathrm{tor}}).$$

3.3 Immersed Curves

If $\dim M = 4$ and $\varphi : \Sigma \to M$ is an immersion then N is a holomorphic line bundle and

$$N \otimes T_\Sigma \simeq \varphi^*(\det T_M).$$

Here T_M is taken as complex vector bundle. From this we are going to deduce a cohomological criterion for the surjectivity of

$$D^N_{\varphi,J} = \bar\partial + R : W^{1,p}(N) \longrightarrow L^p(N \otimes \Lambda^{0,1}).$$

By elliptic theory the cokernel of $D^N_{\varphi,J}$ is dual to the kernel of its formal adjoint operator

$$(D^N_{\varphi,J})^* : W^{1,p}(N^* \otimes \Lambda^{1,0}) \longrightarrow L^p(N^* \otimes \Lambda^{1,1}).$$

Note that $(D^N_{\varphi,J})^* = \bar\partial - R^*$ is also of Cauchy–Riemann type. Now in dimension 4 the bundle N is a holomorphic line bundle over Σ. In a local holomorphic trivialization $(D^N_{\varphi,J})^*$ therefore takes the form $f \mapsto \bar\partial f + \alpha f + \beta \bar f$ for some functions α, β. Solutions of such equations are called *pseudo-analytic* [Ve]. While related this notion predates pseudo-holomorphicity and should not be mixed up with it.

Lemma 3.3 *Let $\alpha, \beta \in L^p(\Delta)$ and let $f \in W^{1,p}(\Delta) \setminus \{0\}$ fulfill*

$$\partial_{\bar z} f + \alpha f + \beta \bar f = 0.$$

Then all zeros of f are isolated and positive.

Proof. This is another application of the Carleman trick, cf. the proof of Lemma 2.9. Replacing α by $\alpha + \beta \cdot \bar f / f$ reduces to the case $\beta = 0$. Note that $\bar f / f$ is bounded, so $\beta \cdot \bar f / f$ stays in L^p. By Proposition 3.1 there exists $g \in W^{1,p}(\Delta)$ solving $\partial_{\bar z} g = \alpha$. Then

$$\partial_{\bar z}(e^g f) = e^g(\partial_{\bar z} g \cdot f + \partial_{\bar z} f) = e^g(\alpha f + \partial_{\bar z} f) = 0.$$

Thus the diffeomorphism $\Psi : (z, w) \mapsto (z, e^{g(z)} w)$ transforms the graph of f into the graph of a holomorphic function. □

Here is the cohomological unobstructedness theorem for immersed curves.

Proposition 3.4 *Let (M, J) be a 4-dimensional almost complex manifold, and let $\varphi : \Sigma \to M$ be an immersed J-holomorphic curve with $c_1(M) \cdot [C] > 0$. Then the moduli space $\mathscr{M}_J$ of J-holomorphic maps to M is smooth at (j, φ, J).*

Proof. By the previous discussion the result follows once we show the vanishing of $\operatorname{kern}(D^N_{\varphi,J})^*$. An element of this space is a section of the holomorphic line bundle $N^* \otimes \Lambda^{1,0}$ over Σ of degree

$$\deg(N^* \otimes \Lambda^{1,0}) = \deg(\det T^*_M|_C) = -c_1(M) \cdot [C] < 0.$$

In a local holomorphic trivialization it is represented by a pseudo-analytic function. Thus by Lemma 3.3 and the degree computation it has to be identically zero. □

3.4 Smoothings of Nodal Curves

A pseudo-holomorphic curve $C \subset M$ is a *nodal curve* if all singularities of C are transversal unions of two smooth branches. It is natural to consider a nodal curve as the image of an injective map from a *nodal Riemann surface*. A nodal Riemann surface is a union of Riemann surfaces Σ_i with finitely many disjoint pairs of points identified. The identification map $\hat{\Sigma} := \coprod_i \Sigma_i \to \Sigma$ from the disjoint union is called *normalization*. (This notion has a more precise meaning in complex analysis.) A map $\varphi : \Sigma \to M$ is pseudo-holomorphic if it is continuous and if the composition $\hat{\varphi} : \hat{\Sigma} \to \Sigma \to M$ is pseudo-holomorphic. Analogously one defines $W^{1,p}$-spaces for $p > 2$.

For a nodal curve C it is possible to extend the above discussion to include topology change by smoothing the nodes, as in $zw = t$ for $(z, w) \in \Delta \times \Delta$ and t the deformation parameter. This follows from the by now well-understood gluing construction for pseudo-holomorphic maps. There are various ways to achieve this, see for example [LiTi, Si]. They share a formulation by a family of non-linear Fredholm operators

$$\prod_i H^1(T_{\Sigma_i}) \times W^{1,p}(\varphi^* T_M) \longrightarrow L^p(\varphi^* T_M \otimes \Lambda^{0,1}),$$

parametrized by $l = \sharp$nodes gluing parameters $(t_1, \ldots, t_l) \in \mathbb{C}^l$ of sufficiently small norm. The definitions of $W^{1,p}$ and L^p near the nodes vary from approach to approach. Thus fixed $(t_1, \ldots, t_l)$ gives the deformation problem with given topological type as discussed above, and putting $t_i = 0$ means keeping the ith node. The linearization $D'_{\varphi,J}$ of this operator for $(t_1, \ldots, t_l) = 0$ and fixed almost complex structures fits into a diagram very similar to the one above:

$$\begin{array}{ccccccccc}
0 & \rightarrow & W^{1,p}(T_\Sigma[A]) & \rightarrow & W^{1,p}(\varphi^* T_M) & \rightarrow & W^{1,p}(N') & \rightarrow & 0 \\
 & & \bar{\partial}_{T_\Sigma} \big\downarrow & & \big\downarrow D'_{\varphi,J} & & \big\downarrow D^{N'}_{\varphi,J} & & \\
0 & \rightarrow & L^p(T_\Sigma[A] \otimes \Lambda^{0,1}) & \rightarrow & L^p(\varphi^* T_M \otimes \Lambda^{0,1}) & \rightarrow & L^p(N' \otimes \Lambda^{0,1}) & \rightarrow & 0.
\end{array}$$

In the right column N' denotes the image of N under the normalization map:

$$N' := \bigoplus_i \varphi_i^* T_M / D\varphi_i(T_{\Sigma_i}).$$

Thus N' is a holomorphic line bundle on Σ only away from the nodes, while near a node it is a direct sum of line bundles on each of the two branches. Note that surjectivity of $W^{1,p}(\Sigma, \varphi^* T_M) \to W^{1,p}(\Sigma, N')$ is special to the nodal case in dimension 4 since it requires the tangent spaces of the branches at a node $P \in \varphi(C)$ to generate $T_{M,P}$. A crucial observation then is that the obstructions to this extended deformation problem can be computed on the normalization:

$$\operatorname{coker} D^{N'}_{\varphi,J} = \operatorname{coker} D^N_{\hat{\varphi},J}.$$

This follows by chasing the diagrams. Geometrically the identity can be understood by saying that it is the same to deform φ as pseudo-holomorphic map or to deform each of the maps $\Sigma_i \to M$ separately. In fact, the position of the identification points of the Σ_i are uniquely determined by the maps to M. In view of the cohomological unobstructedness theorem and the implicit function theorem relative $\mathscr{J} \times \mathbb{C}^l$ we obtain the following strengthening of Proposition 3.4.

Proposition 3.5 *[Sk] Let $C \subset M$ be a nodal J-holomorphic curve on an almost complex 4-manifold (M,J), $C = \bigcup C_i$. Assume that $c_1(M) \cdot C_i > 0$ for every i. Then the moduli space of J-holomorphic curves homologous to C is a smooth manifold of real dimension $c_1(M) \cdot C + C \cdot C$. The subset parametrizing nodal curves is locally a transversal union of submanifolds of real codimension 2. In particular, there exists a sequence of smooth J-holomorphic curves $C_n \subset M$ with $C_n \to C$ in the Hausdorff sense (C can be smoothed), and such smoothings are unique up to isotopy through smooth J-holomorphic curves.*

4 The Theorem of Micallef and White

4.1 Statement of Theorem

In this section we discuss the theorem of Micallef and White on the holomorphicity of germs of pseudo-holomorphic curves up to $\mathscr{C}^1$-diffeomorphism. The precise statement is the following.

Theorem 4.1 *(Micallef and White [MiWh], Theorem 6.2.) Let J be an almost complex structure on a neighborhood of the origin in $\mathbb{R}^{2n}$ with $J_{|0} = I$, the standard complex structure on $\mathbb{C}^n = \mathbb{R}^{2n}$. Let $C \subset \mathbb{R}^{2n}$ be a J-holomorphic curve with $0 \in C$.*

Then there exists a neighborhood U of $0 \in \mathbb{R}^{2n}$ and a $\mathscr{C}^1$-diffeomorphism $\Phi : U \to V \subset \mathbb{C}^n$, $\Phi(0) = 0$, $D\Phi_{|0} = \mathrm{id}$, such that $\Phi(C) \subset V$ is defined by complex polynomial equations. In particular, $\Phi(C)$ is a holomorphic curve.

The proof in loc. cit. might seem a bit computational on first reading, but the basic idea is in fact quite elegant and simple. As it is one substantial ingredient in our proof of the isotopy theorem, we include a discussion here. For simplicity we restrict to the two-dimensional, pseudo-holomorphically fibered situation, just as in Sect. 2.4,I. In other words, there are complex coordinates z, w with $(z,w) \mapsto z$ pseudo-holomorphic. Then w can be chosen in such a way that $T_M^{1,0} = \mathbb{C} \cdot (\partial_{\bar{z}} + b\partial_w) + \mathbb{C} \cdot \partial_{\bar{w}}$ for a $\mathbb{C}$-valued function b with $b(0,0) = 0$. This will be enough for our application and it still captures the essentials of the fully general proof.

4.2 The Case of Tacnodes

Traditionally a *tacnode* is a higher order contact point of a union of smooth holomorphic curves, its branches. The same definition makes sense pseudo-holomorphically. We assume this tangent to be $w = 0$. Then the ith branch of our pseudo-holomorphic tacnode is the image of

$$\Delta \longrightarrow \mathbb{C}^2, \quad t \longmapsto (t, f_i(t)),$$

with $f_i(0) = 0$, $Df_{i|0} = 0$. The pseudo-holomorphicity equation takes the form $\partial_{\bar{t}} f_i = b(t, f_i(t))$. For $i \neq j$ this gives the equation

$$\begin{aligned} 0 &= \big(\partial_{\bar{t}} f_j - b(t, f_j)\big) - \big(\partial_{\bar{t}} f_i - b(t, f_i)\big) \\ &= \partial_{\bar{t}}(f_j - f_i) - \frac{b(t, f_j) - b(t, f_i)}{f_j - f_i} \cdot (f_j - f_i). \end{aligned}$$

Now $\big(b(t, f_j) - b(t, f_i)\big)/(f_j - f_i)$ is bounded, and hence $f_j - f_i$ is another instance of a pseudo-analytic function. The Carleman trick in Lemma 3.3 now implies that f_i and f_j osculate only to finite order. (This also follows from Aronszajn's unique continuation theorem [Ar].)

On the other hand, if $f_j - f_i = O(|t|^n)$ then $|\partial_{\bar{t}}(f_j - f_i)| = \big|b(t, f_i(t)) - b(t, f_j(t))\big| = O(|t|^n)$ by pseudo-holomorphicity and hence, by Lemma 2.11

$$f_j(t) - f_i(t) = at^n + o(|t|^n). \tag{4.1}$$

The polynomial leading term provides the handle to holomorphicity. The diffeomorphism Φ will be of the form

$$\Phi(z, w) = (z, w - E(z, w)).$$

To construct $E(z, w)$ consider the approximations

$$f_{i,n} = M.V.\big\{f_j \,\big|\, f_j - f_i = o(|t|^n)\big\}$$

to f_i, for every i and $n \geq 1$. $M.V.$ stands for the arithmetic mean. Because we are dealing with a tacnode, $f_{i,1} = f_{j,1}$ for every i, j, and by finiteness of osculation orders, there exists N with $f_{i,n} = f_i$ for every $n \geq N$. Now (4.1) gives $a_{i,n} \in \mathbb{C}$ and functions $E_{i,n}$ with

$$f_{i,n} - f_{i,n-1} = a_{i,n} t^n + E_{i,n}(t), \quad E_{i,n}(t) = o(|t|^n).$$

Summing from $n = 1$ to N shows

$$f_i = \sum_{n=2}^{N} a_{i,n} t^n + f_{i,1}(t) + \sum_{n=2}^{N} E_{i,n}(t). \tag{4.2}$$

The rest is a matter of merging the various $E_{i,n}$ into $E(z, w)$ to achieve $\Phi(t, f_i(t)) = (t, \sum_{n=2}^{N} a_{i,n} t^n)$. We are going to set $E = \sum_{i=1}^{N} E_i$ with

$E_i(z,w) = E_{i,n}(z)$ in a strip osculating to order n to the graph of $f_{i,n}$. More precisely, choose a smooth cut-off function $\rho : \mathbb{R}_{\geq 0} \to [0,1]$ with $\rho(s) = 1$ for $s \in [0, 1/2]$ and $\rho(s) = 0$ for $s \geq 1$. Then for $n = 1$ define

$$E_1(z,w) = \rho\left(\frac{|w|}{|z|^{3/2}}\right) \cdot f_{i,1}(z).$$

The exponent $3/2$ may be replaced by any number in $(1,2)$. For $n \geq 2$ take

$$E_n(z,w) = \begin{cases} \rho\left(\frac{|w - f_{i,n}(z)|}{\varepsilon|z|^n}\right) \cdot E_{i,n}(z), & |w - f_{i,n}(z)| \leq \varepsilon|z|^n \\ 0, & \text{otherwise.} \end{cases}$$

To see that this is well-defined for ε and $|z|$ sufficiently small, note that by construction $f_{i,n}$ and $f_{j,n}$ osculate polynomially to dominant order. If this order is larger than n then $f_{i,n} = f_{j,n}$. Otherwise $|f_{j,n}(z) - f_{i,n}(z)| > \varepsilon|z|^n$ for ε and $|z|$ sufficiently small. See Fig. 4.1 for illustration. The distinction between the cases $n = 1$ and $n > 1$ at this stage could be avoided by formally setting $E_{i,1} = f_{i,1}$. However, to also treat branches with different tangents later on, we want Φ to be the identity outside a region of the form $|z| \leq |w|^a$ with $a > 1$.

Finally note that by construction $E_n(t, f_i(t)) = E_{i,n}(t)$, and hence

$$\begin{aligned} \Phi(t, f_i(t)) &= \big(t, f_i(t) - \sum_n E_n(t, f_i(t))\big) \\ &= \big(t, f_i(t) - \sum_{n=2}^{N} E_{i,n}(t) - f_{i,1}(t)\big) \overset{(4.2)}{=} (t, \sum_{n=1}^{N} a_{i,n} t^n). \end{aligned}$$

This finishes the proof of Theorem 4.1 for the case of tacnodes under the made simplifying assumptions.

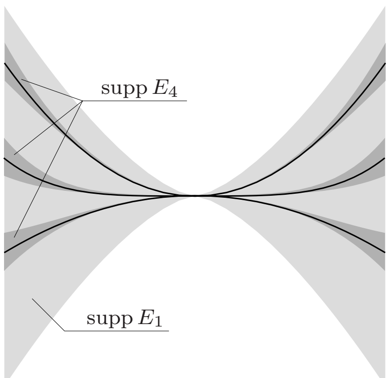

Fig. 4.1. The supports of E_n

As for differentiability it is clear that Φ is smooth away from $(0,0)$. But $\partial_w E_n$ involves the term

$$\left|\partial_{\bar w}\rho\Big(\frac{|w-f_{i,n}|}{\varepsilon\cdot|z|^n}\Big)\right| \;=\; \frac{1}{|z|^n}\cdot o(|z|^n)=o(1).$$

Thus φ may only be $\mathscr{C}^1$.

4.3 The General Case

In the general case the branches of C are images of pseudo-holomorphic maps

$$\Delta\longrightarrow\mathbb{C}^2,\quad t\longmapsto(t^{Q_i},f_i(t)),$$

for some $Q_i\in\mathbb{N}$. Note our simplifying assumption that the projection onto the first coordinate be pseudo-holomorphic. By composing with branched covers $t\mapsto t^{m_i}$ we may assume $Q_i=Q$ for all i. The pseudo-holomorphicity equation then reads $\partial_{\bar t}f_i=Q\bar t^Q b(t,f_i(t))$. The proof now proceeds as before but we deal with multi-valued functions $E_{i,n}$ of z. A simple way to implement this is by enlarging the set of functions f_i by including compositions with $t\mapsto\zeta t$ for all Qth roots of unity ζ. The definition of $E_n(z,w)$ then reads

$$E_n(z,w)=\begin{cases}\rho\Big(\dfrac{|w-f_{i,n}(t)|}{\varepsilon|t|^n}\Big)\cdot E_{i,n}(t), & |w-f_{i,n}(t)|\le\varepsilon|z|^n\\ 0, & \text{otherwise,}\end{cases}$$

for any t with $t^Q=z$. This is well-defined as before since the set of functions f_i is invariant under composition with $t\mapsto\zeta t$ whenever $\zeta^Q=1$.

Finally, if C has branches with different tangent lines do the construction for the union of branches with given tangent line separately. The diffeomorphisms obtained in this way are the identity outside of trumpet-like sets osculating to the tangent lines as in Fig. 4.2. Hence their composition maps each branch to an algebraic curve as claimed in the theorem.

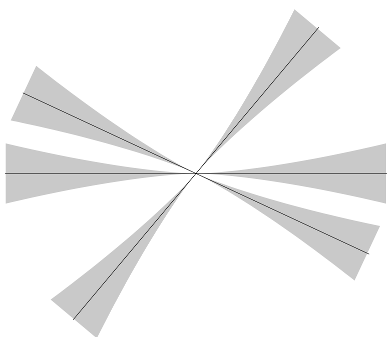

Fig. 4.2. Different tangent lines

5 Unobstructedness II: The Integrable Case

5.1 Motivation

We saw in Sect. 3 that if $c_1(M)$ evaluates strictly positively on a smooth pseudo-holomorphic curve in a four-manifold then this curve has unobstructed deformations. The only known generalizations to singular curves rely on parametrized deformations. These deformations preserve the geometric genus and, by the genus formula, lead at best to a nodal curve. Unobstructedness in this restricted class of deformations is a stronger statement, which thus requires stronger assumptions. In particular, the types of the singular points enter as a condition, and this limits heavily the usefulness of such results for the isotopy problem. Note these problems do already arise in the integrable situation. For example, not every curve on a complex surface can be deformed into a curve of the same geometric genus with only nodes as singularities. This is a fact of life and can not be circumvented by clever arguments.

Thus we need to allow an increase of geometric genus. There are two points of views for this. The first one looks at a singular pseudo-holomorphic curve as the image of a pseudo-holomorphic map from a nodal Riemann surface, as obtained by the Compactness Theorem (Sect. 7). While this is a good point of view for many general problems such as defining Gromov-Witten invariants, we are now dealing with maps from a *reducible domain* to a complex two-dimensional space. This has the effect that (total) degree arguments alone as in Sect. 3 do not give good unobstructedness results. For example, unobstructedness fails whenever the limit stable map contracts components of higher genus. Moreover, it is not hard to show that not all stable maps allowed by topology can arise. There are subtle and largely ununderstood analytical obstructions preventing this. Again both problems are inherited from the integrable situation. A characterization of holomorphic or algebraic stable maps that can arise as limit is an unsolved problem in algebraic geometry. There is some elementary discussion on this in the context of stable reduction in [HaMo]. If a good theory of unobstructedness in the integrable case from the point of view of stable maps was possible it would likely generalize to the pseudo-holomorphic setting. Unfortunately, such theory is not in sight.

The second point of view considers deformations of the limit as a cycle. The purpose of this section is to prove unobstructedness *in the integrable situation* under the mere assumption that every component evaluates strictly positively on $c_1(M)$. This is the direct analogue of Proposition 3.4. Integrability is essential here as the analytic description of deformations of pseudo-holomorphic cycles becomes very singular under the presence of multiple components, see [SiTi2].

5.2 Special Covers

We now begin with the proper content of this section. Let M be a complex surface, not necessarily compact but without boundary. In our application M will be a tubular neighborhood of a limit pseudo-holomorphic cycle in a rational ruled almost complex four-manifold M, endowed with a *different*, integrable almost complex structure. Let $C = \sum_i m_i C_i$ be a compact holomorphic cycle on M of complex dimension one. We assume C to be *effective*, that is $m_i > 0$ for all i. There is a general theory on the existence of moduli spaces of compact holomorphic cycles of any dimension, due to Barlet [Bl]. For the case at hand it can be greatly simplified and this is what we will do in this section.

The approach presented here differs from [SiTi3], Sect. 4 in being strictly local around C. This has the advantage to link the linearization of the modeling PDE to cohomology on the curve very directly. However, such a treatment is not possible if C contains fibers of our ruled surface $M \to S^2$, and then the more global point of view of [SiTi3] becomes necessary.

The essential simplification is the existence of special covers of a neighborhood of $|C| = \bigcup_i C_i$.

Hypothesis 5.1 *There exists an open cover $\mathscr{U} = \{U_\mu\}$ of a neighborhood of C in M with the following properties:*

1. $\mathrm{cl}(U_\mu) \simeq A_\mu \times \bar{\Delta}$*, where A_μ is a compact, connected Riemann surface with non-empty boundary and $\bar{\Delta} \subset \mathbb{C}$ is the closed unit disk.*
2. $(A_\mu \times \partial\Delta) \cap |C| = \emptyset$.
3. $U_\kappa \cap U_\mu \cap U_\nu = \emptyset$ *for any pairwise different κ, μ, ν. The fiber structures given by projection to A_μ are compatible on overlaps.*

The symbol "$\simeq$" denotes biholomorphism. The point of these covers is that on U_μ there is a holomorphic projection $U_\mu \to \mathrm{inn}(A_\mu)$ with restriction to $|C|$ proper and with finite fibers, hence a branched covering. Such cycles in $A_\mu \times \mathbb{C}$ of degree b over A_μ have a description by b holomorphic functions via Weierstraß polynomials, see below. Locally this indeed works in any dimension. But it is generally impossible to make the projections match on overlaps as required by (3), not even locally and for smooth C. The reason is that the analytic germ of M along C need not be isomorphic to the germ of the holomorphic normal bundle along the zero section. This almost always fails in the positive ("ample") case that we are interested in, see [La]. For our application to cycles in ruled surfaces, however, we can use the bundle projection, provided C does not contain fiber components, see Lemma 5.1 below. Without this assumption the arguments below still work but require a more difficult discussion of change of projections, as in [Bl]. As the emphasis of these lectures is on explaining why things are true rather than generality we choose to impose this simplifying assumption.

Another, more severe, simplification special to one-dimensional C is that we do not allow triple intersections. This implies that A_μ can not be contractible for all μ unless all components of C are rational, and hence our charts have a certain global flavor. Under the presence of triple intersections it seems impossible to get our direct connection of the modeling of the moduli space with cohomological data on C.

