

The symplectic sum formula for Gromov-Witten invariants

By ELENY-NICOLETA IONEL and THOMAS H. PARKER*

Abstract

In the symplectic category there is a ‘connect sum’ operation that glues symplectic manifolds by identifying neighborhoods of embedded codimension two submanifolds. This paper establishes a formula for the Gromov-Witten invariants of a symplectic sum $Z = X \# Y$ in terms of the relative GW invariants of X and Y . Several applications to enumerative geometry are given.

Gromov-Witten invariants are counts of holomorphic maps into symplectic manifolds. To define them on a symplectic manifold (X, ω) one introduces an almost complex structure J compatible with the symplectic form ω and forms the moduli space of J -holomorphic maps from complex curves into X and the compactified moduli space, called the space of stable maps. One then imposes constraints on the stable maps, requiring the domain to have a certain form and the image to pass through fixed homology cycles in X . When the correct number of constraints is imposed there are only finitely many maps satisfying the constraints; the (oriented) count of these is the corresponding GW invariant. For complex algebraic manifolds these symplectic invariants can also be defined by algebraic geometry, and in important cases the invariants are the same as the curve counts that are the subject of classical enumerative algebraic geometry.

In the past decade the foundations for this theory were laid and the invariants were used to solve several long-outstanding problems. The focus now is on finding effective ways of computing the invariants. One useful technique is the method of ‘splitting the domain’, in which one localizes the invariant to the set of maps whose domain curves have two irreducible components with the constraints distributed between them. This produces recursion relations relating the desired GW invariant to invariants with lower degree or genus. This paper establishes a general formula describing the behavior of GW invariants under the analogous operation of ‘splitting the target’. Because we

*The research of both authors was partially supported by the NSF. The first author was also supported by a Sloan Research Fellowship.

work in the context of symplectic manifolds the natural splitting of the target is the one associated with the symplectic cut operation and its inverse, the symplectic sum.

The symplectic sum is defined by gluing along codimension two submanifolds. Specifically, let X be a symplectic $2n$ -manifold with a symplectic $(2n-2)$ -submanifold V . Given a similar pair (Y, V) with a symplectic identification between the two copies of V and a complex anti-linear isomorphism between the normal bundles $N_X V$ and $N_Y V$ of V in X and in Y , we can form the symplectic sum $Z = X \#_V Y$. Our main theorem is a ‘Symplectic Sum Formula’ which expresses the GW invariants of the sum Z in terms of the relative GW invariants of (X, V) and (Y, V) introduced in [IP4].

The symplectic sum is perhaps more naturally seen not as a single manifold but as a family depending on a ‘squeezing parameter’. In Section 2 we construct a family $Z \rightarrow D$ over the disk whose fibers Z_λ are smooth and symplectic for $\lambda \neq 0$ and whose central fiber Z_0 is the singular manifold $X \cup_V Y$. In a neighborhood of V , the total space Z is $N_X V \oplus N_Y V$ and the fiber Z_λ is defined by the equation $xy = \lambda$ where x and y are coordinates in the normal bundles $N_X V$ and $N_Y V \cong (N_X V)^*$. The fibration $Z \rightarrow D$ extends away from V as the disjoint union of $X \times D$ and $Y \times D$, and the entire fibration Z can be given an almost Kähler structure. The smooth fibers Z_λ , depicted in Figure 1, are symplectically isotopic to one another; each is a model of the symplectic sum.

The overall strategy for proving the symplectic sum formula is to relate the holomorphic maps into Z_0 , which are simply maps into X and Y which match along V , with the holomorphic maps into Z_λ for λ close to zero. This strategy involves two parts: limits and gluing. For the limiting process we consider sequences of stable maps into the family Z_λ of symplectic sums as the ‘neck size’ $\lambda \rightarrow 0$. In particular, these are stable maps into a compact region of the almost Kähler manifold Z , so that the compactness theorem for stable maps applies, giving limit maps into the singular manifold Z_0 obtained by identifying X and Y along V . Along the way several things become apparent.

First, the limit maps are holomorphic only if the almost complex structures on X and Y match along V . To ensure this we impose the “ V -compatibility” condition (1.10) on the almost complex structure. But there is a price to pay for that specialization. In the symplectic theory of Gromov-Witten invariants we are free to perturb (J, ν) without changing the invariant; this freedom can be used to ensure that intersections are transverse. After imposing the V -compatibility condition, we can no longer perturb (J, ν) along V at will, and hence we cannot assume that the limit curves are transverse to V . In fact, the images of the components of the limit maps meet V at points with well-defined multiplicities and, worse, some components may be mapped entirely into V .

To count stable maps into Z_0 we look first on the X side and ignore the maps which have marked points, double points, or whole components mapped into V . The remaining “ V -regular” maps form a moduli space which is the union of components $\mathcal{M}_s^V(X)$ labeled by the multiplicities $s = (s_1, \dots, s_\ell)$ of the intersection points with V . We showed in [IP4] how these spaces $\mathcal{M}_s^V(X)$ can be compactified and used to define relative Gromov-Witten invariants GW_X^V . The definitions are briefly reviewed in Section 1.




Figure 1. Limiting curves in $Z_\lambda = X \#_\lambda Y$ as $\lambda \rightarrow 0$.

Second, as Figure 1 illustrates, connected curves in Z_λ can limit to curves whose restrictions to X and Y are not connected. For that reason the GW invariant, which counts stable curves from a connected domain, is not the appropriate invariant for expressing a sum formula. Instead one should work with the ‘Gromov-Taubes’ invariant GT, which counts stable maps from domains that need not be connected. Thus we seek a formula of the general form

$$(0.1) \quad \text{GT}_X^V * \text{GT}_Y^V = \text{GT}_Z$$

where $*$ is some operation that adds up the ways curves on the X and Y sides match and are identified with curves in Z_λ . That necessarily involves keeping track of the multiplicities s and the homology classes. It also involves accounting for the limit maps with nontrivial components in V ; such curves are not counted by the relative invariant and hence do not contribute to the left side of (0.1). We postpone this issue by first analyzing limits of curves which are δ -flat in the sense of Definition 3.1.

A more precise analysis reveals a third complication: the squeezing process is not injective. In Section 5 we again consider a sequence of stable maps f_n into Z_λ as $\lambda \rightarrow 0$, this time focusing on their behavior near V , where the f_n do not uniformly converge. We form renormalized maps \hat{f}_n and prove that both the domains and the images of the renormalized maps converge. The images converge nicely according to the leading order term of their Taylor expansions, but the domains converge only after we fix certain roots of unity.

These roots of unity are apparent as soon as one writes down formulas. Each stable map $f : C \rightarrow Z_0$ decomposes into a pair of maps $f_1 : C_1 \rightarrow X$ and $f_2 : C_2 \rightarrow Y$ which agree at the nodes of $C = C_1 \cup C_2$. For a specific example, suppose that f is such a map that intersects V at a single point p

with multiplicity three. Then we can choose local coordinates z on C_1 and w on C_2 centered at the node, and coordinates x on X and y on Y so that f_1 and f_2 have expansions $x(z) = az^3 + \cdots$ and $y(w) = bw^3 + \cdots$. To find maps into Z_λ near f , we smooth the domain C to the curve C_μ given locally near the node by $zw = \mu$ and require that the image of the smoothed map lie in Z_λ , which is locally the locus of $xy = \lambda$. In fact, the leading terms in the formulas for f_1 and f_2 define a map $F : C_\mu \rightarrow Z_\lambda$ whenever

$$\lambda = xy = az^3 \cdot bw^3 = ab(zw)^3 = ab\mu^3$$

and conversely every family of smooth maps which limit to f satisfies this equation in the limit (cf. Lemma 5.3). Thus λ determines the domain C_μ up to a cube root of unity. Consequently, this particular f is, at least *a priori*, close to *three* smooth maps into Z_λ — a ‘cluster’ of order three.

Other maps f into Z_0 have larger associated clusters (the order of the cluster is the product of the multiplicities with which f intersects V). The maps within a cluster have the same leading order formula but have different smoothings of the domain. As $\lambda \rightarrow 0$ the cluster coalesces, limiting to the single map f .

This clustering phenomenon greatly complicates the analysis. To distinguish the curves within each cluster and make the analysis uniform in λ as $\lambda \rightarrow 0$, it is necessary to use ‘rescaled’ norms and distances which magnify distances as the clusters form. With the right choice of norms, the distances between the maps within a cluster are bounded away from zero as $\lambda \rightarrow 0$ and become the fiber of a covering of the space of limit maps. Sections 4–6 introduce the required norms, first on the space of curves, then on the space of maps.

For maps we use a Sobolev norm weighted in the directions perpendicular to V ; the weights are chosen so the norm dominates the C^0 distance between the renormalized maps \hat{f} . On the space of curves we require a stronger metric than the usual complete metrics on $\overline{\mathcal{M}}_{g,n}$. In Section 4 we define a complete metric on $\overline{\mathcal{M}}_{g,n} \setminus \mathcal{N}$ where \mathcal{N} is the set of all nodal curves. In this metric the distance between two sequences that approach \mathcal{N} from different directions (corresponding to the roots of unity mentioned above) is bounded away from zero; thus this metric separates the domain curves of maps within a cluster. This construction also leads to a compactification of $\overline{\mathcal{M}}_{g,n} \setminus \mathcal{N}$ in which the stratum \mathcal{N}_ℓ of ℓ -nodal curves is replaced by a bundle over \mathcal{N}_ℓ whose fiber is the real torus T^ℓ .

The limit process is reversed by constructing a space of approximately holomorphic maps and showing it is diffeomorphic to the space of stable maps into Z_λ . The space of approximate maps is described in Section 6, first intrinsically, then as a subset $\text{Model}_s(Z_\lambda)$ of the space of maps. For each s and λ it is a covering of the space $\mathcal{M}_s(Z_0)$ of the δ -flat maps into Z_0 that meet

V at points with multiplicities s . The fibers of this covering are the clusters described above; they are distinct maps into Z_λ which converge to the same limit as $\lambda \rightarrow 0$.

From there the analysis follows a standard technique that goes back to Taubes and Donaldson: correct the approximate maps to true holomorphic maps by constructing a partial right inverse to the linearization D and applying a fixed point theorem. This involves (a) showing that the operator D^*D is uniformly invertible as $\lambda \rightarrow 0$, and (b) proving *a priori* that every solution is close to an approximate solution, close enough to be in the domain of the fixed point theorem. Proposition 9.3 shows that (b) follows from the renormalization analysis of Section 4. But the eigenvalue estimate (a) proves to be surprisingly delicate and seems to succeed only with a very specific choice of norms.

The difficulty, of course, is that Z_λ becomes singular along V as $\lambda \rightarrow 0$. However, for small λ the bisectional curvature in the neck region is negative; a Bochner formula then shows that eigenfunctions with small eigenvalue cannot be concentrating in the neck. One can then reason that since the cokernel of D vanishes on Z_0 for generic V -compatible (J, ν) it should also vanish on Z_λ for small λ . We make that reasoning rigorous by introducing exponential weight functions into the norms, thereby making the linearizations D_λ a continuous family of Fredholm maps. That in turn necessitates further work on the Bochner formula, bounding the additional term that arises from the derivative of the weight functions. These estimates are carried out in Section 8.

The upshot of the analysis is a diffeomorphism between the approximate moduli space and the true moduli spaces

$$\text{Model}_s(Z_\lambda)_s \xrightarrow{\cong} \mathcal{M}_s(Z_\lambda)$$

which intertwines with the attaching map of the domains and the evaluation map into the target (Theorem 10.1). We then pass to homology, comparing and keeping track of the homology classes of the maps, the domains, and the constraints. This involves several difficulties, all ultimately due to the fact that $H_*(Z_\lambda)$ is different from both $H_*(Z_0)$ and $H_*(X) \oplus H_*(Y)$. This is sorted out in Section 10, where we define the convolution operation and prove a first Symplectic Sum Theorem: a formula like (0.1) holds when all stable maps are δ -flat.

In Sections 11 and 12 we remove the δ -flatness assumption by partitioning the neck into a large number of segments and using the pigeon-hole principle as in Wiczeorek [W]. For that we construct spaces $Z_\lambda^N(\mu_1, \dots, \mu_{2N+1})$, each symplectically isotopic to Z_λ . As $(\mu_1, \dots, \mu_{2N+1}) \rightarrow 0$ these degenerate to the singular space obtained by connecting X to Y through a series of $2N$ copies of the rational ruled manifold \mathbb{P}_V obtained by adding an infinity section to the normal bundle to V . An energy bound shows that for large N each map into $Z_\lambda(\mu_1, \dots, \mu_{2N+1})$ must be flat in most necks. Squeezing some or all of the flat

necks decomposes the curves in Z_λ into curves in X joined to curves in Y by a chain of curves in intermediate spaces \mathbb{P}_V . The limit maps are then δ -flat, so that formula (0.1) applies to each. This process counts each stable map many times (there are many choices of where to squeeze) and in fact gives an open cover of the moduli space. Working through the combinatorics and inverting a power series, we show that the total contribution of the entire neck region between X and Y is given by a certain GT invariant of \mathbb{P}_V — the S -matrix of Definition 11.3.




Figure 2. $Z_\lambda(\mu, \mu, \mu)$ for $|\mu| \ll |\lambda|$

The S -matrix keeps track of how the genus, homology class, and intersection points with V change as the images of stable maps pass through the middle region of Figure 2. Observing this back in the model of Figure 1, one sees these quantities changing abruptly; as the maps pass through the neck, they are “scattered” by the neck. The scattering occurs when some of the stable maps contributing to the GT invariant of Z_λ have components that lie entirely in V in the limit as $\lambda \rightarrow 0$. Those maps are not V -regular, so are not counted in the relative invariants of X or Y . But by moving to the spaces of Figure 2 this complication can be analyzed and related to the relative invariants of the ruled manifold \mathbb{P}_V .

The S -matrix is the final subtlety. With it in hand, we can at last state our main result.

SYMPLECTIC SUM THEOREM. *Let Z be the symplectic sum of (X, V) and (Y, V) and suppose that $\alpha \in \mathbb{T}^*(Z)$ splits as (α_X, α_Y) as in Definition 10.5. Then the GT invariant of Z is given in terms of the relative invariants of X and Y by*

$$(0.2) \quad \text{GT}_Z(\alpha) = \text{GT}_X^V(\alpha_X) * S_V * \text{GT}_Y^V(\alpha_Y)$$

where $*$ is the convolution operation (10.6) and S_V is the S -matrix (11.3).

A detailed statement of this theorem is given in Section 12 and its extension to general constraints α is discussed in Section 13. We actually state and prove (0.2) as a formula for the *relative* invariants of Z in terms of the relative invariants of X and Y (Theorem 12.3). In that more general form the formula can be iterated.

Of course, (0.2) is of limited use unless we can compute the relative invariants of X and Y and the associated S -matrix. That turns out to be perfectly feasible, at least for simple spaces. In Section 14 we build a collection of two and four dimensional spaces whose relative GT invariants we can compute. We also prove that the S -matrix is the identity in several cases of particular interest.

The last section presents applications. The examples of Section 14 are used as building blocks to give short proofs of three recent results in enumerative geometry: (a) the Caporaso-Harris formula for the number of nodal curves in \mathbb{P}^2 [CH], (b) the formula for the Hurwitz numbers counting branched covers of \mathbb{P}^1 ([GJV] [LZZ]), and (c) the “quasimodular form” expression for the rational enumerative invariants of the rational elliptic surface ([BL]). In hindsight, our proofs of (a) and (b) are essentially the same as those in the literature; using the symplectic sum formula makes the proof considerably shorter and more transparent, but the key ideas are the same. Our proof of (c), however, is completely different from that of Bryan and Leung in [BL]. It is worth outlining here.

The rational elliptic surface E fibers over \mathbb{P}^1 with a section s and fiber f . For each $d \geq 0$ consider the invariant GW_d which counts the number of connected rational stable maps in the class $s + df$. Bryan and Leung showed that the generating series $F_0(t) = \sum_d \mathrm{GW}_d t^d$ is

$$(0.3) \quad F_0(t) = \left(\prod_d \frac{1}{1-t^d} \right)^{12}.$$

This formula is related to the work of Yau-Zaslow [YZ] and is one of the simplest instances of some general conjectures concerning counts of nodal curves in complex surfaces — see [Go].

While the intriguing form (0.3) appears in [BL] for purely combinatorial reasons, it arises in our proof because of a connection with elliptic curves. In fact, our proof begins by relating F_0 to a similar series H which counts *elliptic* curves in E . We then regard E as the fiber sum $E \# (T^2 \times S^2)$ and apply the symplectic sum formula. The relevant relative invariant on the $T^2 \times S^2$ side is easily seen to be the generating function $G(t)$ for the number of degree d coverings of the torus T^2 by the torus. The symplectic sum formula reduces to a differential equation relating $F_0(t)$ with $G(t)$, and integration yields the quasimodular form (0.3). The details, given in Section 15.3, are rather formal; the needed geometric input is mostly contained in the symplectic sum formula.

All three of the applications in Section 15 use the idea of ‘splitting the target’ mentioned at the beginning of this introduction. Moreover, all three follow from rather simple cases of the Symplectic Sum Theorem — cases where the S -matrix is the identity and where at least one of the relative invariants

in (0.2) is readily computed by elementary methods. The full strength of the symplectic sum theorem has not yet been used.

This paper is a sequel to [IP4]; together with [IP4] it gives a complete detailed exposition of the results announced in [IP3]. Further applications have already appeared in [IP2] and [I]. A. M. Li and Y. Ruan also have a sum formula [LR]. Y. Eliashberg, A. Givental, and H. Hofer are developing a general theory for invariants of symplectic manifolds glued along contact boundaries [EGH].

Contents

1. GW and GT invariants
2. Symplectic sums
3. Degenerations of symplectic sums
4. The space of curves
5. Renormalization at the nodes
6. The space of approximate maps
7. Linearizations
8. The eigenvalue estimate
9. The gluing diffeomorphism
10. Convolutions and the sum formula for flat maps
11. The space \mathbb{P}_V and the S-matrix
12. The general sum formula
13. Constraints passing through the neck
14. Relative GW invariants in simple cases
15. Applications of the sum formula
- Appendix: Expansions of relative GT invariants

1. GW and GT invariants

For stable maps and their associated invariants we will use the definitions and notation of [IP4], which build on those of Ruan-Tian [RT1] and [RT2]. Thus we work in the context of a closed symplectic manifold (X, ω) with an almost complex structure J and Riemannian metric g which are compatible in the sense that

$$(1.1) \quad g(v, w) = \omega(v, Jw) \quad \forall v, w \in TX.$$

In summary, the definitions are as follows. A *bubble domain* B is a finite connected union of smooth oriented 2-manifolds B_i joined at nodes together with n marked points p_i , none of which are nodes. Collapsing the unstable components to points gives a connected domain $\text{st}(B)$. Let $\overline{\mathcal{U}}_{g,n} \rightarrow \overline{\mathcal{M}}_{g,n}$ be the universal curve over the Deligne-Mumford space of genus g curves with n

marked points. We can put a complex structure j on B , well-defined up to diffeomorphism, by specifying an orientation-preserving map $\phi_0 : \text{st}(B) \rightarrow \overline{\mathcal{U}}_{g,n}$ which is a diffeomorphism onto a fiber of $\overline{\mathcal{U}}_{g,n}$ and taking the standard complex structure on the unstable (rational) components. We will often write C for the curve (B, j, p_1, \dots, p_n) and use the notation (f, C) or $f : C \rightarrow X$ instead of (f, ϕ) .

A (J, ν) -holomorphic map from B is a map $(f, \phi) : B \rightarrow X \times \overline{\mathcal{U}}_{g,n}$ where $\phi = \phi_0 \circ \text{st}$ and whose restriction to each irreducible component B_i of B satisfies

$$(1.2) \quad \bar{\partial}_J f = (f, \phi)^* \nu.$$

Here $\bar{\partial}_J f$ means $\frac{1}{2}(df + J_f df j_\phi)$ and, after an embedding $\overline{\mathcal{U}}_{g,n} \subset \mathbb{P}^N$ is fixed, ν is a section of the bundle $\text{Hom}(\pi_1^* T\mathbb{P}^N, \pi_2^* TX)$ over $X \times \mathbb{P}^N$ satisfying $J\nu = -\nu J_{\mathbb{P}^N}$. Alternatively, we can define an almost complex structure \hat{J} on $X \times \overline{\mathcal{U}}_{g,n}$ by $\hat{J}(u, v) = J(u - 2\nu(v)) + J_{\mathbb{P}^N}$. Equation (1.2) is the \hat{J} -holomorphic map equation for the map (f, ϕ) ; it holds if and only if

$$(1.3) \quad F = (f, \phi) \text{ satisfies } \bar{\partial}_{\hat{J}} F = 0.$$

Furthermore, when ν is small \hat{J} is tamed by the symplectic form $\hat{\omega} = \omega \oplus \omega_{\mathbb{P}^N}$; specifically, (1.1) implies that $\hat{\omega}((u, v), J(u, v)) \geq \frac{3}{4}|(u, v)|^2$ whenever $\|\nu\|_\infty \leq 1/4$. With this tamed condition there are standard elliptic estimates on $F = (f, \phi)$; see [PW] and [RT1]. In particular, the energy

$$(1.4) \quad E(f, \phi) = \frac{1}{2} \int |df|^2 + |d\phi|^2$$

is related to the topological quantity

$$(1.5) \quad E(F) = \frac{1}{2} \int |dF|^2 = \int_{F(B)} \hat{\omega} = \omega([f(B)]) + \omega_{\mathbb{P}^N}([\phi(B)])$$

by $\frac{3}{4}E(F) \leq E(f, \phi) \leq \frac{4}{3}E(F)$. Consequently, we will assume throughout that $\|\nu\|_\infty \leq 1/4$.

A *stable map* is a (J, ν) -holomorphic map for which the energy $E(f, \phi)$ is positive on each component B_i ; this means that each B_i is either a stable curve or the restriction of f to B_i is nontrivial in homology. A stable map $F = (f, \phi)$ is *irreducible* if $F^{-1}(F(x)) = \{x\}$ for generic points x (this is automatically true if all the domain components are stable). More generally, a stable map $F = (f, \phi)$ is *admissible* if it is irreducible when restricted to the union of its unstable domain components B_i with $0 < K_X[f(B_i)]$.

For generic (J, ν) the moduli space $\overline{\mathcal{M}}_{g,n}(X, A)^*$ of irreducible stable (J, ν) -holomorphic maps representing a class $A \in H_2(X)$ is a smooth orbifold of (real) dimension

$$(1.6) \quad -2K_X[A] + \frac{1}{2}(\dim X - 6)\chi + 2n$$

where $\chi = 2 - 2g$ is the Euler characteristic of the domain. For all (J, ν) the moduli space $\overline{\mathcal{M}}_{g,n}(X, A)$ of all stable maps is compact, and stabilization and evaluation at the marked points define a continuous map

$$\overline{\mathcal{M}}_{g,n}(X, A) \xrightarrow{\text{st} \times \text{ev}} \overline{\mathcal{M}}_{g,n} \times X^n$$

where we make the conventions that X^n is a single point when $n = 0$ and that in the unstable range (i.e. when $2g - 3 + n < 0$), we take

$$\overline{\mathcal{M}}_{0,n} = \overline{\mathcal{M}}_{0,3} \quad \text{for } n \leq 2 \quad \text{and } \overline{\mathcal{M}}_{1,0} = \overline{\mathcal{M}}_{1,1}.$$

Note that, with this convention, the complex dimension of $\overline{\mathcal{M}}_{g,n}$ in the stable range is $3g - 3 + n$ while in the unstable range it is g .

When all maps in the moduli space are admissible for generic (J, ν) the image of $\overline{\mathcal{M}}_{g,n}(X, A)$ specifies a homology class called the Gromov-Witten invariant

$$\text{GW}_{X,A,g,n} \in H_*(\overline{\mathcal{M}}_{g,n} \times X^n).$$

These invariants can be assembled into a single invariant by setting $\overline{\mathcal{M}} = \bigcup_{g,n} \overline{\mathcal{M}}_{g,n}$, and introducing variables λ to keep track of the Euler class and t_A satisfying $t_A t_B = t_{A+B}$ to keep track of A . The total GW invariant of (X, ω) is then the formal series

$$(1.7) \quad \text{GW}_X = \sum_{A,g,n} \frac{1}{n!} \text{GW}_{X,A,g,n} t_A \lambda^{2g-2}$$

whose coefficients are multilinear functions on $H^*(\overline{\mathcal{M}}) \otimes \mathbb{T}^*(X)$ where $\mathbb{T}^*(X)$ denotes the total tensor algebra $\mathbb{T}(H^*(X))$. This in turn defines the ‘‘Gromov-Taubes’’ invariant

$$\text{GT}_X = e^{\text{GW}_X}$$

whose coefficients count holomorphic curves whose domains need not be connected (as occur in [T]).

In summary, the results of [RT1] and [RT2] show that generically the moduli spaces of irreducible stable maps are orbifolds and the GW invariants are defined when all maps in $\overline{\mathcal{M}}_{g,n}(X, A)$ are admissible for generic (J, ν) . In practice it is convenient to assume that (X, ω) is *semipositive*: there are no spherical classes $A \in H_2(X)$ with $\omega(A) > 0$ and $0 < 2K_X A \leq \dim X - 6$. For semipositive manifolds all moduli spaces are generically admissible, so that the GW and GT invariants are defined for all n , g and A . On the other hand, one can sometimes show that $\overline{\mathcal{M}}_{g,n}(X, A)$ is admissible for specific n , g and A even though the manifold is not semipositive.

Remark 1.1 (Local stabilization). If f is a stable map, its domain might have unstable components B_i . However, we can stabilize the domains of stable maps close to f as follows. Introduce k additional marked points on the domain such that each B_i has at least three special points, and so it is stable. Then,

for each of the new points on each B_i , choose cycles Γ_{ij} whose homological intersection d_{ij} with $f_*([B_i])$ is nonzero, and which are transverse to f and intersect $f(B_i)$ at distinct points. Each stable map close to f then lifts to $d = \prod_{i,j} d_{ij}$ maps (counted with sign) that take the new marked points onto the corresponding cycles Γ_{ij} . The lift of a (J, ν) -holomorphic map is a $(J, \pi^*\nu)$ -holomorphic map with a stable domain, where π is the map $\overline{\mathcal{M}}_{g,n+k} \rightarrow \overline{\mathcal{M}}_{g,n}$ that forgets the last k marked points.

In the analytic arguments of Sections 5–9, which are local on the space of maps, we will always assume that the domains of the maps have been locally stabilized by this lifting procedure. This allows us to work locally as if all domains were stable, with one important caveat: because the lifted maps are only $(J, \pi^*\nu)$ -holomorphic, the moduli spaces are generically orbifolds only near lifts of *irreducible* maps f . For that reason, irreducibility appears as a technical assumption in some results in Sections 5–10.

It is likely that this irreducibility assumption could be removed by turning on a further generic perturbation ν on $\overline{\mathcal{U}}_{g,n+k}$. There are several approaches to doing that; see for example in [LT].

The dimension (1.6) is the index of the linearization of the (J, ν) -holomorphic equation, which is obtained as follows. A variation of a map f is specified by a $\xi \in \Gamma(f^*TX)$, thought of as a vector field along the image, and a variation in the curve $C = (B, j, p_1, \dots, p_n)$ is specified by

$$(1.8) \quad h \in T_C \overline{\mathcal{M}}_{g,n} \cong H_j^{0,1} \left(TB \otimes \mathcal{O} \left(-\sum p_i \right) \right)$$

(tensoring with $\mathcal{O}(-p)$ accounts for the variation in the marked point p ; the correspondence between h and the variation of the map ϕ is described in Section 4). Let $\Lambda^{01}(f^*TX)$ be the vector bundle of all anti- (J, j) linear homomorphisms from TC to f^*TX . Calculating the variation in the path

$$(1.9) \quad (f_t, C_t) = (\exp_f(t\xi), (j, p_1, \dots, p_n) + th)$$

one finds that the linearization at (f, C) is the operator

$$(1.10) \quad D_{f,C} : \Gamma(f^*TX) \oplus T_C \overline{\mathcal{M}}_{g,n} \rightarrow \Gamma(\Lambda^{01}(f^*TX))$$

given by $D_{f,C}(\xi, h) = L_{f,C}(\xi) + \frac{1}{2}Jf_*h$ with

$$(1.11) \quad L_{f,C}(\xi)(w) = \frac{1}{2} \left[\nabla_w \xi + J \nabla_{jw} \xi + (J \nabla_\xi J)(f_*w - \Phi_f(w) - 2\nu(w)) \right] - (\nabla_\xi \nu)(w)$$

where w is a vector tangent to the domain, ∇ is the pullback connection on f^*TX , $\nu(w)$ means $\nu(\phi_*w)$, and $\Phi_f = \overline{\partial}_J f - \nu$. Writing L as the sum of its J -linear component $\frac{1}{2}(L - J L J)$ and its J -antilinear component, we have

$$(1.12) \quad L_{f,C}(\xi)(w) = \overline{\partial}_{f,C}^J \xi(w) + T_{f,C}(\xi, w).$$

Here $\bar{\partial}_{f,C}^J$ is a J -linear first order operator; so for complex valued functions ϕ

$$(1.13) \quad L_{f,C}(\phi\xi) = \bar{\partial}\phi \cdot \xi + \phi L_{f,C}(\xi) + (\bar{\phi} - \phi)T_{f,C}(\xi, w).$$

The term $T_{f,C}(\xi, w)$ is given in terms of the Nijenhuis tensor N_J of J on X and the tensors $T_\nu(\xi, w) = J(\nabla_{\nu(w)}J)\xi - (\nabla_\xi\nu)(w) - (J\nabla_{J\xi}\nu)(w)$ and $\widehat{\nabla J}(\xi, \Phi) = J(\nabla_\Phi J)\xi + (\nabla_{J\Phi}J)\xi$ by

$$(1.14) \quad \frac{1}{8}N_J(\xi, f_*w - 2\nu(w) - \Phi_f(w)) + \frac{1}{2}T_\nu(\xi, w) + \frac{1}{4}\widehat{\nabla J}(\xi, \Phi_f(w)).$$

The invariant GW_X was generalized in [IP4] to an invariant of (X, ω) relative to a codimension 2 symplectic submanifold V . To define it, we fix a pair (J, ν) which is V -compatible, meaning that along V the normal components of ν and the tensors N_J and T_ν satisfy

$$(1.15)$$

(a) TV is J -invariant and $\nu^N = 0$, and

(b) $N_J^N(\xi, \zeta) = 0$ and $T_\nu^N(\xi, w) = 0$ for all $\xi \in TX$, $\zeta \in TV$ and $w \in TC$.

(This is the same as Definition 3.2 in [IP4] after one observes that (b) is automatic for $\xi \in TV$ since N_J and $T_\nu(\xi, w) = J[\nu, J\xi] + [\nu, \xi]$ are brackets of vector fields in V and hence lie in V .)

A stable map into X is called V -regular if $f^{-1}(V)$ consists of finitely many points x_1, \dots, x_ℓ none of which are equal to the marked points p_i or the nodes. After the x_j are numbered, the orders of contact of f with V at the x_j define a multiplicity vector $s = (s_1, \dots, s_\ell)$ and three associated integers:

$$(1.16) \quad \ell(s) = \ell, \quad \deg s = \sum s_i, \quad |s| = \prod s_i.$$

By convention, when $\ell = 0$ (which corresponds to $f^{-1}(V) = \emptyset$) we take $\deg s = 0$ and $|s| = 1$.

The space of all V -regular maps is the union of components

$$(1.17) \quad \mathcal{M}_{\chi, n, s}^V(X, A) \subset \overline{\mathcal{M}}_{\chi, n + \ell(s)}(X, A)$$

labeled by vectors s of length $\ell = \ell(s)$ (here χ is the Euler characteristic of the domain). This has a compactification that comes with evaluation maps

$$(1.18) \quad \varepsilon_V : \overline{\mathcal{M}}_{\chi, n, s}^V(X, A) \rightarrow \widetilde{\mathcal{M}}_{\chi, n + \ell(s)} \times X^n \times \mathcal{H}_{X, A, s}^V.$$

Here $\widetilde{\mathcal{M}}_{\chi, n}$ is the space of curves with finitely many components, Euler class χ and n marked points, and $\mathcal{H}_{X, A, s}^V$ is the ‘intersection-homology’ space described in Section 5 of [IP4]. There is a covering map $\varepsilon : \mathcal{H}_{X, A, s}^V \rightarrow H_2(X) \times V_s$ whose first component records the class A and whose component in the space $V_s \cong V^{\ell(s)}$ records the image of the last $\ell(s)$ marked points. This covering is a necessary complication to the definition of relative GW invariants.

The complication occurs because of “rim tori”. A rim torus is an element of

$$(1.19) \quad \mathcal{R} = \ker (\iota_* : H_2(X \setminus V) \rightarrow H_2(X))$$

where ι is the inclusion. Each such element can be represented as $\pi^{-1}(\gamma)$ where π is the projection $S_V \rightarrow V$ from the boundary of a tubular neighborhood of V (the “rim of V ”) and $\gamma : S^1 \rightarrow V$ is a loop in V . The group \mathcal{R} is the group of deck transformations of the covering

$$(1.20) \quad \begin{array}{ccc} \mathcal{R} & \longrightarrow & \mathcal{H}_X^V = \bigsqcup_{A,s} \mathcal{H}_{X,A,s}^V \\ & & \downarrow \varepsilon \\ & & H_2(X) \times \bigsqcup_s V_s. \end{array}$$

When there are no rim tori (as is the case if V is simply connected), $\mathcal{H}_{X,A,s}^V$ reduces to V_s and the evaluation map (1.18) is more easily described.

The tangent space to $\mathcal{M}_{\chi,n,s}^V(X, A)$ is modeled on $\ker D_s$ where D_s is the restriction of (1.10) to the subspace where ξ^N has a zero of order s_i at the marked points x_i , $i = 1, \dots, \ell$. For generic (J, ν) as above, $\text{coker } D_s = 0$ at irreducible maps and therefore the irreducible part $\mathcal{M}_{\chi,n,s}^V(X, A)^*$ of the moduli space is an orbifold with

$$(1.21) \quad \dim \mathcal{M}_{\chi,n,s}^V(X, A) = -2K_X A + \frac{\chi}{2} (\dim X - 6) + 2n - 2(\deg s - \ell(s)).$$

With this understood, the definition of the relative GW invariant parallels the above definition of GW_X . A stable V -regular map is called (X, V) -*admissible* if it is irreducible when restricted to the union of its unstable domain components B_i with

$$(1.22) \quad 0 < K_X[f(B_i)] + V \cdot [f(B_i)].$$

In this context we say that the symplectic pair (X, V) is *semipositive* if there are no spherical classes $A \in H_2(X)$ with $A \cdot V \geq 0$, $\omega(A) > 0$, and $0 < K_X A + A \cdot V \leq \frac{1}{2}(\dim X - 6) + \min\{A \cdot V, 2\}$.

Any map from a connected unstable domain B_i satisfying (1.22) is the composition of a covering map with an irreducible map f_0 with unstable rational domain which represents a class A satisfying $0 < K_X A + A \cdot V$. If f_0 does not lie in V then f_0 , marked only with its ℓ intersection points with V , lies in a moduli space of dimension $-2K_X A + \dim X - 6 + 2\ell - 2A \cdot V$ with $\ell \leq A \cdot V$ and, since the domain is unstable, $\ell \leq 2$. When (X, V) is semipositive that dimension is negative, so such f_0 do not exist for generic (J, ν) . If the irreducible map f_0 lies in V then the dimension of its moduli space is $-2K_V A + \dim V - 6 = -2K_X A - 2V \cdot A + \dim X - 8$ which is also negative when (X, V) semipositive. Thus semipositivity implies that all maps in each moduli space $\mathcal{M}_{\chi,n,s}^V(X, A)$ are admissible for generic (J, ν) .

When the moduli spaces are generically admissible the image moduli space under (1.18) carries a homology class which, after summing on χ, n and s , can be thought of as a map

$$(1.23) \quad \mathrm{GW}_{X,A}^V : \mathbb{T}^*(X) \longrightarrow H_*(\overline{\mathcal{M}} \times \mathcal{H}_X^V; \mathbb{Q}[\lambda]).$$

This gives the expansion

$$(1.24) \quad \mathrm{GW}_X^V = \sum_{A, g} \sum_{\substack{\text{ordered seqs } s \\ \deg s = A \cdot V}} \frac{1}{\ell(s)!} \mathrm{GW}_{X,A,g,s}^V t_A \lambda^{2g-2}$$

whose coefficients are (multi)-linear maps $\mathbb{T}^*(X) \rightarrow H_*(\overline{\mathcal{M}} \times \mathcal{H}_{X,A,s}^V)$ (dividing by $\ell(s)!$ eliminates the redundancy associated with renumbering the last ℓ marked points). The corresponding relative Gromov-Taubes invariant is again given by

$$(1.25) \quad \mathrm{GT}_X^V = e^{\mathrm{GW}_X^V}.$$

After imposing constraints one can expand GT_X^V in power series. That is done in the appendix under the assumption that there are no rim tori.

2. Symplectic sums

Assume X and Y are $2n$ -dimensional symplectic manifolds each containing symplectomorphic copies of a codimension two symplectic submanifold (V, ω_V) . Then the normal bundles are oriented, and we assume they have opposite Euler classes:

$$(2.1) \quad e(N_X V) + e(N_Y V) = 0.$$

We then fix, once and for all, a symplectic bundle isomorphism $\psi : (N_X V)^* \rightarrow N_Y V$.

These data determine a family of symplectic sums $Z_\lambda = X \#_{V,\lambda} Y$ parametrized by λ near 0 in \mathbb{C} ; these have been described in [Gf] and [MW]. In fact, this family fits together to form a smooth $2n + 2$ -dimensional symplectic manifold Z that fibers over a disk. In this section we will construct Z and describe its properties.

