Tropical Fukaya algebras

Sushmita Venugopalan and Chris Woodward

Abstract. We introduce a tropical version of the Fukaya algebra of a Lagrangian submanifold and use it to show that tropical Lagrangian tori are weakly unobstructed. A multiple cut operation on a symplectic manifold produces a collection of cut spaces each containing relative normal crossing divisors. As in works of Ionel [50] and Brett Parker [68, 69], tropical graphs arise as large-scale behavior of pseudoholomorphic disks under a multiple cut operation. Given a Lagrangian submanifold in the complement of the relative divisors in one of the cut spaces, the structure maps of the *broken Fukaya algebra* count broken disks associated to rigid tropical graphs. We introduce a further degeneration of the matching conditions (similar in spirit to Bourgeois' version of symplectic field theory [11]) which results in a *tropical Fukaya algebra* whose structure maps are, in good cases, sums of products over vertices of tropical graphs. We show the tropical Fukaya algebra is homotopy equivalent to the original Fukaya algebra. In the case of toric Lagrangians contained in a toric component of the degeneration, an invariance argument implies the existence of projective Maurer-Cartan solutions. We also give various computations of potentials, such as those of Lagrangians in cubic surfaces and flag varieties.

Contents

CHAPTER 1

Statement of results

In this paper we study the behavior of holomorphic curves under a multiple symplectic cut. In particular, we study the behavior of the Fukaya algebra, which is an invariant associated to a Lagrangian submanifold of a symplectic manifold, under the multiple cut operation. In the first part of the paper, we prove a homotopy equivalence (in the A_∞ sense) between two versions of the Fukaya algebra: the standard version defined on the symplectic manifold and the 'broken' version obtained by cutting the manifold. To simplify matters, we assume that the cuts are disjoint from the Lagrangian. In the second part of the paper we degenerate the matching conditions at nodes of broken maps to a split form, so that the curve counts in the total space can be expressed as sums of products of curve counts in the pieces of the cut manifold. As applications of our results we give a new proof of the unobstructedness of toric Lagrangians and a computation of the potential of toric Lagrangians in the cubic surface. In this introductory chapter we give a low-tech tour of the paper, interwoven with motivations, context, and also, limitations of our results.

1.1. Multiple cuts

The *cut* operation introduced by Lerman [**54**] cuts a symplectic manifold along a regular level set of a moment map for a circle action and quotients the boundary by the circle action to produce a smooth symplectic manifold. The inverse operation is called a *symplectic sum*. The *multiple cut* is a generalization where the symplectic manifold is cut along intersecting hypersurfaces. The set-up for a multiple cut is on the lines of a tropical manifold with a polyhedral decomposition.

Definition 1.1. (Polyhedral decomposition) Let $n > 0$ be an integer and $T \simeq (S^1)^n$ a torus with Lie algebra $\mathfrak{t} \cong \mathbb{R}^n$. A *simplicial polyhedral decomposition* of $\mathfrak{t}^\vee$ is a collection

$$\mathcal{P} = \{P \subset \mathfrak{t}^\vee\}$$

of simple polytopes[1] such that

(a) (Covering property) the interiors P° of the polytopes $P \in \mathcal{P}$ cover $\mathfrak{t}^\vee$; that is, $\mathfrak{t}^\vee = \cup_{P \in \mathcal{P}} P^\circ$;

(b) (Face property) for any $\sigma_1, \ldots, \sigma_k \in \mathcal{P}$, $\sigma_1 \cap \cdots \cap \sigma_k$ is a polytope in $\mathcal{P}$ and it is a face of each of the polytopes $\sigma_1, \ldots, \sigma_k$.

[1] Polytopes are intersections of half-planes, see Definition 3.2. They are closed but not necessarily compact.

Any polytope P in a simplicial polyhedral decomposition $\mathcal{P}$ corresponds to a sub-torus

$$T_P \subseteq T, \quad \text{defined by} \quad \mathfrak{t}_P := \operatorname{ann}(TP). \tag{1.1}$$

Thus P and T_P have complementary dimensions, and $T_P = \{\mathrm{Id}\}$ for top-dimensional polytopes $P \in \mathcal{P}$.

DEFINITION 1.2. (Tropical manifold) A *tropical manifold* $(X, \mathcal{P}, \Phi)$ consists of a

(a) a simplicial polyhedral decomposition $\mathcal{P}$ of a polytope $\mathfrak{t}^\vee$,
(b) a compact symplectic manifold (X, ω) with a *tropical moment map*

$$\Phi : X \to \mathfrak{t}^\vee,$$

such for any $P \in \mathcal{P}$ there is a neighborhood $U_P \subset X$ of $\Phi^{-1}(P)$ on which the projection

$$U_P \to \mathfrak{t}^\vee \to \mathfrak{t}_P^\vee$$

is a T_P-moment map.

Given a tropical manifold X the output of a multiple cut is a collection of *cut spaces*

$$X_P := \Phi^{-1}(P)/\sim \tag{1.2}$$

where $\Phi^{-1}(P)$ is a manifold with corners and the equivalence relation quotients any codimension one boundary $\Phi^{-1}(Q) \subset \Phi^{-1}(P)$, $\operatorname{codim}_P(Q) = 1$ by the action of $S^1 \simeq T_Q/T_P$. The spaces X_P, $P \in \mathcal{P}$ are orbifolds, whose local structure is given by an iterative application of Lerman's construction [**54**]. For any pair of facets Q_1, Q_2 of a polytope P, X_{Q_1} and X_{Q_2} are embedded as divisors in X_P, called *relative divisors*, that intersect each other normally along $X_{Q_1 \cap Q_2}$. See Figure 1.4. In the special case of a single cut shown in Figure 1.1, $n = 1$, $\mathcal{P}$ partitions $\mathfrak{t}^\vee \simeq \mathbb{R}$ into two semi-infinite lines intersecting at a point.

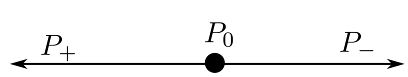

FIGURE 1.1. A single cut

Our definition of tropical manifolds is similar to Gross-Siebert [**44**]. A limitation of our set-up is that we require integral affine singularities to lie away from the cut locus. We expect that this requirement can be weakened by replacing a single cut passing through a singularity by two parallel cuts straddling the singularity, see below in Section 2.4. On the other hand, our definitions include manifolds that are not integral affine, because we require a toric structure only in the neighborhood of the cut locus, rather than globally.

1.2. Neck-stretching

To build a correspondence between curves in the manifold and curves in the cut spaces, we construct a sequence of 'large almost complex structures', which degenerate to almost complex structures in the cut spaces. We call this limiting process *neck stretching*, because it amounts to enlarging the complex structure in the neighborhoods of the cut loci, which we call *necks*.

As an example of neck-stretching, consider a single cut. Let $\Phi : X \to \mathbb{R}$ be a moment map for a Hamiltonian circle action, so that the zero level set

$$Z := \Phi^{-1}(0)$$

is a separating hypersurface with a tubular neighborhood

$$U_Z \subset X, \quad U_Z \cong Z \times I$$

where $I \subset \mathbb{R}$ is an interval. By a sequence of neck-stretched almost complex structures $J^\nu \in \mathcal{J}(X)$ we mean a sequence so that the fibers of the projection $U_Z \to Z/S^1$ are holomorphic cylinders $[\frac{-\nu}{2}, \frac{\nu}{2}] \times S^1$ of increasing length, see Figure 1.2. To define

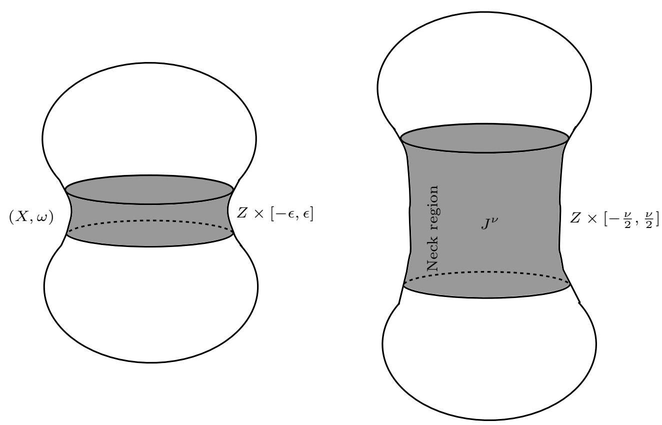

FIGURE 1.2. Neck-stretching for a single cut

neck-stretching for a multiple cut, we need an additional datum of a dual complex, which encodes the proportion in which the neck is stretched in different directions. The dual complex consists of a complementary dimensional polytope denoted $P^\vee$ for every $P \in \mathcal{P}$.

EXAMPLE 1.3. We give some examples. For the single cut in Figure 1.1, the dual complex shown in Figure 1.3 is a line segment. The simplest example of a multiple

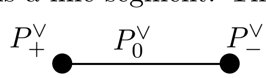

FIGURE 1.3. A dual complex

cut consists of two single cuts along hypersurfaces

$$Z_1 := \Phi^{-1}(P_{12} \cup P_{30}), \quad Z_2 := \Phi^{-1}(P_{01} \cup P_{23}),$$

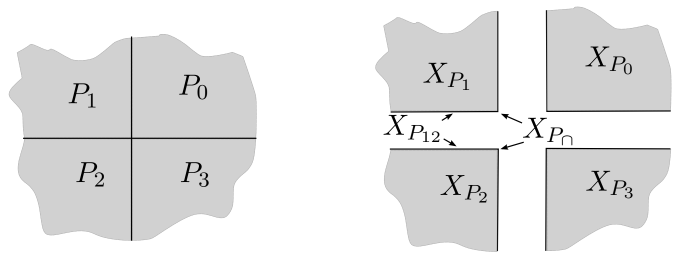

FIGURE 1.4. A multiple cut of $\mathbb{R}^2$

intersecting along $\Phi^{-1}(P_\cap)$ which is a $(S^1)^2$ bundle over $X_{P_\cap}$. See Figure 1.4. The dual complex in this case is a rectangle with side lengths, say l_1, l_2. In the neck-stretched almost complex manifold (X, J^ν), a neighborhood $U_{P_\cap}$ of $\Phi^{-1}(P_\cap)$ fibers over $X_{P_\cap}$. The fibers of $U_{P_\cap} \to X_{P_\cap}$ are holomorphic to the product of cylinders $((\frac{-\nu l_1}{2}, \frac{\nu l_1}{2}) \times S^1) \times ((\frac{-\nu l_2}{2}, \frac{\nu l_2}{2}) \times S^1)$, see Figure 1.6. The ratio $\frac{l_1}{l_2}$ is not required to be rational.

Curve counts will depend on the choice of dual complex, but as with other choices (such as tamed almost complex structures), the curve counts corresponding to any two dual complexes will produce Fukaya algebras that are A_∞ -homotopy equivalent.

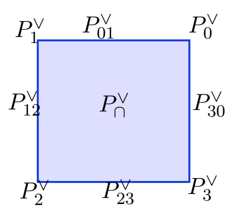

FIGURE 1.5. Dual complex $B^\vee$ for the cut in Figure 1.4

A *broken manifold* corresponding to a multiple cut $(X, \mathcal{P})$ is a disjoint union of cut spaces and their thickenings

$$\mathfrak{X} = \bigsqcup_{P \in \mathcal{P}} X_{\overline{P}}.$$

For top-dimensional polytopes $P \in \mathcal{P}$, $X_{\overline{P}} := X_P$ is the cut space from (1.2), and for lower dimensional polytopes P, $X_{\overline{P}}$ is a thickening of the cut space X_P and is called a *neck piece*. Broken maps lie in the complements of relative divisors, which are denoted by

$$X_P^\square := X_P - \text{relative divisors}, \quad X_{\overline{P}}^\square := X_{\overline{P}} - \text{relative divisors},$$

and these spaces are equipped with almost complex structures that are cylindrical on the ends.

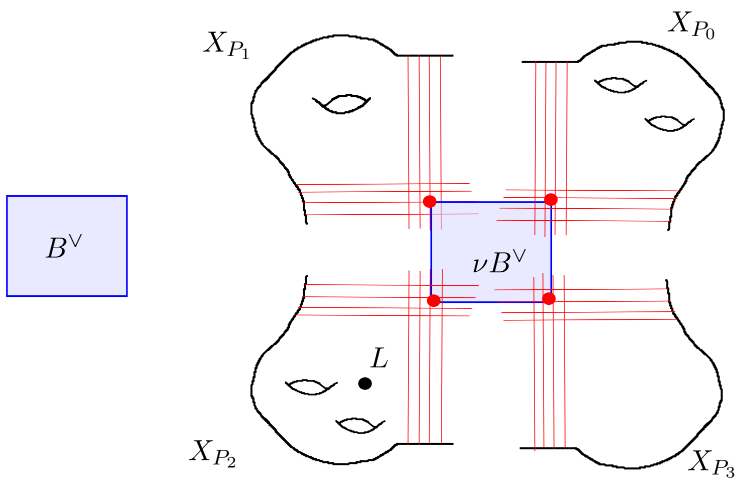

FIGURE 1.6. Neck-stretched manifold (X, J_ν) for the double cut in Figure 1.4 using the dual complex $B^\vee$.

For example, in the case of a single cut as in (1.1), the neck piece denoted by $X^\square_{\overline{P}_0}$ arises as the limit of the neighborhoods of the separating hypersurface and it corresponds to a zero-dimensional polytope $P_0 \in \mathcal{P}$. The space $X^\square_{\overline{P}_0}$ is a $(\mathbb{R} \times S^1)$-bundle over the relative divisor X_{P_0}, and its compactification $X_{\overline{P}_0}$ is a $\mathbb{P}^1$-bundle over X_{P_0}. We view $X^\square_{\overline{P}_0}$ and $X_{\overline{P}_0}$ as thickenings of the divisor X_{P_0}.

Analogously for multiple cuts, for a polytope $P \in \mathcal{P}$, the neck piece $X^\square_{\overline{P}}$ is a torus bundle

$$T_{P,\mathbb{C}} \to X^\square_{\overline{P}} \to X^\square_P$$

over the cut space $X^\square_P$, whose fiber $T_{P,\mathbb{C}}$ is the complexification of the compact torus T_P, which is part of the data of the tropical manifold X, see (1.1). See Figure 1.7. The compactification of $X^\square_{\overline{P}}$ is a fibration $X_{\overline{P}} \to X_P$ with toric orbifold fibers. However, we do not encounter the technical issues associated with moduli spaces of holomorphic curves in orbifolds, since the pieces of the broken map have target the non-compact space $X^\square_{\overline{P}}$ and the compactification is just a useful technical tool in some places.

We will study the behavior of holomorphic disks bounding a Lagrangian submanifold under a family of neck-stretching almost complex submanifolds. To simplify the situation, which is already quite involved, we assume that the Lagrangian submanifold $L \subset (X, \omega)$ is disjoint from the cuts. The Lagrangian therefore descends to a Lagrangian submanifold, also denoted by L, in a cut space $X^\square_P$ for a top-dimensional polytope P.

1.3. Broken maps

A broken map is modelled on a graph and consists of a 'map part' and a 'tropical part'. For the map part, the domain is the normalization of a nodal curve C whose irreducible components $C_v, v \in \mathrm{Vert}(\Gamma)$ correspond to vertices of a graph Γ, and

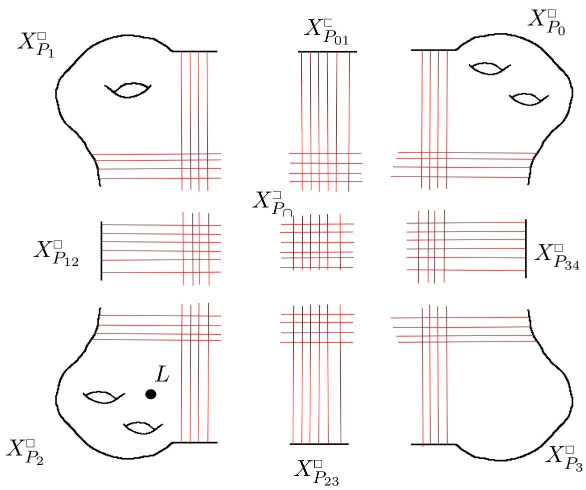

FIGURE 1.7. Broken manifold arising from the neck-stretching in Figure 1.6.

nodes $w(e) \in C$ correspond to edges $e \in \text{Edge}(\Gamma)$ of Γ. The broken map is a collection of holomorphic maps on punctured curves

$$u_v : C_v^\circ \to X^{\square}_{\overline{P}(v)}, v \in \text{Vert}(\Gamma),$$

satisfying certain matching conditions (explained later in the paragraph) on the lifts of nodal points. Each of the domain components C_v° is an irreducible curve component $C_v \subset C$ (possibly with boundary) punctured at interior nodal points, that is,

$$C_v^\circ := C_v \backslash \{\text{interior nodes}\},$$

the target space $X^{\square}_{\overline{P}(v)}$ is a piece of the broken manifold corresponding to the polytope

$$P(v) \in \mathcal{P},$$

and the punctures in the domain are removable singularities when u_v is viewed as a map to the compactification $X_{\overline{P}(v)}$.[2] Thus the map u can be evaluated at the nodal lift $w_+ \in C_{v_+}$, and the point $u(w_+)$ lies on the intersection

$$Y_+ = \cap_i D_i^+$$

of a collection of relative divisors D_i^+ in $X_{P(v_+)}$. Assume D_i^+ is the fixed point set of a one-dimensional torus generated by $\mu_i \in \mathfrak{t}$, and n_i is the intersection multiplicity of the map u_{v_+} at w_+. Then the sum

$$\mathcal{T}(w_+) := \sum n_i \mu_i \in \mathfrak{t}_{\mathbb{Z}}$$

[2]Here for the sake of exposition we assume that $X_{\overline{P}(v)}$ does not have orbifold singularities, the general case is treated in Chapter 4, where the quantity $\mathcal{T}(e)$ is defined using the asymptotic behavior of the map near the nodal point.

lies in the integer lattice $\mathfrak{t}_{\mathbb{Z}} \subset \mathfrak{t}$. Similarly we define $\mathcal{T}(w_-) \in \mathfrak{t}_{\mathbb{Z}}$ for the lift w_-. For any holomorphic coordinate $z_\pm$ on a neighborhood of $w_\pm$ in $C_{v_\pm}$, the limit

$$\lim_{z_\pm \to 0} z_\pm^{-\mathcal{T}(w_\pm)} u$$

exists. (See Section 4.3 for more details.) The *matching condition* at the node w says that

- $\mathcal{T}(w_+) = \mathcal{T}(w_-)$,
- and there exist holomorphic coordinates z_+, z_- such that

$$\lim_{z_+ \to 0} z_+^{-\mathcal{T}(w_+)} u = \lim_{z_- \to 0} z_-^{-\mathcal{T}(w_-)} u. \tag{1.3}$$

The quantity $\mathcal{T}(w_+) = \mathcal{T}(w_-)$ is called the *slope of the node* w or the *slope of the edge* e in Γ corresponding to the node w; and is denoted by

$$\mathcal{T}(e). \tag{1.4}$$

The quantities in the left-hand side and rhs of (1.3) are called the *tropical evaluations* at w_+ and w_- respectively. An equivalent formulation of the matching condition says that for any node the evaluations of the lifts are equal modulo the action of the torus generated by the slope of the node: For a node $w \in C$ corresponding to an edge e, the evaluations of the lifts $w_\pm$ lie in $X^{\square}_{\overline{P}(e)}$ which has the structure of a $T_{P(e),\mathbb{C}}$-bundle, where

$$P(e) := P(v_+) \cap P(v_-) \in \mathcal{P},$$

is the polytope assigned to the edge $e = (v_+, v_-)$. The slope of the node $\mathcal{T}(w)$ lies in the lattice $\mathfrak{t}_{P(e),\mathbb{Z}}$ and generates a one-dimensional torus $T_{\mathcal{T}(e),\mathbb{C}}$. The matching condition is then that the images of the evaluations match in the base of the $T_{\mathcal{T}(e),\mathbb{C}}$-fibration:

$$(\pi^{\perp}_{\mathcal{T}(e)} \circ u)(z_+) = (\pi^{\perp}_{\mathcal{T}(e)} \circ u)(z_-) \in X^{\square}_{\overline{P}(e)}/T_{\mathcal{T}(e),\mathbb{C}}, \tag{1.5}$$

where $\pi^{\perp}_{\mathcal{T}(e)} : X^{\square}_{\overline{P}(e)} \to X^{\square}_{\overline{P}(e)}/T_{\mathcal{T}(e),\mathbb{C}}$ is the projection to the quotient. The quantities in the left-hand side and rhs of (1.5) are called *projected tropical evaluations*. Thus the space $X_{\overline{P}(e),\mathbb{C}}/T_{\mathcal{T}(e),\mathbb{C}}$ in which the matching condition of an edge e is defined is dependent on e. The matching condition is simpler in the special case of a single cut, because the space $X_{\overline{P}(e),\mathbb{C}}/T_{\mathcal{T}(e),\mathbb{C}}$ is the relative divisor X_{P_0} (using notation from (1.1)). The matching condition at a node w is

$$u_{v_+}(w_+) = u_{v_-}(w_-) \in X_{P_0}.$$

A *broken map* is a collection of maps $u_v : C_v \to X_{\overline{P}(v)}$ described above together with a *tropical structure* on the graph Γ underlying the domain nodal curve. A *tropical structure* is a map of Γ to the dual complex $B^\vee \subset \mathfrak{t}^\vee$ of the neck-stretching in a manner that respects slopes of nodes. More precisely, a tropical structure on a graph Γ is a collection of edge slopes

$$\mathcal{T}(e) \in \mathfrak{t}_{P(e),\mathbb{Z}} \subset \mathfrak{t} \simeq \mathfrak{t}^\vee,$$

and polytope assignments $P(v) \in \mathcal{P}$ for vertices $v \in \mathrm{Vert}(\Gamma)$ so that there exist *tropical positions* of the vertices in the dual complex

$$\mathcal{T} : \mathrm{Vert}(\Gamma) \to B^\vee \subset \mathfrak{t}^\vee,$$

that satisfy a *slope condition*, namely that for any edge $e = (v_+, v_-)$ the slope of the line segment joining $\mathcal{T}(v_+)$ to $\mathcal{T}(v_-)$ is equal to the slope $\mathcal{T}(e)$ of the node w corresponding to e, which is given by the intersection multiplicities of the map with relative divisors at the node, see (1.4). That is,

$$\mathcal{T}(e) \in \mathbb{R}_{>0}(\mathcal{T}(v_+) - \mathcal{T}(v_-)). \tag{1.6}$$

The image of $\mathcal{T}(\Gamma)$ under the map to $B^\vee$ induced by $v \mapsto \mathcal{T}(v)$ is called an *embedded tropical graph*, and the underlying graph keeping only the edge slopes $\{\mathcal{T}(e)\}_e$ and vertex polytopes $\{P(v)\}_v$ is called a *tropical graph*. Thus changing the tropical vertex positions $\{\mathcal{T}(v)\}_v$ of a tropical graph is a change of embeddings, see Figure 4.2 for an example. Broken maps may also contain nodes corresponding to the standard nodal degeneration encountered in Gromov-Witten theory. The edges corresponding to such nodes do not appear in the tropical graph.

