The Virtual Moduli Cycle

Dusa McDuff*
State University of New York at Stony Brook
(dusa@math.sunysb.edu)
for the Northern California Symplectic Geometry seminar

December 3, 1998, revised Jan 14, 1999, with added footnotes Dec 1999 and a few additional remarks in 2004

This article is an attempt to describe one possible construction of the virtual moduli cycle that is used as a tool in the construction of Gromov–Witten invariants for a general symplectic manifold. There are many different versions of this construction. Here I will in the main follow Liu–Tian [LiuT1,2] since their approach (when modified by an idea of Seibert's) seems to involve the least amount of analysis. However, it does involve quite a bit of topology, some of which they only outline. The aim here is to flesh out their picture and to explain the different ingredients that are needed to make the construction work. We do not try to give full proofs, nor do we work out all the ideas in full generality.

Liu–Tian work in the category of partially smooth spaces (spaces with two topologies), and represent the virtual moduli cycle by the zero set of a suitable multi-section of a multi-bundle. We show explicitly how to assemble this zero set into an object that we call a branched labelled pseudomanifold. Our theory is quite general, and suggests that every finite dimensional orbifold has a "resolution" of this form. An example is worked out in §4.4.

Other approaches to this question have been developed by Fukaya–Ono [FO], Li–Tian [LiT], Ruan [R], Seibert [Sb] and Hofer–Salamon [HS]. We shall restrict here to curves of genus 0 but similar considerations apply for curves of higher genus.

I wish to thank Salamon for explaining to me that the virtual moduli cycle should be thought of as a branched manifold, Polterovich for his many penetrating questions and useful ideas, Siebert, Tian, Fukaya, Liu and Ruan for informative discussions about general Gromov–Witten invariants, and Olga Buse and Haydee Herrera–Guzman for helpful comments on an earlier draft.

1 Gromov–Witten invariants

Given a symplectic manifold (X, ω) with compatible almost complex structure J, denote by $\mathcal{M}_{0,k}(X, J, A)$ the space of (unparametrized) J-holomorphic spheres in class A with k marked points. Elements of this space are equivalence classes $[h, z_1, \ldots, z_k]$ of elements $(h, z_1, \ldots, z_k)$, where $h : S^2 \to X$ is a somewhere injective (i.e. nonmultiply covered) J-holomorphic map and the z_i are distinct points in S^2. Elements $(h, z_1, \ldots, z_k)$ and $(h', z'_1, \ldots, z'_k)$ are equivalent if there is $\gamma \in \mathrm{PSL}(2, \mathbb{C})$ such that $h' = h \circ \gamma$, and $z_j = \gamma(z'_j)$ for all j.

*Partially supported by NSF grant DMS 9704825.

When J is generic, $\mathcal{M}_{0,k}(X,J,A)$ is a manifold of dimension

$$\dim(X) + 2c_1(TX)(A) + 2k - 6.$$

It is usually not compact. However in many cases one can show that the evaluation map

$$ev : \mathcal{M}_{0,k}(X,J,A) \to X^k : \qquad [h, z_1, \ldots, z_k] \to (h(z_1), \ldots, h(z_k))$$

does represent a well defined element in homology $H_*(X^k)$ that is independent of the choice of J and can be used to define the Gromov–Witten invariants by the following procedure.

Take H_* to mean integral homology modulo torsion. Then, given homology classes $v_1, \ldots, v_k \in H_*(X)$ satisfying the dimension condition

$$\dim(X) + 2c_1(TX)(A) + 2k - 6 + \sum_j \dim(v_j) = k \dim(X),$$

the Gromov–Witten invariant $n_X(v_1, \ldots, v_k; A)$ is defined to be the intersection number of the class represented by ev with the homology class $v_1 \times \cdots \times v_k$ in $H_*(X^k)$.

This approach is described in Ruan–Tian [RT] and McDuff–Salamon [MS], for example, and basically works when X does not contain J-holomorphic spheres in classes B with $c_1(B) < 0$. The difficulty with it is that the presence of these "negative" spheres could in principle cause there to be so many elements "on the boundary" of $\mathcal{M}_{0,k}(X,J,A)$ that ev does not represent a homology class. (For a more detailed discussion of this, see [S].) To deal with this problem, one must first understand how to compactify $\mathcal{M}_{0,k}(X,J,A)$, and then try to construct from what might be a rather complicated compactification something that can represent a homology class. This class should of course be independent of all choices made and in particular of J. Here I will try to show how to construct from the compactification $\overline{\mathcal{M}}_{0,k}(X,J,A)$ a cycle (called the virtual moduli cycle) that represents the desired homology class. Sometime this cycle is called a *regularization* of the moduli space. Other authors, for example Seibert, just construct a homology class. But in all approaches, one constructs an element of *rational* rather than integral homology.

The paper is organized as follows. First I will describe the compactification of $\mathcal{M}_{0,k}(X,J,A)$, i.e. the space $\overline{\mathcal{M}}_{0,k}(X,J,A)$ of J-holomorphic stable maps. §2 describes its elements, and §3 shows how one can use the process of gluing in the domain to give it a topology. Roughly speaking, the idea now is that one can give an orbifold structure to a neighborhood $\mathcal{W}$ of $\overline{\mathcal{M}}$ in the space of all stable maps. The delbar operator $\overline{\partial}_J$ is a Fredholm section of a bundle over $\mathcal{W}$ and what one has to do is define a finite dimensional space R of allowed perturbations of $\overline{\partial}_J$. Then the zero set of a generic perturbation of $\overline{\partial}_J$ should represent the desired cycle.

However, it is not easy to put an orbifold structure on $\mathcal{W}$, at least in the category of Banach manifolds. Seibert [Sb] has succeeded in doing this,[1] but all other authors make do with less structure. Nor is it easy to get one's hands on the perturbation space R: this can be defined locally on $\mathcal{W}$ but there is a problem with fitting the local perturbations together. What we do in §4 is describe explicitly how to use the weaker structure developed by Liu–Tian (of multi-bundles and multi-folds) to define a geometric object that represents a rational homology class. This object is called a branched pseudo-manifold and is constructed from a Fredholm section of a multi-bundle over a multi-fold. Finally in §5 we sketch Liu–Tian's arguments showing that the constructions described in §4 are indeed possible in the present context.

This whole process should be thought of as an existence result for these invariants. It is extremely difficult to use the definition given here for direct calculations. Rather one uses it to establish formal properties of the invariants which are then exploited to do calculations by other means. We sketch one nontrivial example in §5.3 below. See also [Mc2].

[1] Added Jan 00: apparently there are still some unresolved problems with his arguments.

2 Stable maps and curves

It is a remarkable recent discovery (essentially due to Kontsevich) that the compactification $\overline{\mathcal{M}}_{0,k}(X,J,A)$ of $\mathcal{M}_{0,k}(X,J,A)$ can be described as the space of all J-holomorphic genus 0 stable maps. The elements of this space are equivalence classes $\tau = [\Sigma, h, z_1, \dots, z_k]$ of tuples $(\Sigma, h, z_1, \dots, z_k)$. Here the domain Σ is a connected union of components $\Sigma_i, i = 0, \dots, p$, each of which has a given identification with S^2. (Note that we consider Σ to be a topological space: the labelling of its components is a convenience and not part of the data.) The intersection pattern of the components can be described by a tree graph T with $p+1$ vertices, where each edge corresponds to an intersection point of the components represented by its vertices. No more than two components meet at any point. There are also k marked points $z_1, \dots, z_k$ placed anywhere on Σ except at an intersection point of two components. The positions of the marked points are often represented on the graph by labelled spokes emanating from the relevant vertex. These spokes are half-open intervals, i.e. they have no second endpoint.

The restriction h_i of the map h to Σ_i is assumed to be J-holomorphic and we impose the *stability condition* that h_i is nonconstant (though possibly a multiple covering) unless Σ_i contains at least 3 special points. (By definition, *special points* are either points of intersection with other components or marked points.) Also $h_*([\Sigma]) = \sum_i (h_i)_*[\Sigma_i] = A$. Finally, we divide out by all holomorphic reparametrizations. In other words if $\gamma : \Sigma \to \Sigma$ is a holomorphic map such that $\gamma(z_j') = z_j$ for all j we consider $(\Sigma, h \circ \gamma, z_1', \dots, z_k')$ and $(\Sigma, h, z_1, \dots, z_k)$ to be equivalent. Note that when $k \le 2$ the group of reparametrizations for which $z_j' = z_j$ for all j has (real) dimension at least 2. (The other reparametrizations — those with $z_j' \ne z_j$ — in some sense do not count since they can be gotten rid of by normalizing the choice of z_j.) This defines the elements τ of $\overline{\mathcal{M}} = \overline{\mathcal{M}}_{0,k}(X,J,A)$. Components of Σ on which h is constant are often called *ghosts*, and those containing fewer than 3 special points are called unstable. The stability condition means that unstable components are never ghosts.

Each such τ has a finite automorphism group $\Gamma_\tau = \mathrm{Aut}([\Sigma, h, z_1, \dots, z_k])$. Given a representative $(\Sigma, h, z_1, \dots, z_k)$ of τ we may identify Γ_τ with the group of all holomorphic maps $\gamma : \Sigma \to \Sigma$ such that $h \circ \gamma = h$, and $\gamma(z_j) = z_j$ for all j. For most elements this group is trivial, and it is always finite because of the stability condition. It may be nontrivial if any of the h_i are multiple coverings, or if two different h_i have the same image. (Note that γ might permute the components of Σ.) The presence of these automorphism groups means that in many cases the nicest structure that $\overline{\mathcal{M}}_{0,k}(X,J,A)$ can have is that of an orbifold: see Example 3.3.

Example 2.1 Let Σ have three components, with Σ_2 and Σ_3 both intersecting Σ_1 and let z_1 be a marked point on Σ_1. Then we can allow h_1 to be constant without violating stability. If in addition h_2, h_3 have the same image curve, there is an automorphism that interchanges Σ_2 and Σ_3. Since nearby stable maps do not have this extra symmetry, $\tau = [\Sigma, h, z_1]$ is a singular point in its moduli space. However, because marked points are labelled, there is no such automorphism if we put one marked point z_2 on Σ_2 and another z_3 at the corresponding point on Σ_3, i.e. so that $h_2(z_2) = h_3(z_3)$. One can also destroy this automorphism by adding just one marked point z_0 to $[\Sigma, h, z_1]$ anywhere on Σ_2 or Σ_3. □

There is a special case in which X is a single point. All maps are now trivial, and the resulting space $\overline{\mathcal{M}}_{0,k}(pt, J, A)$ (usually written $\overline{\mathcal{M}}_{0,k}$) is known as the Mumford–Deligne compactification of the moduli space of marked spheres. Its elements are called *stable curves* of genus zero. The stability condition now says that every (nondegenerate) component of Σ contains at least 3 special points. When $k < 3$ the space $\overline{\mathcal{M}}_{0,k}$ reduces to a single point, that for consistency is taken to be represented by the degenerate Riemann surface consisting of a single point. Since the Mobius group $\mathrm{PSL}(2, \mathbb{C})$ acts triply transitively on S^2, the space $\overline{\mathcal{M}}_{0,3}$ also is a single point, but now this is represented by a sphere with 3 fixed marked points $0, 1, \infty$.

Example 2.2 $\overline{\mathcal{M}}_{0,4}$ is the compactification of the space $\mathcal{M}_{0,4}$ of 4-tuples $\{0, 1, \infty, w\}$ in S^2. The question here is to understand how to represent the limit of a sequence $\{0, 1, \infty, w_n\}$ as w_n converges to one of the other marked points, say 0. Since the marked points are assumed distinct, it is not allowed for w to be superimposed on 0. But these two points can occupy their own component of Σ. Thus the limit is $[\Sigma, z_1, z_2, z_3, z_4]$ where Σ is the one point union of two copies of S^2, one of which contains the points z_1, z_4 and the other z_2, z_3. (The points are paired up this way because $z_1 = 0, z_4 = w$.) In §3.1 below we show how to use gluing to generate a neighborhood of this end in $\mathcal{M}_{0,4}$ starting from the limit point $[\Sigma, z_1, z_2, z_3, z_4]$. It will follow that $\overline{\mathcal{M}}_{0,4}$ is a manifold diffeomorphic to S^2.[2]

Observe that in this limit the pairs z_1, z_4 and z_2, z_3 appear symmetrically. This seems wrong at first glance because, after all, it is $w_n = z_4$ that is moving towards $0 = z_1$ while the other points are fixed. However, the set $\{0, 1, \infty, w_n\}$ is equivalent to $\{\gamma(0), \gamma(1), \gamma(\infty), \gamma(w_n)\}$ for any $\gamma \in \mathrm{PSL}(2, \mathbb{C})$, and if we choose γ_n so that $\gamma_n(0) = 0$, $\gamma_n(1) = 1$ and $\gamma_n(w_n) = \infty$ then $z_3 = \gamma_n(\infty)$ will converge to $z_2 = 1$. □

It is usually not hard to see how to parametrize the limit of a converging sequence of stable maps when these have genus 0. For example, the limit of a sequence $\tau_n = [\Sigma, h, z_1, z_{2,n}]$ in which $z_{2,n}$ is converging to z_1 (which is fixed) is a stable map with a new bubble component Σ_b attached to the point of Σ where z_1 was, and with the two marked points on Σ_b. Since h is independent of n the new component is a ghost, i.e. h maps it to a single point.

Identifying the limit gets more tricky when one considers maps of higher genus. The following example in genus 1 is taken from Fukaya–Ono [FO] Example 10.5.

Example 2.3 Consider a sequence of curves C_n in $\mathbb{C}P^2$ each made up of a conic and a line intersecting at points x_n, y_n, and suppose that x_n converges to y_n as $n \to \infty$. Each C_n is the image of a stable map $\tau_n = [\Sigma, h_n]$, where Σ has two components of genus 0 that intersect each other in two distinct (and fixed) points w_1, w_2. Since $x_n = h_n(w_1)$ and $y_n = h_n(w_2)$ have the same limit there must be a point on each component that is different from w_1, w_2 at which the derivative of h_n is blowing up. In other words, a bubble forms on each component at some point other than the intersection points w_1, w_2. (For a description of how this happens, see [MS] §4.3.) Therefore the domain of the limit $[\Sigma_\infty, h]$ has 4 components, one mapping to the line, another (disjoint from the first) mapping to the conic, and another two that are ghosts that each intersect the other twice and meet one of the other components. □

2.1 The space $\mathcal{W}$

The virtual moduli cycle is made up of stable maps that satisfy a perturbed Cauchy–Riemann equation. Therefore its elements belong not to $\overline{\mathcal{M}}$ itself but to a suitably small neighborhood $\mathcal{W}$ of $\overline{\mathcal{M}}$ in the space of all (not necessarily holomorphic) stable maps. Below we will define some open neighborhoods of the points τ in $\overline{\mathcal{M}}$, and $\mathcal{W}$ will be taken to be their union. Observe that $\mathcal{W}$ has a natural *coarse* stratification in which each stratum consists of elements $\tau = [\Sigma_\tau, h_\tau, z_1, \dots, z_k]$ where the domain has fixed topological type, where the marked points $z_1, \dots, z_k$ are on specified components of Σ_τ, and where the decomposition $A = \sum A_i$ is given. Such a stratum $\mathcal{W}^D$ can be described by the decorated tree D in which the vertices of the corresponding spoked tree are labelled with the classes A_i. We will also sometimes use the *fine* stratification in which the isomorphism class of Γ_τ is also fixed on each stratum.