THEOREM 2.1. *Given the above data, there exists a $2n + 2$ -dimensional symplectic manifold (Z, ω) and a fibration $\lambda : Z \rightarrow D$ over a disk $D \subset \mathbb{C}$. The center fiber Z_0 is the singular symplectic manifold $X \cup_V Y$, while for $\lambda \neq 0$, the fibers Z_λ are smooth compact symplectic submanifolds — the symplectic connect sums.*

This displays the Z_λ as deformations, in the symplectic category, of the singular space $X \cup_V Y$ as in Figure 1. For $\lambda \neq 0$ these are symplectically isotopic to one another and to the sums defined in [Gf] and [MW].

The proof of Theorem 2.1 involves the following construction. Given a complex line bundle $\pi : L \rightarrow V$ over V , fix a hermitian metric on L , set $\rho(x) = \frac{1}{2}|x|^2$ for $x \in L$, and choose a compatible connection on L . The connection defines a real-valued 1-form α on $L \setminus \{\text{zero section}\}$ with $\alpha(\partial/\partial\theta) = 1$ (identify the principal bundle with the unit circle bundle and pull back the connection form by the radial projection) and the curvature of α defines a 2-form F on V with $\pi^*F = d\alpha$. Then the form

$$(2.2) \quad \omega = \pi^*\omega_V + \rho\pi^*F + d\rho \wedge \alpha$$

extends across the zero section and is S^1 -invariant, closed, and nondegenerate for small ρ . The moment map for this S^1 action $x \mapsto e^{i\theta}x$ is the function $-\rho$ because $i\frac{\partial}{\partial\theta}\omega = i\frac{\partial}{\partial\theta}(d\rho \wedge \alpha) = -d\rho$.

The dual bundle L^* has a dual metric, a radial function $\rho^*(y) = \frac{1}{2}|y|^2$, and connection α^* with $d\alpha^* = -\pi^*F$. This gives a symplectic form similar to (2.2) on L^* and hence one on $\pi : L \oplus L^* \rightarrow V$, namely

$$(2.3) \quad \omega = \pi^*\omega_V + (\rho - \rho^*)\pi^*F + d\rho \wedge \alpha + d\rho^* \wedge \alpha^*.$$

Below, we will denote points in $L \oplus L^*$ by triples (v, x, y) where $v \in V$ and (v, x, y) is a point in the fiber of $L \oplus L^*$ at v . This space has

- (2.4) (a) a circle action $(v, x, y) \mapsto (v, e^{i\theta}x, e^{-i\theta}y)$ with Hamiltonian $t(v, x, y) = \rho^* - \rho$,
 (b) a natural S^1 -invariant map $L \oplus L^* \rightarrow \mathbb{C}$ by $\lambda(z, x, y) = xy \in \mathbb{C}$.

Proof of Theorem 2.1. Fix (ω, J, g) on $L = N_X V$ as above. Using ψ and the Symplectic Neighborhood Theorem, we symplectically identify a neighborhood of V in X with the disk bundle of radius $\varepsilon \leq 1$ in L and a neighborhood of V in Y with the ε -disk bundle in L^* . We assume that $\varepsilon = 1$; the general case then follows by rescaling (ω, J, g) .


Let D denote the disk of radius $\delta < 1/2$ in \mathbb{R}^2 with the symplectic form $\omega_D = r dr d\theta$. The space Z is constructed from three open pieces: two ends $\text{End}_X = (X \setminus V) \times D$, and $\text{End}_Y = (Y \setminus V) \times D$ and a “neck” modeled on the open set

$$(2.5) \quad U = \{ (v, x, y) \in L \oplus L^* \mid |x| \leq 1, |y| \leq 1 \}.$$

These are glued together by the diffeomorphisms

$$(2.6) \quad \begin{aligned} \psi_X : U \setminus L^* &\rightarrow \text{End}_X & \text{by } (v, x, y) &\mapsto ((v, x), \lambda(v, x, y)), \\ \psi_Y : U \setminus L &\rightarrow \text{End}_Y & \text{by } (v, x, y) &\mapsto ((v, y), \lambda(v, x, y)). \end{aligned}$$

This defines Z as a smooth manifold. The function λ extends over the ends as the coordinate on the D factor, giving a projection $\lambda : Z \rightarrow D$ whose fibers are smooth submanifolds Z_λ for small $\lambda \neq 0$.

Figure 3. Construction of Z_λ

In the overlap region of U where $(1 - \delta) \leq |x| \leq 1$ and $|y| \leq \delta$ the form (2.3) is $\omega_X + d(\rho^* \alpha^*)$ (this ω_X is the pull-back of the symplectic form on X to L^*), and $\psi_X^*(\omega_X \oplus \omega_D) = \omega_X + \lambda^* \omega_D$ because of the symplectic neighborhood identification. But $2\lambda^*(rdrd\theta) = d(|\lambda|^2 \lambda^* d\theta) = 4d(\rho\rho^* \lambda^* d\theta)$ and, since $d\theta$ is the connection on the trivial bundle $L \otimes L^*$, $\lambda^* d\theta$ is the connection form on $L \oplus L^*$, namely $\alpha + \alpha^*$. Thus

$$\lambda^* \omega_D = d(\rho^* \eta) \quad \text{where} \quad \eta = 2\rho (\alpha + \alpha^*).$$

We can then smoothly merge $\lambda^* \omega_D$ into $d(\rho^* \alpha^*)$ by replacing η by $\hat{\eta} = \beta\eta + (1 - \beta)\alpha^*$ where $\beta = \beta(|x|)$ is an appropriate cutoff function with $|d\beta| < 2/\delta$. In this overlap region 2ρ lies between $(1 - \delta)^2$ and 1, $d\alpha$ and $d\alpha^*$ are bounded, and $|\alpha| \leq 2$ and $|\alpha^*| \leq 2/\delta$. It follows that $\alpha^* - \eta$ satisfies $|\alpha^* - \eta| \leq C$ and $|d(\alpha^* - \eta)| \leq C/\delta$, and hence

$$|d(\rho^* \hat{\eta}) - d(\rho^* \eta)| = |d[\rho^*(1 - \beta)(\alpha^* - \eta)]| \leq C\delta.$$

Because $\omega_X + d(\rho^* \eta)$ is nondegenerate and δ is as small as desired, the above inequality shows that on this overlap region $\omega_X + d(\rho^* \hat{\eta})$ is closed and nondegenerate for small δ . Thus we have a specific formula extending (2.3) over End_X as a symplectic form whose restriction to the part of $Z_\lambda \subset \text{End}_X$ with $|x| \geq 1$ is the original symplectic form ω_X on $X \times \{\lambda\}$. Repeating the construction on the Y side yields a global symplectic form ω on Z whose restriction on the Y side is ω_Y . Finally, along $Z_\lambda \cap U$ we have $\alpha^* = -\alpha$, so that ω restricts to

$$\omega_\lambda = \pi^* \omega_V - t \pi^* F - dt \wedge \alpha$$

with t as in (2.4). Thus after possibly making δ smaller, we have a fibration $\lambda : Z \rightarrow D$ with symplectic fibers. \square

This construction shows that the neck region U of Z has a symplectic S^1 action with Hamiltonian $t = \frac{1}{2}(|y|^2 - |x|^2)$. This action preserves λ , and so restricts to a Hamiltonian action on each Z_λ . In fact, t gives a parameter along the neck, splitting each Z_λ into manifolds with boundary

$$Z_\lambda = Z_\lambda^- \cup Z_\lambda^+$$

where Z_λ^- is $Z_\lambda \cap \text{End}_X$ together with the part of $Z_\lambda \cap U$ with $t \leq 0$. From this decomposition we can recover the symplectic manifolds X and Y up to isotopy in two ways:

- (1) as $\lambda \rightarrow 0$, the interior of Z_λ^- (respectively Z_λ^+) converges to $X \setminus V$ (respectively $Y \setminus V$) as symplectic submanifolds of Z , or
- (2) X (resp. Y) is the symplectic cut of Z_λ^- (resp. Z_λ^+) at $t = 0$.

Thus we have collapsing maps

$$(2.7) \quad \begin{array}{ccc} X \sqcup Y & & Z_\lambda \\ \pi_0 \searrow & & \swarrow \pi_\lambda \\ & Z_0 & \end{array}$$

and π_λ is a deformation equivalence on the set where $t \neq 0$.

The next step is to define a Riemannian metric and an almost complex structure on (Z, ω) so that the triple (ω, J, g) is compatible in the sense of (1.1). We begin by specifying such a triple on the total space of $L \rightarrow V$. For each small $\rho > 0$ we can extend the symplectic form $\omega_\rho = \omega_V + \rho F$ to a compatible triple $(\omega_\rho, J_\rho, g_\rho)$ on V with $g_\rho = g_V + O(\rho)$. At each $p = (v, x) \in L$ with $\rho = \frac{1}{2}|x|^2 \neq 0$, there is a splitting $T_p L = H \oplus L_p$ into a horizontal subspace $H = \ker d\rho \cap \ker \alpha$ and a vertical subspace $\ker \pi_*$ identified with L_p , and by (2.2) this splitting is ω -orthogonal. Using this splitting, define a metric on $T_p L$ by $g_0 = \pi^* g_\rho \oplus g_L$. Starting from this g_0 and ω , the polarization procedure described in the appendix of [IP4] produces a pair (J, g) compatible with ω which respects the splitting and which extends across the zero section. In fact, around each point of V there is a local trivialization of L in which $g = g_V \oplus g_{\mathbb{C}} + O(|x|^2)$.

Applying the same construction on L^* with the dual connection and taking the direct sum, we have a compatible structure (ω, J, g) on a neighborhood of the zero section $V \subset L \oplus L^*$. That structure locally agrees with the product structure on $V \times \mathbb{C} \times \mathbb{C}$ to second order along V , so is V -compatible as in (1.15), the second fundamental form h of V is zero, and the Nijenhuis tensor satisfies $\langle N_J(\xi, v), \eta \rangle = 0$ for all $v \in TV$ and ξ, η normal to V . In particular, the pair (J, ν) with $\nu = 0$ satisfies condition (i) of the following definition.

Definition 2.2. Let $\mathcal{J}(Z)$ be the set of extensions of the symplectic structure ω on Z to a compatible data (ω, J, g, ν) with $\|\nu\|_\infty < 1/4$ as in (1.4) and so that (i) along V we have (1.15ab) and the second fundamental form of $V \subset Z$ vanishes, and (ii) Z_λ is J -invariant for all λ .

Now consider compatible structures $(\omega_X, J_X, g_X, \nu_X)$ on X and $(\omega_Y, J_Y, g_Y, \nu_Y)$ on Y which satisfy condition (i) above. If these agree along V , they define a compatible structure on Z_0 and the following lemma constructs corresponding elements of $\mathcal{J}(Z)$.

LEMMA 2.3. *Given $\varepsilon > 0$ and structures on X and Y as above, there exists an element $(\omega, J, g, \nu) \in \mathcal{J}(Z)$ which agrees with the given data on $Z_0 \setminus U(\varepsilon)$ and along V . In particular, $\mathcal{J}(Z)$ is nonempty.*

Proof. We have just seen that for small ε there are data $(\omega, J, g, 0)$ on $U(\varepsilon)$ satisfying (i) of Definition 2.2. These data also satisfy (ii) at points $p = (v, x, y) \in U(\varepsilon)$, as follows. Each path v_t in V starting at v has a horizontal lift (v_t, x_t, y_t) in $L \oplus L^*$ with initial point p . Along the lift, $\lambda = x \cdot y$ and $\lambda' = \nabla_{\dot{v}} x \cdot y + x \cdot \nabla_{\dot{v}} y = 0$, so λ is constant. Thus each horizontal lift lies in Z_λ for $\lambda = \lambda(p)$. It follows that $T_p Z_\lambda$ is the sum of the horizontal subspace at p and the tangent space to the complex curve $xy = \lambda$ in the fiber $(L \oplus L^*)_v = \mathbb{C}^2$. Both of those subspaces are J -invariant, so (ii) holds in $U(\varepsilon)$.

We can extend g on $U(\varepsilon/2)$ to a Riemannian metric g_0 on Z which agrees with the given product metrics $Z_0 \setminus U(\varepsilon)$. By Lemma A.1 of [IP4], (ii) is equivalent to the condition that the symplectic normal N_p to $T_p Z_\lambda$ is equal to the metric normal. This already holds on $U(\varepsilon)$ and on $Z_0 \setminus V$, and we can choose g_0 so that it holds everywhere. As above, applying the polarization procedure to g_0 and ω yields a pair (J, g) compatible with ω which still respects the decomposition $TZ_\lambda \oplus N$ and which has $g = g_0$ in the regions where g was already compatible with ω . Finally, we can extend ν on $Z_0 \setminus B(\varepsilon)$ arbitrarily. \square

We will work with structures in $\mathcal{J}(Z)$ throughout the analytical sections of this paper.

We conclude this section with a useful lemma comparing the canonical class of the symplectic sum with the canonical classes K_X and K_Y of X and Y .

LEMMA 2.4. *If $A \in H_2(Z_\lambda; \mathbb{Z})$, $\lambda \neq 0$, is homologous in Z to the union $C_1 \cup C_2 \subset X \cup_V Y$ of cycles C_1 in X and C_2 in Y , then*

$$K_{Z_\lambda}[A] = K_Z[A] = K_X[C_1] + K_Y[C_2] + 2\beta$$

where β is the intersection number $V \cdot [C_1] = V \cdot [C_2]$. In particular, $K_{Z_\lambda}[R] = 0$ for any rim torus R (cf. (1.19)).

Proof. For $\lambda \neq 0$, the normal bundle to Z_λ has a nowhere-vanishing section $\partial/\partial\lambda$. Thus the canonical bundle of Z_λ is the restriction of the canonical bundle of Z , giving

$$K_{Z_\lambda}[A] = K_Z[A] = K_Z[C_1] + K_Z[C_2].$$

Along $X \subset Z_0$ the tangent bundle to Z decomposes as

$$TZ = TX \oplus \pi^* \psi^* N_Y V \cong TX \oplus \pi^* (N_X V)^{-1}$$

where π is the projection $N_X V \rightarrow V$. But the Poincaré dual of V in X , regarded as an element of $H^2(X)$, is the Chern class $c_1(\pi^* N_X V)$. Since the canonical class is minus the first Chern class of the tangent bundle we conclude that

$$K_Z[C_1] = K_X[C_1] + V \cdot [C_1]$$

and similarly on the Y side. \square

3. Degenerations of symplectic sums

The Gromov-Witten invariants of the symplectic sum Z_λ are defined in terms of stable pseudo-holomorphic maps from complex curves into the Z_λ . The basic idea of our symplectic sum formula is to approximate the maps in Z_λ by certain maps into the singular space Z_0 . The first step is a limiting argument. The key point is that, by the construction in Section 2, the spaces $\{Z_\lambda\}$ are embedded in a compact almost Kähler manifold Z — the closure of a neighborhood of the central fiber of the family Z . Hence the ‘‘Gromov Compactness Theorem’’ implies that sequences of (J, ν) -holomorphic maps into Z_λ limit to maps into Z_0 (after passing to subsequences). This section gives a description of the maps into Z_0 which arise as limits of δ -flat stable maps into the Z_λ as $\lambda \rightarrow 0$. The ‘ δ -flat’ condition, defined below, ensures that the limit has no components mapped into V .

Fix a small $\delta > 0$. Given a map f into Z_λ , we can restrict attention to that part of the image that lies in the ‘ δ -neck’

$$(3.1) \quad Z_\lambda(\delta) = \{z = (v, x, y) \in Z_\lambda \mid ||x|^2 - |y|^2| \leq \delta\}.$$

This is a narrow region symmetric about the middle of the neck in Figure 3. The energy of f (more precisely of (f, ϕ)) in this region is

$$(3.2) \quad E^\delta(f) = \frac{1}{2} \int |df|^2 + |d\phi|^2$$

where the integral is over $f^{-1}(Z_\lambda(\delta))$.

By Lemma 1.5 of [IP4] there is a constant $\alpha_V < 1$, depending only on (J_V, ν_V) such that every component of every stable (J_V, ν_V) -holomorphic map into V has energy

$$(3.3) \quad E(f) \geq \alpha_V.$$

Definition 3.1 (δ -flat map). A stable (J, ν) -holomorphic map f into Z is δ -flat if the energy in the δ -neck is at most half α_V , that is

$$(3.4) \quad E^\delta(f) \leq \alpha_V/2.$$

Note that a stable δ -flat map into Z_0 cannot have any components (not even ghosts) lying in V and is thus V -regular in the sense of [IP4]. For each small λ , let

$$\mathcal{M}_{\chi,n}^{V,\delta}(Z_\lambda, A)$$

denote the set of δ -flat maps in $\mathcal{M}_{\chi,n}(Z_\lambda, A)$. These are a family of subsets of the space of stable maps into Z and we write

$$(3.5) \quad \lim_{\lambda \rightarrow 0} \mathcal{M}_{\chi,n}^{V,\delta}(Z_\lambda, A)$$

for the set of limits of sequences of δ -flat maps into Z_λ as $\lambda \rightarrow 0$. Because (3.4) is a closed condition this limit set is a *closed* subspace of $\overline{\mathcal{M}}_{\chi,n}(Z_0, A)$. The remainder of this section is devoted to a precise description of the space (3.5).

LEMMA 3.2. *Each element of (3.5) is a stable map f to $Z_0 = X \cup_V Y$ with no components of the domain mapped entirely into V .*

Proof. By the compactness theorem for (J, ν) -holomorphic maps (cf. §1 of [IP4]) each sequence in (3.5) has a subsequence f_k converging in the space of stable maps $\overline{\mathcal{M}}_{\chi,n}(Z, A)$ to a limit $f : C \rightarrow Z$. In particular, the images converge pointwise, and thus lie in Z_0 .

Suppose that the image of some component C_i of C lies in V . Then the restriction f_i of f to that component satisfies $E(f_i) \leq E^\delta(f)$. Furthermore, by Theorem 1.6 of [IP4] the sequence f_k (after precomposing with diffeomorphisms) converges in C^0 and in energy, so $E^\delta(f) = \lim E^\delta(f_k) \leq \alpha_V/2$. This contradicts (3.3). \square

We can be very specific about how the images of the maps in (3.5) hit V . By Lemma 3.2 and Lemma 3.4 of [IP4], at each point $p \in f^{-1}(V)$ the normal component of f has a local expansion $a_0 z^d + \dots$. This defines a local ‘degree of contact’ with V

$$(3.6) \quad d = \deg(f, p) \geq 1$$

and implies that $f^{-1}(V)$ is a finite set of points. By restricting f to one component C_i of C and removing the points $f^{-1}(V)$, one obtains a map from a connected domain to the disjoint union of $X \setminus V$ and $Y \setminus V$. Thus the components of C are of two types: those components C_i^X whose image lies in X , and those components C_i^Y whose image lies in Y . We can therefore split f into two parts: the union of the components whose image lies in X defines a map $f_1 : C_1 \rightarrow X$, from a (possibly disconnected) curve C_1 , and the remaining components define a similar map $f_2 : C_2 \rightarrow Y$.

LEMMA 3.3. $f^{-1}(V)$ consists of nodes of C . For each node $x = y \in f^{-1}(V)$

$$\deg(f_1, x) = \deg(f_2, y).$$

Proof. The local degree (3.6) is a linking number. Specifically, let $N_X V$ be a tubular neighborhood of V in X and let μ_X be the element of $H_1(N_X V \setminus V)$ represented by the oriented boundary of a holomorphic disk normal to V . If μ_Y is the corresponding element on the Y side, then $\mu_X = -\mu_Y$ in H_1 of the neck $Z_\lambda(\delta)$. For each point x in $f_1^{-1}(V)$ and each small circle S_ε around x , the local degree d satisfies

$$d \cdot \mu = [f_1(S_\varepsilon)].$$

If x is not a node of C then by Theorem 1.6 of [IP4] f_k converges to f_1 in C^1 in a disk D around x . But then for large k , $d \cdot \mu = [f(S_\varepsilon)] = [f_k(S_\varepsilon)] = [f_k(\partial D)] = 0$, contradicting (3.6).

Next consider a node $x = y$ of C which is mapped into V . Choose holomorphic disks $D_1 = D(x, \varepsilon)$ and $D_2 = D(y, \varepsilon)$ that contain no other points of $f^{-1}(V)$ and let $S_i = \partial D_i$. Then $S_1 \cup S_2$ bounds in C , so that $[f_k(S_1)] + [f_k(S_2)] = 0$ in H_1 of the neck $Z_\lambda(\delta)$. Again, $f_k \rightarrow f$ in C^0 , which implies that $0 = [f(S_1)] + [f(S_2)] = d_1 \mu_1 + d_2 \mu_2$ where μ_i is either μ_X or μ_Y , depending on which side $f(S_i)$ lies. Since $d_i > 0$ the only possibility is that $x = y$ is a node between a component in X and one in Y and $d_1 = d_2$. \square

Lemmas 3.2 and 3.3 show that each map f in the limiting set (3.5) splits into (J, ν) -holomorphic maps $f_1 : C_1 \rightarrow X$ and $f_2 : C_2 \rightarrow Y$. Ordering the nodes in $f^{-1}(V)$ gives extra marked points x_1, \dots, x_ℓ on C_1 and matched y_1, \dots, y_ℓ on C_2 with $s_i = \deg x_i = \deg y_i$. Furthermore, the Euler characteristics χ_1 of C_1 and χ_2 of C_2 satisfy

$$(3.7) \quad \chi_1 + \chi_2 - 2\ell = \chi.$$




Figure 4. The map $f_0 = (f_1, f_2)$ into $Z_0 = X \cup_Y Y$

We can now give a global description of how the limit maps f in (3.5) are assembled from their components f_1 and f_2 . First, consider how the domain curves fit together in accordance with (3.7). Given bubble domains C_1 and C_2 , not necessarily connected or stable, with Euler characteristics χ_i and $n_i + \ell$ marked points, we can construct a new curve by identifying the last ℓ marked points of C_1 with the last ℓ marked points of C_2 , and then forgetting the marking of these new nodes. After possibly adding more marked points to stabilize any unstable components, this process is the same as the standard attaching map

$$(3.8) \quad \xi_\ell : \widetilde{\mathcal{M}}_{\chi_1, n_1 + \ell} \times \widetilde{\mathcal{M}}_{\chi_2, n_2 + \ell} \longrightarrow \widetilde{\mathcal{M}}_{\chi_1 + \chi_2 - 2\ell, n_1 + n_2}$$

whose image is a subvariety of complex codimension ℓ . Taking the union over all χ_1, χ_2, n_1 and n_2 gives a (stabilized) attaching map $\xi_\ell : \widetilde{\mathcal{M}} \times \widetilde{\mathcal{M}} \rightarrow \widetilde{\mathcal{M}}$ for each ℓ .

Second, consider how the maps fit together along V . The evaluation map

$$\text{ev}_s s : \mathcal{M}_{\chi_1, n_1, s}^V(X) \times \mathcal{M}_{\chi_2, n_2, s}^V(Y) \xrightarrow{\varepsilon_V \times \varepsilon_V} \mathcal{H}_X^V \times \mathcal{H}_Y^V \xrightarrow{\varepsilon_2 \times \varepsilon_2} V_s \times V_s.$$

records the intersection points with V and the pair (f_1, f_2) lies in the space

$$(3.9) \quad \mathcal{M}^V(X) \times_{\text{ev}_s} \mathcal{M}^V(Y) \stackrel{\text{def}}{=} \text{ev}_s s^{-1}(\Delta_s).$$

where Δ_s is the diagonal

$$\Delta_s \subset V_s \times V_s.$$

Denote by $\mathcal{H}_X^V \times_\varepsilon \mathcal{H}_Y^V = (\varepsilon_2 \times \varepsilon_2)^{-1}(\Delta)$ the fiber sum of \mathcal{H}_X^V and \mathcal{H}_Y^V along the evaluation map ε_2 , where $\Delta = \bigsqcup_s \Delta_s$. Then we have a well defined map

$$(3.10) \quad g : \mathcal{H}_X^V \times_{\varepsilon} \mathcal{H}_Y^V \rightarrow H_2(Z)$$

which describes how the homology-intersection data of f_1 and f_2 determine the homology class of f .

LEMMA 3.4. *For generic (J, ν) the irreducible part of the space (3.9) is a smooth orbifold of the same dimension as $\mathcal{M}_{\chi, n}^{V, \delta}(Z_\lambda, A)$, given by (1.6).*

Proof. The dimensions of $\mathcal{M}_{\chi_1, s}^V(X, A_1)^*$ and $\mathcal{M}_{\chi_2, s}^V(Y, A_2)^*$ are given by (1.21). A small modification of the proof of Lemma 8.5 of [IP4] shows that the evaluation map at the last $\ell = \ell(s)$ marked points (i.e. the intersection points with V) is transversal to the diagonal $\Delta \subset V^\ell \times V^\ell$, imposing $\ell \dim V = \ell(\dim X - 2)$ conditions. Thus the irreducible part of (3.9) is a smooth orbifold of dimension

$$-2K_X[A_1] - 2K_Y[A_2] - 4 \deg s - \frac{1}{2}(\dim X - 6)(\chi_1 + \chi_2 - 2\ell) + 2n.$$

The lemma follows by comparing this with (1.6) using (3.7), Lemma 2.4, and the fact that $\deg s = A_1 \cdot V = A_2 \cdot V$. \square

Finally, note that renumbering the pairs (x_i, y_i) of marked points defines an action of the symmetric group S_ℓ on (3.9) and the limit maps in (3.5) correspond to elements in the quotient. Moreover, after ordering the double points along V the limit set (3.5) is a subset of the set

$$(3.11) \quad \mathcal{K}_\delta \subset \bigsqcup_s \mathcal{M}^V(X) \times_{\text{ev}_s} \mathcal{M}^V(Y)$$

consisting of δ -flat maps into Z_0 with fixed data (χ, n, A) .

Remark 3.5. Observe that the set \mathcal{K}_δ is compact and that the multiplicity vector $s = (s_1, \dots, s_\ell)$ is constant on components of \mathcal{K}_δ . Indeed, as in the proof of Lemma 3.2, any sequence $\{f_n\} \in \mathcal{K}_\delta$ has a subsequence converging to a V -regular map f_0 into Z_0 . Because the energy (3.2) is continuous in the stable map topology the limit is δ -flat, and thus lies in \mathcal{K}_δ . Furthermore, after lifting to the normalization of the domain, stable map convergence implies that the marked points $x_{k,n}$ of $f_n^{-1}(V)$ converge to points $x_k \in f_0^{-1}(V)$ distinct from the nodes. This in turn means that f_n converges to f_0 in C^∞ in a neighborhood of each x_k . Consequently the order of contact s'_k of f_0 with V at x_k is at least s_k . Since the total intersection $\sum s_k = A \cdot V$ is preserved in the limit and all multiplicities are positive, we conclude that $s = s'$. In particular, no new intersection points with V arise in the limit.

Since the δ -flat maps in $\mathcal{M}^{V,\delta}(Z_\lambda)$ are C^0 close to δ -flat maps into Z_0 for small λ there is a decomposition

$$\mathcal{M}^{V,\delta}(Z_\lambda) = \bigsqcup_s \left(\mathcal{M}_s^{V,\delta}(Z_\lambda) \right) / S_{\ell(s)}$$

as a union of components labeled by ordered sequences $s = (s_1, s_2, \dots)$. As in the proof of Lemma 3.3, these s_i are local winding numbers of the $\ell(s)$ vanishing cycles S_ε . In that form the labeling extends to all continuous maps C^0 close to δ -flat maps into Z_0 . Thus for small λ

$$\mathcal{M}_s^{V,\delta}(Z_\lambda) \subset \text{Map}_s(Z_\lambda)$$

where $\text{Map}_s(Z_\lambda)$, the “space of labeled maps”, is the set of labeled continuous maps into Z_λ which are C^0 close to δ -flat maps into Z_0 .

Thus with this notation, the statements of Lemmas 3.2 and 3.3 translate into the commutative diagram

$$(3.12) \quad \begin{array}{ccc} \bigsqcup_s \mathcal{M}^V(X) \times_{\text{ev}_s} \mathcal{M}^V(Y) & \longleftarrow & \lim_{\lambda} \left(\bigsqcup_s \mathcal{M}_s^{V,\delta}(Z_\lambda) \right) \\ \downarrow & & \downarrow \\ (\widetilde{\mathcal{M}} \times \widetilde{\mathcal{M}}) \times \left(\mathcal{H}_X \times_{\varepsilon} \mathcal{H}_Y \right) & \xrightarrow{\xi \times g} & \widetilde{\mathcal{M}} \times H_2(Z). \end{array}$$

The top arrow shows how the maps that arise as limits of δ -flat maps decompose into pairs (f_1, f_2) of V -regular maps into X and Y , while the bottom arrow keeps track of the domains and homology classes (the vertical maps arise from (1.18) in the obvious way).

One then expects the top arrow in (3.12) to be a diffeomorphism for each s and each side to be a model for the stable maps into Z_λ for that s . The analysis of the next six sections will show that this is true after passing to a finite cover.

The necessity of passing to covers is dictated by the clustering phenomenon mentioned in the introduction: when $s > 1$ each curve in Z_0 is close (in the stable map topology) to a cluster of curves in Z_λ for small λ , and these coalesce as $\lambda \rightarrow 0$. To distinguish the curves within a cluster and, indeed even to verify this statement about clustering, it is necessary to use stronger norms and distances — strong enough that the distances between the maps within a cluster are bounded away from zero as $\lambda \rightarrow 0$. The maps in a cluster can then be thought of as the fiber of a covering of the space of limit maps. The next three sections introduce the required norms and construct a first version of the covering. The first step is to define an appropriate distance function on the space of stable curves.

4. The space of curves

One can measure the distance between stable curves using a metric on the moduli space $\overline{\mathcal{M}}_{g,n}$ of curves. However, it is often more convenient to fix a diffeomorphism of the curves, regard the two curves as two complex structures on a single 2-manifold, and measure the distance between the complex structures using a Sobolev norm. In this section we take that approach to define a metric and distance function on $\overline{\mathcal{M}}_{g,n}$. Our metric is designed so that a neighborhood of the image of the attaching map (3.8) is obtained by gluing cylindrical ends of the spaces $\mathcal{M}_{g,n}$. It is a complete metric on $\overline{\mathcal{M}}_{g,n} \setminus \mathcal{N}$ where \mathcal{N} is the set of all nodal curves; in particular it is stronger than the Weil-Petersson metric.

To simplify the exposition we will assume in Sections 4–9 of this paper that stable domains either have no nontrivial automorphisms or come with Prym structures as defined in [Loo]. Prym structures define finite covers of the Deligne-Mumford spaces $\overline{\mathcal{M}}_{g,n}$ which are *smooth* projective varieties. In particular, the corresponding universal curves are smooth and projective, so can be used to define ν as in (1.1).

The construction starts by fixing a Riemannian metric $g_{\mathcal{U}}$ on the universal curve $\overline{\mathcal{U}}_{g,n} \xrightarrow{\pi} \overline{\mathcal{M}}_{g,n}$ compatible with the complex structure. In $\overline{\mathcal{U}}_{g,n}$ the ‘special points’ (marked points and nodes) are distinct and hence, by compactness, are separated by a minimum distance. After rescaling the metric we can assume that the separation distance is at least 4. We also fix a smooth positive function

$\hat{\rho}$ on $\overline{\mathcal{U}}_{g,n}$ equal to the distance to the node in a neighborhood of each node. Finally, we replace $g_{\mathcal{U}}$ by a conformal metric that is singular along the nodal points locus, namely

$$(4.1) \quad g = \hat{\rho}^{-2} g_{\mathcal{U}}.$$

To understand the geometry of this metric we focus attention on a small ball U in the set \mathcal{N}_{ℓ} of ℓ nodal curves and construct a local model for a neighborhood of $\pi^{-1}(U)$ in $\overline{\mathcal{U}}_{g,n}$. The curves C_0 in U have normalizations \tilde{C}_0 with n marked points plus ℓ pairs of marked points x_k, y_k with C_0 obtained by identifying each x_k with y_k . For each k we fix local coordinates $\{z_k\}$ around x_k and $\{w_k\}$ around y_k on the $\tilde{C}_0 \in U$ in which the metric (4.1) is Euclidean. We can use the construction of Section 2 to form a family of symplectic sums of curves; for details see [Ma]. The result is a holomorphic fibration \mathcal{F} with maps

$$(4.2) \quad \begin{array}{ccc} \mathcal{F} & \longrightarrow & \overline{\mathcal{U}}_{g,n} \\ \downarrow & & \downarrow \\ U \times D^{\ell} & \longrightarrow & \overline{\mathcal{M}}_{g,n} \end{array}$$

where D^{ℓ} is the unit disk in \mathbb{C}^{ℓ} . The fiber of \mathcal{F} over (C_0, μ) is a curve $C_0(\mu)$ given by $z_k w_k = \mu_k$ in disjoint balls B_k centered on the nodes. Outside the union of the B_k we can fix a trivialization of \mathcal{F} which respects the marked points. The horizontal arrows in (4.2) are biholomorphic away from the curves with nontrivial automorphisms and biholomorphic everywhere for curves with Prym structures.

Remark 4.1. The parameters $\{\mu_k\}$ are intrinsically elements of the bundle

$$(4.3) \quad \bigoplus_{k=1}^{\ell} (\mathcal{L}_k \otimes \mathcal{L}'_k)^*$$

where \mathcal{L}_k and \mathcal{L}'_k are the relative cotangent bundles to \tilde{C}_0 at x_k and y_k respectively. Thus the fibration (4.3) models a tubular neighborhood of $U \subset \mathcal{N}_{\ell}$ in $\overline{\mathcal{M}}_{g,n}$.

Fix a metric on \mathcal{F} whose restriction to each fiber is Euclidean in the coordinates (z_k, w_k) on each B_k , scaling the metric so that each B_k has radius at least 4. Inside B_k , the induced metric on $C_0(\mu)$ is

$$(4.4) \quad g_{\mu} = \operatorname{Re} (dz^2 + dw^2) \Big|_{zw=\mu} = \left(1 + \frac{|\mu|^2}{r^4}\right) (dr^2 + r^2 d\theta^2)$$

where $r = |z|$ and the distance to the node in B_k is $\rho^2 = |z|^2 + |w|^2 = r^2 + |\mu|^2/r^2$. Switching to the conformal metric $g = \rho^{-2} g_{\mu}$ as in (4.1), and assuming that $\mu_k \neq 0$, we can identify $C_0(\mu) \cap B_k(2)$ with $[-T_k, T_k] \times S^1$

by writing $r = \sqrt{|\mu_k|} e^t$ with $T_k = |\log(2/\sqrt{|\mu_k|})|$. In these coordinates $\rho^2 = 2|\mu_k| \cosh(2t)$ and

$$(4.5) \quad g = \rho^{-2} g_\mu = (r^{-1} dr)^2 + d\theta^2 = dt^2 + d\theta^2.$$

Thus with this metric on \mathcal{F} the curves $C_0(\mu)$ have necks which are isometric to cylinders of radius one and length $2T_k$ with $T_k \rightarrow \infty$ as $\mu_k \rightarrow 0$.

Because the top map in (4.2) is holomorphic with bounded differential, its restriction to each fiber is conformal and the conformal factor is bounded. Consequently, the PDE results of the next several sections, all of which involve only local considerations in the space $\overline{\mathcal{U}}_{g,n}$, can be done in the model space \mathcal{F} using the metric (4.5) and the results will apply uniformly on $\overline{\mathcal{U}}_{g,n}$. We will henceforth consistently use this metric (4.5) on the domains of holomorphic curves and will no longer distinguish between (4.1) and (4.5). Note that the flatness condition (3.4) continues to hold (after a uniform change of constants) because the energy density is conformally invariant.

We next define a metric on $U \times D^\ell$ in terms of a Sobolev metric on the fibers of \mathcal{F} . In the directions tangent to U this will be the Weil-Petersson metric $\|\cdot\|_{WP}$. To describe the metric in the D^ℓ directions, we fix an ℓ -nodal curve C_0 and consider the restriction \mathcal{F}_0 of \mathcal{F} over $\{C_0\} \times D^\ell$. The first step is to construct a diffeomorphism between smooth fibers of \mathcal{F}_0 . Recall that inside $B_k(2)$ the fibers of \mathcal{F}_0 are given by $zw = \mu_k$, while outside the union of the $B_k(1)$ \mathcal{F}_0 has a trivialization which identifies $(z, \mu/z)$ with $(z, \mu'/z)$ when $1 \leq |z| \leq 2$ and $(\mu/w, w)$ with $(\mu'/w, w)$ when $1 \leq |w| \leq 2$; see Figure 3 above and [Ma, p. 626].

For each smooth fiber $C_\mu = C_0(\mu)$ we can parametrize the necks $C_\mu \cap B_k(2)$ by writing $z = \sqrt{|\mu_k|} e^{t+i\theta}$ as above with $(t, \theta) \in [-T_k, T_k] \times S^1$ for $T_k = |\log(2/\sqrt{|\mu_k|})|$. Given a second smooth fiber $C_{\mu'}$ we can similarly parametrize $C_{\mu'} \cap B_k(2)$ by $[-T'_k, T'_k] \times S^1$ and define a map $\psi_{\mu\mu'}^k : C_\mu \cap B_k(2) \rightarrow C_{\mu'} \cap B_k(2)$ by

$$\psi_{\mu\mu'}^k(t, \theta, \mu) = \left(t + \left(\eta(t) - \frac{1}{2} \right) \log \left| \frac{\mu'}{\mu} \right|, \theta - \eta(t) \arg \left(\frac{\mu'}{\mu} \right), \mu' \right)$$

where $\eta(t)$ is a cutoff function equal to 1 for $t \leq -1$ and 0 for $t \geq 1$. This diffeomorphism between the necks, one readily checks, agrees with the given trivialization on $B_k(2) \setminus B_k(1)$ and so defines a diffeomorphism $C_\mu \rightarrow C_{\mu'}$.