Broken maps arise naturally as limits of sequences of holomorphic maps in neck-stretched manifolds. A converging sequence of maps consists of pockets of high symplectic area area separated by long cylinders in neck regions. Each such sequence u_ν of long cylinders maps to a neck piece

$$u_\nu : [\tfrac{-\nu}{2}, \tfrac{\nu}{2}] \times S^1 \to X_{\overline{P}}^{\square}$$

for some $P \in \mathcal{P}$ in the polytopal decomposition, and is asymptotically close to a 'trivial cylinder'. A trivial cylinder is a holomorphic cylinder which lies in a fiber of the projection

$$T_{P,\mathbb{C}} \to X_{\overline{P}}^{\square} \to X_P^{\square},$$

and is thus a sub-torus $T_{\mu,\mathbb{C}} \subset T_{P,\mathbb{C}}$ generated by a rational element $\mu \in \mathfrak{t}_{P,\mathbb{Z}}$. The pockets of high area converge to spheres or disks, and (roughly speaking) the long cylinders connecting these pockets converge to trivial cylinders. We drop the rational cylinders from our description of the 'map part' of a broken map, and instead encode it as part of the tropical graph, see Figure 1.8. The generator

$$\mathcal{T}(e) \in \mathfrak{t}_{P(e),\mathbb{Z}}$$

of the trivial cylinder corresponding to an edge e is the slope of the edge in the tropical graph. This explains why the edge slope is integral. Every component C_v, $v \in \mathrm{Vert}(\Gamma)$ of the limit map has a natural position $\mathcal{T}(v) \in B^\vee$ in the dual complex as follows: For a limit curve component mapping to a neck piece $X_{\overline{P}}^{\square}$, $P \in \mathcal{P}$ with fibers $T_{P,\mathbb{C}} \simeq (\mathbb{C}^\times)^n$, the convergence is modulo $\mathbb{R}^n$-translation in target space $X_{\overline{P}}^{\square}$. More specifically, for a limit curve component $u : C \to X_{\overline{P}}^{\square}$ there is a sequence $t_\nu \in \nu P^\vee \subset \mathbb{R}^n \subset \mathfrak{t}_{P,\mathbb{C}}$ such that u is the limit of the translated maps $e^{-t_\nu} u_\nu$. Thus, the component u inherits a *tropical coordinate* $\lim_{\nu\to\infty} \frac{t_\nu}{\nu}$. For an edge $e = (v_+, v_-)$, the line segment connecting $\mathcal{T}(v_+)$ and $\mathcal{T}(v_-)$ has slope $\mathcal{T}(e)$.

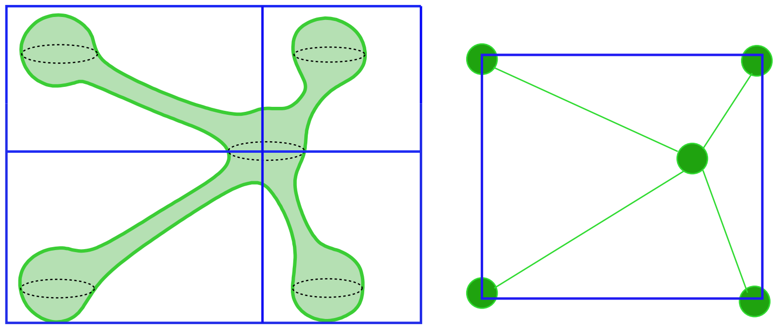

FIGURE 1.8. A holomorphic curve and a tropical graph

Viewing the broken map as a limit of maps in neck-stretched manifolds also explains the matching condition at nodes. A node w corresponding to an edge e in a limit broken map arises from the convergence of a sequence of long cylinders

$$u_\nu : S^1 \times [-l_\nu, l_\nu] \to X, \quad l_\nu \to \infty$$

with small area, which are asymptotically close to a trivial cylinder

$$u_{\text{triv}} : S^1 \times [-T_\nu, T_\nu], \quad (s,t) \mapsto e^{\mathcal{T}(e)(s+it)} x_0, \quad x_0 \in X$$

generated by an integral element $\mathcal{T}(e) \in \mathfrak{t}_{P(e),\mathbb{Z}}$. As a result, the evaluation of both lifts of the node $w_e^{\pm}$ are equal modulo the action of the one-dimensional complex torus $T_{\mathcal{T}(e),\mathbb{C}}$, which is a re-statement of the matching condition, see (1.5).

The set of broken maps of a fixed type Γ has a free action of a *tropical symmetry* group $T_{\text{trop}}(\Gamma)$ arising out of the torus action on neck pieces. The tropical symmetry group of a broken map $u : C \to \mathfrak{X}$ of type Γ is generated by the degrees of freedom of the tropical graph $\mathcal{T}$: these are ways of moving the vertex positions $\{\mathcal{T}(v)\}_v$ without changing edge slopes $\{\mathcal{T}(e)\}_e$, see Figure 4.4. In particular, the symmetry group $T_{\text{trop}}(\Gamma)$ is finite if in the tropical graph vertex positions $\mathcal{T}(v)$ are uniquely determined by the edge slope $\mathcal{T}(e)$ and such tropical graphs called *rigid*. In the second graph in Figure 4.4, there is one degree of freedom in moving the edges of the tropical graph, and this generates a one-dimensional complex torus $T_{\text{trop}}(\Gamma_2)$.

A feature of broken maps is that 'nodes do not lower the index of a map', which is in contrast to stable maps in smooth manifolds. Here the *index* of a broken map u[3] refers to the expected dimension of the moduli space component containing u, that is, the index of the linearization of the Fredholm operator cutting out the moduli space, and is denoted by $i^{\text{brok}}(u)$. The index is equal to the actual dimension

$$i^{\text{brok}}(u) = \dim T_u \mathcal{M}(\mathfrak{X}, L)$$

if the moduli space $\mathcal{M}(\mathfrak{X}, L)$ is regular. Nodes do not lower the index of a map because for a broken map $u = (u_v)_{v \in \text{Vert}(\Gamma)}$ the matching condition at an edge $e \in \text{Edge}(\Gamma)$ has codimension $(\dim X - 2)$, since it is a condition on the quotient

[3]We use the notation i for the index of a map, which is equal to the expected dimension of the moduli space containing the map; and the notation I for Maslov index of a disk.

of $X^{\square}_{\overline{P(e)}}$ by a complex one-dimensional torus $T_{\mathcal{T}(e),\mathbb{C}}$, see (1.5). In contrast the codimension of the matching condition for stable maps is $\dim(X)$. Therefore nodes occur only in strata whose codimension is at least two. A consequence is that broken maps u with components in neck pieces may have index zero, $i^{\mathrm{brok}}(u) = 0$, though that can happen only if the tropical graph is rigid. Indeed if for a type Γ the tropical graph is not rigid, then the tropical symmetry group $T_{\mathrm{trop}}(\Gamma)$ is at least two-dimensional, and since $T_{\mathrm{trop}}(\Gamma)$ has a free action on the moduli space $\mathcal{M}_\Gamma(\mathfrak{X})$ of maps of type Γ, the dimension of the moduli space $\mathcal{M}_\Gamma(\mathfrak{X})$ must be at least 2.

The case of a single cut was studied by Ionel-Parker [**51**] and J. Li [**55**]. They obtained a symplectic sum formula for Gromov-Witten invariants on X in terms of relative Gromov-Witten invariants of the cut spaces, relative to the relative divisor. Their work is a special case of symplectic field theory of Bourgeois-Eliashberg-Hofer-Wysocki-Zehnder [**11**], in which the analogs of broken maps are known as 'holomorphic buildings'. Eleny Ionel [**50**] first studied compactifications of moduli spaces of maps relative to a normal crossing divisor. She used a generalization of holomorphic buildings with levels and chambers to describe the maps appearing in the moduli space. Brett Parker [**68, 69, 70, 71, 72, 73, 74, 75, 76**] used the tropical approach to study this problem. Parker defined a category of exploded manifolds, which combined the map part and tropical part into a single space, in a way that convergence of broken maps is continuous. The corresponding Gromov-Witten invariants in algebraic geometry are studied in, for example, Abramovich-Chen-Gross-Siebert [**2**]. Tehrani gives an alternate compactification of holomorphic curves relative to a normal crossing symplectic divisor [**29, 31**] and uses it to give a degeneration formula [**30**] for Gromov-Witten invariants in the almost Kähler category.

1.4. Broken Fukaya algebras

Counts of disks in the broken symplectic manifold lead to a definition of a broken Fukaya algebra. We use Cieliebak-Mohnke perturbations [**22**] to regularize the various moduli spaces occurring in the paper. In order to use these perturbations we assume the Lagrangian $L \subset X$ and the symplectic manifold (X, ω) are compact, connected, and *rational* in the sense that X admits a line bundle whose curvature is the symplectic form and some tensor power is flat over L. The perturbation datum is a collection

$$\underline{\mathfrak{p}} = \{\mathfrak{p}_\Gamma = (J_\Gamma, F_\Gamma)\}$$

of domain-dependent perturbations $\mathfrak{p}_\Gamma$ for each type Γ of disk, each consisting of a domain-dependent almost complex structure J_Γ and Morse function F_Γ on the Lagrangian L. Using these perturbations one may construct the Fukaya algebra $CF(L)$ of the Lagrangian L in the Morse model constructed in, for example, Seidel [**82**] in the exact case and Charest-Woodward [**18**] in the rational case. The underlying vector space generated by critical points of the Morse function $f : L \to \mathbb{R}$ so that

$$CF(L) = \bigoplus_{x \in \mathrm{crit}(f)} \Lambda x.$$

The structure maps

$$m_d : CF(L)^{\otimes d} \to CF(L)$$

are defined by counting rigid elements in the moduli space $\mathcal{M}(L)$ of holomorphic treed disks

$$\mathcal{M}(L) = \{u : C \to X, \quad u(\partial C) \subset L, (4.5)\}/\sim .$$

Here C is a union of disks, spheres, and line segments. The conditions in Definition 4.5 require the map to be pseudoholomorphic on the disk and sphere components and satisfy a gradient flow equation on the segments, see Figure 1.9. Analogously

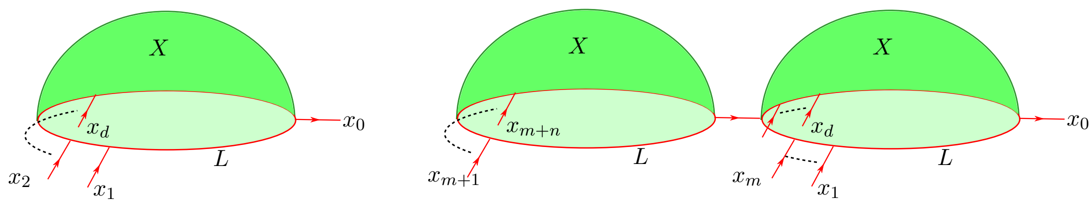

FIGURE 1.9. Holomorphic treed disks counted by the composition map m_d. Here segments with arrows are gradient flow lines of f, and $x_0, \ldots, x_d \in \mathrm{crit}(f)$.

counts of rigid elements of the moduli space of broken disks with boundary in L leads to a broken Fukaya algebra denoted $CF_{\mathrm{brok}}(L)$ with A_∞ structure maps

$$m_d^{\mathrm{brok}} : CF_{\mathrm{brok}}(L)^{\otimes d} \to CF_{\mathrm{brok}}(L), \quad d \geq 0,$$

see Chapter 10. The single cut version of this A_∞ algebra has been constructed by Charest-Woodward [**18**]. The composition maps m_d^{brok} and proof of the A_∞ relations for them only depend on zero or one-dimensional moduli spaces of treed disks. On these components the tropical symmetry group is finite, and thus the composition maps m_d^{brok} only count broken maps with a rigid tropical graph. Our main result is the following, proved in Chapter 10:

THEOREM 1.4. *For a rational Lagrangian submanifold $L \subset X$ as above, the unbroken Fukaya algebra $CF(X, L)$ admits a curved A_∞ homotopy equivalence to the broken Fukaya algebra $CF_{\mathrm{brok}}(\mathfrak{X}, L)$.*

The ingredients of this result for broken maps, proved in Proposition 10.22 below, are a convergence result and its converse, which is a gluing result. The convergence result is a generalization of sft compactness ([**11**], [**23**]) for a single cut and is proved in Chapter 8. The statement is that given a sequence of maps $u_\nu : C \to X$ holomorphic with respect to almost structures J_ν that are stretched along multiple necks, there is a subsequence of u_ν that converges to a broken map. The limit is unique up to the action of the tropical symmetry group. The gluing result in Chapter 9 is proved only for broken maps of index zero: an index zero regular broken map can be glued to produce a family of J^ν-holomorphic maps $u_\nu : C \to X$.

We emphasize that the bijection involves broken maps, and not tropical symmetry orbits of broken maps. The distinction is significant even for rigid maps because rigid broken maps may have a finite non-trivial tropical symmetry group, see Example 4.38. This phenomenon has also been observed by Abramovich-Chen-Gross-Siebert [**2**] and Tehrani [**30**].

We point out that the A_∞ homotopy equivalence in Theorem 1.4 is given by curved A_∞ functors

$$\mathcal{F} : CF(X, L) \to CF_{\text{brok}}(\mathfrak{X}, L), \quad \mathcal{G} : CF_{\text{brok}}(\mathfrak{X}, L) \to CF(X, L).$$

Here 'curved' means that the zero-th terms $\mathcal{F}^0$, $\mathcal{G}^0$ are non-vanishing, which implies that the homotopy equivalence does not preserve disk counts. Therefore, the *potential*, which is a count of disks with no input and a single output, may not be the same for (X, L) and $(\mathfrak{X}, L)$. We prove a limited result (Proposition 12.9) describing conditions under which neck-stretching does not alter the disk count for certain homology classes.

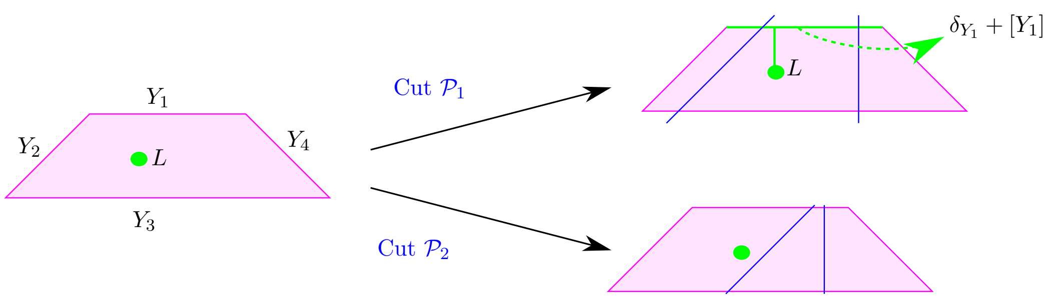

FIGURE 1.10. Moving a cut in the second Hirzebruch surface changes the disk count.

EXAMPLE 1.5. We give an example where a multiple cut changes the potential of a Lagrangian brane. Consider the second Hirzebruch surface $X := H_2$ with two single cuts as in Figure 1.10 and a toric Lagrangian L. In the first case $\mathcal{P}_1$, the broken manifold has the same potential

$$W_1(y_1, y_2) = y_2 y_1^{-1} + y_2 y_1 + 2y_2 + y_2^{-1}$$

as the unbroken manifold, whose terms correspond to

- for each torus invariant divisor $X_1, \ldots, X_4$, a Blaschke product with a single term, which is a disk δ_{X_i} intersecting the divisor X_i;
- and a disk in the homology class $\delta_{X_1} + [X_1]$.

In the second cut $\mathcal{P}_2$, the potential

$$W_2(y_1, y_2) = y_2 y_1^{-1} + y_2 + y_2^{-1}$$

does not have disks in two out of the five classes, namely δ_{Y_4} and $\delta_{Y_1} + [Y_1]$. For example, if $\mathfrak{X}_{P_2}$ were to contain a regular broken disk u in the class $\delta_{Y_1} + [Y_1]$, the homology classes of its pieces would be as shown in Figure 1.10, and u would satisfy a matching condition along the relative divisor Y_2 between two curves of self-intersection -1. But for a generic perturbation the two (-1)-curves will not be incident at the same point on Y_2. By the same argument, broken disks of class δ_{X_4}

are also ruled out in the broken manifold $\mathfrak{X}_{\mathcal{P}_2}$. See Remark 12.10 for more discussion on why the multiple cut $\mathcal{P}_2$ alters the potential.

1.5. Split matching conditions

Our main interest is a further degeneration of the matching condition in broken maps to a split form. By 'splitting the matching condition', we mean producing a homotopy from the moduli space of broken maps to a moduli space of a different kind of map that does not have matching conditions at certain nodes called *split nodes*. In the latter kind of map, the matching condition on a node is 'split' into independent constraints on each of the lifts of the node.

To split the matching condition at nodes we use a deformation of the moduli space of broken maps that has similarities to the Bourgeois version of symplectic field theory [**11**]; as well as the Fulton-Sturmfels degeneration of the diagonal in a toric variety [**33**]. Bourgeois' degeneration, where the degeneration is via the flow of a Morse function, is used by Charest-Woodward [**18**] to give a split form in case of a single cut. We remark that after the first version of this paper appeared, an analogous result in the context of logarithmic Gromov-Witten invariants appeared by Wu [**88**].

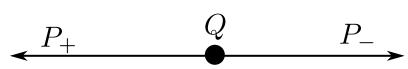

FIGURE 1.11. The polyhedral decomposition for a single cut.

To motivate our construction, we first describe the Morse deformation for broken maps in case of a single cut following Charest-Woodward [**18**]. We recall that for a single cut shown in Figure 1.11 the matching condition at a node w is given by

$$u_+(w_+) = u_-(w_-) \in X_Q$$

where X_Q is the relative divisor. Let

$$H : X_Q \to \mathbb{R}, \quad \phi_t^H : X_Q \to X_Q$$

be a Morse function with time t flow ϕ_t^H. Given a deformation parameter $t \in (0, \infty)$, a t-deformed map is defined by replacing the matching condition at a node by a condition that the lifts of the node are the end-points of a length t Morse trajectory:

$$u_+(w_+) = \phi_t^H u_-(w_-).$$

As $t \to \infty$ the Morse trajectory at any node degenerates to a broken Morse trajectory. The matching condition degenerates to the following condition: The lifts of the node lie on two transversely intersecting Morse cycles, namely the stable and unstable submanifold of some critical point of the Morse function.

The Morse deformation is thus based on a deformation of the diagonal $\Delta_{X_Q} \subset X_Q \times X_Q$ to a 'split' form, which means a union of products

$$\bigcup_{p \in \mathrm{crit}(H)} W^u(p) \times W^s(p) \subset X_Q \times X_Q. \tag{1.7}$$

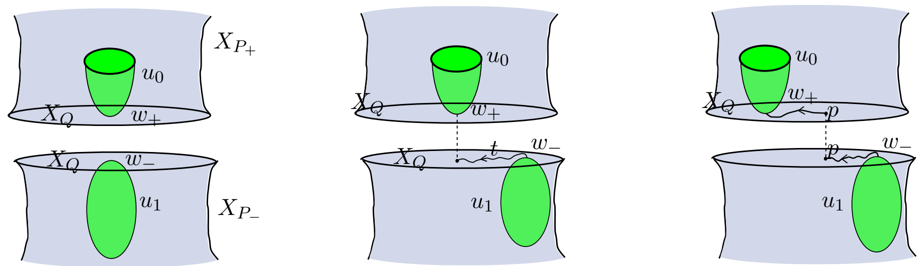

FIGURE 1.12. Morse deformation: From a broken map to a t-deformed map to an ∞-deformed map.

In the special case that X_Q is a T-toric variety, with $T \simeq (S^1)^n$, and the Morse-Smale pair is T-invariant, the critical points are T-fixed points and the stable and unstable manifolds are $T_{\mathbb{C}}$-invariant submanifolds. For many toric varieties, a generic component of the moment map is a T-invariant Morse function. For a generic vector $\xi \in \mathfrak{t}$, the gradient flow of the function $\langle \Phi, \xi \rangle : M \to \mathbb{R}$ is

$$(t, m) \mapsto e^{it\xi} m.$$

In [**18**], Morse deformation by a torus-invariant Morse-Smale pair is used to prove unobstructedness of certain Lagrangian tori.

In this paper we study the behavior of holomorphic curves under degeneration of the matching condition of nodes that lie in toric regions of the broken manifold. The set of edges, called *split edges*, that are degenerated satisfy the following conditions:

(Relative divisor assumption) If $e \in \mathrm{Edge}(\Gamma)$ is a split edge then every facet bounding the polytope $P(e)$ is shared with another polytope $P' \in \mathcal{P}$.

The subset of split edges is denoted by $\mathrm{Edge}_s(\Gamma) \subset \mathrm{Edge}(\Gamma)$.

Similar to the case of Morse deformation with a torus-invariant function in [**18**], we define deformed maps using the action of the complex torus. We choose a generic vector $\eta_0 \in \mathfrak{t}^\vee$ called the *cone direction* and deform the matching condition at the node w_e in the direction

$$\eta_e := \pi^{\perp}_{\mathcal{T}(e)}(\eta_0) \in \mathfrak{t}/\mathfrak{t}_{\mathcal{T}(e)}$$

We say that for any $\tau \in \mathbb{R}$ the *τ-deformed matching condition* at a node w_e is the condition that the projections of the evaluations match up to translation by the group element $e^{i\tau\eta_e}$:

$$(\pi^{\perp}_{\mathcal{T}(e)} \circ u)(z_+) = e^{i\tau\eta_e}(\pi^{\perp}_{\mathcal{T}(e)} \circ u)(z_-) \in X_{\overline{P}(e),\mathbb{C}}/T_{\mathcal{T}(e),\mathbb{C}}, \tag{1.8}$$

where the notation is carried forward from (1.5).

The limit objects in the degeneration of the matching conditions are called *split maps*. A split map is a version of a broken map with no matching condition on split nodes, and whose tropical graph satisfies a 'cone condition', which is explained in the following paragraph. Since there are no matching conditions at split nodes, the moduli space of split maps is a product. For a split tropical graph $\tilde{\Gamma}$ and a set of split edges $\mathrm{Edge}_s(\tilde{\Gamma}) \subset \mathrm{Edge}(\tilde{\Gamma})$, the moduli space of split maps is a product

$$\mathcal{M}^{\mathrm{split}}_{\tilde{\Gamma}}(L, \eta_0) = \prod_{i=1}^{m} \mathcal{M}_{\tilde{\Gamma}_i},$$

where $\tilde{\Gamma}_1, \dots, \tilde{\Gamma}_m$ are the connected components of $\tilde{\Gamma} \setminus \mathrm{Edge}_s(\tilde{\Gamma})$ and $\mathcal{M}_{\tilde{\Gamma}_i}$ is the moduli space of broken maps modelled on the graph $\tilde{\Gamma}_i$.

The cone condition on split tropical graphs is a genericity condition that ensures that the tropical symmetry group of a split map has the maximum possible dimension. We describe this condition: The tropical graph $\tilde{\Gamma}$ of a split map does not satisfy the slope condition on split edges. As a consequence the space $\mathcal{W}(\tilde{\Gamma})$ of tropical vertex positions of the graph $\tilde{\Gamma}$ is positive-dimensional. The amount by which a tropical vertex position map $\mathcal{T} \in \mathcal{W}(\tilde{\Gamma})$ fails to satisfy the slope condition at split edges is given by its *discrepancy* defined as

$$\mathrm{Diff}(\mathcal{T}) := (\mathrm{Diff}_e(\mathcal{T}))_{e \text{ is a split edge}}, \quad \mathrm{Diff}_e(\mathcal{T}) = (\mathcal{T}(v_+) - \mathcal{T}(v_-)) \mod \langle \mathcal{T}(e) \rangle,$$

where $\mathcal{T}(e) \in \mathfrak{t}_{\mathbb{Z}}$ is the slope of the edge $e = (v_+, v_-)$. Note that $\mathrm{Diff}_e(\mathcal{T}) = 0$ for the tropical graph of a broken map. The *discrepancy cone* is the positive linear span of the set of discrepancies

$$\mathrm{Disc}(\tilde{\Gamma}) := \mathbb{R}_+ \langle \mathrm{Diff}(\mathcal{W}(\tilde{\Gamma})) \rangle,$$

where $\quad \mathrm{Diff}(\mathcal{W}(\tilde{\Gamma})) = \{\mathrm{Diff}(\mathcal{T}) : \mathcal{T} \text{ is a vertex position map of the tropical graph } \tilde{\Gamma}\}$,

and $\mathrm{Disc}(\tilde{\Gamma})$ is a cone in the vector space

$$\oplus_{e \text{ is a split edge}} \mathfrak{t}/\langle \mathcal{T}(e) \rangle.$$

In the case of a single split edge e the *cone condition* for a split map says that the discrepancy cone contains the deformation direction η_0:

$$\mathrm{Disc}(\tilde{\Gamma}) \ni \pi^{\perp}_{\mathcal{T}(e)}(\eta_0).$$

If the direction of deformation η_0 is generic, the cone condition implies that the discrepancy cone $\mathrm{Disc}(\tilde{\Gamma})$ is top-dimensional. In case of multiple split edges, an ordering of the split edges $e_1 \prec \dots \prec e_k$ is part of the datum of the graph. The cone condition says that if for a tuple $(c_1, \dots, c_n) \in (\mathbb{R}_+)^n$ the ratios $\frac{c_i}{c_{i-1}}$ are large enough then

$$\mathrm{Disc}(\tilde{\Gamma}) \ni (c_i \pi^{\perp}_{\mathcal{T}(e_i)}(\eta_0))_{1 \le i \le n}. \tag{1.9}$$

The condition in (1.9) is equivalent to saying that if the splitting of the matching condition were to be carried out one edge at a time in the order $e_1, \dots, e_n$, then the cone direction η_0 is contained in the discrepancy cone at each of the steps. We explain this statement more precisely in Example 11.55. The cone condition (1.9) implies that the discrepancy cone $\mathrm{Disc}(\tilde{\Gamma})$ is a top-dimensional cone in the vector space $\oplus_{e \text{ is a split edge}} \mathfrak{t}/\langle \mathcal{T}(e) \rangle$, ensuring that the dimension of the tropical symmetry group is

$$\dim(T_{\mathrm{trop}}(\tilde{\Gamma})) = 2(\dim(T) - 1) * \text{Number of split edges},$$

which is exactly the codimension of the matching condition on split edges.