In order to put a topology on $\mathcal{W}^D$ we must choose a local Γ_τ-equivariant slice (or cross-section) for the action of the reparametrization group on the space of parametrized maps so that the elements of

[2] In fact, all the spaces $\overline{\mathcal{M}}_{0,k}$ are manifolds: see [MS2] for example.

small subsets in $\mathcal{W}^D$ can be considered as parametrized maps. To this end, given a representative[3] $\tilde{\tau} = (\Sigma, \tilde{h}, z_1, \dots, z_k)$ of τ, choose a minimum number of extra points $w_1, \dots, w_\ell$ on Σ so that the domain $[\Sigma, z_1, \dots, z_k, w_1, \dots, w_\ell]$ is stable, i.e. so that each component of Σ has at least 3 special points, and then choose for each j a small open codimension-2 disc $\mathbf{H}_j$ in X that meets $\tilde{h}(\Sigma)$ transversally in $\tilde{h}(w_j)$. For this to be possible the w_j should be generic points on Σ. We also require that the set of added points $\{w_1, \dots, w_\ell\}$ be as invariant as possible under the action of the finite group Γ_τ. This requirement will be spelled out later. Then define

$$\widetilde{U}_\tau^D = \{\tilde{\tau}' = (\Sigma, \tilde{g}, z_1', \dots, z_k') : \tilde{g}(w_j) \in \mathbf{H}_j, ||\tilde{g} - \tilde{h}|| + \sum_i d(z_i, z_i') < \varepsilon\},$$

where d is the distance function in Σ.

The norm $||\cdot||$ used on the space of maps should be stronger than the C^1-topology, since we need the elements in $\widetilde{U}_\tau^D$ to be C^1-close to $\tilde{h}$ to ensure that they meet $\mathbf{H}_j$ transversally. For example the Sobolev norm $L^{2,p}$ with $p > 2$ would do.[4] Moreover, we assume that ε is so small that $[\tilde{g}(\Sigma_j)] = A_j$ for all j.

If the group Γ_τ is trivial, the projection $\pi_\tau : \widetilde{U}_\tau^D \to \mathcal{W}^D$ that takes the point $\tilde{\tau}'$ to its equivalence class τ' is injective (for small enough ε) so that τ has a neighborhood in $\mathcal{W}^D$ modelled on an open subset in a Banach space. If Γ_τ is nontrivial, we must extend its action to a linear action on $\widetilde{U}_\tau^D$ in such a way that a neighborhood U_τ of τ in $\mathcal{W}^D$ can be identified with the quotient $\widetilde{U}_\tau^D/\Gamma_\tau$.

First of all, consider the case when $\Sigma = S^2$ and there are two marked points z_1, z_2, so that we have to choose one point w_1. In this case, $\tilde{h}$ must factor through a holomorphic map $\tilde{f} : S^2 \to S^2$ of degree $n > 1$ and Γ_τ is a cyclic group generated by a rotation γ that fixes z_1, z_2 and such that $\tilde{f} \circ \gamma = \tilde{f}$. Choose w_1 so that the set $\tilde{f}^{-1}(\tilde{f}(w_1))$ contains n distinct points and choose disjoint little discs about each of these points that are permuted by γ. By construction, any element $\tilde{g} \in \widetilde{U}_\tau^D$ is such that $\tilde{g}(w_1) \in \mathbf{H}_1$. Moreover, for small enough ε, $\tilde{g}^{-1}(\mathbf{H}_i)$ is a collection of n points, one in each of the little discs. (cf [LiuT1] Lemma 2.2). Therefore there is unique point w_1' in the little disc containing $\gamma(w_1)$ such that

$$g(w_1') \;\in\; \mathbf{H}_1,$$

and we define $\phi_g^\gamma \in \mathrm{PSL}(2, \mathbb{C})$ to be the unique element that fixes z_1, z_2 and takes w_1 to w_1'. Then set

$$\gamma \cdot \tilde{g} = \tilde{g} \circ \phi_g^\gamma.$$

It is not hard to check that

$$\gamma \cdot \tilde{g} \in \widetilde{U}_\tau^D, \qquad \phi_g^\gamma \circ \phi_{\gamma \cdot \tilde{g}}^\gamma = \phi_g^{\gamma^2},$$

so that this does give an action of Γ_τ. It is not hard to see that it is linear (for small enough ε.)[5] Moreover, a neighborhood of τ in $\mathcal{W}^D$ can be identified with the quotient

$$U_\tau^D = \widetilde{U}_\tau^D/\Gamma_\tau.$$

A triple $(\widetilde{U}_\tau^D, \Gamma_\tau, \pi_\tau)$ where $\pi_\tau : \widetilde{U}_\tau^D \to U_\tau^D$ is the projection is called a *local uniformizer* for τ in $\mathcal{W}^D$. Note that the automorphism group $\Gamma_{\tau'}$ of an element $\tau' \in U_\tau^D$ is isomorphic to the stabilizer subgroup $\mathrm{Stab}(\tilde{\tau}') \subset \Gamma_\tau$ of any lift $\tilde{\tau}'$ of τ' to $\widetilde{U}_\tau$.

[3] For clarity we will here write $\tilde{h} : \Sigma \to X$ to denote a particular map, reserving h to denote an equivalence class of maps.

[4] One could get away with a weaker norm (eg the $L^{1,p}$ norm) if one used Siebert's integral method of fixing the parametrizations.

[5] Added in Jan 00: I claimed more than necessary here: all we need is that the action is continuous. A sketch of the proof is given in §6 below.

A similar argument works if $\Sigma = S^2$ has one special point since Γ_τ is again a rotation group: one just has to pick w_1, w_2 to be generic points on different Γ_τ-orbits. If there are no special points, Γ_τ need no longer be a rotation group but it is not hard to adapt the argument. Again one picks the three points w_j to be generic points on different Γ_τ-orbits, chooses a Γ_τ-invariant set of little disjoint discs containing the w_j and, for each $\gamma \in \Gamma_\tau$, chooses ϕ_g^γ to be an appropriate approximation to γ.

Finally, consider arbitrary τ. The group Γ_τ is the semidirect product of a permutation group with a subgroup $\Gamma_{\tau,0}$ that acts trivially on the set of components of Σ. Further $\Gamma_{\tau,0}$ is a product $\prod \Gamma_i$ where Γ_i is a group of reparametrizations of Σ_i. We have seen how to choose the points w_j on any component Σ_i with nontrivial Γ_i. If the permutation group is nontrivial, we require further that the set $\{w_j\}$ be invariant under this group. Thus, we have the following result: see [LiuT1] Lemma 2.5.

Lemma 2.4 *There is a neighborhood $\mathcal{W}^D$ of $\overline{\mathcal{M}}^D$ in the space of all stable maps of type D that is covered by local uniformizers $(\widetilde{U}_\tau^D, \Gamma_\tau, \pi_\tau)$.*

In fact, these local uniformizers give $\mathcal{W}^D$ the structure of an orbifold: see §4.

3 Gluing

The problem now is to describe how the strata $\mathcal{W}^D$ fit together as the topological type of the domain changes. Here one must use gluing. If one forgets the maps and just considers the finite-dimensional spaces of stable curves, then a stratum $\mathcal{S}$ on which Σ has $k+1$ components has (real) codimension $2k$. We shall see in §3.1 that in this case there is a tubular neighborhood theorem that identifies a neighborhood of $\mathcal{S}$ with a neighborhood of the zero section in a certain orbibundle of gluing parameters $a_1, \dots, a_k$, where each $a_i \in \mathbb{C}$ is attached to one of the intersection points $\Sigma_j \cap \Sigma_k$ of the components of Σ. (There are precisely k of these intersection points since the components of Σ are attached according to a tree graph.) We describe the orbifold structure of $\mathcal{W}$ in §3.2.

3.1 Gluing the domains

We now describe how to glue in the space of stable curves, i.e. we will only consider the domain Σ together with its marked points. The discussion here is essentially taken from Fukaya–Ono [FO] §9. They assume that each domain Σ is provided with a metric that is flat in some neighborhood of the double points and on each component is in the conformal class corresponding to the given complex structure.

First consider the simplest situation when Σ has two components with $x_1 \in \Sigma_1$ identified to $x_2 \in \Sigma_2$, and with two marked points on each component. Then the gluing parameter space $\mathbb{C}$ should be identified with the vector space $T_{x_1}\Sigma_1 \otimes T_{x_2}\Sigma_2$. Let us write $\tau_i, 1 \le i \le 3$, for the stable curve of this form with z_i and z_4 on Σ_2 and the other two points z_ℓ, z_m on Σ_1. As explained in Example 2.2 the three elements τ_i should correspond to the three ends of the open stratum $\mathcal{M}_{0,4} = S^2 - \{0, 1, \infty\}$ in $\overline{\mathcal{M}}_{0,4}$. Therefore, for each i what we must do is show how to identify a small deleted neighborhood $V - \{0\}$ of 0 in $\mathbb{C}$ with a neighborhood of the corresponding end in $\mathcal{M}_{0,4}$, i.e. our aim is to construct an injective map

$$\Psi_i : V - \{0\} \longrightarrow \mathcal{M}_{0,4} : \qquad a \mapsto [\Sigma_a, z_1^a, \dots, z_4^a]$$

whose image is a neighborhood of the corresponding end of $\mathcal{M}_{0,4}$.

Given a gluing parameter $a \in V - \{0\} \subset \mathbb{C}$, set $r = \sqrt{|a|}$. For small r and $j = 1, 2$, let D_j be the 2-disc $D_{x_j}(r)$ with center x_j and radius r in Σ_j and put $U_j = \Sigma_j - \operatorname{int} D_j$. Think of Σ_j as $D_j \cup_{\iota_j} U_j$, where D_j and U_j are positively oriented discs (i.e. their orientation as a subset of Σ_j agrees with their orientation as a disc) and ι_j is an orientation reversing identification of their boundaries. For each gluing parameter $a \in V - \{0\}$ the new domain Σ_a should be $U_1 \cup U_2$. In order to make this precise we have to describe the way ∂U_1 is identified with ∂U_2 and then put a suitable metric on Σ_a.

Since the given metric on Σ is flat near the double points x_j we can use the exponential map to identify the discs D_j isometrically with the discs of radius r in the tangent spaces $T_{x_j}\Sigma_j$. The gluing map is then the restriction of a map

$$\Phi_a : T_{x_1}\Sigma_1 - \{0\} \longrightarrow T_{x_2}\Sigma_2 - \{0\}$$

that is defined by the requirement that

$$u \otimes \Phi_a(u) = a \in \mathbb{C} = T_{x_1}\Sigma_1 \otimes T_{x_2}\Sigma_2,$$

for all $u \in T_{x_1}\Sigma_1$. Thus in local coords Φ_a is the inversion $u \mapsto a/u$. This defines the space Σ_a, and the metric (and hence complex structure) is given as follows. Since

$$\Phi_a^*(|du|^2) = \left|\frac{a}{u^2}\right|^2 |du|^2,$$

$\Phi_a(|du|^2) = |du|^2$ on the circle $|u| = r$. Hence, we may choose a function $\chi_r : (0,\infty) \to (0,\infty)$ so that the metric $\chi_r(|u|)|du|^2$ is invariant by Ψ_a and so that $\chi_r(s) = 1$ when $s \geq 2r$, and then patch together the given metrics on $\Sigma_1 - D_{x_1}(2r)$ and $\Sigma_2 - D_{x_2}(2r)$ via $\chi_r(|u|)|du|^2$.

Finally, we must look at what happens to the marked points. Observe that Σ_a has four marked points, z_ℓ, z_m on U_1 and z_i, z_4 on U_2. We can identify the sets U_j with the disc D^2 so that $z_i = z_\ell = 0, z_4 = z_m = 1/2$. Our normalisation of $\mathcal{M}_4$ assumes that the points z_1, z_2, z_3 are fixed and equal to $0, 1, \infty$ respectively. For each a let $\gamma_a : \Sigma_a \to S^2$ be the unique holomorphic diffeomorphism that takes the marked points z_1, z_2, z_3 on Σ_a to $0, 1, \infty$. The question now is to understand what happens to $\gamma_a(z_4)$ as a varies in a neighborhood of 0 in the gluing parameter space $\mathbb{C}$.

Proposition 3.1 *The map $a \mapsto \{0, 1, \infty, \gamma_a(z_4)\}$ defines an orientation preserving diffeomorphism between a deleted neighborhood of 0 in $\mathbb{C}$ and a neighborhood of an end in $\mathcal{M}_{0,4}$.*

We will not attempt to prove this here but will give some plausibility arguments. Consider the case when z_1 and z_4 are on the same component of Σ. As $|a| \to 0$ the metric on Σ_a has an increasingly pinched neck which implies that $\gamma_a(z_4)$ tends to $\gamma(z_1) = 0$. Inspection shows that the image of this map is a full neighborhood of the end in $\mathcal{M}_{0,4} = S^2 - \{0, 1, \infty\}$ corresponding to the point 0. Hence all one has to check is that the map is injective and orientation preserving.

Here is an argument to show that orientation is preserved. For simplicity let us suppose that the induced metrics on U_j are radially symmetric and that $\ell = 1, m = 2, i = 3$. Then γ_a takes the points $0, 1/2$ on $U_1 = D^2$ to $0, 1$ on S^2 and the point 0 in $U_2 = D^2$ to ∞ on S^2 and is also radially symmetric. Let p be the point 1 on ∂U_1. This is identified to $a/p = a$ on ∂U_2. As a moves around 0 positively as seen from the deleted point x_2, it moves around ∂U_2 negatively with respect to the orientation of U_2. Therefore U_2 must be rotated positively to match a to the fixed point p. It follows that the point z_4 rotates positively round z_3.

This shows how to glue two components together. One can obviously use the same method to glue any number of components.

Proposition 3.2 *An element $\tau = [\Sigma, z_1, \ldots, z_k]$ in $\overline{\mathcal{M}}_{0,k}$ whose domain has $p + 1$ components has a neighborhood consisting of all stable curves obtained by first slightly moving the points $z_1, \ldots, z_k$ and then gluing with an arbitrary small set of gluing parameters $a_1, \ldots, a_p$.*

This is again a nontrivial result that we will not prove. Note that if any $a_j = 0$ one simply leaves that intersection alone. Therefore the induced stratification on a neighborhood of τ corresponds to the stratification of $\mathbb{C}^p$ given by the subspaces $V_I = \{a : a_j = 0, j \in I\}$ where $I \subset \{1, \dots, p\}$. Observe also that because the automorphism group Γ_τ acts on the space of gluing parameters $\overline{\mathcal{M}}_{0,k}$ has the structure of an orbifold.

In nice cases, the space $\overline{\mathcal{M}}_{0,k}(X, J, A)$ is also a finite dimensional orbifold. In the next example we show that the compactification $\overline{\mathcal{M}}_{0,0}(S^2, J, 2[S^2])$ of the space of self-maps of S^2 of degree 2 is a (degenerate) orbifold that is diffeomorphic to $\mathbb{C}P^2$ but all its points have automorphism group $\Gamma_\tau = \mathbb{Z}/2\mathbb{Z}$.