The corresponding infinitesimal diffeomorphism defines lifts of vectors $v = (v_1, \dots, v_\ell) \in T_\mu D^\ell$ to \mathcal{F}_0 : v defines a path $\mu(s) = \mu + sv$ and a vector field

$$\tilde{v} = \frac{d}{ds} \psi_{\mu\mu(s)} \Big|_{s=0} = \sum_k \left(\left(\eta - \frac{1}{2} \right) \operatorname{Re} \left(\frac{v_k}{\mu_k} \right), -\eta \cdot \operatorname{Im} \left(\frac{v_k}{\mu_k} \right), v \right)$$

along C_μ . Going the other way, given any path $\mu(s)$ with $\mu_k(s)$ nonzero for all k and s , we can lift the vectors $\dot{\mu}$ as above and integrate the lifted vector

fields to get diffeomorphisms $\psi_s : C_{\mu(0)} \rightarrow C_{\mu(s)}$. The variation in the complex structure is $h_1 = \frac{d}{ds} \psi_s^* j|_{s=0}$. Define the metric by

$$(4.6) \quad \|h_1\|^2 = \int_{C_\mu} |\nabla^2 h_1|^2 + |\nabla h_1|^2 + |h_1|^2 dv_g$$

using the norms and volume form of the cylindrical metric (4.5). With this norm the Sobolev inequality $\sup |h_1| \leq c \|h_1\|$ holds with a constant uniform in μ (the cylindrical metric on C_μ has an upper bound for curvature and a lower bound on the injectivity radius independent of μ).

LEMMA 4.2. *On the complement of the nodal set $\mathcal{N} = \{\mu \mid \text{some } \mu_k \text{ is zero}\}$ the Riemannian metric (4.6) on $\{C_0\} \times D^\ell$ is uniformly equivalent to the metric*

$$(4.7) \quad \sum_k \frac{|d\mu_k|^2}{|\mu_k|^2}.$$

Proof. Calculating $h_1 = \mathcal{L}_{\tilde{v}} j = \mathcal{L}_{\tilde{v}}(\partial_\theta \otimes dt - \partial_t \otimes d\theta)$ by computing the Lie derivatives $\mathcal{L}_{\tilde{v}} \partial_\theta = [\tilde{v}, \partial_\theta] = 0$, $\mathcal{L}_{\tilde{v}} dt = d\tilde{v}_t$, $\mathcal{L}_{\tilde{v}} \partial_t = [\tilde{v}, \partial_t]$, and $\mathcal{L}_{\tilde{v}} d\theta = d\tilde{v}_\theta$, one finds that

$$h_1 = \eta'(B\partial_t + A\partial_\theta) \otimes dt - \eta'(A\partial_t - B\partial_\theta) \otimes d\theta \quad \text{where} \quad \begin{cases} A = \operatorname{Re} \left(\frac{v_k}{\mu_k} \right) \\ B = \operatorname{Im} \left(\frac{v_k}{\mu_k} \right). \end{cases}$$

Because $d\eta$ has support in $[-1, 1]$ and the integrals of $|d\eta|$, $|\nabla d\eta|$, and $|\nabla^2 d\eta|$ are independent of μ we then have

$$\|h_1\|^2 = \sum_k \frac{2|v_k|^2}{|\mu_k|^2} \int_{-1}^1 \int_0^{2\pi} |\nabla^2 d\eta|^2 + |\nabla d\eta|^2 + |d\eta|^2 dt d\theta = c \sum_k \frac{|v_k|^2}{|\mu_k|^2}.$$

□

Definition 4.3. For $h = (h_0, h_1) \in T(U \times D^\ell)$ in the chart (4.2) set

$$(4.8) \quad \|(h_0, h_1)\|^2 = \|h_0\|_{WP}^2 + \|h_1\|^2.$$

After covering $\overline{\mathcal{M}}_{g,n}$ by such charts (with $\ell \geq 0$) and using a partition of unity, we see that (4.8) gives a Riemannian metric on $\overline{\mathcal{M}}_{g,n} \setminus \mathcal{N}$, well-defined up to uniform equivalence. We fix a metric in that equivalence class.

By Lemma 4.2 the metric (4.8) on $h = (h_0, \dot{\mu}) \in T_{C_0(\mu)}(U \times D^\ell)$ is uniformly equivalent to

$$(4.9) \quad \|(h_0, \dot{\mu})\|^2 \cong \|h_0\|_{WP}^2 + \sum_k \left| \frac{\dot{\mu}_k}{\mu_k} \right|^2.$$

Furthermore, the metric (4.7) is cylindrical in each coordinate: writing $\mu_k = e^{t+i\theta}$ the terms of (4.7) gives $|\mu_k|^{-2} \operatorname{Re} (d\mu_k)^2 = |d \log \mu_k|^2 = dt^2 + d\theta^2$. The

corresponding distance function is also cylindrical, so the distance squared between $\mu = e^{t+i\theta}$ and $\mu' = e^{s+i\theta'}$ is $|t - t'|^2 + |\theta - \theta'|^2 = \log |\mu'_k / \mu_k|^2$. Thus for ℓ -nodal curves C and C' the distance function of the metric (4.8) is given up to uniform equivalence by

$$(4.10) \quad \text{dist}^2(C(\mu), C'(\mu')) \cong \text{dist}^2(C, C') + \sum_k \left| \log \left(\frac{\mu'_k}{\mu_k} \right) \right|^2.$$

Thus the metric (4.8) on $\mathcal{M}_{g,n} = \overline{\mathcal{M}}_{g,n} \setminus \mathcal{N}$ is complete; near the stratum \mathcal{N}_ℓ of curves with ℓ nodes it is asymptotic to a cylinder $W_\ell \times (0, \infty)^\ell$ where W_ℓ is a bundle over \mathcal{N}_ℓ whose fiber is the real torus T^ℓ corresponding to the bundle (4.3).

Finally, observe that this geometry leads to a nonstandard compactification of $\mathcal{M}_{g,n}$: identify the end $W_\ell \times (0, \infty)^\ell$ with $W_\ell \times (0, 1)^\ell$ and compactify to $W_\ell \times (0, 1]^\ell$. This “cylindrical end compactification” projects down to the Deligne-Mumford compactification $\overline{\mathcal{M}}_{g,n}$ so that the fiber along the nodal stratum \mathcal{N}_ℓ is a copy of W_ℓ .

5. Renormalization at the nodes

In this section we will consider a sequence of δ -flat (J, ν) -holomorphic maps

$$(5.1) \quad f_n : C_n \rightarrow Z_{\lambda_n} \quad \text{with} \quad \lambda_n \rightarrow 0.$$

By the Compactness Theorem for holomorphic maps into Z these converge to a limit map f_0 from a nodal curve C_0 to Z_0 as described in Section 3; the convergence is in $L^{1,2}$, in C^0 , and in C^∞ on compact sets in the complement of the nodes. We will refine this by constructing renormalized maps \hat{f}_n around each node that is mapped into V , and proving convergence results for the renormalized maps. This gives detailed information about how the original maps f_n are converging in a neighborhood of the nodes.

As in Section 3, the limit domain C_0 is the union of (not necessarily connected) curves C_1 and C_2 which intersect at nodes, and f_0 decomposes into maps $f_1 : C_1 \rightarrow X$ and $f_2 : C_2 \rightarrow Y$ whose images meet along V with contact vector $s = (s_1, \dots, s_\ell)$. Thus, after regarding C_1 and C_2 as disjoint curves, there are points $x_k \in C_1$ and $y_k \in C_2$, $k = 1, \dots, \ell$, so that f_1 and f_2 contact V of order s_k at the point $q_k = f_1(x_k) = f_2(y_k)$ in V . For short, we simply write

$$f_n \rightarrow f_0 = (f_1, f_2) \in \mathcal{K}_\delta \subset \mathcal{M}_s \times_{\text{ev}} \mathcal{M}_s.$$

where \mathcal{K}_δ is the compact set introduced in (3.11). All the estimates in the next several sections will be uniform on \mathcal{K}_δ .

For the rest of this section we will focus attention on the restrictions of the maps $(f_n, \phi_n) : C_n \rightarrow Z \times \overline{\mathcal{U}}$ to the neck near a fixed node $x_k = y_k$ of $C_0 \subset \overline{\mathcal{U}}$.

There we have coordinates (z, w) centered on the node described before (4.4). However, it is often more convenient to parametrize the graphs $zw = \mu_n$ by the cylindrical coordinates (t, θ) used in the previous section. Thus we write

$$(5.2) \quad \phi_n(t, \theta) = (z, w) = (z, \mu_n/z) \quad \text{where} \quad z = \sqrt{|\mu_n|} e^{t+i\theta}$$

and, as in (1.2), regard (f_n, ϕ_n) as a map

$$(5.3) \quad (f_n, \phi_n) : [-T_n, T_n] \times S^1 \rightarrow Z_{\lambda_n} \times \overline{U}_{g,n}$$

with $T_n = |\log(2/\sqrt{|\mu_n|})| \rightarrow \infty$ as $n \rightarrow \infty$ and with the cylindrical metric (4.5) on the domain. We will use the cylindrical metric in all of our PDE estimates. Since (4.5) is a conformal change of metric we still have $f_n \rightarrow f_0$ in C^0 and in energy. Our goal is to bootstrap from there.

We also choose coordinates on Z around the image point $q = q_k \in V$ of the node as follows. Fix a J -equivariant identification of $T_q V$ with $\mathbb{R}^{2n-2} = \mathbb{C}^{n-1}$ and extend that to normal coordinates $\{v^i\}$ around q in V . Then identifying L_q with \mathbb{C} , taking the direct sum with the dual, and parallel translating along radial lines in the v coordinates, we obtain coordinates $(x, y) : L \oplus L^* \rightarrow \mathbb{C} \oplus \mathbb{C}$. This gives the desired coordinates

$$(5.4) \quad (v, x, y)$$

because, as in Section 2, $L \oplus L^*$ is the normal bundle to V in Z . In these coordinates the almost complex structure J on Z agrees with the standard complex structure J_0 on $\mathbb{C}^{n-1} \oplus \mathbb{C} \oplus \mathbb{C}$ at the origin, and J has the form $J_V \oplus J_{\mathbb{C}} \oplus J_{\mathbb{C}}$ along V . Notice that the v^i are real-valued while x and y are complex-valued, and in these coordinates $xy = \lambda$.

In these coordinates our maps have components $f_n = (v_n, x_n, y_n)$ where v_n is a map into V , $x_n = x \circ f_n$ is the coordinate normal to V in X , y_n is the coordinate of f_n normal to V in Y , and Z_λ is locally the graph of $xy = \lambda$. The expansions of f_0 provided by Lemma 3.4 of [IP4] and the matching condition of Lemma 3.3 show that

$$(5.5) \quad f_0(z, w) = (h^v, a_k z^{s_k}(1 + h^x), b_k w^{s_k}(1 + h^y))$$

where $|(h^v, h^x, h^y)| \leq c\rho$.

In the next lemma we fix $F = (f, \phi)$ as in (5.3) and estimate the neck energy

$$E(F, T) = \frac{1}{2} \int_{-T}^T \int_0^{2\pi} |dF|^2 dt d\theta.$$

LEMMA 5.1. *For each $(J, \nu) \in \mathcal{J}(Z)$, there are constants R_0, c_1, c_2 and E_0 such that if $F = (f, \phi) : [-T, T] \times S^1 \rightarrow Z$ is a (J, ν) -holomorphic map with $\text{diam } F(A_0) < R_0$ then*

$$(5.6) \quad E(F, t) \leq C E(F, T) \rho^{\frac{2}{3}}(t) \quad \forall |t| \leq T$$

with $C = 2\rho^{2/3}(T) \leq 2R_0^{2/3}$. If $E(F, T) \leq E_0$ then there are pointwise bounds of the form

$$(5.7) \quad |df|^2 \leq c_1 |dF|^2 \leq c_2 E_0 \rho^{\frac{2}{3}}(t) \quad \forall |t| \leq \frac{1}{2}T.$$

Proof. When R_0 is small the image of F lies in a coordinate ball in $Z \times \mathbb{P}^N$ and we can choose coordinates which identify that ball with the ball $B(0, R_0)$ in \mathbb{C}^{n+1} so that the complex structure \hat{J} of (1.3) agrees with the standard structure J_0 at the origin. On this ball the Euclidean metric is uniformly equivalent to the metric on $Z \times \mathbb{P}^N$, so it suffices to prove the lemma using the Euclidean norms on \mathbb{C}^{n+1} -valued maps.

Writing the (J, ν) -holomorphic map equation (1.3) in terms of the standard operator $\bar{\partial} = \bar{\partial}_{J_0}$ in those coordinates, we have $\bar{\partial}F = (J - J_0)dF j$. But $|J - J_0| \leq c|F| \leq cR_0$ in our coordinates, giving the pointwise bound

$$(5.8) \quad |\bar{\partial}F| \leq cR_0 |dF|.$$

By writing $F = u + iv$ as the sum of its real and imaginary parts, one finds that

$$4|\bar{\partial}F|^2 dt d\theta = |dF|^2 dt d\theta - 2d(u \cdot dv).$$

Integrating over $A = [-t, t] \times S^1$ and using Stokes' theorem gives

$$\frac{1}{2} \int_A |dF|^2 = 2 \int_A |\bar{\partial}F|^2 + \int_{\partial A} u \cdot v_\theta d\theta.$$

The boundary term is an integral over two circles. On each, we can replace u by $\tilde{u} = u - \frac{1}{2\pi} \int_0^{2\pi} u d\theta$ and apply the Hölder and Poincaré inequalities on the circle

$$(5.9) \quad \int u \cdot v_\theta d\theta = \int \tilde{u} \cdot v_\theta d\theta \leq \|\tilde{u}\|_{L^2} \|v_\theta\|_{L^2} \leq \|\tilde{u}_\theta\|_{L^2} \|v_\theta\|_{L^2} \leq \int |F_\theta|^2.$$

From the definition $2\bar{\partial}F = F_t + iF_\theta$ and the inequality $(a - 2b)^2 \leq 2a^2 + 8b^2$ we also obtain

$$3|F_\theta|^2 = 2|F_\theta|^2 + |F_t - 2\bar{\partial}F|^2 \leq 2|dF|^2 + 8|\bar{\partial}F|^2.$$

Together, the previous three displayed equations and (5.8) imply that

$$\int_A (1 - 4c^2 R_0^2) |dF|^2 \leq \frac{4}{3} \int_{\partial A} (1 + 4c^2 R_0^2) |dF|^2.$$

Taking R_0 small enough that $4c^2 R_0^2 < 1/44$ and adding the previous two equations lead to

$$\frac{2}{3}E(t) \leq E'(t).$$

Integrating this differential inequality from t to T yields (5.6).

On the cylinder $[-T, T] \times S^1$, each point lies in a unit disk with Euclidean metric, and $F = (f, \phi)$ is \hat{J} -holomorphic as in (1.3). Standard elliptic estimates then bound $|dF|^2$ at the center point in terms of the energy of F in that unit disk; see [PW, Th. 2.3]. Thus (5.6) implies (5.7). \square

In the next several sections we repeatedly use the facts that, because $\rho^2 = 2|\mu| \cosh(2t)$ is essentially exponential in t , the integrals of its powers in the cylindrical metric satisfy

$$(5.10) \quad \int_{\rho \leq \rho_0} \rho^\alpha dt d\theta \leq c_\alpha \rho_0^\alpha \quad \text{and} \quad \int_{\rho \geq \rho_0} \rho^{-\alpha} dt d\theta \leq c_\alpha \rho_0^{-\alpha} \quad \text{for } \alpha > 0.$$

We will also use the bump functions defined as follows. Fix a smooth function $\beta : \mathbb{R} \rightarrow [0, 1]$ supported on $[0, 2]$ with $\beta \equiv 1$ on $[0, 1]$. The function $\beta_\varepsilon(z, w) = \beta(\rho/\varepsilon)$ has support where $\rho^2 = |z|^2 + |w|^2 \leq 4\varepsilon^2$. When restricted to the curve $zw = \mu$, $\beta_\varepsilon \equiv 1$ on the neck region $A(\varepsilon) = \{\rho \leq \varepsilon\}$, and $d\beta_\varepsilon$ is supported on two annular regions where $\varepsilon \leq \rho \leq 2\varepsilon$. We can choose β so that its norm in the cylindrical metric (4.5) satisfies

$$(5.11) \quad |d\beta_\varepsilon| \leq 2.$$

We next consider the renormalized maps obtained from $f_n = (v_n, x_n, y_n)$ by centering v_n and rescaling x_n and y_n .

Definition 5.2. In the region $\rho \leq 1$ around each node define *renormalized maps* \hat{f}_n by

$$\hat{f}_n = (\hat{v}_n, \hat{x}_n, \hat{y}_n) = \left(v_n^1 - \bar{v}_n^1, \dots, v_n^k - \bar{v}_n^k, \frac{x_n}{az^s}, \frac{y_n}{bw^s} \right)$$

where \bar{v}_n^i is the average value of v_n^i on the center circle $\gamma_n = \{\rho = \sqrt{|\mu_n|}\}$ of the neck of C_n .

Whenever $\lambda_n = x_n y_n$ is nonzero x_n has no zeros and has (local) winding number s . Hence each \hat{x}_n has winding number zero, so the functions $\log \hat{x}_n$, and similarly $\log \hat{y}_n$, are well-defined. The convergence (5.5) shows that on each set $|z| \geq r$ we have $\hat{x}_n \rightarrow f_0^x / az^s = 1 + O(r)$ in C^1 , and similarly for \hat{y}_n . Thus there is a constant c so that

$$(5.12) \quad \sup_{r \leq |z| \leq 1} |\log \hat{x}_n| + \sup_{r \leq |w| \leq 1} |\log \hat{y}_n| \leq cr \quad \forall n \geq N = N(r).$$

LEMMA 5.3. *For each sequence f_n as in (5.1) we have $\lim_{n \rightarrow \infty} \frac{\lambda_n}{\mu_n^s} = ab$ at each node.*

Proof. For $G_n = \log \hat{x}_n$ the integral

$$\overline{G}_n(\rho) = \frac{1}{2\pi} \int_0^{2\pi} G_n d\theta$$

over the circles with fixed ρ — or equivalently fixed t — satisfies

$$2\pi \frac{d}{dt} \overline{G}_n = \int_0^{2\pi} \partial_t G_n d\theta = \int_0^{2\pi} (\partial_t + i\partial_\theta) G_n d\theta = 2 \int_0^{2\pi} x_n^{-1} \bar{\partial} x_n d\theta.$$

Here x_n is $x \circ f_n$ where x is the x -coordinate on $Z_\lambda = \{xy = \lambda\}$. A calculation similar to (4.4) shows that the differential $dx : TZ_\lambda \rightarrow \mathbb{C}$ satisfies $|dx|^2 = (1 + |\lambda|^2/|x|^4) = |x|^2/R^2$ where $R^2 = |x|^2 + |y|^2$. Also note that $\nu^N = O(R)$ by (1.15a) and that along V , in the coordinates (5.4), $J - J_0$ is acting on normal vectors is $O(R)$. Hence equation (1.2) gives $|\bar{\partial} f_n^N|^2 \leq |(J - J_0)df_n^N j + \nu^N \cdot d\phi_n|^2 \leq cR^2 |dF_n|^2$ where $F_n = (f_n, \phi_n)$. Thus

$$(5.13) \quad |x_n^{-1} \bar{\partial} x_n|^2 = |x_n^{-1} dx_n \circ \bar{\partial} f_n^N|^2 \leq |x_n|^{-2} |dx_n|^2 \cdot cR^2 |dF_n|^2 \leq c |dF_n|^2.$$

These equations and Lemma 5.1 imply that $|\frac{d}{dt} \overline{G}_n| \leq c_1 \rho^{1/3}$. Hence for $\rho \leq r$ and $n > N(r)$

$$(5.14) \quad |\overline{G}_n(\rho)| \leq |\overline{G}_n(r)| + c_1 \int_\rho^r \rho^{1/3} dt \leq |\overline{G}_n(r)| + c_2 r^{1/3} \leq c_3 r^{1/3}$$

where the last inequality uses (5.12). Since we are free to take r arbitrarily small, this implies that the average of $\log \hat{x}_n$ on the center circle γ_{μ_n} satisfies

$$\left| \frac{1}{2\pi} \int_{\gamma_{\mu_n}} \log \hat{x}_n \right| = \left| \overline{G}_n(\sqrt{|\mu_n|}) \right| \rightarrow 0$$

as $n \rightarrow \infty$. Symmetrically, the same limit statement holds with G replaced by $\log \hat{y}$. For each n we can then integrate the constant

$$\log \left(\frac{\lambda_n}{ab\mu_n^s} \right) = \log \left(\frac{x_n y_n}{az^s bw^s} \right) = \log(\hat{x}_n \hat{y}_n)$$

over γ_{μ_n} to see that

$$2\pi \log \left(\frac{\lambda_n}{ab\mu_n^s} \right) = \int_{\gamma_{\mu_n}} \log \left(\frac{\lambda_n}{ab\mu_n^s} \right) = \int_{\gamma_{\mu_n}} \log \hat{x}_n + \log \hat{y}_n \rightarrow 0.$$

The lemma follows. \square

Lemma 5.3 allows us to improve (5.12). For each r , $\sqrt{\mu} \leq r \leq 1$, the region where $\rho = \sqrt{|z|^2 + |w|^2}$ is at least r consists of two components, one of the form $|z| \geq r'$ and one of the form $|w| \geq r'$ where r' is essentially r . On the first, $\log \hat{x}_n$ is bounded as in (5.12). In the second region (5.12) gives a bound on $\log \hat{y}_n$, which we can now parlay into a bound on $\log \hat{x}_n$ using Lemma 5.3, the equations $zw = \mu_n$ and $x_n y_n = \lambda_n$, and the last four displayed equations of the above proof. Specifically, in the second region we have

$$|\log \hat{x}_n| = \left| \log \left(\frac{\lambda_n}{ab\mu_n^s} \right) - \log \hat{y}_n \right| \leq \left| \overline{G}_n(\sqrt{|\mu_n|}) \right| + cr \leq cr^{1/3}.$$

A similar bound holds for $\log \hat{y}_n$ in the first region. Thus there is a constant c so that

$$(5.15) \quad \sup_{r \leq \rho \leq 1} |\log \hat{x}_n| + |\log \hat{y}_n| \leq c r^{1/3} \quad \forall n \geq N = N(r).$$

We conclude this section by translating these bounds on the renormalized maps into bounds on weighted Sobolev norms of the original maps f_n .

LEMMA 5.4. *Given a sequence of δ_0 -flat (J, ν) -holomorphic maps as in (5.3) which satisfy the hypotheses of Lemma 5.1, write $f_n = (v_n, x_n, y_n)$ in the coordinates (5.4) and set $f_n^V = v_n - \bar{v}_n$ and $f_n^N = (x_n, y_n)$. Then for each $p \geq 2$ there are constants C_p and $N = N(r)$ so that whenever $0 < \delta \leq \frac{1}{3}$ and $n \geq N$ the restriction of f_n to the cylinder $A(r) = \{\rho(t) \leq r\} \subset [-T, T] \times S^1$ satisfies*

$$(5.16) \quad \int_{A(r)} (|\nabla f_n^V|^p + |f_n^V|^p + |\rho^{1-s} \nabla f_n^N|^p + |\rho^{1-s} f_n^N|^p) \rho^{-p\delta/2} dt d\theta \leq C_p r^{p/6}.$$

Proof. Write $A(r)$ as the union of cylinders $A_k = \{k \leq t \leq k+1\}$ of unit size and let ρ_k be the value of ρ at one end of A_k . Using the Sobolev inequality which bounds oscillation by the L^4 norm of the derivative and noting that $|df_n^V| = |dv_n| \leq c\rho^{1/3}$ by (5.7), we have

$$(5.17) \quad \sup_{A(r)} |f_n^V| \leq \sum_k |\text{osc}_{A_k} \hat{v}_n| \leq C \sum_k \|dv_n\|_{4, A_k} \leq C \sum_k \rho_k^{1/3} \leq C r^{1/3}$$

where the last inequality comes from the Riemann sum for $\int \rho^{1/3} dt$. Thus $|\nabla f_n^V|^p + |f_n^V|^p \leq c\rho^{p/3}$ pointwise. Integration via (5.10) then gives the first two terms of (5.16).

Next, the Calderon-Zygmund inequality of [IS] shows that the L^p norm of $G = \log \hat{x}_n$ satisfies

$$\begin{aligned} \|dG\|_{p, A(r)} &\leq C \|\bar{\partial}(\beta_r G)\|_{p, A(2r)} \\ &\leq C (\|d\beta_r \cdot G\|_{p, A(2r) \setminus A(r)} + \|x_n^{-1} \bar{\partial} x_n\|_{p, A(2r)}). \end{aligned}$$

We can estimate the last term by integrating using (5.13), (5.7), and (5.10), and can estimate the $d\beta_r \cdot G$ term using (5.11), the bound (5.15) in the region $r \leq \rho \leq 2r$ where $d\beta_r \neq 0$, and (5.10). These imply that the L^p norm of dG is bounded by $c r^{1/3}$. But then for each cylinder $A \subset A(r)$ with unit diameter we can use (5.14) and the Sobolev inequality as in (5.17) to obtain

$$\begin{aligned} \sup_A |G| &\leq |\text{avg}_{\partial A} G| + |\text{osc}_A G| \\ &\leq c r^{1/3} + C \|dG\|_{4, A(r)} \leq c r^{1/3} \quad \text{for all } n \geq N(r). \end{aligned}$$

Exponentiating this bound on G shows that $|\hat{x}_n - 1| \leq c r^{1/3}$ in $A(r)$, and that in turn gives $|d\hat{x}_n| = |\hat{x}_n dG| \leq c |dG|$. Consequently $x_n = \hat{x}_n \cdot a z^s$ satisfies

$$(5.18) \quad |x_n| \leq c \rho^{s+1/3} \quad \text{and} \quad |dx_n| \leq c \rho^s (1 + |dG|)$$

since $|dz/z|$ is bounded in the cylindrical metric. The same bounds hold for the y_n , so integration, combined with the L^p bound on dG , gives the remaining part of (5.16). \square

6. The space of approximate maps

The limit arguments of Section 3 show that as $\lambda \rightarrow 0$ sequences of holomorphic maps f_n into Z_λ have subsequences which converge to maps into $X \cup Y$ with matching conditions along V , i.e. to maps in $\mathcal{M}_s^V(X) \times_{\text{ev}} \mathcal{M}_s^V(Y)$. The results of Section 5 give further information about the convergence near the matching points: they show that for small λ the maps f_n are closely approximated by maps $g(z, w) = (\bar{v}, az^s, bw^s)$ in local coordinates. Over the next four sections we will reverse this process, showing how one can use $\mathcal{M}_s^V(X) \times_{\text{ev}} \mathcal{M}_s^V(Y)$ to construct a model $\text{Model}_s(Z_\lambda)$ for the space of stable maps into Z_λ . The final result is stated as Theorem 10.1.

The construction has two main steps. In the first, maps f in a compact set $\mathcal{K}_\delta \subset \mathcal{M}_s^V(X) \times_{\text{ev}} \mathcal{M}_s^V(Y)$ are smoothed in a canonical way to construct maps F into Z_λ which are approximately holomorphic. The second step corrects those approximate maps F to make them truly holomorphic. This section describes the canonical smoothing and the resulting space of approximate maps and introduces norms on the space of maps which capture the convergence of the renormalized maps. Those norms lead to a precise statement that the approximate maps are nearly (J, ν) -holomorphic.

The maps alone cannot be canonically smoothed — more data are needed. This harks back to the comment at the end of Section 3 that each f will generally be the limit of many maps into Z_λ . Recall that an element of $\mathcal{M}_s^V(X) \times_{\text{ev}} \mathcal{M}_s^V(Y)$ is a map $f : C_0 \rightarrow Z_0$ from a bubble domain C_0 whose last $\ell(s)$ nodes $x_k = y_k$ are mapped into V with contact of order s_k . By the construction of Section 4 each such C_0 determines an ℓ -dimensional family $C_0(\mu)$, $\mu = (\mu_1, \dots, \mu_\ell)$, which smooths those ℓ nodes (and leaves other nodes unaffected). Lemma 5.3 indicates that $f_0 : C_0 \rightarrow Z_0$ is the limit of maps into Z_λ which satisfy

$$(6.1) \quad a_k b_k \mu_k^{s_k} = \lambda$$

to highest order. That leaves $|s| = s_1 s_2 \cdots s_\ell$ possibilities for μ corresponding to the different choices of root for each μ_k .

As in (5.5), the coefficient a_k is the s_k -jet of the component of f_1 normal to V at x_k modulo higher order terms, and so, intrinsically,

$$a_k \in (T_{x_k}^* C)^{s_k} \otimes (N_X V)_{f(x_k)}.$$

Globally on the space of relative stable maps there are two complex line bundles associated with the marked point x_k : (i) the bundle $\text{ev}_s k^* N_X V$ obtained by pulling back $N_X V$ by the evaluation map at x_k and (ii) the relative cotangent

bundle \mathcal{L}_k to the domain (over the maps with stable domains, \mathcal{L}_k is the pullback st^*L_k of the bundle $L_k \rightarrow \overline{\mathcal{M}}_{g,n}$ that appears in (4.3)). The leading coefficients in (5.5) are then sections

$$(6.2) \quad a_k \in \Gamma(\mathcal{L}_k^{s_k} \otimes \text{ev}_s k^* N_X V) \quad \text{and} \quad b_k \in \Gamma((\mathcal{L}'_k)^{s_k} \otimes \text{ev}_s k^* N_Y V).$$

Furthermore, λ is a constant section of $N_X V \otimes N_Y V \cong \mathbb{C}$ via the trivialization fixed at the beginning of Section 2. Thus (6.1) implies that at each node which is mapped into V the coefficients a_k, b_k determine a section

$$\frac{\lambda}{a_k b_k} \in \Gamma\left((\mathcal{L}_k \otimes \mathcal{L}'_k)^{-s_k}\right)$$

over $\mathcal{M}_s^V(X) \times_{\text{ev}} \mathcal{M}_s^V(Y)$. The s_k^{th} root of this section is a multisection of $\mathcal{L}_k^* \otimes (\mathcal{L}'_k)^*$; considering all k at once defines a multisection of the direct sum of the $\mathcal{L}_k^* \otimes (\mathcal{L}'_k)^*$. This gives an intrinsic model for our space $\text{Model}_s(Z_\lambda)$ of approximate maps.

Definition 6.1 (Model space). For each s and $\lambda \neq 0$, the *model space* $\text{Model}_s(Z_\lambda)$ is the multisection of

$$\bigoplus_{k=1}^{\ell} [\mathcal{L}_k^* \otimes (\mathcal{L}'_k)^*] \rightarrow \mathcal{M}_s^V(X) \times_{\text{ev}} \mathcal{M}_s^V(Y)$$

whose fiber over a map f consists of those $\mu = (\mu_1, \dots, \mu_\ell)$ which satisfy

$$(6.3) \quad \mu_k^{s_k} = \frac{\lambda}{a_k b_k} \quad \text{for each } k.$$

This model space is an $|s|$ -fold cover of $\mathcal{M}_s^V(X) \times_{\text{ev}} \mathcal{M}_s^V(Y)$, and hence its irreducible part is an orbifold for generic V -compatible (J, ν) . Elements of the model space are triples (f, C_0, μ) where $f : C_0 \rightarrow Z_0$ and $\mu = (\mu_1, \dots, \mu_\ell)$ satisfies (6.3). Each such element gives rise to an approximate holomorphic map as follows.

From (4.2) and (5.4), we have coordinates (z_k, w_k) centered on the k^{th} node x_k defined on the region $B_k(2)$ where $|z_k|^2 + |w_k|^2 \leq 4$, and coordinates (v, x, y) centered on the image $f(x_k)$ whenever $f(x_k) \in V$. Let β_{μ_k} be the bump function (5.11) with $\varepsilon = |\mu_k|^{1/4}$ determined from λ by (6.3). Using the notation of (5.5), set $\tilde{x}(z_k) = a_k z_k^{s_k} (1 + (1 - \beta_{\mu_k})h^x)$ and $\tilde{y}(w_k) = b_k w_k^{s_k} (1 + (1 - \beta_{\mu_k})h^y)$.

Definition 6.2 (Approximate maps). For each $(f, C_0, \mu) \in \text{Model}_s(Z_\lambda)$, $\lambda \neq 0$, define an *approximate holomorphic map* $F = F_{f, C_0, \mu} : C_\mu \rightarrow Z_\lambda$ by taking C_μ to be the smoothing $C_0(\mu)$, defining F on $C_\mu \cap B_k(2)$ by

$$(6.4) \quad F(z_k, w_k) = \begin{cases} \left((1 - \beta_{\mu_k})h^v(z_k), \tilde{x}(z_k), \frac{\lambda}{\tilde{x}(z)} \right) & \text{if } |z_k| \geq |w_k| \\ \left((1 - \beta_{\mu_k})h^v(w_k), \frac{\lambda}{\tilde{y}(w)}, \tilde{y}(w_k) \right) & \text{if } |z_k| \leq |w_k|, \end{cases}$$

whenever $f(x_k) \in V$, and extending by $F = f$ outside the support of $\sum_k \beta_{\mu_k}$.

Notice that (6.4) is a smooth map into Z_λ which is equal to $(0, a_k z^{s_k}, b_k w^{s_k})$ in the center region where $\beta_{\mu_k} = 1$. In the region where $\beta_{\mu_k} = 0$ and $|z_k| \geq |w_k|$, C_μ is identified with C_1 by $(z, w) \mapsto (z, 0)$ as in Section 4 and (6.4a) is simply $(f_1(z), \lambda)$ in the chart (2.6a). Symmetrically, (6.4b) reduces to $(f_2(w), \lambda)$ in the chart (2.6b). Hence the maps (6.4) do indeed extend as f over the rest of C_μ .

Remark 6.3. In general, C_0 has nodes of two types: those mapped into V and those not mapped into V . When all nodes are of the second type, the gluing problem is relatively easy and has been carefully treated in the literature (see for example [MS] or [RT1]). In fact, the estimates around the two types of nodes are isolated from one another, as follows. Recall that $f^{-1}(V)$ is a finite set of nodes and, as in Section 4, distinct nodes are separated by distance at least 4. Hence each node x with $f(x) \notin V$ lies in a region $B = B_k(2)$ with $\text{dist}(B, f^{-1}(V)) \geq 2$. Lemma 6.8a below then shows that B is mapped into the complement of a fixed tubular neighborhood of $V \subset Z$; there the geometry of Z_λ is uniform in λ and the local estimates involved in the Ruan-Tian gluing apply.

Thus near each node $x_k \in C_0$ which is not mapped into V we can smooth C_0 to C_μ , define approximate maps on $C_\mu \cap B_k(2)$ as Ruan-Tian do, and use their estimates on $C_\mu \cap B_k(2)$. That effectively reduces the analysis to the case where there are no such nodes. Bearing that in mind, in the next several sections we will simplify the analysis by assuming that all nodes of C_0 lie in $f^{-1}(V)$, leaving the conflation of estimates to the reader.

Definition 6.4 (Gluing map). The association $(f, C_0, \mu) \mapsto (F_{f, C_0, \mu}, C_\mu)$ defines a *gluing map*

$$(6.5) \quad \Gamma_\lambda : \mathcal{M}\text{odel}_s(Z_\lambda) \rightarrow \text{Maps}(B, Z_\lambda \times \mathcal{U})$$

where B is the 2-manifold underlying C_μ . The image of Γ_λ is the *space of approximate maps*

$$\mathcal{A}\text{pprox}_s(Z_\lambda).$$

This gluing map is injective for small μ as follows. If $\Gamma_\lambda(f, C_0, \mu) = \Gamma_\lambda(f', C'_0, \mu')$ then the (stabilized) curve $C_\mu = C'_\mu$ lies in the family (4.2), $C_0 = C'_0$ and $\mu = \mu'$. But then f and f' are (J, ν) -holomorphic maps which agree outside the balls $B_k(2)$ and therefore, by the unique continuation property of elliptic equations, agree everywhere.

In Section 9 we will show that Γ_λ is an embedding. Here, as a preliminary, we introduce norms which make the space of maps in (6.5) into a Banach manifold.

We will use weighted Sobolev norms tailored for our problem. On the domain we continue to use the cylindrical metric (4.5) and to use (4.10) to

measure distance between curves. In the target, in a neighborhood of V in Z , the tangent bundle TZ_λ splits as in the proof of Lemma 2.3 as the direct sum of the tangent space to the complex curve $xy = \lambda$ in the fiber $(L \oplus L^*)_v = \mathbb{C}^2$, which is a complex subbundle N_λ of TZ_λ , and an orthogonal subspace, which we can identify with TV . Thus in a neighborhood of V each vector field ξ on Z_λ decomposes orthogonally as

$$(6.6) \quad \xi = (\xi^V, \xi^N) \in \Gamma(TV \oplus N_\lambda).$$

Fix $\rho_1 > 0$, let $B_k(\rho_1)$ denote the unit ball centered on the k^{th} node in the family \mathcal{F} of (4.2), and let $q_k \in V$ be the image of that node (we will specify a precise ρ_1 at the end of this section). For small μ the neck $C_\mu \cap B_k(1)$ of C_μ is an annulus whose center circle γ_k is defined by $\rho = \sqrt{|\mu_k|}$. Given ξ , define the average value of the V components on γ_k

$$(6.7) \quad \bar{\xi}_k = \frac{1}{2\pi} \int_{\gamma_k} \xi^V \in T_{q_k} V$$

and assemble these averaged vectors at the different nodes into a single vector $\bar{\xi} = (\bar{\xi}_1, \dots, \bar{\xi}_\ell) \in T_q V^\ell$ where $q = (q_1, \dots, q_\ell)$. We can extend $\bar{\xi}_k$ to a vector field (also denoted by $\bar{\xi}_k$) in a neighborhood of q_k by parallel translation along geodesics in V emanating from q_k and then lifting to a vector field on TZ_λ perpendicular to N_λ ; that extension has the form $(\bar{\xi}_k^V, 0)$ in the decomposition (6.6). Each ξ then determines a global vector field on C_μ :

$$(6.8) \quad \zeta = \xi - \sum \beta_k \bar{\xi}_k$$

with zero average value, where β_k is the bump function (5.11) with $\varepsilon = \rho_1/2$ centered on the k^{th} node.