The cone condition on split maps is similar to the condition in Fulton-Sturmfels [**33**]. The discrepancy cone $\mathrm{Disc}(\tilde{\Gamma})$ is generated by a set of cones

$$\mathcal{W}_i, \quad i = 1, \ldots, k, \tag{1.10}$$

each corresponding to a connected component $\tilde{\Gamma}_i \subset \tilde{\Gamma} \setminus \mathrm{Edge}_s(\tilde{\Gamma})$. We recall that the Fulton-Sturmfels formula expresses the homology class of the diagonal in a toric variety as a sum of products of classes, where each product has an underlying decomposition of a cone in $\mathfrak{t}$. The degeneration to split maps has a similar flavor since the moduli space of each 'type' of split map is a product of moduli spaces of maps, accompanied by a decomposition of the discrepancy cone $\mathrm{Diff}(\mathcal{W}(\tilde{\Gamma}))$ into the cones in (1.10).

For the purpose of producing A_∞-homotopy equivalences, we apply the deformation process on broken maps with index 0 or 1, i.e. maps with finite tropical symmetry groups. The limit of deformed maps then produces split maps, whose tropical symmetry group is isomorphic to

$$T_{\mathrm{trop}}(\tilde{\Gamma}) \cong (\mathbb{C}^\times)^{(\dim(T)-1)*\text{Number of split edges}} \times \text{ a finite group.}$$

We define the composition maps of the tropical Fukaya algebra

$$CF_{\mathrm{trop}}(\mathfrak{X}, L, \eta_0)$$

via counts of tropical symmetry orbits of split maps, whose split tropical graph contains the deformation direction η_0. Counts of orbits of split maps satisfy an A_∞-relation and the resulting A_∞-algebra $CF_{\mathrm{trop}}(\mathfrak{X}, L)$ is called the *tropical Fukaya algebra*. Our second main result is as follows :

THEOREM 1.6. *Let $\eta_0 \in \mathfrak{t}^\vee$ be a generic direction of deformation. The broken Fukaya algebra $CF_{\mathrm{brok}}(\mathfrak{X}, L)$ is A_∞-homotopy equivalent to the tropical Fukaya algebra $CF_{\mathrm{trop}}(\mathfrak{X}, L, \eta_0)$.*

The proof of the Theorem for split maps is completed in Proposition 11.65 below. We conjecture that the matching condition can be degenerated at all nodes, including the nodes which map to divisors that are not toric. In that case the composition maps in the tropical Fukaya algebra can be expressed as a sum of products, where the sum is over all rigid tropical graphs, and the product is over all vertices in the graph. A limited version of such a factorization formula is given in Remark 11.49.

1.6. Unobstructedness of toric Lagrangians

Using the tropical Fukaya algebra, we prove unobstructedness of certain Lagrangian tori, generalizing a result of Fukaya-Oh-Ohta-Ono result [**35**] on Lagrangian torus orbits in toric varieties. We call a Lagrangian $L \subset X$ a *tropical torus* if there exists a polytope $P \in \mathcal{P}$ of maximal dimension such that X_P is a toric manifold (that is, is a compact symplectic manifold with a completely integrable Hamiltonian torus action) and $L \subset X_P$ is a Lagrangian torus orbit. A tropical torus appears to be a somewhat stronger requirement than the notion of moment fiber in a toric degeneration used in mirror symmetry, where "discriminant" singularities are allowed as in Gross [**42**]. However, since singular cuts can be transformed to smooth straddling cuts (see Section 2.4) our result in fact holds for any Lagrangian torus in a toric degeneration that does not intersect relative divisors.

Homotopy equivalence with the tropical Fukaya algebra is used to prove that tropical tori are weakly unobstructed in the following sense. First, we point out that the geometric part of the Fukaya algebra $CF^{\text{geom}}(L)$ generated by the critical points of a Morse function on the Lagrangian may not have a strict unit. The geometric Fukaya algebra $CF^{\text{geom}}(L)$ can be enlarged to $CF(L)$ in order to include a strict unit $1_L \in CF(L)$ by the *homotopy unit construction* as in Fukaya-Oh-Ohta-Ono [**34**, (3.3.5.2)] (see Section 10.3 for an exposition on the lines of [**18**]). The cohomology $H(A)$ of an A_∞ -algebra A is well-defined if the first order composition map m_1 satisfies $m_1^2 = 0$. The condition $m_1^2 = 0$ may fail to hold if the *curvature* $m_0(1) \in A$ is not a Λ-multiple of the unit 1_L, which is an 'obstruction' to the definition of Floer cohomology. *Weak unobstructedness* is a more general condition under which the Floer cohomology of a Lagrangian brane can be defined. A Lagrangian brane L is *weakly unobstructed* if the projective Maurer-Cartan equation

$$m_0(1) + m_1(b) + m_2(b,b) + \ldots = W(b)1_L \quad \text{for some } W(b) \in \Lambda. \tag{1.11}$$

has an odd solution $b \in CF(L)$. Any odd solution to the Maurer-Cartan equation is called a *weakly bounding cochain* and the set of all the odd solutions is denoted $MC(L)$. Given a weakly bounding cochain, the Fukaya algebra $CF(L)$ may be 'deformed' by b (see (10.6)) to yield an A_∞ -algebra $CF(L,b)$ with composition maps $(m_n^b)_{n\geq 0}$ satisfying

$$m_0^b = W(b)1_L, \quad \text{for some} \quad W(b) \in \Lambda.$$

Consequently, $(m_1^b)^2 = 0$, and the Floer cohomology

$$HF(L,b) := \ker(m_1^b)/\operatorname{im}(m_1^b)$$

is well-defined. The function

$$W : MC(L) \to \Lambda, \quad b \mapsto W(b)$$

is called the *potential* of the curved A_∞ algebra $CF(L,b)$. The following result shows that a tropical Lagrangian is unobstructed, and it has a distinguished solution of the Maurer-Cartan equation.

COROLLARY 1.7. (Unobstructedness of a tropical torus) *Suppose that $L \subset X_{P_0}$ is a tropical torus equipped with a brane structure and that all the facets of P_0 are elements of $\mathcal{P}$. For any generic cone direction $\eta_0 \in \mathfrak{t}^\vee$, we have*

$$m^0_{CF_{\text{trop}}(\mathfrak{X},L,\eta_0)} = wx^\blacktriangledown$$

where $x^\blacktriangledown$ is the unique maximum point of the Morse function on L and $w \in \Lambda_{>0}$ is an element of the positive part of the Novikov ring. In the tropical Fukaya algebra $CF_{\text{trop}}(\mathfrak{X},L,\eta_0)$,

$$b := wx^{\triangledown} \in MC(L)$$

is a solution of the Maurer-Cartan equation, and the potential of the b-deformed A_∞ algebra $CF_{\text{trop}}(\mathfrak{X},L,\eta_0,b)$ is

$$W(b) := w1_L.$$

Consequently $CF_{\text{trop}}(\mathfrak{X}, L, \eta_0)$*, and hence* $CF(X, L)$*, are weakly unobstructed.*

Corollary 1.7 is re-stated and proved in Section 12.1. Corollary 1.7 gives a solution $Wx^{\blacktriangledown}$ of the Maurer-Cartan equation for the tropical Fukaya algebra in this case. The corresponding solution for the unbroken Fukaya algebra $CF(L)$ is $\mathcal{F}(Wx^{\blacktriangledown})$ where $\mathcal{F} : CF_{\text{trop}}(\mathfrak{X}, L) \to CF(X, L)$ is a A_∞-homotopy equivalence.

For a tropical Lagrangian L, the weakly bounding cochain produced by Corollary 1.7 is a good choice. Indeed, for a toric variety X we describe a multiple cut below for which Corollary 1.7 is applicable to yield a weakly bounding cochain b, and the leading order terms in the potential $CF_{\text{trop}}(\mathfrak{X}, L, b)$ coincide with the leading order terms in the Batyrev-Givental potential (12.1); thus re-proving the following result (Corollary 1.8) of Fukaya-Oh-Ohta-Ono [**35**]. We now describe the multiple cut $\mathcal{P}$ of a toric variety X used in our proof. The multiple cut $\mathcal{P}$ consists of a collection of orthogonally intersecting single cuts, with one cut parallel to each facet of the moment polytope. In particular, if X is a T-toric variety with moment map $\Phi : X \to \mathfrak{t}^\vee$ and moment polytope

$$\Delta := \Phi(X) = \{x \in \mathfrak{t}^\vee : \langle \mu_i, x \rangle \le c_i, i = 1, \dots, N\}$$

where $\mu_i \in \mathfrak{t}$ is the primitive outward pointing normal of the i-th facet of Δ, and $c_i \in \mathbb{R}$; the cuts are along the hypersurfaces

$$\langle x, \mu_i \rangle = c_i - \epsilon_i, \quad i = 1, \dots, N$$

where $\epsilon_i > 0$ is a small constant. See Figure 1.13. Thus one of the cut spaces

$$X_{P_0} := \Phi^{-1}(\{x : \langle \mu_i, x \rangle \le c_i - \epsilon_i, i = 1, \dots, N\} \tag{1.12}$$

is diffeomorphic to X. We assume that the Lagrangian torus lies in X_{P_0}. Our

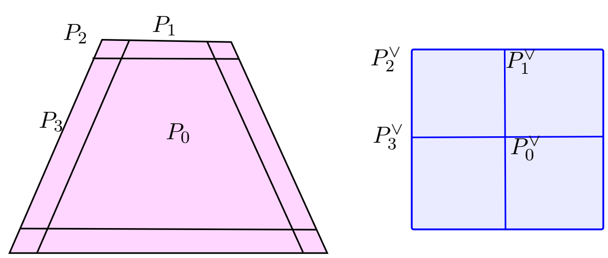

FIGURE 1.13. The multiple cut and its dual complex on a symplectic toric manifold

motivation in considering this multiple cut is that it has a regular perturbation for which the almost complex structure on X_{P_0} is standard. Indeed, the standard almost complex structure on a toric manifold fails to be regular if the torus-invariant divisors contain holomorphic spheres of non-positive Chern number. Such spheres are absent in our setting because all the torus-invariant divisors are relative divisors, and therefore intersect maps at isolated points. Of course there are other multiple cuts for which the Lagrangian is contained in a toric piece, all whose torus-invariant divisors are relative divisors. But we shall see that the multiple cut we have chosen makes the calculation of the potential easy, and also the count of least area disks is unaffected by neck-stretching.

Corollary 1.8. (Disk potentials in a toric manifold) *Let X be a symplectic toric manifold, and let $\mathfrak{X}$ be the broken manifold obtained by applying the multiple cut described above. Suppose $X_{P_0} \subset \mathfrak{X}$ is the top-dimensional component diffeomorphic to X (see* (1.12)*), and $L \subset X$ is a toric Lagrangian. Then the conclusions of Corollary 1.7 hold for L. For the distinguished solution b of the Maurer-Cartan equation on $CF_{\mathrm{trop}}(\mathfrak{X}, L, \eta_0)$, the leading order terms in the potential $W(b)$ coincide with the leading order terms of the Batyrev-Givental potential* (12.1)*.*

Corollary 1.8 is restated and proved in Section 12.2.

Example 1.9. We examine how a disk of Maslov index two in a toric variety deforms to a split disk. See Figure 1.14 for a two-dimensional example. One possibility is that the broken disk $u = (u_+, u_-)$ consists of a disk u_+ of Maslov index two in X_{P_0} and a sphere u_- in X_{P_1}. The space X_{P_1} is a $\mathbb{P}^1$-fibration, and u_- is a fiber. For any deformation parameter $\tau \in \mathfrak{t}/\langle \mathcal{T}_e \rangle$, in the τ-deformed map, the disk u_+ stays the same and u_- shifts to u_-^τ, which is a different $\mathbb{P}^1$-fiber in X_{P_1}. If τ approaches infinity in a generic direction, the limit of the maps u_-^τ lies in $\Phi^{-1}(P')$ where P' is an edge of the polytope P_1. As a result, in the limit split map, the limit u_-^∞ of $(u_-^\tau)_\tau$ lies in a neck piece which is a thickening of $X_{P'}$. The component u_-^τ has a $2(\dim(T) - 1)$-dimensional tropical symmetry group, since the corresponding vertex in the dual polytope is free to move in the interior of $(P')^\vee$. The details are part of the proof of Proposition 12.3. In Figure 1.14, $P' = P_1 \cap P_2$ and $\dim((P')^\vee) = 1$.

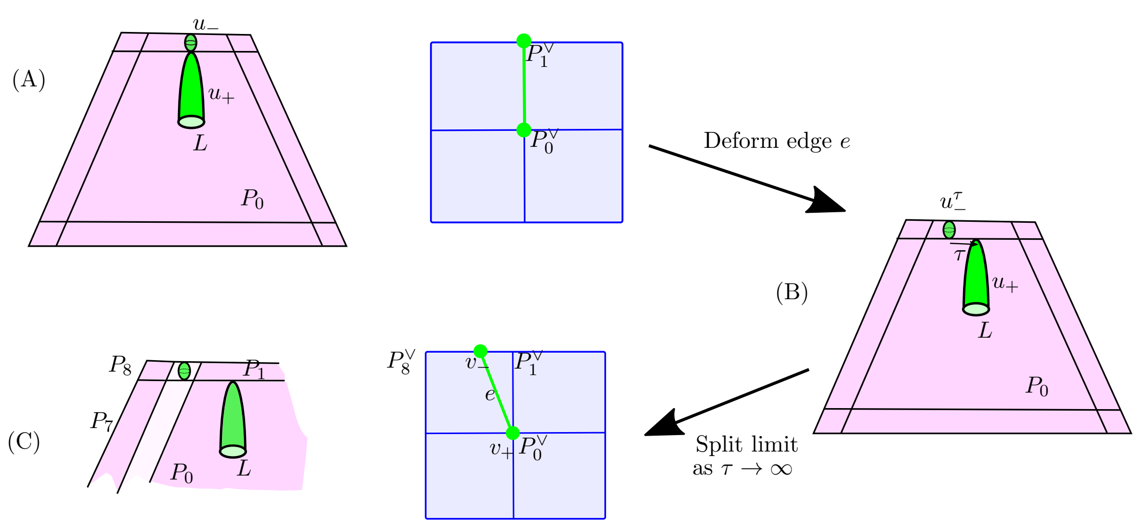

Figure 1.14. (A) A broken disk of Maslov index 2 and its tropical graph, (B) a τ-deformed disk, (C) the split disk obtained in the limit $\tau \to \infty$ and its tropical graph.

CHAPTER 2

Applications to disk counting

Although our main goal is the existence of a tropical version of the Fukaya algebra, rather than development of computational tools, the results of this paper can in some cases be used to compute disk counts, and in particular, prove formulas for disk potentials of various Lagrangians. In this chapter, we describe several such examples that the reader might keep in mind during the reading of the theory.

2.1. Counting disks in cubic surfaces

Broken maps can be used to compute the number of lines in a cubic surface. We also compute its "open" analog, namely the disk potential of a toric Lagrangian in the cubic surface – this number is independent of perturbations because the Lagrangian is monotone. We show that there are twenty one disks contributing to the potential as conjectured by Sheridan [**83**, Appendix B]; this example also appears in Pascaleff-Tonkonog [**77**] and can also be addressed using the method of Chan-Lau-Cho-Tseng [**15**]. The paper Bardwell-Evans-Hong-Lin [**6**] describes the disks in terms of scattering diagrams.

First we deform the cubic surface to a semi-Fano toric surface, and then we describe a multiple cut on the deformation. Recall that a toric manifold X is *semi-Fano* if for any torus invariant sphere $S \subset X$, the first Chern number $c_1(TX|S)$ is non-negative. We have seen in Example 2.24 that a cubic surface has a toric degeneration to an almost toric manifold X^0 with 9 singular points. The toric fibration on X^0 may be altered via a nodal slide, so that we may assume there is a continuous map Φ from X^0 to a triangle

$$\Delta = \text{hull}\{(-1,-1),(2,1),(1,2)\}.$$

The map Φ is a moment map for a Hamiltonian torus action on the complement

$$X^0 \setminus \bigcup_{v \text{ is a vertex of } \Delta} \Phi^{-1}(v).$$

The inverse image of each vertex $v \in \Delta$ is a union $\Phi^{-1}(v) \cong S^2 \cup S^2$ of two Lagrangian spheres, thus each of these is the resolution of an A_2 singularity. In Examples 2.1, 2.4 we count the number of holomorphic spheres or disks of minimal Chern number resp. Maslov index. To do the count, we first deform the symplectic form on X^0 to X^ϵ, $\epsilon > 0$, which is a toric manifold whose moment polytope has facets

$$y = \pm 1, \quad x = \pm 1, \quad x - y = \pm 1, \tag{2.1}$$

and

$$x + y = 1 + \epsilon, \quad y - 1 - \epsilon = 2x, \quad x - 1 - \epsilon = 2y. \tag{2.2}$$

The divisors in X^ϵ corresponding to the facets (2.1) are called *short divisors* , and those corresponding to (2.2) are called *long divisors*. The short resp. long divisors are denoted by E_i resp. D_i in Figure 2.2. Note that as a symplectic manifold $X^0 = \lim_{\epsilon \to 0} X^\epsilon$, and the short facets are collapsed to a single point in the moment polytope of X_0. We fix a small $\epsilon > 0$, and denote

$$X := X^\epsilon. \tag{2.3}$$

As a second step to do the curve count, we apply a multiple cut on X that splits it into Fano pieces as shown in Figure 2.1. The resulting broken manifold is called

$$\mathfrak{X}. \tag{2.4}$$

We count disks and spheres in $\mathfrak{X}$, and prove that the count is the same in the cubic surface.

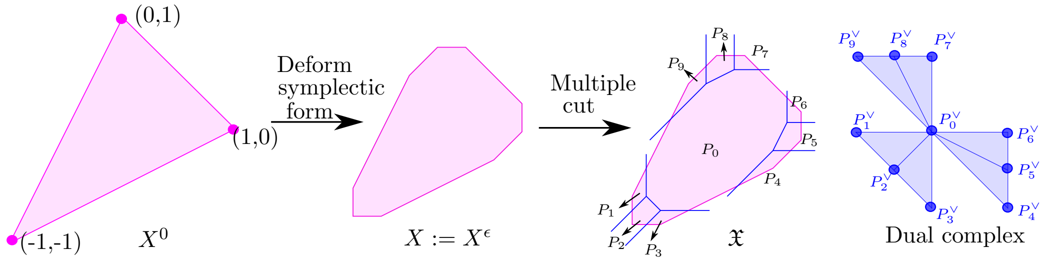

FIGURE 2.1. Multiple cut on the cubic surface.

An alternate way of viewing the cubic surface as an almost toric manifold is as follows: A cubic surface is isomorphic as a Kähler manifold to the projective plane $\mathbb{P}^2$ blown up at generic points $z_1, \dots, z_6$. Let M_1 be the cubic surface given by blowing up 3 generic points z_4, z_5, z_6 in the toric del Pezzo $\mathrm{Bl}_3(\mathbb{P}^2)$, and equipped with the Fubini-Study form. By moving the three points to toric fixed points in $\mathrm{Bl}_3(\mathbb{P}^2)$, we obtain a path (M_t, J_t), $t \in [0,1]$ of Kähler manifolds such that $M_0 = X^0$. The symplectic form on M_t is monotone for $t \in (0,1]$, and by taking the limit, it is also monotone for $X^0 = M_0$. (Monotonicity means that $\omega = \lambda c_1$ for some $\lambda > 0$ and $\omega, c_1 : \pi_2(X) \to \mathbb{R}$ are homomorphisms.)

EXAMPLE 2.1. (Twenty-seven lines on a cubic surface) The number of lines in a cubic surface was proved by Salmon and Cayley to equal twenty-seven, see Mumford [**62**, Section 8D], and is easily seen to be the Gromov-Witten invariant. We re-prove this result by counting broken maps in $\mathfrak{X}$. Indeed, the Gromov-Witten invariant is not altered by deformation of the symplectic form or neck-stretching. The multiple cut in Figure 2.1 splits X into Fano surfaces, so that the homology class of any lowest area sphere degenerates to the homology class of a union of spheres with self-intersection -1. In Proposition 2.2 we show that a broken sphere in $\mathfrak{X}$ with Chern number 1 consists of one long divisor and a non-self crossing path of short divisors.

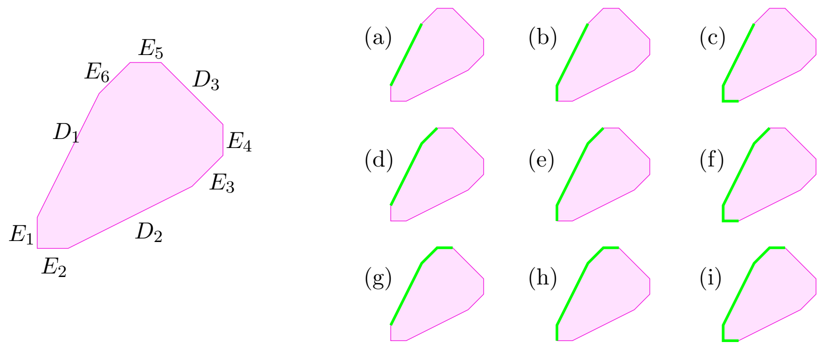

FIGURE 2.2. Nine of the twenty-seven lines on a (deformation of) a cubic surface, in the homology classes $[D_1] + \sum_{i=1,2,5,6} n_i[E_i]$.

There is one configuration corresponding to each of the pictures in Figure 2.2. By symmetry there are twenty-seven such configurations in total.

The following result is proved later in the section.

PROPOSITION 2.2. *Any broken sphere of index zero in the broken manifold $\mathfrak{X}$ is of the form shown in Figure 2.2, and each such broken sphere has multiplicity one. Therefore, a cubic has* 27 *lines.*

REMARK 2.3. (Disks versus spheres) We point out that the results in this paper are about Fukaya algebras (defined using counts of disks), whereas Proposition 2.2 uses analogous results in Gromov-Witten theory (defined using counts of spheres). Broken Gromov-Witten invariants are well-defined in cases when markings are constrained to lie on homology classes that are represented by cycles that are contained in a single top-dimensional cut space. In the case of lines in the cubic surface there are no markings. Whenever the broken (and split) Gromov-Witten invariants are well-defined, they are equal to the ordinary GW invariants on the unbroken manifold. The proofs are easier than the A_∞ -counterpart because the coherence condition on the perturbation datum is weaker. We point out this difference in Remarks 6.8 and 11.7.

EXAMPLE 2.4. (Twenty-one disks in the cubic surface) The Fukaya category of a cubic surface was computed by Sheridan [**83**, Appendix B], and it was shown to have a summand corresponding to twenty-one as an eigenvalue of the first Chern class. Recall that the cubic surface X is a toric fibration

$$\Phi : X \to \Delta := \mathrm{hull}\{(-1,-1),(2,1),(1,2)\} \tag{2.5}$$

in the complement of $\Phi^{-1}(v)$ where v ranges over the three vertices of Δ. The Lagrangian L is the monotone fiber

$$L := \Phi^{-1}(0,0). \tag{2.6}$$

The Lagrangian torus L is expected to generate the sub-category of the Fukaya category corresponding to the "small eigenvalue" in Sheridan's language in [**83**].