Example 3.3 The space $\overline{\mathcal{M}}_0 = \overline{\mathcal{M}}_{0,0}(S^2, J, 2[S^2])$ has two strata. The first, $\mathcal{S}_1$, consists of self-maps of S^2 of degree 2, and the second, $\mathcal{S}_2$, consists of maps whose domain has two components, each taken into S^2 by a map of degree 1. The equivalence relation on each stratum is given by precomposition with a holomorphic self-map of the domain. lt is not hard to check that each equivalence class of maps in $\mathcal{S}_1$ is uniquely determined by its two critical values (or branch points). Since these can be any pair of distinct points, $\mathcal{S}_1$ is diffeomorphic to the set of unordered pairs of distinct points in S^2. On the other hand there is one element σ_w of $\mathcal{S}_2$ for each point $w \in S^2$, the correspondence being given by taking w to be the image under h of the point of intersection of the two components.

If $\sigma_{\{x,y\}}$ denotes the element of $\mathcal{S}_1$ with critical values $\{x, y\}$, we claim that $\sigma_{\{x,y\}} \to \sigma_w$ when x, y both converge to w. To see this, let $h_{\{x,y\}} : S^2 \to S^2$ be a representative of $\sigma_{\{x,y\}}$ and let $\alpha_{\{x,y\}}$ be the shortest geodesic from x to y. (We assume that x, y are close to w.) Then $h^{-1}_{\{x,y\}}(\alpha_{\{x,y\}})$ is a circle $\gamma_{\{x,y\}}$ through the critical points of $h_{\{x,y\}}$. This is obvious if $h_{\{x,y\}}$ is chosen to have critical points at $0, \infty$ and if $x = 0, y = \infty$ since $h_{\{x,y\}}$ is then a map of the form $z \mapsto az^2$. It follows in the general case because Mobius transformations take circles to circles. Hence $h_{\{x,y\}}$ takes each component of $S^2 - \gamma_{\{x,y\}}$ onto $S^2 - \alpha_{\{x,y\}}$. If we now let x, y converge to w, we see that $\sigma_{\{x,y\}}$ converges to σ_w.

Hence $\overline{\mathcal{M}}_0$ is the quotient of $S^2 \times S^2$ by the involution $(x, y) \mapsto (y, x)$. This is well known to be $\mathbb{C}P^2$. Note that every element of $\mathcal{S}_1$ can be represented by a conjugate of $z \mapsto z^2$ and so have automorphism group of order 2. This remains true for elements of Σ_2: here the automorphism interchanges the two components of the domain.

Suppose now that one adds three marked points to the domain and looks at the subset of stable maps where these points z_0, z_1, z_∞ are taken to $\{0, 1, \infty\}$. This space, that we shall call $\overline{\mathcal{M}}_3$ is an 8-fold branched cover over $\overline{\mathcal{M}}_0$, where the inverse image of $\tau \in \overline{\mathcal{M}}_0$ is the set of ways of choosing which point in $h^{-1}(0)$, resp. $h^{-1}(1), h^{-1}(\infty)$, is called z_0, resp. z_1, z_∞. In [Mc1] Lemma 3.7 it is shown that this space can be identified with $\mathbb{C}P^2$ and that the branched covering map is:

$$[x, y, z] \mapsto [x^2, y^2, z^2].$$

Moreover the inverse image of the line $\mathcal{S}_2$ is a set of 4 lines in general position. These lines correspond to the 4 different ways in which 3 points can be arranged on the two components of the domain. Now all the groups Γ_τ are trivial. We will meet this space again in the example discussed in §5.3. □

3.2 The topology on $\mathcal{W}$.

The above process of gluing can also be used to describe a topology on $\mathcal{W}$. We have already seen how to define local uniformizers $(\widetilde{U}^D_\tau, \Gamma_\tau, \pi_\tau)$ for τ in its stratum $\mathcal{W}^D$. To deal with a full neighborhood, add a minimum number of marked points $(w_1, \dots, w_\ell)$ to Σ so that the curve $[\Sigma, z_1, \dots, z_k, w_1, \dots, w_\ell]$ is stable, and let V be a small neighborhood of 0 in the space of gluing parameters $\mathbb{C}^p$ for Σ as in Proposition 3.2.

Then for each $a \in V$ there is a well defined stable curve $[\Sigma_a, z_1^a, \dots, z_k^a]$ that is the result of gluing with parameters $a = (a_1, \dots, a_p)$, and a map

$$\psi_a : (\Sigma_a, z_1^a, \dots, z_k^a) \to (\Sigma, z_1, \dots, z_k)$$

that is injective except that it collapses the boundary of each disc ∂U_j to the intersection point of the relevant components in Σ. The neighborhood $\widetilde{U}_\tau$ of $\tilde{\tau} = (\Sigma, \tilde{h}, z_1, \dots, z_k)$ now contains all points $\tilde{\tau}' = (\Sigma_a, \tilde{h}', z_1', \dots, z_k')$ where $a \in V$, $\tilde{h}'$ is close to $\tilde{h} \circ \psi_a$, and the $\psi_a(z_j')$ are close to z_j. This space has an obvious topology. It is now not hard to extend the action of the symmetry group Γ_τ to $\widetilde{U}_\tau$ and to show that the quotient

$$U_\tau = \widetilde{U}_\tau / \Gamma_\tau$$

can be identified with a subset of $\mathcal{W}$. The set $\mathcal{W}$ is defined to be the union of all the U_τ, and its topology is generated by these sets. [6] Liu–Tian show that this topology is Hausdorff. Moreover $\overline{\mathcal{M}}$ is a compact subset of $\mathcal{W}$: see [LiuT1] Prop 4.1.

We will not try to put a smooth structure on $\mathcal{W}$ that is transverse to the strata. Instead we will, following Liu–Tian, consider $\mathcal{W}$ to be a space with two topologies (or a *partially smooth* space): see §4.1 below. Thus $\mathcal{W}$ is a (partially smooth) orbifold.

4 Branched pseudomanifolds

In this section we describe the topological nature of the virtual cycle. §4.1 discusses the partially smooth category and introduces the idea of a branched, labelled pseudomanifold. In §4.2 the concepts of multifold, multi-bundle and multi-section are defined and illustrated by examples. §4.3 shows how the zero set of a Fredholm multi-section can be assembled into branched, labelled pseudomanifold. Though the main application is infinite dimensional, the definitions all apply in finite dimensions. We show in §4.4 how to calculate the "resolution" and the Euler number of the tear drop orbifold.

4.1 Spaces with two topologies

We will work in the category of *partially smooth* topological spaces Y. Thus Y is a Hausdorff topological space that is the image of a continous bijection

$$i_Y : Y_{sm} \to Y,$$

where Y_{sm} is a finite union of open disjoint subsets, each of which is a smooth Banach manifold. (Another way to think of Y is as a space with two topologies.) A morphism $f : Y \to X$ between two objects in this category is a continuous map $f : Y \to X$ such that the induced map $f : Y_{sm} \to X_{sm}$ is smooth. Usually, we will be interested in the case when X is a manifold, i.e. when i_X is a homeomorphism, so that X can be identified with the manifold X_{sm}. Morphisms $f : Y \to X$ will be called partially smooth maps, or, for short, simply maps. The path components of Y_{sm} are sometimes called *strata* and sometimes *components*. If Y is infinite dimensional, strata that are open subsets of Y are called top strata.

Definition 4.1 A *closed oriented pseudomanifold* of dimension d is a compact Hausdorff partially smooth topological space Y such that one component of Y_{sm} is an oriented smooth d-dimensional manifold that is mapped by i_Y onto a dense open subset Y^{top} of Y. Moreover, all other components of Y_{sm} have dimension $\leq d - 2$. We write Y^{sing} for the image $Y - Y^{top}$ of these lower dimensional submanifolds.

[6] See §6 for more details.

It is easy to see that any partially smooth map f from such Y to a closed manifold X defines an element $f_*([Y])$ of $H_d(X, \mathbb{Q})$ by intersection. Indeed since we are working rationally, the class $f_*([Y])$ is defined by its intersections with classes given by smooth maps $g : Z \to X$ from oriented closed manifolds Z into X. If $\dim Z + d = \dim X$, then we can jiggle g to make it meet each component of the image of Y_{sm} transversally. Hence it will be disjoint from Y^{sing} (for reasons of dimension), and will meet $f(Y^{top})$ transversally in a finite number of points. Moreover, if $g_0, g_1 : Z \to X$ are both transverse to f in this way, and are joined by a homotopy $G : Z \times [0,1] \to X$ then we can jiggle G (fixing $g_i = G|_{Z\times i}$) so that it is also transverse to f. The important point now is that the image of G still does not meet $f(Y^{sing})$ since Y^{sing} has codimension ≥ 2. It follows easily that the intersection number $f \cdot g$ of f with g is well-defined and independent of choices. We now define $f_*([Y])$ to be the unique rational class such that

$$f_*([Y]) \cdot g_*([Z]) = f \cdot g, \quad \text{for all } \; g : Z \to X.$$

Thus, in the language of [MS1], the partially smooth map $f : Y \to X$ is a pseudocycle.

Definition 4.2 A *branched pseudomanifold* Y of dimension d is a partially smooth space $Y_{sm} \to Y$ where the components of Y_{sm} have dimension at most d. The components of dimension d are called M_i, those of dimension $(d-1)$ are called B_j, and $Y_{\leq k}$ denotes the union of all components of Y_{sm} of dimension $\leq k$. Write:

$$Y^{top} = \bigcup_i M_i, \quad B = \bigcup_j B_j, \quad Y^{sing} = Y - (Y^{top} \cup B) = Y_{\leq d-2}.$$

We assume that it is possible to give the set $M_i \cup_{j \in J_i} B_j$ the structure of a smooth manifold with boundary compatibly with its two topologies.[7] Here $j \in J_i$ if the closure $\overline{M}_i$ of M_i in Y meets B_j. Similarly, we assume $\overline{B_j} - B_j$ lies in Y^{sing}. The set B is called the branch locus.

One more ingredient is needed in order for a compact branched pseudomanifold to be a cycle, namely a suitable labelling of the top components. This is spelled out in the next definition.

Definition 4.3 We say that a branched pseudomanifold Y of dimension d is *labelled* if Y is compact and if its top components M_i are oriented and have positive rational labels $\lambda_i \in \mathbb{Q}$, assigned so that the following condition holds:

> for each $x \in B$, pick an orientation of $T_x B$ and divide the components M_i that have x in their closure into two groups I^+, I^- according to whether the chosen orientation on $T_x B$ agrees with the boundary orientation. Then we require:
>
> $$\sum_{i \in I^+} \lambda_i = \sum_{i \in I^-} \lambda_i.$$

For example, if Y has dimension 0, it is just a collection of oriented labelled points (which means in practice that the labels can be both positive and negative), and there are no compatibility conditions at the branch locus since this is empty. However, even if one is considering the 0-dimensional case, one needs to look at 1-dimensional branched cobordisms between them in order to prove that the homology class represented by Y is independent of choices. This is nicely explained in [S] §5.4, but note that Salamon's point of view on branched manifolds is somewhat different from ours.

[7] By "compatibly" we mean that the inclusion of this set into Y is continuous, and that the inclusion of any stratum M_i or B_j into this manifold with boundary is smooth.

Lemma 4.4 *Let Y be a (closed, oriented) branched and labelled pseudomanifold of dimension d. Then every partially smooth map f from Y to a closed manifold X defines a rational class $f_*([Y]) \in H_d(X)$.*

Proof: As above, let Z be a smooth manifold of dimension complementary to Y and $g : Z \to X$ be a smooth map. Then we can jiggle g so that it does not meet $f(Y^{sing})$ and so that it meets $\cup f(M_i)$ transversally in a finite number of points. We then set

$$f \cdot g = \sum_i \varepsilon_i \lambda_i,$$

where $\varepsilon_i = \pm 1$ is the appropriate sign. It is not hard to check that the condition imposed on the labels in Definition 4.3 guarantees that this number is independent of the jiggling: cf. [S] Lemma 5.11. The class $f_*[Y]$ is then defined as before. □

Our basic aim is to use the minimum of structure for the purpose at hand. Never-the-less, the definition of branched pseudomanifold includes more information on the strata of codimension ≥ 2 than is strictly necessary for the preceding lemma to hold. We could get by with something more like the pseudocycles in [MS] §7.1, i.e. we could simply ask that Y^{sing} be the smooth image of a (probably disconnected) manifold of dimension $\leq d-2$, This definition has its own problems since one would have to specify what "smooth" means. For instance, as is the case in our main example there might be an evaluation map *ev* from Y to a smooth manifold X, and one would ask that the composite map be smooth. Since in [Mc2] we also need information on the strata of codimensions 2 and 3, we chose the present approach.

Another point is that, as much as is possible, we are trying to avoid specifying exactly how the strata of Y_{sm} fit together. If Y is branched then one does need some information of the normal structure to the codimension 1 components, but as is shown by the previous lemma, one can often get away without any other restrictions. However, later we will need to consider the intersection of two partially smooth spaces (see §4.3), and in order for this to be well-behaved one does need more structure. This is the reason for the following definition.

Definition 4.5 We will say that a partially smooth space Y has *normal cones* if every stratum S in Y has a neighborhood $\mathcal{N}(S)$ in Y whose induced stratification is that of a cone bundle over a link. More precisely, there is a commutative diagram

$$\begin{array}{ccc} \mathcal{N}(S)_{sm} & \to & \mathcal{N}(S) \\ \pi' \downarrow & & \pi \downarrow \\ S_{sm} & \to & S \end{array}$$

where the maps π' and π are oriented locally trivial fibrations with fiber equal to a cone over a link L. Here L is partially smooth, and the cone $C(L)$ is just the quotient $L \times [0,1]/L \times \{0\}$, stratifed so that it is the union of the vertex (the image of $L \times \{0\}$) with the product strata in $L_{sm} \times (0,1]$. Moreover each stratum in $\mathcal{N}(S)_{sm}$ projects onto the whole of S and so is a locally trivial bundle over S with fiber equal to a component of $C(S)_{sm}$. In particular, we identify S with the section of $\mathcal{N}(S)$ given by the vertices of the cones.

This definition implies that any stratum S' whose closure intersects S must contain the whole of S in its closure. In fact, near S it must look like the cone over some stratum of L_{sm}. Any finite dimensional Whitney stratified space has this normal structure. For example, any finite dimensional orbifold has a stratification such that the isomorphism class of the automorphism group Γ_x at the point x is constant as x varies over each stratum, and it is easy to see that this stratification has normal cones as defined here. It is also not hard to use arguments similar to those in §3.2 to show that the neighborhood $\mathcal{W}$ of $\overline{\mathcal{M}}$ in the space of all stable maps has normal cones with respect to the fine stratification. Finally, the process of gluing J-holomorphic maps shows that the finite dimensional local solution sets $\widetilde{Z}^\nu_\tau$ considered in Lemma 5.1 are stratified with normal cones.

4.2 Multi-folds and multi-sections

In the example of interest to us, the space $\mathcal{W}$ itself can be given the structure of an orbifold in the partially smooth category. [8] However, it is hard to define a suitable global perturbation space, and hence the virtual moduli cycle, in this context. Different authors propose different solutions to this problem. Liu–Tian choose to minimize the needed analysis at the cost of introducing more topological complexity.