Finally, we fix $0 < \delta < 1/6$ and define a weighted Sobolev norm by

$$(6.9) \quad \|\zeta\|_{k,p,s}^p = \int_{C_\mu} \left(|\nabla^k \zeta|^p + |\zeta|^p \right) \rho^{-\delta p/2}$$

using the norms, covariant derivative, and volume form associated to the cylindrical metric (4.5) on the domain and the metric induced on Z_λ from Z . Thus near each node in the coordinates (5.3) the metric on the domain is the flat cylindrical metric and the weighting function ρ is exponential in t .

Definition 6.5. Given a tangent vector (ξ, h) to the space $\text{Map}_s(B, Z_\lambda \times \mathcal{U})$ we form the triple $(\zeta, \bar{\xi}, h)$ as in (6.8) and define the *weighted Sobolev m norm*

$$(6.10) \quad \|(\xi, h)\|_m = \|\zeta\|_{m,2,s} + \|\zeta\|_{m,4,s} + |\bar{\xi}| + \|h\|$$

where $\|h\|$ is given by (4.8). For 1-forms $\eta \in \Gamma(\Lambda^{01}(f^*TZ_\lambda))$ we do the same without averaging:

$$(6.11) \quad \|\eta\|_m = \|\eta\|_{m,2,s} + \|\eta\|_{m,4,s}.$$

By the Sobolev embedding $L_{\text{loc}}^{1,4} \hookrightarrow C_{\text{loc}}^0$ the norm $\|(\zeta, \bar{\xi}, k)\|_1$ dominates the C^0 norm. Hence we can use it to complete the space of C^∞ maps, making $\text{Map}_s(B, Z_\lambda \times \mathcal{U})$ a Banach manifold with a Finsler metric given by (6.10).

Remark 6.6. The norms defined above make sense also at $\lambda = 0$, where the average value $\bar{\xi}$ is equal to the value of ξ^V at the nodes. The choice of norms is guided by the need to dominate the C^0 norm and the need to have uniform estimates as $\lambda \rightarrow 0$. Note that as $\lambda \rightarrow 0$ the domain becomes, near each node, isometric to cylinders $[-T, T] \times S^1$ with $T \rightarrow \infty$. Thus the L^4 norm does not dominate the L^2 norm uniformly in λ , and the $L^{1,4}$ norm does not give uniform bounds on energy. The choice (6.10) with $m = 1$ uniformly controls both C^0 norms and energy and is also strong enough to initiate the bootstrap argument establishing the regularity of the fixed point at the end of Proposition 9.4.

Each $(f_0, C_0) \in \mathcal{K}_\delta$ can be stabilized as in Remark 1.1 by choosing submanifolds transverse to the image. This procedure stabilizes the domains of all maps in a neighborhood of (f_0, C_0) in \mathcal{K}_δ . Because C_0 is then stable, we can choose a local chart \mathcal{F} of the form (4.2) for a neighborhood of C_0 in the universal curve. Elements of \mathcal{K}_δ close to (f_0, C_0) can be regarded as maps f from a fiber of \mathcal{F} into Z . That provides a local “chart” which identifies a neighborhood of (f_0, C_0) in \mathcal{K}_δ with the set of maps from the fibers of \mathcal{F} into Z . In that chart $\phi = \text{Id}$ in all formulas, and a sequence (f_n, C_n) converges if and only if the C_n converge as fibers of \mathcal{F} and the f_n converge in C^0 , in energy, and in C^∞ on compact sets in the complement of the nodes.

Observation 6.7. In the next four sections we will work in the charts just described, seeking estimates which are uniform for (f, C) in the set \mathcal{K}_δ of (3.11). Of course, \mathcal{K}_δ is compact, so that *any estimate which holds locally in a chart at each $(f, C) \in \mathcal{K}_\delta$ holds uniformly in \mathcal{K}_δ .*

Here is one application of this observation. For every $(f, C) \in \mathcal{K}_\delta$ let $A(\rho_0) = \bigcup_k D(x_k, \rho_0)$ denote the region within distance $\rho_0 > 0$ of $f^{-1}(V) = \{x_k\}$.

LEMMA 6.8. *There are constants $c_i, c'_i > 0$ such that for every $(f, C) \in \mathcal{K}_\delta$,*

- (a) *for each ε there is a $\rho_0 > 0$ such that $\text{dist}(f(A(\rho_0)), V) \leq \varepsilon$, and for each $\rho_0 > 0$ there is a $c > 0$ such that $\text{dist}(f(C \setminus A(\rho_0)), V) \geq c$.*
- (b) *$|df| \leq c_1 \rho^{1/3} \leq c_2$ on C and $\|f\|_{C^\ell} \leq c'_\ell \rho^\ell$ on some $A(\rho_0)$ for each ℓ .*
- (c) *The coefficients in (6.3) satisfy $c_3 \leq |a_k|, |b_k| \leq c_4$.*
- (d) *The components of (5.5) satisfy $|h| + |dh| \leq c_5 \rho$ on some $A(\rho_0)$.*

Proof. By Observation 6.7 it suffices to show (a) and (b) on a neighborhood of a fixed $(f_0, C_0) \in \mathcal{K}_\delta$. If the first assertion of part (a) fails, then there exists an $\varepsilon_0 > 0$ and sequences (f_n, C_n) of maps and p_n of points such that (f_n, C_n) converges to (f_0, C_0) in the stable map topology, p_n converges to one of the nodes x_k , but the distance from $f_n(p_n)$ to V is bounded below by ε_0 . Now, since f_n converges in C^0 to f and $f(x_k) \in V$ this gives a contradiction. Similarly, if the second part of statement (a) fails on every neighborhood of (f_0, C_0) we can find a sequence (f_n, C_n) converging to (f_0, C_0) and points $p_n \in C_n \setminus A$ converging to $p \in C_0$ with $f(p) \in V$. Since the p_n are uniformly bounded away from the points of $f_n^{-1}(V)$, p must be a new intersection point, contradicting the invariance of s described in Remark 3.5. Thus (a) holds.

Now recall from Section 4 that the special points on any C are separated by distance at least 4. Thus for $\rho_0 < 1$ consider the region $B(\rho_0)$ defined as the disjoint union of the ρ_0 -neighborhoods of the nodes not mapped into V . We can find a $\rho_0 > 0$ small enough so that all maps (f, C) in a small neighborhood of (f_0, C_0) satisfy the hypotheses of Lemma 5.1 on $B(\rho_0)$. That Lemma then gives the bound $|df| \leq c_1 \rho^{1/3}$ on $B(\rho_0)$ (in cylindrical metric on the domain).

The bound $\|f\|_{C^\ell} \leq c_\ell \rho^\ell$ in the cylindrical metric (4.1) is equivalent to $\|f\|_{C^\ell} \leq c_\ell$ in the metric $g_{\mathcal{U}}$ on the universal curve. Using the latter metric, suppose there is no such bound on $C \setminus B(\rho_0)$ on every neighborhood of (f_0, C_0) . Then there exists $(f_n, C_n) \in \mathcal{K}_\delta$ converging to (f_0, C_0) and points p_n converging to $p \in C_0$ with $\rho(p_n) \geq \rho_0$ and $\|f\|_{C^\ell} \rightarrow \infty$. But restricting to $f_n^{-1}(X)$, then to $f_n^{-1}(Y)$, the sequence p_n is bounded away from the nodes, so stable map convergence implies that $f_n \rightarrow f_0$ in C^∞ , a contradiction. The bounds (b) follow after we note that $A(\rho_0) \subset C \setminus B(\rho_0)$.

The coefficients a_k and b_k , as sections over \mathcal{K}_δ as in (6.2), are smooth by Lemma 7.1 below and cannot vanish because of the invariance of s (Remark 3.5). Since \mathcal{K}_δ is compact they are uniformly bounded above and below. Finally, when f is expanded as in (5.5) its x component satisfies $|f^x - az^s| \leq c_s'' |z|^{s+1}$ by the C^s bound in (b). There are similar bounds on other components which, together with the bounds of part (c) yield $|h| \leq c\rho$. Essentially similar estimates give $|dh| \leq c|dz| \leq c\rho$ in the cylindrical metric. \square

We conclude this section by showing that the approximate maps are nearly holomorphic. The specific statement is that the quantity $\bar{\partial}_{J_F} F - \nu_F$, which measures the failure of the approximate map to be (J, ν) -holomorphic, is small in the norm (6.11).

LEMMA 6.9. *There exist constants c and λ_0 , uniform on \mathcal{K}_δ , such that for $|\lambda| < \lambda_0$ and $(f, C_0) \in \mathcal{K}_\delta$ each approximate map $F = F_{f, C_0, \mu}$ satisfies*

$$\|\bar{\partial}_{J_F} F - \nu_F\|_0 \leq c |\lambda|^{\frac{1}{2|s|}}.$$

Proof. Let A be the union of the annuli A_k in the domain C_μ near the node $x_k = y_k \in f^{-1}(V)$ defined by $\rho \leq \rho_0$ where ρ_0 is the constant of Lemma 6.8a for ε small enough to ensure that $f(A_k)$ lies in a geodesic ball in which we have coordinates (v, x, y) of (5.4). First consider the set $C_\mu \setminus A$.

By Lemma 6.8a the image of $C_\mu \setminus A$ lies in the region of Z_λ with distance at least c to V . In that region the projection gives an identification of Z_λ with Z_0 under which $F = f$ for small λ , while the distance in Z between $F(z) \in Z_\lambda$ and $f(z) \in Z_0$ is at most $c|\lambda|$. That gives the estimate $|J_F - J_f| + |\nu_F - \nu_f| \leq c|\lambda|$. We can then bound the quantity $\Phi_F = \bar{\partial}_{J_F} F - \nu_F$ by noting that $\bar{\partial}_{J_f} f - \nu_f = 0$ and $\bar{\partial}_{J_F} = \bar{\partial}_{J_f} + (J_F - J_f) \circ dF \circ j$:

$$(6.12) \quad |\Phi_F| \leq (|J_F - J_f| + |\nu_F - \nu_f|) |dF| \leq c|\lambda| |df|$$

where $|df| \leq c_2$ by Lemma 6.8b. Then (5.10) immediately gives

$$(6.13) \quad \left(\int_{C_\mu \setminus A} \rho^{-\delta p/2} |\Phi|^p \right)^{1/p} \leq c|\lambda|.$$

Now we focus on the half A_+ of one A_k where $|w_k| \leq |z_k|$ (the estimates in the other half are symmetric). Omitting the subscripts k , we have local coordinates z, w in which, from (6.4a) and (6.3),

$$(6.14) \quad F = ((1 - \beta)h^v, az^s(1 + (1 - \beta)h^x), bw^s(1 + (1 - \beta)h^x)^{-1})$$

and f is given by the same formula with the last entry replaced by zero. Lemma 6.8d shows that, after possibly making ρ_0 smaller, we can assume that $|h^x| \leq 1/2$ on A_k . We then have $|F - f| \leq c|bw^s|$. Differentiating, noting that $w = \sqrt{|\mu|} \exp(t + i\theta)$ satisfies $|dw| = |w|$ in the cylindrical metric on the domain, and using (5.11) and Lemma 6.8cd yields $|dF - df| \leq c|b|(|dw^s| + |w^s h^x d\beta| + |w^s dh^x|) \leq c|w^s|$. But $|\rho w|^2 \leq 2|\mu|^2$ in A_+ and $|\mu^s| \leq c|\lambda|$ by (6.3) and Lemma 6.8d. Thus

$$(6.15) \quad |F - f| + |dF - df| + |J_F - J_f| \leq c \min\{\rho^s, |\lambda|\rho^{-s}\}$$

on A_k . Using the bound of Lemma 6.8b we also have $|dF| \leq |dF - df| + |df| \leq c\rho$.

Let J_0 be the complex structure of the coordinate system on the target used in (6.14); this agrees with J at the origin, so that $|J_F - J_0| \leq c|F| \leq c\rho$. As above, the J -holomorphic map equation for f leads to the estimate

$$(6.16) \quad |\Phi_F| \leq |\bar{\partial}_0(F - f)| + |(J_F - J_0) \circ (dF - df)| + |(J_F - J_f) \circ dF| + |\nu_F - \nu_f|.$$

Also as above, the term $|\bar{\partial}_0(F - f)| \leq |bw^s \bar{\partial}_0(1 + (1 - \beta)h^x)^{-1}|$ is bounded by $c\rho|w^s|$. Furthermore, the one-form ν satisfies $|\nu_F - \nu_f| \leq c|F - f| \leq c|w^s|$ in the metric on the universal curve, so $|\nu_F - \nu_f| \leq c\rho|w^s|$ in the cylindrical metric (4.1). With these facts, (6.16) reduces to

$$(6.17) \quad |\Phi_F| \leq c \min\{\rho^{1+s}, |\lambda|\rho^{1-s}\}$$

on A_+ and by symmetry on all of A_k , with this constant c uniform on \mathcal{K}_δ . Now integrate over A_k using (5.10) and note that $\rho^2 = |z|^2 + |w|^2 \geq 2|zw| = 2|\mu|$ on A_k . Recalling that $\delta < 1/6$, $|\mu_k^{s_k}| \leq c|\lambda|$, and $s_k \leq |s|$, and using Lemma 6.8c we obtain

$$(6.18) \quad \left(\int_{A_k} \rho^{-\delta p/2} |\Phi_F|^p \right)^{1/p} \leq c |\lambda| |\mu_k|^{(2-2s_k-\delta)/4} \leq c |\lambda|^{\frac{1}{2|s|}}.$$

The lemma follows by combining (6.13) and (6.18) and summing on k and on $p = 2, 4$. \square

We can now be precise about the choice of the radius ρ_1 of support of the cutoff functions β_k which appear in (6.8). Specifically, we take ρ_1 to be the ρ_0 of the first sentence of the proof of Lemma 6.9. That ensures that for each $(f, C_0) \in \mathcal{K}_\delta$ the image of the support of β_k lies in a region where we have the coordinates (5.4) and where the extension $\bar{\xi}_k$ and the vector field (6.8) are well-defined.

7. Linearizations

This section describes the linearization of the (J, ν) -holomorphic map equation as an operator on the Sobolev spaces of Definition 6.5. We describe first the moduli space $\mathcal{M}_s^V(X)$ which defines the relative invariants, then the space $\mathcal{M}_s(Z_0) = \mathcal{M}_s^V(X) \times_{\text{ev}} \mathcal{M}_s^V(Y)$ of maps into the singular space Z_0 . This serves as background for our main purpose: describing the linearization operator $\mathbf{D}_{f,C}$ at an approximate map into Z_λ and its adjoint $\mathbf{D}_{f,C}^*$. These operators are used to be used in Section 9 to correct the approximate maps into holomorphic maps.

To begin, consider a V -regular map $(f, \phi) : C \rightarrow X \times \bar{\mathcal{U}}_{g,n}$ in the space (1.17). For simplicity we usually omit mentioning χ, n and A , and will suppress the ϕ by considering C to be a curve with the complex structure induced by ϕ . Thus $(f, C) \in \mathcal{M}_s^V(X)$ means that $f : C \rightarrow X$ is a map from a smooth complex curve marked by the $\ell(s)$ points $\{x_k\}$ of $f^{-1}(V)$ with corresponding multiplicity vector $s = (s_1, \dots, s_\ell)$. Our first aim is to describe the linearization (1.10) with the norms (6.10) and (6.11). Thus we decompose ξ into $(\zeta, \bar{\xi})$ and regard ζ as an element of the Banach space $L_{1;s,0}(f^*TX)$ obtained by completing the smooth sections with average value $\bar{\zeta} = 0$ in the norm $\|\zeta\|_{1,2,s} + \|\zeta\|_{1,4,s}$ as in (6.10). One can show, as we prove in the context of Lemma 7.3 below, that the linearization extends to a bounded linear map

$$(7.1) \quad \mathbf{D}_{f,C} : L_{1;s,0}(f^*TX) \oplus T_q V^\ell \oplus T_C \bar{\mathcal{M}}_{g,n+\ell} \rightarrow L_s(\Lambda^{01}(f^*TX))$$

defined in terms of the linearization (1.10) by

$$(7.2) \quad \mathbf{D}_{f,C}(\zeta, \bar{\xi}, h) = D_{f,C} \left(\zeta + \sum \beta_k \bar{\xi}_k, h \right)$$

(the last space in (7.1) is the completion in the norm (6.11) with $m = 0$). Thus $\mathbf{D}_{f,C}$ and $D_{f,C}$ are essentially the same operator; the boldface indicates that we have separated out the average value vector $\bar{\xi}$ and completed in the weighted norms.

Using results from [Loc] it is routine to verify (cf. Lemma 3.4 and Lemma 7.3 below) that for generic $0 < \delta < 1$, (7.1) is Fredholm with respect to these norms. Furthermore, infinitesimal deformations of $f : C \rightarrow X$ in $\text{Ker } \mathbf{D}_{f,C}$ preserve the multiplicity vector $s = (s_1, \dots, s_\ell)$ (because s can be expressed as a winding number at infinity and the norm (6.9) dominates the C^0 norm). Thus (7.1) locally models the space $\mathcal{M}_s^V(X)$. In particular, for generic V -compatible (J, ν) the irreducible part $\mathcal{M}_s^V(X)^*$ of $\mathcal{M}_s^V(X)$ is an orbifold of dimension $2 \text{ind}_{\mathbb{C}} \mathbf{D}_{f,C}$. We also have the following two facts.

LEMMA 7.1. (a) $\text{st} \times \text{ev} : \mathcal{M}_s^V(X)^* \rightarrow \overline{\mathcal{M}}_{g,n+\ell} \times V^\ell$ is a smooth map of Banach manifolds.

(b) The leading coefficient (6.2) is a smooth section of the bundle $\mathcal{L}_k^{s_k} \otimes \text{ev}_s k^* N_X V$ over $\mathcal{M}_s^V(X)^*$.

Proof. (a) It suffices to show that the linearization is a smooth map everywhere. Let (ξ, h) be a tangent vector to $\mathcal{M}_s^V(X)$ at $f : C \rightarrow X$ and decompose $\xi = \zeta + \sum \beta_k \bar{\xi}_k$ with $\bar{\xi}_k = \xi(x_k) \in T_{f(x_k)} V$. The linearization of the map $\text{st} \times \text{ev}$ is $(\xi, h) \rightarrow (h, \bar{\xi})$ which is obviously smooth with our norms.

(b) Choose a path (f_t, C_t) in \mathcal{M}_s^V and let (ξ, h) be its tangent vector at $t = 0$. Assume for simplicity that $\ell = 1$ and let $f_t^N = a_t z^s + O(|z|^{s+1})$ be the expansion near the single point x with $q = f(x) \in V$. Writing $\xi = \zeta + \beta \bar{\xi}$ with $\bar{\xi} = \xi(x) \in T_q V$ and differentiating, we see that the tangential component $\zeta^V(x) \in T_q V$ at x vanishes and the normal component is $\zeta^N = \dot{a}_t z^s + O(|z|^{s+1})$. Consequently, $|\dot{a}_t| \leq c |\zeta^N|_{C^s} \leq c \|(\xi, h)\|_{C^s}$, and that is dominated by $c \|(\xi, h)\|_1$ by elliptic bootstrapping for the equation $D_{f,C}(\xi, h) = 0$. Thus the differential of the section (6.2) is bounded. \square

For maps into the singular space Z_0 the linearization is essentially two copies of (7.1), as follows. Regarding $f : C_0 \rightarrow Z_0$ as a pair of maps $(f_1, C_1; f_2, C_2) \in \mathcal{M}_s^V(X) \times_{\text{ev}} \mathcal{M}_s^V(Y)$, a variation ξ of f consists of smooth sections ξ_1 of $f_1^* TX$ and ξ_2 of $f_2^* TY$ with $\xi_1(x_k) = \xi_2(y_k)$ at each node $x_k = y_k$ which is mapped into V . Thus the domain of the linearization consists of sections $(\zeta_1, \bar{\xi}_1, h_1; \zeta_2, \bar{\xi}_2, h_2)$ with the matching condition $\bar{\xi}_1 = \bar{\xi}_2$ at the nodes mapped into V , and the linearization extends to a bounded linear map

$$(7.3) \quad \mathbf{D}_{f,C_0} : L_{1;s,0}(f_0^* TZ_0) \oplus T_q V^\ell \oplus T_{C_1} \widetilde{\mathcal{M}} \oplus T_{C_2} \widetilde{\mathcal{M}} \rightarrow L_s(\Lambda^{01}(f_0^* TZ_0)).$$

Again, as in Lemma 3.4, one can verify that for generic V -compatible (J, ν) in $\mathcal{J}(Z)$, Coker $\mathbf{D}_{f,C_0} = 0$ at irreducible maps (f, C_0) . In fact, for $f \in \mathcal{K}_\delta$ this can be done, as in Lemma 4.2 of [IP4], by a perturbation in (J, ν) supported

outside the $\delta/2$ neighborhood of V (maps in \mathcal{K}_δ , stabilized as in Observation 6.7 have no components mapped into the δ neighborhood of V), and hence by perturbations in the class $\mathcal{J}(Z)$ of Lemma 2.3. Similarly, the evaluation map $ev : \mathcal{M}_s^V \times \mathcal{M}_s^V \rightarrow V \times V$ is smooth and its image is transverse to the diagonal Δ for generic (J, ν) in $\mathcal{J}(Z)$. Therefore the irreducible part of the space $\mathcal{M}_s^V \times_{ev} \mathcal{M}_s^V = ev^{-1}(\Delta)$ is generically a smooth orbifold and its tangent space at an irreducible map (f, C_0) is identified with $\text{Ker } \mathbf{D}_{f, C_0}$.

We next turn to the linearization D_F at an approximate map $F = F_{f, C_0, \mu}$, which is defined by (7.2) with (f, C) replaced by the approximate map F of Definition 6.4. Our goal in the remainder of this section is to show that D_F extends to a Fredholm operator

$$(7.4) \quad \mathbf{D}_F : L_{1;s,0}(F^*TZ_\lambda) \oplus T_q V^\ell \oplus T_{C_\mu} \overline{\mathcal{M}}_{g,n} \rightarrow L_s(\Lambda^{01}(F^*TZ_\lambda)).$$

We can also consider the formal adjoint of D_F with respect to the weighted L^2 inner products, that is, the operator \mathbf{D}_F^* determined by the relation

$$\langle (\zeta, \bar{\xi}, h), \mathbf{D}_F^* \eta \rangle = \langle \mathbf{D}_F(\zeta, \bar{\xi}, h), \eta \rangle \quad \text{where} \quad \langle \zeta_1, \zeta_2 \rangle = \int_{C_\mu} \rho^{-\delta} \langle \zeta_1, \zeta_2 \rangle.$$

Using (7.2) and (1.11) one finds that

$$(7.5) \quad \mathbf{D}_F^* \eta = (D_F^* \eta, A\eta, -F_*^t J\eta)$$

$$\text{where} \quad \begin{cases} D_F^* \eta &= \rho^\delta L_F^*(\rho^{-\delta} \eta) \\ A\eta &= \int_{C_\mu} \rho^{-\delta} \sum_k (\bar{\partial} \beta_k) \eta^V + \beta_k \langle \nabla J df j, \eta^V \rangle. \end{cases}$$

Here F_*^t is the adjoint of dF , and L_F^* is the formal L^2 adjoint of the operator L_F of (1.11). We will see that this also extends to a Fredholm operator

$$(7.6) \quad \mathbf{D}_F^* : L_{1;s}(\Lambda^{01}(F^*TZ_\lambda)) \rightarrow L_{s,0}(F^*TZ_\lambda) \oplus T_q V^\ell \oplus T_{C_\mu} \overline{\mathcal{M}}_{g,n}.$$

The proof that D_F is bounded requires a preliminary lemma which is based on our assumption from Definition 2.2 that the second fundamental form of $V \subset Z$ vanishes.

LEMMA 7.2. *Let $\zeta = (\zeta^v, \zeta^x, \zeta^y)$ be a vector field on Z_λ in a coordinate chart (5.4). Then along the image of F , $\zeta^V = (\zeta^v, 0, 0)$ satisfies $|(\nabla \zeta^V)^N| \leq c\rho|\zeta^V|$ and $|L_F^N(\zeta^V)| \leq c\rho|\zeta|$.*

Proof. Each coordinate chart (5.4) is foliated by the submanifolds $V_{x,y}$ of points (v, x, y) with fixed x and y . Those submanifolds are deformations of $V = V_{0,0}$ so that the assumption of Definition 2.2 implies that the second fundamental form of $V_{x,y}$ is $O(R)$ where $R^2 = |x|^2 + |y|^2$. In particular, the $V_{x,y}$ foliate $Z_\lambda = \{xy = \lambda\}$ and the second fundamental form of $V_{x,y}$ in Z_λ is

also $O(R)$, which means that $|(\nabla \zeta^V)^N| \leq cR|\zeta^V|$. But we know that $R \leq c\rho^s$ from Lemma 6.8, and so $|(\nabla \zeta^V)^N| \leq c\rho|\zeta^V|$.

Expanding L_F by (1.11), we then see that the first two terms of $L_F^N \zeta^V$ are dominated by $c\rho|\zeta^V|$. Along V , the remaining terms of L_F can be expressed in terms of the second fundamental form h : since $\nu^N = 0$ along V the term $(\nabla_{\zeta^V} \nu)^N$ is $h(\zeta^V, \nu)$, and since J preserves the normal direction to V along V , the term $(J\nabla_{\zeta^V} J(X))^N = J[(\nabla_{\zeta^V} (JX))^N - J(\nabla_{\zeta^V} X)^N]$ is $Jh(\zeta^V, JX) + h(\zeta^V, X)$. These terms vanish along V by our assumption $h_V = 0$, and hence along $V_{x,y}$ they are bounded by $cR|\zeta^V|$. We conclude that $|L_F^N(\zeta^V)| \leq c\rho|\zeta|$ after again noting that $R \leq c\rho^s$. \square

PROPOSITION 7.3. *For approximate maps $F = F_{f,C_0,\mu}$ with $(f, C_0) \in \mathcal{K}_\delta$ the operators (7.4) and (7.6) are bounded uniformly in λ . The index of \mathbf{D}_F is independent of s and*

$$\text{index}_{\mathbb{C}} \mathbf{D}_F = \frac{1}{2} \dim \mathcal{M}_s^V(X) \times_{\text{ev}} \mathcal{M}_s^V(Y).$$

Proof. By (7.2), (6.8), and (1.11), $\mathbf{D}_{F,C}(\zeta, \bar{\xi}, h)$ is $L_F(\xi) + \frac{1}{2}J_F dF h$ with $\xi = \zeta + \sum \beta_k \bar{\xi}_k$. In the formula (1.11) for L_F , the one-forms ν and $\nabla \nu$ on the domain C_μ are bounded in the metric on the universal curve, so are bounded by $c\rho$ in the cylindrical metric (4.1); hence

$$(7.7) \quad |L_F \xi| \leq c(|\nabla \xi| + B|\xi|) \quad \text{with} \quad B = |dF| + |\Phi_F| + c\rho.$$

To proceed, write C_μ as the union of $C_\mu \setminus A$ and A exactly as in the proof of Lemma 6.9. On $C_\mu \setminus A$ there we have the bound $|dF| + |\Phi_F| \leq c$ as in (6.12). On each A_k , $|dF| \leq c\rho$ from after (6.15), and $|\Phi_F| \leq c\rho$ by (6.17). Thus $B \leq c\rho$ globally on C_μ . With that, (7.7) yields the pointwise inequality:

$$(7.8) \quad |\mathbf{D}_{F,C}(\zeta, \bar{\xi}, h)| \leq c \left(|\nabla \zeta| + \rho|\zeta| + \sum_k |L_F(\beta_k \bar{\xi}_k)| + |J_F dF h| \right).$$

Next observe that $|d\beta_k| \leq c\rho$ by the bound (5.11) and the fact, from the last sentence of Section 6, that $d\beta$ has support where $\rho_0/2 \leq \rho \leq \rho_0$. Consequently, $|L(\beta_k \bar{\xi}_k)| = |\bar{\partial} \beta_k \bar{\xi}_k + \beta_k L \bar{\xi}_k|$ is less than $c\rho|\bar{\xi}_k| + |L^V \bar{\xi}_k| + |L^N \bar{\xi}_k|$ which, in turn, is less than $c(|\nabla \bar{\xi}_k| + \rho|\bar{\xi}_k|)$ after one uses Lemma 7.2 and the estimate (7.7). Here $\nabla \bar{\xi}_k$ is the covariant derivative of Z_λ , which differs from covariant derivative $\tilde{\nabla}$ of the Riemannian metric on A_k . However, as in the proof of Lemma 7.2, $\bar{\xi}_k$ is tangent to the submanifold $V_{x,y}$ whose second fundamental form is $O(R)$, so that $|(\tilde{\nabla} - \nabla)\bar{\xi}_k| \leq cR|\bar{\xi}_k| \leq c\rho^s|\bar{\xi}_k|$. Furthermore, because $\bar{\xi}_k$ is a constant vector field in the coordinates (5.4) we know that $|\tilde{\nabla} \bar{\xi}_k| \leq \Gamma_k|\bar{\xi}_k|$ where Γ_k is a bound for the Christoffel symbols; in fact $|\Gamma_k(v, x, y)| \leq c|(v, x, y)|$ by the construction of those coordinates (normal coordinates along V extended off V by the exponential map). In particular, at

an image point $F(z)$, Γ_k is bounded by $c|F| \leq c\rho$ because $|F - f| \leq c\rho$ by (6.15) and $|f| \leq c\rho$ by (5.5) and Lemma 6.8d. Thus

$$(7.9) \quad |L_F(\beta_k \bar{\xi}_k)| \leq c\rho |\bar{\xi}|.$$

The quantity $J_F dFh$ in (7.8) expands to $(J_F - J_f)dFh + J_f(dF - df)h + J_f dfh$. By writing $f = (h^v, \tilde{x}, 0)$ in the notation of Definition 6.4 and applying 6.8d, one sees that $|J_f dfh| \leq c\rho|h|$. Using (6.15) we then have

$$|J_F dFh| \leq c\rho|h|.$$

With these bounds, the right-hand side of (7.8) simplifies to $|\nabla\zeta| + \rho|\zeta| + \rho|\bar{\xi}| + \rho|h|$. After integrating as in (6.9), comparing with (6.10), and using (5.10) we obtain

$$(7.10) \quad \|\mathbf{D}_F(\zeta, \bar{\xi}, h)\|_0 \leq c \left(\|\zeta\|_1 + |\bar{\xi}| + \|\rho^{1-\delta/2}h\|_{L^2} + \|\rho^{1-\delta/2}h\|_{L^4} \right).$$

Finally, to bound the h terms, consider vectors $h = (h_0, h_1)$ at a curve C in a fixed chart $U \times D^\ell$ as in Definition 4.3. After lifting to the normalization of C at the nodes of $f^{-1}(V)$, the components $h_0 \in T\mathcal{N}_\ell$ satisfy elliptic equations which are uniform for C in U . Standard elliptic theory then implies that $\sup|h_0| \leq c\|h_0\|_{WP}$ with a constant c uniform on the chart. Adding the similar bound on $\sup|h_1|$ after (4.6) and comparing with (4.8), we see that $\sup|h| \leq c\|h\|$. Thus the last two terms of (7.10) are bounded by

$$(7.11) \quad \int_A |\rho h|^p \rho^{-\delta p/2} \leq c \sup|h|^p \int_A \rho^{p(1-\delta/2)} \leq c \|h\|^p,$$

with c uniform on $f \in \mathcal{K}_\delta$ by Observation 6.7.

The right-hand side of (7.10) is now the norm $\|(\zeta, \bar{\xi}, h)\|_1$ of (6.10). We conclude that \mathbf{D}_F is bounded uniformly in $\lambda \neq 0$ and in $f \in \mathcal{K}_\delta$. The proof for \mathbf{D}_F^* is similar, and the index is given by Lemma 3.4. \square

8. The eigenvalue estimate

We now come to the key analysis step: obtaining estimates on the linearization D_F of the holomorphic map equation along the space of approximate maps. We establish a lower bound for the eigenvalues of $D_F D_F^*$ and construct a right inverse P_F for D_F . In the next section P_F will be used to correct approximate maps to true holomorphic maps.

To get uniform estimates we fix (J, ν) generic in the sense of Lemma 3.4. We continue to work with δ -flat maps, which we call δ_0 -flat in this section to avoid confusion with the exponential weight δ of the norm (6.9), which will also appear. As in (3.11) this δ_0 defines a compact set $\mathcal{K}_{\delta_0} \subset \mathcal{M}_s^V(X) \times_{\text{ev}} \mathcal{M}_s^V(Y)$ of (3.11) and corresponding subsets

$$(8.1) \quad \text{Model}_s^{\delta_0}(Z_\lambda) \subset \text{Model}_s(Z_\lambda) \quad \text{and} \quad \text{Approx}_s^{\delta_0}(Z_\lambda) \subset \text{Approx}_s(Z_\lambda)$$

of the model space and the space of approximate maps. Thus $\text{Model}_s^{\delta_0}(Z_\lambda)$ is the inverse image of \mathcal{K}_{δ_0} under the covering map of Definition 6.1 and $\text{Approx}_s^{\delta_0}(Z_\lambda)$ is the image of $\text{Model}_s^{\delta_0}(Z_\lambda)$ under the gluing map (6.5). For the maps (f_1, f_2) in \mathcal{K}_{δ_0} Lemma 6.8c implies that $|\lambda|$, defined by (6.3), is uniformly equivalent to $|\mu_k|^{s_k}$ for each k .

In Sections 8 and 9 we also use the notation

$$(8.2) \quad \text{Model}_s^{\delta_0}(Z_\lambda)^* \subset \mathcal{M}\text{odel}_s^{\delta_0}(Z_\lambda) \quad \text{and} \quad \mathcal{A}\text{pprox}_s^{\delta_0}(Z_\lambda)^* \subset \mathcal{A}\text{pprox}_s^{\delta_0}(Z_\lambda)$$

for the subsets constructed from the maps in $\mathcal{M}_s^V(X) \times_{\text{ev}} \mathcal{M}_s^V(Y)$ which are irreducible as defined after equation (1.4). Recall from Lemma 3.4 that $\text{coker } D_s = 0$ at irreducible maps and hence the spaces (8.2) are orbifolds. That understood, the aim of this section is to prove the following analytic result.

PROPOSITION 8.1. *For each generic $(J, \nu) \in \mathcal{J}(Z)$, there is a constant $E > 0$ independent of λ such that the linearization \mathbf{D}_F at an approximate map $(F, C_\mu) = F_{f, C_0, \mu} \in \mathcal{A}\text{pprox}_s^{\delta_0}(Z_\lambda)^*$ has a right inverse*

$$P_F : L_s(\Lambda^{01}(F^*TZ_\lambda)) \rightarrow L_{1,s}(F^*TZ_\lambda) \oplus T_{C_\mu} \overline{\mathcal{M}}_{g,n}$$

such that

$$(8.3) \quad E^{-1} \|\eta\|_0 \leq \|P_F \eta\|_1 \leq E \|\eta\|_0$$

Proof. Lemma 8.5 below shows that $\mathbf{D}_F \mathbf{D}_F^*$ is uniformly invertible. Therefore $P_F = \mathbf{D}_F^* (\mathbf{D}_F \mathbf{D}_F^*)^{-1}$ is a right inverse for \mathbf{D}_F . Since \mathbf{D}_F and \mathbf{D}_F^* are bounded by Proposition 7.3 we have $\|\eta\|_0 = \|\mathbf{D}_F P_F \eta\|_0 \leq E \|P_F \eta\|_1$ and $\|P_F \eta\|_1 \leq c \|(\mathbf{D}_F \mathbf{D}_F^*)^{-1} \eta\|_2 \leq E \|\eta\|_0$. \square

For small λ the domain of an approximate map $F = F_{f, C_0, \mu} : C_\mu \rightarrow Z_\lambda$ has a neck $N_k = B_k(1) \cap C_\mu$ around each node of C_0 . These necks are isometric to cylinders using the metric and coordinates of (4.5). The following lemma and its corollary give *a priori* estimates for the formal L^2 adjoint L_F^* of L_F on the necks $N_k(\varepsilon) = B_k(\varepsilon) \cap C_\mu$.