Lemma 5.1 *Assume there exists a map $p : M \to S$ with $\dim S = 1$ such that p is a holomorphic submersion near $|C|$ and no component of $|C|$ contains a connected component of a fiber of p. Then Hypothesis 5.1 is fulfilled.*

Proof. Denote by $Z \subset M$ the finite set of singular points of $|C|$ and of critical points of the projection $|C| \to S$. Add finitely many points to Z to make sure that Z intersects each irreducible component of $|C|$. For $P \in |C|$ let $F = p^{-1}(p(P))$ be the fiber through P. Then $F \cap |C|$ is an analytic subset of F that by hypothesis does not contain a connected component of F. Since F is complex one-dimensional P is an isolated point of $F \cap C$. Because p is a submersion there exists a holomorphic chart $(z, w) : U(P) \to \mathbb{C}^2$ in a neighborhood of P in M with $(z, w)(P) = 0$ and $z = \text{const}$ describing the fibers of p. We may assume $U(P)$ to be so small that $|C|$ is the zero locus of a holomorphic function h defined on all of $U(P)$. Let $\varepsilon > 0$ be such that $F \cap w^{-1}(B_\varepsilon(0)) = \{P\}$. Then $\min\big\{|f(0, w)| \,\big|\, |w| = \varepsilon\big\} > 0$. Hence for $\delta > 0$ sufficiently small $\min\{|f(z, w)| \,|\, |w| = \varepsilon, |z| \le \delta\}$ remains nonzero. This shows $|C| \cap (z, w)^{-1}(\overline{B}_\delta(0) \times \partial B_\varepsilon(0)) = \emptyset$. This verifies Hypothesis 5.1,1 and 2 for $U(P)$, while the compatibility of fiber structures in 5.1,3 holds by construction. This defines elements $U_1, \ldots, U_{\sharp Z}$ of the open cover $\mathscr{U}$ intersecting Z.

To finish the construction we define one more open set U_0 as follows. This construction relies on Stein theory, see e.g. [GrRe] or [KpKp] for the basics. (With some effort this can be replaced by more elementary arguments, but it does not seem worthwhile doing here for a technical result like this.) Choose a Riemannian metric on M. Let δ be so small that $B_{2\delta}(Z) \subset \bigcup_{\nu \ge 1} U_\nu$. Since we have chosen Z to intersect each irreducible component of $|C|$, the complement $|C| \setminus \operatorname{cl} B_\delta(Z)$ is a union of open Riemann surfaces. Thus $|C| \setminus \operatorname{cl} B_\delta(Z)$ is a Stein submanifold of $M \setminus \operatorname{cl} B_\delta(Z)$, hence has a Stein neighborhood $W \subset M \setminus \operatorname{cl} B_\delta(Z)$. Now every hypersurface in a Stein manifold is defined by one global holomorphic function, say f in our case. So f is a global version of the fiber coordinate w before. Then for δ sufficiently small the projection

$$U_0 := \big\{P \in W \setminus \operatorname{cl} B_{3\delta/2}(Z) \,\big|\, |f(P)| < \delta\big\} \xrightarrow{p} S$$

factors holomorphically through $\pi : U_0 \to |C| \setminus \operatorname{cl} B_{3\delta/2}(Z)$. This is an instance of so-called Stein factorization, which contracts the connected components of the fibers of a holomorphic map. Choosing δ even smaller this factorization gives a biholomorphism

$$U_0 \xrightarrow{(\pi, f)} |C| \setminus \operatorname{cl} B_\delta(Z) \times \varDelta$$

extending to $\operatorname{cl}(U_0)$. Hence $U_0, U_1, \ldots, U_{\sharp Z}$ provides the desired open cover $\mathscr{U}$. $\square$

5.3 Description of the Deformation Space

Having a cover fulfilling Hypothesis 5.1 it is easy to describe the moduli space of small deformations of C as the fiber of a holomorphic, non-linear Fredholm map. We use Čech notation $U_{\mu\nu} = U_\mu \cap U_\nu$, $U_{\kappa\mu\nu} = U_\kappa \cap U_\mu \cap U_\nu$. Write $V_\mu = \mathrm{inn}(A_\mu)$. Fix $p > 2$ and denote by $\mathcal{O}^p(V_\mu)$ the space of holomorphic functions on $\mathrm{inn}(V_\mu)$ of class L^p. This is a Banach space with the L^p-norm defined by a Riemannian metric on A_μ, chosen once and for all. Similarly define $\mathcal{O}^p(V_{\mu\nu})$ and $\mathcal{O}^p(V_{\kappa\mu\nu})$.

To describe deformations of C let us first consider the local situation on $V_\mu \times \Delta \subset M$. For this discussion we drop the index μ. Denote by w the coordinate on the unit disk. For any holomorphic cycle C' on $V \times \Delta$ with $|C'| \cap (V \times \partial\Delta) = \emptyset$ the Weierstraß preparation theorem gives bounded holomorphic functions $a_1, \ldots, a_b \in \mathcal{O}^p(V)$ with $w^b + a_1 w^{b-1} + \ldots + a_b = 0$ describing C', see e.g. [GrHa]. Here b is the relative degree of C' over V, and everything takes multiplicities into account. The tuple $(a_1, \ldots, a_b)$ should be thought of as a holomorphic map from V to the b-fold symmetric product $\mathrm{Sym}^b\,\mathbb{C}$ of $\mathbb{C}$ with itself.

Digression on $\mathrm{Sym}^b\,\mathbb{C}$. By definition $\mathrm{Sym}^b\,\mathbb{C}$ is the quotient of $\mathbb{C} \times \cdots \times \mathbb{C}$ by the permutation action of the symmetric group S_b on b letters. Quite generally, if a finite group acts on a complex manifold or complex space X then the topological space X/G has naturally the structure of a complex space by declaring a function on X/G holomorphic whenever its pull-back to X is holomorphic. For the case of the permutation action on the coordinates of $\mathbb{C}^b$ we claim that the map

$$\Phi : \mathrm{Sym}^d\,\mathbb{C} \longrightarrow \mathbb{C}^d$$

induced by $(\sigma_1, \ldots, \sigma_b) : \mathbb{C}^b \to \mathbb{C}^b$ is a biholomorphism. Here

$$\sigma_k(w_1, \ldots, w_b) = \sum_{1 \le i_1 < i_2 < \ldots < i_k \le b} w_{i_1} w_{i_2} \ldots w_{i_b}$$

is the ith elementary symmetric polynomial. In fact, set-theoretically $\mathrm{Sym}^b\,\mathbb{C}$ parametrizes unordered tuples of b not necessarily disjoint points in $\mathbb{C}$. By the fundamental theorem of algebra there is precisely one monic (leading coefficient equal to 1) polynomial of degree b having this zero set, with multiplicities. The coefficients of this polynomial are the elementary symmetric functions in the zeros. This shows that Φ is bijective. Now any symmetric holomorphic function in $w_1, \ldots, w_b$ is a holomorphic function in $\sigma_1, \ldots, \sigma_b$. Thus Φ is a biholomorphism. By the same token, if w, w' are two holomorphic coordinates on an open set $W \subset \mathbb{C}$ then the induced holomorphic coordinates σ_i, σ_i', $i = 1, \ldots, b$, are related by a biholomorphic transformation. Note however that something peculiar is happening to the differentiable structure: Not every S_b-invariant smooth function on $\mathbb{C}^b$ leads to a smooth function on $\mathrm{Sym}^b\,\mathbb{C}$ for the differentiable structure coming from holomorphic geometry. As an example consider $b = 2$ and the function $f(w_1, w_2) = (w_1 - w_2)(\bar{w}_1 - \bar{w}_2) = |w_1 - w_2|^2$. This is the pull-back of $|\sigma_1^2 - \sigma_2|$, which is only a Lipschitz function on $\mathrm{Sym}^2\,\mathbb{C}$. Another, evident feature of symmetric products is that a neighborhood of a point $\sum_i m_i P_i$ of $\mathrm{Sym}^b\,\mathbb{C}$ with the P_i pairwise disjoint, is canonically biholomorphic to an open set in $\prod_i \mathrm{Sym}^{m_i}\,\mathbb{C}$. End of digression.

From this discussion it follows that we can compare Weierstraß representations with compatible projections. Let $(a_\mu : V_\mu \to \mathrm{Sym}^{b_\mu} \mathbb{C}) \in \mathcal{O}^p(V_\mu)^{b_\mu}$ be the Weierstraß representations for the given cycle C. Let $V^b_{\mu\nu} \subset V_{\mu\nu}$ be the union of connected components where the covering degree is equal to b. Note that $V^b_{\mu\nu} = \emptyset$ whenever $b > \min\{b_\mu, b_\nu\}$. For a'_μ sufficiently close to a_μ in L^p there is a cycle C'_μ in U_μ with Weierstraß representation a'_μ. Denote by $\mathscr{F}_\mu$ a sufficiently small neighborhood of a_μ in $\mathcal{O}^p(V_\mu)^{b_\mu}$ where this is the case. For every μ, ν, b let $\mathscr{F}^b_{\mu\nu} = \mathcal{O}^p(V^b_{\mu\nu})^b$, viewed as space of maps $V^b_{\mu\nu} \to \mathrm{Sym}^b \mathbb{C}$. The above discussion gives comparison maps

$$\Theta^b_{\mu\nu} :\ \mathscr{F}_\nu \longrightarrow \mathscr{F}^b_{\mu\nu}.$$

Define the gluing map by

$$\Theta : \prod_\mu \mathscr{F}_\mu \longrightarrow \prod_{\mu<\nu, b} \mathscr{F}^b_{\mu\nu}, \quad \Theta(a'_\mu) = (a'_\mu - \Theta^b_{\mu\nu}(a'_\nu))_{\mu\nu}. \tag{5.1}$$

Clearly $(a'_\mu)_\mu$ glues to a holomorphic cycle iff $\Theta(a'_\mu) = 0$. We will see that Θ is a holomorphic Fredholm map with kernel and cokernel canonically isomorphic to the first two cohomology groups of the normal sheaf of C in M. To follow this plan, two more digressions are necessary. One to explain the notion of normal sheaf and one for the needed properties of sheaf cohomology.

5.4 The Holomorphic Normal Sheaf

One definition of tangent vector of a differentiable manifold M at a point P works by emphasizing its property of *derivation*. Let $\mathscr{C}^\infty_{M,P}$ be the space of germs of smooth functions at P. An element of this space is represented by a smooth function defined on a neighborhood of P, and two such functions give the same element if they agree on a common neighborhood. Applying a tangent vector $X \in T_{M,P}$ on representing functions defines an $\mathbb{R}$-linear map $D : \mathscr{C}^\infty_{M,P} \to \mathbb{R}$. For $f, g \in \mathscr{C}^\infty_{M,P}$ Leibniz' rule gives

$$X(fg) = fX(g) + gX(f).$$

In particular, $X(f^2) = 0$ if $f(P) = 0$ and $X(1) = X(1^2) = 2X(1) = 0$. Thus the interesting part of D is the induced map on $\mathfrak{m}^\infty_P/(\mathfrak{m}^\infty_p)^2$ where $\mathfrak{m}^\infty_P = \{f \in \mathscr{C}^\infty_{M,P} \,|\, f(P) = 0\}$ is the maximal ideal of the ring $\mathscr{C}^\infty_{M,P}$. We claim that the map $T_{M,P} \to \mathrm{Hom}_\mathbb{R}(\mathfrak{m}^\infty_P/(\mathfrak{m}^\infty_p)^2, \mathbb{R})$ is an isomorphism. In fact, let $x_1, \ldots, x_n$ be coordinates of M around P. Then $X = \sum_i a_i \partial_{x_i}$ applied to x_i yields a_i, so the map is injective. On the other hand, if $f \in \mathfrak{m}^\infty_P$ then by the Taylor formula $f - \sum_i \partial_{x_i} f(P) \cdot x_i \in (\mathfrak{m}^\infty)^2$, and hence any linear map $\mathfrak{m}^\infty_P/(\mathfrak{m}^\infty_p)^2 \to \mathbb{R}$ is determined by its values on x_i. This gives the well-known canonical identification

$$T_{M,P} = \mathrm{Hom}_\mathbb{R}(\mathfrak{m}^\infty_P/(\mathfrak{m}^\infty_p)^2, \mathbb{R}).$$

If M is a complex manifold then the holomorphic tangent space at P can similarly be described by considering holomorphic functions and $\mathbb{C}$-linear maps.

Now if $Z \subset M$ is a submanifold the same philosophy applies to the normal bundle. For this and what follows we will inevitably need some elementary sheaf theory that we are also trying to explain briefly.

An (abelian) *sheaf* is an association of an abelian group $\mathcal{F}(U)$ (the *sections* of $\mathcal{F}$ over U) to every open subset $U \subset M$, together with restriction homomorphisms $\rho_{VU} : \mathcal{F}(U) \to \mathcal{F}(V)$ whenever $V \subset U$. The restriction maps must be compatible with composition. Moreover, the following sheaf axioms must hold for every covering $\{U_i\}$ of an open set $U \subset M$. Write $U_{ij} = U_i \cap U_j$.

(S1) (*local-global principle*) If $s \in \mathcal{F}(U)$ and $\rho_{U_i U}(s) = 0$ for all i then $s = 0$.
(S2) (*gluing axiom*) Given $s_i \in \mathcal{F}(U_i)$ with $\rho_{U_{ij}U_i}(s_i) = \rho_{U_{ij}U_j}(s_j)$ for every i, j there exists $s \in \mathcal{F}(U)$ with $s_i = \rho_{U_i U}(s)$.

The following are straightforward examples:

1. $\mathscr{C}^\infty_M : U \mapsto \{f : U \to \mathbb{R} \text{ smooth}\}$ with restriction of functions defining the restriction maps.
2. For $E \downarrow M$ a fiber bundle the sheaf $\mathscr{C}^\infty(E)$ of smooth sections of E.
3. For M a complex manifold, $\mathcal{O}_M : U \mapsto \{f : U \to \mathbb{C} \text{ holomorphic}\}$. This is a subsheaf of the sheaf of complex valued smooth functions on M.
4. For G an abelian group the constant sheaf
$$G_M : U \mapsto \{f : U \to G \text{ locally constant}\},$$
ρ_{VU} the restriction. The sections of this sheaf over any *connected* open set are identified with elements of G.

A homomorphism of sheaves $\mathcal{F} \to \mathcal{G}$ is a system of homomorphisms $\mathcal{F}(U) \to \mathcal{G}(U)$ compatible with restriction.

Returning to the normal bundle of $Z \subset M$ let $\mathcal{I}^\infty_Z$ be the sheaf of smooth functions vanishing along Z. Then a normal vector field ν along Z on $U \subset Z$ induces a well-defined map
$$\mathcal{I}^\infty_Z(U) \longrightarrow \mathscr{C}^\infty_Z(U), \quad f \longmapsto \big(\tilde{\nu}(f)\big)\big|_Z,$$
where $\tilde{\nu}$ is a lift of ν to a vector field on M defined on a neighborhood of U in M. As we restrict to Z after evaluation and since f vanishes along Z the result does not depend on the choice of lift. Now as with $T_{M,P}$ one checks that $(\mathcal{I}^\infty_Z(U))^2$ maps to zero and that the map of sheaves
$$\mathscr{C}^\infty(N_{Z|M}) \longrightarrow \mathcal{H}om_{\mathscr{C}^\infty_Z}(\mathcal{I}^\infty_Z/(\mathcal{I}^\infty_Z)^2, \mathscr{C}^\infty_Z)$$
is an isomorphism. A section over U of the Hom-sheaf on the right is a $(\mathscr{C}^\infty_Z)|_U$-linear sheaf homomorphism $(\mathcal{I}^\infty_Z/(\mathcal{I}^\infty_Z)^2)|_U \to (\mathscr{C}^\infty_Z)|_U$. Note that multiplication of sections $\mathcal{I}^\infty_Z/(\mathcal{I}^\infty_Z)^2$ by sections of $\mathscr{C}^\infty_Z$ is well-defined, so this makes sense.

Generally one has to be careful here where to put the brackets because the notions of quotient and Hom-sheaves are a bit delicate. For example, if $\mathcal{F} \to \mathcal{G}$ is a sheaf homomorphism the sheaf axioms need not hold for $\mathcal{Q} : U \mapsto \mathcal{G}(U)/\mathcal{F}(U)$. The standard example for this is the inclusion $\mathbb{Z}_M \to \mathscr{C}^\infty_M$ with $M = \mathbb{R}^2 \setminus \{0\}$ or any other non-simply connected space. On the complements of the positive and negative real half-axes $1/2\pi$ of the angle in polar coordinates give sections of $\mathscr{C}^\infty_M$ agreeing on $\mathbb{R}^2 \setminus \mathbb{R}$ modulo integers, but they do not glue to a single-valued function on M.

In any case, there is a canonical procedure to force the sheaf axioms, by taking the sheaf of continuous sections of the space of stalks

$$\text{Ét}(\mathcal{Q}) := \coprod_{P \in M} \varinjlim_{U \ni P} \mathcal{Q}(U)$$

(étale space) of $\mathcal{Q}$. Every $s \in \mathcal{Q}(U)$ induces a section of $\text{Ét}(\mathcal{Q})$ over U, and the images of these sections are taken as basis for the topology of $\text{Ét}(\mathcal{Q})$. In writing down quotient or Hom-sheaves it is understood to apply this procedure. This is the general definition that we will need for the holomorphic situation momentarily, but it is in fact not needed for sheaves that allow partitions of unity (fine sheaves) such as $\mathcal{I}^\infty_Z$. So a section of $\text{Hom}_{\mathscr{C}^\infty_Z}(\mathcal{I}^\infty_Z/(\mathcal{I}^\infty_Z)^2, \mathscr{C}^\infty_Z)$ over U is indeed just a $\mathscr{C}^\infty_Z(U)$-linear map of section spaces $\mathcal{I}^\infty_Z(U)/(\mathcal{I}^\infty_Z(U))^2 \to \mathscr{C}^\infty_Z(U)$.

Now if M is a complex manifold and $Z \subset M$ is a complex submanifold then $N_{Z|M}$ is a holomorphic vector bundle. Let $\mathcal{I}_Z \subset \mathcal{O}_M$ be the sheaf of holomorphic functions on M vanishing along Z. Then the natural map

$$\mathcal{O}(N_{Z|M}) \longrightarrow \text{Hom}_{\mathcal{O}_Z}(\mathcal{I}_Z/\mathcal{I}_Z^2, \mathcal{O}_Z)$$

is an isomorphism. Explicitly, if Z is locally given by one equation $f = 0$ then f generates $\mathcal{I}_Z/\mathcal{I}_Z^2$ as an $\mathcal{O}_Z$-module. Hence a section φ of the Hom-sheaf on the right is uniquely defined by $\varphi(f)$, a holomorphic function on Z. This provides an explicit local identification of $\mathcal{O}(N_{Z|M})$ with $\mathcal{O}_Z$, which is nothing but a local holomorphic trivialization of $N_{Z|M}$.

Now the whole point of this discussion is that it generalizes well to the non-reduced situation. Let us discuss this in the most simple situation of a multiple point mP in $M = \mathbb{C}$. Taking P the origin and w for the coordinate on $\mathbb{C}$ (z will have a different meaning below) we have $\mathcal{I}_{mP} = \mathcal{O}_\mathbb{C} \cdot w^m$ and $\mathcal{O}_{mP} = \mathcal{O}_\mathbb{C}/\mathcal{I}_{mP}$ is an m-dimensional complex vector space with basis $1, w, \ldots, w^{m-1}$. A homomorphism $\mathcal{I}_{mP}/(\mathcal{I}_{mP})^2 \to \mathcal{O}_{mP}$ is uniquely defined by $(\alpha_0, \ldots, \alpha_m) \in \mathbb{C}^m$ via

$$w^m \longmapsto \alpha_0 + \alpha_1 w + \ldots + \alpha_{m-1} w^{m-1}.$$

This fits well with limits as follows. Consider mP as the limit of m pairwise different points $Z_t := \{P_1(t), \ldots, P_m(t)\}$, $t > 0$, given by the vanishing of $f_t := w^m + a_1(t) w^{m-1} + \ldots + a_m(t)$ where $a_i(t) \xrightarrow{t \to 0} 0$. Then $\mathcal{I}_{Z(t)}/(\mathcal{I}_{Z(t)})^2 = \bigoplus_i \mathcal{I}_{P_i(t)}/\,(\mathcal{I}_{P_i(t)})^2$ is the sheaf with one copy of $\mathbb{C}$ at each point of $Z(t)$ (a "skyscraper sheaf"). Thus $\bigoplus_i T_{P_i(t)} = \text{Hom}(\mathcal{I}_{Z(t)}/(\mathcal{I}_{Z(t)})^2, \mathcal{O}_{Z(T)})$. Now

$\mathcal{I}_{Z(t)}/(\mathcal{I}_{Z(t)})^2$ is globally generated over $\mathcal{O}_{Z(T)}$ by f_t, and $1, w, \ldots, w^{m-1}$ are a basis for the sections of $\mathcal{O}_{Z(T)}$ as a complex vector space. This gives an identification of $\bigoplus_i T_{P_i(t)}$ with polynomials $\alpha_0 + \alpha_1 w + \ldots + \alpha_{m-1} w^{m-1}$, so this is compatible with the description at $t = 0$! Note that a family of vector fields along $Z(t)$ has a limit for $t \to 0$ if and only if it extends to a continuous family of holomorphic vector fields in a neighborhood of P. The limit is then the limit of $(m-1)$-jets of this family.

It therefore makes sense to define, for any subspace $Z \subset M$ defined by an ideal sheaf $\mathcal{I}_Z \subset \mathcal{O}_M$, reduced or not, the *holomorphic normal sheaf*

$$\mathcal{N}_{Z|M} := \mathcal{H}om_{\mathcal{O}_Z}(\mathcal{I}_Z/\mathcal{I}_Z^2, \mathcal{O}_Z),$$

where $\mathcal{O}_Z = \mathcal{O}_M/\mathcal{I}_Z$ is the sheaf of holomorphic functions on Z. In our application we have Z the generally non-reduced subspace of M defined by the one-codimensional cycle C. By abuse of notation we use C both to denote the cycle and this subspace. Explicitly, $\mathcal{I}_C$ is the sheaf locally generated by $\prod_i f_i^{m_i}$ if $C = \sum_i m_i C_i$ and f_i vanishes to first order along C_i. Note that such a choice of generator of $\mathcal{I}_C$ gives a local identification of $\mathcal{N}_{C|M}$ with $\mathcal{O}_C$.

The importance of the normal sheaf for us comes from its relation with local deformations of holomorphic cycles.

Lemma 5.2 *Let V be a complex manifold and consider the open subset $\mathcal{M} \subset \mathcal{O}^p(V, \mathrm{Sym}^b \mathbb{C})$ of tuples $(a_1, \ldots, a_b)$ such that the zero set of $f_{(a_1,\ldots,a_b)}(z, w) := w^b + a_1(z) w^{b-1} + \ldots + a_b(z)$ is contained in $V \times \Delta$. Then there is a canonical isomorphism*

$$T_{(a_1,\ldots,a_b)}\mathcal{M} \simeq \Gamma(C, \mathcal{N}_{C|V\times\Delta}),$$

for C the holomorphic cycle in $V \times \Delta$ defined by $f_{(a_1,\ldots,a_b)}$.

Proof. The map sends $\frac{d}{dt}\big|_{t=0} (a_1(t), \ldots, a_b(t)) = (\alpha_1, \ldots, \alpha_b)$ to the section

$$f_{(a_1,\ldots,a_b)} \longmapsto w^b + \alpha_1 w^{b-1} + \ldots + \alpha_b.$$

By the above discussion every global holomorphic function on C has a unique representative of the form $w^b + \alpha_1 w^{b-1} + \ldots + \alpha_b$. Hence this map is an isomorphism. □

5.5 Computation of the Linearization

If $\mathcal{F}$ is an abelian sheaf on a topological space X and $\mathscr{U} = \{U_i\}$ is an open cover of X the Čech cohomology groups $\check{H}^k(\mathscr{U}, \mathcal{F})$ are the cohomology groups of the Čech complex $(C^\bullet(\mathscr{U}, \mathcal{F}))$ with cochains

$$C^k(\mathscr{U}, \mathcal{F}) = \prod_{i_0 < i_1 < \ldots < i_k} \Gamma(U_{i_0} \cap \ldots \cap U_{i_k}, \mathcal{F}),$$

and differentials

$$\check{d}(s_{i_0}\dots s_{i_k})_{i_0\dots i_k} = \Big(\sum_l (-1)^l s_{i_0\dots\hat{i}_l\dots i_k}\Big)_{i_0\dots i_k}.$$

By the gluing axiom (S2) it holds $\check{H}^0(\mathscr{U},\mathcal{F}) = \mathcal{F}(X)$.