We begin by defining a generalization of the usual concept of orbifold that seems to be familiar to algebraic geometers as a "stack". We will call it a *multi-fold.* It is a certain kind of atlas (or covering) of a space $\mathcal{W}$ with two topologies that locally is an orbifold in the partially smooth category, in which the inclusions that relate one uniformizer to another in an orbifold are replaced by fiber products. We then show how to define multi-sections of a multi-bundle over a multi-fold. Our construction is taken from [LiuT1,2] but is actually fairly close to that of Fukaya–Ono: for example, they also consider something they call a multi-section. However, Liu–Tian's construction is topologically more transparent because they can restrict attention to multi-sections that are single-valued when pulled back to the desingularization of the multi-fold.

The theory described here could be generalized, but this would require quite a bit of work that we shall not do, since our aim is to understand the special case of the virtual moduli cycle. In this application, $\mathcal{W}$ is a neighborhood of the moduli space $\overline{\mathcal{M}}$ in the space of all stable maps. We assume that it is given the fine stratification so that the isomorphism class of the group Γ_τ is fixed on each stratum. This means that the strata are manifolds. If instead one used the coarse stratification, then a stratum could have some orbifold singularities.

All maps, spaces and group actions considered below are in this partially smooth category. Branching will not be an issue until we construct the finite-dimensional virtual cycle in §4.3. Recall that a stratum (or component of Y_{sm}) is called a top stratum if it is an open subset of Y.

Multi-folds

Consider a space $\mathcal{W}$ that is a finite union of open sets $U_i, i = 1, \dots, k,$, each with uniformizers $(\widetilde{U}_i, \Gamma_i, \pi_i)$ with the following properties. Each Γ_i is a finite group acting on $\widetilde{U}_i$ and the projection π_i is the composite of the quotient map $\widetilde{U}_i \to \widetilde{U}_i/\Gamma_i$ with an identification $\widetilde{U}_i/\Gamma_i = U_i$. The inverse image in $\widetilde{U}_i$ of each stratum in U_i is an open subset of a (complex) Banach space on which Γ_i acts complex linearly. For simplicity, we will assume that Γ_i acts freely on the points of the top strata in $\widetilde{U}_i$, and that the isomorphism class of the stabilizer subgroup $\mathrm{Stab}_i(\tilde{x}_i)$ of Γ_i is fixed as $\tilde{x}_i$ varies over a stratum of $\widetilde{U}_i$.[9]

To keep our construction as simple as possible we will also assume that the set of local uniformizers $\{(\widetilde{U}_i, \Gamma_i, \pi_i)\}$ can be enlarged to an orbifold atlas. In an orbifold, the local uniformizers satisfy the following compatibility conditions:

(i) for every $x \in U_i \cap U_j$ there is an index ℓ such that $x \in U_\ell \subset U_i \cap U_j$;

(ii) for $m = i, j$, there are injections $\iota_{m\ell} : \widetilde{U}_\ell \to \widetilde{U}_m$ and group homomorphisms $\lambda_m : \Gamma_\ell \to \Gamma_m$ such that

$$\lambda_m(\gamma) \cdot \iota_{m\ell}(\tilde{y}) = \iota_{m\ell}(\gamma(\tilde{y})),$$

[8] Added in Dec 99: as we point out in §6 this is not quite correct.

[9] The stabilizer group $\mathrm{Stab}_i(\tilde{x}_i)$ of a point $\tilde{x}_i \in \widetilde{U}_i$ is the subgroup of Γ_i that fixes $\tilde{x}_i$.

for all $\tilde{y} \in \widetilde{U}_\ell, \gamma \in \Gamma_\ell$.

Thus, for example, we are not permitting a situation in which $U_i = \widetilde{U}_i = S^1$ for $i = 1, 2$ and $\pi_i : S^1 \to S^1$ quotients out by the action of $\mathbb{Z}/2\mathbb{Z}$ when $i = 1$ and by $\mathbb{Z}/3\mathbb{Z}$ when $i = 2$, though one could extend the theory to cover this case. Note that we do not assume that the cover $\{U_i\}_{1 \le i \le k}$ itself satisfies condition (i) above.

For each subset $I = \{i_1, \dots, i_p\}$ of $\{1, \dots, k\}$ set

$$U_I = \cap_{j \in I} U_j, \quad U_\emptyset = \emptyset.$$

Let $\mathcal{N}$ be the set of all I for which $U_I \ne \emptyset$. For each $I \in \mathcal{N}$ define the group Γ_I to be the product $\prod_{j \in I} \Gamma_j$ and then define $\widetilde{U}_I$ to be the fiber product

$$\widetilde{U}_I = \{\tilde{x}_I = (\tilde{x}_j)_{j \in I} : \pi_j(\tilde{x}_j) = \pi_\ell(\tilde{x}_\ell) \in U_I \text{ for all } j, \ell\} \subset \prod_{j \in I} \widetilde{U}_j,$$

with the two topologies induced from $\prod \widetilde{U}_j$. Clearly, Γ_I acts on $\widetilde{U}_I$ and the quotient $\widetilde{U}_I/\Gamma_I$ can be identified with U_I. Further, it is not hard to check that each $\widetilde{U}_I$ is a partially smooth space, with strata that project to strata in U_I. Again, the isomorphism class of the stabilizer subgroup

$$\mathrm{Stab}_I(\tilde{x}_I) = \prod_{i \in I} \mathrm{Stab}_i(\tilde{x}_i)$$

of $\tilde{x}_I$ in Γ_I is constant on each stratum, and trivial at points of top strata.

In the restricted case we are now considering, the topological structure of the fiber product $\widetilde{U}_I$ can best be understood in terms of the notion of "local component" that was introduced by Liu–Tian in [LiuT2]. Given a point $\tilde{x}_I \in \widetilde{U}_I$, choose $i_0 \in I$ and a small open neighborhood $\widetilde{C}$ of $\tilde{x}_{i_0}$ in $\widetilde{U}_{i_0}$. Then a neighborhood of $\tilde{x}_I$ in $\widetilde{U}_I$ can be identified with the set

$$\{(\gamma_j \circ \iota_{j i_0}(\tilde{y}))_{i \in I} \ : \ \tilde{y} \in \widetilde{C}, \gamma_{i_0} = id, \gamma_j \in \mathrm{Stab}_j(\iota_{j i_0}(\tilde{y})), j \ne i_0\}.$$

This is a finite union of sets $\widetilde{C}(\tilde{x}_I, \gamma)$, one for each element γ in the group

$$\mathrm{Stab}_{I - i_0}(\tilde{x}_I) = \prod_{j \in I - i_0} \mathrm{Stab}_j(\tilde{x}_j).$$

Each such set is homeomorphic (in the partially smooth category) to $\widetilde{C}$. It is also easy to check that the germs[10] at $\tilde{x}_I$ of the sets $\widetilde{C}_\gamma(\tilde{x}_I)$ are independent of the choices of i_0 and $\widetilde{C}$, as is the isomorphism class of the "reduced" group $\mathrm{Stab}'_I(\tilde{x}_I) = \mathrm{Stab}_{I-i_0}(\tilde{x}_I)$: see [LiuT2] Lemma 4.3. These germs are called the local components of $\widetilde{U}_I$ at $\tilde{x}_I$ and we will denote them by $\langle \widetilde{C}_\gamma(\tilde{x}_I) \rangle$. Note that points $\tilde{x}_I$ in the top strata on $\widetilde{U}_I$ have a single local component since their stabilizer groups are trivial. Liu–Tian define the *desingularization* $\widehat{U}_I$ of $\widetilde{U}_I$ to be the union of all such local components:

$$\widehat{U}_I = \{(\tilde{x}_I, \langle \widetilde{C}_\gamma(\tilde{x}_I) \rangle) \ : \ \tilde{x}_I \in \widetilde{U}_I, \gamma \in \mathrm{Stab}'_I(\tilde{x}_I)\}.$$

[10] Two subsets C_1, C_2 that both contain the point x are said to have the same germ at x if there is a neighborhood U of x such that $C_1 \cap U = C_2 \cap U$.

$\widehat{U}_I$ is topologized so that a germ of neighborhood of $\widehat{U}_I$ at the point $(\tilde{x}_I, \langle \widetilde{C}_\gamma(\tilde{x}_I)\rangle)$ is homeomorphic to the local component $\langle \widetilde{C}_\gamma(\tilde{x}_I)\rangle$ itself. Thus the projection

$$proj\ :\ \ \widehat{U}_I \to \widetilde{U}_I$$

is locally a homeomorphism onto its image. The effect of considering $\widehat{U}_I$ is to get rid of the extra singularities of $\widetilde{U}_I$ that are introduced by passing to the fiber product: locally $\widehat{U}_I$ is homeomorphic to the initial sets $\widetilde{U}_j$.

Example 4.6 For $j = 1, 2$ let $U_j = [0, \infty) = \mathcal{W}$ be uniformized by

$$(\widetilde{U}_j = \mathbb{R}, \Gamma_j = \mathbb{Z}/2\mathbb{Z}, \pi_j)$$

where $\Gamma_j = \mathbb{Z}/2\mathbb{Z}$ acts by multiplication by -1. Then, $\widetilde{U}_{12}$ can be identified with the union of the two lines $\{y = \pm x\}$ in $\mathbb{R}^2$, with the 5 strata

$$\{(\pm x, \pm x) : x > 0\}, \ \ \{(0, 0)\}.$$

Moreover, $\widehat{U}_{12}$ is the disjoint union of these two lines. □

If $J \subset I$, there are projections

$$\pi^I_J : \widetilde{U}_I \to \widetilde{U}_J, \quad \lambda^I_J : \Gamma_I \to \Gamma_J,$$

where π^I_J quotients out by the action of the product group Γ_{I-J}. Observe that if $\pi^I_J(\tilde{y}_I) = \tilde{x}_J$ is in a top stratum, then the inverse image $(\pi^I_J)^{-1}(\tilde{x}_J)$ has $|\Gamma_{I-J}|$ points.

More generally, suppose that $\{V_I\}$ is a family of open sets that covers $\mathcal{W}$ and is such that $V_I \subset U_I$ for all I. (In the application V_I is in general much smaller than U_I and the sets $V_j, j = 1, \dots, k$, no longer cover $\mathcal{W}$: see Example 4.10.) Then define $\widetilde{V}_I \subset \widetilde{U}_I$ to be the inverse image of V_I under the map $\pi : \widetilde{U}_I \to U_I$. Clearly Γ_I acts on each $\widetilde{V}_I$, and the map $\pi : \widetilde{V}_I \to V_I$ simply quotients out by this action. Further the projections $\pi^I_J : \widetilde{U}_I \to \widetilde{U}_J$ restrict to give partially defined maps $\widetilde{V}_I \to \widetilde{V}_J$ with domain $(\pi^I_J)^{-1}(\widetilde{V}_J)$. We will say that such a collection

$$\widetilde{\mathcal{V}} = \{(\widetilde{V}_I, \Gamma_I, \pi^I_J, \lambda^I_J) : I \in \mathcal{N}\}$$

is a *multi-fold atlas* for $\mathcal{W}$ or a multi-fold, for short. The motivation for considering the subcover $\{V_I\}$ will not be apparent until we start constructing multi-sections as in Example 4.7.

Multi-bundles

We now want to define a *multi-bundle* $\tilde{p} : \widetilde{\mathcal{E}} \to \widetilde{\mathcal{V}}$.[11] The example of interest to us is induced from a collection of local orbibundles over the cover U_i and so we will restrict to that case here, though one could make a more general definition that refers only to the atlas $\widetilde{\mathcal{V}}$. We start with a space

$$\mathcal{E} = \bigcup_i E_i$$

[11]Both the definition presented here and the example are essentially due to Polterovich. He realised that at this point the original presentation of Liu–Tian, that appeared in the early versions of [LiuT1] on which this paper is mostly based, was inadequate.

that has two topologies, and a map $p : \mathcal{E} \to \mathcal{W}$ with the property that each set $E_i = p^{-1}(U_i)$ has a local uniformizer $(\widetilde{E}_i, \Gamma_i, \Pi_i)$ such that the following diagram commutes:

$$\begin{array}{ccc} \widetilde{E}_i & \stackrel{\Pi_i}{\to} & E_i \\ \tilde{p}\downarrow & & p\downarrow \\ \widetilde{U}_i & \stackrel{\pi_i}{\to} & U_i. \end{array}$$

Here we require $\tilde{p} : \widetilde{E}_i \to \widetilde{U}_i$ to be a Γ_i-equivariant map that restricts over each stratum of $\widetilde{U}_i$ to a locally trivial vector bundle.[12] Thus the fiber $\widetilde{F}(\tilde{x}_i)$ of $\tilde{p}$ at each point $\tilde{x}_i$ is a vector space, but there is no natural way of identifying one with another if they lie over points in different strata. In the application, the points of $\widetilde{U}_i$ are parametrized stable maps $\tilde{\tau} = (\Sigma, \tilde{h})$ and the fiber $F(\tilde{\tau})$ is the space $L^{1,k}(\Lambda^{0,1}(\Sigma, \tilde{h}^*TP))$ of sections of a bundle over Σ. When $\tilde{\tau}$ moves from one stratum to another, the topological type of Σ may change, and so there is no easy way to identify these fibers.[13]

We assume further that the orbifold structure on $\mathcal{W}$ lifts to one on $\mathcal{E}$. Thus, $\mathcal{E}$ has the same local structure as $\mathcal{W}$, and we can define a multi-fold atlas $\widetilde{\mathcal{E}} = \{\widetilde{E}_I\}$ for it as above. By slight abuse of notation, we write $\widetilde{E}_I$ for the restriction to $\widetilde{V}_I$ of the fiber product of the $\widetilde{E}_i, i \in I$, over $\widetilde{U}_I$. More precisely, the fiber $\widetilde{F}(\tilde{x}_I)$ of $\widetilde{E}_I$ at $\tilde{x}_I = (\tilde{x}_i)_{i\in I} \in V_I$ has elements $(\tilde{x}_i, \tilde{v}_i)_{i\in I}$ where $\tilde{v}_i \in \widetilde{F}(\tilde{x}_i)$ and, for all i, j,

$$\Pi_i(\tilde{v}_i) = \Pi_j(\tilde{v}_j) \in \mathcal{E}.$$

Note that the desingularization $\widehat{E}_I \to \widehat{V}_I$ of $\widetilde{E}_I \to \widetilde{V}_I$ is an honest vector bundle since locally it has the same structure as the maps $\widetilde{E}_{i_0} \to \widetilde{U}_{i_0}$. Thus $\widetilde{E}_I \to \widetilde{V}_I$ is the finite union of vector bundles, and its fibers are finite unions of vector spaces.