In Proposition 2.5 below we show that a disks in the cubic surface X with boundary in L and of Maslov index $I(u) = 2$ either

(a) intersects a long divisor D_i (there are 3 such disks), or
(b) is a nodal disk consisting of a disk $u_0 : S \to X$ intersecting a short divisor E_i and a non-intersecting path of spheres $u_1, \ldots, u_k : \mathbb{P}^1 \to X$, each mapping to a short divisor E_j, and thus having Maslov index zero. There are 18 such nodal disks. See Figure 2.3.

To prove the Proposition we use a multiple cut on X and count broken disks, and then show that neck-stretching preserves the disk count. These computations can also be addressed using the method of Chan-Lau-Cho-Tseng [**15**]. The paper Bardwell-Evans-Hong-Lin [**6**] describes the disks in terms of scattering diagrams.

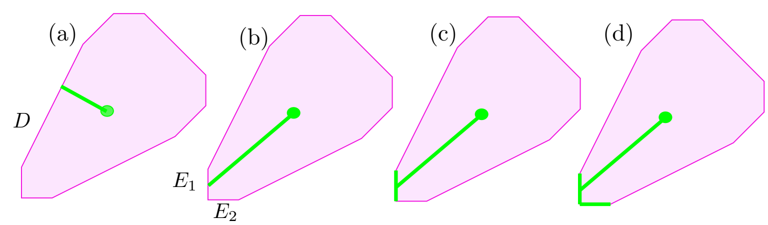

FIGURE 2.3. There are twenty-one Maslov-index-two-disks on a cubic surface: 3 of type (a), and 6 each of type (b), (c) and (d).

PROPOSITION 2.5. *Let X be the cubic surface with the Fubini-Study form, and let L be the monotone toric Lagrangian from* (2.6). *Any disk contributing to the potential of $CF(X, L)$ is of the form shown in Figure 2.3, and each disk has a multiplicity of one.*

PRELUDE TO THE PROOF OF PROPOSITIONS 2.2, 2.5. We describe the perturbation used for counting broken spheres and broken disks on the multiply cut cubic surface. We use a perturbation that depends both on the domain curve and the tropical graph. Such perturbations, called *split perturbations*, are defined rigorously in Section 10.7. In Section 10.7 we also show that split perturbations yield the same Gromov-Witten invariants and A_∞-homotopy equivalent Fukaya algebras as coherent domain dependent almost complex structures. We recall that a domain-dependent perturbation has an underlying *background almost structure* which is perturbed in a compact subset of the complement of the nodal points in the domain. In a split perturbation, the background almost complex structure is allowed to vary across components C_v, $v \in \mathrm{Vert}(\Gamma)$ of the domain. For a tropical graph Γ we denote the background almost complex structures by

$$J_v^b \in \mathcal{J}^{\mathrm{cyl}}(X_{\overline{P}(v)}), \quad v \in \mathrm{Vert}(\Gamma),$$

which, for the broken cubic surface, is defined as

$$J_v^b := \phi_v^* J_{std}, \tag{2.7}$$

where J_{std} is the standard almost complex structure on the toric manifold $X_{P(v)}$; and $\phi_v : X_{P(v)} \to X_{P(v)}$ is a Hamiltonian diffeomorphism that is supported away from the relative divisors. We observe that the tropical graph of a broken map allows us to detect when the component of a map is homologous to a multiple of a (-1)-curve, such as E_1', E_1'', E_2', E_2'' etc. in Figure 2.5. For example, a map component u_v with $P(v) = P_1$ (see Figure 2.5 for notation) is homologous to kE_1' exactly if all edges e incident on v satisfy $P(e) = P_{12} := P_1 \cap P_2$, and the sum of multiplicities of the edges is k. We choose perturbations that

> (Constant on simple components) vanish on components where the map is not a multiple cover of a (-1)-sphere.

That is, on such a component C_v the perturbed almost complex structure is equal to the background J_v^b. On components where the map is a multiple cover of a (-1)-sphere, the background J_v^b is perturbed in a domain-dependent way. The proof of the sphere count and disk count is carried out separately below. □

NOTATION 2.6. (Pseudoholomorphic (-1)-spheres) Suppose $P \in \mathcal{P}$ is a top-dimensional polytope in the multiple cut of the cubic surface, and $P \neq P_0$. Suppose $E \subset X_P$ is a (-1)-sphere that is J_{std}-holomorphic. For any tamed almost complex structure J on X_P, there is a J-holomorphic sphere homologous to E which we denote by

$$E_J \subset X_P. \tag{2.8}$$

Observe that in this notation E is the same as $E_{J_{std}}$.

REMARK 2.7. (There are no broken (-2)-spheres) A split perturbation as above ensures that, with respect to the background almost complex structure, there are no pseudoholomorphic spheres homologous to short divisors (whose self-intersection is -2) by the following reason: Suppose E is a short divisor whose broken version has components $E' \subset X_{P_0}$, $E'' \subset X_{P_1}$. Let $Q = P_0 \cap P_1$. For any tropical graph Γ, and $e = (v_0, v_1) \in \mathrm{Edge}(\Gamma)$, with $P(v_0) = P_0$, $P(v_1) = P_1$, for generic background almost complex structures $J_{v_0}^b$, $J_{v_1}^b$ the spheres $E'_{J_{v_0}}$ and $E''_{J_{v_1}}$ (using notation from (2.8)) intersect X_Q at different points. Therefore there is no broken pseudoholomorphic sphere homologous to $E' \cup E''$ for the background almost complex structures. We may assume the domain-dependent perturbations of $J_{v_0}^b$, $J_{v_1}^b$ are small enough that the evaluations $\mathrm{ev}_{w_e^-}(u_{v_0})$, $\mathrm{ev}_{w_e^+}(u_{v_1}) \in X_Q$ lie in disjoint neighborhoods of the points $E'_{J_{v_0}} \cap X_Q$ and $E''_{J_{v_1}} \cap X_Q$, thereby ruling out perturbed (-2)-spheres also.

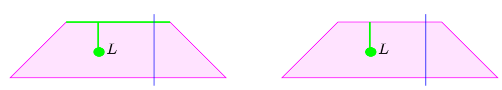

FIGURE 2.4. Disks in the second Hirzebruch surface that intersect the exceptional divisor.

REMARK 2.8. As an aside we remark that in a resolution of an A_1-singularity, a symplectic cut can be used to count disks. As in the case of the cubic surface, the cut is useful because it splits up the surface into Fano surfaces. For example, the resolution of the weighted projective plane $\mathbb{P}(1,1,2)$ is the second Hirzebruch surface, and in this case we may use a single cut to count disks. Given a toric Lagrangian, there are two disks of Maslov index two that intersect the exceptional locus (that is, the inverse image of the A_2-singularity in the resolution). The homology classes of the disks and the cut are shown in Figure 2.4.

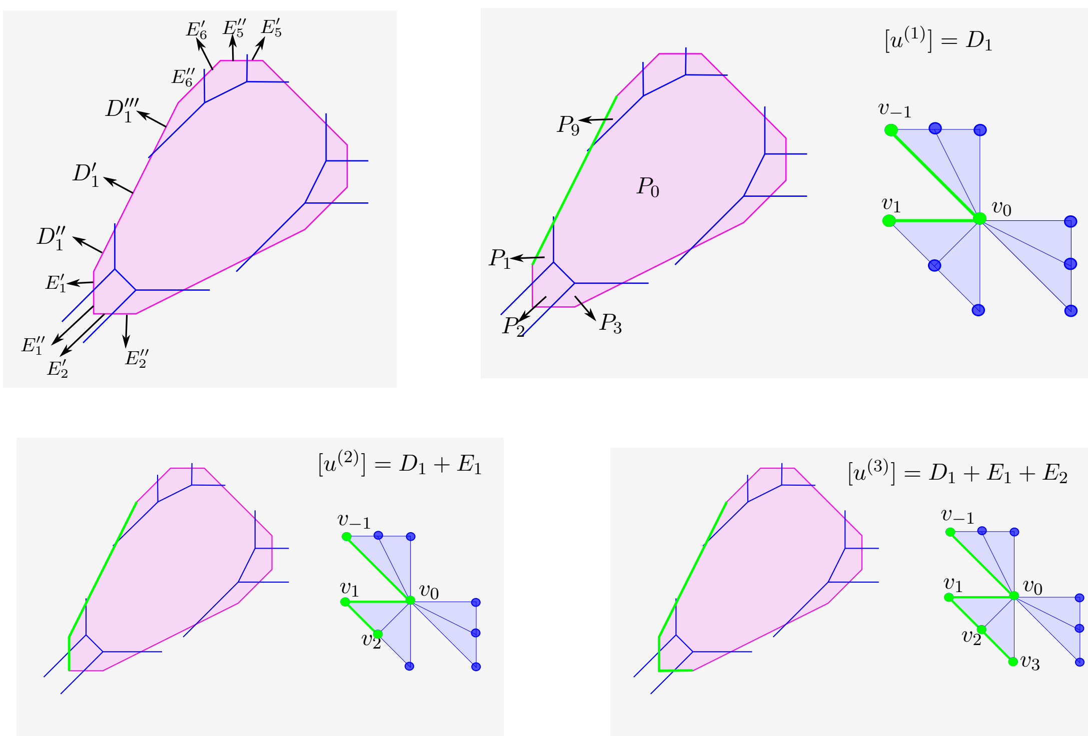

FIGURE 2.5. Top left: Broken versions of some divisors in the cubic surface. Other figures: Broken maps u^1, u^2, u^3 whose gluing is homologous to $[D_1] + k_1[E_1] + k_2[E_2]$ for some $k_1, k_2 \geq 0$.

We now prove the result on spheres on the cubic surface.

PROOF OF PROPOSITION 2.2. We recall from the discussion at the beginning of the section that the homology class of a sphere of index zero in X is the sum of a long divisor class $[D_i]$ and an arbitrary number of short divisor classes. Consider a broken sphere u such that its gluing has homology class

$$[u_{\text{glue}}] = [D_1] + n_1[E_1] + n_2[E_2] + n_5[E_5] + n_6[E_6],$$

where D_1, E_i are defined in Figure 2.2. We aim to prove that u is of the form shown in Figure 2.2. We first assume that that $n_5 = n_6 = 0$, and so,

$$[u_{\text{glue}}] = [D_1] + n_1[E_1] + n_2[E_2]. \tag{2.9}$$

We will show that a regular configuration is one of the types (a)-(c) in Figure 2.2. The conclusion for the general case ($n_5, n_6 \neq 0$) follows by symmetry.

Consider a broken map u with homology class as in (2.9) and a tropical graph Γ, we will show that $n_1, n_2 \leq 1$. The proof is carried out in steps, where in every step, we uncover the possible maps u_v corresponding to a vertex v of the tropical graph. In the following analysis we refer to Figure 2.5 for notation.

Step 1 (v_0): Since the long divisor is a summand in the homology class represented by the map u, there is a component of u_{v_0} of the broken map that maps to D_1' (see Figure 2.5 for notations). By the assumption that $n_5 = n_6 = 0$,

- v_0 has an edge to v_{-1} with $P(v_{-1}) = P_9$, and the homology class of the map $u_{v_{-1}}$ is $[D_1''']$;
- and an edge to v_1 with $P(v_1) = P_1$,

and no other edges.

Step 2 (v_1): So far we have shown that the broken map has vertices v_0, v_{-1} and further, v_0 has an edge of multiplicity one to a vertex v_1 with $P(v_1) = P_1$. Given that the homology class of the gluing of u is given by (2.9), the homology class of u_{v_1} is

$$[D_1''] + k[E_1'] \quad \text{for some } k \geq 0. \tag{2.10}$$

We proceed to show that $k = 0$ or 1: We recall that $E'_{1,J_{v_1}}$ is a J_{v_1}-holomorphic sphere in X_{P_1} that is homologous to the (-1)-sphere $E_1' \subset X_{P_1}$, see notation in (2.8). Since the map u_{v_1} intersects the relative divisor $X_{P_{01}}$, the homology class $[u_{v_1}]$ can not be a multiple of the class $[E_1']$. The property (Constant on simple components) then implies that the domain-dependent almost complex structure for u_{v_1} is equal to the constant J_{v_1}. In (2.10) if $k > 1$, then $[u_{v_1}].[E'_{1,J_{v_1}}] < 0$. This implies that the image of u_{v_1} is contained in $E'_{1,J_{v_1}}$, contradicting the fact that u_{v_1} intersects the relative divisor $X_{P_{01}}$. In case $k = 0$, the broken map u is u^1 represented in Figure 2.2 (a) and Figure 2.5 (a), $[u^1_{\text{glue}}] = [D_1]$ and there are no more vertices in the tropical graph. If $k = 1$, v_1 has an edge to a vertex v_2 with $P(v_2) = P_2$.

Step 3 (v_2): So far we have shown that for a broken map u whose gluing has homology class (2.9), either $[u_{\text{glue}}] = [D_1]$ or the tropical graph of u has vertices v_{-1}, v_0, v_1 and v_1 has an edge of multiplicity 1 to a vertex v_2 with $P(v_2) = P_2$. In the latter case, the homology class (2.9) of the glued map $[u_{\text{glue}}]$ implies that the homology of u_{v_2} is

$$[u_{v_2}] = k_1[E_1''] + k_2[E_2'],$$

and since v_2 has an edge to v_1, $k_1 \geq 1$.

First consider the case $k_2 = 0$. If $k_1 = 1$, then there are no more vertices, $u = u^2$, and the map is represented in Figure 2.2 (b) and Figure 2.5 (b). Now suppose $k_2 = 0$

and $k_1 > 1$. Then u_{v_2} is a multiple cover of the sphere $E''_{1,J_{v_2}}$ (using notation from (2.8)). Besides the edge e to v_1, v_2 has at least one more edge to a vertex, say $\overline{v}_3$, with $P(\overline{v}_3) = P_1$, see Figure 2.6. Indeed, at the node $w_{e'}$ corresponding to the edge $e' := (v_1, v_2)$, the map u_{v_2} has an intersection of order 1 with the relative divisor $X_{P_{12}}$, and more edges are required to account for the k intersections that u_{v_2} has with $X_{P_{12}}$. The homology class of u_{glue} in (2.9) implies that the homology class of the map $u_{\overline{v}_3}$ is a multiple of $[E'_1]$. For generic background almost complex structures $J^b_{v_2}$, $J^b_{\overline{v}_3}$ and small enough domain-dependent perturbations, the evaluations

$$\text{ev}_{w_e}(u_{v_2}), \quad \text{ev}_{w_e}(u_{\overline{v}_3}) \in X_{P_{12}}$$

will not agree, see Remark 2.7 which explains a similar result. Therefore matching conditions can not be simultaneously satisfied at the node (v_1, v_2) and $(v_1, \overline{v}_3)$. Thus the possibility that $k_1 > 1$ is ruled out when $k_2 = 0$.

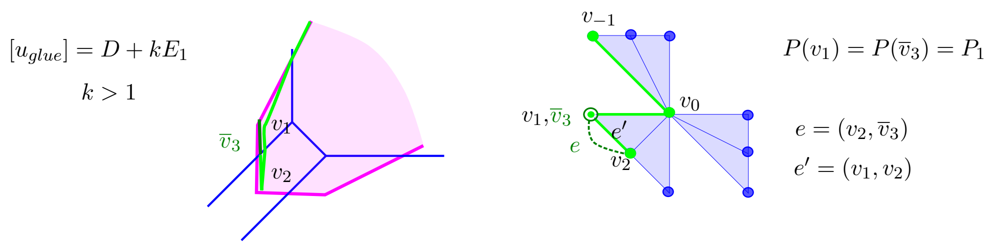

Figure 2.6. A hypothetical broken map (ruled out in proof of Proposition 2.2).

Next, assuming $k_2 > 0$, we show that $k_1 = k_2$. The argument is similar to the one used in Step 2. For example, if $k_2 > k_1$, then $[u_{v_2}].E'_2 < 0$. So the image of u_{v_2} lies in the J_{v_2}-holomorphic sphere homologous to E'_2, contradicting the fact that $k_1 \geq 1$. The case $k_1 > k_2$ is ruled out by a symmetric argument.

Next, we rule out the case $k_1 = k_2 > 1$ using the genericity of J. Gluing the (-1)-spheres $E''_{1,J_{v_2}}$ and $E'_{2,J_{v_2}}$ produces a J_{v_2}-holomorphic sphere with a trivial normal bundle in X_{P_2}. Then there is a fibration

$$\pi : X_{P_2} \backslash (X_{P_{02}} \cup E''_{1,J_{v_2}} \cup E'_{2,J_{v_2}}) \to \mathbb{C}^\times$$

where the fibers are J_{v_2}-holomorphic spheres. When $k_1 = k_2$, the map u_{v_2} is a cover of one of the fibers of π. We define a diffeomorphism

$$f : X_{P_{12}} \backslash X_{P_{012}} \to X_{P_{23}} \backslash X_{P_{023}}$$

by the condition that x and $f(x)$ lie on the same fiber of π. As we argued in the previous paragraph, $k_1 > 1$ implies that v_2 has an edge to a vertex, say, $\overline{v}_3$ with $P(\overline{v}_3) = P_1$ and $u_{\overline{v}_3}$ is a cover of the sphere $E'_{1,J_{\overline{v}_3}}$ (see Figure 2.7). Since $k_2 > 0$, v_2 also has an edge to a vertex, say, $\overline{v}_4$ with $P(\overline{v}_4) = P_3$ and $u_{\overline{v}_4}$ is a cover of the sphere $E''_{2,J_{\overline{v}_4}}$. For generic background almost complex structures $J_{\overline{v}_3}$, $J_{\overline{v}_4}$,

$$f(E'_{1,J_{\overline{v}_3}} \cap X^{\square}_{P_{12}}) \neq E''_{2,J_{\overline{v}_4}} \cap X^{\square}_{P_{23}}.$$

The domain-dependent perturbations are small enough that the evaluations $\mathrm{ev}_{w_{e''}}(u_{v_2})$, $\mathrm{ev}_{w_{e''}}(u_{\overline{v}_4}) \in X_{P_{23}}$ lie in disjoint neighborhoods of $f(E'_{1,J_{\overline{v}_3}} \cap X_{P_{12}})$, $E''_{2,J_{\overline{v}_4}} \cap X_{P_{23}}$. Consequently there is no broken map u for which $[u_{\mathrm{glue}}] = D_1 + k(E_1 + E_2)$, $k > 1$.

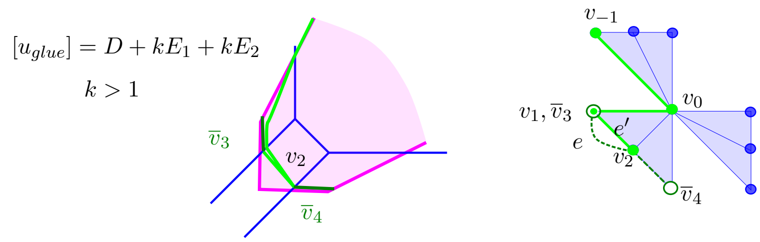

FIGURE 2.7. A hypothetical broken map (ruled out in proof of Proposition 2.2).

Finally we are left with the case $k_1 = k_2 = 1$. In this case u_{v_2} is a simple map to a fiber of π, and v_2 has an edge to a vertex v_3 with $P(v_3) = 0$.

STEP 4 (v_3) : So far we have shown that if a broken map u whose gluing has homology class (2.9), then either $[u_{\mathrm{glue}}] = [D_1]$ or $[u_{\mathrm{glue}}] = [D_1] + [E_1]$ or the tropical graph of u has vertices v_{-1}, v_0, v_1, v_2, and v_2 has an edge of multiplicity 1 to a vertex v_3 with $P(v_3) = P_3$. In this last case, the only possibility is that u_{v_3} is a simple map of class $[E''_2]$. The resulting broken map, of class $[D_1] + [E_1] + [E_2]$ is shown in Figure 2.2 (c) and Figure 2.5 (c). Multiple covers of $[E''_2]$ are ruled out using arguments similar to those in Step 2: If u_{v_3} is a multiple cover of E''_2 then there is a path P in Γ emanating out of v_3 (which does not contain v_2) whose end-point is a (-1)-curve. Let $e \ni v_3$ be the first edge in the path P. Since the path ends in a (-1)-curve, there is a point constraint on u_{v_3} at the node w_e corresponding to e. For a generic background J and a small enough perturbation, the point constraint will be distinct from $E''_{2,J_{v_3}} \cap X_{P_{23}}$, contradicting the existence of such a multiple cover $m[E''_2]$ in the broken map.

Finally we show that each tropical graph in Figure 2.2 contributes $+1$ to the curve count. Since the sphere u_v corresponding to each vertex $v \in \mathrm{Vert}(\Gamma)$ is pseudoholomorphic with respect to an integrable almost complex structure J_v^b (see (2.7)), and matching conditions are cut out by diagonals on the relative divisors, each of the broken maps has positive orientation in the moduli space. In each of the cases, there is no automorphism of the tropical graph. Therefore, a broken map with n interior markings contributes $\frac{1}{n!}$ to the curve count. Assuming that the components of the broken map have $n_1, \ldots, n_m$ markings (that add up to n), permuting the labellings of marked points within each component gives $(n_1! \ldots n_m!)$ curves. Finally there are $\frac{n!}{n_1! \ldots n_m!}$ ways of assigning interior markings to components – each of these would give rise to a different almost complex structure. Our proof shows that in all the cases, the number of broken maps corresponding to each tropical graph is unaffected. This finishes the proof of the curve count in broken maps. Since Gromov-Witten invariants are not altered by neck-stretching and deformation of the symplectic form, we have shown that there are 27 lines on a cubic surface. □

The following is the proof of the disk count on a cubic surface.

Proof of Proposition 2.5. As in the proof of the sphere count, we first count the number of broken disks of Maslov index 2 in the broken manifold $\mathfrak{X}$, and then argue that the disk count is the same for the cubic surface. We recall that the broken manifold $\mathfrak{X}$ is obtained by first deforming the symplectic form on the cubic surface to produce a toric manifold $X := X^\epsilon$, and then by multiply-cutting X.

The homology classes of the disks of Maslov index two in X (equipped with a torus-invariant almost complex structure) must consist of a single Blaschke product (i.e. a disk having exactly one intersection with one of the toric divisors) and a collection of short divisors, each of which have Chern number zero:

$$H_2(X,L) \ni [u] = \begin{cases} \delta_{long} & \text{or} \\ \delta_{short} + \sum_i a_i[E_i], \end{cases} \tag{2.11}$$

where δ_{long} resp. $\delta_{short} \in H_2(X,L)$ is the class of a disk which is a single Blaschke product intersecting a long resp. short divisor, $a_i \in \mathbb{Z}_+$, and any $E_i \in H_2(X)$ is the homology class of a short divisor.

A semi-Fano toric surface has a well-defined potential (see Proposition 12.7). Further the potential is unchanged under neck-stretching if the broken manifold, with the standard almost complex structure (componentwise torus-invariant) does not have disks of Maslov index zero, see Proposition 12.9. This condition is indeed satisfied in the multiple cut we consider on the cubic surface, see Proposition 12.11. Therefore, if J_0 resp. $\mathfrak{J}_0$ is the standard almost complex structures on X resp. $\mathfrak{X}$, for any regular domain-dependent almost complex structure J resp. $\mathfrak{J}$ close enough to J_0 resp. $\mathfrak{J}_0$, the potential on (X,J) and $(\mathfrak{X},\mathfrak{J})$ is the same. In particular, the contribution to the potential from any disk homology class $\beta \in H_2(X,L)$ is the same for (X,J) and $(\mathfrak{X},\mathfrak{J})$. Thus we may assume that the homology class of any broken disk in $\mathfrak{X}$ is as in (2.11).