Theorem 5.3 *The gluing map Θ from (5.1) is a holomorphic map with* $\operatorname{kern} D\Theta = \check{H}^0(\mathscr{V},\mathcal{N}_{C|M})$ *and* $\operatorname{coker} D\Theta = \check{H}^1(\mathscr{V},\mathcal{N}_{C|M})$.

Proof. The holomorphicity claim is evident. For the linearization we remark that the components of Θ factor through

$$(a'_\mu, a'_\nu) \longmapsto (a'_\mu, \Theta^b_{\mu\nu}(a'_\nu)) \longmapsto (a'_\mu - \Theta^b_{\mu\nu}(a'_\nu)).$$

In view of Lemma 5.2 the linearization at (a_μ, a_ν) of the components of the first map are canonically the restriction maps $\Gamma(C\cap U_\mu,\mathcal{N}_{C|M}) \to \Gamma(C\cap U_{\mu\nu},\mathcal{N}_{C|M})$, $\Gamma(C\cap U_\nu,\mathcal{N}_{C|M}) \to \Gamma(C\cap U_{\mu\nu},\mathcal{N}_{C|M})$. Hence $D\Theta$ is canonically isomorphic to the Čech complex

$$\prod_\mu \Gamma(C\cap U_\mu,\mathcal{N}_{C|M}) \longrightarrow \prod_{\mu<\nu}\Gamma(C\cap U_{\mu\nu},\mathcal{N}_{C|M}),\quad (\alpha_\mu)_\mu \longmapsto (\alpha_\mu - \alpha_\nu)_{\mu\nu}.$$

Note that $C^k(C\cap\mathscr{U},\mathcal{N}_{C|M}) = 0$ for $k>1$ because triple intersections in $\mathscr{U}$ are empty. □

5.6 A Vanishing Theorem

For this paragraph we assume some familiarity with sheaf cohomology and Serre duality on singular curves. So this section will be (even) harder to read for somebody without training in complex geometry. Unfortunately we could not find a more elementary treatment.

It is well-known that the Čech cohomology groups $\check{H}^i(C\cap\mathscr{V},\mathcal{N}_{C|M})$ are finite dimensional and canonically isomorphic to the *sheaf cohomology groups* $H^i(C,\mathcal{N}_{C|M})$, see [GrHa] and the references given there. In particular, Θ is a non-linear Fredholm map. Our aim in this paragraph is to prove surjectivity of its linearization by the following result.

Proposition 5.4 *Let $C = \sum_{i=0}^r m_i C_i$ be a compact holomorphic cycle on a complex surface M. Assume that $c_1(M)\cdot C_i > 0$ and $C_i\cdot C_i \ge 0$ for every i. Then $H^1(C,\mathcal{N}_{C|M}) = 0$.*

Proof. In view of the identification $\mathcal{N}_{C|M} = \mathcal{O}_C(C)$ a stronger statement is the vanishing of $H^1(C,\mathcal{O}_{C'}(C))$ for every effective subcycle $C'\subset C$. This latter formulation allows an induction over the sum of the multiplicities of C'.

As an auxiliary statement we first show the vanishing of $H^1(\mathcal{O}_{C_i}(C''))$ for every i and every subcycle C'' of C containing C_i. Serre duality (see e.g. [BtPeVe], Theorem II.6.1) shows that $H^1(\mathcal{O}_{C_i}(C''))$ is dual to $H^0(\mathcal{H}om_{\mathcal{O}_{C_i}}(\mathcal{O}_{C_i}(C''),\mathcal{O}_{C_i})\otimes\omega_{C_i})$. Here ω_{C_i} is the *dualizing sheaf* of C_i. If $\mathcal{K}_M$ denotes the sheaf of holomorphic sections of $\det T^*_M$ then it can be computed as $\omega_{C_i}=\mathcal{K}_M\otimes\mathcal{N}_{C_i|M}=\mathcal{K}_M\otimes\mathcal{O}_{C_i}(C_i)$. Therefore

$$\mathcal{H}om(\mathcal{O}_{C_i}(C''),\mathcal{O}_{C_i})\otimes\omega_{C_i}\simeq\mathcal{O}_{C_i}(-C'')\otimes\omega_{C_i}\simeq\mathcal{K}_M\otimes\mathcal{O}_{C_i}(C_i-C'').$$

But $\mathcal{K}_M\otimes\mathcal{O}_{C_i}(C_i-C'')$ is the sheaf of sections of a holomorphic line bundle over C_i of degree $c_1(T^*_M)\cdot C_i-(C''-C_i)\cdot C_i\le -c_1(M)\cdot C_i<0$. Here we use that C'' contains C_i and $C_i\cdot C_i\ge 0$. Because C_i is reduced and irreducible this implies that any global section of this line bundle is trivial. Hence $H^1(\mathcal{O}_{C_i}(C''))=0$.

Setting $C'=C_i$ for some i starts the induction. For the induction step assume $H^1(C,\mathcal{O}_{C'}(C))=0$ and let i be such that $C'+C_i$ is still a subcycle of C. Let $\mathcal{I}_{C'|C'+C_i}$ be the ideal sheaf of C' in $C'+C_i$. Because $\mathcal{I}_{C'+C_i}=\mathcal{I}_{C'}\cdot\mathcal{I}_{C_i}$, multiplication induces an isomorphism $\mathcal{O}_{C_i}(-C')=\mathcal{I}_{C'}\otimes(\mathcal{O}_M/\mathcal{I}_{C_i})\simeq\mathcal{I}_{C'|C'+C_i}$. Thus we have a restriction sequence

$$0\longrightarrow\mathcal{O}_{C'+C_i}(C)\otimes\mathcal{O}_{C_i}(-C')\longrightarrow\mathcal{O}_{C'+C_i}(C)\longrightarrow\mathcal{O}_{C'}(C)\longrightarrow 0.$$

Observing $\mathcal{O}_{C'+C_i}(C)\otimes\mathcal{O}_{C_i}(-C')\simeq\mathcal{O}_{C_i}(C-C')$ the long exact cohomology sequence reads

$$\ldots\longrightarrow H^1(\mathcal{O}_{C_i}(C-C'))\longrightarrow H^1(\mathcal{O}_{C'+C_i}(C))\longrightarrow H^1(\mathcal{O}_{C'}(C))\longrightarrow\ldots$$

The term on the right vanishes by induction hypothesis, while the term on the left vanishes by the auxiliary result applied to $C''=C-C'$. Hence $H^1(\mathcal{O}_{C'+C_i}(C))=0$ proving the induction step. □

5.7 The Unobstructedness Theorem

Under the assumptions of Proposition 5.4 and Hypothesis 5.1 we now have a description of deformations of C by the fiber of a holomorphic map between complex Banach manifolds whose linearization is surjective with finite dimensional kernel. Applying the implicit function theorem gives the main theorem of this lecture.

Theorem 5.5 *Let M be a complex surface and $C=\sum_i m_iC_i$ a compact holomorphic 1-cycle with $c_1(M)\cdot C_i>0$ and $C_i\cdot C_i\ge 0$ for all i. Assume that a covering $\mathscr{U}$ of a neighborhood of $|C|$ in M exists satisfying Hypotheses 5.1, 1–3. Then the space of holomorphic cycles in M is a complex manifold of dimension $\Gamma(\mathcal{N}_{C|M})$ in a neighborhood of C. Moreover, analogous statements hold for a family of complex structures on M preserving the data described in Hypothesis 5.1.*

Remark 5.6 The hypotheses of the theorem do not imply that the cycle C is the limit of smooth cycles. For example, $\Gamma(\mathcal{N}_{C|M})$ may still be trivial and then C does not deform at all. Smoothability only follows with the additional requirement that $\mathcal{N}_{C|M}$ is *globally generated*. This statement can be checked by a transversality argument inside symmetric products of $\mathbb{C}$, cf. [SiTi3], Lemma 4.8.

The proof of global generatedness at some $P \in M$ in our fibered situation $p : M \to S^2$ follows from comparing dimensions of $\Gamma(\mathcal{N}_{C|M})$ and $\Gamma(\mathcal{N}_{C|M}(-F))$ where $F = p^{-1}(p(P))$. These dimensions differ maximally, namely by the sum of the fibers of $\mathcal{N}_{C|M}$ over $|C| \cap F$ if also $H^1(\mathcal{N}_{C|M}(-F)) = 0$. This is true by the same method as seen because also $(c_1(M) - H) \cdot C_i > 0$. See [SiTi3], Lemma 4.4 for details.

In any case, if C is the limit of smooth curves then the subset of the moduli space parametrizing singular cycles is a proper analytic subset of the moduli space, which is smooth, and hence does not locally disconnect the moduli space at C. In particular, any two smoothings of C are isotopic through a family of smooth holomorphic curves staying close to C.

6 Application to Symplectic Topology in Dimension Four

One point of view on symplectic topology is as an area somewhere between complex geometry and differential topology. On one hand symplectic constructions sometimes have the same or almost the flexibility as constructions in differential topology. As an example think of Gompf's *symplectic normal sum* [Go]. It requires two symplectic manifolds M_1, M_2 with symplectic hypersurfaces $D_1 \subset M_1$, $D_2 \subset M_2$ (real codimension two) and a symplectomorphism $\Phi : D_1 \to D_2$ with a lift to an isomorphism of symplectic line bundles $\tilde{\Phi} : N_{D_1|M_1} \to \Phi^*(N^*_{D_2|M_2})$. The result is a well-defined one-parameter family of symplectic manifolds $M_1 \amalg_{\Phi,\varepsilon} M_2$; each of its elements is diffeomorphic to the union of $M_i \setminus U_i$, where U_i is a tubular neighborhood of D_i and the boundaries ∂U_i are identified via $\tilde{\Phi}$. So the difference to a purely differential-topological construction is that (1) the bundle isomorphism Φ needs to preserve the symplectic normal structure along D_i and (2) there is a finite-dimensional parameter space to the construction.

Compare this with the analogous problem in complex geometry. Here $D_i \subset M_i$ is a divisor and one can form a singular complex space $M_1 \amalg_\Phi M_2$ by gluing M_1 and M_2 via an isomorphism $\Phi : D_1 \simeq D_2$. The singularity looks locally like D_i times the union of coordinate lines $zw = 0$ in $\mathbb{C}^2$. However, even if the holomorphic line bundles $N_{D_1|M_1}$ and $\Phi^* N^*_{D_2|M_2}$ are isomorphic there need not exist a smoothing of this space [PsPi]. A smoothing would locally replace $zw = 0$ by $zw = \varepsilon$ in appropriate holomorphic coordinates. Should this smoothing problem be unobstructed, it has as local parameter space the product of a complex disk for the smoothing parameter ε and some finite-dimensional space dealing with deformations of the singular space $M_1 \amalg_\Phi M_2$. Note how the deformation parameter ε reappears on the symplectic side as gluing parameter. The most essential difference to the symplectic situation is the appearance of obstructions.

The correspondence between complex and symplectic geometry can be expected to be especially interesting in dimension four, where a great deal is known classically on the complex side and where differential topology is so rich. In this context it is quite natural to consider the question when a symplectic submanifold in $\mathbb{CP}^2$ is isotopic to a holomorphic curve, which is the main topic of these lectures. An even stronger motivation is explained in the contribution by Auroux and Smith to this volume [AuSm], where they discuss how closely related the classification of symplectic manifolds is to the classification of symplectic surfaces in $\mathbb{CP}^2$. (These surfaces can have classical singularities, that is, nodes and cusps.)

The purpose of this section is to give a slightly different view on the relation between complex and symplectic geometry via Lefschetz fibrations. We will see that a certain class of Lefschetz fibrations called hyperelliptic arises as two-fold covers of rational ruled surfaces, and that our isotopy theorem for symplectic submanifolds of S^2-bundles over S^2 gives a classification of a subclass of such Lefschetz fibrations. This point of view also has an interpretation via representations of the braid group, a topic of independent interest.

6.1 Monodromy Representations – Hurwitz Equivalence

A *symplectic Lefschetz fibration* of an oriented four-manifold (M, ω) is a proper differentiable surjection $q : M \to S^2$ with only finitely many critical points in pairwise disjoint fibers with local model $\mathbb{C}^2 \to \mathbb{C}$, $(z, w) \mapsto zw$. Here z, w and the coordinate on S^2 are complex-valued and compatible with the orientations. With the famous exception of certain genus-one fiber bundles without sections, for example a Hopf-surface $S^3 \times S^1 \to S^2$, see [McSa1], Expl. 6.5 for a discussion, M then has a distinguished deformation class of symplectic structures characterized by the property that each fiber is symplectic [GoSt]. Note that if ω has this property then this is also the case for $q^*\omega_{S^2} + \varepsilon\omega$ for any $\varepsilon > 0$. In particular, this deformation class of symplectic structures has $q^*\omega_{S^2}$ in its closure.

For a general discussion of Lefschetz fibrations we refer to the lectures of Auroux and Smith. From this discussion recall that any symplectic four-manifold (M, ω) with $[\omega] \in H^2(M, \mathbb{Q})$ arises as total space of a Lefschetz fibration after blowing up finitely many disjoint points and with fibers Poincaré dual to the pull-back of $k[\omega]$ for $k \gg 0$ [Do]. The fibration structure is unique up to isotopy for each $k \gg 0$. In other words, for each ray $\mathbb{Q}_{>0}\omega$ of symplectic structures with rational cohomology one can associate a sequence of Lefschetz fibrations, which is unique up to taking subsequences. However, the sequence depends heavily on the choice of $[\omega]$, and it also seems difficult to control the effect of increasing k on the fibration structure. Conversely, it is also difficult to characterize Lefschetz fibrations arising in this way. Necessary conditions are the existence of sections with self-intersection number -1 and irreducibility of all singular fibers for $k \gg 0$ [Sm2], but these conditions are certainly not sufficient. So at the moment the use of this point of view for an effective classification of symplectic four-manifolds is limited.

On the other hand, in algebraic geometry Lefschetz fibrations have been especially useful for low degrees, that is, for low genus of the fibers. This is the point of view taken up in this section symplectically.

6.2 Hyperelliptic Lefschetz Fibrations

Auroux and Smith explain in their lectures that a symplectic Lefschetz fibration $\pi : M \to S^2$ with singular fibers over $s_1, \ldots, s_\mu \in S^2$ is characterized by its monodromy representation into the mapping class group

$$\rho : \pi_1(S^2 \setminus \{s_1, \ldots, s_\mu\}, s_0) \longrightarrow \pi_0 \operatorname{Diff}^+(\pi^{-1}(\Sigma)).$$

Here $s_0 \in S^2 \setminus \{s_1, \ldots, s_\mu\}$ is some fixed non-critical point in the base and $\Sigma = \pi^{-1}(s_0)$. For each loop running around only one critical point once the monodromy is a Dehn twist. There is even a one-to-one correspondence between isomorphism classes of Lefschetz fibrations with μ singular fibers along with a diffeomorphism $\pi^{-1}(s_0) \simeq \Sigma$, and such representations [Ka].

If one chooses a generating set of μ loops $\gamma_1, \ldots, \gamma_\mu$ intersecting only in s_0 and each encircling one of the critical points then ρ is uniquely determined by the tuple $(\tau_1, \ldots, \tau_\mu)$ of μ Dehn twists $\tau_i = \rho(\gamma_i)$ of Σ. Conversely, any such tuple with the property $\prod_i \tau_i = 1$ arises from such a representation. This gives a description of symplectic Lefschetz fibrations up to isomorphism by finite algebraic data, namely by the word $\tau_1 \ldots \tau_\mu$ of Dehn twists in the genus-g mapping class group $\mathrm{MC}_g \simeq \pi_0 \operatorname{Diff}^+(\pi^{-1}(\Sigma))$. This description is unique up to an overall conjugation (coming from the choice of isomorphism $\mathrm{MC}_g \simeq \pi_0 \operatorname{Diff}^+(\pi^{-1}(\Sigma))$) and up to so-called *Hurwitz equivalence*. The latter accounts for the choice of $\gamma_1, \ldots, \gamma_\mu$. It is generated by transformations of the form

$$\tau_1 \ldots \tau_r \tau_{r+1} \ldots \tau_\mu \longrightarrow \tau_1 \ldots \tau_{r+1} (\tau_r)_{\tau_{r+1}} \ldots \tau_\mu,$$

(Hurwitz move) where $(\tau_r)_{\tau_{r+1}} = \tau_{r+1}^{-1} \tau_r \tau_{r+1}$. Note that the set of Dehn twists is stable under conjugation, and hence the word on the right-hand side still consists of Dehn twists.

We now want to look at a special class of Lefschetz fibrations called *hyperelliptic*. By definition their monodromy representations take values in the hyperelliptic mapping class group $\mathrm{HMC}_g \subset \mathrm{MC}_g$. Recall that a hyperelliptic curve of genus g is an algebraic curve that admits a two-fold cover $\kappa : \Sigma \to \mathbb{CP}^1$ branched in $2g+2$ points. The hyperelliptic mapping class group is the subgroup of MC_g of isotopy classes of diffeomorphisms of Σ respecting κ. So each $\sigma \in \mathrm{HMC}_g$ induces a diffeomorphism of S^2 fixing the branch set, well-defined up to isotopy. This defines a homomorphism

$$\mathrm{HMC}_g \longrightarrow \mathrm{MC}(S^2, 2g+2)$$

to the mapping class group of S^2 marked with a set of $2g+2$ points. The kernel is generated by the hyperelliptic involution that swaps the two points in the fibers of κ. For genus two it happens that $\mathrm{HMC}_g = \mathrm{MC}_g$, otherwise the inclusion $\mathrm{HMC}_g \subset \mathrm{MC}_g$ is strict. For all this a good reference is the book [Bi].

Of course, given a closed surface Σ of genus g there are many involutions with $2g+2$ fixed points exhibiting Σ as a two-fold cover of S^2, and these give different copies of HMC_g in MC_g. The definition of hyperelliptic Lefschetz fibrations requires that $\operatorname{im}\rho \subset \mathrm{HMC}_g$ for one such choice of involution.

One method to produce a hyperelliptic Lefschetz fibration $M \to S^2$ is as composition of a two-fold cover with an S^2-bundles $p : P \to S^2$, with branch locus a so-called *Hurwitz curve* $B \subset P$ of degree $2g+2$ over S^2. Recall from the lectures of Auroux and Smith that a smooth submanifold $B \subset P$ is a Hurwitz curve if near the critical points of the composition $B \to P \to S^2$ there are local complex coordinates (z, w) on P such that $p(z, w) = z$ and B is locally given by $w = z^2$. Then $B \to S^2$ is a branched cover of degree $2g+2$ with only simple branch points. The critical points of this projection produce singular fibers as follows. In the local coordinates (z, w) introduced above, M is locally the solution set of $v^2 - w^2 + z$. This four-dimensional manifold has complex coordinates $v' = v - w$, $w' = v + w$ because one can eliminate z. The projection to the z-coordinate then has the standard description $(v', w') \mapsto v'w'$ of a Lefschetz fibration near a singular point.

This construction yields hyperelliptic Lefschetz fibrations with only irreducible singular fibers. But there are relatively minimal hyperelliptic Lefschetz fibrations with reducible fibers. (To be relatively minimal means that we have not introduced reducible fibers artificially by blowing up the total space. Technically this leads to spheres contained in fibers with self-intersection number -1, and "relatively minimal" means there are no such spheres.) After a slight perturbation of the Lefschetz fibration any fiber contains at most one critical point (we took this as a condition in the definition of Lefschetz fibrations), and then any reducible fiber is a union of two surfaces, of genera h and $g-h$ for some $0 < h < [g/2]$.

Now how can one construct hyperelliptic Lefschetz fibrations with reducible fibers? Here is the construction of the model for a neighborhood of a singular fiber with irreducible components of genera h and $g-h$. Let $\Delta \subset \mathbb{C}$ be the unit disk, and consider in $\Delta \times \mathbb{CP}^1$ the holomorphic curve $\bar{B}$ given by

$$(w-\alpha_1)\cdot\ldots\cdot(w-\alpha_{2(g-h)+1})\cdot(w-\beta_1 z^2)\cdot\ldots\cdot(w-\beta_{2h+1}z^2) = 0. \quad (6.1)$$

Here z is the coordinate on Δ, w is the affine coordinate on $\mathbb{C} \subset \mathbb{CP}^1$, and both the α_i and the β_j are pairwise disjoint and non-zero. So $\bar{B}$ consists of $2g+2$ irreducible components, each projecting isomorphically to Δ. There is one singular point at $(0,0)$, a tacnode with tangent line $w = 0$ and contained in $2h+1$ branches of $\bar{B}$. Figure 6.1, left, depicts the case $g = 2$, $h = 1$. Now let $\rho_1 : P_1 \to \Delta \times \mathbb{CP}^1$ be the blow up at $(0,0)$. This replaces the fiber F over $z = 0$ by two (-1)-spheres, the strict transform F_1 of F and another one E contracted under ρ_1. By taking the strict transform B_1 of $\bar{B}$ (the closure

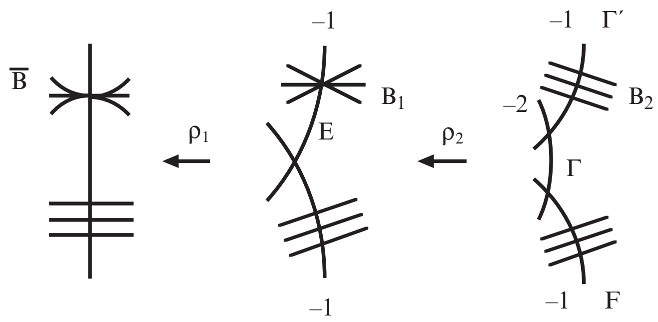

Fig. 6.1. Producing reducible fibers ($g = 2, h = 1$)

of $\rho_1^{-1}(\bar{B}) \setminus E$) the tacnode of $\bar{B}$ transforms to an ordinary singular point of multiplicity $2h + 1$ lying on $E \setminus F$. Another blow up $\rho_2 : \tilde{P} \to P_1$ in this singular point desingularizes B_1. The strict transform Γ of E is a sphere of self-intersection -2. So the fiber over 0 of $\tilde{P} \to \Delta$ is a union of two (-1)-spheres and one (-2)-sphere intersecting as depicted in Fig. 6.1 on the right. Denote by B_2 the strict transform of B_1 under the second blow-up.

It is not hard to check that, viewed as divisor, $B_2 + \Gamma$ is divisible by 2 up to rational equivalence, just as $\bar{B}$. Hence there exists a holomorphic line bundle L on $\tilde{P}$ with a section s with zero locus $B_2 + \Gamma$. The solution set to $u^2 = s$ with $u \in L$ is a two-fold cover of $\tilde{P}$ branched over $B_2 + \Gamma$. This is an instance of the standard construction of cyclic branched covers with given branch locus.

The projection $\tilde{M} \to \Delta$ is a local model for a genus-g Lefschetz fibration. The singular fiber over $0 \in \Delta$ is a chain of three components, of genera $g - h$, 0 and h, respectively. Since Γ is a (-2)-sphere in the branch locus the rational component $\tilde{\kappa}^{-1}(\Gamma)$ has self-intersection number -1 and hence can be contracted. The result is a manifold M with the projection to Δ the desired local model of a relatively minimal Lefschetz fibration with a reducible fiber with components of genera h and $g - h$. It is obviously hyperelliptic by construction.