Example 4.6 continued. Suppose that $\tilde{p} : \widetilde{E}_j \to \widetilde{U}_j$ is a trivial line bundle, on which Γ_j acts by $(x, v) \mapsto (-x, -v)$, and identify the quotients $\widetilde{E}_j/\Gamma_j = E_j$ in the obvious way. Then the fiber over a point (y, y) of $\widetilde{U}_{12}$ with $y \neq 0$ consists of the points $(y, v; y, v)$ while the fiber over $(y, -y)$ consists of the points $(y, v; -y, -v)$. When $y = 0$ the fiber is the union of the two lines $(0, v; 0, v)$ and $(0, v; 0, -v)$. Thus, the bundle $\widetilde{E}_{12} \to \widetilde{V}_{12}$ is the union of two line bundles, one over each line $y = \pm x$ in $\widetilde{U}_{12}$ with distinct fibers over the intersection point $(0, 0)$ of these lines that intersect only at the zero element. Observe also that $\mathrm{Stab}_{12}(0, 0) = \Gamma_{12} = \mathbb{Z}/2\mathbb{Z} \oplus \mathbb{Z}/2\mathbb{Z}$, and acts via

$$(\gamma_1, \gamma_2) \cdot (y, v; \pm y, \pm v) \mapsto (\gamma_1(y), \gamma_1(v), \gamma_2(\pm y), \gamma_2(\pm y)). \qquad \Box$$

Multi-sections

This is the heart of the construction. Everything we have done so far is aimed towards making it possible to define a large enough family of multi-sections to get transversality for the delbar operator. Intuitively a *multi-section* $\tilde{s}$ of $\widetilde{\mathcal{E}} \to \widetilde{\mathcal{V}}$ is a compatible collection $\{\tilde{s}_I\}$ of multi-valued sections. In other words, for each point $\tilde{x}_I$ of $\widetilde{V}_I$ we are given a nonempty finite subset $\tilde{s}_I(\tilde{x}_I)$ of $\widetilde{F}(\tilde{x}_I)$. In the situation at hand, it is enough to consider multi-sections that are single-valued when lifted to the desingularization $\widehat{E}_I \to \widehat{V}_I$. Thus we will assume that there are sections $\widehat{s}_I : \widehat{V}_I \to \widehat{E}_I$ such that

$$\tilde{s}_I(\tilde{x}_I) = \{proj \widehat{s}_I(\tilde{x}_I, \langle\widetilde{C}_\gamma(\tilde{x}_I)\rangle) : \gamma \in \mathrm{Stab}'_I(\tilde{x}_I)\}.$$

[12] Observe that if the stabilizer $\mathrm{Stab}_i(\tilde{x}_i)$ of any point $\tilde{x}_i \in \widetilde{U}_i$ acts nontrivially on the fiber $\widetilde{F}(\tilde{x}_i)$ then the quotient object $E_i \to U_i$ is not a vector bundle but rather a (local) orbibundle.

[13] Siebert [Sb] actually makes such an identification, thereby simplifying the topology at the cost of some extra analysis.

In particular, when $I = \{j\}$, the section $\widehat{s}_j = \tilde{s}_j$ is a single-valued but not necessarily equivariant section of the bundle $\widetilde{E}_j \to \widetilde{U}_j$. Note also that the sets $\tilde{s}_I(\tilde{x}_I)$ contain a single element when $\tilde{x}_I$ is in a top stratum.

If $J \subset I$, we define the pullback $(\Pi^I_J)^{-1}(\tilde{s}_J)$ to be the multi-valued section that associates to $\tilde{x}_I$ the full inverse image $(\Pi^I_J)^{-1}(\tilde{s}_J(\tilde{x}_J))$ where $\tilde{x}_J = \pi^I_J(\tilde{x}_I)$. The compatibility condition is that $\tilde{s}_I$ restricts over $(\pi^I_J)^{-1}(\widetilde{V}_J)$ to the pullback of $\tilde{s}_J$.

As the following example shows, the compatibility conditions are quite strong. In fact, if the sets U_i overlap too much there are no interesting (i.e. nonequivariant) multi-sections, which is precisely the reason why we consider the subcover $\{V_I\}$ of the $\{U_I\}$.

Example 4.7 Suppose that in the situation of the previous example we are given a section $\tilde{s}_1$ of $\widetilde{E}_1 \to \widetilde{U}_1$. Then, if $\tilde{s}_{12}$ denotes its pullback to $\widetilde{U}_{12}$,

$$\tilde{s}_{12}(y, y) = (y, \tilde{s}_1(y); y, \tilde{s}_1(y)), \quad \tilde{s}_{12}(y, -y) = (y, \tilde{s}_1(y); -y, \tilde{s}_1(-y)).$$

Thus $\tilde{s}_{12}$ is single-valued over the top strata in $\widetilde{U}_{12}$ since $y \neq 0$ there. Similarly, the pullback $\tilde{s}'_{12}$ of a section $\tilde{s}'_2$ over $\widetilde{U}_2$ has the form

$$\tilde{s}'_{12}(y, y) = (y, \tilde{s}'_2(y); y, \tilde{s}'_2(y)), \quad \tilde{s}'_{12}(-y, y) = (-y, -\tilde{s}'_2(y); y, \tilde{s}'_2(y)).$$

Therefore, in order for $\tilde{s}_{12}$ to be such a pullback, we need $\tilde{s}_1(y) = \tilde{s}'_2(y)$ and $-\tilde{s}_1(y) = \tilde{s}'_2(-y)$. This forces $\tilde{s}_1$ to be equivariant, i.e. $\tilde{s}_1(-y) = -\tilde{s}_1(y)$. In particular $\tilde{s}_j(0) = 0$. Thus $\tilde{s}_1$ extends to a multi-section of $\widetilde{\mathcal{E}} \to \widetilde{\mathcal{U}}$ only under this equivariance condition.

Suppose now that $V_1 = [0, 2), V_2 = (4, \infty) = \widetilde{V}_2$ and $V_{12} = [0, 5)$, and let $\tilde{s}_1$ be any section of $\widetilde{E}_1 \to \widetilde{U}_1$ that has support in $(-3, 3)$. It is easy to see that now $\tilde{s}_1$ does extend to a multi-section over $\widetilde{\mathcal{V}}$ since $\tilde{s}_1(\tilde{y}) = 0$ if $\tilde{y}$ lies above the overlap $V_{12} \cap V_2$. Observe also that if $\tilde{s}_1(0) \neq 0$ then $\tilde{s}_{12}$ has two values at $(0, 0) \in \widetilde{V}_{12}$. □

Remark 4.8 The graph $gr(\tilde{s}_I)$ is simply the union of the sets $\tilde{s}_I(\tilde{x}_I), \tilde{x}_I \in \widetilde{V}_I$, and it is often useful to impose geometric conditions on it, e.g. to require that both it and the projection map $gr(\tilde{s}_I) \to \widetilde{V}_I$ be an object in our category. For instance, such conditions are needed when discussing the intersection of the graph with the zero section as in Definition 4.11 below. This amounts to requiring that for each j the graphs $gr(\tilde{s}_j)$ have a stratification that is compatible with the projection to $\widetilde{V}_j$, that is preserved by the action of the groups Γ_j and is such that the rank of the stabilizer is constant on strata. Here one is allowed to refine the stratification of $\widetilde{V}_I$ provided that one does not subdivide its top strata. (Subdividing the top strata would introduce codimension 1 strata, which cause problems when one wants to define cycles.)

However, if one incorporates such geometric requirements into the definition of a multi-section, it is not possible to add an arbitrary pair of multi-sections since, for example, the sum $\tilde{s}_I + \tilde{s}'_I$ might equal zero on a very nasty set. Therefore we will not impose these conditions initially, though the sections that we will consider in practise will all satisfy such geometric conditions: see Example 4.10 below. □

Constructing multi-sections

Any family of compatible Γ_i-equivariant sections $\tilde{s}_i$ of the local bundles $\tilde{p} : \widetilde{E}_i \to \widetilde{U}_i$ gives rise to a multi-section of $\widetilde{\mathcal{E}}$. For example, the Cauchy–Riemann operator $\overline{\partial}_J$ is a multi-section of this kind. However, one cannot regularize the moduli space just by considering equivariant perturbations of the Cauchy–Riemann equation: see §5.1. The whole point of the present construction is to give a way of extending nonequivariant sections of the bundles $\widetilde{E}_j \to \widetilde{U}_j$ to multi-sections of $\widetilde{\mathcal{E}} \to \widetilde{\mathcal{V}}$, so that they can form global perturbations of the Cauchy–Riemann operator. As is suggested by the previous example, such an extension is possible provided that the covering $\{V_I\}$ is suitably chosen. The next lemma is adapted from [LiuT1] Lemma 4.3. The notation $A \subset\subset B$ means that the closure $cl(A)$ of A is contained in B.

Lemma 4.9 *Given any finite open covering $\{U_i\}$ of the compact subset $\overline{\mathcal{M}}$ of $\mathcal{W}$ there are open subsets $U_i^0 \subset\subset U_i$ and $V_I \subset U_I$ with the following properties:*
(i) $\overline{\mathcal{M}} \subset \cup_i U_i^0$, $\overline{\mathcal{M}} \subset \cup_I V_I$;
(ii) $V_I \cap U_i^0 = \emptyset$ if $i \notin I$;
(iii) if $V_I \cap V_J \neq \emptyset$ then one of the sets I, J contains the other.

Proof: If there are N sets U_i, choose for $n = 0, \ldots, N$ open coverings $\{U_i^n\}$, $\{W_i^n\}$ of $\overline{\mathcal{M}}$ such that

$$U_i^0 \subset\subset W_i^1 \subset\subset U_i^1 \subset\subset W_i^2 \subset\subset \ldots \subset\subset U_i^N = U_i.$$

Then, if $\operatorname{card} I = |I| = \ell$, define

$$V_I = W_I^\ell - \bigcup_{J:|J|>\ell} cl(U_J^{\ell+1}).$$

It is easy to check that the required properties hold. □

The next example is taken from [LiuT1] Lemma 4.5, and explains how it is possible to construct a sufficiently large family of multi-sections.

Example 4.10 Let the U_i^0 and V_I be as above and (by shrinking $\mathcal{W}$) suppose that $\mathcal{W} = \cup_i U_i^0 = \cup_I V_I$. Consider a multi-bundle $\widetilde{\mathcal{E}} \to \widetilde{\mathcal{V}}$ as above and suppose that for some j we are given a section $\sigma(j) : \widetilde{U}_j \to \widetilde{E}_j$ of the bundle $\widetilde{U}_j \to \widetilde{E}_j$ with support in $\widetilde{U}_j^0$. Then there is an induced multi-section $\tilde{s}(j)$ of $\widetilde{\mathcal{E}} \to \widetilde{\mathcal{V}}$ defined as follows:

> $\tilde{s}(j)_I$ is the zero section if $j \notin I$, and, if $j \in I$, is the restriction to $\widetilde{V}_I$ of the pullback to $\widetilde{U}_I$ of the graph of $\sigma(j)$.

Suppose further that $\sigma(j)$ is generic in the sense that, for all points $\tilde{x}_j \in \widetilde{U}_j$ no element of the stabilizer group $\operatorname{Stab}_j(\tilde{x}_j)$ fixes $\sigma(j)(\tilde{x}_j)$ unless $\sigma(j)(\tilde{x}_j) = 0$. We require also that the boundary of the support of $\sigma(j)$ is a union of strata. Then, it is not hard to see that its graph can be given a stratification compatible with the projections to the $\widetilde{V}_I$ as discussed in Remark 4.8. The essential point to check is that each set on which the number $|\tilde{s}(j)_I(\tilde{x}_I)|$ is constant is a union of strata. By construction this number is 1 over the top strata of $\widetilde{V}_I$, so there is no problem there. Note also that the genericity condition implies that, if $\pi_j^I(\tilde{x}_I)$ is not in the zero set of $\sigma(j)$, then $|\tilde{s}(j)_I(x_I)|$ is precisely equal to the number $|\operatorname{Stab}_{I-i_0}(\tilde{x}_I)|$ of local components at $\tilde{x}_I$ (i.e. is the maximum possible) and so is constant on strata.

The multi-sections that we will use in §5 have the form

$$\tilde{s} = \overline{\partial}_J + \sum_j \tilde{s}(j)$$

where $\overline{\partial}_J$ is the section defined by the Cauchy–Riemann operator. □

4.3 Construction of the virtual cycle

Finally, we need to see how multi-sections give rise to branched pseudomanifolds. From now on we will assume that the covering $\{V_I\}$ satisfies the conditions of Lemma 4.9. The following definition is very much *ad hoc*: we are gathering together a collection of useful properties of $\tilde{s}$ rather than formulating the best and most general definition. In particular, normally a Fredholm object has finite dimension kernel and cokernel. Here we retain the information on the kernel plus the vestiginal information on the cokernel that can be gleaned from condition (i).

Definition 4.11 We say that $\tilde{s}$ is *Fredholm* of index d if the following conditions are satisfied:

(i) for each I, $gr(\tilde{s}_I)$ meets the zero section of $\widetilde{E}_I \to \widetilde{V}_I$ transversally in a d-dimensional (open) pseudomanifold $\widetilde{Z}_I$, whose top stratum coincides with its intersection with the top stratum of $\widetilde{V}_I$.

(ii) The top strata S of $\widetilde{Z}_I$ are provided with orientations that are preserved under the partially defined projections π^I_J.

(iii) The union $Z_{\mathcal{W}} = \cup_I \pi_I(\widetilde{Z}_I) \subset \mathcal{W}$ is compact.

Before showing how to assemble the local zero sets $\widetilde{Z}_I$ into a pseudomanifold, let us discuss the question of the existence of Fredholm sections. In the finite dimensional case, our assumption that $\widetilde{Z}_I$ has no codimension 1 strata implies that the sets $\widetilde{V}_I$ themselves have no codimension 1 strata (or at least no such strata that are anywhere near the zero set of $\tilde{s}$.) In our main (infinite dimensional) example all bundles etc are complex and so all strata in $\widetilde{Z}_I$ have even codimension. Much of the work in [LiuT1] is taken up with showing that a suitable perturbation of the delbar operator does give rise to a Fredholm section. They construct a finite dimensional family of sections $\tilde{s}^\nu$, for $\nu \in R$, which are Fredholm in our sense for generic ν. This is discussed further in §5 below.

One would also hope that in the finite, even dimensional context, generic sections are Fredholm. Here $\widetilde{V}_I$ would have dimension $2n$ say and the fibers of $\tilde{p} : \widetilde{E}_I \to \widetilde{V}_I$ would have dimension $2k$ so that $\tilde{s}$ would have index $d = 2n - 2k$. The idea would be to show that any multi-section can be perturbed so that it intersects the zero section of $\tilde{p} : \widetilde{E}_I \to \widetilde{V}_I$ transversally. In fact, the strata of $\widetilde{E}_I$ have the form $\tilde{p}^{-1}(S)$, where S is a stratum of $\widetilde{V}_I$ and, by construction, each stratum S' of $gr(\tilde{s}_I)$ lies inside a stratum $\tilde{p}^{-1}(S)$ in $\widetilde{E}_I$. Hence it would suffice to perturb $gr(\tilde{s}_I)$ so that the intersection of each S' with such S is transverse inside $\tilde{p}^{-1}(S)$. This can be arranged, but only in the presence of suitable normal cones.[14]

We now show how to glue the local pseudomanifolds $\widetilde{Z}_I$ together to get a closed branched labelled pseudomanifold Y that projects onto $Z_{\mathcal{W}} = \cup_I \pi_I(\widetilde{Z}_I)$. First of all we must replace the open sets $\widetilde{Z}_I$ by closed pseudomanifolds that we will call $\widetilde{Y}_I$. We will do this by choosing manifolds with boundary $\overline{V}'_I \subset V_I$ that cover $\mathcal{W}$ and then setting

$$\widetilde{Y}_I = \widetilde{Z}_I \cap \pi_I^{-1}(\overline{V}'_I).$$

Here $\widetilde{Y}_I$ is given the obvious first topology, and the components of $(\widetilde{Y}_I)_{sm}$ either lie on its boundary (i.e. in $\pi_I^{-1}(\partial\overline{V}'_I)$), or are the intersections of strata in $\widetilde{Z}_I$ with its interior. In order for there to be such a stratification, it suffices that the stratification of $\mathcal{W}$ can be refined so that each boundary $\partial\overline{V}'_I$ is a union of strata. Thus the boundaries $\partial\overline{V}'_I$ must intersect transversally and be in general position with respect

[14] The point here is that with stratified spaces one works inductively over the strata. If one perturbs $\tilde{s}_I$ over one stratum so that it is transverse to the zero section there, one needs to be able to extend this perturbation to nearby strata, and hence one needs information on how that strata fit together.

to the original strata of $\mathcal{W}$. Strictly speaking, one can arrange this for general $\mathcal{W}$ only in the presence of suitable normal cones.