PROPOSITION 8.2. *For $\delta > 0$ small there are constants ε_0 and c such that for all λ sufficiently small, all approximate maps $F \in \mathcal{A}\text{pprox}_s^{\delta_0}(Z_\lambda)^*$ and each neck $N = N_k(\varepsilon)$ with $\varepsilon \leq \varepsilon_0$, each $\eta \in \Omega^{01}(F^*TZ_\lambda)$ satisfies*

$$(8.4) \quad \int_N \rho^\delta (|\nabla \eta|^2 + |\eta|^2) \leq c \int_N \rho^\delta |L_F^* \eta|^2 + c \int_{\partial N} \rho^\delta (|\nabla \eta|^2 + |\eta|^2).$$

Proof. Write ρ^δ as the derivative of $\psi(t) = \int_0^t \rho^\delta(\tau) d\tau$ and integrate by parts:

$$\int_N \rho^\delta |\eta|^2 = \int_N \psi' |\eta|^2 dt d\theta \leq \int_N |\psi| \cdot 2\langle \eta, \nabla \eta \rangle + \int_{\partial N} |\psi| \cdot |\eta|^2.$$

Because $\rho^2 = 2|\mu| \cosh(2t)$ satisfies $\rho^2 \leq 2|\mu|e^{2t} \leq 2\rho^2$, we have $|\psi| \leq c\rho^\delta/\delta$, so that $|\psi| \cdot 2\langle \eta, \nabla \eta \rangle$ is bounded by $\frac{1}{2}\rho^\delta |\eta|^2 + c_\delta \delta^{-2} \rho^\delta |\nabla \eta|^2$. Rearranging gives

$$(8.5) \quad \int_N \rho^\delta |\eta|^2 \leq C_\delta \int_N \rho^\delta |\nabla \eta|^2 + c_\delta \int_{\partial N} \rho^\delta |\eta|^2.$$

Now on the cylinder N every $(0,1)$ -form η can be written $\eta = \eta_1 dt - (J\eta_1) d\theta$ where η_1 is a section of F^*TZ_λ . Taking the formal L^2 adjoint of (1.11) and using the bounds $|dF| \leq c\rho$, $|\nu| + |(\nabla \nu)^*| \leq c\rho$ from the proof of Proposition 7.3 and the obvious bound $|(\nabla J)^*| < c$, one sees that

$$L_F^* \eta = -\nabla_t \eta_1 + J\nabla_\theta \eta_1 + O(\rho|\eta|).$$

Therefore

$$\begin{aligned} c(|\rho\eta|^2 + \rho|\eta||\nabla\eta|) + |L_F^* \eta|^2 &\geq |\nabla_t \eta_1|^2 + |J\nabla_\theta \eta_1|^2 - 2\langle \nabla_t \eta_1, J\nabla_\theta \eta_1 \rangle \\ &= \frac{1}{2}|\nabla \eta|^2 - 2\langle \nabla_t \eta_1, J\nabla_\theta \eta_1 \rangle. \end{aligned}$$

Differentiating the 1-form $\omega = \langle \eta_1, J\nabla_t \eta_1 \rangle dt + \langle \eta_1, J\nabla_\theta \eta_1 \rangle d\theta$ and moving J past ∇ we also have

$$\begin{aligned} *d\omega &= (2\langle \nabla_t \eta_1, J\nabla_\theta \eta_1 \rangle + \langle \eta_1, \nabla_t (J\nabla_\theta \eta_1) - \nabla_\theta (J\nabla_t \eta_1) \rangle) \\ &\geq 2\langle \nabla_t \eta_1, J\nabla_\theta \eta_1 \rangle + \langle \eta_1, J\mathcal{R}(\partial_t, \partial_\theta)\eta_1 \rangle - c|\nabla J||dF||\eta||\nabla\eta| \\ &\geq 2\langle \nabla_t \eta_1, J\nabla_\theta \eta_1 \rangle + \langle \eta_1, J\mathcal{R}(\partial_t, \partial_\theta)\eta_1 \rangle - c\rho|\eta||\nabla\eta| \end{aligned}$$

where \mathcal{R} is the curvature of ∇ . Combining the last two displayed equations, multiplying by ρ^δ , integrating by parts, and using the bound $2\rho|\eta||\nabla\eta| \leq \rho|\nabla\eta|^2 + \rho|\eta|^2$ we obtain

$$(8.6) \quad \begin{aligned} \frac{1}{2} \int_N \rho^\delta |\nabla \eta|^2 &\leq \int_N \rho^\delta [|L_F^* \eta|^2 + \langle \mathcal{R}(\partial_t, \partial_\theta)\eta_1, J\eta_1 \rangle] \\ &\quad - d(\rho^\delta) \wedge \omega + \int_{\partial N} \rho^\delta \omega + c \int_N \rho^{\delta+1} (|\nabla \eta|^2 + |\eta|^2). \end{aligned}$$

Because the domain metric is flat, \mathcal{R} is the (pulled back) curvature of Z_λ . Setting $U = F_* \partial_t$ and $V = F_* \partial_\theta$ and applying the Gauss equations we obtain

$$(8.7) \quad \begin{aligned} \langle \mathcal{R}(U, V)\eta_1, J\eta_1 \rangle &= \langle R^Z(U, V)\eta_1, J\eta_1 \rangle \\ &\quad - \langle h(\eta_1, V), h(J\eta_1, U) \rangle + \langle h(J\eta_1, V), h(\eta_1, U) \rangle \end{aligned}$$

where R^Z is the curvature of Z and h is the second fundamental form of $Z_\lambda \subset Z$. Since R^Z is bounded, the term containing it is dominated by

$c|dF|^2|\eta|^2 \leq c\rho^2|\eta|^2$. Likewise, the necks around any nodes not mapped into V are uniformly bounded away from V by Lemma 6.8a. For those nodes, h is bounded along the image $F(N_k)$. Consequently, the entire right-hand side of (8.7) is bounded by $c|\eta|^2|dF| \leq c\rho|\eta|^2$.

For nodes mapped into V , $F(N_k)$ lies in one of the sets A_k used in the proof of Lemma 6.9. In the coordinates (v, x, y) on A_k defined by (5.4) Z_λ is the set $xy = \lambda$ and its second fundamental form is $h(X, Y) = \langle \nabla_X Y, \nu \rangle$. We can compare h with the second fundamental form h_0 of the complex curve $xy = \lambda$ in \mathbb{C}^2 . To calculate h_0 , note that at each point (x, y) with $xy = \lambda$ the tangent space T and the normal space N are, respectively, the complex span of the unit tangent vector $\tau = (x/R, -y/R)$ and the unit normal vector $\nu = (\bar{y}/R, \bar{x}/R)$ where $R^2 = x^2 + y^2$. Then h_0 is the symmetric complex bilinear map $T \otimes_{\mathbb{C}} T \rightarrow N$ given by $h_0(\tau, \tau) = \langle \tau \cdot \tau, \nu \rangle \nu = 2\lambda R^{-3} \nu$.

Because the Christoffel symbols of the coordinate system (v, x, y) are bounded, the 2-tensor $h - h_0$ on $Z_\lambda \cap A_k$ is bounded uniformly in λ . Examining the proof Lemma 6.9, we see that the normal component F^N of F — see (6.14) — satisfies $|F^N| + |dF^N| \leq c\rho^{s_k} \leq cR$ on A_k . Because ν^N vanishes along V by condition (1.15a) we have $|\nu^N| \leq c\rho|F^N|$ as in the sentence preceding (6.17); that gives $|\nu^N| \leq c\rho R$ and together with (6.17) shows that $U - JV = (dF + JdFj)(\partial_\theta) = 2(\nu + \Phi_F)$ satisfies $|(U - JV)^N| \leq c\rho R$. It follows that $|(h - h_0)(F_*v, \cdot)| \leq c\rho|v|$ for any v and that $|h_0(F_*v, \cdot)| \leq c|xy|R^{-3}|dF^N||v| \leq c|v|$. We can therefore replace h by h_0 in (8.7), and then replace V by JU , each time making an error of at most $c\rho|\eta|^2$. The result is

$$\begin{aligned} \langle \mathcal{R}(U, V)\eta_1, J\eta_1 \rangle &\leq -\langle h_0(\eta_1, JU), h_0(J\eta_1, U) \rangle \\ &\quad + \langle h_0(J\eta_1, JU), h_0(\eta_1, U) \rangle + c\rho|\eta|^2. \end{aligned}$$

But h_0 is complex linear and J preserves the normal direction, so this reduces to

$$(8.8) \quad \langle \mathcal{R}(\partial_t, \partial_\theta)\eta_1, J\eta_1 \rangle \leq -2|h_0(\eta_1, U)|^2 + c\rho|\eta|^2 \leq c\rho|\eta|^2;$$

that is, the sign of the curvature is compatible with our inequalities.

It remains to bound the ω term in (8.6). As in (4.5) we can introduce cylindrical coordinates (v^i, τ, Θ) on the necks of Z_λ by choosing normal coordinates $\{v^i\}$ on V centered at $f(x_k)$, extending by parallel translation into a neighborhood of V , and writing $x = \sqrt{|\lambda|} \exp(\tau + i\theta)$. Then the metric on Z_λ is $g_\lambda = R^2(d\tau^2 + d\Theta^2) + g^V$ where $R^2 = |x|^2 + |y|^2 = 2|\lambda| \cosh(2\tau)$ and g^V is the metric of V . Direct computations with the formula (6.4) for F show that in these coordinates (i) $F_*\partial_\theta = s_k\partial_\Theta + O(\rho)$ and (ii) the Christoffel symbols are all bounded and those in the Θ direction are

$$\Gamma_{\Theta\Theta}^\Theta = \Gamma_{\tau\tau}^\Theta = 0, \quad \Gamma_{\Theta\tau}^\Theta = -\Gamma_{\tau\Theta}^\Theta = \tanh(2\tau).$$

Thus $\nabla_\theta = \partial_\theta + \tanh(2\tau)J + O(\rho)$. Recalling that $\rho^2 = 2|\mu| \cosh(2t)$ we also have

$$(8.9) \quad -d(\rho^\delta) \wedge \omega = -\partial_t \rho^\delta \langle \eta_1, J \nabla_\theta \eta_1 \rangle dt \wedge d\theta = \delta \rho^\delta \tanh(2t) \langle J \eta_1, \nabla_\theta \eta_1 \rangle dt \wedge d\theta.$$

Because F^*g_λ is independent of θ in these coordinates, the method of (5.9) and the inequality $|\tanh(2\tau)| \leq 1$ gives the bound

$$-\tanh(2\tau) \int_{S^1} \langle J \eta_1, \partial_\theta \eta_1 \rangle d\theta \leq \int_{S^1} |\partial_\theta \eta_1|^2 d\theta.$$

Replacing ∂_θ by $\nabla_\theta - \tanh(2\tau)J + O(\rho)$ on the right-hand side once gives

$$0 \leq \int_{S^1} \langle \nabla_\theta \eta_1, \partial_\theta \eta_1 \rangle + c\rho \int_{S^1} |\nabla \eta|^2 + |\eta|^2$$

and again gives

$$\tanh(2\tau) \int_{S^1} \langle J \eta_1, \nabla_\theta \eta_1 \rangle d\theta \leq \int_{S^1} |\nabla_\theta \eta_1|^2 + c\rho \int_{S^1} |\nabla \eta|^2 + |\eta|^2.$$

But $\tanh(2\tau) = \tanh(2s_k t) + O(\rho)$ and $0 \leq \tanh(2t)/\tanh(2s_k t) \leq 1$ so that this last inequality and (8.9) yield

$$-\int_N d(\rho^\delta) \wedge \omega \leq \delta \int_N \rho^\delta |\nabla \eta|^2 + c \int_N \rho^{1+\delta} (|\nabla \eta|^2 + |\eta|^2).$$

Insert this and (8.8) into (8.6) and multiply by the constant C_δ of (8.5). Adding (8.5) and noting that $|\omega| \leq |\eta|^2 + |\nabla \eta|^2$ then gives (8.4) for small δ and ε . \square

Write $\nabla^\delta \eta = \rho^\delta \nabla(\rho^{-\delta} \eta)$ where ∇ is as usual the covariant derivative of the cylindrical metric on the domain and the metric induced on Z_λ from Z . Note that when $\delta > 0$ is small, the $L^{1,2}$ weighted norm (6.9) defined using ∇^δ is equivalent, uniformly in λ , to the one using ∇ . Then Proposition 8.2 implies the following:

COROLLARY 8.3. *For $\delta > 0$ small there are constants ε_0 and c such that for all λ sufficiently small and all approximate maps $F \in \mathcal{A}pprox_s^{\delta_0}(Z_\lambda)^*$ and each neck $N = N_k(\varepsilon)$ with $\varepsilon \leq \varepsilon_0$, each $\eta \in \Omega^{01}(F^*TZ_\lambda)$ satisfies*

$$(8.10) \quad \|\eta\|_{1,2,N} \leq c \|D_F^* \eta\|_{0,2,N} + c \|\eta\|_{1,2,\partial N};$$

that is,

$$\int_N \rho^{-\delta} (|\nabla^\delta \eta|^2 + |\eta|^2) \leq c \int_N \rho^{-\delta} |D_F^* \eta|^2 + c \int_{\partial N} \rho^{-\delta} (|\nabla^\delta \eta|^2 + |\eta|^2).$$

Proof. Replace η by $\rho^{-\delta} \eta$ in (8.4) and use (7.5) to replace L_F^* with D_F^* . \square

From now on we will fix the weight $\delta > 0$ in our norms (6.9) small and generic. We will extend the above estimates on the necks to global estimates, first for the weighted L^2 norms, then for the weighted $L^2 \cap L^4$ norms of Definition 6.5.

LEMMA 8.4. *For each generic $(J, \nu) \in \mathcal{J}(Z)$, there is an $E > 0$ such that for all λ sufficiently small and all $F \in \text{Approx}_s^{\delta_0}(Z_\lambda)^*$, the first eigenvalue of $\mathbf{D}_F \mathbf{D}_F^*$ is bounded below by E .*

Proof. Suppose the claim is false. Then there are sequences $\lambda_n \rightarrow 0$, maps $F_n : C_n = C_{\mu_n} \rightarrow Z_{\lambda_n}$ in $\text{Approx}_s^{\delta_0}(Z_\lambda)$ and $(0, 1)$ forms η_n along F_n with $\mathbf{D}_n \mathbf{D}_n^* \eta_n = \alpha_n \eta_n$ and $\alpha_n \rightarrow 0$. Taking the inner product with η_n and taking the adjoint (7.5) using our norms (6.9) we obtain

$$(8.11) \quad \alpha_n \|\eta_n\|_{0,2,s}^2 \geq \|D_n^* \eta_n\|_{0,2,s}^2$$

on each C_n . We may normalize the η_n so that $\|\eta_n\|_{1,2,s} = 1$. By the Bubble Tree Convergence Theorem there is a subsequence of the F_n that converges to a stable map f_0 from $C_0 = C_1 \cup C_2$ into Z_0 , and this convergence is in C^∞ away from the nodes of C_0 . On each compact set K in the complement of the nodes, the $L_s^{1,2}$ norm (6.9) in the cylindrical metric is uniformly equivalent to the usual $L^{1,2}$ norm. Standard elliptic theory implies that there is a subsequence of the η_n that converges in C^∞ on K to an $L_s^{1,2}$ section η with $\|\eta\|_{1,2,s} \leq 1$ and $\mathbf{D}_0^* \eta = 0$ along K (as in the sentence before Observation 6.7, convergence is C^∞ away from nodes). Doing this for the sequence $K_m = \rho^{-1}([\frac{1}{m}, \infty))$ which exhausts the complement of the nodes and then passing to a diagonal subsequence yields a limit η defined on $C_0 \setminus \{\text{nodes}\}$ with $L_s^{1,2}$ norm at most one and $\mathbf{D}_0^* \eta = 0$ weakly on $C_0 \setminus \{\text{nodes}\}$.

We next show that $\mathbf{D}_0^* \eta = 0$ weakly on all of C_0 ; that is, $\langle \mathbf{D}_0 X, \eta \rangle = 0$ for all $X = (\zeta, \bar{\xi}, h)$ where $\zeta \in L_s^{1,2}(f_0^* T Z_0)$, $\bar{\xi} \in T_q V$ and $h \in T_{C_1} \mathcal{M} \times T_{C_2} \mathcal{M}$. Given X and $\varepsilon > 0$, we can then choose $\delta < \varepsilon$ so that the norm of X on the region $N = \rho \leq \delta$ satisfies $\|(\zeta, 0, h)\|_{1,N} \leq \varepsilon$. Then fix a cutoff function β as in (5.11) supported on N with $|d\beta| \leq 1$ and write $X = \beta X + X_1$. Since $X_1 = (1 - \beta)X$ has support on the outside region where $\rho \geq \delta/2$ we have $\langle \mathbf{D}_0 X_1, \eta \rangle = \langle X_1, \mathbf{D}_0^* \eta \rangle = 0$. On the other hand, βX is $(\beta \zeta, \bar{\xi}, \beta h)$ because (6.7) is unaffected by the cutoff function. Integrating (7.9) over $\rho \leq \delta$ using (5.10) gives $\|\mathbf{D}_0(0, \bar{\xi}, 0)\|_{0,2,s} \leq c\delta|\bar{\xi}| \leq c\varepsilon\|X\|_1$, and it is easy to see from (6.10) and (5.11) that $\|(\beta \zeta, 0, \beta h)\|_1 \leq 2\|(\zeta, 0, h)\|_{1,N} \leq 2\varepsilon$. With those facts, Hölder's inequality, the bound $\|\eta\|_{0,2,s} \leq 1$, and Proposition 7.3 give

$$|\langle \mathbf{D}_0(\beta X), \eta \rangle| \leq \|\mathbf{D}_0(\beta \zeta, 0, \beta h) + \mathbf{D}_0(0, \bar{\xi}, 0)\|_{0,2,s} \cdot \|\eta\|_{0,2,s} \leq c\varepsilon\|X\|_1.$$

We conclude that $|\langle \mathbf{D}_0 X, \eta \rangle| = 0$, so that $\mathbf{D}_0^* \eta = 0$ weakly. Pairing against $\mathbf{D}_0^* \eta$, which lies in the image of the bounded map (7.6), then shows that $\langle \mathbf{D}_0^* \eta, \mathbf{D}_0^* \eta \rangle = 0$, and therefore $\mathbf{D}_0^* \eta = 0$.

Now for generic (J, ν) , as observed after equation (7.3), we have $\mathbf{D}_0 = 0$. Thus our solution of $\mathbf{D}_0^* \eta = 0$ must be $\eta = 0$. Consequently, for each small fixed δ_0 we have $\eta_n \rightarrow 0$ in C^∞ on the complement of $N = \bigcup_k N_k(\delta_0)$ and in particular on ∂N . Then $\|\eta_n\|_{1,2,C_n \setminus N} \rightarrow 0$ and (8.10) bounds $\|\eta_n\|_{1,2,N} - c\|\eta_n\|_{1,2,\partial N}$, so that (8.11) and the normalization $\|\eta_n\|_{1,2,C_n} = 1$ imply the inequalities

$$\begin{aligned} \frac{1}{2} &\leq \|\eta_n\|_{1,2,N} - \frac{1}{4} \leq C \|D_n^* \eta_n\|_{0,2,N} \\ &\leq C \|D_n^* \eta_n\|_{0,2,C_n} \leq C \alpha_n \|\eta_n\|_{0,2,C_n} \leq C \alpha_n \end{aligned}$$

for large n . That contradicts the assumption that $\alpha_n \rightarrow 0$, completing the proof. \square

LEMMA 8.5. *There is a constant C such that for all λ sufficiently small and all $F \in \text{Approx}_s^{\delta_0}(Z_\lambda)^*$, each $\eta \in \Omega^{01}(F^*TZ_\lambda)$ satisfies*

$$\|\eta\|_2 \leq C \|\mathbf{D}_F \mathbf{D}_F^* \eta\|_0$$

in the norms of Definition 6.5.

Proof. Cover the domain C_μ of F by disks of radius 1 in the cylindrical metric so that each point lies in at most 10 disks. Since ρ varies by a bounded factor across each unit interval in the neck, we can apply the basic elliptic estimate on each disk, multiply by $\rho^{-\delta/2}$ and sum to get

$$\|\eta\|_{2,p,s} \leq c (\|\mathbf{D}_F \mathbf{D}_F^* \eta\|_{0,p,s} + \|\eta\|_{0,2,s})$$

for a constant $c = c(p)$ independent of λ . Combining the $p = 4$ and $p = 2$ inequalities we get

$$\|\eta\|_2 \leq c (\|\mathbf{D}_F \mathbf{D}_F^* \eta\|_0 + \|\eta\|_{0,2,s}).$$

Using Lemma 8.4 and applying Holder's inequality for the weighted L^2 norm, we obtain

$$\begin{aligned} c \|\eta\|_{0,2,s}^2 &\leq \|\mathbf{D}_F^* \eta\|_{0,2,s}^2 = \langle \eta, \mathbf{D}_F \mathbf{D}_F^* \eta \rangle_{0,2,s} \\ &\leq \|\eta\|_{0,2,s} \|\mathbf{D}_F \mathbf{D}_F^* \eta\|_{0,2,s} \leq \|\eta\|_2 \|\mathbf{D}_F \mathbf{D}_F^* \eta\|_0 \end{aligned}$$

which combined with the previous inequality give the lemma. \square

9. The gluing diffeomorphism

Recall that the norm (6.10) induces a topology on the space $\text{Maps}_s(B, Z_\lambda \times \mathcal{U})$. Specifically, for C^0 close maps with the same label s in a chart as in Observation 6.7 we can write $(f', C') = \exp_{(f,C)}(\xi, h)$ and set

$$(9.1) \quad \text{dist}((f, C), (f', C')) = \|\xi, h\|_1.$$

This defines a distance (the inf of the lengths over all paths piecewise of the above type) and hence a topology on $\text{Maps}(B, Z_\lambda \times \mathcal{U})$. Using this distance, we will show that the moduli space of stable maps into Z_λ is close to the space of approximate maps, and that these spaces are in fact isotopic.

As in Observation 6.7 we work in a chart in which all domains have been stabilized. In that chart we use coordinates normal to the stratum \mathcal{N}_ℓ of ℓ -nodal curves in the Deligne-Mumford space as in Section 4. Thus a nearly nodal curve $C = (B, j) \in \overline{\mathcal{M}}_{g,n}$ is written as $(C_0, \mu) = C_0(\mu)$ where $C_0 \in \mathcal{N}_\ell$ and $\mu = (\mu_1, \dots, \mu_\ell)$ are coordinates in the normal direction to \mathcal{N}_ℓ .

We start by describing a parametrization for a neighborhood of the space of approximate maps $\text{Approx}_s(Z_\lambda)$ of (8.1). Consider the Banach space bundle Λ^{01} over $\text{Model}_s(Z_\lambda)^*$ whose fiber at an approximate map F is the completion of $\Gamma(\Lambda^{01}(F^*TZ_\lambda))$ in the weighted norm $\|\cdot\|_0$ of (6.11). We write elements of Λ^{01} as quadruples (f, C_0, μ, η) with $(f, C_0) \in \mathcal{M}_s^V(X) \times_{\text{ev}} \mathcal{M}_s^V(Y)$, μ determined from λ by (6.3), and $\eta \in \Lambda^{01}(F^*TZ_\lambda)$ with $F = F_{(f, C_0, \mu)}$. The map

$$(9.2) \quad \Phi_\lambda : \Lambda^{01}(\varepsilon) \rightarrow \text{Maps}_s(B, Z_\lambda \times \mathcal{U}), \quad \Phi_\lambda(f, C_0, \mu, \eta) = \exp_{F, C_\mu}(P_F \eta)$$

(with P_F as in Proposition 8.1) is defined on an ε neighborhood of the zero section of Λ^{01} and agrees with the gluing map Γ_λ along the zero section. The following proposition shows that Φ_λ coordinatizes a neighborhood of $\text{Approx}_s^\delta(Z_\lambda)^*$.

PROPOSITION 9.1. *There is a constant $c > 0$ so that, for all small λ , Φ_λ is a diffeomorphism from an ε -neighborhood of the zero section in Λ^{01} onto a neighborhood of $\text{Approx}_s^\delta(Z_\lambda)^*$ in $\text{Maps}_s(B, Z_\lambda \times \mathcal{U})$ that contains at least a $c\varepsilon$ neighborhood of $\text{Approx}_s^\delta(Z_\lambda)^*$.*

Proof. By Lemma 3.4 the tangent space of $\mathcal{A}_s = \text{Approx}_s^\delta(Z_\lambda)^*$ at $F = F_{f, C_0, \mu} : C_\mu \rightarrow Z_\lambda$ has the same dimension as $\text{Ker } \mathbf{D}_F = (\text{Im } P_F)^\perp$. In fact,

$$(9.3) \quad T_F \Lambda^{01} = T_F \mathcal{A}_s \oplus \text{Im } P_F$$

because any $P_F \eta$ which lies in $T_F \mathcal{A}_s$ satisfies, by (8.3), Lemma 5.3 and Lemma 9.2 below,

$$\|P_F \eta\|_1 \leq E \|\eta\|_0 = E \|\mathbf{D}_F P_F \eta\|_0 \leq CE |\lambda|^{1/(8|s|)} \|P_F \eta\|_1,$$

so that for small λ , $P_F \eta$ is zero.

Next fix a path (f_t, C_t, μ_t) in $\text{Model}_s(Z_\lambda)$ starting at (f_0, C_0, μ_0) and let $(\xi, h) \in T\mathcal{A}_s$ be the tangent vector at $t = 0$ of the corresponding path of approximate maps $F_t = \Phi_\lambda(f_t, C_t, \mu_t, 0)$. Each element τ in the fiber of Λ^{01} over (f_0, C_0, μ_0) determines a vector field $P_{F_0} \tau$ along the image of F_0 in TZ_λ . After extending τ along F_t by parallel translation we calculate

$$d\Phi_\lambda \Big|_{(f_0, C_0, \mu_0, \eta_0)}(\xi, h, \tau) = \frac{d}{dt} \exp_{F_t}(t P_{F_t} \tau) \Big|_{t=0} = \xi + P_{F_0} \tau.$$

Thus $d\Phi_\lambda$ is an isomorphism by (9.3). Consequently, Φ_λ is a local diffeomorphism near the zero section of Λ^{01} .

To show injectivity, we suppose that injectivity fails on the disk bundle $\Lambda^{01}(\varepsilon) = \{\eta \in \Lambda^{01} : \|\eta\|_0 \leq \varepsilon\}$ for every ε . Then for each n there exist elements $(f_n, C_{0,n}, \mu_n, \eta_n) \neq (f'_n, C'_{0,n}, \mu'_n, \eta'_n)$ in $\Lambda^{01}(1/n)$ which have the same image under Φ_λ . After passing to subsequences, we can assume that the $\{(f_n, C_{0,n}, \mu_n)\}$ and $\{(f'_n, C'_{0,n}, \mu'_n)\}$ converge in the stable map topology to limits $f : C_0 \rightarrow Z_0$ and $f' : C'_0 \rightarrow Z_0$ in $\mathcal{K}_\delta \subset \mathcal{M}_s^V \times_{\text{ev}} \mathcal{M}_s^V$. Furthermore, convergence in the metric (4.10) implies convergence $C_{0,n} \rightarrow C$ and $C'_{0,n} \rightarrow C'$ to elements on the boundary of the cylindrical end compactification of $\mathcal{M}_{g,n}$ defined at the end of Section 4. These limits C and C' consist of the nodal curves C_0 and C'_0 together with, for each, an element of the real torus T^ℓ .

Now fix a compact region R in B which contains no nodes. Then for small λ we have $F_n \rightarrow f$ and $F'_n \rightarrow f'$ in C^1 on R . Since our $\|\cdot\|_1$ norm dominates the C^0 norm on maps, the triangle inequality gives

$$\begin{aligned} & \lim_{n \rightarrow \infty} \text{dist} (C_{0,n}(\mu_n), C'_{0,n}(\mu'_n)) + \sup_{x \in R} \text{dist} (f(x), f'(x)) \\ & \leq c \lim_{n \rightarrow \infty} (\|P_{F_n} \eta_n\|_1 + \|P_{F'_n} \eta'_n\|_1) \leq c \lim_{n \rightarrow \infty} (\|\eta_n\|_0 + \|\eta'_n\|_0) = 0 \end{aligned}$$

by Proposition 8.1. The convergence to zero of the first term implies convergence in the topology of the cylindrical end compactification. Thus (i) $C = C'$, and (ii) f and f' agree on R and therefore, as in the argument after (6.5), agree everywhere. Consequently, for large n $(f_n, C_{0,n}, \mu_n, \eta_n)$ and $(f'_n, C'_{0,n}, \mu'_n, \eta'_n)$ lie in the region where Φ_λ is a local diffeomorphism and are therefore equal. That establishes injectivity. Surjectivity onto a $c\varepsilon$ -neighborhood follows from the first inequality in (8.3). \square

LEMMA 9.2. *There are constants C and λ_0 such that for $0 < |\lambda| < \lambda_0$, the tangent vectors (ξ, h) to $\text{Approx}_s^\delta(Z_\lambda)$ at an approximate map $F = F_{f, C_0, \mu}$ satisfy*

$$(9.4) \quad \|\mathbf{D}_F(\xi, h)\|_0 \leq C|\lambda|^{1/(8|s|)} \|(\xi, h)\|_1.$$

Proof. Choose a path $F_t = F_{(f_t, C_{0t}, \mu_t)} \in \text{Approx}_s^\delta(Z_\lambda)$ starting at F with initial tangent vector (ξ, h) . Thus F_t are approximate maps constructed from the path $(f_t, C_{0t}) \in \mathcal{M}_s^V \times_{\text{ev}} \mathcal{M}_s^V$ starting at (f, C_0) with initial vector (ξ_0, h_0) . We can regard vector fields on C_μ as living on C_0 using the map $C_\mu \rightarrow C_0$ defined by $(z, w) \mapsto z$ for $|w| \leq |z|$ and $(z, w) \mapsto w$ for $|z| \leq |w|$ near each node and extended over C_μ as in Section 4. In the notation of Definition 4.3 the variation in the complex structure of C_μ decomposes as $h = (h_0, h_1)$ while the variation of C_0 is $h_0 = (h_0, 0)$. In particular, $h - h_0$ is the variation h_1 of Lemma 4.2, which is supported on $B = \bigcup_k B_k$, where B_k is the region of

C_0 near the k^{th} node where $\rho \leq 2|\mu_k|^{1/4}$. With that understood, we will separately estimate the left-hand side of (9.4) on $C_0 \setminus B$ and on each B_k .

On $C_0 \setminus B$ we can regard vector fields on Z_λ as living on Z_0 by the map $(v, x, y) \mapsto (v, x)$ for $|y| \leq |x|$, and $(v, x, y) \mapsto (v, y)$ for $|x| \leq |y|$, extended over Z_λ as in (2.6). With these identifications, D_F acts on sections of f^*TZ_0 along C_0 . Since $D_f(\xi_0, h_0) = 0$ and $D_F(\xi, h) = L_F(\xi) + \frac{1}{2}J_F dFh$ by (1.10), we have

$$(9.5) \quad D_F(\xi, h) = [L_F(\xi) - L_f(\xi_0)] \\ + [(J_F - J_f)dFh + J_f(dF - df)h + J_f df(h - h_0)]$$

with $h_0 = h$ as above. Under our identifications we also have $F_t = f_t$ in Definition 6.2, so taking the t -derivative shows that $\xi = \xi_0$. Using (6.15), we have

$$|D_F(\xi, h)| \leq |(L_F - L_f)\xi| + (|J_F - J_f| + |dF - df|)|h| \\ \leq |(L_F - L_f)(\xi)| + c|\lambda|\rho^{1-s}|h|.$$

But on $C_0 \setminus B$ the bounds in the proof of Lemma 6.9 (on both $C_\mu \setminus A$ and A) show that the C^1 distance from F_t to f_t is bounded by $c|\lambda|^{1/4}$. It follows that $(L_F - L_f)(\xi)$ is dominated by $c|\lambda|^{1/4}(|\nabla\xi| + |\xi|)$. After expanding ξ as in (6.8) and noting that $|\nabla(\beta_k \bar{\xi}_k)| \leq c|\bar{\xi}|$ by the estimates preceding (7.9), the above bound simplifies to

$$|D_F(\xi, h)| \leq c|\lambda|^{1/4} (|\nabla(\zeta)| + |\zeta| + |\bar{\xi}|) + c|\lambda|\rho^{1-s}|h|.$$

We can then integrate over $C_0 \setminus B$, bounding $\sup |h|$ as in (7.11) and integrating the power of ρ over the sets where $\rho \geq 2|\mu|^{1/4}$ as in (6.18). The result is

$$(9.6) \quad \|\mathbf{D}_F(\xi, h)\|_{0, C_0 \setminus B} \leq c|\lambda|^{1/4} \|(\xi, h)\|_1.$$

Now restrict attention to one B_k and again write ξ as $(\zeta, \bar{\xi})$. The estimates of Lemma 7.3 show that

$$(9.7) \quad |D_F(\xi, h)| \leq |\nabla\zeta| + c\rho|\zeta| + c\rho|\bar{\xi}| + c\rho|h|.$$

To proceed, we will explicitly find ζ in the region $B_+ \subset B_k$ where $|z| \geq |w|$ by differentiating formula (6.4a). That formula is given in the normal coordinate system (5.4) centered on $q_t = f_t(x_k)$. If we parallel translate the trivialization of $T_{q_0}Z$ along the path q_t , the corresponding coordinate systems are related by diffeomorphisms ϕ_t with $|\phi_t(v, x, y) - (v + q_t, x, y)| \leq ct|(v, x, y)|$ where c depends only on the local geometry of Z near V . Hence on the region where $\rho \leq |\mu_k|^{1/4}$, where $\beta = 1$ in (6.4a), we can take the t -derivative of $\phi_t^* F_t$ at $t = 0$ to obtain

$$\xi = (\bar{\xi}_0, \dot{a}z^s, \dot{b}w^s) (1 + O(\rho)).$$

In particular, the average value (6.7) is $\bar{\xi} = \bar{\xi}_0$, so under the decomposition (6.8) $\xi = (\zeta, \bar{\xi}_0)$ where

$$\zeta = (0, \dot{a}z^s, \dot{b}w^s) (1 + O(\rho))$$

and therefore $|\nabla\zeta| \leq c\rho(|\dot{a}| + |\dot{b}|)$. But $|\dot{a}| + |\dot{b}| \leq c\|(\xi, h)\|_1$ by Lemma 7.1b, the term $\rho|\zeta|$ in (9.7) is bounded by $|\mu_k|^{1/4}|\zeta|$ on B_k , and $\sup|h| \leq c\|h\| \leq c\|(\xi, h)\|_1$ as in (7.11). We can then multiply (9.7) by $\rho^{-\delta}$ and use (5.10) to integrate over the region over B_k where $\rho \leq |\mu_k|^{1/4}$. That yields the bound

$$\begin{aligned} \|\mathbf{D}_F(\xi, h)\|_{0, B_k} &\leq c|\mu_k|^{1/4} \|\zeta\|_0 + (\|(\xi, h)\|_1 + |\bar{\xi}| + \sup|h|) \\ &\quad \times \left(\|\rho^{1-\delta/2}\|_{L^2(B_k)} + \|\rho^{1-\delta/2}\|_{L^4(B_k)} \right) \\ &\leq c|\mu_k|^{(1-\delta/2)/4} \|(\xi, h)\|_1 \end{aligned}$$

with $\delta \leq 1/6$ and $|\mu_k| \leq c|\lambda|^{1/|s|}$. The lemma follows after summing on k and adding (9.8). \square

PROPOSITION 9.3. *For small $\delta, \varepsilon > 0$ there is a λ_0 so that $\mathcal{M}_s^{V, \delta}(Z_\lambda)$ lies in an ε -neighborhood of $\text{Approx}_s^\delta(Z_\lambda)$ whenever $|\lambda| < \lambda_0$.*

Proof. We will prove this by contradiction. Assume that, for some $\delta > 0$ there exist $\varepsilon_0 > 0$, a sequence $\lambda_n \rightarrow 0$, and $(f_n, C_n) \in \mathcal{M}_s^{V, \delta}(Z_{\lambda_n})$ such that the distance from (f_n, C_n) to $\text{Approx}_s^\delta(Z_{\lambda_n})$ is at least ε_0 for all n . Then (f_n, C_n) has a subsequence which converges as in (5.1) to a limit $f_0 : C_0 \rightarrow Z_0$ from an ℓ -nodal curve C_0 . Using coordinates in the normal bundle to the ℓ -nodal strata in the Deligne-Mumford space as in (4.2), we can write each C_n as $(C_{0n}, \mu_{n,1}, \dots, \mu_{n,\ell})$ where C_{0n} is an ℓ -nodal curve close to C_0 . Choose $\hat{\mu}_n = (\hat{\mu}_{n,1}, \dots, \hat{\mu}_{n,\ell})$ with $\lambda_n = a_k b_k (\hat{\mu}_{n,k})^{s_k}$ for each k ; there are $|s|$ choices for each $\hat{\mu}_n$ which differ by roots of unity. The data $(f_0, C_0, \hat{\mu}_n)$ specifies approximate maps $F_n : C_0(\hat{\mu}_n) \rightarrow Z_{\lambda_n}$ in $\text{Approx}_s^\delta(Z_{\lambda_n})$ as prescribed in (6.4). We will show that for some choice of the roots of unity a subsequence satisfies

$$(9.8) \quad \text{dist}(C_{0n}(\mu_n), C_0(\hat{\mu}_n)) + \text{dist}(f_n, F_n) < \varepsilon_0$$

where the left-hand side is the Finsler distance between maps defined in (9.1). Then (9.8) contradicts our assumption and proves the proposition.

Lemma 5.3 shows that $(\mu_n/\hat{\mu}_n)^s \rightarrow 1$ at each node. After passing to a subsequence and modifying our choice of $\hat{\mu}_n$ we have $\mu_n/\hat{\mu}_n \rightarrow 1$. But then $\text{dist}(C_{0n}(\mu_n), C_0(\hat{\mu}_n)) \rightarrow 0$ by (4.10).

Let $A_k(\rho_0)$ denote the k^{th} neck region $\{\rho \leq \rho_0\}$ of C_{0n} . Outside the union of the $A_k(\rho_0)$ both the approximate maps F_n and the maps f_n converge uniformly to f_0 in C^∞ , so on this region $\text{dist}(f_n, F_n) \rightarrow 0$ as $n \rightarrow \infty$. Inside each $A_k(\rho_0)$, we can write $F_n - f_n = (\hat{\zeta}_n, \bar{\xi}_n)$ in the notation of (6.7) and (6.8) with $\bar{\xi}_n \rightarrow 0$ because $f_n \rightarrow f_0$ in C^0 . Write $\hat{\zeta}_n = \zeta_n + (F_n - f_0)$ where $\zeta_n = f_0 - f_n$. Then $\|\zeta_n\|_1 \leq c\rho_0^{1/6}$ on $A_k(\rho_0)$ by Lemma 5.4, while $|F_n - f_0| + |\nabla(F_n - f_0)| \leq c\rho$ by (6.15). Integration using (5.10) then gives the bound $\|F_n - f_0\|_{1,p,s} \leq c_p \rho_0^{1-\delta/2}$ on the integrals (6.9) over each $A_k(\rho_0)$. Combining

these facts and taking ρ_0 small enough, we conclude that $\text{dist}(f_n, F_n)$, which is uniformly equivalent to $\|\zeta_n\|_1 + |\bar{\zeta}|$ by definition, is less than $\varepsilon_0/2$ for large n . Thus (9.8) holds. \square

The next step is to correct each approximate map $(F, C_\mu) \in \mathcal{A}\text{pprox}_s^\delta(Z_\lambda)^*$ to get a true (J, ν) -holomorphic map (F', C') . Specifically, (F', C') will be a solution of the equation

$$(9.9) \quad \bar{\partial}_j f = \nu_f \quad \text{where} \quad (f, C) = \exp_{F, C_\mu}(P_F \eta),$$

where P_F is the operator of Proposition 8.1, and where η lies in the Banach space $L_s(\Lambda^{01}(F^*TZ_\lambda))$ obtained by completing $\Lambda^{01}(F^*TZ_\lambda)$ in the norm (6.11) with $m = 0$.