One interesting remark is that the covering involution of $\tilde{M}$ descends to M. This action has one fixed point at the unique critical point of the fibration $M \to \Delta$. In local holomorphic coordinates it looks like $(u, v) \mapsto (-u, -v)$, and the ring of invariant holomorphic functions is generated by $x = u^2$, $y = v^2$, $z = uv$ (this can be chosen to agree with the z from before). The generators fulfill the relation $z^2 = xy$, so the quotient is isomorphic to the two-fold cover of $\mathbb{C}^2$ branched over the coordinate axes. This is called the A_1-singularity, and we have just verified the well-known fact that this singularity is what one obtains locally by contracting the (-2)-curve Γ on $\tilde{P}$. Alternatively, and maybe more appropriately, one should view the singular space P obtained by this contraction as the orbifold $M/(\mathbb{Z}/2)$.

6.3 Braid Monodromy and the Structure of Hyperelliptic Lefschetz Fibrations

We have now found a way to construct a hyperelliptic Lefschetz fibration starting from a certain branch surface $\bar{B}$ in S^2-bundles $\bar{p} : \bar{P} \to S^2$. This surface may have tacnodal singularities with non-vertical tangent line, and these account for reducible singular fibers in the resulting Lefschetz fibration. Otherwise the projection $\bar{B} \to S^2$ is a simply branched covering with the simple branch points leading to irreducible singular fibers.

Theorem 6.1 *[SiTi1] Any hyperelliptic Lefschetz fibration arises in this way.*

The case of genus-two has also been proved in [Sm1], and a slightly more topological proof is contained in [Fu]. The key ingredients of our proof are to describe B by its monodromy in the braid group $B(S^2, 2g+2)$ of the sphere on $2g+2$ strands, and to observe that $B(S^2, 2g+2)$ is also a $\mathbb{Z}/2$-extension of $\mathrm{MC}(S^2, 2g+2)$, just as HMC_g. While nevertheless $B(S^2, 2g+2)$ and HMC_g are not isomorphic, there is a one-to-one correspondence between the set of half-twists in the braid group on one side and the set of Dehn-twists in HMC_g on the other side. This correspondence identifies the two kinds of monodromy representations. We now give some more details.

For a topological space X the *braid group on d strands* $B(X, d)$ can be defined as the fundamental group of the *configuration space*

$$X^{[d]} := (X \times \ldots \times X \setminus \Delta)/S_d.$$

Here $\Delta = \{(x_1, \ldots, x_d) \in X^d \mid \exists i \neq j, x_i = x_j\}$ is the generalized diagonal and S^d acts by permutation of the components. So a braid takes a number of fixed points on X, moves them with $t \in [0,1]$ such that at no time t two points coincide, and such that at $t = 1$ we end up with a permutation of the tuple we started with.

If X is a two-dimensional oriented manifold and $\gamma : [0,1] \to X$ is an embedded path connecting two different P_j, P_k and disjoint from $\{P_1, \ldots, P_d\}$ otherwise, there is a braid exchanging P_j and P_k as in the following Fig. 6.2. Any such braid is called a *half-twist*. Note that the set of half-twists is invariant under the action of the group of homeomorphisms on X fixing $\{P_1, \ldots, P_k\}$.

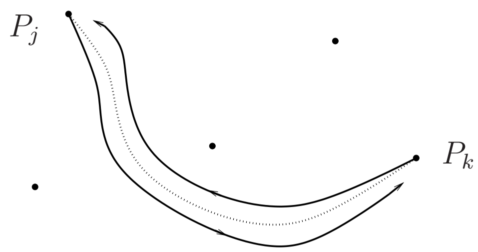

Fig. 6.2. A half-twist

The classical *Artin braid group* $B_d := B(\mathbb{R}^2, d)$ can be explicitly described as follows. Take as points $P_j := e^{2\pi\sqrt{-1}j/d}$ the dth roots of unity. Then define σ_j as the half-twist associated to the line segment connecting P_j and P_{j+1} for $1 \le j < d$. In the following we take the index j modulo d. Then B_d is the group generated by the σ_j subject to the famous braid relations

$$\sigma_j\sigma_k = \sigma_k\sigma_j \quad \text{for } |j-k| \ge 2,$$
$$\sigma_j\sigma_{j+1}\sigma_j = \sigma_j\sigma_{j+1}\sigma_j \quad \text{for all } j.$$

It is important to note that there are infinitely many different sets of generators such as the σ_j, one for each self-intersection free path running through all the P_j. This is responsible for some of the complications when dealing with the braid group.

Now given one of our branch surfaces $p : \bar{B} \subset \bar{P}$ with critical set $\{s_1, \ldots, s_\mu\} \subset S^2$, a closed path γ in $S^2 \setminus \{s_1, \ldots, s_\mu\}$ defines a braid in S^2 by trivializing $\bar{P}$ over γ and by interpreting the pull-back of B over γ as the strands of the braid. This defines the monodromy representation

$$\rho' : \pi_1(S^2 \setminus \{s_1, \ldots, s_\mu\}, s_0) \longrightarrow B(S^2, 2g+2),$$

which characterizes $\bar{B}$ uniquely up to isotopy. Note that the braid group on the right is really the braid group of the fiber $\bar{p}^{-1}(s_0)$ with the point set $\bar{B} \cap \bar{p}^{-1}(s_0)$.

The following possibilities arise for the monodromy around a loop γ enclosing only one of the s_i. If $\bar{B}$ is smooth over s_i then $\bar{B} \to S^2$ has a simple branch point over s_i. In this case $\rho'(\gamma)$ is the half-twist swapping the two branches of B coming together at the branch point. Otherwise $\bar{B}$ has a tacnode mapping to s_i. Then the local standard form (6.1) gives the following description of $\rho'(\gamma)$. There is an embedded loop $S^1 \hookrightarrow S^2 = \bar{p}^{-1}(s_0)$ passing through a subset $P_{i_1}, \ldots, P_{i_{2h+1}}$ of $\bar{B} \cap \bar{p}^{-1}(s_0)$ and *not enclosing any other* P_j. Now $\rho'(\gamma)$ is given by a full counterclockwise rotation of these points along the loop, and by the identity on all other points.

The point now is that in the hyperelliptic case any Dehn twist arises as a two-fold cover of a distinguished braid of the described form once a choice of a north pole $\infty \in S^2$ has been made. In fact, the three groups HMC_g, $\mathrm{MC}(S^2, 2g+2)$ and $B(S^2, 2g+2)$ all have $2g+2$ generators $\sigma_1, \ldots, \sigma_{2g+1}$ fulfilling the Artin braid relations, and in addition:

1. $\mathrm{MC}(S^2, 2g+2)$: $I = 1$, $T = 1$
2. HMC_g: $I^2 = 1$, $T = 1$, and I is central ($I\sigma_i = \sigma_i I$ for all i)
3. $B(S^2, 2g+2)$: $I = 1$, and this implies $T^2 = 1$ and T central

where

$$I = \sigma_1 \ldots \sigma_{2g+1}\sigma_{2g+1} \ldots \sigma_1, \quad T = (\sigma_1 \ldots \sigma_{2g+1})^{2g+2}$$

Geometrically, $I \in \mathrm{HMC}_g$ is the hyperelliptic involution; it induces the trivial element in $\mathrm{MC}(S^2, 2g+2)$, and it can not be produced from braids via two-fold covers. On the other hand, T is the full-twist along a loop passing through all the points; its square is the trivial braid as one sees by "pulling the bundle of strands across $\infty \in S^2$," and again it induces the trivial element in both $\mathrm{MC}(S^2, 2g+2)$ and in HMC_g.

Thus given a hyperelliptic Lefschetz fibration one can produce a branch surface $\bar{B}$ uniquely up to isotopy by going via braid monodromy. There are two minor global issues with this, one being the extension of $\bar{B}$ over one point at infinity, the other homological two-divisibility of $\bar{B}$. The latter follows by an a priori computation of the possible numbers of singular fibers of each type, of the critical points of $\bar{B}$ and their relation to the homology class of $\bar{B}$. The first issue can be resolved by closing up $\bar{B}$ either in the trivial or in the non-trivial S^2-bundle over S^2.

6.4 Symplectic Noether–Horikawa Surfaces

A simple, but important observation is that any of the branch surfaces $\tilde{B} \subset \tilde{P}$ are symplectic with respect to $\omega_{\tilde{P}} + k\tilde{p}^*\omega_{S^2}$, for $k \gg 0$. Here $\omega_{\tilde{P}}$ and ω_{S^2} are any Kähler structures on $\tilde{P}$ and on S^2, and $\tilde{P} \to P$ resolves the tacnodes of $\bar{B} \subset P$. Thus the question whether a hyperelliptic Lefschetz fibrations is isomorphic (as a Lefschetz fibration) to a holomorphic one, is equivalent to asking if $\tilde{B}$ can be deformed to a holomorphic curve within the class of branch surfaces in $\tilde{P}$.

For the understanding of symplectic Lefschetz fibrations this point of view is certainly limited for the following two reasons. First, it is not true that any hyperelliptic Lefschetz fibration is isomorphic to a holomorphic one. For example, [OzSt] shows that fiber sums of two copies of a certain genus-2 Lefschetz fibration produce infinitely many pairwise non-homeomorphic symplectic 4-manifolds of which only finitely many can be realized as complex manifolds. And second, the general classification of holomorphic branch curves up to isotopy, hence of hyperelliptic holomorphic Lefschetz fibrations, is complicated, see e.g. [Ch].

On the other hand, the complex geometry becomes regular in a certain stable range, when the deformation theory of the branch curve is always unobstructed. This is the case when the total number μ of singular fibers is much larger than the number t of reducible singular fibers. In the genus-2 case, the discussions in [Ch] suggest $\mu > 18t$ as this stable range. The example of [OzSt] has $\mu = 4t$. So in our opinion, the holomorphic point of view is appropriate for a classification in a certain stable range.

Conjecture 6.2 *For any g there exists an integer N_g such that any hyperelliptic symplectic genus-g Lefschetz fibration with μ singular fibers of which t are reducible, and such that*

$$\mu > N_g t,$$

is isomorphic to a holomorphic one.

The holomorphic classification in the stable range should in turn be simple. We expect that there is only a very small number of deformation classes of holomorphic genus-g Lefschetz fibrations with fixed numbers and types of singular fibers (given by the genera of its irreducible components), distinguished by topological invariants of the total space such as Euler characteristic and signature.

The conjecture in particular says that any hyperelliptic Lefschetz fibration with reducible fibers is holomorphic. By the discussion above this is equivalent to saying that each smooth branch surface (no tacnodes) in a rational ruled surface is isotopic as branch surface to a holomorphic curve. The main theorem of these lectures Theorem 9.1 says that this is true for connected $\bar{B}$ provided $\deg(\bar{B} \to S^2) \leq 7$. If $\bar{B}$ is disconnected it is either a product, or it has precisely two components and one of them is a section. In the disconnected case the monodromy representation does not act transitively on the set of strands, while this is true in the case with connected $\bar{B}$. (With hindsight we will even see that in the connected case the monodromy representation is surjective.) In any case, we say the case with connected $\bar{B}$ *has transitive monodromy.* Then we have the following:

Theorem 6.3 *[SiTi3] A Lefschetz fibration with only irreducible singular fibers of genus two, or of genus one with a section, and with transitive monodromy is isomorphic to a holomorphic Lefschetz fibration.*

By the standard technique of degeneration to nodal curves (see [Te] for a survey) it is not hard to compute the braid monodromy for smooth algebraic branch curves in $\mathbb{CP}^1$-bundles over $\mathbb{CP}^1$, i.e. in Hirzebruch surfaces $\mathbb{F}_k$, see [Ch].

Proposition 6.4 *The braid monodromy word of a smooth algebraic curve $\bar{B} \subset \mathbb{F}_k$ of degree d and with μ simple critical points is Hurwitz-equivalent to one of the following:*

1. $(\sigma_1 \ldots \sigma_{d-1}\sigma_{d-1} \ldots \sigma_1)^{\frac{\mu}{2d-2}}$ *($\bar{B}$ connected and k even)*
2. $(\sigma_1 \ldots \sigma_{d-1}\sigma_{d-1} \ldots \sigma_1)^{\frac{\mu-d(d-1)}{2d-2}}(\sigma_1 \ldots \sigma_{d-1})^d$ *($\bar{B}$ connected and k odd)*
3. $(\sigma_1 \ldots \sigma_{d-2})^{\frac{\mu}{d-2}}$ *($\bar{B}$ disconnected; $k = 2d$)*

Taken together this gives a complete classification of symplectic Lefschetz fibrations with only irreducible singular fibers and transitive monodromy in genus two.

In the non-hyperelliptic case it is not clear what an analogue of Conjecture 6.2 should be. Any symplectic manifold arises as total space of a symplectic Lefschetz fibration without reducible singular fibers, after blowing up finitely many points [Do, Sm2]. Thus the absence of reducible singular fibers alone certainly does not suffice as obstruction to holomorphicity.

By purely braid-theoretic methods Auroux very recently achieved the following beautiful stable classification result:

Theorem 6.5 *[Au2] For each g there exists a universal genus-g Lefschetz fibration π_g^0 with the following property: Given two genus-g Lefschetz fibrations $\pi_i : M_i \to S^2$, $i = 1, 2$, with the same numbers of reducible fibers then the fiber connected sums of π_i with sufficiently many copies of π_g^0 are isomorphic Lefschetz fibrations, provided*

1. *there exist sections $\Sigma_i \subset M_i$ with the same self-intersection numbers, and*
2. *M_1 and M_2 have the same Euler number.*

This refines a previous, slightly simpler result by the same author for the genus 2 case [Au1], building on Theorem 6.3.

7 The $\mathscr{C}^0$-Compactness Theorem for Pseudo-Holomorphic Curves

In this section we discuss a compactness theorem for J-holomorphic maps in the case that J is only assumed to be continuous. Such a compactness theorem was first due to Gromov [Gv] and was further discussed by Parker-Wolfson, Pansu, Ye and Ruan-Tian [PrWo, Pn, Ye, RuTi]. For the reader's convenience, we will present a proof of this compactness theorem and emphasize that it depends only on the $\mathscr{C}^0$-norm of the involved almost complex structures. Our proof basically follows [Ti], where further smoothness was discussed. We should point out that the dependence on a weaker norm for the almost complex structures is crucial in our study of the symplectic isotopy problem.

7.1 Statement of Theorem and Conventions

First we note that in this section by a J-holomorphic map we mean a Hölder continuous ($\mathscr{C}^{0,\alpha}$, $0 < \alpha < 1$) map from a Riemann surface Σ into M whose derivative is L^2-bounded and which satisfies the J-holomorphicity equation in the distributional sense. Explicitly, the last phrase says that for any smooth vector field X on M with compact support in a neighborhood of $f(\Sigma)$ and any smooth vector field v with compact support in Σ,

$$\int_\Sigma g\big(X, Df(v) + J \cdot Df(j_\Sigma(v))\big)dz = 0,$$

where j_Σ denotes the conformal structure of Σ. This coincides with the standard J-holomorphicity equation whenever f is smooth. By our assumption on f, any L^2-section of f^*T_M over Σ can be approximated in the L^2-topology by the pull-back of a locally constant vector field on M, so it follows that the above equation for f also holds when X is replaced by any L^2-section of f^*T_M.

As before, we denote by (M, ω) a compact symplectic manifold and by g a fixed Riemannian metric. Let J_i be a sequence of *continuous* almost complex structures on M converging to J_∞ in the $\mathscr{C}^0$-topology and uniformly tamed in the following sense: There exists a constant $c > 0$ such that for any $X \in T_M$ and any i

$$cg(X,X) \le \omega(X, J_i X) \le c^{-1} g(X,X). \tag{7.1}$$

Here is the main result of this section.

Theorem 7.1 *Let (M,ω) and g be as above. Assume that Σ_i is a sequence of Riemann surfaces of fixed genus and $f_i : \Sigma_i \to M$ are J_i-holomorphic with uniformly bounded homology classes $f_{i*}[\Sigma_i] \in H_2(M,\mathbb{Z})$.*

Then there is a connected singular Riemann surface Σ_∞ with finitely many irreducible components $\Sigma_{\infty,a}$, and smooth maps $\phi_i : \Sigma_i \to \Sigma_\infty$ such that the following holds:

1. *ϕ_i is invertible on the pre-image of the regular part of Σ_∞*
2. *A subsequence of $f_i \circ \phi_i^{-1}$ converges to a J_∞-holomorphic map f_∞ on the regular part of Σ_∞ in the $\mathscr{C}^0$-topology*
3. *Each $f_\infty|_{\Sigma_{\infty,a}}$ extends to a J_∞-holomorphic map from $\Sigma_{\infty,a}$ into M, and the homology classes $f_{i*}[\Sigma_i]$ converge to $f_{\infty *}[\Sigma_\infty]$ in $H_2(M,\mathbb{Z})$*

The rest of this section is devoted to the proof. We start with the monotonicity formula for pseudo-holomorphic maps.

7.2 The Monotonicity Formula for Pseudo-Holomorphic Maps

For notational simplicity we will denote by J one of the almost complex structures J_i or J_∞. Let I and g_{stan} denote the standard almost complex structure and standard flat Riemannian metric on $\mathbb{R}^{2n}$, respectively. By (7.1), for any $\eta > 0$ there is a uniform δ_η such that for any geodesic ball $B_R(p)$ ($p \in M$ and $R \le \delta_\eta$), there is a $\mathscr{C}^1$-diffeomorphism $\phi : B_R(p) \to B_R(0) \subset \mathbb{R}^{2n}$ such that

$$\|J - \phi^* I\|_{\mathscr{C}^0(B_R(p))} \le \eta, \tag{7.2}$$

where norms are taken with respect to g. We may further assume that $\|g - \phi^* g_{\mathrm{stan}}\|_{\mathscr{C}^0} \le C$ for some uniform constant C.

Denote by Δ_r the disk in $\mathbb{C}$ with center at the origin and radius r, and $\Delta = \Delta_1$. Throughout the proof c will be a uniform constant whose actual value may vary.

Lemma 7.2 *There is an $\epsilon > 0$ such that for any $\alpha \in (0,1)$ and any J-holomorphic map $f : \Delta_r \to M$ ($r > 0$) with $\int_{\Delta_r} |Df|_g^2 dz \le \epsilon$, we have*

$$\int_{B_{r'}(q)} |Df|_g^2 dz \le c_\alpha r'^{2\alpha}, \quad \forall q \in \Delta_{r/2} \ \text{ and } \ r' \le r/4, \tag{7.3}$$

Moreover, we have

$$\operatorname{diam} f(\Delta_{r'/2}) \le c_\alpha \sqrt{\epsilon} \left(\frac{r'}{r}\right)^\alpha . \tag{7.4}$$

Here c_α is a uniform constant which may depend on α.

Proof. Since all estimates are scaling-invariant, we may assume $r = 1$. Let $\eta > 0$ be a sufficiently small positive number and let δ_η be given as in (7.2). For simplicity, we will write δ for δ_η and identify $B_R(p)$, $p = f(0)$, with $B_R(0) \subset \mathbb{R}^{2n}$ by the diffeomorphism ϕ in (7.2). Because g and $\phi^* g_{\mathrm{stan}}$ are uniformly equivalent we may then also replace $|\,.\,|_g$ in the statement by the standard norm $|\,.\,|$ in $\mathbb{R}^n$. Choose any $\rho_0 \le \frac{1}{2}$ so that $f(B_{2\rho_0}(0)) \subset B_\delta(p)$. As input to Morrey's Lemma (see e.g. [GiTr], Lemma 12.2) we now derive a growth condition for the local L^2-norm of Df at a fixed $y \in B_{\rho_0}(0)$, see (7.6).

In polar coordinates (r, θ) centered at y, the Cauchy–Riemann equation becomes

$$\frac{\partial f}{\partial r} + \frac{1}{r} J \frac{\partial f}{\partial \theta} = 0,$$

and $|Df|^2 = |\partial_r f|^2 + |r^{-1}\partial_\theta f|^2$. In particular, both $|\partial_r f|^2$ and $|r^{-1}\partial_\theta f|^2$ are close to $\frac{1}{2}|Df|^2$:

$$\begin{aligned} \big|2|\partial_r f|^2 - |Df|^2\big| &= \big||I\partial_r f|^2 - |r^{-1}\partial_\theta f|^2\big| \\ &\le \big|(J-I)\partial_r f\big|^2 + \big||J\partial_r f|^2 - |r^{-1}\partial_\theta f|^2\big| \;\le\; \eta |Df|^2, \\ \big|2|r^{-1}\partial_\theta f|^2 - |Df|^2\big| &\le \eta |Df|^2. \end{aligned} \tag{7.5}$$

We also obtain the pointwise estimate

$$\begin{aligned} 0 = \big|\partial_r f + r^{-1}J\partial_\theta f\big|^2 \;&=\; \big|\partial_r f + r^{-1}I\partial_\theta f + r^{-1}(J-I)\partial_\theta f\big|^2 \\ &= \big|\partial_r f + r^{-1}I\partial_\theta f\big|^2 + 2\big\langle \partial_r f + r^{-1}I\partial_\theta f, r^{-1}(J-I)\partial_\theta f\big\rangle + \big|r^{-1}(J-I)\partial_\theta f\big|^2 \\ &\ge \big|\partial_r f\big|^2 + 2\big\langle \partial_r f, r^{-1}I\partial_\theta f\big\rangle + \big|r^{-1}I\partial_\theta f\big|^2 - 2\eta\cdot\big(|\partial_r f| + |r^{-1}\partial_\theta f|\big)|r^{-1}\partial_\theta f| \\ &\ge (1-4\eta)\big|Df\big|^2 + 2\big\langle \partial_r f, r^{-1}I\partial_\theta f\big\rangle. \end{aligned}$$

Then integrating by parts twice, we have for $\rho \le \rho_0$ and any constant vector $\lambda \in \mathbb{R}^n$,

$$\begin{aligned} 0 &\ge (1-4\eta)\int_{B_\rho(y)} \big|Df\big|^2 r\,dr\,d\theta + 2\int_0^\rho \int_0^{2\pi} \Big\langle \frac{\partial f}{\partial r}, I\Big(\frac{\partial f}{\partial \theta}\Big)\Big\rangle dr\,d\theta \\ &= (1-4\eta)\int_{B_\rho(y)} \big|Df\big|^2 r\,dr\,d\theta + 2\int_0^{2\pi} \Big\langle f-\lambda, I\Big(\frac{\partial f}{\partial \theta}\Big)\Big\rangle(\rho,\theta)\,d\theta \\ &\quad -2\int_0^\rho \int_0^{2\pi} \Big\langle f-\lambda, I\Big(\frac{\partial^2 f}{\partial r\partial \theta}\Big)\Big\rangle dr\,d\theta \\ &= (1-4\eta)\int_{B_\rho(y)} \big|Df\big|^2 r\,dr\,d\theta + 2\int_0^{2\pi} \Big\langle f-\lambda, I\Big(\frac{\partial f}{\partial \theta}\Big)\Big\rangle(\rho,\theta)\,d\theta \\ &\quad +2\int_0^\rho \int_0^{2\pi} \Big\langle \frac{\partial f}{\partial \theta}, I\Big(\frac{\partial f}{\partial r}\Big)\Big\rangle dr\,d\theta. \end{aligned}$$

The last term gives another $(1-2\eta)\int_{B_\rho(y)} \big|Df\big|^2 r\,dr\,d\theta$ by the following:

$$\begin{aligned}2r^{-1}\big\langle\partial_\theta f, I\partial_r f\big\rangle &= \big\langle r^{-1}\partial_\theta f, I\partial_r f\big\rangle - \big\langle r^{-1}I\partial_\theta f, \partial_r f\big\rangle\\ &\geq -2\big|Df\big|^2\cdot\eta + \big\langle r^{-1}\partial_\theta f, J\partial_r f\big\rangle + \big\langle r^{-1}J\partial_\theta f, \partial_r f\big\rangle\\ &= -2\big|Df\big|^2\cdot\eta + \big|r^{-1}\partial_\theta f\big|^2 + \big|\partial_r f\big|^2 \;=\; (1-2\eta)\big|Df\big|^2.\end{aligned}$$

It follows that

$$\begin{aligned}(1-3\eta)\int_{B_\rho(y)}\big|Df\big|^2 rdrd\theta &\leq -\int_0^{2\pi}\langle f-\lambda, I(\frac{\partial f}{\partial\theta})\rangle d\theta.\\ &\leq \Big(\int_0^{2\pi}|f-\lambda|^2 d\theta\Big)^{1/2}\cdot\Big(\int_0^{2\pi}\Big|\frac{\partial f}{\partial\theta}\Big|^2 d\theta\Big)^{1/2}.\end{aligned}$$

Now choose

$$\lambda = \frac{1}{2\pi}\int_0^{2\pi} f d\theta.$$

Then by the Poincaré inequality on the unit circle, we have

$$\int_0^{2\pi}|f-\lambda|^2 d\theta \leq \int_0^{2\pi}\Big|\frac{\partial f}{\partial\theta}\Big|^2 d\theta \leq \rho^2\int_0^{2\pi}\Big|r^{-1}\frac{\partial f}{\partial\theta}\Big|^2 d\theta.$$

Moreover, $|r^{-1}\partial_\theta f|^2 \leq \frac{1+\eta}{2}|Df|^2$ by (7.5) and $\frac{1-3\eta}{1+\eta} \geq 1-4\eta$. Plugging all this into the previous inequality gives

$$(1-4\eta)\int_{B_\rho(y)}|Df|^2 rdrd\theta \leq \frac{\rho^2}{2}\int_0^{2\pi}|Df|^2 d\theta.$$

But

$$\frac{\partial}{\partial\rho}\int_{B_\rho(y)}|Df|^2 rdrd\theta = \rho\int_0^{2\pi}|Df|^2 d\theta,$$

so the above is the same as

$$2(1-4\eta)\int_{B_\rho(y)}|Df|^2 rdrd\theta \leq \rho\frac{\partial}{\partial\rho}\int_{B_\rho(y)}|Df|^2 rdrd\theta,$$

that is,

$$\frac{\partial}{\partial\rho}\left(\rho^{-2(1-4\eta)}\int_{B_\rho(y)}|Df|^2 rdrd\theta\right) \geq 0.$$

This implies, for any $\rho<\rho_0$ and $y\in B_{\rho_0}(0)$,

$$\int_{B_\rho(y)}|Df|^2 rdrd\theta \leq c\Big(\frac{\rho}{\rho_0}\Big)^{2(1-4\eta)}\int_\Delta |Df|^2 rdrd\theta, \tag{7.6}$$

where c is a uniform constant. It follows from this and Morrey's lemma that

$$\sup_{x,y\in\Delta_{\rho_0}}\frac{|f(x)-f(y)|}{|x-y|^{1-4\eta}} \leq c_\eta\rho_0^{-1+4\eta}\left(\int_\Delta |Df|^2 rdrd\theta\right)^{1/2},$$

where c_η is some uniform constant depending only on η. In particular, choosing $\eta \le 1/8$, we obtain for $x, y \in \Delta_{\rho_0}$

$$|f(x) - f(y)| \;\le\; c_\eta \Big(\frac{|x-y|}{\rho_0}\Big)^{1/2} \cdot \sqrt{\varepsilon} \;\le\; c_\eta \sqrt{2}\sqrt{\varepsilon}.$$

Thus the diameter of $f(\Delta_{\rho_0})$ is bounded by $c\sqrt{\epsilon}$.