If one can find such sets $\widetilde{Y}_I$ we will say that they form a *shrinking* of the $\widetilde{Z}_I$. For now, we have stratified the $\widetilde{Y}_I$ so that their only codimension 1 strata lie in the boundary $\partial \overline{V}'_I$. These will form the branching locus of Y.

The next step is to construct a topological space Y such that there are continuous maps

$$\coprod_I \widetilde{Y}_I \longrightarrow Y \longrightarrow \mathcal{Z}_{\mathcal{W}}.$$

Points of Y are equivalence classes $\langle z \rangle$ under the equivalence relation that is generated by setting

$$\tilde{y}_I \sim \tilde{z}_J, \quad \text{if } J \subset I, \ \pi^I_J(\tilde{y}_I) = \tilde{z}_J.$$

The first topology on Y its the quotient topology, and the strata of its second (smooth) topology are formed by the images of strata in the $\widetilde{Y}_I$, subdivided as necessary. It will be convenient now to refine the stratifications of the $\widetilde{Y}_I$ so that the projections $q_Y : \widetilde{Y}_I \to Y$ take strata to strata. This introduces new codimension 1 strata in the $\widetilde{Y}_I$ coming from the boundaries of the $\widetilde{Y}_K$.

Example 4.12 As an instructive example consider the bundle $\widetilde{\mathcal{E}} \to \widetilde{\mathcal{V}}$ of Example 4.7 and suppose that $\widetilde{Z}_I$ is simply its zero section. Thus $\widetilde{Z}_I = \widetilde{V}_I$ for all i. Take $\overline{V}'_1 = [0,1], \overline{V}'_2 = [4.1, \infty)$ and $\overline{V}'_{12} = [.5, 4.5]$. Then Y projects to the base $\mathcal{W} = [0, \infty)$ by a map that is injective over $\overline{V}'_1 \cup \overline{V}'_2$ but is 2 to 1 elsewhere. In other words, the branching in Y comes from the properties of the cover $\overline{V}'_I$. Although this example does not have all the properties we require (eg it is not even-dimensional), this feature remains. Our multi-sections $\tilde{s}$ are single valued over top strata so that $\widetilde{Z}_I$ has no branching. When one makes identifications to form Y, the branching does not come from properties of the multi-section $\tilde{s}$ that generates $\widetilde{Z}_I$ but from the properties of the cover $\{\overline{V}'_I\}$. □

It remains to define the labelling function. If S is a top stratum of Y lying in the image of $q_I : \widetilde{Y}_I \to Y$, we set

$$\lambda_I(y) = \frac{|\{q_I^{-1}(y)\}|}{|\Gamma_I|}, \quad y \in S.$$

Proposition 4.13 *Suppose that $\{V_I\}$ is as in Lemma 4.9, that $\tilde{s}$ is Fredholm and that the zero sets $\widetilde{Z}_I$ have a shrinking $\widetilde{Y}_I$. Then the labelling function λ_I descends to a function λ on Y, and gives (Y, λ) the structure of a compact branched labelled pseudo-manifold.*

Proof: To prove the first statement it suffices to check that $\lambda_I(y) = \lambda_J(y)$ if $J \subset I$ and $y \in q_J(\widetilde{Y}_J) \cap q_I(\widetilde{Y}_I)$. This holds because $\tilde{s}_I$ is the pullback of $\tilde{s}_J$ over $(\pi^I_J)^{-1}(\widetilde{V}_J)$, so that

$$|\{q_I^{-1}(y)\}| = |\{q_J^{-1}(y)\}| \cdot |\Gamma_{I-J}|, \quad \Gamma_I = \Gamma_J \cdot |\Gamma_{I-J}|.$$

To prove the second statement, we have to check the branching condition of Definition 4.3 on the codimension 1 components of Y. Thus we have to understand the effect of one of the generating equivalences $\tilde{y}_J \sim \tilde{z}_J$ where $J \subset I$ on this condition. Note first that a class $\langle y \rangle$ lies in a codimension 1 strata B_ℓ in Y if for one K the representatives of $\langle y \rangle$ in $\widetilde{Y}_K$ lie on boundary strata B_K and if for all other J the representatives in $\widetilde{Y}_J$ lie in the interior of $\widetilde{Y}_J$.

We will call the side of B_K that meets the interior of $\widetilde{Y}_K$ the positive side, and will write $\{S^+_{K,\alpha}\}$ for the set of all top strata in $\widetilde{Y}_K$ whose closures contain representatives of $\langle y\rangle$. Similarly, for all other J we write $\{S^\pm_{J,\alpha}\}$ for the set of top strata in $\widetilde{Y}_J$ that contain representatives of $\langle y\rangle$, with signs $\pm$ assigned in the obvious way. By Lemma 4.9, the set of J in question have nonempty intersection I_y with K and everything is essentially pulled back from $\widetilde{V}_{I_y}$. For example, two strata in $\{S^+_{J,\alpha}\}$ are identified in Y if and only if they have the same image in $\widetilde{V}_{I_y}$. Therefore, we only need to consider the effect of inclusions of the form $I_y \subset I$ on the branching condition. There are three cases, the first in which $K \neq I_y, I$, the second in which $K = I$, the third in which $K = I_y$. For instance, in Example 4.12 the codimension 1 strata of Y are the points $.5, 1, 4.1$ and 4.5. At the point $y = 4.1$, $K = \{2\} = I_y$ and $J = \{1, 2\}$, while at the point $y = 4.5$, $K = \{1, 2\}$ and $J = \{2\} = I_y$. Here inclusions of the first kind do not occur, inclusions of the second kind have no branching effect, while inclusions of the third kind give rise to the branching at the points $.5, 4.5$.

This illustrates what happens in general. Inclusions of the first kind are essentially neutral. For if J is either I_y or I there is a bijective correspondence between the components $\{S^+_{J,\alpha}\}$ and $\{S^-_{J,\alpha}\}$ on both sides that commutes with the identifications coning from $\sim$. So whatever happens is the same on both sides. In the second case, each stratum in $\{S^+_{K,\alpha}\}$ is mapped bijectively by $\pi^K_{I_y}$ onto a stratum in $\{S^+_{I_y,\alpha}\}$, and each of the latter strata is covered exactly $|\Gamma_{K-I_y}|$ times. Hence both sides of B in Y are the same, and no real branching occurs here either.

However, the situation is different if $K = I_y$. Then the components in $\{S^+_{I,\alpha}\}$ are identified via π^I_K with components in $\{S^+_{K,\alpha}\}$, while no identification is put on the components $\{S^+_{I,\alpha}\}$ on the other side. Since π^I_K is $|\Gamma_{I-K}|$ to 1, $|\Gamma_{I-K}|$ components on the negative side of B in Y will correspond to each component on the positive side. It is not hard to check that the sum of the labels on the two sides is the same. □

Obviously, one would now like to prove a statement to the effect that the homology class represented by the pseudo-manifold (Y_s, λ) defined by a Fredholm multi-section $\tilde{s}$ is independent of the choice of multi-section and of coverings $\{U_j\}, \{V_J\}$ and so on. For such a statement to make sense, one first has to assume that there is a map $ev : \mathcal{W} \to X$, so that the classes one is comparing are given by the composite maps

$$Y_s \xrightarrow{\pi} \mathcal{W} \xrightarrow{ev} X.$$

Secondly, one has to set up a context in which there is a well defined range of possible choices. This is no problem as far as the coverings are concerned, but does raise questions about what Fredholm multi-sections to consider. In the main example considered in §5 below, there is an evaluation map $ev : \mathcal{W} \to X$ and also a well defined family of Fredholm sections, namely those of the form $\overline{\partial}_J + \sum_j \tilde{s}^\nu(j)$. Since this is a connected family it is not too hard to show that the homology class defined by (Y, λ) is independent of choices.

Another natural framework to consider is the finite dimensional case. As an illustration of our method we show below how to calculate the Euler number of a teardrop T using multi-sections.

4.4 Finite dimensional orbifolds

Resolutions

As is shown by Example 4.12 the method used in the last section to piece the zero sets $\widetilde{Z}_I$ together into a branched, labelled pseudomanifold can be applied to the zero section of any finite dimensional multi-bundle over a finite dimensional multi-fold, provided that this is oriented. (Note that any finite dimensional oriented orbifold in the usual smooth category has no codimension 1 strata, since the strata come from points with nontrivial stabilizers and these always have codimension ≥ 2 in the oriented case.) In other words, as observed by Polterovich, one can construct from an oriented orbifold M a branched pseudocycle Y_M that maps to M and can be thought of as a *resolution* of M. Note that when constructing this we do not have to pass to the cover $\{\widetilde{V}_I\}$ but can simply use the $\{\widetilde{U}_I\}$, since the zero section is equivariant.

To illustrate this, let us look at the tear drop T. This is the simplest 2-dimensional orbifold. Topologically T is a 2-sphere but it has one singular point at the north pole p_N where the group Γ_N is cyclic of order k. Therefore it can be covered by two open discs, U_1 containing p_N and U_2 containing the south pole p_S, that intersect in an open annulus U_{12}. Then $\widetilde{U}_2 = U_2$ and $\pi_1 : \widetilde{U}_1 \to U_1$ is given in polar coordinates (r, θ) by $(r, \theta) \mapsto (r, k\theta)$. The sets $\overline{V}'_I$ should be slightly smaller closed discs and annuli. The resulting branched pseudomanifold Y_T is the union of an open disc $D_2 = \operatorname{int}(\overline{V}'_2)$ with its bounding circle C and another disc $D_1 = U_1 - \operatorname{int}(\overline{V}'_2)$, whose boundary is attached to C by the k-fold covering $\theta \mapsto k\theta$. Points in D_1 have labels $1/k$ while those in D_2 have labels 1.

Calculating the Euler number

The general idea is this. Let M be a finite dimensional oriented orbifold with local uniformizers $(\widetilde{U}_j, \pi_j, \Gamma_j)$. Here we assume that $\widetilde{U}_j$ is an open subset of some oriented vector space and that Γ_j acts linearly, preserving orientation. It follows that any subspace with nontrivial stabilizer subgroup has codimension at least 2. Hence there are no codimension 1 strata and we are in a situation in which Fredholm sections can exist. (See the discussion after Definition 4.11.) Moreover, all strata have normal cones.

One can construct a multi-fold atlas $\mathcal{V}_M$ for M from any subcovering $\{V_I\}$ of the sets $\{U_I\}$ as in §4.2. The group Γ_j acts on the tangent bundle $\widetilde{E}_j = T\widetilde{U}_j \to \widetilde{U}_j$ of $\widetilde{U}_j$ and the quotient by these local actions is an orbibundle $\mathcal{E} \to M$.[15] From this we construct a multi-bundle $\widetilde{\mathcal{E}} \to \mathcal{V}_M$ as before. Given a suitable compactly supported section $\sigma(j)$ of $\widetilde{E}_j \to \widetilde{U}_j$, one can construct a multi-section $\tilde{s}(j)$ of $\widetilde{\mathcal{E}}$ as in Example 4.10. Then the section

$$\tilde{s} = \sum_j \tilde{s}(j)$$

should be well-defined and Fredholm. Hence we can construct a 0-dimensional branched, labelled pseudomanifold $Y_{\tilde{s}}$ from its zero set. The Euler number of M is defined to be the number of points in $Y_{\tilde{s}}$, weighted by their labels. The claim is that this is independent of the choice of $\tilde{s}$. This looks very plausible: further details are left to the reader.

In the case of the tear drop, with U_1, U_2 as above, choose polar coordinates on U_1 so that the annulus $U_{12} = V_{12}$ is given by

$$(t, \theta) \; : \; t \in (1, 4), \; \theta \in S^1 = \mathbb{R}/\mathbb{Z}.$$

Then let V_1, V_2 be smaller discs in U_1, U_2 respectively, such that

$$V_1 \cap V_{12} = (1, 2) \times S^2, \quad V_2 \cap V_{12} = (3, 4) \times S^1.$$

The tangent bundle $\widetilde{E}_1 = T\widetilde{U}_1 \to \widetilde{U}_1$ can be trivialised away from p_N by the vector fields $\tilde{\partial}_r, \tilde{\partial}_\theta$ and similarly, if (s, ϕ) are polar coordinates in U_2 centered at p_S, $\widetilde{E}_2$ can be trivialised by $\tilde{\partial}_s, \tilde{\partial}_\phi$. Note that the sections $\tilde{\partial}_r, \tilde{\partial}_\theta$ are invariant under the action of Γ_N and so descend to sections $\partial_r, \partial_\theta$ of $E_1 \to U_1 - p_N$. Moreover, on the overlap V_{12}, $\partial_r = -\partial_s$.

Let $\sigma(1) : \widetilde{U}_1 \to \widetilde{E}_1$ be a section of the form $\beta(r)\tilde{\partial}_r$ where the function β has an isolated zero at $r = 0$ and is everywhere ≤ 0, and let $\sigma(2) : \widetilde{U}_2 \to \widetilde{E}_2$ have the form $\gamma(s)\tilde{\partial}_s$ where the function γ has an isolated zero at $s = 0$ and is everywhere ≥ 0. Then it is not hard to see that the only zeros of $\tilde{s} = \tilde{s}(1) + \tilde{s}(2)$ are

[15] As an exercise the reader might like to give a formal definition of an orbibundle: cf [FO]. The point is to formulate exactly how the elements of the different local bundles E_i should be identified with each other.

at p_N and p_S respectively. However, the point p_N is not in the top stratum of $\widetilde{U}_1$, and so condition (i) in Definition 4.11 is not satisfied. By slightly perturbing $\tilde{s}$ near p_N we can achieve transversality. Now there is a one zero near p_N that counts with weight $1/k$ and one at p_S that counts with weight 1. Thus the Euler number is $1 + 1/k$.