Let u be a broken disk in $\mathfrak{X}$ whose gluing has relative homology class

$$[u_{\text{glue}}] = [\delta_{E_1}] + \sum_i n_i[E_i] \in H_2(X,L),$$

where $[\delta_{E_1}]$ is the class of a Maslov index two disk in (X,L) that intersects the short divisor E_1 (using notation in Figure 2.3). The other cases where δ_{E_1} is replaced by δ_{E_i}, $i = 2,\dots,6$, can be analyzed similarly. Now we need to show that the gluing of u is either (b), (c) or (d) in Figure 2.3. Since the class $[\partial u] \in H_1(L)$ is preserved by neck stretching, we conclude that the disk component of u is a Maslov index two disk in X_{P_0} that intersects the relative divisor $X_{P_{01}}$. The rest of the proof that u is of the expected form is exactly same as Steps 2, 3, 4 in the proof of the sphere count (Proposition 2.2). Each of the tropical graphs contributes $+1$ to the curve count: Using the standard spin structure, isolated Blaschke disks have a positive orientation sign (see [**20**, p22]), the maps u_v are pseudoholomorphic with respect to an integrable almost complex structure J_v^b (see (2.7)), and matching conditions are cut out by diagonals on the relative divisors, each of the broken maps has positive orientation in the moduli space. The combinatorics involving marked points carries over from the sphere case.

We have so far proved that for any regular perturbation $\underline{\mathfrak{p}}$ of the standard almost complex structure J_0 on the toric manifold X, the potential of $CF(X, L, \mathfrak{p})$ has twenty one terms. It remains to show that the potential is the same for the cubic surface. There is a family of Kähler manifolds $\{(M_t, \omega_t, J_t^M)\}_{t\in[0,1]}$, each of whose elements is the monotone blow-up of the toric del Pezzo $\mathrm{Bl}_3\,\mathbb{P}^2$ at three points $z_{4,t}, z_{5,t}, z_{6,t} \in \mathrm{Bl}_3\,\mathbb{P}^2$. For $t > 0$ the points $z_{4,t}, z_{5,t}, z_{6,t}$ are generic, and so, (M_t, J_t^M) is a cubic surface. For $t = 0$ the points $z_{4,0}, z_{5,0}, z_{6,0}$ are torus fixed-points, corresponding to alternate vertices on the moment polytope of $\mathrm{Bl}_3\,\mathbb{P}^2$. Since the elements of the family are symplectomorphic to each other, we may view J_t^M as a family of compatible complex structures on (M_0, ω_0). On the other hand, there is a family of semi-Fano toric manifolds $\{X^\epsilon\}_{\epsilon>0}$ (with varying symplectic forms) in Figure 2.1 such that the limit X^0 is almost toric and isomorphic to (M_0, ω_0, J_0^M) as a Kähler manifold. Via this isomorphism, the monotone toric Lagrangian $L \subset X^0$ is identified to a Lagrangian $L \subset M_0$. The potential $CF(M_0, L, J_t^M)$ is t-independent by monotonicity of L, and is equal to the semi-Fano toric potential of $CF(X^\epsilon, L, J_0)$ in the limit $\epsilon \to 0$. Finally we recall from (2.3) that X was defined to be X^ϵ for a small $\epsilon > 0$. This finishes the proof of Proposition 2.5 on the count of disks on a cubic surface. □

REMARK 2.9. Proposition 2.5 implies that the potential of the monotone toric Lagrangian on the cubic surface has twenty one terms coming from the disks shown in Figure 2.3. We obtain the potential

$$W(y_1, y_2) = y_1y_2 + y_1/y_2^2 + y_2/y_1^2 + 3(y_1 + y_1/y_2 + y_2 + y_2/y_1 + 1/y_1 + 1/y_2)$$

originally obtained by Pascaleff-Tonkonog [**77**] using mutations. This formula is equivalent to the one in [**77**, Table 1] by the change of variables

$$z_1 = y_1^{-1}, z_2 = y_2^{-1}$$

after which the formula becomes

$$W(z_1, z_2) = (1 + z_1 + z_2)^3/z_1z_2 - 6.$$

The formula is also obtained in Galkin-Usnich [**36**]; c.f. [**7**, Table 5.1].

2.2. Counting disks in flag varieties

In this section we use our results to compute the disk potential of partial flag varieties, without using small resolutions which were the original approach of Nishinou-Nohara-Ueda [**63**] and Nohara-Ueda [**64**]. Given a tuple $\underline{n} := (n_1, \dots, n_r)$ with $\sum_i n_i = n$, a *flag variety* $\mathrm{Fl}(n_1, \dots, n_r)$ is the set of flags

$$\mathrm{Fl}(\underline{n}) := \mathrm{Fl}(n_1, \dots, n_r) := \{\{0\} \subset V_{n_1} \subset V_{n_1+n_2} \subset \dots V_{n_1+\dots+n_r} = \mathbb{C}^n\}.$$

A flag variety may be identified with a co-adjoint orbit of the action of $U(n)$ on $u(n)^\vee$. In particular, $\mathrm{Fl}(\underline{n})$ is the coadjoint orbit of an element

$$\mathrm{diag}(\iota(\underbrace{\Lambda_1, \dots, \Lambda_1}_{n_1 \text{ times}}, \underbrace{\Lambda_2, \dots, \Lambda_2}_{n_2 \text{ times}}, \dots, \underbrace{\Lambda_r, \dots, \Lambda_r}_{n_r \text{ times}})) \in \mathfrak{u}(n)^\vee,$$

where

$$\Lambda_i \in \mathbb{R}, \quad \Lambda_1 > \cdots > \Lambda_r,$$

and $\mathfrak{u}(n)$ is identified to its dual via the $\mathrm{Ad}_{U(n)}$-invariant inner product

$$(\cdot,\cdot) : \mathfrak{u}(n) \times \mathfrak{u}(n) \to \mathbb{R}, \quad (A, B) \mapsto \mathrm{trace}(A^* B).$$

We recall the construction of the Gelfand Cetlin system on a flag variety from Guillemin-Sternberg [**47**]. The *Gelfand-Cetlin system* is a collection of functions that form a completely integrable system on an open set of $\mathrm{Fl}(\underline{n})$. For any $m = 1, \dots, n$, let

$$\pi_m : \mathrm{Fl}(\underline{n}) \to \mathfrak{u}(m)^\vee$$

be the projection obtained by taking the upper left $m \times m$ submatrix. The eigenvalues of π_m, which are necessarily imaginary, are denoted by $\iota\Phi_{m,k}$, $1 \le k \le m$, where

$$\Phi_{m,1} \ge \Phi_{m,2} \ge \ldots \ge \Phi_{m,m}.$$

Together, the collection of maps $\{\iota\Phi_{i,k}\}_{1 \le k \le i \le n}$ define a map

$$\Phi : \mathrm{Fl}(\underline{n}) \to \mathfrak{t}^\vee := (\iota\mathbb{R})^{n(n-1)/2}, \tag{2.12}$$

For any i, the eigenvalue functions $\Phi_{i,j}$ are smooth functions whenever they are distinct and are continuous everywhere. The eigenvalues satisfy the interlacing inequalities

$$\Phi_{i,k} \in [\Phi_{i+1,k}, \Phi_{i,k+1}], \quad \forall i, k < n$$

by the min-max theorem for Hermitian matrices [**89**, Theorem 8.10], and they fit into the following diagram.

$$\begin{array}{ccccccccccc}
\Phi_1 & & \Phi_2 & & \Phi_3 & & \cdots & \Phi_{n-1} & & \Phi_n \\
& \rotatebox{-45}{$\geq$} \; \rotatebox{45}{$\leq$} & & \rotatebox{-45}{$\geq$} \; \rotatebox{45}{$\leq$} & & & & & \rotatebox{-45}{$\geq$} \; \rotatebox{45}{$\leq$} & \\
& \Phi_{n-1,1} & & \Phi_{n-1,2} & & & & & \Phi_{n-1,n-1} & \\
& & \rotatebox{-45}{$\geq$} \; \rotatebox{45}{$\leq$} & & & & & \rotatebox{45}{$\leq$} & & \\
& & \Phi_{n-2,1} & & & & & \Phi_{n-2,n-2} & & \\
& & & \rotatebox{-45}{$\geq$} & & & \rotatebox{45}{$\leq$} & & & \\
& & & & \ddots & & \iddots & & & \\
& & & & \rotatebox{-45}{$\geq$} & & \rotatebox{45}{$\leq$} & & & \\
& & & & & \Phi_{1,1} & & & &
\end{array} \tag{2.13}$$

We give a proof of the min-max theorem for completeness.

Lemma 2.10. (Interlacing inequality) *Let A be an $n \times n$ Hermitian matrix with eigen values $a_1 \ge \cdots \ge a_n$. Suppose the upper left $(n-1) \times (n-1)$ submatrix $\pi_{n-1}(A)$ has eigenvalues $b_1 \ge \cdots \ge b_{n-1}$. Then,*

$$a_1 \ge b_1 \ge a_2 \ge \cdots \ge a_{n-1} \ge b_n \ge a_n.$$

PROOF. It is enough to prove the result assuming that $b_1 > \cdots > b_n$, since the general result follows by continuity. The $\mathrm{Ad}_{U(n)}$-orbit of A contains a matrix

$$A' := \begin{pmatrix} b_1 & & 0 & \bar{z}_1 \\ & \ddots & & \vdots \\ 0 & & b_n & \bar{z}_n \\ z_1 & \dots & z_n & r_{n+1}, \end{pmatrix}.$$

$$(2.14)\quad \det(A - t\,\mathrm{Id}) = \det(A' - t\,\mathrm{Id}) = \prod_{i=1}^{n}(b_i - t) f(t), \quad f(t) = (r_{n+1} - t - \sum_{j=1}^{n} \frac{|z_j|^2}{b_j - t}),$$

and $f(t)$ is decreasing in t and discontinuous at $b_1, \dots, b_n$. Since the b_js are distinct $f(t)$ has exactly one root each in the intervals $(-\infty, b_n)$, (b_{i+1}, b_i), $i = 1, \dots, n-1$, (b_1, ∞). □

For a partial flag manifold, certain values in the Gelfand-Cetlin system are fixed by the interlacing inequalities. In particular, for $k \le i \le n-1$,

$$\Phi_k = \Phi_{n+k-i} = \Lambda_j \implies \Phi_{i,k} = \Lambda_j.$$

The Gelfand Cetlin system thus consists of exactly

$$N := \sum_{1 \le i < j \le r} n_i n_j$$

variables, which is half the dimension of $\mathrm{Fl}(\underline{n})$. See Figure 2.8. We denote the set of variables in the Gelfand-Cetlin system of the partial flag variety $\mathrm{Fl}(\underline{n})$ by

$$\mathrm{Free}(\underline{n}) := \{(i,j) : \Phi_{i,j} \text{ is not a constant on } \mathrm{Fl}(\underline{n})\}.$$

The image $\Phi(\mathrm{Fl}(\underline{n}))$ is contained in a polytope

$$\Delta_{\mathrm{Fl}(\underline{n})} \subset (\iota\mathbb{R})^N$$

cut out by the Gelfand-Cetlin inequalities involving free variables $\Phi_{i,j}$, $(i,j) \in \mathrm{Free}(\underline{n})$, which we call the *Gelfand-Cetlin polytope of* $\mathrm{Fl}(\underline{n})$. We denote by

$$F^{reg} \subset \mathrm{Fl}(\underline{n})$$

the set of points x for which for every $m < n$ the set of eigenvalues of $\pi_m(x)$ has the maximum cardinality. In other words, $x \in F^{reg}$ exactly when the following holds: For all positive integers $i \le n-1$, $k < i$,

$$\Phi_{i,k}(x) = \Phi_{i,k+1}(x) \Leftrightarrow (i,k), (i,k+1) \notin \mathrm{Free}(\underline{n}).$$

REMARK 2.11. In the Gelfand-Cetlin system corresponding to the flag variety $\mathrm{Fl}(n_1, \dots, n_r)$ the non-free variables can be partitioned into r disjoint sets

$$\{(i,k) : 1 \le k \le i \le n-1\} \setminus \mathrm{Free}(\mathrm{Fl}(\underline{n})) = N_1 \cup \dots N_r$$

where N_j consists of $\Phi_{i,k}$ which are fixed to be equal to Λ_j by the interlacing inequalities. We call N_j a *block* of non-free variables. A block N_j is non-empty if $n_j > 1$.

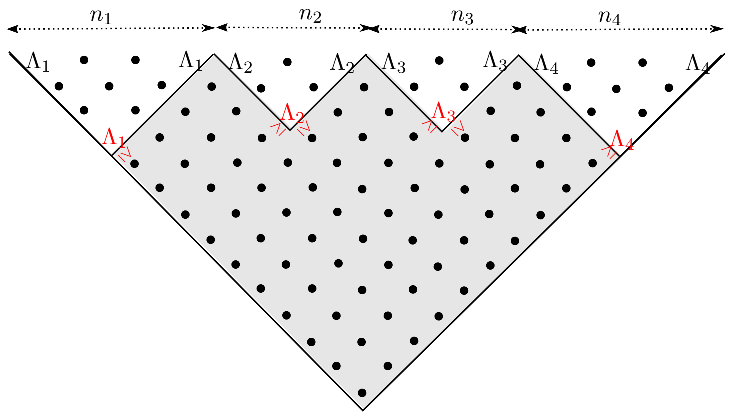

FIGURE 2.8. For a partial flag, the variables of the Gelfand-Cetlin system lie in the shaded region.

Each non-empty block is triangular and has a *lowest* element whose location in the Gelfand-Cetlin diagram is the lowest. The lowest element in the block N_j is

$$\Phi_{\mathfrak{i}_j,\mathfrak{k}_j}, \quad \text{where } (\mathfrak{i}_j,\mathfrak{k}_j) = \left(n - n_j + 1, \sum_{\ell=1}^{j} n_\ell + 1\right). \tag{2.15}$$

In Figure 2.8 the lowest position in each block is in red.

PROPOSITION 2.12. ([**47**]) *Let $\Phi : \mathrm{Fl}(\underline{n}) \to \mathbb{R}^N$ be the Gelfand-Cetlin system on the flag variety.*

(a) The restriction $\Phi|F_{reg}$ is the moment map for the action of a torus $\mathbb{T}^N$.
(b) The map Φ surjects onto the polytope $\Delta_{\mathrm{Fl}(\underline{n})}$.
(c) For any point λ in the interior of $\Delta_{\mathrm{Fl}(\underline{n})}$, $\Phi^{-1}(\lambda)$ is a Lagrangian torus.

PROOF. Any pair of functions in the Gelfand-Cetlin system Poisson commute :

$$\{\Phi_{i,j}, \Phi_{k,l}\} = 0$$

by [**63**, Lemma 3.3]. So, the system $\{\Phi_{i,k}\}_{i,k}$ generates a Hamiltonian $\mathbb{R}^n$-action on $\mathrm{Fl}(\underline{n})$. Any non-constant Hamiltonian trajectory generated by $\Phi_{i,k}$ has a period of 2π by [**47**, Lemma 3.4]. Consequently the system $\{\Phi_{i,k}\}_{i,k}$ generates an action of the torus $(\mathbb{R}/2\pi\mathbb{Z})^N$ on F^{reg}.

An inductive application of Corollary 2.14 shows that Φ surjects onto $\Delta_{\mathrm{Fl}(\underline{n})}$. Further, the rank of the differential $d\Phi_x$ is equal to N minus the number of Gelfand-Cetlin inequalities involving variables in $\mathrm{Free}(\underline{n})$ that are equalities for x. In particular, if $\lambda := \Phi(x)$ is in the interior of $\Delta_{\mathrm{Fl}(\underline{n})}$ the rank of $d\Phi(x)$ is N, and therefore the action of $\mathbb{T}^N$ is free. Consequently the orbit $\mathbb{T}^N x$ is a Lagrangian torus. Corollary 2.14, when inductively applied, shows that the inverse image $\Phi^{-1}(\lambda)$ is a torus $\mathbb{T}^N$. We may conclude that $\Phi^{-1}(\lambda)$ coincides with $\mathbb{T}^N x$, and is therefore, Lagrangian. □

The fact that the fibers of a complete integrable system are Lagrangian tori is the content of the Arnol'd-Liouville theorem and is proved in [**27**]. In the case of

Gelfand-Cetlin systems, we explicitly show that the fibers of the moment map are tori via Lemma 2.13 and Corollary 2.14 below. Lemma 2.13 is also stated without proof in [**63**] (Lemma 3.5).

LEMMA 2.13. *Let*

$$b_1 > b_2 > \cdots > b_n \tag{2.16}$$

be a set of real numbers and let

$$\Delta_b := [b_1, \infty) \times [b_2, b_1] \times \cdots \times [b_n, b_{n-1}] \times (-\infty, b_n].$$

There is a continuous map

$$r : \Delta_b \to (\mathbb{R}_{\geq 0})^n \times \mathbb{R}, \quad (a_1, \ldots, a_{n+1}) \mapsto (r_1, \ldots, r_n, r_{n+1})$$

such that for any $a = (a_1, \ldots, a_{n+1}) \in \Delta_b$*, a matrix*

$$A := \begin{pmatrix} b_1 & & 0 & \bar{z}_1 \\ & \ddots & & \vdots \\ 0 & & b_k & \bar{z}_k \\ z_1 & \ldots & z_k & r_{k+1} \end{pmatrix}$$

satisfying $|z_i| = r_i$ *for* $i = 1, \ldots, n$ *has eigenvalues* $a_1, \ldots, a_{n+1}$*. The map* r *is smooth on each stratum of* Δ_b*, and* $r_i = 0$ *exactly when either* $b_i = a_i$ *or* $b_i = a_{i+1}$*.*

PROOF. We first consider the open stratum of Δ_b, that is, we assume that

$$a_1 < b_1 < a_2 < \cdots < a_n < b_n < a_{n+1}.$$

We observe that

$$\det(A - t\,\mathrm{Id}) = \prod_{i=1}^{k}(b_i - t)\left(r_{k+1} - t - \sum_{j=1}^{k}\frac{|z_j|^2}{b_j - t}\right). \tag{2.17}$$

We set

$$r_{k+1} := \sum_i a_i - \sum_i b_i,$$

and solve the system of equations

$$\det(A - a_i\,\mathrm{Id}) = 0 \quad \text{for } i = 1, \ldots, k \tag{2.18}$$

for $z_1, \ldots, z_k$. Substituting $|z_i|^2 = r_i$, the system (2.17) is linear in $r_1, \ldots, r_n$, and the coefficient matrix for the variables r_i, $i = 1, \ldots, n$ is

$$\left(\prod_{1 \leq k \leq n : k \neq j}(b_k - a_i)\right)_{i,j}$$

whose determinant is $\pm\prod_{i<j}(a_i - a_j)\prod_{i<j}(b_i - b_j) \neq 0$. Therefore, we can uniquely solve (2.17) for $(r_1, \ldots, r_n)$.

It remains to show that $r_1, \ldots, r_n$ are non-negative. Firstly, we observe that $r^{-1}(\partial((\mathbb{R}_{\geq 0})^n \times \mathbb{R})) \subset \partial\Delta_b)$. Indeed, in the matrix A, if $z_i = 0$ then b_i is an eigenvalue, and so b_i is equal to either a_i or a_{i+1}. Secondly, by the interlacing property the image $r(\Delta_b^\circ)$ of the interior Δ_b° of Δ_b intersects the interior of $(\mathbb{R}_{\geq 0})^n \times \mathbb{R}$. We may then conclude that $r(\Delta_b^\circ)$ is in fact contained in $(\mathbb{R}_+)^n \times \mathbb{R}$.

Finally we define the map r on lower dimensional strata of Δ_b. if $b_i = a_i$ resp. $b_i = a_{i+1}$, then we set $r_i = 0$. We apply the above construction of r with b_i deleted from the list of eigenvalues, and obtain a smooth map that takes $(a_1, \dots, \hat{a}_i, \dots, a_{n+1})$ resp. $(a_1, \dots, \hat{a}_{i+1}, \dots, a_{n+1})$ to $(r_1, \dots, \hat{r}_i, \dots, r_{n+1})$. □

Corollary 2.14. *Let B be an $(n-1)\times(n-1)$ Hermitian matrix with eigenvalues $b_1 \geq \cdots \geq b_{n-1}$. Consider the map*

$$\Phi^B : \{A \text{ is } n \times n \text{ Hermitian } : \pi_{n-1}(A) = B\} \to \mathbb{R}^n$$

that maps A to its eigenvalues $\Phi_1^B(A), \dots, \Phi_n^B(A)$ arranged in decreasing order. For any A in the domain of Φ^B, let

$$n^B(A) := \#\{1 \leq i \leq n : b_i \text{ is not an eigenvalue of } A\}.$$

Then,

- *(a) The map Φ^B surjects onto $\Delta_b := [b_1 \times \infty) \times [b_2, b_1] \times \cdots \times [b_{n-1}, b_{n-2}] \times (-\infty, b_{n-1}] \subset \mathbb{R}^n$.*
- *(b) For any A in the domain of Φ, $\operatorname{rank}(d\Phi^B(A)) = n^B(A)$,*
- *(c) and $(\Phi^B)^{-1}(\Phi^B(A)) = \mathbb{T}^{n^B(A)}$.*

Proof. It is enough to prove the result assuming that B is a diagonal matrix. If the eigenvalues of B are distinct then the result is a consequence of Lemma 2.13 Indeed, $\Phi^B(A)$ lies on an $n^B(A)$-dimensional stratum S of Δ_b, and $\Phi^B|(\Phi^B)^{-1}(S)$ has a smooth right inverse. Therefore $\operatorname{rank}(d\Phi^B(A)) = n^B(A)$. In the matrix A the number of z_is that are non-zero is exactly $n^B(A)$, and therefore the inverse image $(\Phi^B)^{-1}(\Phi^B(A)$ is $\mathbb{T}^{n^B(A)}$. Next consider the case when B has repeated eigenvalues. If $b_i = b_{i+1}$, then $\Phi_{i+1}^B(A) = b_i$ and in the matrix A, $z_i = 0$. The result can therefore be obtained by applying Lemma 2.13 to $B' = \operatorname{diag}\{b_1, \dots, \hat{b}_i, \dots, b_{n-1}\}$. □

We recall some features of $\Delta_{\mathrm{Fl}(\underline{n})}$ the Gelfand-Cetlin polytope of the flag variety $\mathrm{Fl}(\underline{n})$:

- (a) (Interior) A point $\Phi \in \Delta_{\mathrm{Fl}(\underline{n})}$ is an interior point exactly if any Gelfand-Cetlin inequality involving at least one free variable $\Phi_{i,j}$ is strict.
- (b) (Facet) Faces of the Gelfand-Cetlin polytope $\Delta_{\mathrm{Fl}(\underline{n})}$ can be characterized by face graphs as in [**3**]. We explicitly list the set of facets (codimension one faces): A subset

$$\{\Phi_{i,k} = \Phi_{i+1,k}\} \quad \text{resp.} \quad \{\Phi_{i,k} = \Phi_{i+1,k+1}\} \tag{2.19}$$

is a facet of $\Delta_{\mathrm{Fl}(\underline{n})}$ if either both variables in the equality are free or the non-free variable is the lowest variable in a block of non-free variables (as defined in Remark 2.11). We remark that if the non-free variable in (2.19) were not the lowest in the block, then the codimension of the subset would be more than 1, because the interlacing inequalities would fix the values of some free variables not occuring in (2.19). We index the facet in (2.19) by

$$(i, k, \searrow) \quad \text{resp.} \quad (i, k, \nearrow)$$

and denote the set of indices of facets by

$$\mathrm{Facets}(\Delta_{\mathrm{Fl}(\underline{n})}) \subset \mathbb{Z}_+^2 \times \{\diagdown, \diagup\}.$$

(c) (Singular points) The complement of the regular set $F^{reg} \subset \mathrm{Fl}(\underline{n})$ is mapped by Φ to non-simplicial points of the polytope $\Delta_{\mathrm{Fl}(\underline{n})}$ as we now explain. Consider $x \in \mathrm{Fl}(\underline{n}) \backslash F^{reg}$ and suppose $\Phi(x)$ lies in the interior of the face $\mathcal{F}$ of $\Delta_{\mathrm{Fl}(\underline{n})}$. Then $\mathcal{F}$ is equal to the intersection of $n_{\mathcal{F}}$ facets where $n_{\mathcal{F}} > \mathrm{codim}(\mathcal{F})$ as follows: The set of Gelfand-Cetlin inequalities that are equalities on $\mathcal{F}$ contains a loop

$$\begin{array}{ccccc} & & \Phi_{i+1,k+1} & & \\ & \rotatebox{45}{$=$} & & \rotatebox{-45}{$=$} & \\ \Phi_{i,k} & & & & \Phi_{i,k+1} \\ & \rotatebox{-45}{$=$} & & \rotatebox{45}{$=$} & \\ & & \Phi_{i-1,k} & & \end{array} \tag{2.20}$$

corresponding to every pair of equal eigenvalues $\Phi_{i,k} = \Phi_{i,k+1}$. Any such loop of equalities gives a subset S of codimension three. However, S is the intersection of four facets corresponding to each of the equalities in the loop. As an aside, we remark that S is the intersection of four faces of codimension two, each given by a pair of equalities

$$\begin{array}{ccccc} & & \Phi_{i+1,k+1} & & \\ & \rotatebox{45}{$\leq$} & & \rotatebox{-45}{$=$} & \\ \Phi_{i,k} & & & & \Phi_{i,k+1} \\ & \rotatebox{-45}{$=$} & & \rotatebox{45}{$\leq$} & \\ & & \Phi_{i-1,k} & & \end{array} \quad \begin{array}{ccccc} & & \Phi_{i+1,k+1} & & \\ & \rotatebox{45}{$=$} & & \rotatebox{-45}{$\geq$} & \\ \Phi_{i,k} & & & & \Phi_{i,k+1} \\ & \rotatebox{-45}{$\geq$} & & \rotatebox{45}{$=$} & \\ & & \Phi_{i-1,k} & & \end{array} \quad \begin{array}{ccccc} & & \Phi_{i+1,k+1} & & \\ & \rotatebox{45}{$\leq$} & & \rotatebox{-45}{$=$} & \\ \Phi_{i,k} & & & & \Phi_{i,k+1} \\ & \rotatebox{-45}{$\geq$} & & \rotatebox{45}{$=$} & \\ & & \Phi_{i-1,k} & & \end{array} \quad \begin{array}{ccccc} & & \Phi_{i+1,k+1} & & \\ & \rotatebox{45}{$=$} & & \rotatebox{-45}{$\geq$} & \\ \Phi_{i,k} & & & & \Phi_{i,k+1} \\ & \rotatebox{-45}{$=$} & & \rotatebox{45}{$\leq$} & \\ & & \Phi_{i-1,k} & & \end{array}$$

EXAMPLE 2.15. The full flag $F^{(3)} = \mathrm{Fl}(1,1,1)$ resp. the Grassmannian $G(2,4) = \mathrm{Fl}(2,2)$ has a singular point resp. a singular line given by a loop of inequalities

$$\begin{array}{ccccc} & & \Lambda_2 & & \\ & \rotatebox{45}{$=$} & & \rotatebox{-45}{$=$} & \\ \Phi_{2,1} & & & & \Phi_{2,2} \\ & \rotatebox{-45}{$=$} & & \rotatebox{45}{$=$} & \\ & & \Phi_{1,1} & & \end{array} \quad \text{resp.} \quad \begin{array}{ccccc} & & \Phi_{3,2} & & \\ & \rotatebox{45}{$=$} & & \rotatebox{-45}{$=$} & \\ \Phi_{2,1} & & & & \Phi_{2,2}. \\ & \rotatebox{-45}{$=$} & & \rotatebox{45}{$=$} & \\ & & \Phi_{1,1} & & \end{array} \tag{2.21}$$

In the rest of the Section, we prove the following result, which was originally proved by Nishinou-Nohara-Ueda [**63**].