It remains to prove that if ϵ is sufficiently small, then $f(\Delta_{1/2})$ is contained in a ball of radius δ. In fact, we can then set $\rho_0 = 1/2$ above and conclude the desired uniform estimates for $\alpha = 1 - 4\eta$. For any $x \in \Delta$, define

$$t(x) = \sup \big\{ t \in [0, 1/2] \,\big|\, \mathrm{diam}(f(B_{t(1-|x|)}(x))) \le \delta \big\}.$$

If the above claim is false then $t(0) < 1/2$. Let $x_0 \in \Delta_{1/2}$ be such that $t(x_0) = \inf t(x) < 1/2$. Set $a(x) = t(x)(1 - |x|)$. Then for any $x \in B_{a(x_0)}(x_0)$, we have

$$a(x) \ge t(x_0)(1 - |x|) \ge a(x_0) - t(x_0)|x - x_0| > a(x_0) - t(x_0)a(x_0) > \frac{1}{2}a(x_0).$$

This implies $B_{a(x)}(x) \supset B_{a(x_0)/2}(x)$ and thus, from the above diameter estimate, for any $x \in B_{a(x_0)}(x_0)$, we have

$$\mathrm{diam}(f(B_{a(x_0)/2}(x))) \le c\sqrt{\epsilon}.$$

It follows that

$$\mathrm{diam}(f(B_{a(x_0)}(x_0))) \le 2c\sqrt{\epsilon}.$$

Since the constant c here depends only on δ, we get a contradiction if ϵ is sufficiently small. The claim is proved.

7.3 A Removable Singularities Theorem

As an application of the Monotonicity Lemma we derive the following sort of Uhlenbeck removable singularity theorem under the condition that J is only continuous as described above.

Proposition 7.3 *Let (M, ω) and J be as above. If $f : \Delta_{r_0} \backslash \{0\} \to M$ is a J-holomorphic map with $\int_{\Delta_{r_0}} |Df|_g^2 dz < \infty$, then f extends to a Hölder continuous map from Δ into M.*

Proof. Fix any $\alpha \in (0, 1)$. By choosing r_0 smaller, we may assume

$$\int_{\Delta_{r_0}} |Df|_g^2 dz < \epsilon,$$

where ϵ is as in Lemma 7.2.

Let $x, y \in \Delta_{r_0/2}$ with $|y| \le |x|$, say. If $|x - y| \le |x|/2$, then by Lemma 7.2 applied to the restriction of f to $B_{r_0/2}(x) \subset \Delta_{r_0}$, we have

$$d(f(x), f(y)) \leq \mathrm{diam}(f(B_{|x-y|}(x))) \leq 2c_\alpha\sqrt{\epsilon}\left(\frac{|x-y|}{r_0}\right)^\alpha.$$

If $|x-y| > |x|/2 \geq |y|/2$, choose z such that it is collinear to x and has $|z| = |y|$. We can cover $\partial B_{|y|}(0)$ by 12 balls of radius $|y|/2$, so applying Lemma 7.2 at most 12 times, we get

$$d(f(z), f(y)) \leq 12c_\alpha\sqrt{\epsilon}\left(\frac{|y|}{r_0}\right)^\alpha \leq 24c_\alpha\sqrt{\epsilon}\left(\frac{|x-y|}{r_0}\right)^\alpha.$$

Next we can find finitely many balls $B_{|x|/2}(x_0), \ldots, B_{|x|/2^k}(x_k)$ such that $x_0 = x$, $x_k = z$ and $x_{i+1} \in B_{|x_{i/2}|}(x_i)$. Then applying Lemma 7.2, we get

$$d(f(z), f(x)) \leq \left(1 + \tfrac{1}{2} + \ldots + \tfrac{1}{2^{k-1}}\right)c_\alpha\sqrt{\epsilon}\left(\frac{|x|}{r_0}\right)^\alpha \leq 4c_\alpha\sqrt{\epsilon}\left(\frac{|x-y|}{r_0}\right)^\alpha.$$

Hence,

$$d(f(x), f(y)) \leq 28c_\alpha\sqrt{\epsilon}\left(\frac{|x-y|}{r_0}\right)^\alpha.$$

It follows that f extends to a Hölder continuous map from $\Delta_{r_0/2}$ to M.

7.4 Proof of the Theorem

Now we are in position to prove Theorem 7.1.

First we observe: There is a uniform constant c depending only on g and $[\omega](f_{i*}[\Sigma_i])$ such that for any f_i,

$$\int_{\Sigma_i} |Df_i|_g^2 dz \leq c.$$

Therefore, the L^2-norm of Df_i is uniformly bounded.

Next we observe: If ϵ in Lemma 7.2 is sufficiently small, say $c_{1/2}\sqrt{\epsilon} \leq \delta_{1/2}$, then for each i, either f_i is a constant map or

$$\int_{\Sigma_i} |Df_i|_g^2 dz \geq \epsilon > 0.$$

This can be seen as follows: If the above inequality is reversed, Lemma 7.2 implies that the image $f_i(\Sigma_i)$ lies in a Euclidean ball; on a ball the symplectic form ω is exact and so the energy is zero, that is, f_i is constant.

Consider the following class of metrics h_i on the regular part of Σ_i. The metrics h_i have uniformly bounded geometry, namely, for each $p \in \Sigma_i$ there is a local conformal coordinate chart (U, z) of Σ_α containing p such that U is identified with the unit ball Δ and

$$h_i|_U = e^\varphi dz d\bar{z}$$

for some $\varphi(z)$ satisfying:

$$\|\varphi\|_{\mathscr{C}^k(U)} \leq c_k, \quad \text{for any } k > 0,$$

where c_k are uniform constants independent of i. We also require that there are finitely many cylinder-like necks $N_{i,a} \subset \Sigma_i$ $(a = 1, \dots, n_i)$ satisfying:
(1) n_i are uniformly bounded independent of α.
(2) The complement $\Sigma_i \backslash \bigcup_a N_{i,a}$ is covered by finitely many geodesic balls $B_R(p_{i,j})$ $(1 \leq j \leq m_i)$ of h_i in Σ_i, where R and m_i are uniformly bounded.
(3) Each $N_{i,a}$ is diffeomorphic to a cylinder of the form $S^1 \times (\alpha, \beta)$ (α and β may be $\pm\infty$) satisfying: If s, t denote the standard coordinates of $S^1 \times [0, \beta)$ or $S^1 \times (\alpha, 0]$, then

$$h_i|_{N_{i,a}} = e^{\varphi}(ds^2 + dt^2),$$

where φ is a smooth function satisfying uniform bounds as stated above.

We will say that such a h_i is admissible. We will call $\{h_i\}$ uniformly admissible if all h_i are admissible with uniform constants R, c_k, etc.

Admissible metrics always exist on any sequence Σ_i of Riemann surfaces of the same genus. We will start with a fixed sequence of uniformly admissible metrics h_i on Σ_i. We will introduce a new sequence of uniformly admissible metrics $\tilde{h}_i$ on Σ_i such that there is a uniform bound on the gradient of f_i. Once this is done, the theorem follows easily.

We will define $\tilde{h}_i$ by induction.

Set

$$r_i = \inf \left\{ r \;\Big|\; \int_{B_r(x,h_i)} |Df_i|^2_{h_i,g} dz \geq \epsilon \text{ for some } x \in \Sigma_i \right\}.$$

Here $|\cdot|_{h,g}$ denotes the norm induced by g on M and h on the domain. If r_i is uniformly bounded from below, the induction stops and we just take $\tilde{h}_i = h_i$. Then our main theorem follows from Lemma 7.2 and standard convergence theory.

Now assume that r_i tends to zero as i goes to infinity. By going over to a subsequence we may assume $r_i \leq 1/2$ for all i. Let p_i^1 be the point where r_i is attained. Let z be a local complex coordinate on Σ_i centered at p_i^1 and with values containing $2\Delta \subset \mathbb{C}$. Define $h_i^1 = h_i$ outside the region where $|z| < 1$ and

$$h_i^1 = \frac{r_i^{-2}}{\chi_i(r_i^{-2}|z|^2)}\, h_i \quad \text{for } |z| < 1.$$

Here $\chi_i : \mathbb{R} \to \mathbb{R}$ is a cut-off function satisfying: $\chi(t) = 1$ for $t \leq 1$, $\chi_i(t) = t - 1/2$ for $t \in [2, r_i^{-2}]$, and $\chi_i(t) = r_i^{-2}$ for $t \geq r_i^{-2} + 1$; we may also assume that $0 \leq \chi_i'(t) \leq 1$. Clearly, we have $h_i^1 \geq h_i$ and it holds $h_i(z) = r_i^{-2} h_i$ for $|z| \leq r_i^2$ and $h_i(z) = h_i$ for $|z| \geq r_i^{-2} + 1$. It is easy to check that the sequence h_i^1 is uniformly admissible. Moreover, we have

$$\int_{B_1(p_i^1, h_i^1)} |Df_i|^2_{h_i^1,g} dz = \epsilon$$

where $B_1(p_i^1, h_i^1)$ denotes the geodesic ball of radius 1 and centered at p_i with respect to the metric h_i^1.

Next we define

$$r_i^1 = \inf\left\{ r \,\Big|\, \int_{B_r(x,h_i^1)} |Df_i|^2_{h_i^1,g} dz \geq \epsilon \text{ for some } x \in \Sigma_i \right\}.$$

If r_i^1 is uniformly bounded from below, the induction stops and we just take $\tilde{h}_i = h_i^1$. Then our main theorem again follows from Lemma 7.2 and standard convergence theory. Otherwise, by taking a subsequence if necessary, we may assume that $r_i^1 \to 0$ as $i \to \infty$ and $r_i^1 \leq 1/2$ for all i. Let p_i^2 be the point where r_i^1 is attained. Then for i sufficiently large, $p_i^2 \in \Sigma \backslash B_2(p_i^1, h_i^1)$. For simplicity, we assume that this is true for all i. Now we can get h_i^2 by repeating the above construction with h_i replaced by h_i^1. Clearly, h_i^2 coincides with h_i^1 on $B_1(p_i, h_i^1)$, so

$$B_1(p_i^1, h_i^2) = B_1(p_i^1, h_i^1).$$

We also have $B_1(p_i^2, h_i^2) \cap B_1(p_i^1, h_i^1) = \emptyset$ and

$$\int_{B_1(p_i^2,h_i^2)} |Df_i|^2_{h_i^2,g} dz = \epsilon > 0.$$

We can continue this process to construct metrics h_i^L $(L \geq 2)$ and find points p_i^α $(\alpha = 1, \ldots, L)$ such that $B_1(p_i^\alpha, h_i^L) \cap B_1(p_i^\beta, h_i^L) = \emptyset$ for any $\alpha \neq \beta$ and

$$\int_{B_1(p_i^\alpha,h_i^L)} |Df_i|^2_{h_i^L,g} dz = \epsilon > 0.$$

It follows that

$$c \geq \int_{\Sigma_i} |Df_i|^2_{h_i^L,g} dz \geq L\epsilon.$$

Hence the process has to stop at some L. We obtain $\tilde{h}_i = h_i^L$ and a uniform $r_0 > 0$ such that for any $x \in \Sigma_i$,

$$\int_{B_{r_0}(x,\tilde{h}_i)} |Df_i|^2_{\tilde{h}_i,g} dz < \epsilon.$$

By uniform admissibility of $\tilde{h}_i$, we may choose m and R such that there are finitely many cylinder-like necks $N_{i,\alpha} \subset \Sigma_i$ $(\alpha = 1, \ldots, l)$ satisfying:
(1) $\Sigma_i \backslash \bigcup_\alpha N_{i,\alpha}$ is covered by geodesic balls $B_R(q_{ij}, \tilde{h}_i)$ $(1 \leq j \leq m)$ in Σ_i
(2) Each $N_{i,\alpha}$ is diffeomorphic to a cylinder of the form $S^1 \times (a_{i,\alpha}, b_{i,\alpha})$ ($a_{i,\alpha}$ and $b_{i,\alpha}$ may be $\pm\infty$)

Now by taking a subsequence if necessary, we may assume that for each j, the sequence $(\Sigma_i, \tilde{h}_i, q_{ij})$ of pointed metric spaces converges to a Riemann surface $\Sigma^0_{\infty,j}$. Such a limit $\Sigma^0_{\infty,j}$ is of the form

$$\Sigma_{\infty,j} \backslash \{p_{j1}, \ldots, p_{j\gamma_j}\},$$

where $\Sigma_{\infty,j}$ is a compact Riemann surface. More precisely, there is a natural admissible metric $\tilde{h}_{\infty,j}$ on each $\Sigma^0_{\infty,j}$ and a point $q_{\infty j}$ in $\Sigma^0_{\infty,j}$, such that for any fixed $r > 0$, when i is sufficiently large, there is a diffeomorphism $\phi_{i,r}$ from $B_r(q_{\infty j}, \tilde{h}_{\infty,j})$ onto $B_r(q_{ij}, \tilde{h}_i)$ satisfying: $\phi_{i,r}(q_{\infty j}) = q_{ij}$ and the pull-backs $\phi^*_{i,r}\tilde{h}_i$ converge to $\tilde{h}_{\infty,j}$ uniformly in the $\mathscr{C}^\infty$-topology over $B_r(q_{\infty j}, \tilde{h}_{\infty,j})$. Note that such a convergence of $\tilde{h}_i$ is assured by uniform admissibility.

Next we put together all these $\Sigma_{\infty,j}$ to form a connected curve Σ'_∞ as follows: For any two components $\Sigma_{\infty,j}$ and $\Sigma_{\infty,j'}$, we identify punctures $p_{js} \in \Sigma_{\infty,j}$ and $p_{j's'} \in \Sigma_{\infty,j'}$ (j may be equal to j') if for any sufficiently large i and r, the boundaries of $B_r(q_{ij}, \tilde{h}_i)$ and $B_r(q_{ij'}, \tilde{h}_i)$ specified above are contained in a cylindrical neck $N_{i,\alpha}$. In this way, we get a connected curve Σ_∞ (not necessarily stable) since each Σ_i is connected.

Since we have

$$\int_{B_{r_0}(x,\tilde{h}_i)} |Df_i|^2_{\tilde{h}_i,g} dz < \epsilon,$$

by taking a subsequence if necessary, we may assume that f_i converge to a J-holomorphic map f_∞ from $\bigcup_j \Sigma^0_{\infty,j}$ into M. By Proposition 7.3, f_∞ extends to a Hölder continuous J-holomorphic map from Σ_∞ into M. There is clearly also a limiting metric $\tilde{h}_\infty$ on Σ_∞, and Σ_∞ has the same genus as Σ_i for large i.

It remains to show that the homology class of f_∞ is the same as that of f_i. By convergence we have

$$\int_{\Sigma_\infty} |Df_\infty|^2_{\tilde{h}_\infty,g} dz = \lim_{r\to\infty} \lim_{i\to\infty} \int_{\bigcup_j B_r(q_{ij},\tilde{h}_i)} |Df_i|^2_{\tilde{h}_i,g} dz.$$

In fact, since the complement of $\bigcup_j B_r(q_{ij}, \tilde{h}_i)$ in Σ_i is contained in the union of cylindrical necks $N_{i,\alpha}$, it suffices to show that for each i, if $N_{i,\alpha} = S^1 \times (a, b)$, then

$$\lim_{r\to\infty} \lim_{i\to\infty} \int_{S^1\times(a+r,b-r)} |Df_i|^2_{\tilde{h}_i,g} dz = 0.$$

This can be seen as follows: By our choice of $\tilde{h}_i$, we know that for any $p \in N_{i,\alpha}$,

$$\int_{B_1(p,\tilde{h}_i)} |Df_i|^2_{\tilde{h}_i,g} dz \le \epsilon.$$

It follows from Lemma 7.2 that

$$\operatorname{diam}(f_i(B_2(p, \tilde{h}_i))) \le c\sqrt{\int_{B_4(p,\tilde{h}_i)} |Df_i|^2_{\tilde{h}_i,g} dz},$$

where c is a uniform constant. Since ϵ is small, both $f_i(S^1 \times \{a + r\})$ and $f_i(S^1 \times \{b - r\})$ are contained in geodesic balls of radius $c\sqrt{\epsilon}$. Moreover, by varying r slightly, we may assume that

$$\int_{S^1\times(a+r,b-r)} |Df_i|^2_{\tilde{h}_i,g} dz \le 10 \int_{B_4(p,\tilde{h}_i)} |Df_i|^2_{\tilde{h}_i,g} dz.$$

It follows that there are two smooth maps $u_{ij} : \Delta_1 \to M$ $(j = 1, 2)$ with

$$u_{i1}|_{\partial\Delta_1} = f_i|_{S^1\times\{a+r\}}, \quad u_{i2}|_{\partial\Delta_1} = f_i|_{S^1\times\{b-r\}}.$$

and such that

$$\int_{\Delta_1} |Du_{i1}|^2_{\tilde{h}_i,g} \le c \int_{S^1\times(a+r-2,a+r+2))} |Df_i|^2_{\tilde{h}_i,g} dz$$

and

$$\int_{\Delta_1} |Du_{i2}|^2_{\tilde{h}_i,g} \le c \int_{S^1\times(b-r-2,b-r+2))} |Df_i|^2_{\tilde{h}_i,g} dz,$$

where c is a uniform constant. The maps $f_i|_{N_{i,\alpha}}$ and u_{ij} can be easily glued together to form a continuous map from S^2 into M. Since each $f_i|_{S^1\times[d-1,d+1]}$, where $d \in (a + r, b - r)$, is contained in a small geodesic ball of M, this map must be null homologous. It follows

$$\int_{S^1\times(a+r,b-r)} |Df_i|^2_{\tilde{h}_i,g} dz = \int_{S^1\times(a+r,b-r)} f_i^*\omega = \int_{\Delta_1} u_{i1}^*\omega - \int_{\Delta_1} u_{i2}^*\omega.$$

Therefore, we have

$$\int_{S^1\times(a+r,b-r)} |Df_i|^2_{\tilde{h}_i,g} dz \le c \int_{S^1\times(a+r-2,a+r+2)\cup(b-r-2,b-r+2)} |Df_i|^2_{\tilde{h}_i,g} dz.$$

This implies that the homology classes of f_i converge to the homology class of f_∞. So Theorem 7.1 is proved. □

Remark 7.4 If (Σ_i, f_i) are stable maps, we may construct a stable limit $(\Sigma_\infty, f_\infty)$. Observe that Σ'_∞ may have components $\Sigma_{\infty,j}$ where f_∞ restricts to a constant map and which are conformal to $\mathbb{CP}^1$ and contain fewer than three other components. There are two possibilities for such $\Sigma_{\infty,j}$'s. If a $\Sigma_{\infty,j}$ attaches to only one other component, we simply drop $\Sigma_{\infty,j}$ from the construction; if $\Sigma_{\infty,j}$ contains exactly two other components, then we contract $\Sigma_{\infty,j}$ and identify those points where $\Sigma_{\infty,j}$ intersects with the other two components. Carrying out this process inductively, we eventually obtain a connected curve Σ_∞ such that the induced $f_\infty : \Sigma_\infty \to M$ is a stable map.

8 Second Variation of the $\bar{\partial}_J$-Equation and Applications

In Sect. 2 we saw that one prime difficulty in proving the isotopy theorem is the existence of a smoothing of a singular pseudo-holomorphic curve. Under positivity assumptions nodal curves can always be smoothed according to

Proposition 3.5. So to solve this problem it remains to find criteria when a pseudo-holomorphic map $\varphi : \Sigma \to M$ can be deformed to an immersion with only transversal branches. This seems generally a difficult problem, but over generic paths of almost complex structures miracles happen.

The content of this section is the technical heart of Shevchishin's work in [Sh]. The purpose of our presentation is to make this work more accessible by specializing to what we actually need.

8.1 Comparisons of First and Second Variations

Any of our moduli spaces $\mathscr{M}$ of pseudo-holomorphic maps is the zero set of a transverse section s of a Banach bundle $\mathscr{E}$ over a Banach manifold $\mathscr{B}$. This ambient Banach manifold also comes with a submersion π to a Banach manifold of almost complex structures. The purpose of this paragraph is to relate the first and second variations of s with those of π.

After choosing a local trivialization of the Banach bundle we have the abstract situation of two submersions of Banach manifolds. For the first variation the following result holds.