One could try this with a different choice of $\sigma(1), \sigma(2)$. For example, suppose that we give U_2 rectangular coordinates (x, y) and let $\sigma(2) = \gamma(s)\partial_x$ for some cutoff function γ. Then $\sigma(2) = \gamma(s)(\cos(\phi)\partial_s - \sin(\phi)\partial_\phi)$ and so, in the coordinates $(r, \theta) = (-s, -\phi)$ of U_1 looks like $\beta(r)(-\cos(\theta)\partial_r - \sin(\theta)\partial_\theta)$. Therefore it lifts to the section

$$\tilde{\sigma} = \beta(r)(-\cos(k\theta)\tilde{\partial}_r - \frac{\sin(k\theta)}{k}\tilde{\partial}_\theta)$$

of $\widetilde{E}_1$ over the circle $r = 4$. Now $\sigma(1)$ vanishes on this circle. Therefore, however we choose $\sigma(1)$ the only zeros of $\tilde{s} = \tilde{s}(1) + \tilde{s}(2)$ are going to be in $\widetilde{U}_1$ and they each will be weighted by $1/k$. Moreover the number of these zeros is the winding number of $\tilde{\sigma}$ where this is considered as a map $S^1 \to \mathbb{R}^2 - \{0\}$. But $\tilde{\sigma}$ is homotopic to the map

$$\begin{aligned} \theta &\mapsto \cos(k\theta)\tilde{\partial}_r + \sin(k\theta)\tilde{\partial}_\theta \\ &= \cos((k+1)\theta)\tilde{\partial}_x + \sin((k+1)\theta)\tilde{\partial}_y, \end{aligned}$$

and so has winding number $k + 1$. Hence we get $1 + 1/k$ as before.

We can check this by calculating the Euler number in a different way. Triangulate T, making p_N one of the vertices, and then to take the sum of the weighted number of vertices and faces minus the weighted number of edges. Clearly, all edges, faces, and vertices other than p_N should have weight 1 as usual. However, intuitively it should be correct to count p_N with weight $1/k$. This method gives an Euler characteristic of $1 + 1/k$. To justify this, note that if there were 2 orbifold points, both of order k, then this calculation would give an Euler number of $2/k$. This is surely correct since this second orbifold has a k-fold covering by the 2-sphere.

5 The virtual moduli cycle

In this section we briefly describe how to construct the virtual moduli cycle $\overline{\mathcal{M}}^\nu = \overline{\mathcal{M}}^\nu_{0,k}(X, J, A)$ as a branched pseudomanifold Y of dimension d, where $d = \dim X + 2c_1(A) + 2k - 6$. For simplicity, we will often ignore the marked points.

Here are the main features of the structure of Y. The elements of Y are parametrized stable maps $(\Sigma, \tilde{h})$ that each satisfy a perturbed Cauchy–Riemann equation

$$\overline{\partial}_J \tilde{h}(w) = \nu_{\tilde{h}}(w), \qquad w \in \Sigma,$$

where $\nu_{\tilde{h}}$ is a C^∞ smooth 1-form in the Sobolev space

$$\widetilde{L}_{\tilde{h}} = L^{1,p}\left(\Lambda^{0,1}(\Sigma, \tilde{h}^*(TX))\right),$$

consisting of all sections of the bundle $\Lambda^{0,1}(\Sigma, \tilde{h}^*(TX))$ that are $L^{1,p}$-smooth on each component. There is a finite-to-one map $Y \to \mathcal{W}$ that forgets the parametrization, such that each stratum in Y_{sm} of codimension $2k$ or $2k+1$ is taken to a stratum in $\mathcal{W}$ consisting of elements whose domain has $\le k+1$ components. (We would have equality here if it were not for the extra strata that are introduced to deal with the singular points of the closure of the branching locus.) In particular, the top strata M_i and the branch components B_j consist of maps with domain a single sphere.

The evaluation map $ev : \overline{\mathcal{M}}^\nu \to X^k$ is given by composing the forgetful map $Y \to \mathcal{W}$ with the evaluation map $\mathcal{W} \to X^k$. $\overline{\mathcal{M}}^\nu$ is called a virtual moduli cycle because the class represented by $ev : \overline{\mathcal{M}}^\nu \to X^k$ can be used to calculate Gromov–Witten invariants. Sometimes $\overline{\mathcal{M}}^\nu$ is called a *regularization* of $\overline{\mathcal{M}}$.

5.1 The local construction

We saw in §3.2 that $\mathcal{W}$ has the local structure assumed in §4.2, i.e. it is a partially smooth orbifold. The next step is to define a bundl $\widetilde{\mathcal{L}}_\tau \to \widetilde{U}_\tau$ whose fiber at an element $\tilde\tau' = (\Sigma', \tilde h')$ is the space $\widetilde{L}_{\tilde h'}$ defined above. Note that the fiber changes as when the topological type of the domain Σ' changes. But it is a locally trivial bundle over each stratum and it is not hard to give it a global topology. Therefore it is a bundle in our category. Clearly, the action of Γ_τ by reparametrization lifts to $\widetilde{\mathcal{L}}_{\tilde h}$.

The Cauchy–Riemann operator $\overline{\partial}_J$ gives rise to a Γ_τ-equivariant section of $\widetilde{\mathcal{L}}_\tau \to \widetilde{U}_\tau$. Moreover, its linearization at $\tilde\tau'$ is:[16]

$$D\tilde h' : L^{2,p}(\Sigma', (\tilde h')^*(TX)) \to \widetilde{L}_{\tilde h'},$$

where $p > 2$. Here, if Σ' has $\ell + 1$ components, the domain of $D\tilde h'$ consists of all $\ell + 1$-tuples

$$(\xi_0, \dots, \xi_\ell), \qquad \xi_i \in L^{2,p}(\Sigma'_i, (\tilde h'_i)^*(TX),$$

such that

(i) $\xi_i(x) = \xi_j(x)$ at all intersection points x of the components Σ'_i, Σ'_j, and

(ii) $\xi_i(w_j) \in (\tilde h')^* T\mathbf{H}_j$ at all added points w_j.

Conditions (i) are needed since this domain is the tangent space of a space of maps $\Sigma' \to X$ with connected domain Σ', while conditions (ii) refer to the added marked points needed to stabilize the domain and correspond to the fact that we are restricting to a slice for the action of the reparametrization group. Observe that there are no similar compatibility conditions on the elements of the range $\widetilde{L}_{\tilde h'}$.

Now look at the "center" point τ of U_τ and let $\tilde\tau = (\Sigma, \tilde h) \in \widetilde{U}_\tau$ be one of its lifts. Since $D\tilde h$ is elliptic, there is a finite dimensional subspace R_τ of the vector space $C^\infty\left(\Lambda^{0,1}(\Sigma \times X, pr^*TX)\right)$ such that the operator

$$D\tilde h \oplus e : L^{2,p}(\Sigma, \tilde h^*(TX)) \oplus R_\tau \to \widetilde{L}_{\tilde h} \tag{1}$$

is surjective, where e is given by restricting the sections in R_τ to the graph of $\tilde h$. In general the elements of R_τ cannot be chosen to be invariant under the automorphism group $\Gamma_{\tilde\tau}$ of $\tilde\tau$. However, we can enlarge R_τ by adding to it all its images under the elements of $\Gamma_{\tilde\tau}$, so that this group does act on it. It will be convenient to suppose that this action is free on the nonzero elements of R_τ. Clearly, this can be arranged by slightly perturbing R_τ.

The next step (see [LiuT1] §3) is to show that we can choose R_τ and a partially smooth family of embeddings $e_{\tilde\tau'} : R_\tau \to \widetilde{L}_{\tilde h'}$, so that the linearized operator

$$D\tilde h_{\tau'} \oplus e_{\tilde\tau'} : L^{2,p}(\tilde h^*_{\tau'}(TX)) \oplus R_\tau \to \widetilde{L}_{\tilde h'} \tag{2}$$

is surjective (with uniform estimates for the inverse) for all $\tilde\tau' \in \widetilde{U}_\tau$. Here we are allowed to shrink U_τ. The fact that this is possible is a deep result. If $\tilde\tau'$ lies in the same stratum as $\tilde\tau$, it holds because of the

[16] If one used Siebert's way of fixing parametrizations one could consider this as a map from $L^{1,p}$-sections to L^p-sections.

openness of the regularity condition. However, to prove this in general one has to show that there is a uniformly bounded family of right inverses to $D\tilde{h}_{\tau'} \oplus e_{\tilde{\tau}'}$ as $\tilde{\tau}'$ varies over a small enough neighborhood $\widetilde{U}_\tau$. This analytic fact is the basis of all gluing arguments: see for example, [RT], [LiuT1] and [MS] Appendix A.

With this done, let $pr : \widetilde{U}_\tau \times R_\tau \to \widetilde{U}_\tau$ denote the projection and consider the pullback bundle $pr^*(\widetilde{\mathcal{L}}_\tau) \to \widetilde{U}_\tau \times R_\tau$. This bundle has a section s defined by

$$s(\tilde{\tau}', \nu) = \overline{\partial}_J(\tilde{h}_{\tau'}) + e_{\tilde{\tau}'}(\nu). \tag{3}$$

By construction, its linearization is surjective at all points $(\tilde{\tau}, 0)$ and so it remains surjective for $|\nu| \le \varepsilon$. One then shows by some variant of gluing that the intersection of s with the zero section is a (partially smooth) open pseudomanifold of dimension equal to $\operatorname{ind} D\tilde{h} + \dim R_\tau$, whose components are given by intersecting the solution space with the different strata in $\widetilde{U}_\tau$. In other words, s is Fredholm. Since R_τ is finite dimensional, the Sard–Smale theorem (cf [MS] §3.4) now implies that for generic fixed $\nu \in R_\tau$ the corresponding section s^ν of $\widetilde{\mathcal{L}}_\tau \to \widetilde{U}_\tau$ is also Fredholm. Hence:

Lemma 5.1 *For generic $\nu \in R_\tau$ the section s^ν is Fredholm and its solution set $\widetilde{Z}^\nu_\tau$ has the structure of an open pseudomanifold of dimension $d = \operatorname{ind} Dh$.*

5.2 The global construction

Because $\overline{\mathcal{M}}$ is compact, it is covered by a finite number $U_j, j = 1, \ldots, k$ of the sets of the form U_τ and $\mathcal{W}$ is defined to be their union. We write $(\widetilde{U}_j, \Gamma_j, \pi_j)$ for the corresponding uniformizers, where

$$\pi_j : \widetilde{U}_j \to \widetilde{U}_j/\Gamma_j = U_j$$

is the obvious projection. Choose subcovers $\{U_i^0\}$ and $\{V_I\}$ as in Lemma 4.9, and then a partition of unity β_i on $\mathcal{W}$ that is subordinate to the covering $\{U_i^0\}$ and such that the boundaries of the sets where $\beta_i \ne 0$ are transversally intersecting submanifolds. Then define the section $\sigma(i)$ of the pullback bundle $pr^*(\widetilde{\mathcal{L}}_i) \to \widetilde{U}_i \times R_i$ by

$$\sigma(i)(\tilde{\tau}', \nu_i) = \beta_i(\tau') \,.\, e_{\tilde{\tau}'}(\nu_i),$$

as in Example 4.10. Next, choose $\varepsilon > 0$ so that, for all i, the linearization of $\overline{\partial}_J + \sigma(i)$ is surjective at all points where $\beta_i \ne 0$ and $|\nu| \le \varepsilon$, and put

$$R = \oplus_i R_i, \quad R_\varepsilon = \{\nu = (\nu_i) \in R : |\nu_i| \le \varepsilon, \text{ for all } i.\}$$

Let $\mathcal{W}_\varepsilon = \mathcal{W} \times R_\varepsilon$. The structure on $\widetilde{\mathcal{L}} \to \widetilde{\mathcal{V}}$ constructed in §4.2 can be pulled back to give a multi-bundle $pr^*(\widetilde{\mathcal{L}}) \to \widetilde{\mathcal{V}}_\varepsilon$. Hence this multi-bundle has a multi-section $\tilde{s}$ given by

$$\tilde{s}(\tilde{\tau}', \nu) = \overline{\partial}_J + \sum_i \tilde{s}(i)(\tilde{\tau}', \nu_i) = \overline{\partial}_J + \sum_i \beta_i(\tau') \,.\, e_{\tilde{\tau}'}(\nu_i).$$

Moreover, if the R_i are chosen so that the elements in $\cup_{i \in I} e_{\tilde{\tau}'}(R_i)$ are linearly independent for each $\tilde{\tau}' \in \widetilde{V}_I$, it is not hard to see that this section $\tilde{s}$ satisfies the geometric conditions of Example 4.10. The proof of Lemma 5.1 then adapts to show that $\tilde{s}$ is Fredholm. Hence the corresponding multi-section $\tilde{s}^\nu$ of $\widetilde{\mathcal{L}} \to \widetilde{\mathcal{V}}$ is Fredholm for generic $\nu \in R$. Therefore we can apply Proposition 4.13 to conclude:

Proposition 5.2 *For generic $\nu \in R$ the zero sets $\widetilde{Z}_I^\nu$ of the multi-section $\tilde{s}^\nu$ of $\widetilde{\mathcal{L}} \to \widetilde{\mathcal{V}}$ fit together to give a compact branched and labelled pseudomanifold $\overline{\mathcal{M}}^\nu$ of dimension d.*

Definition 5.3 Any d dimensional branched pseudomanifold $Y = \overline{\mathcal{M}}^\nu$ constructed as above from some covering of a neighborhood $\mathcal{W}$ of $\overline{\mathcal{M}}_{0,k}(X, J, A)$ will be called an *regularization* of $\overline{\mathcal{M}}_{0,k}(X, J, A)$.

By considering a common refinement of two covers, it is not hard to see that the rational bordism class of the map $ev : \overline{\mathcal{M}}^\nu \to X^k$ is independent of the choice of covers $\{U_j\}$ and $\{V_I\}$. Further, by considering a 1-parameter version of this construction, Liu–Tian show that, for any regularization $\overline{\mathcal{M}}^\nu$ of $\overline{\mathcal{M}}_{0,k}(X, J, A)$, this is independent of the choice of J and of perturbation term ν. In fact, it depends only on the deformation class of the symplectic form ω on X, on the homology class A and on the number of marked points k. The Gromov–Witten invariants are now defined by the method given in §1. It is not hard to show that, in all cases where the previous definition in terms of the uncompactified space $\mathcal{M}_{0,k}(X, J, A)$ makes sense, the new invariants equal the old. (On the other hand, I am not sure that anyone has yet checked that all the different ways of defining general Gromov–Witten invariants give the same result.)

Observe that by this method we construct a cycle that represents a *rational* homology class. Thus in general the Gromov–Witten invariants take rational values. In the case when the class A is primitive (i.e. not a multiple of any other integral class) one would expect that the invariant should be integral. However, this has not yet been proven.

Remark 5.4 (i) Liu–Tian's definitions in [LiuT1] are actually slightly different from those given here. Their aim is to prove the Arnold conjecture, and so their basic domain, i.e. the domain of the uncompactified space $\mathcal{M}$, is a cylinder rather than a sphere. They consider this cylinder to be a sphere with two special points at the poles and a marked arc between them, and perturb the equation on this component by an appropriate Hamiltonian term so that the elements of $\mathcal{M}$ are Floer trajectories between periodic orbits. This means that they must replace the reparametrization group $\mathrm{PSL}(2, \mathbb{C})$ by a subgroup that respects this extra structure. Further, there are now two kinds of bubbling: the bubbling off of spheres and the lengthwise splitting of the cylinder into two or more pieces. Thus the stable curves that they consider have a more complicated structure than ours: there is a connected series of "principal components" that join the two poles — corresponding to a series of Floer trajectories — with attached bubble components. Moreover, the equation on all the principal components has a Hamiltonian perturbation term. But apart from these rather minor extra complications, their construction is exactly as described here.