THEOREM 2.16. *Let $X := \mathrm{Fl}(\underline{n})$ be a flag variety. There is a symplectic form on X for which the Gelfand-Cetlin system of X has a unique monotone Lagrangian fiber L_0. The A_∞ algebra $CF(X, L_0)$ is unobstructed and the potential of $CF(X, L_0)$ consists of one term for every facet of the Gelfand-Cetlin polytope of X.*

As part of the proof of Theorem 2.16 we compute the potential of $\mathrm{Fl}(\underline{n})$ as

$$W(y_{i,k})_{i,k} = \left(\sum_{k\leq i:(i+1,k)\in \mathrm{Free}(\mathrm{Fl}(\underline{n}))} \frac{y_{i+1,k}}{y_{i,k}} + \sum_{k\leq i:(i+1,k+1)\in \mathrm{Free}(\mathrm{Fl}(\underline{n}))} \frac{y_{i,k}}{y_{i+1,k+1}} + \sum_{1\leq j\leq r-1} y_{\mathfrak{i}_j,\mathfrak{k}_j} + \sum_{2\leq j\leq r} \frac{1}{y_{\mathfrak{i}_j,\mathfrak{k}_j}} \right) q. \tag{2.22}$$

where the terms $\mathfrak{i}_j$, $\mathfrak{k}_j$ are introduced later.

Remark 2.17. In the case of a full flag variety ($r = n$), the least order terms of the potential in (2.22) coincide with the potential introduced by Givental in [**39**]. In the case of the Grassmannian ($r = 2$), the analysis of the potential (2.22) by Castronovo in [**14**] implies that the Lagrangian L_0 has non-trivial Floer homology.

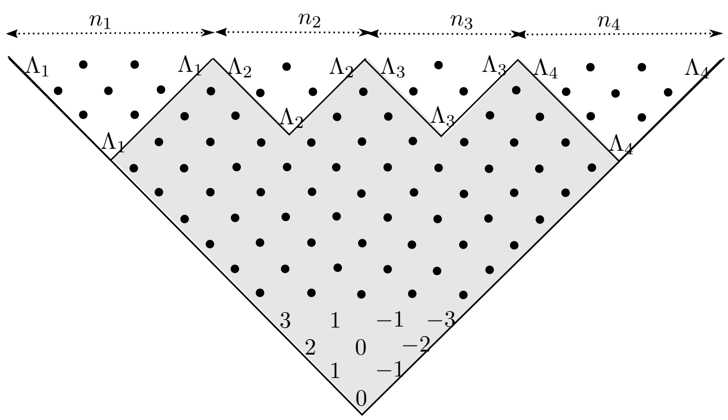

Figure 2.9. Center of the Gelfand-Cetlin polytope $\Delta_{\mathrm{Fl}(\underline{n})}$.

Proof of Theorem 2.16. First we describe the monotone symplectic form and the monotone Lagrangian in the flag variety $\mathrm{Fl}(\underline{n})$. We define $\mathrm{Fl}(\underline{n})$ to be the coadjoint $U(n)$-orbit of the element

$$\mathrm{diag}(\iota(\underbrace{\Lambda_1,\dots,\Lambda_1}_{n_1 \text{ times}},\underbrace{\Lambda_2,\dots,\Lambda_2}_{n_2 \text{ times}},\dots,\underbrace{\Lambda_r,\dots,\Lambda_r}_{n_r \text{ times}})) \in \mathfrak{u}(n)^\vee,$$

where

$$\Lambda_j := \mathfrak{i}_j - 1 - 2(\mathfrak{k}_j - 1),$$

and $(\mathfrak{i}_j, \mathfrak{k}_j)$ is the index of the lowest element in the j-th block of non-free variables, see (2.15). The Kostant-Kirillov form on $\mathrm{Fl}(\underline{n})$ is monotone (see [**63**, p653-654]). The moment polytope $\Delta_{\mathrm{Fl}(\underline{n})}$ has a center λ (see Figure 2.9) given by

$$\lambda_{i,k} = (i-1) - 2(k-1)$$

with Gelfand-Cetlin coordinates satisfying

$$\lambda_{i,k} - \lambda_{i,k-1} = 1, \quad \forall i,k.$$

By Cho-Kim [21, Theorem B], the center of the moment polytope corresponds to a monotone Lagrangian

$$L_\lambda := \Phi^{-1}(\lambda).$$

We remark that the monotonicity constant is $\frac{1}{2}$, that is,

$$\forall u : (\mathbb{D}^2, \partial\mathbb{D}^2) \to (X, L)\ \ \omega(u) = \tfrac{1}{2} I(u). \tag{2.23}$$

Indeed, for a disk of Maslov index two that has a single intersection with a divisor corresponding to a facet of the Gelfand-Cetlin polytope, the symplectic area is 1.

We consider a multiple cut $\mathcal{P}$ of the flag variety given by a collection of transversely intersecting single cuts along the hypersurfaces

$$\{\Phi_{i,k} - \Phi_{i+1,k+1} = \epsilon_{i,k,\diagup}\} \quad \text{resp.} \quad \{\Phi_{i,k} - \Phi_{i+1,k} = \epsilon_{i,k,\diagdown}\} \tag{2.24}$$

for each $(i, k, \diagup)$ resp. $(i, k, \diagdown) \in \text{Facets}(\text{Fl}(\underline{n}))$, where the parameters

$$(\epsilon_{i,k,\diagdown})_{i,k}, (\epsilon_{i,k,\diagup})_{i,k}$$

are positive and small enough that the cuts in (2.24) bound a polytope P_0

- that contains the Lagrangian L_λ,
- and the facets of P_0 are bijective to facets of the Gelfand-Cetlin polytope $\Delta_{\text{Fl}(\underline{n})}$, and that each cut is parallel to the corresponding facet of $\Delta_{\text{Fl}(\underline{n})}$.

The cuts are indeed well-defined because the hypersurfaces in (2.24) lie in F^{reg}, where the torus action is well-defined. Further the constants $(\epsilon_{i,k,\diagdown})_{i,k}, (\epsilon_{i,k,\diagup})_{i,k}$ are generic which ensures that P_0 does not have non-simplicial faces. Therefore, P_0 is a Delzant polytope and the cut space X_{P_0} is a toric manifold. For any $\theta \in \text{Facets}(\text{Fl}(\underline{n}))$, let

$$P_\theta \in \mathcal{P} \tag{2.25}$$

denote the top-dimensional polytope whose intersection with P_0 is the facet parallel to θ. We also denote $P_{0\theta} := P_0 \cap P_\theta \in \mathcal{P}$. For a generic tamed almost complex structure, the component X_{P_θ} is a fibration

$$\mathbb{P}^1 \to X_{P_\theta} \xrightarrow{\pi_\theta} X_{P_{0\theta}} \tag{2.26}$$

by J-holomorphic spheres, and π_θ is such that for any $x \in X_{P_{0\theta}}$, $\pi^{-1}(x)$ is the unique J-holomorphic sphere through x.

Unobstructedness is a consequence of the monotonicity of the Lagrangian: Let J^ν be a family of neck-stretching almost complex structures on the flag manifold $X = \text{Fl}(\underline{n})$. Let u_ν be a J^ν-holomorphic disk with a single output $x \in \text{crit}(F : L \to \mathbb{R})$. By monotonicity of (X, L_λ) the disk part of u_ν has Maslov index $I(u_\nu) \geq 2$. The index of u_ν (including the treed trajectory) is

$$\dim(L_\lambda) + 1 - (\dim(L_\lambda) - i(x)) + I(u_\nu) - \text{Aut}(\mathbb{D}^2) = 0.$$

Here $i(x)$ is the Morse index of the critical point x of a Morse function $F : L_\lambda \to \mathbb{R}$ on the Lagrangian. The only possibility then is $I(u_\nu) = 2$ and $i(x) = \dim(L)$, that is, x is the maximum point of F. Therefore

$$\mathfrak{m}^0_{CF(X,L,J_\nu)} = W x^\blacktriangledown, \quad \text{for some } W \in \Lambda_{>0}.$$

Unobstructedness now follows in a standard way with $b = W1^{\blacktriangledown}$ as the solution of the Maurer-Cartan equation (see for example the proof of Corollary 1.7 in Section 1.6). Note that since (X, L) is monotone, the set of disks contributing to $\mathfrak{m}^0$ is the same for all J^ν, and it also stays the same in the limit broken almost complex structure $\mathfrak{J}$ on $\mathfrak{X}_\mathcal{P}$.

We compute the potential by counting $\mathfrak{J}$-holomorphic broken disks in $\mathfrak{X}$. For any $\theta \in \mathrm{Facets}(\mathrm{Fl}(\underline{n}))$, there is a broken holomorphic disk u_θ of Maslov index 2, that intersects the divisor $\Phi^{-1}(\theta)$. Indeed,

- by Cho-Oh [20], there is a disk $u_{\theta,0}$ in the toric manifold X_{P_0} with $I(u_{\theta,0}) = 2$, $\omega(u_{\theta,0}) = 1 - \epsilon_\theta$, and that intersects the relative divisor $X_{P_{0\theta}}$ at a single point, say p;
- and there is a single holomorphic sphere $u_{\theta,1} : \mathbb{P}^1 \to X_{P_\theta}$ through p of area $1 - \epsilon_\theta$ (see (2.26));

and therefore $u_\theta = (u_{\theta,0}, u_{\theta,1})$ is a broken disk.

Next we show that there are no other broken disks whose gluing has Maslov index two by using the fact that the symplectic area of the broken disk is half the Maslov index of its gluing. Let u be any broken disk whose gluing has Maslov index 2. The disk component $u_0 : \mathbb{D}^2 \to X_{P_0}$ has Maslov index two, otherwise the area of the disk is at least $2-\epsilon_{\theta_1}-\epsilon_{\theta_2}$ for some $\theta_1, \theta_2 \in \mathrm{Facets}(\mathrm{Fl}(\underline{n}))$. Assuming $\epsilon_{\theta_1}, \epsilon_{\theta_2} \ll 1$, we get $\omega(u) \geq \omega(u_{\theta_0}) > 1$, which contradicts monotonicity (2.23). If u_0 intersects the relative divisor $X_{P_{0\theta}}$, then the sphere in X_{P_θ} has to be $u_{\theta,1}$ as above. Indeed if ϵ_θ is small enough, all other $\mathfrak{J}$-holomorphic spheres in X_{P_θ} have larger area, and then $\omega(u) > 2$ contradicting monotonicity.

The potential of $CF(\mathfrak{X}, L)$ is

$$W(y_{i,k})_{i,k} = \left(\sum_{k \leq i : (i+1,k) \in \mathrm{Free}(\mathrm{Fl}(\underline{n}))} \frac{y_{i+1,k}}{y_{i,k}} + \sum_{k \leq i : (i+1,k+1) \in \mathrm{Free}(\mathrm{Fl}(\underline{n}))} \frac{y_{i,k}}{y_{i+1,k+1}} \right. \\ \left. + \sum_{1 \leq j \leq r-1} y_{\mathfrak{i}_j,\mathfrak{k}_j} + \sum_{2 \leq j \leq r} \frac{1}{y_{\mathfrak{i}_j,\mathfrak{k}_j}} \right) q. \tag{2.27}$$

Here we recall that $\mathrm{Free}(\mathrm{Fl}(\underline{n}))$ is the set of variables in the Gelfand-Cetlin system of $\mathrm{Fl}(\underline{n})$ that are not fixed by the interlacing inequalities, and for any j, $(\mathfrak{i}_j, \mathfrak{k}_j)$ is the lowest element in the j-th block of non-free variables in the Gelfand-Cetlin system, see (2.15). This finishes the proof of Theorem 2.16. □

2.3. Counting curves in the plane

We relate our curve counts to those of Mikhalkin's tropical curves [61]. The correspondence is very natural and explains why the objects in this paper are called 'tropical', see Remark 2.23 at the end of this section. We start by reviewing some definitions from [61].

Definition 2.18. A *tropical curve* is a map

$$h : \Gamma \to \mathbb{R}^2$$

from a graph Γ (some of whose edges $e \in \text{Edge}(\Gamma)$, called *ends*, are incident on just one vertex $v \in \text{Vert}(\Gamma)$ instead of two) to $\mathbb{R}^2$ such that

- every non-leaf edge $e \in \text{Edge}(\Gamma)$ maps to a line segment with a rational slope $\mathcal{T}_e \in \mathbb{Z}^n$;
- every leaf edge $e \in \text{Edge}(\Gamma)$ maps to a semi-infinite line of rational slope $\mathcal{T}(e) \in \mathbb{Z}^n$, and for both leaf and non-leaf edges e the quantity $\mathcal{T}(e)$ is called the *slope of the edge*;
- and at any vertex $v \in \text{Vert}(\Gamma)$ a *balancing condition* is satisfied, namely that the sum of the slopes $\mathcal{T}(e)$ of the edges $e \ni v$ emanating out of v is equal to 0 :

$$\sum_{v \in e} \mathcal{T}_e = 0. \tag{2.28}$$

The *multiplicity* μ_e *of an edge* e is a positive integer such that $\frac{\mathcal{T}(e)}{\mu_e}$ is a primitive vector $w_e \in \mathbb{Z}^n$, which is called the *primitive slope*.[1]

Let $\Delta \subset \mathbb{R}^2$ be a simple polytope. A tropical curve $h : \Gamma \to \mathbb{R}^2$ is *generically adapted to* Δ if the slope $\mathcal{T}_e$ of any leaf $e \in \text{Edge}(\Gamma)$ is an outward normal of a facet of Δ. The discussion in Mikhalkin focusses on $\mathbb{P}^2$ with moment polytope $\Delta_{\mathbb{P}^2}$ whose outward normals are $\nu = (-1, 0), (0, -1), (1, 1)$. For a tropical curve generically adapted to $\Delta_{\mathbb{P}^2}$ the *degree* is defined as the number of ends (counted with multiplicity) that are normal to a fixed facet. Note that by the balancing condition the number is the same for any of the three facets. The *genus* of a tropical curve is the first Betti number $b_1(\Gamma)$ of the domain graph Γ.

Definition 2.19. A tropical curve $h : \Gamma \to \mathbb{R}^2$ is *simple* if Γ is trivalent, h is an immersion, and for any $y \in \mathbb{R}^2$, if $|h^{-1}(y)| > 1$ then y as exactly two pre-images, neither of which is a vertex of Γ.

Definition 2.20. A set of points $x_1, \ldots, x_k$ is *tropically generic* if any genus g tropical curve $h : \Gamma \to \mathbb{R}^2$ whose image $h(\Gamma)$ contains $x_1, \ldots, x_k$ and $|\,\text{Ends}(\Gamma)| \leq k + 1 - g$ satisfies the following:

(a) The curve h is simple,
(b) the images of vertices $h(\text{Vert}(\Gamma))$ are disjoint from $x_1, \ldots, x_k$,
(c) and $|\,\text{Ends}(\Gamma)| = k + 1 - g$.

Mikhalkin [**61**] shows that the number of curves in $\mathbb{P}^2$ of degree d and genus g passing through $3d - 1 + g$ tropically generic points can be computed by counting tropical curves with multiplicity.

Remark 2.21. Let $x_1, \ldots, x_{3d-1+g}$ be tropically generic points. A tropical curve adapted to $\mathbb{P}^2$ of degree d passing through those points must have the most number of ends, i.e. it must have $3d$ ends. Consequently each of the ends has primitive slope.

---

[1] In Mikhalkin's notation in [**61**], the primitive slope w_e is called 'slope'.

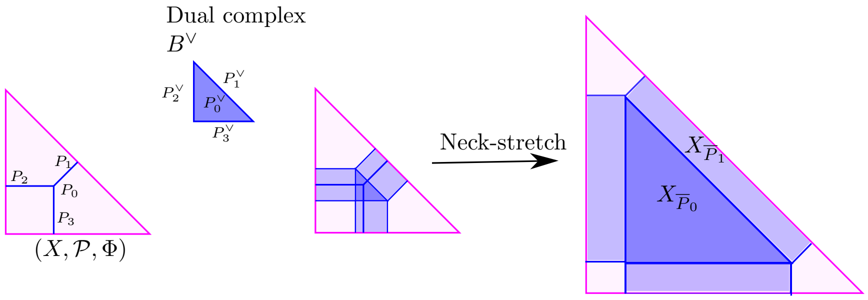

FIGURE 2.10. A multiple cut of $\mathbb{P}^2$

The following result says that in the genus zero case, there is a bijective correspondence between Mikhalkin graphs adapted to $\Delta_{\mathbb{P}^2}$ and broken maps in $\mathbb{P}^2$ with respect to the multiple cut shown in Figure 2.10. In the bijective correspondence described below, the Mikhalkin graph corresponding to a broken map u is the tropical graph of u with slight modifications. We view the Mikhalkin graph as lying in $\Delta_{\mathbb{P}^2} \subset \mathbb{R}^2$ and the tropical graph as lying in $P_0^\vee$. We chose the multiple cut to be the one in Figure 2.10 so that there is an isomorphism

$$\Delta_{\mathbb{P}^2} \simeq P_0^\vee. \tag{2.29}$$

We set up some notation needed to state the result. Let $\mathfrak{X}$ be the broken manifold obtained by applying the multiple cut in Figure 2.10 to $\mathbb{P}^2$. A 'point' in a broken manifold $\mathfrak{X}$ is given by a polytope $P \in \mathcal{P}$, a tropical position $\mathcal{T}_x \in P^\vee$ and the point itself, which is $x \in X_{\overline{P}}$. We say that a broken map $u : (C, z) \to \mathfrak{X}$ satisfies $u(z) = x$ if the component $C_v \subset C$ containing the marking z satisfies

$$P(v) = P_x, \quad \mathcal{T}(v) = \mathcal{T}_x.$$

Note that the tropical position $\mathcal{T}_x$ of a point x in $\mathfrak{X}$ corresponds to a point in $\Delta_{\mathbb{P}^2}$ by (2.29).

PROPOSITION 2.22. (Tropical curves as broken maps) *Consider a set of points $x_1, \dots, x_{3d-1} \in \mathfrak{X}$ lying in the piece $X_{\overline{P}_0}$. There is a bijective correspondence between the set of genus zero broken maps passing through the points $\{x_i\}_i$ and the set of genus zero Mikhalkin graphs in $\mathbb{R}^2$ passing through the points $\{\mathcal{T}_{x_i}\}_i$.*

PROOF. Given a Mikhalkin graph $h : \Gamma \to \mathbb{R}^2$ generically adpated to $\Delta_{\mathbb{P}^2}$, we first construct a tropical graph, which is an augmentation of Γ denoted by $\Gamma_{\text{aug}} \supset \Gamma$. For all vertices $v \in \text{Vert}(\Gamma)$, we assign $P(v) := P_0$. For every leaf edge $e \in \text{Edge}(\Gamma)$, the augmented graph Γ_{aug} contains an extra vertex v_e on which e is incident. If e in Γ intersects the facet $F_i \subset \Delta_{\mathbb{P}^2}$, then the polytope $P(v_e)$ is P_i. The set of vertices thus added is denoted by

$$\text{Vert}_1(\Gamma_{\text{aug}}).$$

Lastly, if an edge $e \in \text{Edge}(\Gamma)$ of the Mikhalkin graph passes through a point constraint $x \in \mathbb{R}^2$ then in Γ_{aug} we subdivide e into two edges e_1 and e_2 by inserting a new vertex v_x with $P(v_x) := P_0$. The new vertex has a marking, and thus has a valence of 3. The set of vertices $v \in \mathrm{Vert}(\Gamma_{\mathrm{aug}})$ with markings is denoted by

$$\mathrm{Vert}_{\to}(\Gamma_{\mathrm{aug}}).$$

Thus,

$$\mathrm{Vert}(\Gamma_{\mathrm{aug}}) = \mathrm{Vert}(\Gamma) \cup \mathrm{Vert}_{\to}(\Gamma_{\mathrm{aug}}) \cup \mathrm{Vert}_1(\Gamma_{\mathrm{aug}}).$$

The slopes of edges in Γ_{aug} are the same as their slopes in Γ. The tropical positions for the vertices in Γ_{aug} are given by the map $h : \Gamma \to \mathbb{R}^2$ on the Mikhalkin graph. See Figure 2.11 for an example.

We describe a way of orienting the edges in Γ_{aug} which is useful in the rest of the proof, called the *marking orientation*, that satisfies the following condition: If $v \in e$ and e points towards v then in the tropical graph $\Gamma_{\mathrm{aug}} \backslash \{e\}$ the vertex v is free to move without altering the slopes of the other edges. A one-sided edge e (corresponding to a marking) that is incident on a vertex v is oriented to point towards v. It is easy to verify that marking orientations can be assigned to all the edges of Γ_{aug}, and that every vertex $v \in \mathrm{Vert}(\Gamma)$ has exactly two incoming edges, see Figure 2.11 for an example.