Proposition 8.1 *Let $\Phi : \mathscr{X} \to \mathscr{Y}$, $\Psi : \mathscr{X} \to \mathscr{Z}$ be locally split submersions of Banach manifolds. For $P \in \mathscr{X}$ let $\mathscr{M} = \Phi^{-1}(\Phi(P))$, $\mathscr{F} = \Psi^{-1}(\Psi(P))$ be the fibers through P and $\bar\Phi = \Phi_{|\mathscr{F}}$, $\bar\Psi = \Psi_{|\mathscr{M}}$ the restrictions.*

Then there exist canonical isomorphisms

$$\begin{aligned}
\operatorname{kern}(D\bar\Phi_{|P}) &= \operatorname{kern}(D\Phi_{|P}) \cap \operatorname{kern}(D\Psi_{|P}) = \operatorname{kern}(D\bar\Psi_{|P}),\\
\operatorname{coker}(D\bar\Phi_{|P}) &= (T_{\mathscr{Y},\Phi(P)} \oplus T_{\mathscr{Z},\Psi(P)})/(D\Phi_{|P}, D\Psi_{|P})(T_{\mathscr{X},P}) = \operatorname{coker}(D\bar\Psi_{|P}).
\end{aligned}$$

Proof. Let $X = T_{\mathscr{X},P}$, $Y = T_{\mathscr{Y},\Phi(P)}$, $Z = T_{\mathscr{Z},\Psi(P)}$, $M = \operatorname{kern} D\Phi = T_{\mathscr{M},P}$, $F = \operatorname{kern} D\Psi = T_{\mathscr{F},P}$. The claim follows from the following commutative diagram with exact rows.

$$\begin{tikzcd}
0 \arrow[r] & \operatorname{kern}(D\bar\Phi_{|P}) \arrow[r] \arrow[d, equal] & F \arrow[r, "D\bar\Phi_{|P}"] \arrow[d] & Y \arrow[r] \arrow[d] & \operatorname{coker}(D\bar\Phi_{|P}) \arrow[r] \arrow[d, "\simeq"] & 0\\
0 \arrow[r] & F \cap M \arrow[r] & X \arrow[r, "(D\Phi_{|P}, D\Psi_{|P})"] & Y \oplus Z \arrow[r] & (Y \oplus Z)/X \arrow[r] & 0\\
0 \arrow[r] & \operatorname{kern}(D\bar\Psi_{|P}) \arrow[r] \arrow[u, equal] & M \arrow[r, "D\bar\Psi_{|P}"] \arrow[u] & Z \arrow[r] \arrow[u] & \operatorname{coker}(D\bar\Psi_{|P}) \arrow[r] \arrow[u, "\simeq"] & 0
\end{tikzcd}$$

□

As an application of this lemma we can detect critical points of the projection $\mathscr{M} \to \mathscr{J}$ by looking at critical points of s for fixed almost complex structure. Note also that the linearization of a section $\mathscr{B} \to \mathscr{E}$ of a Banach bundle at any point P of its zero set is a well-defined map $T_{\mathscr{B},P} \to \mathscr{E}_P$. In fact, if $Q \in \mathscr{E}_P$ lies on the zero section then $T_{\mathscr{E},Q} \simeq T_{\mathscr{B},P} \oplus \mathscr{E}_P$ canonically.

The intrinsic meaning of the second variation is less apparent. In the notation of the proposition we are interested in situations when $\Psi|_{\mathscr{M}}$ is totally degenerate at $\overline{P}$, that is, if $T_{\mathscr{M},P} = T_{\mathscr{M},P} \cap T_{\mathscr{F},P}$. Now in any case the second variations of $\bar{\Phi}$ and $\bar{\Psi}$ induce two bilinear maps

$$\begin{aligned} \beta_1 :\ & (T_{\mathscr{M},P} \cap T_{\mathscr{F},P}) \times (T_{\mathscr{M},P} \cap T_{\mathscr{F},P}) \longrightarrow \operatorname{coker}(D\bar{\Phi}_{|P}) \\ \beta_2 :\ & (T_{\mathscr{M},P} \cap T_{\mathscr{F},P}) \times (T_{\mathscr{M},P} \cap T_{\mathscr{F},P}) \longrightarrow \operatorname{coker}(D\bar{\Psi}_{|P}), \end{aligned}$$

as follows. For $v, w \in T_{\mathscr{M},P} \cap T_{\mathscr{F},P}$ let $\tilde{v}, \tilde{w}$ be local sections around P of $T_{\mathscr{F}}$ and $T_{\mathscr{M}}$, respectively, with $\tilde{v}(P) = v$, $\tilde{w}(P) = w$. Then $D\bar{\Phi} \cdot \tilde{v}$ is a section $\alpha_{\tilde{v}}$ of $\bar{\Phi}^* T_{\mathscr{Y}}$ with $\alpha_{\tilde{v}}(P) = 0$ since $\tilde{v}(P) = v$ lies in $T_{\mathscr{M},P} = \operatorname{kern} D\Phi$. If

$$\mathrm{pr}_{\mathscr{Y}} : T_{\bar{\Phi}^* T_{\mathscr{Y}},P} = T_{\mathscr{F},P} \oplus T_{\mathscr{Y},\Phi(P)} \ \longrightarrow\ T_{\mathscr{Y},\Phi(P)}$$

denotes the projection define

$$\beta_1(v, w) = \mathrm{pr}_{\mathscr{Y}}(D\alpha_{\tilde{v}} \cdot w),$$

viewed modulo $D\bar{\Phi}(T_P\mathscr{F})$. This definition does not depend on the choice of extension $\tilde{v}$ by applying the following lemma with $\tilde{v}$ the difference of two extensions.

Lemma 8.2 *If $\tilde{v}(P) = 0$ then $\mathrm{pr}_{\mathscr{Y}}(D\alpha_{\tilde{v}} \cdot w) \in \operatorname{im}(D\bar{\Phi})$.*

Proof. In the local situation of open sets in Banach spaces $\mathscr{X} \subset X = T_{\mathscr{X},P}$, $\mathscr{Y} \subset Y = T_{\mathscr{Y},\Phi(P)}$, $\mathscr{Z} \subset Z = T_{\mathscr{Z},\Psi(P)}$ we have

$$\mathrm{pr}_{\mathscr{Y}}(D\alpha_{\tilde{v}} \cdot w) = \partial_w(\partial_{\tilde{v}}\Phi) = \partial^2_{w\tilde{v}(P)}\Phi + D\Phi \cdot \partial_w \tilde{v}.$$

The claim follows because $\tilde{v}(P) = 0$ and $\partial_w \tilde{v} \in T_P\mathscr{F}$. □

The analogous definition with Φ and Ψ swapped defines β_2.

Proposition 8.3 *Let $\Phi : \mathscr{X} \to \mathscr{Y}$, $\Psi : \mathscr{X} \to \mathscr{Z}$ be submersions of Banach manifolds with splittable differentials. For $P \in \mathscr{X}$ let $\mathscr{M} = \Phi^{-1}(\Phi(P))$, $\mathscr{F} = \Psi^{-1}(\Psi(P))$ be the fibers through P and $\bar{\Phi} = \Phi|_{\mathscr{F}}$, $\bar{\Psi} = \Psi|_{\mathscr{M}}$ the restrictions. Let $\Lambda : \operatorname{coker}(D\bar{\Phi}_{|P}) \to \operatorname{coker}(D\bar{\Psi}_{|P})$ denote the canonical isomorphism of Proposition 8.1 and β_1, β_2 the bilinear maps introduced above. Then*

$$\beta_2 = \Lambda \circ \beta_1.$$

Proof. In the local situation of the proofs of Proposition 8.1 and Lemma 8.2, by the definition of $\tilde{v}$, $\tilde{w}$ it holds $\partial_{\tilde{v}}\Psi = 0$, $\partial_{\tilde{w}}\Phi = 0$. Hence

$$\begin{aligned} (D\Phi_{|P}, D\Psi_{|P})[\tilde{w}, \tilde{v}] &= \big(D\Phi_{|P}[\tilde{w}, \tilde{v}], D\Psi_{|P}[\tilde{w}, \tilde{v}]\big) \\ &= \big(\partial_w(\partial_{\tilde{v}}\Phi) - \partial_v(\partial_{\tilde{w}}\Phi), \partial_w(\partial_{\tilde{v}}\Psi) - \partial_v(\partial_{\tilde{w}}\Psi)\big) \\ &= (\beta_1(v, w), 0) - (0, \beta_2(v, w)). \end{aligned}$$

Hence $\beta_1(v, w)$ and $\beta_2(v, w)$ induce the same element in $(Y \oplus Z)/X$, that is, $\beta_2 = \Lambda \circ \beta_1$. □

8.2 Moduli Spaces of Pseudo-Holomorphic Curves with Prescribed Singularities

Let $\varphi : \Sigma \to (M, J)$ be a pseudo-holomorphic map with $D\varphi$ vanishing at P of order $\mu - 1 > 0$. Then $C = \operatorname{im}\varphi$ has a singular point at $\varphi(0)$. The number μ is the *multiplicity* of the singularity, which agrees with the degree of the composition of φ with a *general* local projection $M \to \mathbb{C}$ with J-holomorphic fibers. (For non-general projections this mapping degree can be larger than the multiplicity.) Choosing charts we may assume $M = \mathbb{C}^2$, $\Sigma = \Delta$, $\varphi(0) = 0$ and J to agree with the standard complex structure I at 0. Let j be the complex structure on Δ. Writing $0 = D\varphi + J\circ D\varphi\circ j = (D\varphi + I\circ D\varphi\circ j) + (J - I)\circ D\varphi\circ j$ gives the estimate

$$\left|\partial_{\bar{t}}\varphi(t)\right| \leq c \cdot |t|^{\mu-1} \cdot |\varphi|.$$

Thus the higher dimensional analogue of Proposition 2.13 shows that φ is polynomial in t up to order $2\mu - 1$. It is not hard to see that this defines a holomorphic $(2\mu - 1)$-jet on $T_{\Delta,0}$ with values in $T_{M,\varphi(0)}$. Note that this jet generally does not determine the embedded topological type of C at $\varphi(0)$. In the integrable situation one needs twice the Milnor number of C at $\varphi(0)$ minus one coefficients to tell the topological type, and the Milnor number can be arbitrarily large for given multiplicity.

The induced jet with values in the normal bundle $N_{\varphi,0}$ (see Sect. 3.2) vanishes either identically or to order $\mu + \nu$, $0 \leq \nu \leq \mu - 1$. Define the *cusp index* of φ at P to be μ in the former case and to equal $\nu \leq \mu - 1$ in the latter case. (In [Sh] the multiplicity and cusp index are called primary and secondary cusp indices, respectively.)

For example, let $\varphi : \Delta \to (\mathbb{C}^2, J)$ be a pseudo-holomorphic singularity of multiplicity 2. Then

$$\varphi(t) = \big(\alpha t^2 + \beta t^3 + a(t), \gamma t^2 + \delta t^3 + b(t)\big)$$

with one of α or β non-zero and $a(t) = o(|t|^3)$, $b(t) = o(|t|^3)$. A linear coordinate change transforms $\begin{pmatrix}\alpha & \beta\\ \gamma & \delta\end{pmatrix}$ to $\begin{pmatrix}1 & 0\\ 0 & \delta\end{pmatrix}$, and hence we may assume $\alpha = 1$, $\beta = \gamma = 0$ and $\delta = 0$ or 1. Then φ defines the 3-jet with values in $T_{M,\varphi(0)}$ represented by $t \mapsto t^2\partial_z + \delta t^3\partial_w$. Going over to N means reducing modulo ∂_z. This leads to the 3-jet represented by $t \mapsto \delta t^3$. Thus φ has cusp index 0 if $\delta \neq 0$ and cusp index 1 otherwise. In analogy with the integrable situation the former singularity is called an *ordinary cusp*. We have seen that in this case

$$\varphi(t) = (t^2, t^3) + o(|t|^3)$$

in appropriate complex coordinates. We will use this below.

We can now define moduli spaces $\mathscr{M}_{\boldsymbol{\mu},\boldsymbol{\nu}}$ with prescribed multiplicities $(\mu_1, \ldots, \mu_m)$ and cusp indices $(\nu_1, \ldots, \nu_m)$, $0 \leq \nu_i \leq \mu_i$, in k marked points $P_1, \ldots, P_m \in \Sigma$, and an immersion everywhere else. A straightforward transversality argument shows that $\mathscr{M}_{\boldsymbol{\mu},\boldsymbol{\nu}}$ is a submanifold of the total moduli space $\mathscr{M}$ (without marked points) of real codimension

$$2(|\boldsymbol{\mu}|n - m(n+1)) + 2(n-1)|\boldsymbol{\nu}|, \tag{8.1}$$

where $n = \dim_{\mathbb{C}} M$. For details see [Sh], Sects. 3.2 and 3.3.

8.3 The Locus of Constant Deficiency

By the implicit function theorem critical points of the projection $\pi : \mathscr{M} \to \mathscr{J}$ have the property that the $\bar\partial_J$-operator for fixed J is obstructed. In fact, according to Proposition 8.1 the cokernels of the respective linearizations

$$D_{(j,\varphi,J)} = D_{\varphi,J} + J \circ D\varphi \circ j'$$

and $D\pi_{|(j,\varphi,J)}$ are canonically isomorphic. It is therefore important to study the stratification of $\mathscr{M}$ into subsets

$$\mathscr{M}^{h_1} := \{(j,\varphi,J) \in \mathscr{M} \mid \dim \operatorname{coker}(D_{(j,\varphi,J)}) \ge h^1\}.$$

To obtain an analytic description note that the discussion in Sect. 3.2 implies the following:

$$\operatorname{coker} D_{(J,\varphi,j)} \;=\; \operatorname{coker} D_{\varphi,J}/H^1(T_\Sigma) \;=\; \operatorname{coker} D^N_{\varphi,J}.$$

So in studying $\operatorname{coker} D\pi$ we might as well study the cokernel of the normal $\bar\partial$-operator $D^N_{\varphi,J}$.

The bundles $N = N_\varphi = \varphi^* T_M / D\varphi(T_\Sigma[A])$ on Σ from Sect. 3.2 do not patch to a complex line bundle on $\mathscr{M} \times \Sigma$ because their degree decreases under the presence of critical points of φ. However, once we restrict to $\mathscr{M}_{\boldsymbol{\mu},\boldsymbol{\nu}}$ the holomorphic line bundles $\mathcal{O}([A])$ encoding the vanishing orders of $D\varphi$ vary differentiably with φ; hence for any $\boldsymbol{\mu},\boldsymbol{\nu}$ there exists a complex line bundle N on $\mathscr{M}_{\boldsymbol{\mu},\boldsymbol{\nu}} \times \Sigma$ with fibers N_φ relative $\mathscr{M}_{\boldsymbol{\mu},\boldsymbol{\nu}}$. For the following discussion we therefore restrict to one such stratum $\mathscr{M}_{\boldsymbol{\mu},\boldsymbol{\nu}} \subset \mathscr{M}$. Denote by $\mathcal{N}$, $\mathcal{F}$ the Banach bundles over $\mathscr{M}_{\boldsymbol{\mu},\boldsymbol{\nu}}$ with fibers $W^{1,p}(N_\varphi)$ and $L^p(N_\varphi \otimes \Lambda^{0,1})$, respectively. The normal $\bar\partial$-operators define a family of Fredholm operators

$$\sigma : \mathcal{N} \longrightarrow \mathcal{F}$$

with the property that for any $x = (j,\varphi,J) \in \mathscr{M}_{\boldsymbol{\mu},\boldsymbol{\nu}}$ there is a canonical isomorphism

$$\operatorname{coker} \sigma_x \;=\; \operatorname{coker} D^N_{\varphi,J} \;=\; \operatorname{coker} D\pi_{|x}.$$

To understand the situation around some $x_0 = (j,\varphi,J) \in \mathscr{M}_{\boldsymbol{\mu},\boldsymbol{\nu}}$ choose a complement $Q \subset \mathcal{F}_{x_0}$ to $\operatorname{im} \sigma_{x_0}$ and let $V \subset \mathcal{N}_{x_0}$ be a complement to $\operatorname{kern} \sigma_{x_0}$. Extend these subspaces to subbundles $\mathcal{V} \subset \mathcal{N}$ and $\mathcal{Q} \subset \mathcal{F}$. (Here and in the following we suppress the necessary restrictions to an appropriate neighborhood of x_0.) Then

$$\mathcal{V} \oplus \mathcal{Q} \longrightarrow \mathcal{F}, \quad (v,q) \longmapsto \sigma(v) + q$$

is an isomorphism around x_0 since this is true at x_0. In particular, $\sigma(\mathcal{V}) \subset \mathcal{F}$ is a subbundle and there are canonical isomorphisms

$$\mathcal{Q} \xrightarrow{\simeq} \mathcal{F}/\sigma(\mathcal{V}), \quad \mathcal{V} \xrightarrow{\simeq} \mathcal{F}/\mathcal{Q}.$$

Having set up the bundles $\mathcal{Q}$ and $\mathcal{V}$ we fit $\operatorname{kern} \sigma_{x_0}$ into a vector bundle by setting

$$\mathcal{K} := \operatorname{kern}(\mathcal{N} \to \mathcal{F}/\mathcal{Q}).$$

By the Fredholm property of σ this is a bundle of finite rank; it contains $\operatorname{kern} \sigma_x$ for any x and it is complementary to $\mathcal{V}$:

$$\mathcal{K} \oplus \mathcal{V} = \mathcal{N}.$$

We claim that σ induces a section $\bar\sigma$ of the finite rank bundle $\operatorname{Hom}(\mathcal{K}, \mathcal{Q})$ with the property that there are canonical isomorphisms

$$\operatorname{kern} \sigma_x \simeq \operatorname{kern} \bar\sigma_x, \quad \operatorname{coker} \sigma_x \simeq \operatorname{coker} \bar\sigma_x,$$

for any $x \in \mathscr{M}_{\boldsymbol{\mu},\boldsymbol{\nu}}$ in the domain of our construction. In fact, this follows readily from the Snake Lemma applied to the following commutative diagram with exact columns and rows.

$$\begin{tikzcd}
 & \operatorname{kern} \bar\sigma_x \arrow[d] & \operatorname{kern} \sigma_x \arrow[d] & & \\
0 \arrow[r] & \mathcal{K}_x \arrow[r] \arrow[d, "\bar\sigma_x"'] & \mathcal{N}_x \arrow[r] \arrow[d, "\sigma_x"'] & \mathcal{N}_x/\mathcal{K}_x \arrow[r] \arrow[d, "\simeq"] & 0 \\
0 \arrow[r] & \mathcal{Q}_x \arrow[r] \arrow[d] & \mathcal{F}_x \arrow[r] \arrow[d] & \mathcal{F}_x/\mathcal{Q}_x \arrow[r] & 0 \\
 & \operatorname{coker} \bar\sigma_x & \operatorname{coker} \sigma_x & &
\end{tikzcd} \tag{8.2}$$

Because σ_x maps $\mathcal{N}_x/\mathcal{K}_x$ isomorphically to $\mathcal{F}_x/\mathcal{Q}_x$ we also see that

$$\mathscr{M}^{h^1}_{\boldsymbol{\mu},\boldsymbol{\nu}} = \left\{ (j, \varphi, J) \in \mathscr{M}_{\boldsymbol{\mu},\boldsymbol{\nu}} \;\middle|\; \dim \operatorname{coker} D_{(j,\varphi,J)} \geq h^1 \right\}, \quad h^1 = \dim \operatorname{coker} \sigma_x$$

equals the zero locus of $\bar\sigma$ viewed as section of $\operatorname{Hom}(\mathcal{K}, \mathcal{Q})$ locally.

Proposition 8.4 *([Sh], Corollary 4.4.2.)* $\mathscr{M}^{h_1}_{\boldsymbol{\mu},\boldsymbol{\nu}}$ *is a submanifold inside* $\mathscr{M}_{\boldsymbol{\mu},\boldsymbol{\nu}}$ *of codimension* $(\operatorname{index} + h^1) \cdot h^1$.

Proof. This follows from the implicit function theorem once we prove that $\bar\sigma$ is a transverse section for $\operatorname{Hom}(\mathcal{K}, \mathcal{Q})$ since

$$\operatorname{rank} \operatorname{Hom}(\mathcal{K}, \mathcal{Q}) = (\operatorname{index} + h^1) \cdot h^1.$$

To this end we look at variations at $x \in \mathscr{M}$ with φ and j fixed and with the variation J_s of J constant along $\operatorname{im}\varphi$. Note that such a path stays inside $\mathscr{M}$. The pull-back to the path of the bundles $\mathcal{N}$ and $\mathcal{F}$ with fibers

$$\begin{aligned} W^{1,p}(N) &= W^{1,p}(\varphi^* T_M)/D\varphi\big(W^{1,p}(T_\Sigma)[A]\big) \\ L^p(N \otimes \Lambda^{0,1}) &= L^p(\varphi^* T_M \otimes \Lambda^{0,1})/D\varphi\big(L^p(T_\Sigma)[A] \otimes \Lambda^{0,1}\big) \end{aligned}$$

are manifestly trivial. Now σ is fiberwise given by $D^N_{\varphi,J}$, which in turn can be computed by lifting a section of N to $\varphi^* T_M$, applying

$$D_{\varphi,J} = \nabla + J \circ \nabla_{j(.)} + \nabla J \circ D_{j(.)}\varphi$$

and reducing modulo $D\varphi(T_\Sigma[A])$. The result of our variation is thus

$$\big(\nabla_{J'}\sigma\big)(v) = (\nabla_v J') \circ D\varphi \circ j,$$

written in the form lifted to $\varphi^* T_M$. Following the discussion above we now want to look at the derivative of the induced section $\bar\sigma$ of $\operatorname{Hom}(\mathcal{K}, \mathcal{Q})$. Write $h^0 = \dim \operatorname{kern} \sigma_x$, in analogy to $h^1 = \dim \operatorname{coker} \sigma_x$. For the construction of Q let $W \subset M$ be an open set such that $\varphi^{-1}(W) \subset \Sigma$ is a unit disk and such that there are complex-valued coordinates z, w on M with

$$\varphi(t) = (t, 0) \quad \text{for } t \in \Delta$$

in these coordinates. Note that $D_\varphi(T_\Sigma[A]) = \langle \partial_z \rangle$, so for the induced section of N only the ∂_w-part matters. Let χ be the characteristic function of $\varphi^{-1}(W)$ in Σ, that is, $\chi|_{\varphi^{-1}(W)} = 1$ and $\operatorname{supp}\chi \subset \operatorname{cl}\varphi^{-1}(W)$. Then

$$\chi \partial_w \otimes d\bar t \in L^p(\Sigma, \varphi^* T_M \otimes \Lambda^{0,1}).$$

Because φ is injective away from finitely many points, open sets of the form $\varphi^{-1}(W)$ span a base for the topology of Σ away from finitely many points. Now characteristic functions span a dense subspace in L^p. We can therefore find pairwise disjoint $W_1, \ldots, W_{h^1} \subset M$ such that the corresponding $\chi_j \partial_w \otimes d\bar t$ span the desired complementary subspace Q of $\operatorname{im}\sigma_x$.

To compute $\nabla_{J'}\bar\sigma \in \operatorname{Hom}(\mathcal{K}_x, \mathcal{Q}_x) = \operatorname{Hom}(\operatorname{kern}\sigma_x, Q)$ it suffices to restrict $\nabla_{J'}\sigma$ to $\operatorname{kern}\sigma_x \subset \mathcal{N}_x$, and to compose with a projection

$$q : \mathcal{F}_x \longrightarrow \mathbb{R}^{h^1}$$

that induces an isomorphism $Q \to \mathbb{R}^{h^1}$. This follows from Diagram (8.2). For q we take the map

$$\mathcal{F}_x \ni \xi \longmapsto \Big(\operatorname{Im} \int_{\varphi^{-1}(W_j)} dt \wedge \langle dw, \xi \rangle\Big)_{j=1,\ldots,h^1}.$$

This maps $\chi_j \partial_w \otimes d\bar{t}$ to a non-zero multiple of the jth unit vector. Hence q is one-to-one on Q.