PROPOSITION 9.4. *There are constants ε, λ_0 and C , uniform on \mathcal{K}_δ , such that for each approximate map $F \in \mathcal{A}\text{pprox}_s^\delta(Z_\lambda)^*$ and $0 < |\lambda| < \lambda_0$ equation (9.9) has a unique solution $\eta \in L_s(\Lambda^{01}(F^*TZ_\lambda))$ in the ball $\|\eta\|_0 \leq \varepsilon$, and that solution is smooth and satisfies $\|\eta\|_0 \leq C|\lambda|^{\frac{1}{2|s|}}$.*

Proof. If we write $(F', C') = \exp_{(F, C_\mu)}(\zeta)$ where $\zeta = (\xi, h)$ then

$$(9.10) \quad \bar{\partial}_{C'} F' - \nu_{(F', C')} = \bar{\partial}_{C_\mu} F - \nu_F + \mathbf{D}_F \zeta + Q_F(\zeta)$$

where \mathbf{D}_F is the linearization at (F, C_μ) and the quadratic Q_F satisfies (as in [F])

$$(9.11) \quad \|Q_F(\zeta_1) - Q_F(\zeta_2)\|_0 \leq C (\|\zeta_1\|_1 + \|\zeta_2\|_1) \cdot \|\zeta_1 - \zeta_2\|_1$$

Taking $\zeta = P_F \eta$ and noting that $\mathbf{D}_F P_F \eta = \eta$, the equation (9.9) that we must solve becomes, via (9.10),

$$(9.12) \quad \eta + Q_F(P_F \eta) = -\Phi_F \quad \text{where} \quad \Phi_F = \bar{\partial}_{C_\mu} F - \nu_F.$$

Now define an operator T_F on the Banach space $L_s(\Lambda^{01}(F^*TZ_\lambda))$ by

$$T_F \eta = -\Phi_F - Q_F(P_F \eta).$$

By (9.11) and (8.3)

$$\begin{aligned} \|T_F \eta_1 - T_F \eta_2\|_0 &\leq C (\|P_F \eta_1\|_1 + \|P_F \eta_2\|_1) \cdot \|P_F(\eta_1 - \eta_2)\|_1 \\ &\leq C E^2 (\|\eta_1\|_0 + \|\eta_2\|_0) \cdot \|\eta_1 - \eta_2\|_0. \end{aligned}$$

Hence whenever $\varepsilon < 1/(4CE^2)$ and $\|T_F(0)\|_0 \leq \varepsilon/2$, the map $T_F : B(0, \varepsilon) \rightarrow B(0, \varepsilon)$ is a contraction. Therefore T_F has a unique fixed point in $B(0, \varepsilon)$, and that fixed point η satisfies (9.12) and $\|\eta\|_0 \leq 2\|T_F(0)\|_0 = 2\|\Phi_F\|_0$ is bounded by Lemma 6.9. Finally, since $\eta \in L_{\text{loc}}^4$ we have $\zeta = P_F \eta \in L_{\text{loc}}^{1,4}$ with $\mathbf{D}_F \zeta + Q_F(\zeta) = -P_F \Phi_F \in C^\infty$. Elliptic regularity then shows that ζ and η are smooth. \square

10. Convolutions and the sum formula for flat maps

We can now assemble the analysis of the previous several sections to show that the approximate moduli space, which is built from maps into Z_0 , is a good model of the moduli space of stable maps into the symplectic sum Z_λ . Recall that in Sections 3 and 5 we showed that as $\lambda \rightarrow 0$ stable maps into Z_λ limit to maps into Z_0 and that the complex structure on their domains is asymptotically determined by λ and the limit map up to a finite ambiguity corresponding to the different solutions of the equations $\mu_k^s = \lambda/a_k b_k$. That led to the definition of the model moduli space $\mathcal{M}odel_s(Z_\lambda)$ in Section 6. On the other hand, each element of $\mathcal{M}odel_s(Z_\lambda)$ defines an approximate holomorphic map by equation (6.4); for each λ that gives a gluing map

$$\Gamma_\lambda : \mathcal{M}odel_s(Z_\lambda)^* \rightarrow \mathcal{A}pprox_s(Z_\lambda)^* \subset \mathcal{M}aps(B, Z_\lambda \times \mathcal{U}).$$

Proposition 9.4 shows that each such approximate map can be uniquely perturbed to be a true (J, ν) -holomorphic map.

In this section we will show that $\mathcal{A}pprox_s(Z_\lambda)$ is isotopic to the space $\mathcal{M}_s^V(Z_\lambda)$ of δ -flat (labeled) maps through an isotopy compatible with the evaluation maps. Thus $\mathcal{M}odel_s(\lambda)$ keeps track of the fundamental homology class $[\mathcal{M}_s^V(Z_\lambda)]$ which defines the GW and GT invariants of Z_λ . Passing to homology, we then define a “convolution” operation and establish a formula of the form

$$(10.1) \quad \mathrm{GT}_X^V * \mathrm{GT}_Y^V = \mathrm{GT}_Z$$

under the assumption that all curves contributing to the invariants are δ -flat along V (this condition will be eliminated in Section 12).

We noted in (3.11) that as $\lambda \rightarrow 0$ the limits of the δ -flat maps into Z_λ lie in the compact set \mathcal{K}_δ of $\mathcal{M}^V(X) \times_{\mathrm{ev}} \mathcal{M}^V(Y)$. We will work on the corresponding compact sets $\mathcal{M}odel_s^\delta(Z_\lambda)$ and $\mathcal{A}pprox_s^\delta(Z_\lambda)$ defined in (8.1). We will also assume that for generic (J, ν)

$$(10.2) \quad \mathrm{coker} D_{f,s} = 0 \quad \text{for all } f \in \mathcal{K}_\delta.$$

In particular, this is true when the moduli space $\mathcal{M}^V(X) \times_{\mathrm{ev}} \mathcal{M}^V(Y)$ consists of only irreducible maps for generic (J, ν) .

THEOREM 10.1. *Fix an ordered sequence s and write $|s| = \prod s_i$. For generic (J, ν) for which (10.2) holds and for small $|\lambda|$, there is an $|s|$ -fold cover $\mathcal{M}odel_s^\delta(Z_\lambda)$ of \mathcal{K}_δ with a diagram*

$$(10.3) \quad \begin{array}{ccc} \bigsqcup_s \mathcal{M}odel_s^\delta(Z_\lambda) & \xrightarrow{\Phi_\lambda^1} & \bigsqcup_s \mathcal{M}_s^{V,\delta}(Z_\lambda) \\ \downarrow st & & \downarrow st \\ \widetilde{\mathcal{M}} \times \widetilde{\mathcal{M}} & \xrightarrow{\xi} & \widetilde{\mathcal{M}} \end{array}$$

where the top arrow is an embedding and is isotopic to the gluing map (6.5) to $\text{Model}_s^\delta(Z_\lambda)$. The diagram commutes up to homotopy. Furthermore, there is a constant $c = c(\delta)$ so that the image of Φ_λ^1 consists of maps which are $(\delta - c|\lambda|)$ -flat, and the image contains all $(\delta + c|\lambda|)$ -flat maps in $\mathcal{M}_s^V(Z_\lambda)$.

Proof. For each $(f, C, \mu) \in \text{Model}_s^\delta(Z_\lambda)$ the gluing map associates a smooth curve C_μ and an approximate map $F = F_{f, C, \mu} : C_\mu \rightarrow Z_\lambda$. By Proposition 9.1 any map that is L_s^1 close to $F = \Gamma_\lambda(f, C, \mu)$ can be uniquely written as

$$(10.4) \quad \Phi_\lambda(f', C', \mu', \eta) = \exp_{(F', C'_{\mu'})}(P_F \eta) \quad \text{for } F' = F_{f', C', \mu'},$$

for some L_s^0 section η of the bundle $\Lambda^{0,1}$ with $\|\eta\|_0 < \varepsilon$. Proposition 9.4 shows that for small $|\lambda|$ there is a unique such $\eta = \eta(f', C', \mu')$ such that (10.4) is (J, ν) -holomorphic. Then

$$\Phi_\lambda^t(f', C', \mu', \eta) = \exp_{(F', C'_{\mu'})}(t P_F \eta(f', C', \mu'))$$

is a smooth 1-parameter family of maps from $\text{Model}_s^\delta(Z_\lambda)$ to $\text{Maps}_s(B, Z_\lambda \times \mathcal{U})$ with $\Phi_\lambda^0 = \Gamma_\lambda$ and the image of Φ_λ^1 lying in the $(\delta - c|\lambda|)$ -flat maps in $\mathcal{M}_s^V(Z_\lambda)$. The uniqueness of η in the fibers of $\Lambda^{0,1}$, combined with Proposition 9.1 and the uniqueness of η in the fibres of $\Lambda^{0,1}$ imply that Φ_λ^1 is injective.

It remains to show that Φ_λ^1 is surjective. But Proposition 9.1 shows that (10.4) is onto at least a $c\varepsilon$ neighborhood of $\text{Approx}_s^{2\delta}(Z_\lambda)$ and Proposition 9.3 implies that $\mathcal{M}_s^{V, \delta}(Z_\lambda)$ lies in that neighborhood when $|\lambda|$ is small enough. Hence for small $|\lambda|$, each element of $\mathcal{M}_s^{V, \delta}(Z_\lambda)$ can be written in the form (10.4) with $(f', C', \mu') \in \text{Model}_s^{\delta+c|\lambda|}(Z_\lambda)$ and η the corresponding fixed point of Proposition 9.4 satisfying $\|\eta\|_0 \leq \varepsilon$. Thus Φ_λ^1 is surjective. \square

Diagram (10.3) leads to our first formula expressing the absolute invariants of a symplectic sum $Z = Z_\lambda$ in terms of the relative invariants of X and Y . Recall that the relative invariant GW_X^V is obtained by forming the space $\overline{\mathcal{M}}_{\chi, n, s}^V(X, A)$ of relatively stable maps and pushing forward its fundamental homology class by the map

$$(10.5) \quad \varepsilon_V : \overline{\mathcal{M}}_{\chi, n, s}^V(X, A) \rightarrow \widetilde{\mathcal{M}}_{\chi, n} \times X^n \times \mathcal{H}_{X, A, s}^V.$$

We can also consider the space of stable maps from compact, not necessarily connected, domains by taking the union of products of $\overline{\mathcal{M}}_{\chi, n, s}^V(X, A)$ and again pushing forward in homology. The resulting class in the homology of $\mathcal{M}_{\chi, n} \times X^n \times \mathcal{H}_{X, A, s}^V$ is the relative GT invariant (1.25). As we observed in the introduction (see Figure 1), it is the GT invariant that will appear in the symplectic sum formula.

To proceed, we should replace the vertical arrows in Diagram (10.3) by the above maps ε_V and pass to homology. We will do that in two steps, first

incorporating the spaces \mathcal{H}_X^V and then including the X^n . In each case we will see that the operation of gluing maps defines an extension of the bottom arrow in Diagram (10.3), which we examine in homology.

The convolution operation. We can glue a map f_1 into X to a map f_2 into Y provided the images meet V at the same points with the same multiplicity. The domains of f_1 and f_2 glue according to the attaching map ξ of (3.8), while the images determine elements of the intersection-homology spaces $\mathcal{H}_{X,A,s}^V$ and $\mathcal{H}_{Y,A,s}^V$ which glue according to the map g of (3.10). The convolution operation records the effect of these gluings at the level of homology.

For each s the attaching map (3.8) defines a bilinear form

$$(\xi_\ell)_* : H_*(\widetilde{\mathcal{M}}; \mathbb{Q}) \otimes H_*(\widetilde{\mathcal{M}}; \mathbb{Q}) \longrightarrow H_*(\widetilde{\mathcal{M}}; \mathbb{Q})$$

for $\ell = \ell(s)$. Similarly, for each s the map g from (3.10) induces a bilinear form on the homology of $\mathcal{H}_Y^V \times \mathcal{H}_Y^V$ with values in $RH_2(Z)$, the (rational) Novikov ring of $H_2(Z)$, namely

$$\langle \ , \ \rangle : H_*(\mathcal{H}_X^V; \mathbb{Q}) \otimes H_*(\mathcal{H}_Y^V; \mathbb{Q}) \longrightarrow RH_2(Z)$$

$$\langle h, h' \rangle_s = g_* \left[h \times h' \Big|_{\varepsilon^{-1}(\Delta_s)} \right] = \sum_{A \in H_2(Z)} g_* [\Delta_{A,s} \cap (h \times h')] t_A.$$

This last equality holds because $\varepsilon^{-1}(\Delta_s)$ is the union of components $\Delta_{A,s} = \varepsilon^{-1}(\Delta_s) \cap g^{-1}(A)$.

Combining the two bilinear forms gives the convolution operator that describes how homology classes of maps combine in the gluing operation.

Definition 10.2. The *convolution operator*

$$* : H_*(\widetilde{\mathcal{M}} \times \mathcal{H}_X^V; \mathbb{Q}[\lambda]) \otimes H_*(\widetilde{\mathcal{M}} \times \mathcal{H}_Y^V; \mathbb{Q}[\lambda]) \longrightarrow H_*(\widetilde{\mathcal{M}}; RH_2(Z)[\lambda])$$

is given by

$$(10.6) \quad (\kappa \otimes h) * (\kappa' \otimes h') = \sum_s \frac{|s|}{\ell(s)!} \lambda^{2\ell(s)} (\xi_{\ell(s)})_* (\kappa \otimes \kappa') \langle h, h' \rangle_s.$$

The right-hand side of (10.6) includes three numerical factors which keep track of how maps glue when we form the symplectic sum. Recall that the powers of λ record the Euler characteristic in the generating series of the invariants (1.7) and (1.24); the factor $\lambda^{2\ell(s)}$ in (10.6) reflects the relation (3.7) between the Euler characteristics when we glue along $\ell(s)$ points. The factor $|s|$ is the degree of the covering in Theorem 10.1; this reflects the fact that each stable map into Z_0 can be smoothed in $|s| = s_1 \cdot \dots \cdot s_\ell$ ways. Finally, note that elements in the space $\mathcal{M}_s^{V,\delta}(Z_\lambda)$ in Diagram 10.3 are *labeled* maps; i.e. they have $\ell(s)$ numbered curves on their domains as explained at the end

of Section 3. But the GW and GT invariants of Z_λ are defined using the space of *unlabeled* stable maps, which is the quotient of the space of labeled maps by the action of the symmetric group. That accounts for the factor $1/\ell(s)!$ in (10.6).

Since \mathcal{H}_X^V is the disjoint union of components $\mathcal{H}_{X,A,s}^V$ with $A \in H_2(X)$ and $\deg s = A \cdot V$, there is an isomorphism

$$H_*(\mathcal{H}_X^V) \cong \sum_A \sum_{\deg s = A \cdot V} H_*(\mathcal{H}_{X,A,s}^V) t_A.$$

Below, we will identify $h \in H_*(\mathcal{H}_X^V)$ with $\sum_A h_A t_A$, where h_A are its components in $H_*(\mathcal{H}_{X,A}^V)$.

Example 10.3. The formula for the convolution simplifies when there are no rim tori in X and Y , and therefore in Z (cf. (1.19)). Then (i) the relative invariants have an expansion of the form (A.3), (ii) the map g of (3.10) is the restriction to the diagonal $\Delta_s \subset V^s \times V^s$, and (iii) the h part of the convolution (10.6) is then given by the cup product with the Poincaré dual of the diagonal:

$$g_* \left[h \times h' \Big|_{\Delta_s} \right] = \text{PD}(\Delta_s) \cup (h \times h').$$

We can then ‘split the diagonal’ by fixing a basis $\{C^p\}$ of $H^*(\bigsqcup_s V^s)$ and writing

$$\text{PD}(\Delta_s) = \sum_{p,q} Q_{p,q}^V C^p \times C^q = \sum_p C^p \times C_p$$

where $Q_{p,q}^V$ is the intersection form of V^s for the basis $\{C^p\}$ and $C_p = \sum Q_{p,q}^V C^q$ is the dual basis (with respect to Q^V). If $\{\gamma^i\}$ is a basis of $H^*(V)$, let $\{\mathbf{C}_m\}$ be the basis (A.4) (in the appendix) of $H^*(\bigsqcup_s V^s)$ corresponding to $\{\gamma^i\}$ and let $\{\mathbf{C}_{m^*}\}$ be the one corresponding to the dual basis $\{\gamma_i\}$ (with respect to Q^V). The convolution then has the more explicit form

$$(10.7) \quad (\kappa \otimes h) * (\kappa' \otimes h') = \sum_m \frac{|m|}{m!} \lambda^{2\ell(m)} (\xi_{\ell(m)})_* (\kappa \otimes \kappa') \mathbf{C}_m(h) \mathbf{C}_{m^*}(h').$$

In passing from s to m , we used the fact that each fixed sequence m corresponds to $\binom{\ell(s)}{(m_{a,i})} = \frac{\ell(s)!}{m!}$ ordered sequences s .

More generally, let X be a symplectic manifold with two disjoint symplectic submanifolds U and V with real codimension two. Suppose that V is symplectically identified with a submanifold of similar triple (Y, V, W) and that the normal bundles of $V \subset X$ and $V \subset Y$ have opposite Chern classes. Let (Z, U, W) be the resulting symplectic sum. In this case, (3.10) is replaced by

$$(10.8) \quad g : \mathcal{H}_X^{U,V} \times_\varepsilon \mathcal{H}_Y^{V,W} \rightarrow \mathcal{H}_Z^{U,W}$$

which combines with the map $\xi_{\ell(s)}$ to give the convolution operator

$$(10.9) \quad * : H_*(\widetilde{\mathcal{M}} \times \mathcal{H}_X^{U,V}; \mathbb{Q}[\lambda]) \otimes H_*(\widetilde{\mathcal{M}} \times \mathcal{H}_Y^{V,W}; \mathbb{Q}[\lambda]) \longrightarrow H_*(\widetilde{\mathcal{M}} \times \mathcal{H}_Z^{U,W}; \mathbb{Q}[\lambda])$$

as in (10.6). It describes how homology classes of maps combine in the gluing operation for the symplectic sum.

Finally, we include the evaluation maps which record the images of the n marked points. These combine with the projections from (2.7) to give the diagram

$$(10.10) \quad \begin{array}{ccccc} & & \bigsqcup_s \mathcal{Model}_s(Z_\lambda) & & \\ & \swarrow & \downarrow \text{ev} & \searrow & \\ \bigsqcup_s \mathcal{M}_s^V(Z_0) & & \bigsqcup_n (X \sqcup Y)^n & & \bigsqcup_s \mathcal{M}_s^V(Z_\lambda) \\ \downarrow \text{ev} & & \downarrow \pi_0 & & \downarrow \text{ev} \\ \bigsqcup_n (X \sqcup Y)^n & \searrow \pi_0 & \bigsqcup_n (Z_0)^n & \swarrow \pi_\lambda & \bigsqcup_n (Z_\lambda)^n \end{array}$$

which commutes up to homotopy. We can also include the spaces $\widetilde{\mathcal{M}}$ of curves from Diagram (10.3). Pushing forward then gives $\pi_{0*}(\text{GT}_X^V * \text{GT}_Y^V) = \pi_{\lambda*}(\text{GT}(Z_\lambda))$ as shown in the next theorem.

THEOREM 10.4. *Assume that for generic (J, ν)*

- (a) *all maps in $\bigsqcup_s \mathcal{M}_s^V(Z_\lambda)$ are δ -flat along V when λ is small, and*
- (b) *all maps in $\bigsqcup_s \mathcal{M}_s^V(Z_0)$ are (X, V) and (Y, V) -admissible as defined in (1.22).*

Then for every $\alpha_0 \in \mathbb{T}^(Z_0)$*

$$(10.11) \quad \text{GT}_Z^{U \cup W}(\pi^* \alpha_0) = (\text{GT}_X^{U \cup V} * \text{GT}_Y^{V \cup W})(\pi_0^* \alpha_0).$$

Note that condition (b) is generically satisfied when (X, V) and (Y, V) are semipositive.

Proof. It suffices to verify this for decomposable elements $\alpha_0 = \alpha_0^1 \otimes \cdots \otimes \alpha_0^n$. Let $\alpha_V^k, \alpha_X^k, \alpha_Y^k$ denote the restriction of α_0^k to V, X and respectively Y . We can then choose geometric representatives B_V^k of the Poincaré dual of α_V^k in V and Poincaré duals B_X^k of α_X^k in X and B_Y^k of α_Y^k in Y which intersect V transversely such that, moreover, $B_X^k \cap V = B_Y^k \cap V = B_V^k$. Then the

inverse image under π_λ of $B_X^k \cup_{B_V^k} B_Y^k$ gives a continuous family of geometric representatives B_λ^k of the Poincaré dual of $\pi_\lambda^* \alpha_0^k$ in $H^*(Z_\lambda)$. The theorem then follows from Theorem 10.1 by cutting down the moduli spaces on the left of Diagram 10.10 by (B_X, B_Y) and the ones on the right by B_λ . Constraints in $H^*(\mathcal{M})$ are handled similarly. The details of such arguments are standard (cf. [RT1]). \square

We should comment on how the assumption that all maps are δ -flat enters the above proof. Notice that in the statement of Theorem 10.1 the δ -flat maps in $\text{Model}_s^\delta(Z_\lambda)$ are paired with maps in $\mathcal{M}_s(Z_\lambda)$ which are not exactly δ -flat — there is a slight variation in δ . But when all contributing maps are δ -flat, the cut-down moduli space $\text{ev}^{-1}(B_\lambda) \subset \overline{\mathcal{M}}(Z_\lambda)$ limits as $\lambda \rightarrow 0$ to a compact subset of the open set $\mathcal{M}_s \times_{\text{ev}} \mathcal{M}_s$ as in (3.11). Hence for sufficiently small δ the set of elements of the limit set which are δ -flat is the same as the set of 2δ -flat elements, so the variation in δ is inconsequential.

Theorem 10.4 is a formula for the GT invariants evaluated on only certain constraints in $H^*(Z_\lambda)$ — those of the form $\pi^*(\alpha_0)$. The following definition characterizes those constraints. It is based on the diagram induced by the collapsing maps of (2.7)

$$(10.12) \quad \begin{array}{ccc} & \mathbb{T}^*(Z_0) & \\ \pi^* \swarrow & & \searrow \pi_0^* \\ \mathbb{T}^*(Z) & & \mathbb{T}^*(X) \oplus \mathbb{T}^*(Y). \end{array}$$

Definition 10.5. We say that a constraint $\alpha \in \mathbb{T}^*(Z)$ separates as (α_X, α_Y) if there exists an $\alpha_0 \in \mathbb{T}^*(Z_0)$ so that $\pi^* \alpha_0 = \alpha$ and $\pi_0^*(\alpha_0) = (\alpha_X, \alpha_Y) \in \mathbb{T}(H^*(X) \oplus H^*(Y))$.

Here are three observations to help clarify which classes $\alpha \in H^*(Z)$ separate. These follow by combining the Mayer-Vietoris sequences for $Z_\lambda = (X \setminus V) \cup (Y \setminus V)$:

$$H^{*-1}(S_V) \xrightarrow{\delta^*} H^*(Z) \xrightarrow{i^*} H^*(X \setminus V) \oplus H^*(Y \setminus V) \rightarrow H^*(S_V) \xrightarrow{\delta^*}$$

and the similar one for Z_0 with the Gysin sequence for $p: S_V \rightarrow V$:

$$(10.13) \quad H^{*-2}(V) \xrightarrow{\cup c_1} H^*(V) \xrightarrow{p^*} H^*(S_V) \xrightarrow{p_*} H^{*-1}(V).$$

- (a) When the first map in (10.13) is injective then all classes α separate. In dimension four, this occurs whenever the normal bundle of V in X is topologically nontrivial.
- (b) In general the separating classes are those α for which $j^*(\alpha) \in H^*(S_V)$ is in the image of the second map in (10.13).

- (c) The decomposition (α_X, α_Y) , if it exists, is unique only up to elements in the image of $\delta_X^* \oplus \delta_Y^* : H^{*-1}(S_V) \rightarrow H^*(X) \oplus H^*(Y)$ (the elements that can be “pushed to either side”).

By Definition 10.5 and for simplicity when U and W is empty, Theorem 10.4 becomes:

THEOREM 10.6. *Under the assumptions of Theorem 10.4 suppose that moreover α separates as (α_X, α_Y) . Then*

$$(10.14) \quad \mathrm{GT}_Z(\alpha) = (\mathrm{GT}_X^V * \mathrm{GT}_Y^V)(\alpha_X, \alpha_Y).$$

Note that when (α_X, α_Y) decomposes as $\alpha = \alpha_X \otimes \alpha_Y$ the right-hand side is $\mathrm{GT}_X^V(\alpha_X) * \mathrm{GT}_Y^V(\alpha_Y)$, but in general (α_X, α_Y) is a sum of tensors of the form $(\alpha_X^1 + \alpha_Y^1) \otimes \cdots \otimes (\alpha_X^k + \alpha_Y^k)$ and the right-hand side of (10.14) is the corresponding sum.

To focus on the decomposable case we make another definition: we say α is *supported off the neck* if the restriction $j^*(\alpha) \in H^*(S_V)$ vanishes. In that case α separates into relative classes $\alpha_X \in H^*(X, V)$ and $\alpha_Y \in H^*(Y, V)$, generally in several ways. For each such decomposition Theorem 10.6 gives

$$(10.15) \quad \mathrm{GT}_Z(\alpha_X, \alpha_Y) = \mathrm{GT}_X^V(\alpha_X) * \mathrm{GT}_Y^V(\alpha_Y).$$

This was the formula described in [IP3].

Example 10.7. Take α to be the Poincaré dual of a point in Z . This constraint is supported off the neck and has two independent decompositions depending on whether the point is in X or Y .

Example 10.8. Suppose $\alpha = \alpha_X \otimes \alpha_Y$ is supported off the neck and there are no rim tori in (X, V) and (Y, V) and that all curves contributing to the invariants are δ -flat along V . Then we can choose a basis of $H^*(V)$ and expand the relative GT invariants as in the appendix. Combining (10.15) with (10.7) gives the explicit formula

$$\begin{aligned} & \mathrm{GT}_{\chi, A, Z}(\alpha_X, \alpha_Y) \\ &= \sum_{\substack{A=A_1+A_2 \\ \chi_1+\chi_2-2\ell(m)=\chi}} \sum_m \lambda^{2\ell(m)} \frac{|m|}{m!} \mathrm{GT}_{\chi_1, A_1, X}^V(\alpha_X; C_m) \cdot \mathrm{GT}_{\chi_2, A_2, Y}^V(C_{m^*}; \alpha_Y). \end{aligned}$$

Note that from the definition of relative invariants, the only terms contributing are those for which $A_1 \cdot V = \ell(m) = A_2 \cdot V$. E. Getzler has pointed out that the formula above can be neatly expressed in terms of the generating series

(A.6) and the intersection matrix Q^V of V , specifically,

$$\begin{aligned} & \text{GT}_Z(\alpha_X, \alpha_Y) \\ &= \exp \left(\sum_{a,i,j} a \lambda^2 Q_{ij}^V \frac{\partial}{\partial z_{a,i}} \frac{\partial}{\partial w_{a,j}} \right) \bigg|_{z=w=0} (\text{GT}_X^V(\alpha_X)(z) \cdot \text{GT}_Y^V(\alpha_Y)(w)). \end{aligned}$$

Because the decomposition of separating constraints α is not unique, we can often choose several different decompositions, and use Theorem 10.14 to get several expressions for the same GT invariant. That yields relations among relative GT invariants. In Section 15 we will use that idea to derive recursive formulas which determine the relative invariants in some interesting cases.

11. The space \mathbb{P}_V and the S-matrix

Starting from the normal bundle $N_X V$ of V in X , we can form the \mathbb{P}^1 bundle

$$\mathbb{P}_V = \mathbb{P}(N_X V \oplus \mathbb{C})$$

over V by projectivizing the sum of the normal bundle $N_X V$ and the trivial complex line bundle. Let $\pi : \mathbb{P}_V \rightarrow V$ be the projection map. In \mathbb{P}_V , the zero section V_0 and the infinity section V_∞ are disjoint symplectic submanifolds, both symplectomorphic to V . Moreover, note that $\mathbb{P}_V \underset{V_0=V_\infty}{\#} \mathbb{P}_V = \mathbb{P}_V$.

Under the natural identification of V_0 with V_∞ , the convolution operation (10.9) defines an algebra structure on $H_*(\widetilde{\mathcal{M}} \times \mathcal{H}_{\mathbb{P}_V}^{V_\infty, V_0}; \mathbb{Q}[\lambda])$. This allows us to multiply by GT invariants. Of particular interest are the invariants with no constraints on the image, that is $\text{GT}_{\mathbb{P}_V}^{V_\infty, V_0}(\alpha)$ with $\alpha = 1$, which give an operator

$$(11.1) \quad \left[\text{GT}_{\mathbb{P}_V}^{V_\infty, V_0}(1) \right] * : H_*(\widetilde{\mathcal{M}} \times \mathcal{H}_{\mathbb{P}_V}^{V_\infty}; \mathbb{Q}[\lambda]) \rightarrow H_*(\widetilde{\mathcal{M}} \times \mathcal{H}_{\mathbb{P}_V}^{V_0}; \mathbb{Q}[\lambda])$$

defined by a power series as in (1.24). This operator is key to the general symplectic sum formula given in the next section. In this section we describe (11.1) and its inverse and develop some examples.

Each (J, ν) -holomorphic bubble map f into \mathbb{P}_V projects to a map $f_V = \pi \circ f$ into V . Although f_V may not be (J, ν) -holomorphic, we can still ask whether f_V is stable, using the second definition of stability given before (1.4); namely, f is stable if its restriction to each unstable domain component is nontrivial in homology.

Definition 11.1. A (V_∞, V_0) -stable map $f : C \rightarrow \mathbb{P}_V$ is \mathbb{P}_V -trivial if each of its components is an unstable rational curve whose image represents a multiple of the fiber F of \mathbb{P}_V .

Thus the \mathbb{P}_V -trivial curves are rational curves representing dF with one marked point on the zero section and one on the infinity section, both intersecting with multiplicity d . Let $\mathcal{M}_{\mathbb{I}}$ denote the set of \mathbb{P}_V -trivial maps in $\mathcal{M}_{\mathbb{P}_V}^{V_\infty, V_0}$ and consider the disjoint union

$$(11.2) \quad \mathcal{M}_{\mathbb{P}_V}^{V_\infty, V_0} = \mathcal{M}_{\mathbb{I}} \sqcup \mathcal{M}_R$$

where \mathcal{M}_R is the set of non- \mathbb{P}_V -trivial maps.

For the next lemma we fix a compatible triple (ω, J, g) on \mathbb{P}_V for which the projection $\pi : \mathbb{P}_V \rightarrow V$ is holomorphic and is a Riemannian submersion; such a triple is described before the proof of Proposition 6.6 of [IP4]. Then each J -holomorphic map (f, j) into \mathbb{P}_V projects to a J -holomorphic map $(\pi \circ f, j)$ into V .

LEMMA 11.2. (a) $\mathcal{M}_{\mathbb{I}}$ is both open and closed. The corresponding decomposition of (11.1) is

$$(11.3) \quad \text{GT}_{\mathbb{P}_V}^{V_\infty, V_0}(1) = \mathbb{I} + R^{V_\infty, V_0};$$

that is, the \mathbb{P}_V -trivial maps contribute the identity to the GT invariant.

(b) The non- \mathbb{P}_V -trivial maps have $E(f_V) \geq \alpha_V$, where α_V is the constant of Definition 3.1.

(c) For each fixed A , n and χ , the corresponding term in the convolution $R^m = R * \cdots * R$ vanishes for m large enough. Therefore, the inverse of GT is well defined by:

$$(11.4) \quad \left(\text{GT}_{\mathbb{P}_V}^{V_\infty, V_0}(1) \right)^{-1} = \sum_{m=0}^{\infty} (-1)^m R^m.$$

Proof. (a) Clearly $\mathcal{M}_{\mathbb{I}}$ is closed. To show that the complement of $\mathcal{M}_{\mathbb{I}}$ is closed, suppose that a sequence (f_i) in the complement converges to a trivial map f in the topology of the space of stable maps. Then the homology classes converge and so, after passing to a subsequence, we can assume that each f_i represents dF . Similarly, the stabilizations of the domains converge in the Deligne-Mumford space, so we can assume that all domain components of each f_i are unstable. But then the f_i lie in $\mathcal{M}_{\mathbb{I}}$. We conclude that $\mathcal{M}_{\mathbb{I}}$ is both open and closed. Finally, the decomposition (11.2) gives splitting (11.3) of the GT invariant because convolution by elements of $\mathcal{M}_{\mathbb{I}}$ is the identity.

(b) If $E(f_V) < \alpha_V$ then, as in the proof of Lemma 1.5b of [IP4], every component of the domain is unstable and f_V is trivial in homology and therefore f represents a multiple of F .

(c) For each (J, ν) , we shall bound the number N for which there are maps in the moduli space defining the convolution R^N . That moduli space consists of maps f from a domain C (whose Euler class χ and number n of

marked points is fixed) to the singular manifold $\mathbb{P}_V \# \cdots \# \mathbb{P}_V$ obtained from N copies of \mathbb{P}_V by identifying the infinity section of one with the zero section of the next. Furthermore, these f decompose as $f = \bigcup f^j$ where f^j is a map from some of the components of C into the j^{th} copy of \mathbb{P}_V .

Fixing such an f , let N_1 be the number of f^j whose domain has at least one stable component C_j . These components appear in the stabilization $\text{st}(C)$. But $\text{st}(C)$ lies in the space $\mathcal{M}_{\chi,n}$ of stable curves, and hence has at most $\dim \mathcal{M}_{\chi,n}$ components. This gives an explicit bound for N_1 in terms of χ and n .

The remaining $N_2 = N - N_1$ of the f^j each have unstable domain C_j with $\pi_*[f^j(C_j)] \in H_2(V)$ nontrivial, and so satisfy $E(\pi \circ f^j) > \alpha_V$ by (b) above. Furthermore, since C_j is unstable f^j is $(J, 0)$ -holomorphic. We therefore have

$$N_2 \alpha_V \leq \sum E(\pi \circ f^j) \leq E(\pi \circ f) \leq C \langle \omega_V, \pi_* A \rangle + C_\nu,$$

where the first sum is over those j contributing to N_2 and the last inequality is as in the proof of Lemma 12.1 below. Since the right-hand side depends only on A , χ , and (J, ν) , this bounds N_2 and hence N . \square

Definition 11.3. The S -matrix is defined to be the inverse of the GT invariant of Lemma 11.2:

$$S_V = \left(\text{GT}_{\mathbb{P}_V}^{V_\infty, V_0}(1) \right)^{-1}.$$

(Note that this depends not just on V but on the normal bundle to V and the 1-jet of (J, ν) along V .)

The symplectic sum of (X, U, V) and $(\mathbb{P}_V, V_\infty, V_0)$ along $V = V_\infty$ is a symplectic deformation of (X, U, V) , and so has the same GT invariant. The convolution then defines an operation

$$H_*(\widetilde{\mathcal{M}} \times \mathcal{H}_X^{U,V}; \mathbb{Q}[\lambda]) \otimes H_*(\widetilde{\mathcal{M}} \times \mathcal{H}_{\mathbb{P}_V}^{V_\infty, V_0}; \mathbb{Q}[\lambda]) \longrightarrow H_*(\widetilde{\mathcal{M}} \times \mathcal{H}_X^{U,V}; \mathbb{Q}[\lambda]).$$

Thus for each choice of constraints $\alpha \in \mathbb{T}^*(\mathbb{P}_V, V_\infty \cup V_0)$, the GT invariant of \mathbb{P}_V relative to its zero and infinity section defines an endomorphism

$$(11.5) \quad \text{GT}_{\mathbb{P}_V}^{V_\infty, V_0}(\alpha) \in \text{End} \left(H_*(\widetilde{\mathcal{M}} \times \mathcal{H}_X^{U,V}; \mathbb{Q}[\lambda]) \right)$$

which describes how families of curves on X are modified — “scattered” — as they pass through a neck modeled on $(\mathbb{P}_V, V_\infty, V_0)$ containing the constraints α .

The identity endomorphism in (11.5) is always realized as the convolution by the element

$$\mathbb{I} \in H_*(\widetilde{\mathcal{M}} \times \mathcal{H}_{\mathbb{P}_V}^{V_\infty, V_0}; \mathbb{Q}[\lambda])$$

corresponding to that part of GT coming from \mathbb{P}_V -trivial maps.

Example 11.4. When $V = \mathbb{P}^1$, $\mathbb{P}_V \rightarrow V$ is one of the rational ruled surfaces with its standard symplectic structure. If we wish to count all pseudo-holomorphic maps, without constraints on the genus or the induced complex structure, the relevant S-matrix is the relative GT invariant with $(\kappa, \alpha) = (1, 1)$. This case works out neatly: Lemma 14.6 below implies that $S_V = \text{Id}$.