(ii) Here we have chosen to equip the virtual moduli cycle with the minimum amount of structure needed for it to represent a homology class. However, one could keep track of more structure, if this is appropriate. For example, in [Mc2] we consider fibrations $P \to S^2$ with symplectic fibers M and we need information on the strata of codimensions 3 and 4 in $\overline{\mathcal{M}}^\nu$ in order to relate the Gromov–Witten invariants of P to those of M. These strata correspond to stable maps whose domains have two components.

5.3 An example

So far there are rather few calculations of Gromov–Witten invariants using this general definition. Some results for manifolds that are symplectically fibered over S^2 are worked out in [Mc2]. We now work out an example in which the unperturbed moduli space $\overline{\mathcal{M}}$ is in fact a manifold, but it has the wrong dimension. In this situation one can often define a global obstruction bundle $E \to \overline{\mathcal{M}}$, and we will see that the cycle represented by the regularization $\overline{\mathcal{M}}^\nu$ is homologous to that represented by the zero set of a generic section of E. Hence one can calculate the desired invariants by investigating this zero set. The resulting global calculations are often best done in the context of algebraic geometry, because the topology of spaces of stable maps is so hard to understand.

Suppose that C is a symplectically embedded sphere in a 6-dimensional manifold X with $c_1(C) = 0$. (For example X might be a Calabi–Yau 3 fold.) We can choose J to be integrable near C, so that C is J-holomorphic and so that the normal bundle ν_C to C splits as the sum $\mathcal{O}(-1) \oplus \mathcal{O}(-1)$ of two line bundles each with Chern class -1. Then any J-holomorphic parametrization $h : S^2 \to X$ of C is regular, and C contributes $+1$ to the count of all J-holomorphic curves in class $A = [C]$. Let us work out the contribution of a k-fold cover of C to the invariant $n_X(H, H, H; kA)$ where H is a hypersurface such that $H \cdot A = 1$, at least in the case $k = 2$. According to Claire Voisin's proof in [V] of the Aspinwall-Morrison formula (cf [MS] §9.3), this contribution is 1 for all k.[17]

Since $c_1(C) = 0$ the formal dimension of the moduli space $\mathcal{M}_{0,3}(X, J, kA)$ is 6 for all k. We are interested in counting the number of elements of this moduli space whose three marked points meet representatives H_1, H_2, H_3 of the hypersurface H, where the H_i meet C in distinct points. Hence, if $\mathcal{M}_3(H, kA)$ denotes this cut down moduli space, the formal dimension of $\mathcal{M}_3(H, kA)$ is 0. However, when $k > 0$ multiple covers of the form $h_k = h \circ f_k$ (where $f_k : S^2 \to S^2$ has degree k) are never regular. One can see this from the usual regularity criterion in [MS] Lemma 3.5.1, or by using the fact that the actual moduli space always has positive dimension, even after the intersection conditions are imposed. For example, when $k = 2$, it is diffeomorphic to the moduli space $\overline{\mathcal{M}}_3$ considered in Example 3.3 and so is diffeomorphic to $\mathbb{C}P^2$. Moreover, because a neighborhood of C in X is biholomorphically equivalent to a neighborhood of the zero section in the bundle $\mathcal{O}(-1) \oplus \mathcal{O}(-1)$ over S^2, there are no other curves in class $2A$ near double covers of C except these double covers themselves. Hence $\mathcal{M}_3(H, 2A)$ has a component $\overline{\mathcal{M}}(C)$ that can be identified with $\mathbb{C}P^2$, and we want to understand what the regularization $\overline{\mathcal{M}}^{\nu}(C)$ of this component is.

According to [MS] Lemma 3.5.1, for each $\tilde{\tau}$ in the open stratum $\mathcal{S}_1$ of $\overline{\mathcal{M}}(C)$ the cokernel of $D\tilde{h}$ can be identified with the space of holomorphic sections of the rank 2-bundle $\mathcal{O}(-2)^* \otimes K \oplus \mathcal{O}(-2)^* \otimes K$ over S^2 where K is the canonical bundle $T^*(S^2) = \mathcal{O}(-2)$. (In general for k-fold maps this cokernel would be the sum of two copies of $\mathcal{O}(-k)^* \otimes K$ since the normal line bundles $\mathcal{O}(-1)$ pull back to $\mathcal{O}(-k)$.) Hence this is a trivial rank 2 complex bundle with a 2-dimensional family of (constant) sections. Hence at each point $\tilde{\tau}$ of $\mathcal{S}_1$ the cokernel R_τ has 2 complex dimensions. Moreover, it clearly splits as the sum of two line bundles one for each factor of ν_C.

Similarly, for $\tilde{\tau} \in \mathcal{S}_2$, the cokernel has two complex dimensions. Now it arises because of the compatibility condition $\xi_1(x) = \xi_2(x)$ put on the elements of the domain of $D\tilde{h}$ at the intersection point of the two components of Σ: see §5.1. Without this condition, the operator $D\tilde{h}$ would be surjective and with kernel equal to the tangent space to the reparametrization group G_R. The component of the equation that is tangential to C does not affect the image of $D\tilde{h}$ since one can alter the tangential component by adding to ξ_1 (or ξ_2) a vector field that is tangent to the action of G_R and hence in the kernel of $D\tilde{h}$. Hence the cokernel at $\tilde{\tau}$ can be identified with the normal space to C at $\tilde{h}(x)$. As τ varies over each line in the stratum $\mathcal{S}_2$, x varies over the curve C and so the cokernel bundle over this line is $\mathcal{O}(-1) \oplus \mathcal{O}(-1)$.

Altogether these fibers fit together to give a complex rank 2 bundle E over $\overline{\mathcal{M}}(C)$ that is called the obstruction bundle. We can use this bundle instead of R in the construction of the regularization. Because E is globally defined and all the groups Γ_τ are trivial, the sets $\widetilde{V}_I$ can be identified with the V_I, and multi-bundles and multi-sections are the same as usual bundles and sections. Hence a perturbation ν is simply a section of E and the corresponding regularization $\overline{\mathcal{M}}^{\nu_E}(C)$ is just the zero set of this section (with no branching and all labels equal to 1.) Hence the Gromov–Witten invariant that we are trying to calculate is simply the Euler number of E.

[17] She uses a different compactification of the moduli space $\mathcal{M}$ which is easier to calculate with, but has to work to show that she really is calculating the same Gromov–Witten invariants as defined here.

Therefore the question boils down to understanding the topological type of this obstruction bundle E. Note that because the image curve C is fixed, the gluing of $\mathcal{S}_2$ to $\mathcal{S}_1$ really takes place in the domain, i.e. in the space of curves (Σ, z_1, z_2, z_3). It is therefore very plausible that E is the sum of line bundles $E = L \oplus L'$, i.e. the fact that the normal bundle to C splits as a sum of line bundles is reflected in the structure of E. To prove this one has to investigate the gluing process in detail. Usually this is done locally near a point $\tilde{\tau}$ and the arguments show that the cokernel of $D\tilde{h}$ can be identified with the cokernel $D\tilde{h}_a$ of the glued map. Here we need to look globally as $\tilde{\tau}$ varies over a line in $\mathcal{S}_2$ in order to understand the structure of E near this line. Similar questions are discussed in [Mc1] §4.2. In the notation used there, the cokernel bundle K over $\mathcal{S}_2$ is precisely $\mathcal{O}(-1) \oplus \mathcal{O}(-1)$ and it maps onto the cokernel of $D\tilde{h}_a$ for small a by the arguments given there. Moreover the map that identifies it with this cokernel can be chosen to respect the splitting of the cokernel bundle over $\mathcal{S}_1$. Hence E is the sum of line bundles $L \oplus L'$. It is then clear that $L = L'$ because of the local automorphism of X near C that interchanges the two factors of its normal bundle. Hence L restricts to $\mathcal{O}(-1)$ over each line and so must be the universal line bundle over $\mathbb{C}P^2$. In particular the Euler number of E is 1.

These arguments should give the flavor of the topological questions that arise in this kind of context. Some related global calculations are carried out in [Mc1].

6 More remarks on the topology of $\mathcal{W}$

After the publication of this article, Hofer pointed out to me that there are subtle topological questions concerning the moduli spaces of stable maps $\mathcal{M}$ and $\mathcal{W}$ that need further discussion. These are caused by the awkward fact that the action of the reparametrization group $\mathrm{PSL}(2, \mathbb{C})$ on the space of $L^{k,p}$-smooth maps $S^2 \to X$ is not continuous. For example, if one takes $k = 0, p = 2$ and defines $R_\theta : S^2 \to S^2$ to be rotation about a fixed axis through angle θ then

$$\sup_{g : \|g\|_{L^2} = 1} \|R_\theta(g) - g\|_{L^2} = \sqrt{2}, \quad \text{for all } \theta > 0.$$

It would be continuous if one thought of R_θ as mapping a space of $L^{k,p}$-functions to a space of $L^{k-1,p}$-functions, i.e. if one could afford to lose a derivative. It is also continuous on spaces of C^∞-maps.

In view of this, it seems worthwhile to give more details of the arguments that establish that $\mathcal{W}$ is indeed a space with two topologies as claimed above.

First of all, let us consider the local uniformizers $(\widetilde{U}^D_\tau, \Gamma_\tau, \pi_\tau)$ defined in §2.1. Here is a sketch of an argument, supplied to me by Liu, that shows that the action of $\gamma \in \Gamma_\tau$ on a sufficiently small neighborhood $\widetilde{U}^D_\tau$ is continuous. Note that $\phi^\gamma_{g'}$ can be made as C^∞-close to ϕ^γ_g as we want, by choosing g' sufficiently $L^{k,p}$-close to g so that the point $w_1'(g')$ is sufficiently C^0-close to $w_1'(g)$. Secondly, in order to make $d(g' \circ \phi^\gamma_{g'}, g \circ \phi^\gamma_g)$ less than some given δ, consider the inequality:

$$d(g' \circ \phi^\gamma_{g'}, g \circ \phi^\gamma_g) \le d(g' \circ \phi^\gamma_{g'}, g \circ \phi^\gamma_{g'}) + d(g \circ \phi^\gamma_{g'}, g \circ \phi^\gamma_g).$$

The first term is easy to make $\le \delta/2$ because $\phi^\gamma_{g'}$ is fixed and is C^∞ close to ϕ^γ_g. To make the second term $\le \delta/2$ we replace g by a C^∞ smooth approximation h, chosen independently of g', and use the inequality:

$$d(g \circ \phi^\gamma_{g'}, g \circ \phi^\gamma_g) \le d(h \circ \phi^\gamma_{g'}, h \circ \phi^\gamma_g) + d(h \circ \phi^\gamma_{g'}, g \circ \phi^\gamma_{g'}) + d(h \circ \phi^\gamma_g, g \circ \phi^\gamma_g).$$

By choosing h sufficiently close to g the last two terms can each be made $\le \delta/6$ for all g' sufficiently close to g. Since h is C^∞ the first term also converges to zero as g' tends to g in the $L^{k,p}$-norm. Therefore we can also make g' so close to g that the first term is $\le \delta/6$.

Moreover, if $U^D_\tau = \widetilde{U}^D_\tau/\Gamma_\tau$ is given the quotient topology, a similar argument shows that the two topologies induced on an overlap $U^D_\tau \cap U^D_{\tau'}$ agree. Hence the stratum $\mathcal{W}^D$ has a nice smooth (i.e. $L^{k,p}$) topology, and we define $\mathcal{W}_{sm}$ to be the disjoint union of these strata $\mathcal{W}^D$.

The next remarks amplify the description given in §3.2 of the topology of $\mathcal{W}$. The neighborhood $\widetilde{U}_\tau$ of $\tilde\tau = (\Sigma, \tilde h, z_1, \dots, z_k)$ contains all points $\tilde\tau' = (\Sigma_a, \tilde h', z_1', \dots, z_k')$ where $a \in V$, $\tilde h'$ is close to $\tilde h \circ \psi_a$, and the $\psi_a(z_j')$ are close to z_j. Here the distance between $\tilde h'$ and $\tilde h \circ \psi_a$ should be the sum of their C^0-distance and the $L_{k,p}$-distance between their restrictions to a set of the form $\psi_a^{-1}(K)$, where $K \subset \Sigma - \{\text{doublepoints}\}$ is a compact subset containing the points z_j and w_j. The set $\widetilde{U}_\tau$ is given the corresponding topology (where K is allowed to range over all possible compact sets).

It is not hard to extend the action of the symmetry group Γ_τ to a continuous action on $\widetilde{U}_\tau$ such that the quotient

$$U_\tau = \widetilde{U}_\tau/\Gamma_\tau$$

can be identified with a subset of $\mathcal{W}$. Moreover the induced topologies on the overlaps $U_\tau \cap U_{\tau'}$ agree. Hence we can take $\mathcal{W}$ to be the union of all the U_τ for $\tau \in \mathcal{M}$ with the induced topology. (Note that in fact we only need to consider a finite number of the U_τ since $\mathcal{M}$ is compact.)

References

[FO] K. Fukaya and K. Ono, Arnold conjecture and Gromov–Witten invariants, *Topology* **38** (1999), 933–1048.

[LiT] Jun Li and Gang Tian, Virtual moduli cycles and Gromov–Witten invariants for general symplectic manifolds, *Topics in Symplectic* 4*-manifolds (Irvine CA 1996),* Internat. Press, Cambridge, MA (1998), pp 47–83.

[LiuT1] Gang Liu and Gang Tian, Floer homology and Arnold conjecture, *Journ. Diff. Geom.*, **49** (1998), 1–74.

[LiuT2] Gang Liu and Gang Tian, On the equivalence of multiplicative structures in Floer Homology and Quantum Homology, *Acta Math. Sinica* **15** (1999).

[Mc1] D. McDuff, Almost complex structures on $S^2 \times S^2$, preprint (1998), to appear in *Duke Math. Journal*

[Mc2] D. McDuff, Quantum homology of fibrations over S^2, preprint (1998). to appear in *Internat. Journ. of Math.*

[MS] D. McDuff and D.A. Salamon, *J-holomorphic curves and quantum cohomology*, Amer Math Soc Lecture Notes #6, Amer. Math. Soc. Providence (1994).

[MS2] D. McDuff and D.A. Salamon, *J-holomorphic curves and symplectic topology*, Amer Math Soc Colloquium Publications, (2004)

[R] Y. Ruan, Virtual neighborhoods and pseudoholomorphic curves, *Topics in Symplectic* 4*-manifolds (Irvine CA 1996),* Internat. Press, Cambridge, MA (1998), pp 101–116.

[RT] Y. Ruan and G. Tian, A mathematical theory of quantum cohomology. Journ Diff Geo **42** (1995), 259–367.

[S] D. Salamon, Lectures on Floer theory, *Proceedings of the IAS/Park City Summer Institute, 1997*, ed Eliashberg and Traynor, Amer. Math. Soc., Providence, RI. (1999), 143–229.

[Sb] B. Siebert, Gromov–Witten invariants for general symplectic manifolds, preprint (1997)

[V] C. Voisin, A mathematical proof of a formula of Aspinwall and Morrison, *Compositio Math.*, **104**, (1996), 135–151.