Then $q \circ \nabla_{J'} \sigma$ maps $v \in \operatorname{kern} \sigma_x$ to

$$\Big(\operatorname{Im} \int_{\varphi^{-1}(W_j)} dt \wedge \langle dw, \nabla_v J' \circ D\varphi \circ j \rangle \Big)_{j=1,\ldots,h^1}.$$

Now consider variations of the form $J' = gw \, d\bar{z} \otimes \partial_w$ in coordinates (z, w), that is, $\begin{pmatrix} 0 & 0 \\ gw & 0 \end{pmatrix}$ in matrix notation. If v is locally represented by $f\partial_w$ then

$$\nabla_v J' \circ D\varphi \circ j \;=\; (fgd\bar{z} \otimes \partial_w) \circ D\varphi \circ j \;=\; \sqrt{-1} fgd\bar{t} \otimes \partial_w,$$

and

$$\Big(q \circ \nabla_{J'} \sigma\Big)(v) = \Big(\operatorname{Im} \int_{\varphi^{-1}(W_j)} \sqrt{-1} fg \, dt \wedge d\bar{t}\Big)_{j=1,\ldots,h^1}.$$

By the identity theorem for pseudo-analytic functions the restriction map $\operatorname{kern} \sigma_x \rightarrow L^p(\varphi^{-1}(W_j), N)$ is injective. Thus for each j there exist $g_{j1}, \ldots, g_{jh^0}$ with support on W_j such that

$$\operatorname{kern} \sigma_x \longrightarrow \mathbb{R}^{h^0}, \quad f\partial_w \longmapsto \Big(\operatorname{Im} \int_{\varphi^{-1}(W_j)} \sqrt{-1} f g_{jk} \, dt \wedge d\bar{t}\Big)_{k=1,\ldots,h^0}$$

is an isomorphism. The corresponding variations J'_{jk} of $\bar{\sigma}_x$ (with support on $\varphi^{-1}(W_j)$) span $\operatorname{Hom}(\operatorname{kern} \bar{\sigma}_x, Q)$. □

Remark 8.5 The proof of the proposition in [Sh] has a gap for $h^1 > 1$, as pointed out to us at the summer school by Jean-Yves Welschinger. In this reference $\mathcal{Q}_x$ is canonically embedded into $L^p(\Sigma, N^* \otimes \Lambda^{1,0})$ as kernel of the adjoint operator. The problem is that the proof of surjectivity of the relevant linear map

$$T_{\mathscr{M}_{\boldsymbol{\mu\nu}},(j,\varphi,J)} \longrightarrow \operatorname{Hom}(\mathcal{K}_x, \mathcal{Q}_x)$$

relies on the fact that $\langle \mathcal{K}_x, \mathcal{Q}_x \rangle$ spans an $h^0 \cdot h^1$-dimensional subspace of $L^p(\Sigma, \Lambda^{1,0})$, where $\langle \, , \, \rangle$ is the dual pairing. Our proof shows that this is indeed the case.

Corollary 8.6 *For a general path* $\{J_t\}_{t \in [0,1]}$ *of almost complex structures any critical point* (j, φ, t) *of the projection* $p : \mathscr{M}_{\{J_t\}} \rightarrow [0, 1]$ *is a pseudo-holomorphic map with only ordinary cusps and such that* $\dim \operatorname{coker} D^N_{\varphi, J_t} = 1$.

Proof. This is a standard transversality argument together with dimension counting. Note that each singular point of multiplicity μ causes $\dim(\operatorname{kern} \sigma_{(j,\varphi,J_t)})$ to drop by $\mu - 1$. □

8.4 Second Variation at Ordinary Cusps

Corollary 8.6 leaves us with the treatment of pseudo-holomorphic maps with only ordinary cusps and such that D^N_{φ,J_t} has one-dimensional cokernel. Then the projection $p : \mathscr{M}_{\{J_t\}} \to [0,1]$ is not a submersion. The maybe most intriguing aspect of Shevchishin's work is that one can see quite clearly how the presence of cusps causes these singularities. They turn out to be quadratic, non-zero and indefinite. In particular, such a pseudo-holomorphic map always possesses deformations *with fixed almost complex structure* J_t into non-critical points of π.

Let us fit this situation into the abstract framework of Sect. 8.1. In the notation employed there $\mathscr{Z} = (-\varepsilon, \varepsilon)$ is the local parameter space of the path $\{J_t\}$, $\mathscr{X}$ is a neighborhood of the critical point $P = (j, \varphi, t)$ in the pull-back (via $\mathscr{Z} = (-\varepsilon, \varepsilon) \to \mathscr{J}$) of the ambient Banach manifold

$$\mathscr{B} := \mathbb{T}_g \times W^{1,p}(\Sigma, M) \times \mathscr{J},$$

and $\mathscr{Y} = L^p(\Sigma, \varphi^* T_M \otimes \Lambda^{0,1})$. The map $\Phi : \mathscr{X} \to \mathscr{Y}$ is a local non-linear $\bar\partial_J$-operator obtained from s via a local trivialization of the Banach bundle $\mathscr{E}$, while $\Psi : \mathscr{X} \to \mathscr{Z}$ is the projection, and $\bar\Psi = p$. The fiber of Φ through P is (an open set in) $\mathscr{M}_{\{J_t\}}$, and the fiber of Ψ through P is the ambient Banach manifold for $\mathscr{M}_{J_t}$.

Because $\mathscr{Z}$ is one-dimensional and (j, φ, J) is a critical point of the projection $p : \mathscr{M}_{\{J_t\}} \to [0,1]$ it holds $T_{\mathscr{M}_{\{J_t\}},P} \subset T_{\mathscr{F},P}$. Proposition 8.3 now says that we can compute the second order approximation of p near P by looking at the second order approximation of Φ restricted to $\mathscr{F}$, composed with the projection to the cokernel of the linearization. In Sect. 8.1 this symmetric bilinear form was denoted β_1. We are going to compute the associated quadratic form. We denote tangent vectors of the relevant tangent space $T_{\mathscr{M}_{\{J_t\}},P} \cap T_{\mathscr{F},P} = T_{\mathscr{M}_{\{J_t\}},P}$ by pairs (j', v) with j' a tangent vector to the space of complex structures on Σ and $v \in W^{1,p}(\Sigma, \varphi^* T_M)$.

Recall that the linearization of the $\bar\partial_J$-operator for fixed almost complex structure J is

$$D\bar\partial_J : (j', v) \longmapsto D_{(j,\varphi,J)}(j', v) = D_{\varphi,J} v + J \circ D\varphi \circ j',$$

where

$$D_{\varphi,J} v = \nabla v + J \circ \nabla \circ j(v) + \nabla_v J \circ D\varphi \circ j = 2\bar\partial_{\varphi,J} + R.$$

Near a cusp choose local coordinates z, w on M and t on Σ such that $J_{|(0,0)}$ equals I, the standard complex structure on $\mathbb{C}^2$, and $\varphi(t) = (t^2, t^3) + o(|t|^3)$ in these coordinates. Let $0 < \varepsilon < 1$ and let $\rho : \Delta \to [0,1]$ be a smooth function with support in $|t| < 3\varepsilon/4$, identically 1 for $|t| < \varepsilon/4$ and with $|d\varphi| < 3/\varepsilon$. Ultimately we will let ε tend to 0, but for the rest of the computation ε is fixed. We consider the variation of σ along (j', v) with

$$v = D\varphi(\rho\, t^{-1}\partial_t) = \rho\, t^{-1}\partial_t\varphi, \quad j' = j \circ \bar\partial(\rho\, t^{-1}\partial_t) = it^{-1}\partial_{\bar t}\rho\, \partial_t \otimes d\bar t.$$

Again we use complex notation for the complex vector bundle $\varphi^* T_M \otimes_{\mathbb{C}} T_\Sigma$. Taking the real part reverts to real notation. Note that j' is smooth and supported in the annulus $\varepsilon/4 < |t| < 3\varepsilon/4$. For any $\mu \in \mathbb{C}$ the multiple $\mu \cdot (j', v)$ is indeed a tangent vector to $\mathscr{M}_{J_t}$ because

$$D_{\varphi,J} \circ D\varphi = \bar\partial_{\varphi,J} \circ D\varphi = D\varphi \circ \bar\partial,$$

since $R \circ D\varphi = 0$ and by definition of the holomorphic structure on $\varphi^* T_M$; taking into account pseudo-holomorphicity $J \circ D\varphi \circ j = -D\varphi$ of φ this implies

$$D_{\varphi,J}(\mu v) + J \circ D\varphi \circ (\mu j') = (D\varphi \circ \bar\partial)(\mu\rho t^{-1}\partial_t) + (J \circ D\varphi \circ j)\big(\bar\partial(\mu\rho\, t^{-1}\partial_t)\big) = 0,$$

as needed.

At this point it is instructive to connect this variation to the discussion of the normal $\bar\partial$-operator $D^N_{\varphi,J}$ in Sect. 3.2. There we identified the tangent space of $\mathscr{M}$ relative $\mathscr{J}$ with $\operatorname{kern} D^N_{\varphi,J} \oplus H^0(\mathscr{N}^{\mathrm{tor}})$, see (3.4). Away from the node the vector field v lies in $D\varphi(W^{1,p}(T_\Sigma))$, and indeed v is a local frame for the complex line bundle $D\varphi(T_\Sigma[A])$. Thus we are dealing with a variation whose part in $\operatorname{kern} D^N_{\varphi,J}$ vanishes and which generates the skyscraper sheaf $\mathscr{N}^{\mathrm{tor}}$ locally at the cusp.

The variation (j', v) is concentrated in $\mathscr{B}_\varepsilon(0) \subset \Sigma$ and we can work in our local coordinates t, z, w. Then

$$v = \rho t^{-1}\partial_t\big((t^2, t^3) + o(|t|^3)\big) = \rho \cdot (2, 3t) + o(|t|),$$

which can be represented by the variation

$$\varphi_s = \varphi + s\rho t^{-1}\partial_t\varphi = \varphi + s\rho \cdot \big((2, 3t) + o(|t|)\big).$$

Similarly, we represent j' by a variation of holomorphic coordinate t_s of t with associated $\bar\partial$-operator

$$\partial_{\bar t_s} = \partial_{\bar t} + a_s\partial_t, \quad a_s = sit^{-1}\partial_{\bar t}\rho.$$

The derivative with respect to s yields $it^{-1}\partial_{\bar t}\rho\partial_t$, and hence $\bar\partial_s$ indeed represents j'.

The non-linear $\bar\partial$-operator for (j_s, φ_s) applied to φ_s yields

$$\bar\partial_s\varphi_s \;=\; \frac{1}{2}\Big(D\varphi_s + J_{|\varphi_s} \circ D\varphi_s \circ j_s\Big) \;=\; (\partial_{\bar t_s}\varphi_s)d\bar t_s + \frac{1}{2}K_s,$$

with

$$K_s = (J_{|\varphi_s} - I) \circ D\varphi_s \circ j_s.$$

Using $d\bar t_s$ to trivialize $\mathscr{E}$ along the path (j_s, φ_s) we now compute for the second variation

$$\begin{aligned}\tfrac{d^2}{ds^2}\big|_{s=0}\partial_{\bar t_s}\varphi_s = {} & \tfrac{d^2}{ds^2}\big|_{s=0}(\partial_{\bar t}+a_s\partial_t)\big(\varphi+s\rho\cdot\big((2,3t)+o(|t|)\big)\big)\\ = {} & it^{-1}\partial_{\bar t}\rho\cdot\big(\rho\cdot 3\partial_w+o(1)\big) \;=\; \frac{3i}{2}t^{-1}\partial_{\bar t}\rho^2\cdot\partial_w+t^{-1}\cdot o(1),\end{aligned} \tag{8.3}$$

and

$$K'':=\tfrac{d^2}{ds^2}\big|_{s=0}K_s \;=\; (J_{|\varphi}-I)\cdot\tfrac{d^2}{ds^2}\big|_{s=0}(D\varphi_s\circ j_s)+\nabla_v J\circ\nabla v\circ j+\nabla_v J\circ D\varphi\circ j'.$$

This is non-zero only for $|t|<\varepsilon$, and the first two terms are bounded pointwise, uniformly in ε. Here ∇ denotes the flat connection with respect to the coordinates z,w. For the last term we have

$$|\nabla_v J\circ D\varphi\circ j'|\le \mathrm{const}\cdot|t||t^{-1}|\cdot|\partial_{\bar t}\rho|\le\mathrm{const}\cdot\varepsilon^{-1},$$

by the choice of ρ. Taken together this gives the pointwise bound

$$|K''|\le\mathrm{const}\cdot\varepsilon^{-1}\cdot\chi_\varepsilon,$$

where χ_ε is the characteristic function for $B_\varepsilon(0)$, that is, $\chi_\varepsilon(t)=1$ for $|t|<\varepsilon$ and 0 otherwise.

Lemma 8.7 *Let $\lambda dt\in\ker(D^N_{\varphi,j})^*\subset W^{1,p}(N^*\otimes\Lambda^{1,0})$ and denote by $\Lambda:\operatorname{coker}D^N_{\varphi,j}\to\mathbb{R}$ the associated homomorphism with kernel $\operatorname{im}D^N_{\varphi,j}$ induced by*

$$L^p(\varphi^*T_M\otimes\Lambda^{0,1})\to\mathbb{R},\quad \gamma\longmapsto\operatorname{Re}\int_\Sigma\langle\lambda dt,\gamma\rangle.$$

Then for $\mu\in\mathbb{C}$ it holds

$$(\Lambda\circ\beta_1)\big(\mu\cdot(j',v),\mu\cdot(j',v)\big)\xrightarrow{\varepsilon\to 0}-3\pi\operatorname{Re}(\mu^2)\lambda(\partial_w)(0).$$

Proof. The formal adjoint of $D^N_{\varphi,j}$ is, just as the operator itself, the sum of a $\bar\partial$-operator and an operator of order zero:

$$(D\bar\partial_J)^*=\bar\partial_{N^*\otimes\Lambda^{1,0}}+R^*.$$

Thus

$$\bar\partial_{N^*\otimes\Lambda^{1,0}}(\lambda dt)=-R^*(\lambda dt) \tag{8.4}$$

is uniformly bounded pointwise. By the definition of β_1 and the discussion above we need to compute

$$\begin{aligned}\int_\Sigma\langle\lambda dt,(\tfrac{d^2}{ds^2}\big|_{s=0}\partial_{\bar t_s}\varphi_s)d\bar t+\tfrac12 K''\rangle = {} & \frac{3i}{2}\int_{B_\varepsilon(0)}\lambda(\partial_w)\partial_{\bar t}\rho^2 t^{-1}dt\wedge d\bar t\\ & +\int_{B_\varepsilon(0)}\big(o(1)t^{-1}+\mathrm{const}\cdot\varepsilon^{-1}\big)dt\wedge d\bar t.\end{aligned}$$

The second integral tends to 0 with ε. The first integral can be rewritten as a sum of

$$\frac{3i}{2}\int_{B_\varepsilon(0)} t^{-1}\rho^2\bar\partial\big(\lambda(\partial_w)dt\big),$$

again tending to zero with ε in view of (8.4), and

$$-\frac{3i}{2}\int_{B_\varepsilon(0)} t^{-1}\bar\partial\big(\rho^2\lambda(\partial_w)dt\big) \;=\; \frac{3i}{2}\cdot(2\pi i)\big(\rho^2\lambda(\partial_w)\big)(0) \;=\; -3\pi\lambda(\partial_w)(0).$$

Here the first equality follows from the Cauchy integral formula.

Our computation is clearly quadratic in rescaling (j', v) by μ. Thus replacing (j', v) by $\mu\cdot(j', v)$ and taking the real part gives the stated formula. □

Proposition 8.8 *Let $\{J_t\}_{t\in[0,1]}$ be a general path of almost complex structures on a four-manifold M. Assume that $P = (j, \varphi, t)$ is a critical point of the projection $p : \mathscr{M}_{\{J_t\}} \to [0,1]$ with φ not an immersion. Then there exists a locally closed, two-dimensional submanifold $Z \subset \mathscr{M}_{\{J_t\}}$ through P with coordinates x, y such that*

$$p(x, y) = x^2 - y^2.$$

Moreover, Z can be chosen in such a way that the pseudo-holomorphic maps corresponding to $(x, y) \neq 0$ are immersions.

Proof. By Corollary 8.6 the critical points of φ are ordinary cusps and $\dim\operatorname{coker} D^N_{\varphi,j} = 1$. Assume first that there is exactly one cusp. Another transversality argument shows that for general paths $\{J_t\}$ a generator λdt of $\operatorname{kern}(D^N_{\varphi,j})^*$ does not have a zero at this cusp. Let (j', v) be as in the discussion above with $\varepsilon > 0$ so small that the quadratic form

$$\mathbb{C} \ni \mu \longmapsto (\Lambda\circ\beta_1)\big(\mu\cdot(j', v), \mu\cdot(j', v)\big)$$

is non-degenerate and indefinite. This is possible by Lemma 8.7 and by what we just said about generators of $\operatorname{kern}(D^N_{\varphi,j})^*$. Let $Z \subset \mathscr{M}_{\{J_t\}}$ be a locally closed submanifold through P with $T_{Z,P}$ spanned by $\operatorname{Re}(j', v)$ and $\operatorname{Im}(j', v)$. The result is then clear by the Morse Lemma because $\beta_1|_{T_{Z,P}}$ describes the second variation of the composition $Z \to \mathscr{M}_{J_t} \to [0,1]$.

In the general case of several cusps, for each cusp we have a tangent vector (j'_l, v_l) with support close to it. Now run the same argument as before but with $(j', v) = \sum_l (j'_l, v_l)$. For ε sufficiently small these variations are supported on disjoint neighborhoods, and hence the only difference to the previous argument is that the coefficient $\lambda(\partial_w)$ for the quadratic form gets replaced by the sum $\lambda\big(\sum_l \partial_{w_l}\big)$. Again, for general paths, this expression is non-zero. □

Remark 8.9 We have chosen to use complex, local notation as much as possible and to neglect terms getting small with ε. This point of view clearly exhibits the holomorphic nature of the critical points in the moduli space near a cuspidal curve and is also computationally much simpler than the full-featured computations in [Sh]. In fact, in the integrable case coordinates can be chosen in such a way that all error terms $o(1)$ etc. disappear and the formula in Lemma 8.7 holds for $\varepsilon > 0$.

9 The Isotopy Theorem

9.1 Statement of Theorem and Discussion

In this section we discuss the central result of these lectures. It deals with the classification of symplectic submanifolds in certain rational surfaces. As a consequence the expected "stable range" for this problem indeed exists. In this range there are no new symplectic phenomena compared to complex geometry.

Theorem 9.1 *(1) Let M be a Hirzebruch surface and $\Sigma \subset M$ a connected surface symplectic with respect to a Kähler form. If* $\deg(p|_\sigma) \leq 7$ *then Σ is symplectically isotopic to a holomorphic curve in M, for some choice of complex structure on M.*

(2) Any symplectic surface in $\mathbb{CP}^2$ of degree $d \leq 17$ is symplectically isotopic to an algebraic curve.

A Hirzebruch surface M is a holomorphic $\mathbb{CP}^1$-bundle over $\mathbb{CP}^1$. They are projectivizations of holomorphic 2-bundles over $\mathbb{CP}^1$. The latter are all split, so $M = \mathbb{F}_k := \mathbb{P}(\mathcal{O} \oplus \mathcal{O}(k))$ for some $k \in \mathbb{N}$. The k is determined uniquely as minus the minimal self-intersection number of a section. If $k = 0$ we have $M = \mathbb{CP}^1 \times \mathbb{CP}^1$ and there is a whole $\mathbb{CP}^1$ worth of such sections; otherwise the section is holomorphically rigid and it is in fact unique. Topologically $\mathbb{F}_k$ is the non-trivial S^2-bundle over S^2 for k odd and $\mathbb{F}_k \simeq S^2 \times S^2$ for k even. It is also worthwhile to keep in mind that for any k, l with $2l \leq k$ there is a holomorphic one-parameter deformation with central fiber $\mathbb{F}_k$ and general fiber $\mathbb{F}_{k-2l}$, but not conversely. So in a sense, $\mathbb{CP}^1 \times \mathbb{CP}^1$ and $\mathbb{F}_1$ are the most basic Hirzebruch surfaces, those that are stable under small deformations of the complex structure. Note also that $\mathbb{F}_1$ is nothing but the blow-up of $\mathbb{CP}^2$ in one point.

The degree bounds in the theorem have to do with the method of proof and are certainly not sharp. For example, it should be possible to treat the case of degree 18 in $\mathbb{CP}^2$ with present technology. We even believe that the theorem should hold without any bounds on the degree.

In Sect. 6.4 we saw the importance of this result for genus-2 Lefschetz fibrations and for Hurwitz-equivalence of tuples of half-twists in the braid group $B(S^2, d)$ with $d \leq 7$.

9.2 Pseudo-Holomorphic Techniques for the Isotopy Problem

Besides the purely algebraic approach by looking at Hurwitz-equivalence for tuples of half-twists, there exists only one other approach to the isotopy problem for symplectic 2-manifolds inside a symplectic manifold (M, ω), namely by the technique of pseudo-holomorphic curves already discussed briefly in Sect. 2.4,III, as explained in Sect. 2.4,II. This works in three steps. (1) Classify pseudo-holomorphic curves up to isotopy for *one particular* almost complex structure I on M; typically I is integrable and the moduli space of holomorphic curves can be explicitly controlled by projective algebraic geometry; this step is quite simple. (2) Choose a general family of almost complex structures $\{J_t\}_{t\in[0,1]}$ with (a) B is J_0-holomorphic (b) $J_1 = I$. By the results on the space of tamed almost complex structures this works without problems as long as the symplectic form ω is isotopic to a Kähler form for I. (3) Try to deform B as pseudo-holomorphic curve with the almost complex structure, that is, find a smooth family $\{B_t\}$ of submanifolds such that B_t is J_t-holomorphic.

The last step (3) is the hardest and most substantial obstruction for isotopy results for two reasons. First, while for general paths of almost complex structures the space of pseudo-holomorphic curves $\mathscr{M}_{\{J_t\}}$ over the path is a manifold, the projection to the parameter interval $[0, 1]$ might have critical points. If $\{B_t\}_{t\le t_0}$ happens to run into such a point it may not be possible to deform B_{t_0} to $t > t_0$ and we are stuck. To avoid this problem one needs an unobstructedness result for deformations of smooth pseudo-holomorphic curves. The known results on this require some positivity of M, such as in Proposition 3.4. And second, even if this is true, as in the cases of $\mathbb{CP}^2$ and $\mathbb{F}_k$ that we are interested in, there is no reason that $\lim_{t\to t_0} B_t$ is a submanifold at all. The Gromov compactness theorem rather tells us that such a limit makes only sense as a stable J_t-holomorphic map or as a J_t-holomorphic 2-cycle. In the sequel we prefer to use the embedded point of view and stick to pseudo-holomorphic cycles. In any case, these are singular objects that we are not allowed to use in the isotopy. So we need to be able to change the already constructed path $\{B_t\}$ to bypass such singular points. This is the central problem for the proof of isotopy theorems in general.

In view of unobstructedness of deformations of smooth curves it suffices to solve the following:

1. Find a J_{t_0}-holomorphic smoothing of $C = \lim_{t\to t_0} B_t$
2. Show that any two pseudo-holomorphic smoothings of C are isotopic

In our situation the smoothing problem (1) of a pseudo-holomorphic cycle $C = \sum_a m_a C_a = \lim_{t\to t_0} B_t$ has the following solution. For a general path $\{J_t\}$ we know by the results of Sect. 8 that each J_{t_0}-holomorphic curve has a deformation into a nodal curve. Now for each a take m_a copies $C_{a,1}, \ldots, C_{a,m_a}$ of C_a. Deform each $C_{a,j}$ slightly in such a way that $\sum_{a,j} C_{a,j}$ is a nodal curve. This is possible by positivity. Finally apply the smoothing result for nodal curves (Proposition 3.5) to obtain a smoothing of C.