Example 11.5. When we put no constraints on either the domain or the image S_V is an operator given in terms of $\text{GT}_{\mathbb{P}_V}^{V_\infty, V_0}$ by the S -matrix expansion (11.2). In cases where there are no rim tori in \mathbb{P}_V , we can expand the GT invariants in the power series (A.6) of the appendix. Let $\text{GT}_{\chi, A}(C_m; C_{m'})$ denote the relative invariant of \mathbb{F} satisfying the contact constraints C_m along V_∞ and $C_{m'}$ along V_0 . Then the S -matrix S_V has an expansion like (A.6) with coefficients

$$\begin{aligned} S_{\chi, A}(C_m; C_{m'}) &= \delta_{m, m'} - \text{GT}_{\mathbb{P}_V, \chi, A}^{V_\infty, V_0}(C_m; C_{m'}) \\ &+ \sum_{\substack{A_1 + A_2 = A \\ \chi_1 + \chi_2 - 2\ell(s_1) = \chi}} \sum_{m_1} \lambda^{2\ell(m_1)} \frac{|m_1|}{m_1!} \text{GT}_{\mathbb{P}_V, \chi_1, A_1}^{V_\infty, V_0}(C_m; C_{m_1}) \\ &\times \text{GT}_{\mathbb{P}_V, \chi_2, A_2}^{V_\infty, V_0}(C_{m_1^*}; C_{m'}) - \dots \end{aligned}$$

12. The general sum formula

In all of our work thus far we have assumed that the (J, ν) -holomorphic maps we are gluing are δ -flat as in Definition 3.1. In this section we remove this flatness assumption and prove the symplectic sum formula in the general case.

The idea is to reduce the general case to the δ -flat case by degenerating along many parallel copies of V . Thus instead of viewing Z_λ as the symplectic sum $X \#_V Y$ along V we regard it as the symplectic sum of $2N + 2$ spaces: X and Y at the ends and $2N$ middle pieces each of which is a copy of the ruled space \mathbb{P}_V associated to V — see Figure 2 of the introduction. The pigeon-hole principle then implies that for large N all holomorphic maps into Z_λ are close to maps which are δ -flat along each ‘seam’ of the $2N$ -fold sum.

LEMMA 12.1. *There is a constant $E = E_{\chi, n, A}(J, \nu)$ such that every (J, ν) -holomorphic map into Z representing a class $A \in H_2(Z)$ has energy at most E .*

Proof. In an orthonormal frame $\{e_1, e_2 = je_1\}$ on the domain, the holomorphic map equation is $f_*e_1 + Jf_*e_2 = 2\nu(e_1)$. Taking the norm squared and noting that $\langle f_*e_1, Jf_*e_2 \rangle = f^*\omega(e_1, e_2)$ give $|df|^2 = 2|\nu|^2 + 2f^*\omega(e_1, e_2)$. The energy of (f, C) is therefore the L^2 norm of ν plus the topological quantity $\langle \omega, A \rangle$ plus a term — the symplectic area of C in the universal curve — which depends only on χ and n . The lemma follows. \square

For the remainder of this section we fix the data χ, n, A which determined the constant E of Lemma 12.1 and fix an integer N with

$$(12.1) \quad N\alpha_V > E$$

where $\alpha_V < 1$ is the constant of Definition 3.1.

Fixing a small λ_0 , we partition the neck (2.5) of Z_{λ_0} (see Figure 3 of §2) into $2N$ segments Z^j using the coordinate t from (2.4):

$$Z^j = \{z \in Z_{\lambda_0} \mid (j - N - 1) \leq Nt(z) \leq (j - N)\}, \quad j = 1, \dots, 2N.$$

Let Z_0 be the singular symplectic manifold obtained by symplectically cutting Z_{λ_0} at the values $t_j(z) = j - N - \frac{1}{2}$ in the middle of each of these segments. The construction of Section 2 defines a family

$$(12.2) \quad \mathcal{Z} \rightarrow D \subset \mathbb{C}^{2N}$$

of smoothings of Z_0 as in Theorem 2.1 but with $2N$ necks. The fiber over (μ_1, \dots, μ_{2N}) , defined for $|\mu| \ll |\lambda|$, is a space $Z(\mu_1, \dots, \mu_{2N})$ with a neck of size μ_j inside each Z^j . The fiber Z_0 over $\mu = 0$ is the singular space obtained by connecting X to Y through a series of $2N - 1$ copies of the rational ruled manifold \mathbb{P}_V associated with V . We symplectically identify Z_{λ_0} with a fixed generic fiber $Z(\mu_0)$ as depicted in Figure 2 of the introduction.

Fix $\delta > 0$ such that $\delta \leq \frac{\varepsilon}{10N}$ and consider the space $\mathcal{M} = \mathcal{M}_{\chi, n, A}(Z(\mu_0))$ of holomorphic maps into $Z(\mu_0)$. Let f^j denote the restriction of $f \in \mathcal{M}$ to $f^{-1}(Z^j)$. We can then define an open cover of \mathcal{M} that keeps track of the values of j for which the energy $E^\delta(f^j)$ on the δ neck around the cut is small as in equation (3.4). Specifically, to each subset $\{i_1, \dots, i_k\}$ of $\{1, \dots, 2N\}$ we associate the open subset of \mathcal{M}

$$(12.3) \quad \mathcal{M}^{i_1, \dots, i_k} = \left\{ f \in \mathcal{M} \mid E^\delta(f^j) < \alpha_V/2 \text{ for } j = i_1, \dots, i_k \right\}.$$

LEMMA 12.2. *The $\mathcal{M}^{i_1, \dots, i_k}$ cover $\mathcal{M} = \mathcal{M}_{\chi, n, A}(Z(\mu_0), A)$ and set theoretically*

$$(12.4) \quad \mathcal{M} = \bigcup \mathcal{M}^i - \bigcup \mathcal{M}^{i_1, i_2} + \bigcup \mathcal{M}^{i_1, i_2, i_3} - \dots$$

Proof. Each $f \in \mathcal{M}$ has $\sum_j E(f^j) \leq E(f) < E$, and so (12.1) implies that $f \in \mathcal{M}^i$ for at least one i . If $E^\delta(f^j) < \alpha_V/2$ for exactly ℓ of the j , then f is counted

$$\ell - \binom{\ell}{2} + \binom{\ell}{3} - \dots \pm \binom{\ell}{\ell} = 1$$

times on the right-hand side of (12.4). □

Now every $f \in \mathcal{M}^{i_1, \dots, i_k}$ has small energy in the segment Z^j for $j = i_1, \dots, i_k$. Replacing μ_j by $\lambda\mu_j$ for those values of j (and keeping the remaining μ_j fixed) defines a 1-parameter subfamily Z_λ of (12.2). That family degenerates in the middle of exactly k of the segments Z^j . At each of those degenerations f is δ -flat in the sense of Definition 3.1. Hence as $\lambda \rightarrow 0$

$$(12.5) \quad \mathcal{M}^{i_1, \dots, i_k} \rightarrow \mathcal{M}_X^V \times_{\text{ev}} (\mathcal{M}_{\mathbb{P}_V}^{V_\infty, V_0})^{k-1} \times_{\text{ev}} \mathcal{M}_Y^V.$$

Since the closure of $\mathcal{M}^{i_1, \dots, i_k}$ is a compact subset of the set of $\delta/2$ -flat maps, we can apply the sum formula (10.6). For fixed χ, n and A we obtain

$$(12.6) \quad \text{GT}_{X \# Y} = \text{GT}_X^V * \left[\sum_{k=1}^{2N} (-1)^{k-1} \binom{2N}{k} \left(\text{GT}_{\mathbb{P}_V}^{V_\infty, V_0} \right)^{k-1} \right] * \text{GT}_Y^V.$$

This formula appears to be dependent on the number of cuts $2N$. However, there is a way to rewrite it to see that it is independent of N . Note that after multiplication by GT the middle sum is a binomial expansion; in fact, by Lemma 11.2c,

$$\sum_{k=1}^{2N} (-1)^{k-1} \binom{2N}{k} (\text{GT})^{k-1} = \frac{1 - (1 - \text{GT})^{2N}}{\text{GT}} = \frac{1 - (-R)^{2N}}{\text{GT}} = \text{GT}^{-1}.$$

Thus the middle part of (12.6) is exactly the S -matrix of Definition (11.3). This gives the symplectic sum formula in the general case.

THEOREM 12.3 (symplectic sum formula). *Let (Z, U, W) be the symplectic sum of (X, U, V) and (Y, V, W) along V . Suppose that $\alpha \in \mathbb{T}^*(Z)$ is supported off the neck as in Example 10.8. For any fixed decomposition (α_X, α_Y) of α the relative GT invariant of Z is given in terms of the invariants of (X, U, V) and (Y, V, W) and the S -matrix (11.3) by*

$$(12.7) \quad \text{GT}_Z^{U, W}(\alpha) = \text{GT}_X^{U, V}(\alpha_X) * S_V * \text{GT}_Y^{V, W}(\alpha_Y).$$

Implicit in this is the assumption that the right-hand side of (12.7) is well-defined. As in Section 10, that will be the case if (X, V) and (Y, V) are semipositive or more generally if the condition (b) in Theorem 10.4 is generically satisfied. As noted in Remark 1.1 these assumptions are probably not essential.

Theorem 12.3 holds more generally when α separates as in Definition 10.5, except that the definition of the S -matrix needs to be enlarged. Instead of restricting $\text{GT}_{\mathbb{P}_V}^{V_\infty, V_0}$ to $\alpha = 1$ we restrict it to the subtensor algebra \mathbb{T}_V^* of $\mathbb{T}^*(\mathbb{P}_V)$ generated by the kernel of the composition

$$H^*(\mathbb{P}_V) \xrightarrow{i^*} H^*(SV_0) \xrightarrow{p_*} H^*(V_0)$$

where SV_0 is the circle bundle of the normal bundle to the zero section V_0 in \mathbb{P}_V , p_* is the integration along its fiber and $i : SV_0 \rightarrow \mathbb{P}_V$ is the inclusion. In that case we get an S -matrix defined by

$$(12.8) \quad S_V = (\mathrm{GT}_{\mathbb{P}_V}^{V_0, V_\infty} \Big|_{\mathbb{T}_V^*})^{-1}$$

In the important case when U and W are empty Theorem 12.3 expresses the absolute invariant of Z in terms of the relative invariants of X and Y .

THEOREM 12.4. *Let Z be the symplectic sum of (X, V) and (Y, V) and suppose that $\alpha \in \mathbb{T}^*(Z)$ separates as (α_X, α_Y) as in Definition 10.5. Then*

$$(12.9) \quad \mathrm{GT}_Z(\alpha) = (\mathrm{GT}_X^V * S_V * \mathrm{GT}_Y^V)(\alpha_X, \alpha_Y).$$

where S_V is the S -matrix (12.8).

If moreover α decomposes as $\alpha = \alpha_X \otimes \alpha_Y$ then (12.9) becomes

$$\mathrm{GT}_Z(\alpha) = \mathrm{GT}_X^V(\alpha_X) * S_V(\alpha_V) * \mathrm{GT}_Y^V(\alpha_Y)$$

where $\alpha_V \in \mathbb{T}_V^*$ is the pullback to \mathbb{P}_V of the restriction of α to V .

As a check, it is interesting to verify the symplectic sum formula in one very simple case where the GW invariant is simply the Euler characteristic.

Example 12.5. Fix an elliptic curve C with one marked point (and fixed complex structure). Consider the (J, ν) -holomorphic maps with domain C and representing the class 0. When $\nu = 0$ all such maps are maps to a single point, so the moduli space is X itself. Furthermore, the fiber of the obstruction bundle at a constant map p is $H^1(C, p^*TX)$, which is naturally identified with T_pX . The (virtual) moduli space for $\nu \neq 0$ consists of the zeros of the generic section $\overline{\nu} = \int_C \nu$ of this obstruction bundle $TX \rightarrow X$. Thus this particular GW invariant is $\chi(X)$.

Similarly, when $\nu = 0$ the moduli space of V -regular curves is $X \setminus V$ and its V -stable compactification, defined in [IP4], is X . To compute the GW invariant relative to V , we need to know how many of these point maps become V -regular after we perturb to a generic V -compatible $\nu \neq 0$. Because any V -compatible ν is tangent to V along V the corresponding section $\overline{\nu}$ has $\chi(X)$ zeros on X , and of those, $\chi(V)$ lie on V . Thus the relative invariant is $\mathrm{GW}_X^V = \chi(X) - \chi(V)$. Moreover, the only contribution of the S -matrix in this case comes from the relative invariant of $(\mathbb{P}_V, V_0 \sqcup V_\infty)$ which counts maps with fixed domain C representing the class 0. But similarly, this invariant is $\mathrm{GW}_{\mathbb{P}_V}^{V_0, V_\infty} = \chi(\mathbb{P}_V) - 2\chi(V) = 0$, and therefore the S matrix does not contribute in this case. The symplectic sum formula then reduces to the formula

$$\chi(X) + \chi(Y) - 2\chi(V) = \chi(X \#_V Y).$$

Much more interesting examples will be given in Section 15.

Finally, one can also include ψ and τ classes as constraints. As in [IP4], $\phi_i \in H^2(\overline{\mathcal{M}}_{g,n})$ is the first Chern class of \mathcal{L}_i , the relative cotangent bundle at the i^{th} marked point p_i . There is a similar bundle $\tilde{\mathcal{L}}_i$ over the space of stable maps whose fiber at a map f is the cotangent space to the (unstabilized) domain curve, and whose Chern class is denoted by ψ_i . It is also useful to pair each ψ_i class with an $\alpha_i \in H^*(Z)$ and consider the ‘descendent’ $\tau_k(\alpha_i) = \text{ev}_s i^*(\alpha_i) \cup \psi_i^k$. It is a straightforward exercise, left to the reader, to incorporate these constraints into Theorems 12.3 and 12.4.

13. Constraints passing through the neck

Not every constraint class $\alpha \in H^*(Z)$ separates as in Definition 10.5. Yet for applications it is useful to have a version of the symplectic sum formula for more general constraints — ones whose Poincaré dual cuts across the neck. Since the Poincaré dual of $\alpha \in H^*(Z)$ restricts to a class in $H_*(X, V)$ such a general symplectic sum formula will necessarily involve relative GT invariants of classes $\alpha \in H^*(X \setminus V)$. That requires generalizing the relative invariant GT_X^V , which was defined in [IP4] only for constraints in $H^*(X)$.

We begin by recalling the ‘symplectic compactification’ of $X \setminus V$ which was used in [IP4]. Let \hat{X} be the manifold obtained from $X \setminus V$ by attaching as boundary a copy of the unit circle bundle $p : SV \rightarrow V$ of the normal bundle of V in X , and let $p : \hat{X} \rightarrow X$ be the natural projection. Suppose that Z is a symplectic sum obtained by gluing \hat{X} to a similar manifold \hat{Y} along SV . We can then consider stable maps in Z constrained by classes B in $H_k(Z)$, i.e. the set of stable maps f with the image $f(x)$ of a marked point lying on a geometric representative of B . If we restrict g to the \hat{X} side, such geometric representatives define constraints associated with classes in $H_*(\hat{X}, \partial\hat{X})$.

Specifically, given a class $B \in H_*(\hat{X}, \partial\hat{X})$, we can find a pseudo-manifold P with boundary Q and a map $\phi : P \rightarrow X$ so that $\phi(Q) \subset \partial\hat{X}$ represents B and use this to cut-down the moduli space. Thus for generic (J, ν)

$$\varepsilon_V \left(\overline{\mathcal{M}}_s^V(X, A) \right) \cap p(\phi(P))$$

defines an orbifold with boundary denoted by

$$(13.1) \quad \text{GT}_{X,A,s}^V(\phi).$$

After we cut down by further constraints of the appropriate dimension, this reduces to a finite set of points, giving numerical invariants constructed using ϕ . This is particularly simple when $B \in H_*(X \setminus V)$, i.e. when B can be represented by a map into $\hat{X} \setminus \partial\hat{X}$. The cobordism argument of Theorem 8.1 of [IP4] then shows that the relative invariants (13.1) are well-defined. Note that these relative invariants depend on $B \in H_*(X \setminus V)$, not on its inclusion

$i_*B \in H_*(X)$. For example, rim tori and the zero class in $H_2(\hat{X}, \partial\hat{X})$ have the same image under $p : \hat{X} \rightarrow X$, but might have different invariants (13.1).

In general the constrained invariant (13.1) will not be well-defined but will depend on the choice of ϕ . The space

$$(13.2) \quad \mathcal{J}^V \times \text{Maps}((P, Q), (\hat{X}, \partial\hat{X}))$$

has a subset

$$W = \bigcup_{i=1}^n \left\{ (J, \nu, \phi) \mid \begin{array}{l} \text{there is a } V\text{-stable } (J, \nu)\text{-holomorphic} \\ \text{map } f \text{ with } f(p_i) \in p(\phi(Q)) \subset V \end{array} \right\}$$

where for some map the marked point p_i lands on the projection of $\phi(Q)$ into V . In general, W will have real codimension one, and thus will form walls which separate (13.2) into chambers.

LEMMA 13.1. *The number (13.1) is constant within a chamber. When $B = [\phi]$ satisfies $p_*[\partial B] = 0$ then there is only one chamber, and therefore (13.1) depends only on B .*

Proof. Any two pairs (f, ϕ) that lie in the same chamber can be connected by a path (f_t, ϕ_t) with $f_t(p_i) \in \phi_t(P \setminus Q)$. The cobordism argument of Theorem 8.1 of [IP4] then proves the first statement.

Each B in the kernel of $p_*\partial$ can be represented by a map ϕ as above with $\phi(Q)$ of the form $p^{-1}(R)$ for some $k-2$ cycle R in V . After restricting the last factor of (13.2) to such ϕ , the wall W has codimension two, giving the second statement. \square

The following lemma relates the invariants associated with different chambers.

LEMMA 13.2. (a) *Two maps $\phi_1, \phi_2 : P \rightarrow X$ that agree on ∂P define a class $a = [\phi_1 \# (-\phi_2)] \in H_*(X \setminus V)$ with Poincaré dual $\alpha = PD(a) \in H^*(X, V)$; the corresponding invariants are related by*

$$(13.3) \quad \text{GT}_X^V(\phi_1) = \text{GT}_X^V(\phi_2) + \text{GT}_X^V(\alpha).$$

(b) *If ϕ_1, ϕ_2 define the same class in $H_*(X, V)$ then there exists $\phi' : R \rightarrow \partial\hat{X}$ with $\partial R = Q_1 \sqcup (-Q_2)$ such that ϕ' agrees with ϕ_1 on Q_1 and agrees with ϕ_2 on Q_2 . Then ϕ_1 and $\phi_2 \# \phi'$ have the same boundary. Moreover,*

$$(13.4) \quad \text{GT}_X^V(\phi_2 \# \phi') = \text{GT}_X^V(\phi_2) + \text{GT}_X^V * \text{GT}_F^{VV}(\phi')$$

This actually means that in order to extend the definition of the relative invariants from [IP4], we only need to pick one geometric representative B (any one) such that $[B] \in H_*(X, V)$, $[\partial B] = \beta$ for each $\beta \in \text{Ker } [H_{*-1}(SV) \rightarrow H_{*-1}(X)]$.

Altogether, the invariants can be thought of as giving a map

$$(13.5) \quad \mathrm{GT}_X^V : \mathbb{T}^*(X \setminus V) \longrightarrow H_*(\widetilde{\mathcal{M}} \times \mathcal{H}_X^V)$$

which is noncanonical: it depends on the actual representatives for the class α as described in Lemma 13.2.

With this extended definition of the relative invariants, the proof of Theorem 10.4 carries through. That proof began by choosing geometric representatives of constraints α which separate. For a general constraint $\alpha \in H^*(Z)$ we can still choose a geometric representative B of the Poincaré dual, and consider its restrictions B_X and B_Y to $(\hat{X}, \partial\hat{X})$ and $(\hat{Y}, \partial\hat{Y})$ respectively. The remainder of the proof still applies, giving a sum formula relating the invariants $\mathrm{GT}_Z(\alpha)$ of Z to the relative GT invariants (13.5) of X and Y cut down by the constraints B_X and B_Y .

14. Relative GW invariants in simple cases

The symplectic sum formula of Corollary 12.4 expresses the invariants of $X \# Y$ in terms of the relative invariants of X and Y . In the next section we will apply that formula to spaces that can be decomposed as symplectic sums where the spaces on one or both sides are simple enough that their relative invariants are computable. That strategy can succeed only if one has a collection of simple spaces with known relative invariants. This section provides four families of such simple spaces.

In some of the examples below the set \mathcal{R} of rim tori is nontrivial. In those cases we will give formulas for the invariants $\overline{\mathrm{GW}}_X^V$ defined in the appendix although, as the examples will show, it is sometimes possible to compute the GW_X^V themselves even though there are rim tori present.

14.1. Riemann surfaces. For Riemann surfaces one can consider the GW invariants as absolute invariants or relative to a finite set of points. These invariants count coverings, and the homology class A is simply the degree d of the covering.

In dimension two the symplectic sum is the same as the ordinary connect sum — one joins two Riemann surfaces by identifying a point on one with a point on another, and then smoothing. Of course, to apply the sum formula one must first find S_V , which in this case is built from the relative invariants of (\mathbb{P}^1, V) where $V = \{p_0, p_\infty\}$ two distinct points and where all the constraints lie on V . In that context, we fix a nonzero degree d and two sequences s, s' that describe the multiplicities of points at the preimages of p_0 and p_∞ respectively.

LEMMA 14.1. *The invariants $\mathrm{GW}_{d,g,s,s'}^V$ with no constraints except those on $V = \{p_0, p_\infty\}$ vanish except when $g = 0$ and s and s' are single points with*

multiplicity d . In that case

$$\mathrm{GW}_{d,0,s,s'}^V = 1/d$$

Moreover, in dimension two the S -matrix is always the identity.

Proof. This invariant is the oriented count of the 0-dimensional components of $\overline{\mathcal{M}}_{d,g,s,s'}^V$. But by (1.21)

$$\begin{aligned} \dim_{\mathbb{C}} \mathcal{M}_{d,g,s,s'}^V &= 2d + 2g - 2 + \ell(s) - \deg s + \ell(s') - \deg s' \\ &= 2g - 2 + \ell(s) + \ell(s') \end{aligned}$$

is zero only if $g = 0$ and $\ell(s) = \ell(s') = 1$, i.e. s and s' specify single points with multiplicity d . If we stabilize, there is only one such map, given by the equation $z \rightarrow z^d$, so its contribution to $\mathrm{GW}_{d,0,s,s'}^V$ is $1/d$. This map is \mathbb{P} -trivial, and hence contributes as the identity to the S -matrix. \square

The same dimension count gives the invariant with one constraint:

LEMMA 14.2. *When $d \geq 2$, the invariants $\mathrm{GW}_{d,g,s,s'}^V(b)$ of maps with one simple branch point over a fixed point in the target \mathbb{P}^1 and no other constraints except those on $V = \{p_0, p_\infty\}$ vanish except when $g = 0$ and $\ell(s) + \ell(s') = 3$, in which case $\mathrm{GW}_{d,0,s,s'}^V(b) = 1$.*

Perhaps the most interesting two-dimensional example is the $g = 1$ invariant of the torus T^2 .

LEMMA 14.3. *Fix an integer $k \geq 0$. Then the $g = 1$ invariants of the torus relative to a set V of k points form a series*

$$\mathrm{GW}_1^V(T^2) = \sum_d \mathrm{GW}_{d,1}^V(T^2) t^d$$

that is equal to the generating function for the sum of the divisors $\sigma(d) = \sum_{l|d} l$, namely

$$(14.1) \quad G(t) = \sum_{d=1}^{\infty} \sigma(d) t^d = \sum_{l=1}^{\infty} \frac{l t^l}{1 - t^l}.$$

Proof. This is a matter of counting the (unbranched) covers of the torus. That was done in [IP1] for $k = 0$. In general, for each degree d cover, each point of V has d inverse images, each with multiplicity one. Following the notation of [IP4] we order the inverse images and divide by $d!$, leaving us with $G(t)$ again. \square

14.2. $T^2 \times S^2$. Next we consider the $g = 1$ invariants of $X = T^2 \times S^2$. Thinking of this as an elliptic fibration over S^2 , we fix a section $S = \{pt\} \times S^2$ and two disjoint fibers F and denote the corresponding homology classes by s and f . Focusing on the classes df and $s + df$ for $d \geq 0$, we can form generating functions for the absolute GW invariants and the GW invariants relative to one or two copies of the fiber.

First consider the classes df , where the invariants $\text{GW}_{df,1}$, $\text{GW}_{df,1}^F$, and $\text{GW}_{df,1}^{F,F}$ have dimension 0 by (1.21). There are no rim tori in $X \setminus F$, and when V is one or two copies of the fiber we have $df \cdot V = 0$, so that $\ell = 0$ and V^ℓ is a point in (1.20). Therefore $\text{GW}_{df,1}^F$ has values in $H_2(X)$ and $\text{GW}_{df,1}^{F,F}$ has values in $\mathcal{H}^V = H_2(X) \times \mathcal{R}$. Thus all three invariants can be written as power series with numerical coefficients.

LEMMA 14.4. *The genus one invariants GW and GW^F in the classes df are given by*

$$\sum_d \text{GW}_{df,1} t_f^d = 2G(t_f) \quad \text{and} \quad \sum_d \text{GW}_{df,1}^F t_f^d = G(t_f)$$

with $G(t)$ as in (14.1). The corresponding relative invariants $\text{GW}^{F,F}$ are indexed by classes $df + R$ for rim tori R and these all vanish:

$$\sum_d \text{GW}_{df,1}^{F,F} t_{df+R} = 0.$$

Proof. The generic complex structure on a topologically trivial line bundle over T^2 admits no nonzero holomorphic sections. After projectivizing, we get a complex structure on $T^2 \times S^2$ for which the only holomorphic curves representing df are multiple covers of the zero section F_0 and the infinity section F_∞ . This is a generic V -compatible structure for $V = F_0$ or $F_0 \cup F_\infty$. As in Lemma 14.3 these contribute $G(t)$ to the power series for these invariants. (Note that for the relative invariant, we compute only the contribution of curves that have no components in V). \square

The invariants for the classes $s + df$ are more complicated. By (1.21) the corresponding moduli spaces have real dimension 4, and so become points in \mathcal{H}_X^V after imposing two point constraints; these constraints can be either points $p \in X \setminus V$, or $C_1(q)$, a contact of order 1 to V at a fixed point $q \in V$. Again rim tori R appear only for the invariant relative to two copies of a fiber.

LEMMA 14.5. *The genus one invariants GW and GW^F in the classes $s + df$, $d > 0$, are*

$$\sum_d \text{GW}_{s+df,1}(p^2) t^d = 2tG'(t) \quad \text{and} \quad \sum_d \text{GW}_{s+df,1}^F(p; C_1(q)) t^d = tG'(t).$$

The corresponding relative invariants $\mathrm{GW}^{F,F}$ can be indexed by classes $s + df + R$ for rim tori R and those with two point constraints on V vanish:

$$\sum_d \mathrm{GW}_{s+df+R,1}^{F,F}(\beta) t_{s+df+R} = \begin{cases} 2tG'(t) & \text{if } \beta = p, p \text{ and } R = 0, \\ tG'(t) & \text{if } \beta = p; C_1(q) \text{ and } R = 0, \\ 0 & \text{if } \beta = C_1(q_0); C_1(q_\infty). \end{cases}$$

Proof. We can compute using the product structure J_0 on $T^2 \times S^2$. Consider a J_0 -holomorphic map representing $s + df$, passing through generic points p_1 and p_2 , whose domain is a genus 1 curve $C = \cup C_i$. The projection onto the second factor gives a degree 1 map $C \rightarrow S^2$, so C must have a rational component C_0 which represents s . The projection of the remaining components is zero in homology, therefore they are multiple covers of the fibers. Because the total genus is one there is only one such component.

In summary, for the product structure J_0 , the only $g = 1$ holomorphic curves representing $s + df$ have two irreducible components, one of them a section S , and the other a multiple cover of a fiber $F \notin V$. A simple computation shows that $H^1(C, TX|_C) = 0$ for these curves, so that J_0 is generic for these classes. The constraints require that S pass through p_1 and F pass through p_2 , or vice versa. For each of those two cases there are d choices of the marked point on the domain of F , so the count is the same as in Lemma 14.4 with $G(t)$ replaced by $tG'(t)$. This gives the first formula.

The count for the second formula is similar. Any V -regular genus 1 holomorphic map through an interior point p and a point $q \in V$ has two components: a section through q and a d -fold cover of a fiber F through p . The fiber domain can be marked in d ways, giving the count $tG'(t)$.

For the invariant relative two copies of the fiber F , there are rim tori, but the discussion above implies that for J_0 the only holomorphic curves in the classes $s + df + R$ appear only for $R = 0$ (where these curves *define* what $R = 0$ means). \square

14.3. Rational ruled surfaces. Here let \mathbb{F}_n be the rational ruled surface whose fiber F , zero section S and infinity section E define homology classes with $S^2 = -E^2 = n$. We will compute some of the relative invariants GW^V with $V = S \cup E$ and with no constraint on the complex structure of the domain ($\kappa = 1$).

Fix a nonzero class $A = aS + bF$ and two sequences s, s' of multiplicities that describe the intersection with S and E respectively. The relative GW invariant with no constraint on the complex structure and k marked points lies in the homology of $X^k \times S^\ell \times E^{\ell'}$ with $\ell = \ell(s)$ and $\ell' = \ell(s')$. After imposing constraints $\alpha = (\alpha_1, \dots, \alpha_k) \in \mathbb{T}^*(X)$, we have

$$\mathrm{GW}_{A,g,s,s'}^{S,E}(\alpha) \in H_*(S^\ell) \otimes H_*(E^{\ell'})$$

where $S \cong E \cong \mathbb{P}^1$. Noting that the canonical class of \mathbb{F}_n is $K = -2S + (n-2)F$ and $\deg s' = E \cdot A = b$ and $\deg s = S \cdot A = b + na$, we have

$$\begin{aligned} \dim_{\mathbb{C}} \mathrm{GW}_{A,g,s,s'}^{S,E}(\alpha) &= (n+2)a + 2b + g - 1 \\ &\quad - (\deg s - \ell(s)) - (\deg s' - \ell(s')) + k - \deg \alpha \\ &= 2a + g - 1 + \ell + \ell' + k - \deg \alpha. \end{aligned}$$

But $S^\ell \times E^{\ell'}$ has complex dimension $\ell + \ell'$, so the pushforward of the moduli space represents zero in homology unless $\dim_{\mathbb{C}} \mathcal{M}_{g,k,s,s'}(\mathbb{F}_n, A) \leq \ell + \ell'$. Thus the invariant vanishes unless

$$(14.2) \quad 2a + g \leq 1 + \deg \alpha - k.$$

LEMMA 14.6. *The invariants $\mathrm{GW}_{A,g,s,s'}^{S,E}$ with no constraints except those on $V = S \cup E$ vanish except when $A = bF$, $g = 0$, and s and s' are single points with multiplicity $b > 0$. In that case*

$$\mathrm{GW}_{bF,0,s,s'}^{S,E} = \frac{1}{b} (S \otimes 1 + 1 \otimes E).$$

Moreover, the S -matrix in \mathbb{F}_n is the identity.

Proof. It suffices to show that the only contributions to GW from classes $A = aS + bF$ come from unstable rational domains with $a = 0$, i.e. from \mathbb{F}_n -trivial maps. Taking $\kappa = \alpha = 1$, (14.2) implies that $A = bF$ and $g = 0$ or 1. Moreover, because every bF curve intersects both E and S , we have $\ell + \ell' \geq 2$, and when $g = 0$, stability of the domain requires that $\ell + \ell' \geq 3$. In these cases the moduli space $\mathcal{M}_{g,s,s'}^V(\mathbb{F}_n, bF)$ is either empty or has dimension ≥ 2 .

Suppose that the moduli space is nonempty and the above stability conditions hold. Since E and S are copies of \mathbb{P}^1 , $H_*(S^\ell) \otimes H_*(E^{\ell'})$ is generated by point or $[\mathbb{P}^1]$ constraints. Then for each generic (J, ν) there are maps f in the moduli space whose images passes through at least two fixed points $p, q \in E \cup S$ in generic position. Take $(J, \nu) \rightarrow (J_0, 0)$ where J_0 is a complex structure with a holomorphic projection $\pi : \mathbb{F}_n \rightarrow \mathbb{P}^1$. In the limit we obtain a connected stable map f_0 through p and q with components representing $a_i S + b_i f$ such that $bF = \sum a_i S + b_i f$. But then each $a_i = 0$, so that the image of $\pi \circ f_0$ is a single point containing $\pi(p)$ and $\pi(q)$. This cannot happen for generic p, q .

Thus $\mathcal{M}_{g,s,s'}^V(\mathbb{F}_n, bF)$ consists of \mathbb{F}_n -trivial maps (cf. Definition 11.1) representing $A = bF$. Such maps contribute the identity to the S -matrix. \square

LEMMA 14.7. *Fix a point $p \in \mathbb{F}_n \setminus V$ with $V = E \cup S$. Then $\mathrm{GW}_{A,g,s,s'}^V(p)$ vanishes except in the following cases:*

- (i) $\mathrm{GW}_{bF,0,s,s'}^{S,E}(p) = 1$ when s and s' are single points with multiplicity $b > 0$.
- (ii) $\mathrm{GW}_{S+bF,0,s,s'}^{S,E}(p) = SV_s \times SV_{s'}$ whenever $\deg s = b$, $\deg s' = b + n$.

Proof. From (14.2) we have $\mathrm{GW}_{aS+bF,g,s,s'}^V(p) = 0$ unless $2a+g \leq 2$. Thus either (i) $a = 0$, or (ii) $a = 1$ and $g = 0$.

In case (i) each map contributing to the invariant represents bF , passes through p , and hits E and S . Hence $\dim \mathcal{M}_{g,s,s'}^V(\mathbb{F}_n, bF) = g - 1 + \ell(s) + \ell(s')$ with $\ell(s) + \ell(s') \geq 2$. The limiting argument used in Lemma 14.6 then shows that $\mathrm{GW}_{bF,g,s,s'}^V(p)$ vanishes unless $g = 0$ and $\ell(s) = \ell(s') = 1$. Thus s and s' are single points of multiplicity b , and the maps pass through p . Moving to the fibered complex structure, one sees that there is a unique such stable map for each $b > 0$. This gives (i).

In case (ii) the moduli space $\mathcal{M}_{0,s,s'}^V(\mathbb{F}_n, S+bF)$ has dimension $\ell(s) + \ell(s')$ and is empty unless $b \geq 0$. That means that the image of $\mathcal{M}_{0,s,s'}^V(\mathbb{F}_n, S+bF)$ is a multiple of $SV_s \times SV_{s'}$, so that the invariant vanishes except when all contact points on E and S are fixed. By the adjunction inequality, any irreducible curve C representing $S+bF$ is rational and embedded, so we can compute the invariant by intersections in $\mathbb{P}(H^0(\mathbb{F}_n, \mathcal{O}_{\mathbb{F}_n}(S+bF)))$ (the standard complex structure on \mathbb{F}_n is generic for these curves C because $h^1(C; \mathcal{O}(S+bF)|_C) = h^1(\mathbb{P}^1; \mathcal{O}(n+2b)) = 0$). But $h^0(\mathbb{F}_n, \mathcal{O}(S+bF)) = n+2+2b$, and each of the conditions imposed (including multiplicities) are linear conditions. Thus the number of curves representing $S+bF$ passing through a point p and meeting E and S at fixed contact points is 1. \square

14.4. *The rational elliptic surface.* As a final example we consider the rational elliptic surface E . Let f and s denote, respectively, the homology classes of a fiber and a fixed section of an elliptic fibration $E \rightarrow \mathbb{P}^1$. The following lemma describes the invariants relative to a fixed fiber F in the classes $A = s + df$ where d is an integer. In this case there are rim tori in $E \setminus F$, suggesting that one use the summed invariant $\overline{\mathrm{GW}}$ defined in the appendix. However, the lemma shows that the sum contains only one nonzero term (as happened in the last case of Lemma 14.5).

LEMMA 14.8. *The genus g relative and absolute invariants of E in the classes $s + df \in H_2(E)$ are related by:*

$$\mathrm{GW}_{s+df,g}(p^g) = \overline{\mathrm{GW}}_{s+df,g}^F(p^g; C_1(f)) = \mathrm{GW}_{s+df,g}^F(p^g; C_1(f))$$

where the second equality means that GW^F can be indexed by classes $s + df + R$ for rim tori R and these vanish whenever $R \neq 0$.

Proof. The first equality holds because generically all maps contributing to the absolute invariant are V -regular. That is true because if some component of a stable map is taken into $V = F$, then that component must have genus at least 1. But then the remaining components have genus less than g , so cannot pass through g generic points.

The second equality follows from a projection argument like the one used for Lemma 14.5. Consider a curve $C = \cup C_i$ representing $s + df$ which is holomorphic for a fibered complex structure J_0 on E . Since the projection to \mathbb{P}^1 gives a degree one composition, C must have a rational component C_0 that intersects each fiber in exactly one point, while the other components are multiple covers of fibers and so represent $df \in H_2(E \setminus F)$. Moreover, C_0 is an embedded section representing s . Since $s^2 = -1$ then C_0 must be the unique holomorphic curve in the class s . Thus the only curves in the class $s + df + R$ appear only for $R = 0$. \square

The invariants of Lemma 14.8 will be explicitly computed in Section 15.3.

14.5. *Rational relative invariants.* Counting rational curves requires only the $g = 0$ relative invariants and the corresponding S -matrix. The following two propositions show that these are particularly simple: the S -matrix is the identity and the relative invariant is the same as the absolute invariant in the absence of rim tori.

PROPOSITION 14.9. *When $g = 0$, $s = (1, \dots, 1)$ and $A \in H_2(X)$, the relative invariant (summed over rim tori as in (A.1)) equals the absolute invariant:*

$$\frac{1}{\ell(s)!} \overline{\text{GW}}_{A,0}^V(\alpha; C_s([V])) = \text{GW}_{A,0}(\alpha)$$

where $\alpha = (\alpha_1, \dots, \alpha_n) \in (H_*(X))^n$.