Next we describe the map at each vertex. Choose an ordering of the vertices $v \in \mathrm{Vert}(\Gamma_{\mathrm{aug}})$ that respects the marking orientation described above, and define the maps $(u_v)_v$ in that order. For a vertex $v \in \mathrm{Vert}(\Gamma)$ whose incident edges have slopes μ_1, μ_2, μ_3 the map is

$$u_v : \mathbb{P}^1 \backslash \{0, 1, 2\} \to X^{\square}_{\overline{P}_0} \simeq (\mathbb{C}^\times)^2, z \mapsto cz^{\mu_1}(z-1)^{\mu_2}(z-2)^{\mu_3},$$

where the domain is parametrized so that $0, 1, 2 \in \mathbb{P}^1$ are lifts of nodal points. The constant $c \in (\mathbb{C}^\times)^2$ is chosen so that the map satisfies the matching constraint at the nodal points corresponding to the two incoming edges. For a vertex $v \in \mathrm{Vert}_{\to}(\Gamma)$ with a marking, and whose incident edges have slope $\mu, -\mu$ the map is a trivial cylinder of slope μ

$$u_v : \mathbb{P}^1 \backslash \{0, \infty\} \to X^{\square}_{\overline{P}_0} \simeq (\mathbb{C}^\times)^2, z \mapsto cz^{\mu},$$

and the constant c is determined by the tropical point constraint at the marking. To define the map u_v for a vertex $v \in \mathrm{Vert}_1(\Gamma_{\mathrm{aug}})$, we first observe that the manifold $X^{\square}_{\overline{P}(v)}$ is a $\mathbb{P}^1$-fibration

$$\mathbb{P}^1 \to X^{\square}_{\overline{P}(v)} \xrightarrow{\pi_{P(v)}} Y \tag{2.30}$$

where the manifold Y is a subset of a torus-invariant divisor of X. The map u_{v_e} corresponding to the vertex v_e is an injective map to the fiber of (2.30).

Conversely, consider a broken map passing through the tropical point constraints. To prove that the broken map arises from a Mikhalkin graph, it is enough to show that

(a) there are no components mapping to X_P with $\dim(P) = 2$,

(b) if a component u_v maps to X_P with $\dim(P) = 1$, then $\Pi_{P(v)} \circ u_v$ is constant.

Both claims follow from dimension reasons and the fact that markings lie on components mapping to $X_{\overline{P}_0}$.

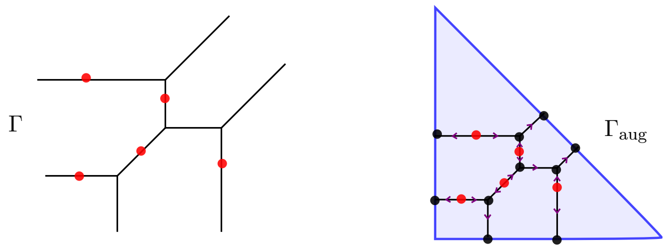

FIGURE 2.11. A Mikhalkin graph Γ of degree 2 in $\mathbb{R}^2$ through 5 generic points (marked in red) and the corresponding tropical graph Γ_{aug} of a broken map in the dual complex $B^\vee$ from Figure 2.10. In Γ_{aug} the red vertices contain marked points mapping to the point constraints, and the arrows on edges indicate the marking orientation.

Finally we show that the multiplicity in Mikhalkin [**61**] corresponds to the size of the tropical symmetry group of the broken map. We recall from [**61**, p11-12] that Mikhalkin's multiplicity is the product of multiplicities $\text{mult}(v)$ of all the vertices v in the tropical curve, and the multiplicity $\text{mult}(v)$ at any trivalent vertex v is the area of the parallelogram spanned by the slopes $\mathcal{T}(e_1)$, $\mathcal{T}(e_2)$ of any of two edges e_1, e_2 incident on v. A tropical symmetry of a broken map with tropical graph Γ_{aug} is a tuple

$$(\underline{g}, \underline{z}) = ((g_v)_{v \in \text{Vert}(\Gamma_{\text{aug}})}, (z_e)_{e \in \text{Edge}(\Gamma_{\text{aug}})}), \quad g_v \in T_{P(v),\mathbb{C}}, \quad z_e \in \mathbb{C}^\times$$

satisfying

$$g_{v_+} g_{v_-}^{-1} = z_e^{\mathcal{T}(e)} \quad \forall e = (v_+, v_-) \in \text{Edge}(\Gamma_{\text{aug}}), \tag{2.31}$$

see Definition 4.33. We count the number of solutions of (2.31). For $v \in \text{Vert}_\to(\Gamma_{\text{aug}})$, $g_v = \text{Id}$ in order to ensure that the evaluation of the marking satisfies the point constraint. We solve for g_v, one at a time, with vertex ordering respecting the edge orientation. Consider a vertex $v \in \text{Vert}(\Gamma)$. Recall that by the genericity of $\{x_i\}_i$, v is trivalent with neighboring vertices v_1, v_2, v_3. Two of the edges, say $e_1 = (v_1, v)$ and $e_2 = (v_2, v)$, are directed towards v. Assuming the values of g_{v_1}, g_{v_2} to be given we solve the equation

$$g_v g_{v_1}^{-1} = z_{e_1}^{\mathcal{T}_{e_1}}, \quad g_v g_{v_2}^{-1} = z_{e_2}^{\mathcal{T}_{e_2}}$$

for $g_v \in T_{P_0,\mathbb{C}} \simeq (\mathbb{C}^\times)^2$, z_{e_1}, $z_{e_2} \in \mathbb{C}^\times$. Use an integral basis of $\mathfrak{t}_{P_0,\mathbb{Z}}$ so that $\mathcal{T}_{e_1} = (p_1, 0)$, $\mathcal{T}_{e_1} = (q, p_2)$ for some integers p_1, p_2, q. In this notation the multiplicity $\text{mult}(v)$ is $|p_1 p_2|$. We may write $g_v = (g_v^0, g_v^1)$. We solve in the order g_v^1, z_{e_2}, g_v^0, z_{e_1}. For a fixed value of g_{v_1}, g_{v_2}, there are $p_1 p_2$ solutions: There are p_1 solutions for z_{e_2}, and for each of these solutions, there are p_2 solutions of z_{e_1}. Thus we have shown that for a vertex $v \in \text{Vert}(\Gamma) \subset \text{Vert}(\Gamma_{\text{aug}})$, the number of solutions for (g_v, z_{e_1}, z_{e_2}) is equal to the multiplicity of v. Finally, given a solution of (2.31) for vertices $v \in \text{Vert}(\Gamma)$ and connecting edges, there is a unique solution g_v for any vertex $v \in \text{Vert}_1(\Gamma_{\text{aug}})$, since the edge incident on v has primitive slope. Consequently, we have shown that

$$|T_{\text{trop}}(\Gamma_{\text{aug}})| = \prod_{v \in \text{Vert}(\Gamma)} \text{mult}(v).$$

□

REMARK 2.23. (Origin of the word 'tropical') The connection between broken maps and Mikhalkin graphs explains why the objects of study in this paper are called 'tropical', as we elaborate in this remark. The *tropical semi-ring* $\mathbb{R}_{trop}$ is the set $\mathbb{R}$ with the addition and multiplication operations defined as max and $+$ respectively. For example, a tropical quadratic polynomial in $\mathbb{R}_{trop}$ in variables x and y has the form

$$f(x,y) = \max\{a_{00}, a_{10}+x, a_{01}+y, a_{11}+x+y, a_{20}+2x, a_{02}+2y\}. \tag{2.32}$$

A *tropical hypersurface* is the zero set of a tropical polynomial in $\mathbb{R}^n$, and is defined to be the set of points where the function is not linear. Thus tropical hypersurfaces are complexes of polytopes. Mikhalkin's tropical curve in $\mathbb{R}^2$ is the parametrized

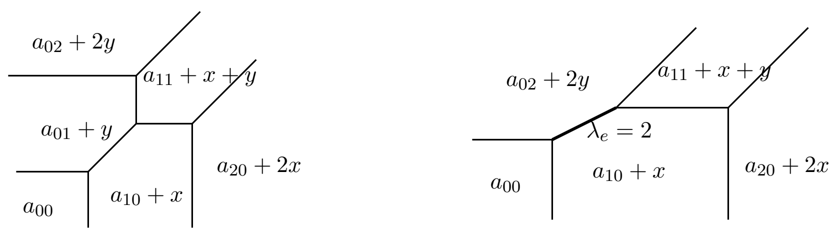

FIGURE 2.12. Tropical curves of degree 2 that arise as zero sets of f in (2.32), and the maximum monomial in regions of $\mathbb{R}^2$.

version of a tropical hypersurface in $\mathbb{R}^2$.

2.4. Relation to toric degeneration

We relate our results to those for toric degenerations as in Gross and Siebert ([**45**], [**46**]), and in particular explain what kinds of toric degenerations are *not* covered by our results. Recall that a *toric degeneration* is a flat family

$$\pi : \mathcal{X} \to \mathbb{C}$$

whose fiber

$$X_t := \pi^{-1}(t)$$

is regular and diffeomorphic to an irreducible projective variety X for $t \in \mathbb{C}\backslash\{0\}$, and whose central fiber X_0 is a union of toric varieties glued pair-wise along torus-invariant divisors. The central fiber X_0 is a tropical manifold, which means the following: The space X_0 is naturally represented by a polyhedral complex $\mathcal{P}$ consisting of Delzant polytopes glued along facets, and there is a tropical moment map $\Phi : X_0 \to B$ to the underlying topological space $B = \cup_{P \in \mathcal{P}} P^\circ$, given by the union of interiors P° of all the polytopes P. The space B is a *singular integral affine manifold*: An *integral affine structure* on B (see for example Gross [**43**]) is a torsion-free flat connection on the tangent bundle TB with monodromy contained in $SL(n, \mathbb{Z})$, and it arises from the Lagrangian torus fibration $\Phi : X_0 \to B$. The singular set of B, known as the *discriminant locus* $\Delta \subset B$, is a cell complex of codimension 2 contained in $B \backslash (\mathcal{P}^{\dim(T)} \cup \mathcal{P}^0)$. Here $\mathcal{P}^n \subset \mathcal{P}$ is the subset consisting of polytopes $P \subset \mathfrak{t}^\vee$ of codimension n. The integral structure of B has a non-trivial monodromy in $SL(n, \mathbb{Z})$ around points in the discriminant locus. In the complement of the discriminant locus, the tropical manifold X_0 looks exactly like a union of cut spaces obtained from a multiple cut. Because of the discriminant locus, the gluing of the toric varieties in the tropical manifold along torus-invariant divisors is not toric.

In the four-dimensional case, *almost toric manifolds* are the symplectic geometric analogues of tropical manifolds. Almost toric manifolds, introduced by Symington [**86**] are generalizations of toric manifolds, where the completely integrable system is allowed to have isolated focus-focus singularities. (See [**28**] for a nice exposition.) An almost toric manifold is equipped with a *base diagram* which is the image of the Hamiltonians, with the additional data of the locations of the singularities on the base diagram, and the directions of the eigenvector of the monodromy. For example the singular point in Figure 2.13, has a monodromy $\begin{pmatrix} 1 & 0 \\ d & 1 \end{pmatrix}$, and the eigen-vector of the monodromy is indicated by a dotted line. There is a toric fibration that projects to the complement of the dotted line, and which is glued along the dotted line by the descent of the map $\begin{pmatrix} 1 & 0 \\ d & 1 \end{pmatrix} : \mathbb{R}^2 \to \mathbb{R}^2$ to the torus $\mathbb{R}^2/\mathbb{Z}^2 \to \mathbb{R}^2/\mathbb{Z}^2$. The singular torus fibration is equivalently represented by the two base diagrams in Figure 2.13. Transitioning from one to the other is called a *mutation*. The fiber above the singular point is a pinched torus.

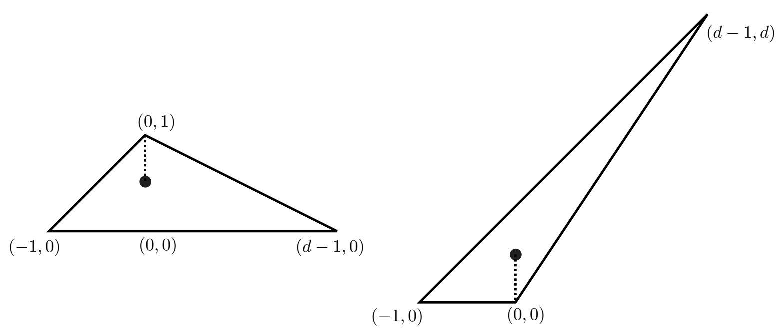

FIGURE 2.13. Two isomorphic almost toric manifolds with a focus-focus singularity.

Gross-Siebert's toric degenerations may be viewed as multiple cuts where the cuts are allowed to pass through singular Lagrangian fibers in an almost toric manifold, such as a cut along $\{x = 0\}$ in the Figure 2.13. (This is the same as Figure 20 in [**44**, p26].) The resulting cut space is a singular manifold, since the cut operation requires us to quotient by S^1, and the cut locus contains a fixed point of the S^1-action. Our techniques do not apply to such singular spaces.

In some cases we may apply our results to these kinds of singular toric degenerations if we replace a single cut by two parallel non-singular cuts *straddling* the singular point as shown in Figure 2.14 (a). In this example, the straddling cut produces three cut spaces, where the middle piece is almost toric. The 'non-toric gluing' in toric degenerations is thus replaced by an extra cut space that is non-toric. Multiple singular cuts may also be replaced by straddling cuts, see Figure 2.16 for an example. The middle piece in Figure 2.14 (a) contains the resolution of an A_d singularity. Indeed, the singularity in the almost toric base diagram may be pushed to the vertex by a 'nodal slide' (Figure 2.14 (b)). The nodal slide operation moves the singular value in the base diagram in the direction of the eigenvector of the monodromy. This operation does not affect the total space (as a symplectic manifold), but changes the toric fibration, see [**86**, Theorem 6.5]. In the limit when the singular value in the base diagram coincides with a vertex v, the resulting fibration has a non-toric singularity at the vertex. The fiber over the vertex v is a path of $d-1$ Lagrangian spheres.[2] Via a deformation of the symplectic form, the Lagrangian spheres may be deformed to symplectic spheres with self-intersection -2, see Figure 2.14 (c), which is a resolution of an A_d-singularity. A_2-singularities occur in toric degenerations of cubic surfaces, which we analyze in Section 2.1. Similar multiple cuts to the one in Section 2.1 can presumably be used to analyze spheres and disks in resolutions of A_d-singularities for $d \geq 2$, which may give alternate proofs of results in Chan-Lau [**16**]. The case $d = 1$ is particularly easy since a single cut suffices to count disks, see Remark 2.8.

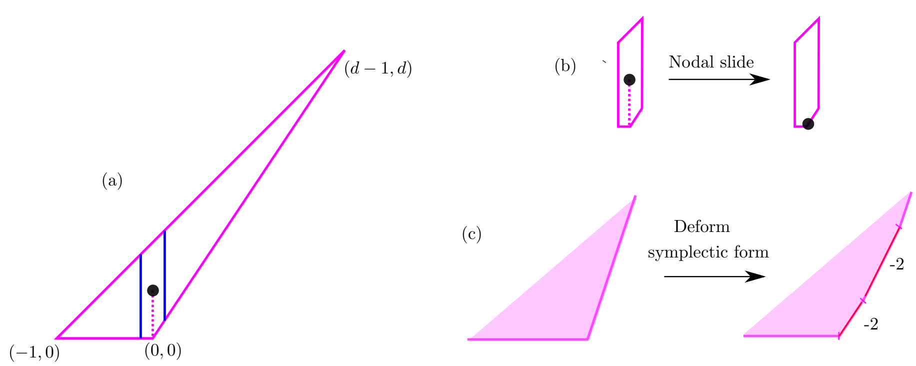

FIGURE 2.14. Straddling cuts

EXAMPLE 2.24. We describe a toric degeneration of a cubic surface in $\mathbb{P}^2$ to 3 hyperplanes from [**46**, Section 2.1]. Let

$$f \in \mathbb{C}[x_0, x_1, x_2, x_3]$$

[2]For $t \in (0, 1]$ the hypersurface $\mathcal{P}_t := \{z_1 z_2 + tP(z_3) = 0\} \subset \mathbb{C}^3$ (where P is a polynomial of degree d) has an almost toric structure with d focus-focus singularities ([**28**, Section 7.3]). There is a path of $d-1$ Lagrangian spheres, each passing through a pair of the singular points ([**28**, Remark 7.7]). The path of spheres gets collapsed to an A_{d-1}-singularity in the limit $t \to 0$.

be a generic homogeneous polynomial of degree 3 in the variables $x_0, \ldots, x_3$. Consider the toric generation given by the family

$$Y := \{tf + x_0x_1x_2 = 0 : t \in \mathbb{C}\} \overset{\pi}{\mapsto} t \in \mathbb{C},$$

The family Y is not smooth at the intersections of $V(f) := \{f = 0\}$ with the lines

$$V(x_i, x_j) := \{x_i = x_j = 0\}, \quad i, j \in \{0, 1, 2\}.$$

Thus each intersection $V(x_i, x_j) \cap Y$ has 3 singular points. The fibers of π are smooth except for $X_0 := \pi^{-1}(0)$, which is the union of 3 hyperplanes

$$X_0 = V(x_0) \cup V(x_1) \cup V(x_2), \quad V(x_i) = \{x_i = 0\}.$$

The fiber $\pi^{-1}(0)$ is a tropical manifold where any pair $V(x_i)$, $V(x_j)$ is glued along a line $V(x_i, x_j)$ containing three points in the discriminant locus. The integral affine manifold B underlying X_0 is shown in Figure 2.15. The monodromy of the integral

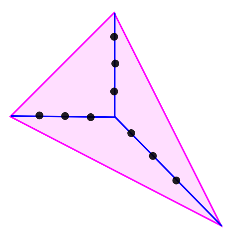

FIGURE 2.15. The integral affine manifold underlying X_0 in Example 2.24.

affine structure at each of the singular points is $\begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}$. Gluing the central fiber X_0 produces an almost toric manifold X_0^{glue} with base diagram as in Figure 2.16, and a focus-focus singularity corresponding to each point in the discriminant locus. In our analysis of disks and spheres on the cubic surface in Chapter 2, we modify the toric fibration on X_0^{glue} by moving the singular values of the moment map to the vertices of the moment polytope by a nodal slide. In the new fibration, denoted by $\Phi : X_0 \to \mathbb{R}^2$, Φ is an $(S^1)^2$-moment map in the complement of the fibers of the 3 vertices. The inverse image of each vertex $\mu \in P$ is a union $\Phi^{-1}(\mu) \cong S^2 \cup S^2$ of two Lagrangian spheres, that is, a resolution of an A_2 singularity. One may also use straddling cuts in this example, see Figure 2.16. However, we do not use this multiple cut in our analysis in Section 2.1.

EXAMPLE 2.25. A quartic surface in $\mathbb{P}^3$ can be degenerated in a similar way to the cubic surface, see Gross [**46**, Section 2.1]. The degenerated variety consists of four copies of $\mathbb{P}^2$ glued pairwise along toric divisors, and there are 4 focus-focus singularities along each of the glued divisors. See Figure 2.17. Via a straddling cut we obtain a multiply cut manifold whose (symplectic) sum is symplectomorphic to the quartic. Our set-up assumes that the polyhedral decomposition corresponding to the tropical manifold is embedded in a vector space, while in the quartic example

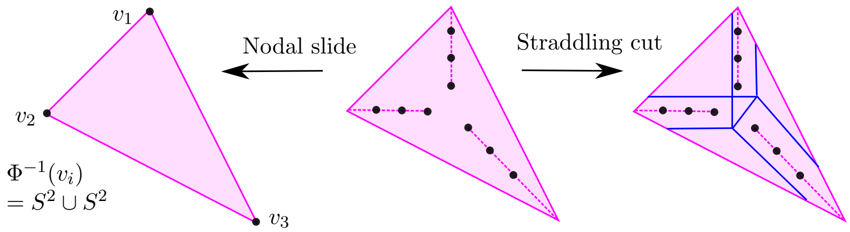

FIGURE 2.16. Center : Base diagram for an almost toric fibration of the cubic surface.

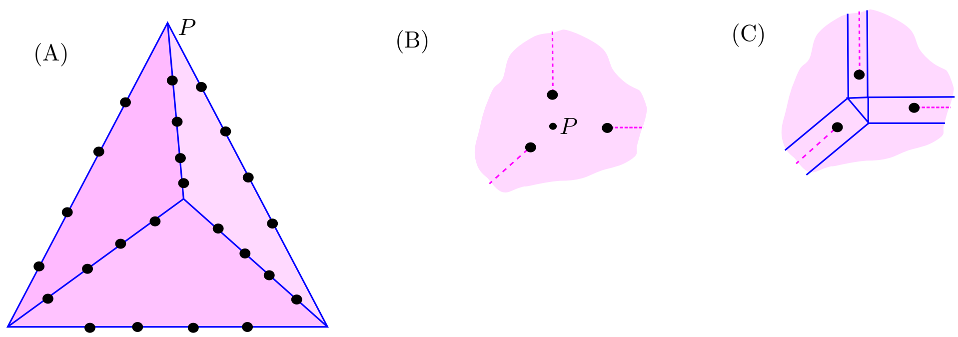

FIGURE 2.17. (a) Integral affine manifold corresponding to the toric degeneration X_0 of a quartic surface. There are 24 singular points on the cut locus. (b) Base diagram for the almost toric fibration on the gluing X_0^{glue} in a neighborhood of the vertex P. (c) A part of the straddling cut near P.

the topological manifold $B = \cup_{P \in \mathcal{P}} P^\circ$ underlying the polyhedral decomposition is a two-sphere. We expect our result to hold in the more general set-up where B is covered by affine charts, which would imply that the Lagrangian fibers in the pieces of the multiply cut quartic are unobstructed. It will be interesting to compute the potential of such a Lagrangian using broken disks. This ends the Example.

CHAPTER 3

Broken manifolds

In this chapter we review the multiple cut construction and the associated degenerations of almost complex structures. Our approach is much less general than, for example, Parker [**68**], but we wish to be completely explicit.

In the first half of the chapter we describe cut spaces and broken manifolds as symplectic manifolds. The multiple cut is a generalization of the symplectic cut operation of Lerman [**54**] and Gompf [**40**], which we refer to as the 'single cut'. A multiple cut is defined on a symplectic manifold equipped with a tropical moment map and a polyhedral decomposition $\mathcal{P}$, and it yields a symplectic cut space X_P corresponding to every polytope P in the decomposition $\mathcal{P}$. The broken manifold consists of top-dimensional cut spaces and thickenings of the lower-dimensional cut spaces, denoted by $X_{\overline{P}}$ (see Section 3.3). The definition of the thickenings $X_{\overline{P}}$ requires the additional datum of a dual complex.

In the second half of the chapter (starting from Section 3.4), we describe broken manifolds and cut spaces as almost complex manifolds with cylindrical almost complex structures. The manifold X is equipped with a family of neck-stretched almost complex structures J^ν. In the infinite neck length limit $\nu \to \infty$, the almost complex manifolds $X^\nu := (X, J^\nu)$ degenerate into broken almost complex manifolds, denoted by $X^\square_{\overline{P}}$. The broken almost complex manifold $X^\square_{\overline{P}}$ is diffeomorphic to a complement of relative divisors of the symplectic broken manifold $X_{\overline{P}}$, but there is no canonical embedding where the symplectic form tames the almost complex structure.

The definition of the broken manifold $X^\square_{\overline{P}}$ as the degenerate limit gives a natural family of identifications between subsets of the neck-stretched manifolds (X, J_ν) and the broken manifold $\sqcup_P X^\square_{\overline{P}}$. These identifications, called *translations*, are defined in Section 3.6, and are analogous to 'target rescalings' of Ionel [**50**].

Finally in Section 3.7, we prove the existence of a cylindrical structure on the symplectic form in the neighborhood of cut loci. The cylindrical structure underlying neck-stretched almost complex structure is chosen to be the same as the cylindrical structure on the symplectic form. Choosing the cylindrical structures in this manner will allow us (later in Chapter 7) to construct families of diffeomorphisms $(X, J^\nu) \to (X, \omega)$ for which ω tames J^ν.

3.1. Symplectic cut

A multiple cut is the generalization of the *symplectic cut* construction of Lerman [**54**] and Gompf [**40**], which we now review. We call this construction a *single* *cut* , to contrast it with a multiple cut. The construction of symplectic cuts uses Hamiltonian circle actions on symplectic manifolds.