Problem (2) concerns the isotopy of smoothings of singular pseudo-holomorphic objects, which boils down to a question about the local structure of the moduli space of pseudo-holomorphic cycles as follows:

Let $C = \sum_a m_a C_a$ be a J_{t_0}-holomorphic cycle. Looking at the space of pairs (C', t) where $t \in [0,1]$ and C' is a J_t-holomorphic cycle, do the points parametrizing singular cycles locally disconnect it?

So we ask if any point has a neighborhood that stays connected once we remove the points parametrizing singular cycles. We believe this question has a positive answer under the positivity assumption that $c_1(M) \cdot C_a > 0$ for each irreducible component C_a of C. In the integrable case this follows from the unobstructedness results of Sect. 5, which say that there is a local complex parameter space for holomorphic deformations of C; the subset of singular cycles is a proper analytic subset, and hence its complement remains connected. However, as already discussed briefly in Sect. 5.1 no such parametrization is known for general almost complex structures except in the nodal case of Sect. 3.4.

9.3 The Isotopy Lemma

Instead of solving the parametrization problem for pseudo-holomorphic cycles we use a method to reduce the "badness" of singularities of $\lim_{t\to t_0} B_t$ by cleverly adding pointwise incidence conditions. This technique has been introduced into symplectic topology by Shevchishin in his proof of the local isotopy theorem for smoothings of a pseudo-holomorphic curve singularity [Sh].

How does this work? The pseudo-holomorphic cycle C consists of irreducible components C_a. Write C_a as image of a pseudo-holomorphic map $\varphi_a : \Sigma_a \to M$. Pseudo-holomorphic deformations of C_a keeping the genus (*equigeneric deformations*) can be realized by deforming φ_a and the complex structure on Σ_a. The moduli space of such maps has dimension

$$k_a := c_1(M) \cdot C_a + g(C_a) - 1 \geq 0.$$

Each imposing of an incidence with a point on M reduces this dimension by one, *provided the number of points added does not exceed $c_1(M) \cdot C_a$*. Thus choosing k_a general points on C_a implies that there are no non-trivial equigeneric deformations of C_a respecting the incidences. Doing this for all a we end up with a configuration of $\sum_a k_a$ points such that any J_{t_0}-holomorphic deformation of C containing all these points must somehow have better singularities. This can happen either by dropping $\sum_a (m_a - 1)$, which measures how multiple C is, or, if this entity stays the same, by the virtual number δ of double points of $|C|$. The latter is the sum over the maximal numbers of double points of local pseudo-holomorphic deformations of the map with image $|C|$ near the singular points. Pseudo-holomorphic deformations where the pair (m, δ) gets smaller in the way just described are exactly the deformations that can not be realized by deforming just the φ_a.

On the other hand, to have freedom to move B_t by keeping the incidence conditions there is an upper bound on the number of points we can add by the excess positivity that we have, which is $c_1(M) \cdot C$. In fact, each point condition decreases this number in the proof of Proposition 3.4 by one. Thus this method works as long as

$$\sum_a \big(c_1(M) \cdot C_a + g(C_a) - 1\big) \;<\; c_1(M) \cdot C. \tag{9.1}$$

If C is reduced ($m_a = 1$ for all a) then this works only if all components have at most genus one and at least one has genus zero. So this is quite useless for application to the global isotopy problem.

The following idea comes to the rescue. Away from the multiple components and from the singularities of the reduced components not much happens in the convergence $B_t \to C$: In a tubular neighborhood B_t is the graph of a function for any t sufficiently close to t_0 and this convergence is just a convergence of functions. So we can safely replace this part by some other (part of a) pseudo-holomorphic curve, for any t, and prove the isotopy lemma with this replacement made. By this one can actually achieve that each reduced component is a sphere, see below. If C_a is a sphere it contributes one less to the left-hand side than to the right-hand side of (9.1). So reduced components do not matter! For multiple components the right-hand side receives an additional $(m_a - 1)c_1(M) \cdot C_a$ that has to be balanced with the genus contribution $g(C_a)$ on the left-hand side.

Here is the precise formulation of the Isotopy Lemma from [SiTi3].

Lemma 9.2 *[SiTi3] Let $p : (M, J) \to \mathbb{CP}^1$ be a pseudo-holomorphic S^2-bundle. Let $\{J_n\}$ be a sequence of almost complex structures making p pseudo-holomorphic. Suppose that $C_n \subset M$, $n \in \mathbb{N}$, is a smooth J_n-holomorphic curve and that*

$$C_n \overset{n\to\infty}{\longrightarrow} C_\infty = \sum_a m_a C_{\infty,a}$$

in the $\mathscr{C}^0$-topology, with $c_1(M) \cdot C_{\infty,a} > 0$ for every a and $J_n \to J$ in $\mathscr{C}^{0,\alpha}_{\mathrm{loc}}$. We also assume:

(∗) If $C' = \sum_a m'_a C'_a$ is a non-zero J'-holomorphic cycle $\mathscr{C}^0$-close to a subcycle of $\sum_{m_a>1} m_a C_{\infty,a}$, with $J' \in \mathscr{J}$, then

$$\sum_{\{a|m'_a>1\}} \big(c_1(M) \cdot C'_a + g(C'_a) - 1\big) < c_1(M) \cdot C' - 1.$$

Then any J-holomorphic smoothing $C^\dagger_\infty$ of C_∞ is symplectically isotopic to some C_n. The isotopy from C_n to $C^\dagger_\infty$ can be chosen to stay arbitrarily close to C_∞ in the $\mathscr{C}^0$-topology, and to be pseudo-holomorphic for a path of *almost complex structures that stays arbitrarily close to J in $\mathscr{C}^0$ everywhere, and in $\mathscr{C}^{0,\alpha}_{\mathrm{loc}}$ away from a finite set.*

In the assumptions $\mathscr{C}^0$-convergence $C_n \to C_\infty$ is induced by $\mathscr{C}^0$-convergence inside the space of stable maps.

Using the genus formula one can show easily that the degree bounds in the theorem imply Assumption $(*)$ in the Isotopy Lemma, see [SiTi3], Lemma 9.1.

9.4 Sketch of Proof

We want to compare two different smoothings of the pseudo-holomorphic cycle C_∞, one given by C_n for large n and one given by some J-holomorphic smoothing, for example constructed via first deforming to a nodal curve and then smoothing the nodal curve, as suggested above. There is only one general case where we know how to do this, namely if J is integrable locally around the cycle, see Sect. 5. But J generally is not integrable and we seem stuck.

Step 1: Make J integrable around $|C|$. On the other hand, for the application of the Isotopy Lemma to symplectic geometry we are free to change our almost complex structures within a $\mathscr{C}^0$-neighborhood of J. This class of almost complex structures allows a lot of freedom! To understand why recall the local description of almost complex structures with fixed fiberwise complex structure along $w = \text{const}$ via one complex-valued function b:

$$T^{0,1}_M = \mathbb{C} \cdot (\partial_{\bar z} + b\partial_w) + \mathbb{C} \cdot \partial_{\bar w}.$$

The graph $\Gamma_f = \{(z, f(z))\}$ of a function f is pseudo-holomorphic with respect to this almost complex structure iff

$$\partial_{\bar z} f(z) \;=\; b(z, f(z)).$$

Thus the space of almost complex structures making Γ_f pseudo-holomorphic is in one-to-one correspondence with functions b with prescribed values $\partial_{\bar z} f$ along Γ_f. On the other hand, the condition of integrability (vanishing of the Nijenhuis tensor) turns out to be equivalent to

$$\partial_{\bar w} b = 0.$$

Thus it is very simple to change an almost complex structure only slightly around a smooth pseudo-holomorphic curve to make it locally integrable; for example, one could take b constant in the w-direction locally around Γ_f.

It is then also clear that if we have a $\mathscr{C}^1$-convergence of smooth J_n-holomorphic curves $C_n \to C$ with $J_n \to J$ in $\mathscr{C}^0$ it is possible to find $\tilde J_n$ integrable in a fixed neighborhood of C such that C_n is $\tilde J_n$-holomorphic and $\tilde J_n \to \tilde J$ in $\mathscr{C}^1$.

For the convergence near multiple components of C a little more care shows that Hölder convergence $J_n \to J$ in $\mathscr{C}^{0,\alpha}$ is enough to assure sufficient convergence of the values of b on the various branches of C_n. Finally, near the singular points of $|C|$ one employs the local holomorphicity theorem of Micallef-White to derive:

Lemma 9.3 *([SiTi3], Lemma 5.4) Possibly after going over to a subsequence, there exists a finite set $A \subset M$, a $\mathscr{C}^1$-diffeomorphism Φ that is smooth away from A, and almost complex structures $\tilde{J}_n$, $\tilde{J}$ on M with the following properties:*

1. *p is $\tilde{J}_n$-holomorphic*
2. *$\Phi(C_n)$ is $\tilde{J}_n$-holomorphic*
3. *$\tilde{J}_n \to \tilde{J}$ in $\mathscr{C}^0$ on M and in $\mathscr{C}^{0,\alpha}_{\mathrm{loc}}$ on $M \setminus A$*
4. *$\tilde{J}$ is integrable in a neighborhood of $|C|$*

Thus we can now assume that J is integrable in a neighborhood of $|C|$, but the convergence $J_n \to J$ is only $\mathscr{C}^0$ at finitely many points.

Note that if the convergence $J_n \to J$ is still in $\mathscr{C}^{0,\alpha}$ everywhere we are done at this point! In fact, in the integrable situation we do have a smooth parametrization of deformations of holomorphic cycles when endowing the space of complex structures with the $\mathscr{C}^{0,\alpha}$-topology. So the whole difficulty in the Isotopy Problem stems from the fact that the theorem of Micallef-White only gives a $\mathscr{C}^1$-diffeomorphism rather than one in $\mathscr{C}^{1,\alpha}$ for some $\alpha > 0$.

Step 2: Replace reduced components by spheres. The next ingredient, already discussed in connection with (9.1), is to make all non-multiple components rational. To this end we use the fact, derived in Proposition 2.8, that any J-holomorphic curve singularity can be approximated by J-holomorphic spheres. Let $U \subset M$ be a small neighborhood of the multiple components of C union the singular set of $|C|$. Then from C_n keep only $C_n \cap U$, while the rest of the reduced part of C_n gets replaced by large open parts of J-holomorphic approximations by spheres of the reduced branches of C at the singular points. For this to be successful it is important that the convergence $J_n \to J$ is in $\mathscr{C}^{0,\alpha}$ rather than in $\mathscr{C}^0$, for the former implies $\mathscr{C}^{1,\alpha}$-convergence $C_n \to C_\infty$ near smooth, reduced points of C_∞. As this is true in our case it is indeed possible to extend $C_n \cap U$ outside of U by open sets inside J-holomorphic spheres.

There are two side-effects of this. First, the result $\tilde{C}_n$ of this process on C_n is not a submanifold anymore, for the various added parts of spheres will intersect each other, and they will also intersect C_n away from the interpolation region. There is, however, enough freedom in the construction to make these intersections transverse. Then the $\tilde{C}_n$ are nodal curves. Second, $\tilde{C}_n$ is neither J_n nor J-holomorphic. But in view of the large freedom in choosing the almost complex structures that we saw in Step 1 it is possible to perform the construction in such a way that $\tilde{C}_n$ is $\tilde{J}_n$-holomorphic with $\tilde{J}_n \to \tilde{J}$ in $\mathscr{C}^0$, and in $\mathscr{C}^{0,\alpha}_{\mathrm{loc}}$ on $M \setminus \tilde{A}$, and $\tilde{J}_n$, $\tilde{J}$ having the other properties formulated above.

Now assume the Isotopy Lemma holds for these modified curves and almost complex structures, so an isotopy exists between $\tilde{C}_n$ for large n and some smoothing of $\tilde{C}_\infty = \lim_{n\to\infty} C_n$. Here "isotopy" means an isotopy of nodal, pseudo-holomorphic curves, with the almost complex structure and the connecting family of pseudo-holomorphic curves staying close to J, in $\mathscr{C}^{0,\alpha}$ away from finitely many points where this is only true in $\mathscr{C}^0$. Then one can revert the process, thus replace the spherical parts by the original ones, and produce a similar isotopy of C_n with the given smoothing of C_∞.

Thus we can also suppose that the reduced parts of C are rational, at the expense of working with nodal curves rather than smooth ones in the isotopy. As we can mostly work with maps rather than subsets of M, the introduction of nodes is essentially a matter of inconvenience rather than a substantial complication. We therefore ignore this for the rest of the discussion and simply add the assumption that the reduced components of C are rational.

Step 3: Break it! Now comes the heart of the proof. We want to change C_n slightly, for sufficiently large n, such that we find a path of pseudo-holomorphic cycles connecting C_n with a J-holomorphic smoothing of C_∞. Recall the pair (m, δ) introduced above as a measure of how singular a pseudo-holomorphic cycle is. By induction we can assume that the Isotopy Lemma holds for every convergence of pseudo-holomorphic curves where the limit has smaller (m, δ). This implies that whenever we have a path of pseudo-holomorphic cycles with smaller (m, δ) then there is a close-by path of smooth curves, pseudo-holomorphic for the same almost complex structure at each time. Thus in trying to connect C_n with a J-holomorphic smoothing of C_n we have the luxury to work with pseudo-holomorphic cycles, as long as they are less singular than labelled by (m, δ). We achieve this by moving C_n along with appropriate point conditions that force an enhancement of singularities throughout the path.

We start with choosing $k \le c_1(M) \cdot C_\infty - 1$ points $x_1, \ldots, x_k$ on $|C_\infty|$ such that $k_a = c_1(M) \cdot C_{\infty,a} + g(C_{\infty,a}) - 1$ of them are general points on the component $C_{\infty,a}$. Then there is no non-trivial *equigeneric* J-holomorphic deformation of $|C_\infty|$ incident to these points, provided J is general for this almost complex structure and the chosen points. One can show that one can achieve this within the class of almost complex structures that we took J from, e.g. integrable in a neighborhood of $|C_\infty|$. With such choices of points and of J any non-trivial J-holomorphic deformation of C_∞ decreases (m, δ). This enhancement of singularities even holds if we perturb J in a general one-parameter family. By applying an appropriate diffeomorphism for each n we may assume that the points also lie on C_n, for each n.

For the rest of the discussion in this step we now restrict to the most interesting case $m > 0$, that is, C_∞ does have multiple components. Then Condition $(*)$ in the statement of the Isotopy Lemma implies that there even exists a multiple component $C_{\infty,b}$ of C_∞ such that

$$c_1(M) \cdot C_{\infty,b} + g(C_{\infty,b}) - 1 < c_1(M) \cdot C_{\infty,b} - 1.$$

Thus we are free to ask for incidence with one more point x without spoiling genericity. Now the idea is to use incidence with a deformation $x(t)$ of x to move C_n *away* from C_∞, uniformly with n but keeping the incidence with the other k points. The resulting C'_n then converge to a J-holomorphic cycle $C'_\infty \neq C_\infty$ incident to the k chosen points and hence, by our choice of points, having smaller (m, δ) as wanted.

The process of deformation of C_n incident to $x(t)$ and to the k fixed points works well if we also allow a small change of almost complex structure along the path to make everything generic – as long as (1) we stay sufficiently close to $|C_\infty|$ and (2) the deformation of C_n does not produce a singular pseudo-holomorphic cycle with the $k+1$ points unevenly distributed. This should be clear in view of what we already know by induction on (m, δ) about deformations of pseudo-holomorphic curves, smoothings and isotopy. A violation of (1) actually makes us happy because the sole purpose was to move C_n away from C_∞ slightly. If we meet problems with (2) we start all over with the process of choosing points etc. but only for one component $\hat{C}_n$ of the partial degeneration of C_n containing less than $c_1(M) \cdot \hat{C}_n - 1$ points. To keep the already constructed rest of the curve pseudo-holomorphic we also localize the small perturbation of J_n away from the other components. Because each time $c_1(M) \cdot \hat{C}_n$ decreases by an integral amount (2) can only be violated finitely many times, and the process of moving C_n away from C_∞ will eventually succeed.

This finishes the proof of the Isotopy Lemma under the presence of multiple components.

Step 4: The reduced case. In the reduced case we do not have the luxury to impose one more point constraint. But along a general path of almost complex structures incident to the chosen points non-immersions have codimension one and can hence be avoided. One can thus try to deform C_n along a general path $J_{n,t}$ of almost complex structures connecting J_n with J and integrable in a fixed neighborhood of $|C_\infty|$. This bridges the difference between $\mathscr{C}^0$-convergence and $\mathscr{C}^{0,\alpha}$-convergence of J_n to J. If successful it leads to a J'-holomorphic smoothing of C_∞ that falls within the smoothings we have a good parametrization for, and which hence are unique up to isotopy. The only problem is if for every n this process leads to pseudo-holomorphic curves moving too far away from C_∞. In this case we can again take the limit $n \to \infty$ and produce a J-holomorphic deformation of C_∞ with smaller (m, δ). As in the non-reduced case we are then done by induction.

References

[Ar] N. Aronszajn: *A unique continuation theorem for elliptic differential equations or inequalities of the second order*, J. Math. Pures Appl. **36** (1957), 235–239.

[Au1] D. Auroux: *Fiber sums of genus 2 Lefschetz fibrations*, Turkish J. Math. **27** (2003), 1–10.

[Au2] D. Auroux: *A stable classification of Lefschetz fibrations*, Geom. Topol. **9** (2005) 203–217.

[AuSm] D. Auroux, I. Smith: *Lefschetz pencils, branched covers and symplectic invariants*, this volume.

[Bl] Barlet, D.: Espace analytique réduit des cycles analytiques complexes compacts d'un espace analytique complexes de dimension finie. Sem. F. Norguet 1974/75 (Lect. Notes Math. **482**), 1–158.

[BtPeVe] W. Barth, C. Peters, A. van de Ven: *Compact complex surfaces*, Springer 1984.

[Bi] J. Birman: *Braids, links and mapping class groups*, Princeton University Press 1974.

[Ch] K. Chakiris: *The monodromy of genus two pencils*, Thesis, Columbia Univ. 1983.

[Do] S.K. Donaldson: *Lefschetz pencils on symplectic manifolds*, J. Differential Geom. **53** (1999) 205–236

[DoSm] S. Donaldson, I. Smith: *Lefschetz pencils and the canonical class for symplectic four-manifolds*, Topology **42** (2003) 743–785.

[Fu] T. Fuller: *Hyperelliptic Lefschetz fibrations and branched covering spaces*, Pacific J. Math. **196** (2000), 369–393.

[GiTr] D. Gilbarg and N. S. Trudinger: *Elliptic partial differential equations of second order*, Springer 1977

[Go] R. Gompf: *A new construction of symplectic manifolds*, Ann. Math. **142** (1995), 527–595.

[GoSt] R. Gompf, A. Stipsicz: 4-*manifolds and Kirby calculus*, Amer. Math. Soc. 1999.

[GrRe] H. Grauert, R. Remmert: *Theory of Stein spaces*, Springer 1979.

[GrHa] P. Griffiths, J. Harris: *Principles of Algebraic Geometry*, Wiley 1978.

[Gv] M. Gromov: *Pseudo-holomorphic curves in symplectic manifolds*, Inv. Math. **82** (1985), 307–347.

[HaMo] J. Harris, I. Morrison: *Moduli of curves*, Springer 1998.

[IvSh1] S. Ivashkovich, V. Shevchishin: *Pseudo-holomorphic curves and envelopes of meromorphy of two-spheres in* $\mathbb{CP}^2$, preprint `math.CV/9804014`.

[IvSh2] S. Ivashkovich, V. Shevchishin: *Structure of the moduli space in a neighbourhood of a cusp curve and meromorphic hulls*, Inv. Math. **136** (1998), 571–602.

[Ka] A. Kas: *On the handlebody decomposition associated to a Lefschetz fibration*, Pacific J. Math. **89** (1980), 89–104.

[KpKp] B. Kaup, L. Kaup: *Holomorphic functions of several variables*, de Gruyter 1983.

[La] F. Lárusson: *Holomorphic neighbourhood retractions of ample hypersurfaces*, Math. Ann. **307** (1997), 695–703.

[LiTi] J. Li, G. Tian: *Virtual moduli cycles and Gromov-Witten invariants of general symplectic manifolds*, in: Topics in symplectic 4-manifolds (Irvine, 1996), R. Stern (ed.), Internat. Press 1998.

[MD] D. McDuff: *The structure of rational and ruled symplectic four-manifolds*, J. Amer. Math. Soc. **3** (1990), 679–712.

[McSa1] D. McDuff and D. Salamon, *Introduction to symplectic topology*, Oxford Univ. Press 1995.

[McSa2] D. McDuff and D. Salamon, *J-holomorphic curves and symplectic topology*, Amer. Math. Soc. 2004.

[MiWh] M. Micallef, B. White: *The structure of branch points in minimal surfaces and in pseudoholomorphic curves*, Ann. of Math. **141** (1995), 35–85.

[OzSt] B. Ozbağçi, A. Stipsicz: *Noncomplex smooth 4-manifolds with genus-2 Lefschetz fibrations*, Proc. Amer. Math. Soc. **128** (2000), 3125–3128.

[Pn] P. Pansu: *Compactness*, in: Holomorphic curves in symplectic geometry, M. Audin, J. Lafontaine (eds.), 233–249, Birkhäuser 1994.

[PrWo] T. Parker and J. Wolfson, *Pseudo-holomorphic maps and bubble trees*, J. Geom. Anal. **3**(1993), 63–98.

[PsPi] U. Persson, H. Pinkham: *Some examples of nonsmoothable varieties with normal crossings*, Duke Math. J. **50** (1983), 477–486.

[RuTi] Y.B. Ruan and G. Tian, *A mathematical theory of quantum cohomology,* J. Diff. Geom. **42**(1995), 259–367.

[Sh] V. Shevchishin: *Pseudoholomorphic curves and the symplectic isotopy problem*, preprint `math.SG/0010262`, 95 pp.

[Si] B. Siebert: *Gromov-Witten invariants for general symplectic manifolds*, preprint `dg-ga 9608005`.

[SiTi1] B. Siebert, G. Tian: *On hyperelliptic symplectic Lefschetz fibrations of four manifolds*, Communications Cont. Math. **1** (1999), 255–280

[SiTi2] B. Siebert, G. Tian: *Weierstraß polynomials and plane pseudo-holomorphic curves*, Chinese Ann. Math. Ser. B **23** (2002), 1–10.

[SiTi3] B. Siebert, G. Tian: *On the holomorphicity of genus-two Lefschetz fibrations*, Ann. of Math. **161** (2005), 955–1016.

[Sk] J.-C. Sikorav: *The gluing construction for normally generic J-holomorphic curves*, in: Symplectic and contact topology: interactions and perspectives (Toronto, 2001), 175–199, Y. Eliashberg et al. (eds.), Amer. Math. Soc. 2003.

[Sm1] I. Smith: *Symplectic geometry of Lefschetz fibrations*, Ph.D. thesis, Oxford University, 1998.

[Sm2] I. Smith: *Lefschetz fibrations and divisors in moduli space*, Geom. Topol. **5** (2001), 579–608.

[SU] J. Sacks and K. Uhlenbeck, *The existence of minimal immersions of 2 spheres*, Ann. of Math. **113** (1981), 1–24.

[Te] M. Teicher: *Braid groups, algebraic surfaces and fundamental groups of complements of branch curves*, in: Algebraic geometry – Santa Cruz 1995, 127–150, J. Kollár et al. (eds.), Amer. Math. Soc. 1997.

[Ti] G. Tian: *Quantum cohomology and its associativity*, in: Current developments in mathematics (Cambridge, 1995), 361–401, Internat. Press 1994.

[Ve] I. N. Vekua: *Generalized analytic functions*, Pergamon 1962.

[Ye] R. Ye, *Gromov's compactness theorem for pseudo-holomorphic curves*, Trans. Amer. Math. Soc. **342** (1994), 671–694.