Proof. Fix a generic V -compatible pair (J, ν) . Recall that (J, ν) is generic for curves that have no components in V , and also its restriction to V gives a generic pair on V . However, for a curve entirely contained in V , even though (J, ν) is generic when the curve is considered in V , it might not be generic when the curve is considered in X .

For any genus g and ordered sequence s , consider the natural inclusion:

$$\mathcal{M}_{g,n,s}^V(X, A) \hookrightarrow \mathcal{M}_{g,n+\ell}(X, A)$$

where $A \in H_2(X)$ and where the left-hand moduli space involves a union over rim tori. When $s = (1, \dots, 1)$, any element in $\overline{\mathcal{M}}_{g,n}(X, A)$ that has no components in V lifts in $(A \cdot V)!$ ways to an element of $\mathcal{M}_{g,n,s}^V(X, A)$. We will show that for generic V -compatible (J, ν) , when $g = 0$ the contribution of the moduli space of curves with some components in V to the absolute invariant vanishes, and therefore the two invariants are equal.

For simplicity, start with the case when f has only one component, and this is entirely contained in V . Then $\ell(s) = A \cdot V = c_1(N_X V) \cdot A$ and the moduli space of such curves has

$$\begin{aligned} \dim \mathcal{M}_{V,A_0,g} &= -2K_V \cdot A_0 + (\dim V - 6)(1 - g) \\ &= \dim \mathcal{M}_{X,A,g}^V + 2g - 2 \end{aligned}$$

as in equation (6.4) of [IP4]. This means that for genus $g = 0$ the dimension of the moduli space of curves entirely contained in V is two less than the (virtual) dimension when considered as curves in X . Therefore if the virtual dimension in X is 0, there are no curves in V which could contribute. The general case of a curve with some components in V and some off V follows similarly. \square

PROPOSITION 14.10. *The $g = 0$ part of the S -matrix is the identity for any V and any normal bundle N .*

Proof. By (11.3) this statement is equivalent to showing that there is no contribution to the $g = 0$ GW-invariant coming from maps into \mathbb{P}_V which are not \mathbb{P}_V -trivial. Consider the 0 dimensional moduli space $\mathcal{M}_{\mathbb{P}_V, A, 0, s}^{V_0, V_\infty}(\gamma)$ constrained only along V_0 and V_∞ , such that the corresponding GW invariant is not zero. By Theorem 1.6 of [IP4] the same moduli space would be nonempty for the submersive structure associated with a generic ν on V (as defined before Lemma 11.3). Then each $f \in \mathcal{M}_R$ would project to a map f_V in $\mathcal{M}_{V, \pi_* A, 0, s}$ that passes through the γ constraints. But counting virtual dimensions using equation (6.4) of [IP4], we see that

$$\dim \mathcal{M}_{V, \pi_* A, 0, s}(\gamma) = \dim \mathcal{M}_{\mathbb{P}_V, A, 0, s}^{V_\infty, V_0}(\gamma) - 2 \operatorname{index} D_s^N = 2g - 2$$

is negative when $g = 0$, so this moduli space is empty for generic ν_V . \square

15. Applications of the sum formula

This last section presents three applications of the sum formula: (a) the Caporaso-Harris formula for the number of nodal curves in \mathbb{P}^2 , (b) the formula for the Hurwitz numbers counting branched covers of \mathbb{P}^1 , and (c) the formula for the number of rational curves representing a primitive homology class in the rational elliptic surface. These formulas have all recently been established using Gromov-Witten invariants in some guise. Here we show that all three follow rather easily from the symplectic sum formula.

15.1. The Caporaso-Harris formula. Our first application is a derivation of the Caporaso-Harris recursion formula for the number $N^{d, \delta}(\alpha, \beta)$ of curves in \mathbb{P}^2 of degree d with δ nodes, having a contact with a line L of order k at α_k fixed points, and at β_k moving points, for $k = 1, 2, \dots$ and passing through the appropriate number r of generic fixed points in the complement of L .

For this we consider the pair (\mathbb{P}, L) , which can be written as a symplectic connect sum:

$$(15.1) \quad (\mathbb{P}^2, L) \#_{L=E} (\mathbb{P}^1, E, L) = (\mathbb{P}^2, L)$$

where (\mathbb{P}_1, E, L) is the rational ruled surface with Euler class one with its zero section L and its infinity section E . We can then get a recursive formula for the GT invariant of (\mathbb{P}^2, L) by moving one point constraint pt to the \mathbb{P} side, and then using the symplectic sum formula.

The splitting (15.1) is along a sphere $V = E = L$, so there are no rim tori. The relative invariant therefore lies in the homology of SV and is invariant under the action of the subgroup of the symmetric group that switches the order of points of the same multiplicity. A basis for this homology is given by (A.4), where $\{\gamma_i\}$ with $\gamma^1 = p$ a point and $\gamma_2 = [\mathbb{P}^1]$ is a basis of $H_*(V)$.

To recover the notation of Caporaso and Harris, for each sequence $(m_{a,i})$, denote $\alpha_a = m_{a,1}$ and $\beta_a = m_{a,2}$, and let $\alpha = (\alpha_1, \alpha_2, \dots)$, $\beta = (\beta_1, \beta_2, \dots)$. Then $m!$ and $|m|$ correspond to $\alpha! = \prod_i \alpha_i$ and $|\alpha| = \prod_i i^{\alpha_i}$. With this change of coordinates,

$$N^{d,\delta}(\alpha, \beta) = \frac{1}{m!} \text{GT}_{\chi, dL, \mathbb{P}^2}^L(p^r, C_m)$$

where $\chi - 2\delta = -d(d-3)$ is the “embedded Euler characteristic”, $r = 3d + g - 1 - \sum \alpha_i - \sum (\beta_j - 1)$, and we are imposing no constraints on the complex structure of the curves. Similarly, let

$$N^{a,b,\chi}(\alpha', \beta'; p; \alpha, \beta) = \frac{1}{m!m'} \text{GT}_{\chi, aL+bF, \mathbb{P}}^{E,L}(C_m; p; C_{m'})$$

denote the number of curves of Euler characteristic χ in \mathbb{P} representing $aL + bF$ that have contact described by (α', β') along E , (α, β) along L and pass through an extra point $p \in \mathbb{P}$ (we prefer to label these numbers using χ rather than the number of nodes).

By Lemma 14.6 the S -matrix vanishes. The symplectic sum theorem then implies:

$$N^{d,\chi}(\alpha, \beta) = \sum |\alpha'| \cdot |\beta'| \cdot N^{d',\chi'}(\alpha', \beta') \cdot N^{d-d',b,\chi''}(\beta', \alpha'; p; \alpha, \beta)$$

where the sum is over all α', β' and all decompositions of (dL, χ) into $(d'L, \chi')$ and $((d-d')L + bF, \chi'')$ such that $\chi = \chi' + \chi'' - 2\ell(\alpha') - 2\ell(\beta')$. Combining Lemmas 14.6 and 14.7 we see that there are exactly two types of curves that contribute to the relative GW invariant $\text{GT}_{\mathbb{P}^1}^{E,L}(C_m; p; C_{m'})$ of \mathbb{P} with one fixed point p .

- (1) Several $g = 0$ unstable domain multiple covers of the fiber, one of them say of multiplicity k passing through the point p , corresponding to the situation $d' = d$ and

$$\beta' = \beta + \varepsilon_k; \alpha' = \alpha - \varepsilon_k$$

where ε_k is the sequence that has a 1 in position k and 0 everywhere else.

- (2) Several $g = 0$ unstable domain multiple covers of the fiber together with one $g = 0$ curve in the class $L + aF$ passing through p and having all contact points with E and L fixed say described by α'_0 and α_0 ; this corresponds to $d' = d - 1$ and the situation

$$\alpha = \alpha_0 + \alpha'; \quad \beta' = \alpha'_0 + \beta; \quad \text{equivalently } \beta' \geq \beta; \quad \alpha \geq \alpha'$$

In each situation above, the number of V -stable curves is 1. In the second case, note that there are $\binom{\alpha}{\alpha'}$ choices of α_0 and $\binom{\beta'}{\beta}$ of α'_0 . Moreover, for each \mathbb{P} -trivial curve its invariant combines with its corresponding multiplicities in $|m'|\ell(m')!$ to give 1. Therefore, the remaining multiplicity in case 1 is k , while in case 2, it is $|\alpha'_0| = |\beta' - \beta|$. Putting all these together, we get:

$$N^{d,\delta}(\alpha, \beta) = \sum k N^{d,\delta'}(\alpha - \varepsilon_k, \beta + \varepsilon_k) + \sum |\beta' - \beta| \binom{\alpha}{\alpha'} \binom{\beta'}{\beta} N^{d-1,\delta'}(\alpha', \beta')$$

where the last sum is over all $\beta' \geq \beta$, $\alpha' \geq \alpha$. This is exactly the Caporaso-Harris formula.

15.2. Hurwitz numbers. The method of Section 15.1. can also be applied for maps into \mathbb{P}^1 . In that case the symplectic sum formula yields the cut and paste formula for Hurwitz numbers that was first proved using combinatorics by Goulden, Jackson and Vainstein in [GJV] (recently Li-Zhao-Zheng [LZZ] have derived a similar formula using [LR]).

The Hurwitz number $N_{d,g}(\alpha)$ counts the number of genus g , degree d covers of \mathbb{P}^1 that have the branching pattern over a fixed point $p \in \mathbb{P}^1$ specified by the unordered partition α of d , while the remaining branch points are simple and located at fixed points in the target \mathbb{P}^1 . We can get at these numbers by regarding the pair (\mathbb{P}^1, p) as a symplectic sum:

$$(15.2) \quad (\mathbb{P}^1, p) = (\mathbb{P}^1, x) \#_{x=y} (\mathbb{P}^1, y, p).$$

We then get a recursive formula for the GW invariant of (\mathbb{P}^1, p) by moving one simple branch point b to the (\mathbb{P}^1, y, p) side and applying the symplectic sum formula.

In fact the Hurwitz numbers are the coefficients, in a specific basis, of the GW invariants of \mathbb{P}^1 relative to a point $V = p$. More precisely, each unordered partition $\alpha = (\alpha_1, \alpha_2, \dots)$ of d defines numbers $m_a = \#\{i \mid \alpha_i = a\}$; let C_m be the corresponding basis (A.4) (in this case the basis $\{\gamma_i\}$ of $H^*(V)$ has only one element). Then

$$N_{d,g}(\alpha) = \text{GW}_{\mathbb{P}^1, d, g}^p(b^r; C_m)$$

is the number of degree d , genus g covers that have the branching pattern over $p \in \mathbb{P}^1$ determined by α , and $r = 2d - 2 + 2g - 2 - \sum (a - 1)m_a$ other fixed,

distinct branch points in the target. (Note that the branching order is the order of contact to $p = V$.) The corresponding generating function (A.6) is

$$F = \mathrm{GW}_{\mathbb{P}^1}^p = \sum \mathrm{GW}_{\mathbb{P}^1, d, g}^p(b^r; C_m) \prod_a \frac{(z_a)^{m_a}}{m_a!} \frac{u^r}{r!} t^d \lambda^{2g-2}.$$

Now apply the symplectic sum formula to the decomposition (15.2), putting $r - 1$ branch points on the first copy of \mathbb{P}^1 and one on the second copy. Since there are no rim tori and the S -matrix vanishes by Lemma 14.1 we obtain

$$(15.3) \quad \mathrm{GW}_{d, g}^p(b^r; C_m) = \sum \frac{|m'|}{m'!} \cdot \mathrm{GT}_{d, \chi_1}^p(b^{r-1}; C_{m'}) \cdot \mathrm{GT}_{d, \chi_2}^{p, p}(C_{m'}; b; C_m)$$

where the sum is over all $m' = (m'_1, m'_2, \dots)$ and all χ_1, χ_2 such that $2 - 2g = \chi_1 + \chi_2 - 2\ell(m')$ and so that the attached domain is connected. But $\mathrm{GT} = \exp \mathrm{GW}$ and Lemma 14.2 implies that the only possibility for the last factor in (15.3) is a union of trivial spheres together with a degree a sphere constrained by C_a at one end and $C_{i, j}$ with $i + j = a$ at the other end (plus the branch point in the middle). Therefore there are only two possibilities for the other factor and for the partition α' corresponding to m' .

- (1) $\alpha = (i, j, \beta)$ and $\alpha' = (i + j, \beta)$ for some i, j, β , so the covering map has genus g and degree d .
- (2) $\alpha = (a, \beta)$ and $\alpha' = (i, j, \beta)$ with $a = i + j$. Then $\chi_1 = 2g - 4$ so the covering map is either genus $g - 1$ and degree d or genus g_1, g_2 of degrees d_1, d_2 .

The sum formula (15.3) can then be written as a relation for the generating function, namely

$$\partial_u F = \frac{1}{2} \sum_{i, j \geq 1} (ij \lambda^2 z_{i+j} [\partial_{z_i} \partial_{z_j} F + \partial_{z_i} F \cdot \partial_{z_j} F] + (i + j) z_i z_j \partial_{z_{i+j}} F).$$

This is the ‘cut-join’ operator equation $DF = 0$ of [GJV]. It clearly determines the Hurwitz numbers recursively. The same formula works to give the ‘Hurwitz numbers’ counting branched covers of higher genus curves.

15.3. Curves in the rational elliptic surface. We next consider the invariants of the rational elliptic surface $E \rightarrow \mathbb{P}^1$. Using the notation of Section 14.4, we focus on the classes $A = s + df$ where d is an integer. The numerical invariants $\mathrm{GW}_{A, g}(p^g)$ then count the number of connected genus g stable maps in the class $s + df$ through g generic points (with no constraints on the complex structure of the domain). For each g these define power series

$$F_g(t) = \sum_{d \geq 0} \mathrm{GW}_{d, g}(p^g) t^d t_s$$

where $t = t_f$. Recently, Bryan-Leung [BL] proved that

$$(15.4) \quad F_g(t) = F_0(t) [G'(t)]^g$$

with G as in (14.1) and

$$(15.5) \quad F_0(t) = t_s \left(\prod_d \frac{1}{1-t^d} \right)^{12}.$$

As mentioned in the introduction, this formula is related to the work of Yau-Zaslow [YZ] and to more general conjectures (such as those stated in [Go]) about counts of nodal curves in complex surfaces.

We will use our symplectic sum theorem to give a short proof of this formula, beginning with the $g = 0$ case. The proof is accomplished by relating $F(t)$ to the similar series of elliptic ($g = 1$) invariants

$$H(t) = \sum_{d \geq 0} \text{GW}_{d,1}(\tau_1[f^*]) t^d t_s$$

where $f^* \in H^2(E)$ is the Poincaré dual of the fiber class and where $\tau_1[f^*] = \text{ev}_s 1^*(f^*) \cup \psi_1$ is the corresponding ‘descendent constraint’ described at the end of Section 12.

We will compute H in two different ways. The first is based on the standard method of ‘splitting the domain’, which yields the following general facts for 4-manifolds.

LEMMA 15.1. *Let X be a symplectic 4-manifold with canonical class K .*

(a) *For $A = 0$ and $g = 1$ the GW invariant with a single constraint $B \in H^2(X)$ is*

$$(15.6) \quad \text{GW}_{0,1}(B) = \frac{1}{24} K \cdot B.$$

(b) *For any classes $A, f \in H_2(X)$ satisfying $A \cdot K = -1$,*

$$\begin{aligned} \text{GW}_{A,1}(\tau_1[f^*]) &= \frac{(f \cdot A)}{24} (A^2 + K \cdot A) \text{GW}_{A,0} \\ &+ \sum_{\substack{A_1 + A_2 = A \\ A_1 \neq 0, A_2 \neq 0}} (f \cdot A_2) (A_1 \cdot A_2) \text{GW}_{A_1,1} \text{GW}_{A_2,0}. \end{aligned}$$

Proof. (a) For $\nu = 0$, $\overline{\mathcal{M}}_{1,1}(X, 0)$ is the space $\overline{\mathcal{M}}_{1,1} \times X$ of ‘ghost tori’ $f : (T^2, j) \rightarrow X$ with $f(z) = p$ a constant map. At such f , the fiber of the obstruction bundle is $H^1(T^2, f^*TX) = H^1(T^2, \mathcal{O}) \otimes TX$. The dual of the bundle $H^1(T^2, \mathcal{O})$ over $\overline{\mathcal{M}}_{1,1}$ is the Hodge bundle. Since the first Chern number of the Hodge bundle is $1/24$, the Euler class of the obstruction bundle is

$$\chi(X)[\overline{\mathcal{M}}_{1,1}] \otimes 1 + \frac{1}{24} 1 \otimes K \in H_2(\overline{\mathcal{M}}_{1,1} \times X).$$

For $\nu \neq 0$, the (virtual) moduli space is the zeros of a generic section of the obstruction bundle, which consists of (i) maps from a torus with any complex structure to $\chi(X)$ specified points of X and (ii) maps from a torus of specified complex structure into some point on the canonical divisor. Generically, the images of the type (i) maps will miss the constraint surface representing B . The maps of type (ii) give the formula (15.6).

(b) The genus 1 topological recursive relation says

$$\begin{aligned} \mathrm{GW}_{A,1}(\tau[f^*]) &= \frac{1}{24} \mathrm{GW}_{A,0}(H_\alpha, H^\alpha, f) \\ &+ \sum_{A_1+A_2=A} \sum_{\alpha} \mathrm{GW}_{A_1,1}(H_\alpha) \mathrm{GW}_{A_2,0}(H^\alpha, f) \end{aligned}$$

where $\{H_\alpha\}$ and $\{H^\alpha\}$ are bases of $H^*(X)$ dual by the intersection form. But for $A \neq 0$ $\mathrm{GW}_{A,0}(H^\alpha, H^\beta, f)$ vanishes by dimension count unless H^α and H^β are two-dimensional, and then each A -curve hits a generic geometric representative of H_α at $H^\alpha \cdot A$ points counted with algebraic multiplicity. A dimension count also shows that the moduli spaces with $A_1 = A$ and $A_2 = 0$ are of the wrong dimension to contribute to the double sum above. Hence the expression above becomes

$$\begin{aligned} &\frac{1}{24} \sum (H^\alpha \cdot A)(H_\alpha \cdot A)(f \cdot A) \mathrm{GW}_{A,0} \\ &+ \sum_{\substack{A_1+A_2=A \\ A_1 \neq 0, A_2 \neq 0}} (H_\alpha \cdot A_1)(H^\alpha \cdot A_2)(f \cdot A_2) \mathrm{GW}_{A_1,1} \mathrm{GW}_{A_2,0} \end{aligned}$$

plus the term with $A_1 = 0$, which by (15.6) is

$$\frac{1}{24} (K \cdot H_\alpha) \mathrm{GW}_{A,0}(H^\alpha, f) = \frac{1}{24} (K \cdot H_\alpha) (A \cdot H^\alpha) (A \cdot f) \mathrm{GW}_{A,0}.$$

The lemma follows because $\sum (H_\alpha \cdot A_1)(H^\alpha \cdot A_2) = A_1 \cdot A_2$. \square

Taking X to be the rational elliptic surface E , we can apply Lemma 15.1 with $A = s + df$. Then $K = -f$, $A \cdot f = 1$ and $A^2 = 2d - 1$. The only possible decompositions are $A_1 = kf$ and $A_2 = s + (d - k)f$ so that:

$$\mathrm{GW}_{s+df,1}(\tau[f^*]) = \frac{d-1}{12} \mathrm{GW}_{s+df,0} + \sum_{k=1}^d k \mathrm{GW}_{kf,1} \mathrm{GW}_{s+(d-k)f,0}.$$

But for the rational elliptic surface the invariant $\mathrm{GW}_{kf,1}(s)$ is $\sigma(k)$ for $k > 0$ (since in \mathbb{P}^2 there is a unique cubic through 9 generic points). As in Section 4 of [IP1], for each k there are $\sigma(k)$ distinct k -fold covers of an elliptic curve with a marked point, all with positive sign. Because the marked point can go to any of $s \cdot kf = k$ points, this means that the unconstrained invariant is

$$\mathrm{GW}_{kf,1} = \sigma(k)/k \quad \text{for } k > 0.$$

It follows that

$$(15.7) \quad H(t) = \frac{1}{12} (t F'_0 - F_0) + F_0 \cdot G.$$

On the other hand, we can calculate $H(t)$ by splitting the target and using the symplectic sum theorem. Let $\mathbb{P} = T^2 \times S^2$, and let F denote both a fiber of the elliptic fibration E and a fixed torus $T^2 \times \{\text{pt}\}$ inside \mathbb{P} . We can apply the sum formula by writing $E = E \#_F \mathbb{P}$ for the class $A = s + df$ with the constraint on the \mathbb{P} side. Since $A \cdot F = 1$, the connected curves representing A split into the union of connected curves in E and in \mathbb{P} ; thus the symplectic sum formula applies for the GW (as well as the GT invariants).

If we have a genus 0 curve on the \mathbb{P} side in the class $s + d_1 F$, then by projecting onto the T^2 factor and noting that there are no maps from S^2 to T^2 of nonzero degree, we conclude that $d_1 = 0$. But the moduli space of genus 0 curves in \mathbb{P} representing s and passing through F is isomorphic to $F = T^2$, and moreover the relative cotangent bundle to them along F is isomorphic to the normal bundle to F . Thus,

$$\text{GW}_{s,0}(\tau_1[f^*]) = \text{GW}_{s,0}(f^* \cup f^*) = 0.$$

We conclude that there is no contribution from genus 0 curves on the \mathbb{P} side or in the neck (which is also a copy of \mathbb{P}). The same argument shows that there are no rational curves in F , and so the $g = 0$ absolute and relative invariants are the same.

With these observations, the only possibility is to have a genus 1 curve on the \mathbb{P} side, genus 0 on the E side, and no contribution from the neck. The symplectic sum formula thus says

$$\text{GW}_{d,1}(\tau_1[f^*]) = \sum_{d_1+d_2=d} \text{GW}_{s+d_1f,0}(E) \cdot \text{GW}_{s+d_2f,1}(\mathbb{P})(\tau_1[f^*]).$$

This last invariant can be computed by applying the topological recursive relation to $X = \mathbb{P}$ just as in Lemma 15.1:

$$\begin{aligned} \text{GW}_{s+df,1}(\tau_1[f^*]) &= \frac{d-1}{12} \text{GW}_{s+df,0} + \sum_{\substack{d_1+d_2=d \\ d_1 \neq 0, d_2 \neq 0}} d_1 \text{GW}_{d_1f,1} \text{GW}_{s+d_2f,0} \\ &\quad + d_2 \text{GW}_{s+d_1f,1} \text{GW}_{d_2f,0}. \end{aligned}$$

But the invariants of \mathbb{P} satisfy $\text{GW}_{df,0} = \text{GW}_{s+df,0} = 0$ for $d \neq 0$ by the projection argument above, while for $d \neq 0$, Lemma 14.4 gives $d_1 \text{GW}_{d_1f,1} = \text{GW}_{d_1f,1}(s) = 2\sigma(d_1)$. We therefore get

$$(15.8) \quad H = 2F_0 \cdot \left(G - \frac{1}{24} \right).$$

Combining (15.7) with (15.8) and noting that $F_0(0) = \text{GW}_{s,0} = 1 \cdot t_s$ we see that F_0 satisfies the ODE

$$t F'_0 = 12 G \cdot F_0$$

with $F_0(0) = 1 \cdot t_s$. Hence

$$F_0(t) = t_s \exp \left(12 \int G(t)/t \, dt \right).$$

Using the Taylor series of $\log(1 - t)$ and some elementary combinatorics, we obtain

$$F_0(t) = t_s \left(\prod_d \frac{1}{1 - t^d} \right)^{12}.$$

It remains to show (15.4) for $g > 0$. This case is different because for genus $g > 0$ the relative invariants are no longer equal to the absolute invariants. We start by fixing a fiber F of E and introducing two generating functions for the genus g relative invariant: one recording the number of curves passing through g points in $E \setminus F$, the other recording the number of curves passing through $g - 1$ points in $E \setminus F$ plus a fixed point on F :

$$\begin{aligned} F_g^V(f) &= \sum_d \overline{\text{GW}}_{s+df, g}^F(p^g; C_1(f)) t^d, \\ F_g^V(p) &= \sum_d \overline{\text{GW}}_{s+df, g}^F(p^{g-1}; C_1(p)) t^d. \end{aligned}$$

Using Lemma 14.8, we can relate the absolute and relative $g = 1$ invariants of E .

LEMMA 15.2. *For $X = E$, the absolute and relative $g = 1$ invariants in the classes $s + df \in H_2(E(1))$ are related by equations*

$$\begin{aligned} \text{(a)} \quad & F_g = F_g^V(p) + F_{g-1}^V(f) \cdot G', \\ \text{(b)} \quad & F_g = F_g^V(f), \\ \text{(c)} \quad & 0 = F_g^V(p) \cdot F_0 + F_{g-1} \cdot F_1^V(p). \end{aligned}$$

Proof. To prove (a), we again write $E = E \#_F \mathbb{P}$ where $\mathbb{P} = T^2 \times S^2$, and put $g - 1$ points on E and the remaining point on \mathbb{P} . If we start with a class $s + df$, the only possible decompositions are $s + af$ and $s + bf$ where $d = a + b$. Since there are $g - 1$ points on the E side, the genus $g_1 \geq g - 1$. There are two possibilities:

- (1) Genus g in class $s + df$ on E and genus 0 in class $s + bf$ on \mathbb{P} . But that forces $b = 0$ so that $a = d$.
- (2) Genus $g - 1$ in class $s + df$ on E and genus 1 in class $s + bf$ on \mathbb{P} .

Putting these together gives (a). Relation (b) is a reformulation of Lemma 14.8.

Relation (c) is seen by applying the symplectic sum formula to the sum $K3 = X_1 \#_F X_2$ where $X_1 = X_2 = E$ (the elliptic surface $K3 = E(2)$ is the

fiber sum of $E = E(1)$ with itself). Because a generic complex structure on $K3$ admits no holomorphic curves, then all relative and absolute invariants of $K3$ vanish. In particular, the genus g invariants through $g - 1$ points in the class $[s + df] \in H_2(K3)/\mathcal{R}$ vanish, where \mathcal{R} is the set of rim tori corresponding to the gluing $K3 = E \#_F E$.

Now, for any $g \geq 1$, put all the $g - 1$ points on X_1 and split as above. A dimension count shows that the genus of the curve on X_1 must be at least $g - 1$, so the only possible decompositions are:

- (1) A genus g curve in the class $s + d_1 F$ on X_1 and a genus 0 curve on X_2 in the class $s + d_2 f$, $d = d_1 + d_2$;
- (2) A genus $g - 1$ curve in the class $s + d_1 F$ on X_1 and a genus 1 curve on X_2 in the class $s + d_2 f$, $d = d_1 + d_2$.

The symplectic sum formula then gives $0 = F_g^V(p) \cdot F_0 + F_{g-1}^V(f) \cdot F_1^V(p)$, which simplifies by (b). \square

Formula (15.4) follows quickly from Lemma 15.2. Taking $g = 1$ in Lemma 15.2c and factoring out $F_0 \neq 0$ yields $F_1^V(p) = 0$. Putting that in Lemma 15.2a and again noting that $F_0 \neq 0$, we see that $F_g^V(p) = 0$ for all $g > 0$. Parts (a) and (b) of Lemma 15.2 then reduce to

$$F_g = F_{g-1} \cdot tG'$$

which gives (15.4) by induction.

16. Appendix: Expansions of relative GT invariants

The Gromov-Witten invariants described in Section 1 are homology elements — the pushforward of the compactified moduli space under (1.18). These can be assembled into a power series (1.24) with coefficients in homology. Often, however, it is convenient to write the GW and GT invariants as power series whose coefficients are *numbers*, preferably numbers with clear geometric interpretations. This appendix describes how this can be done for the relative GT invariants which appear in the symplectic sum formula.

Such series expansions are easiest when we can ignore the complications caused by the covering (1.20), replacing the space $\mathcal{H}_{X,A,s}^V$ by the more easily understood space $V_s \cong V^{\ell(s)}$. This can be done by pushing the homology class of the invariant down under the projection ε of (1.20), obtaining a ‘summed’ GW series

$$(A.1) \quad \overline{\text{GW}}_X^V = \varepsilon_* (\text{GW}_X^V) = \sum_{A \in H_2(X)} \overline{\text{GW}}_{X,A}^V t_A$$

whose coefficients are homology classes in $\sqcup_s V_s$. This is a less refined invariant, but has the advantage that its coefficients become numbers after choosing a basis of $H^*(V)$.

Of course (A.1) is the same as the original GW invariant when the set \mathcal{R} of (1.19) vanishes, that is, when there are no rim tori. This occurs whenever $H_1(V) = 0$ or more generally when every rim torus represents zero in $H_2(X \setminus V)$. We will describe the numerical expansion under that assumption; the same discussion applies in general to $\overline{\text{GW}}_X^V$.

When there are no rim tori, $\mathcal{H}_{X,A}^V$ is the union of those $V_s \cong V^{\ell(s)}$ with $\deg s = A \cdot V$. Fix a basis γ_i of $H_*(V; \mathbb{Q})$. Then a basis for the tensor algebra on $\mathbb{N} \times H^*(V)$ is given by elements of the form

$$(A.2) \quad C_{s,I} = C_{s_1, \gamma_{i_1}} \otimes \cdots \otimes C_{s_\ell, \gamma_{i_\ell}}$$

where $s_i \geq 1$ are integers. Let $\{C_{s,I}^\vee\}$ denote the dual basis. When $\kappa \in H^*(\overline{\mathcal{M}})$ and $\alpha \in \mathbb{T}^*(X)$, we can expand

$$(A.3) \quad \text{GT}_X^V(\kappa, \alpha) = \sum_{s,I} \frac{1}{\ell(s)!} \text{GT}_{X,A,\chi}^V(\kappa, \alpha; C_{s,I}) C_{s,I}^\vee t_A \lambda^{-\chi}.$$

The coefficients in (A.3) have a direct geometric interpretation. Choose generic pseudomanifolds $K \subset \overline{\mathcal{M}}_{g,n}$, $A_i \subset X$, and $\Gamma_j \subset V$ representing the Poincaré duals of κ , α , and the γ_j in their respective spaces. Then $\text{GT}_{X,A,\chi}^V(\kappa, \alpha; C_{s,I})$ is the oriented number of genus g (J, ν) -holomorphic, V -regular maps $f : C \rightarrow X$ with $C \in K$, $f(p_i) \in A_i$ and having a contact of order s_j with V along Γ_j at points x_j . Because of that interpretation, the $C_{s,I}$ are called “contact constraints”.

While for the analysis it is important to work with *ordered* sequences s , in applications it is more convenient to forget the ordering. The symmetries of the GW invariants allow us to replace the basis (A.2) with the one having elements of the form

$$(A.4) \quad \mathbf{C}_m = \prod_{a,i} (C_{a,\gamma_i})^{m_{a,i}}$$

where $m = (m_{a,i})$ is a finite sequence of nonnegative integers. Generalizing (1.16), we write

$$(A.5) \quad |m| = \prod_{a,i} a^{m_{a,i}}, \quad m! = \prod_{a,i} m_{a,i}!, \quad \ell(m) = \sum_{a,i} m_{a,i}, \quad \deg m = \sum_{a,i} a \cdot m_{a,i}.$$

Let $\{z_{a,i}\}$ denote the dual basis; these generate a (super-)polynomial algebra with the relations $z_{a,i} z_{b,j} = \pm z_{b,j} z_{a,i}$ where the sign is $+$ if and only if $(\deg \gamma_i)(\deg \gamma_j)$ is even. Then the generating series of the relative GT invariant

is

$$(A.6) \quad \text{GT}_X^V(\kappa, \alpha) = \sum_{A,g} \sum_m \text{GT}_{X,A,\chi}^V(\kappa, \alpha; \mathbf{C}_m) \prod_{a,i} \frac{(z_{a,i})^{m_{a,i}}}{m_{a,i}!} t_A \lambda^{-\chi}$$

where the sum is over all sequences $m = (m_{a,i})$ as above and where the coefficients $\text{GT}_{X,A,\chi}^V(\kappa, \alpha; \mathbf{C}_m)$ vanish unless $\deg m = A \cdot V$. Formally, this generating series (A.6) is given by

$$\text{GT}_X^V(\kappa, \alpha) = \sum_{A,g} \text{GT}_{X,A,g}^V \left(\kappa, \alpha; \exp \left(\sum_{a,i} C_{a,\gamma_i} z_{a,i} \right) \right) t_A \lambda^{-\chi}.$$

UNIVERSITY OF WISCONSIN, MADISON, WI

E-mail address: ionel@math.wisc.edu

MICHIGAN STATE UNIVERSITY, EAST LANSING, MI

E-mail address: parker@math.michigan.edu

REFERENCES

- [BL] J. BRYAN and N.-C. C. LEUNG, The enumerative geometry of $K3$ surfaces and modular forms, *J. Amer. Math. Soc.* **13** (2000), 371–410.
- [CH] L. CAPORASO and J. HARRIS, Counting plane curves in any genus, *Invent. Math.* **131** (1998), 345–392.
- [DK] S. DONALDSON and P. KRONHEIMER, *The Geometry of Four-Manifolds*, *Oxford Math. Monographs*, Oxford Univ. Press, New York, 1990.
- [EGH] Y. ELIASBERG, A. GIVENTAL, and H. HOFER, Introduction to symplectic field theory, in *GAGA 2000* (Tel Aviv, 1999), *Geom. Funct. Anals.* **2000** (Special Vol., Part II), 560–673.
- [F] A. FLOER, The unregularized gradient flow of the symplectic action, *Comm. Pure Appl. Math.* **41** (1988), 775–813.
- [GJV] I. P. GOULDEN, D. M. JACKSON, and A. VAINSTEIN, The number of ramified coverings of the sphere by the torus and surfaces of higher genera, *Ann. Comb.* **4** (2000), 27–46.
- [Gf] R. GOMPF, A new construction of symplectic manifolds, *Ann. of Math.* **142** (1995), 527–595.
- [Go] L. GÖTTSCHE, A conjectural generating function for numbers of curves on surfaces, *Comm. Math. Phys.* **196** (1998), 523–533.
- [I] E. IONEL, Topological recursive relations in $H^{2g}(\mathcal{M}_{g,n})$, *Invent. Math.* **148** (2002), 627–658.
- [IP1] E. IONEL and T. PARKER, The Gromov invariants of Ruan-Tian and Taubes, *Math. Res. Lett.* **4** (1997), 521–532.
- [IP2] ———, Gromov invariants and symplectic maps, *Math. Ann.* **314** (1999), 127–158.
- [IP3] ———, Gromov-Witten invariants of symplectic sums, *Math. Res. Lett.* **5** (1998), 563–576.
- [IP4] ———, Relative Gromov-Witten invariants, *Ann. of Math.* **157** (1993), 45–96.

- [IS] S. IVASHKOVICH and V. SHEVCHISHIN, Complex curves in almost-complex manifolds and meromorphic hulls, preprint; <http://front.math.ucdavis.edu/math.CV/9912046>.
- [L] E. LERMAN, Symplectic cuts, *Math. Res. Lett.* **2** (1995) 247–258.
- [LR] A.-M. LI and Y. RUAN, Symplectic surgery and Gromov-Witten invariants of Calabi-Yau 3-folds I, *Invent. Math.* **145** (2001), 151–218.
- [LZZ] A.-M. LI, G. ZHAO, and Q. ZHENG, The number of ramified covering of a Riemann surface by Riemann surface, *Comm. Math. Phys.* **213** (2000), 685–696.
- [LT] J. LI and G. TIAN, Virtual moduli cycles and Gromov-Witten invariants of general symplectic manifolds, in *Topics in Symplectic 4-Manifolds* (Irvine, CA, 1996), 47–83, *First Internat. Press Lect. Ser. I*, Internat. Press, Cambridge, MA, 1998.
- [Loc] R. LOCKHART, Fredholm, Hodge and Liouville theorems on noncompact manifolds, *Trans. Amer. Math. Soc.* **301** (1987), 1–35.
- [Loo] E. LOOIJENGA, Smooth Deligne-Mumford compactifications by means of Prym level structures, *J. Algebraic Geom.* **3** (1994), 283–293.
- [Ma] H. MASUR, The extension of the Weil-Petersson metric to the boundary of Teichmüller space, *Duke Math J.* **43** (1976), 623–635.
- [MS] D. McDUFF and D. SALAMON, *J-Holomorphic Curves and Quantum Cohomology*, *University Lecture Series* **6**, A. M. S., Providence, RI (1994).
- [MW] J. MCCARTHY and J. WOLFSON, Symplectic normal connect sum, *Topology* **33** (1994) 729–764.
- [PW] T. PARKER and J. WOLFSON, Pseudo-holomorphic maps and bubble trees, *J. Geom. Anal.* **3** (1993), 63–98.
- [RT1] Y. RUAN and G. TIAN, A mathematical theory of quantum cohomology, *J. Differential Geom.* **42** (1995), 259–367.
- [RT2] ———, Higher genus symplectic invariants and sigma model coupled with gravity, *Invent. Math.* **130** (1997), 455–516.
- [T] C. H. TAUBES, Counting pseudo-holomorphic curves in dimension 4, *J. Differential Geom.* **44** (1996), 818–893.
- [V] R. VAKIL, Counting curves of any genus on rational ruled surfaces, preprint; <http://front.math.ucdavis.edu/alg-geom/9709003>.
- [W] W. WIECZOREK, The Donaldson invariant and embedded 2-spheres, *J. reine angew. Math. Geom.* **489** (1997), 15–51.
- [YZ] S. T. YAU and E. ZASLOW, BPS states, string duality, and nodal curves on $K3$, *Nuclear Phys. B* **471** (1996), 503–512.

(Received October 25, 2000)

(Revised July 22, 2001)