DEFINITION 3.1. (Lerman's symplectic cut construction)

(a) (Hamiltonian circle actions) Let (X, ω) be a compact symplectic manifold. Let

$$S^1 = \{z \in \mathbb{C} \mid |z| = 1\}$$

denote the circle group; we identify its Lie algebra $\mathrm{Lie}(S^1)$ with $\mathbb{R}$ by division by i. A *Hamiltonian action* of the circle group S^1 on X is an action $S^1 \times X \to X$ generated by the Hamiltonian flow of a *moment map*

$$\Phi : X \to \mathbb{R}, \quad \omega(\xi_X, \cdot) = -\mathrm{d}\Phi$$

where the generating vector field of an element $\xi \in \mathbb{R}$

$$\xi_X \in \mathrm{Vect}(X), \quad \xi_X(x) = \frac{d}{dt}|_{t=0} \exp(it\xi)x. \tag{3.1}$$

In particular the affine line $\mathbb{C}$ has symplectic form

$$\omega_{\mathbb{C}} = \frac{-i}{2} \mathrm{d}z \wedge \mathrm{d}\overline{z} \in \Omega^2(\mathbb{C}).$$

The Hamiltonian action of S^1 is given by scalar multiplication and has moment map

$$\Phi_{\mathbb{C}} : \mathbb{C} \to \mathbb{R}, z \mapsto \frac{-|z|^2}{2}.$$

(b) (Global symplectic cuts) Let X be a symplectic manifold with symplectic form ω_X and a free Hamiltonian S^1-action with moment map Φ. The product $\hat{X} = X \times \mathbb{C}$ has product symplectic form $\hat{\omega} = \pi_1^* \omega_X + \pi_2^* \omega_{\mathbb{C}}$. The diagonal action of S^1 has moment map

$$\hat{\Phi} : \hat{X} \to \mathbb{R}, \quad (x, z) \mapsto \Phi(x) - \frac{|z|^2}{2}.$$

The zero level set is the union

$$\hat{\Phi}^{-1}(0) = (\Phi^{-1}(0) \times \{0\}) \sqcup \left\{ (x, z) \,\middle|\, \Phi(x) = \frac{|z|^2}{2} > 0 \right\}$$

of two pieces where both Φ and z are zero and the piece where both Φ and z are non-zero. The action on $z \neq 0$ has a natural slice given by $z \in \mathbb{R}_{>0}$ so that

$$\{(x, z) | \Phi(x) = |z|^2/2 > 0\} \cong \Phi^{-1}(\mathbb{R}_{>0}).$$

The symplectic quotient $\hat{\Phi}^{-1}(0)/S^1$ is called the *symplectic cut space*, and is alternately viewed as

$$X_+ := \hat{\Phi}^{-1}(0)/S^1 \simeq \{\Phi^{-1} \geq 0\}/ \sim . \tag{3.2}$$

where $\sim$ mods out the boundary $\Phi^{-1}(0)$ by the S^1-action. Thus the cut space is the union of $\{\Phi > 0\}$ and the symplectic quotient $\Phi^{-1}(0)/S^1$. One has a similar construction of a cut space

$$X_- := \{\Phi \leq 0\}/ \sim,$$

which is the union of $\Phi^{-1}(\mathbb{R}_{>0})$ and the symplectic quotient $\Phi^{-1}(0)/S^1$. The symplectic manifolds X_-, X_+ both contain a copy of $\Phi^{-1}(0)/S^1$ via the embeddings

$$i_- : \Phi^{-1}(0)/S^1 \to X_-, \quad i_+ : \Phi^{-1}(0)/S^1 \to X_+$$

with opposite normal bundles $N_\pm \to X_\pm$, so that $N_-^{-1} \cong N_+$.

(c) (Local symplectic cuts) Given an open subset $U \subset X$ with a free Hamiltonian S^1-action with moment map $\Phi : U \to \mathbb{R}$, such that $X \backslash U$ is disconnected, gluing together the cut $U_+ \cup U_-$ with $X - \Phi^{-1}(0)$ produces cut spaces X_+, X_-.

3.2. Multiple cuts in a symplectic manifold

We recall that the input datum for a single cut consists of a hypersurface and a Hamiltonian S^1-action in the neighborhood of the hypersurface. The input datum for a multiple cut consists of a collection of intersecting hypersurfaces with Hamiltonian S^1-actions in their neighborhoods. Neighborhoods of intersections of hypersurfaces have Hamiltonian torus actions, whose restrictions coincide with the S^1-actions corresponding to individual hypersurfaces. The various Hamiltonian actions are recorded by a *tropical moment map* with target space $\mathfrak{t}^\vee$, and a *polyhedral decomposition* of $\mathfrak{t}^\vee$. These polyhedral decompositions appeared in, for example, Meinrenken [57].

Definition 3.2. (Simple resp. Delzant polytopes) Let T be a torus with Lie algebra $\mathfrak{t}$. Let $\mathfrak{t}_\mathbb{Z} \subset \mathfrak{t}$ denote the coweight lattice of points that map to the identity under the exponential map, so that $T \cong \mathfrak{t}/\mathfrak{t}_\mathbb{Z}$. A convex polytope P in $\mathfrak{t}^\vee$ is described by a collection of linear inequalities determined by constants $c_F \in \mathbb{R}$ and normal vectors $\nu_F \in \mathfrak{t}$:

$$P = \{\lambda \in \mathfrak{t}^\vee \mid \langle \lambda, \nu_F \rangle \geq c_F, \quad \forall F \subset P \text{ facets}\}.$$

We allow polytopes to be unbounded. The polytope P is a *simple* if for each vertex point $\lambda \in P$, the normal primitive vectors $\nu_F \in \mathfrak{t}_\mathbb{Z}$ to the facets $F \subset P$ containing λ form a basis of $\mathfrak{t}$. For a simple polytope P, a face $\lambda \subset P$ which is the intersection of facets $F_1, \ldots, F_m$ is *smooth* if

$$\operatorname{span}(\nu_{F_i}, i = 1, \ldots, m) \cap \mathfrak{t}_\mathbb{Z} = \operatorname{span}_\mathbb{Z}(\nu_{F_i}, i = 1, \ldots, m). \tag{3.3}$$

A simple polytope is *Delzant* if each of its vertices is a smooth face (or equivalently all its faces are smooth).

By Delzant [25], there is a bijection between Delzant polytopes and compact symplectic toric manifolds. A Delzant polytope P corresponds to a symplectic manifold V_P with an effective Hamiltonian action of a torus $T \simeq (S^1)^{\frac{1}{2}(\dim(V_P)}$ and moment map and polytope

$$\Psi : V_P \to \mathfrak{t}^\vee, \quad \Psi(V_P) = P.$$

We remark that if in a simple polytope P collection of facets $F_1, \ldots, F_m$ are not smooth then, $\mathfrak{t}_\mathbb{Z} \subset \operatorname{span}_\mathbb{Q}(\nu_{F_i})_i$. In this case V_P is an *orbifold*, that is, it is covered by charts that are finite quotients of $\mathbb{R}^n$.

Definition 3.3. (Tropical Hamiltonian action) A *tropical Hamiltonian action* of a torus T with Lie algebra $\mathfrak{t}$ is a triple $(X, \Phi, \mathcal{P})$ consisting of a

(a) compact symplectic manifold X;

(b) (Polyhedral decomposition) a decomposition

$$\mathfrak{t}^\vee = \bigcup_{P \in \mathcal{P}} \operatorname{int}(P), \quad \mathcal{P} = \{P \subset \mathfrak{t}^\vee\}$$

of $\mathfrak{t}^\vee$ into the disjoint union of the interiors of simple polytopes $P \in \mathcal{P}$ such that

- if $P_0, P_1 \in \mathcal{P}$ have non-empty intersection, then $P_0 \cap P_1 \in \mathcal{P}$ and $P_0 \cap P_1$ is a face of both P_0 and P_1,

and $\mathcal{P}^0 \subset \mathcal{P}$ is the subset of top-dimensional polytopes in $\mathfrak{t}^\vee$;

(c) a *tropical moment map* compatible with $\mathcal{P}$

$$\Phi : X \to \mathfrak{t}^\vee$$

in the following sense. For any $P \in \mathcal{P}$, we denote by

$$\mathfrak{t}_P := \operatorname{ann}(TP) \subset \mathfrak{t}$$

the annihilator of the tangent space of P at any point $p \in P$, and by

$$T_P = \exp(\mathfrak{t}_P)$$

the torus whose Lie algebra is $\mathfrak{t}_P$. Let $\mathfrak{t}_{P,\mathbb{Z}}$ be the coweight lattice in $\mathfrak{t}_P$ so that

$$T_P \cong \mathfrak{t}_P / \mathfrak{t}_{P,\mathbb{Z}}.$$

For any $P \in \mathcal{P}$, there exists an open neighbourhood U_P of $\Phi^{-1}(P)$ such that the composition

$$\pi_P \circ \Phi : U_P \to \mathfrak{t}_P^\vee$$

is a moment map for a free action of T_P on U_P.

Remark 3.4. (Tropical manifold for a single cut) In the single breaking case, the tropical manifold datum consists of a map $\Phi : X \to \mathbb{R}$ and a decomposition $\mathbb{R} := (-\infty, c] \cup [c, \infty)$, such that Φ generates a free S^1-action in the neighborhood of $\Phi^{-1}(c)$. Thus the set of polytopes is

$$\mathcal{P} = \{(-\infty, c], \{c\}, [c, \infty)\}.$$

In the case of a single cut, we denote the tropical manifold by the triple (X, Φ, c).

Definition 3.5. (Cut space for a multiple cut) Given a tropical symplectic manifold $(X, \Phi, \mathcal{P})$ for every polytope $P \in \mathcal{P}$ the *cut space* is a symplectic manifold (or orbifold, if P is not Delzant)

$$X_P := \Phi^{-1}(P)/\sim, \tag{3.4}$$

where the equivalence $\sim$ mods out by the following torus actions:

$$x \sim tx, \quad \forall x \in \Phi^{-1}(Q^\circ), t \in T_Q \tag{3.5}$$

for all polytopes $Q \subseteq P$ contained in $\mathcal{P}$. Thus the open set $\Phi^{-1}(P^\circ) \subset \Phi^{-1}(P)$ is modded out by the T_P-action and the quotient

$$\Phi^{-1}(P^\circ)/T_P$$

is an open subset of X_P. The space X_P is a manifold (or orbifold if P is not Delzant) by an iterative application of Lerman's cut. Cut spaces have natural inclusions

$$X_Q \subset X_P, \quad Q \subset P.$$

For a face $Q \subset P$ with $\operatorname{codim}_P(Q) = 1$, the corresponding subset X_Q is called a *relative divisor of* X_P. The relative divisors X_Q intersect ω-orthogonally, and the intersections correspond to a cut space X_R for some polytope $R \in \mathcal{P}$. This ends the Definition.

Example 3.6. For a tropical manifold (X, Φ, c) with a single cut (using notation as in Remark 3.4), the cut spaces are

$$\{\Phi \geq c\}/\sim, \quad \Phi^{-1}(c)/\sim, \quad \{\Phi \leq c\}/\sim,$$

and in all three spaces, the relation $\sim$ quotients $\Phi^{-1}(c)$ by the Hamiltonian S^1-action.

Example 3.7. The multiple cut in Figure 3.1 is made up of two simultaneous single cuts along hypersurfaces that intersect ω-orthogonally. The set of polytopes is

$$\mathcal{P} = \{P_i, 0 \leq i \leq 3, P_{ij}, j = (i+1) \mod 4, P_\cap\}.$$

The manifolds $X_{P_{i(i+1)}}$, $X_{P_{(i-1)i}}$ are relative divisors of X_{P_i}.

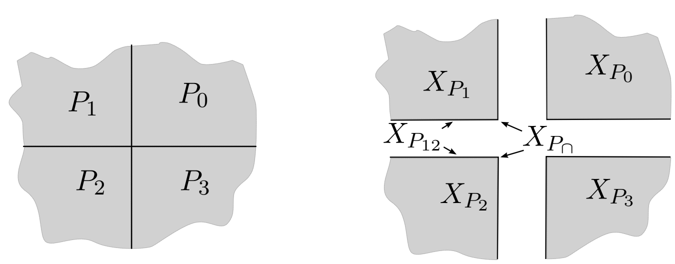

Figure 3.1. A multiple cut of $\mathbb{R}^2$

3.3. Symplectic broken manifolds

In this section we describe the broken manifold as a symplectic space. The symplectic broken manifold consists of top-dimensional cut spaces and thickenings of the lower dimensional cut spaces into toric fibrations. The toric fibrations are defined by considering neighborhoods of cut loci in X and modding out the boundaries as in Lerman's construction. This construction requires the additional datum of a dual complex associated to the polyhedral decomposition $\mathcal{P}$. The dual complex $B^\vee$

(defined later in the section) is a union of dual polytopes $\cup_{P\in\mathcal{P}} P^\vee$, and $P^\vee$ is the moment polytope of the fibers of the toric fibration $X_{\overline{P}}$ over the cut space X_P.

Definition 3.8. (A complex of polytopes) Let V be an affine space. A *complex of polytopes* in V is a collection B of polytopes satisfying the following:

(a) (Faces) If $P \in B$ then all the faces of P are also elements of P.
(b) (Intersection) If $P, Q \in B$, then $P \cap Q$ is a face of both P and Q and hence is an element of B.

We fix an inner product

$$\mathfrak{t} \times \mathfrak{t} \to \mathbb{R} \tag{3.6}$$

that restricts to an integral product $\mathfrak{t}_{\mathbb{Z}} \times \mathfrak{t}_Z \to \mathbb{Z}$.

Definition 3.9. (Dual complex) The *dual complex* for a tropical action $(X, \Phi, \mathcal{P})$ is a complex

$$B^\vee := \{P^\vee : P \in \mathcal{P}\}$$

of polytopes in $\mathfrak{t}$ (in the sense of Definition 3.8) consisting of a *dual polytope* $P^\vee$ for each polytope P in $\mathcal{P}$ satisfying

(a) $Q, P \in \mathcal{P}$, $Q \subset P \implies P^\vee \subset Q^\vee$,
(b) for any $P \in \mathcal{P}$, $P^\vee$ is a simple polytope of dimension $\dim(\mathfrak{t}) - \dim(P)$ lying in $c + \mathfrak{t}_P \subset \mathfrak{t}$ for some constant $c \in \mathfrak{t}$. With respect to the identification $\mathfrak{t}^\vee \simeq \mathfrak{t}$ from the inner product (3.6) on $\mathfrak{t}$, the polytopes $P, P^\vee \subset \mathfrak{t}^\vee$ are orthogonal.

Example 3.10. In the case of a single cut (X, Φ, c) (see Remark 3.4 for notation), the set of polytopes is

$$\mathcal{P} = \{P_- := (-\infty, c], P_0 := \{c\}, P_+ := [c, \infty)\},$$

and the dual polytopes are

$$P_-^\vee = \{-\epsilon\}, \quad P_0^\vee = [-\epsilon, \epsilon], \quad P_+^\vee = \{\epsilon\},$$

for any $\epsilon > 0$, with identifications $P_\pm^\vee \hookrightarrow P_0^\vee$ given by inclusion of the endpoints. Thus the dual complex is an interval: $B^\vee \cong [-\epsilon, \epsilon]$.

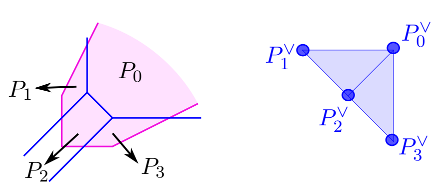

Figure 3.2. A polyhedral decomposition (left) and its dual complex (right).

Example 3.11. Figure 3.2 shows the dual complex corresponding to a multiple cut on a toric surface. We recall that we used three such multiple cuts to study disks in a cubic surface in Section 2.1.

The broken manifold is a collection of manifolds $X_{\overline{P}}$ corresponding to polytopes $P \in \mathcal{P}$, each of which is a toric fibration over the cut space X_P. The manifold $X_{\overline{P}}$ is modelled on a fibered polytope defined next.

Definition 3.12. (Fibered polytope) Let P, $P^\vee$ be complementary dimensional polytopes in $\mathfrak{t}^\vee$ that intersect transversely. A polytope $\overline{P} \subset \mathfrak{t}$ is *fibered over P with fiber $P^\vee$* if it is equipped with a submersion

$$\overline{P} \xrightarrow{\pi_P} P$$

whose fibers are polytopes, and a map

$$\overline{P} \xrightarrow{\pi_{P^\vee}} P^\vee$$

which is the restriction of the linear projection $\mathfrak{t}^\vee \to \mathfrak{t}_P^\vee$ such that

- on any fiber $\pi_P^{-1}(x)$, $\pi_{P^\vee}$ is a bijection,
- and $P = \pi_{P^\vee}^{-1}(c)$ for some interior point $c \in P^\vee$.

The facets of the fibered polytope $\overline{P}$ are

$$\text{Facets}(\overline{P}) = \{\pi_P^{-1}(Q) : Q \in \text{Facets}(P), Q \in \mathcal{P}\} \cup \{\pi_{P^\vee}^{-1}(Q) : Q \in \text{Facets}(P^\vee)\}. \tag{3.7}$$

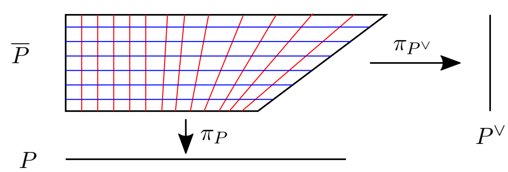

Figure 3.3. A fibered polytope $\overline{P}$. The fibers of π_P are in red, and those of $\pi_{P^\vee}$ are in blue.

Definition 3.13. (Cutting datum) Given a tropical manifold $(X, \mathcal{P}, \Phi)$ a *cutting datum* consists of

- a dual complex $B^\vee$,
- and for each polytope P a *fibered neighborhood* $\overline{P} \subset \mathfrak{t}^\vee$ given by the inclusion

$$\mathrm{i}_P : \overline{P} \hookrightarrow \mathfrak{t}^\vee. \tag{3.8}$$

By a fibered neighborhood we mean that $\overline{P} \to P$ is a fibered polytope (as in Definition 3.12) with fibers $P^\vee$:

$$\overline{P} \xrightarrow{\pi_P} P$$

such that

- the projection $\mathfrak{t}^\vee \to \mathfrak{t}_P^\vee$ generates a free T_P-action on $\Phi^{-1}(\overline{P})$,
- and for $x \in P$ in a neighborhood of a face $Q \subset P$ the fiber $\pi_P^{-1}(x)$ is contained in $\pi_Q^{-1}(x)$ and is orthogonal to $P \cap \pi_Q^{-1}(x)$ with respect to the fixed metric (3.6) on $\mathfrak{t}^\vee$.

Indeed, such fibered neighborhoods exist possibly after scaling down the dual complex.

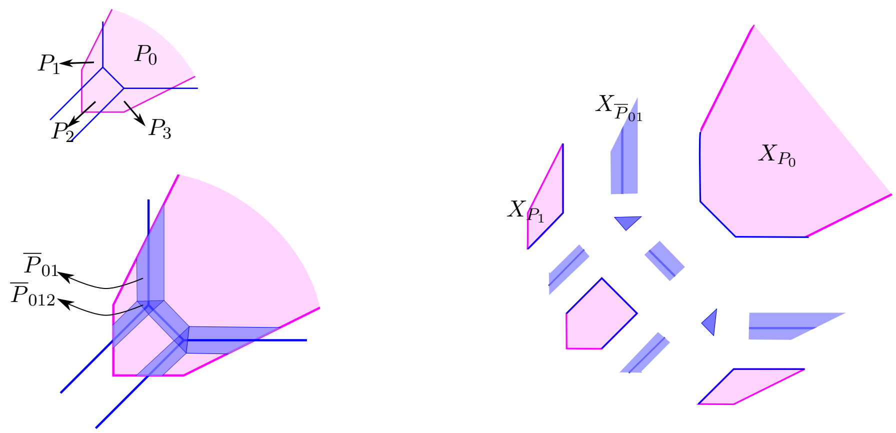

FIGURE 3.4. Top left: Polyhedral decomposition from Figure 3.2. Bottom left: Cutting datum. Right: The broken symplectic manifold is isomorphic to the inverse image of these polytopes under the tropical moment map.

DEFINITION 3.14. (Symplectic broken manifold) Suppose $(X, \mathcal{P}, \Phi)$ is a tropical Hamiltonian manifold. Suppose each of the polytopes $P \in \mathcal{P}$ possesses a fibered neighborhood

$$\mathfrak{t} \supset \overline{P} \xrightarrow{\pi_P} P$$

with fibers $P^\vee$ such that for $x \in P$ in a neighborhood of a face $Q \subset P$ the fibers $\pi_P^{-1}(x)$ are translates of $P^\vee \subset \mathfrak{t}_P^\vee \subset \mathfrak{t}_Q^\vee$.

(a) The *symplectic broken manifold* corresponding to the tropical manifold $(X, \Phi, \mathcal{P})$ is a disjoint union

$$\mathfrak{X}_{\mathcal{P}} \quad \text{or} \quad \mathfrak{X} := \sqcup_{P \in \mathcal{P}} X_{\overline{P}},$$

where we use the notation $\mathfrak{X}$ when $\mathcal{P}$ is obvious from the context. Here

$$X_{\overline{P}} := \Phi^{-1}(\overline{P})/\sim, \tag{3.9}$$

$\sim$ mods out any boundary component $\Phi^{-1}(Q)$, $Q \in \text{Facets}(\overline{P})$ by the S^1-action generated by the vector $\nu_Q \in \mathfrak{t}_{\mathbb{Z}}$ normal to the hyperplane $Q \in \mathfrak{t}^\vee$. The space $X_{\overline{P}}$ is a smooth symplectic manifold if all faces $F \subset \overline{P}$ are smooth (as in Definition 3.2). Otherwise $X_{\overline{P}}$ is a symplectic orbifold (see Remark 3.15 below).

(b) (Relative divisors of a broken manifold) For any facet Q of $\overline{P}$ the quotient

$$Y_Q := \Phi^{-1}(Q)/\exp(\mathbb{R}\nu_Q) \subset X_{\overline{P}}$$

is a *relative divisor* of $X_{\overline{P}}$. Here the normal $\nu_Q \in \mathfrak{t}_{\mathbb{Z}}$ to the facet Q generates a subgroup $\exp(\mathbb{R}\nu_Q) \simeq S^1$. The relative divisor Y_Q is *horizontal* resp. *vertical* if the facet $Q \subset \overline{P}$ corresponds to a facet of P resp. $P^\vee$ (see (3.7)).

Thus

$$X^{\square}_{\overline{P}} = X_{\overline{P}} - \bigcup_{Q \subset \overline{P}} Y_Q$$

is the complement of all the relative divisors $Y_Q, Q \subset \overline{P}$ of $X_{\overline{P}}$.

REMARK 3.15. Orbifold singularities in the broken manifolds do not pose an issue in the analysis of holomorphic curves. Indeed, we treat our holomorphic curves as lying in manifolds with cylindrical ends $X^{\square}_{\overline{P}}$ and not in their possibly-orbifold compactifications. Whenever compactifications are used, we provide work-arounds to avoid dealing with holomorphic curves in orbifolds.

EXAMPLE 3.16. In the case of a single cut (X, Φ, c) the broken manifold consists of the components

$$\{\Phi > c\}/\sim, \quad \{\Phi < c\}/\sim, \quad \Phi^{-1}([c-\epsilon, c+\epsilon])/\sim,$$

where $\sim$ mods out the boundaries $\Phi^{-1}(c-\epsilon)$ and $\Phi^{-1}(c+\epsilon)$ by S^1-actions.

REMARK 3.17. (Components of the broken manifold as thickenings) Components of the broken manifold are fibrations over cut spaces whose fibers are complex tori. Therefore they can be regarded as thickenings of the cut spaces. The compactification $X_{\overline{P}}$ of $X^{\square}_{\overline{P}}$ is a toric fibration

$$V_{P^\vee} \to X_{\overline{P}} \xrightarrow{\pi_P} X_P, \tag{3.10}$$

whose fiber $V_{P^\vee}$ is a T_P-toric manifold with moment polytope $P^\vee$.

DEFINITION 3.18. (Torus bundles on cut spaces) For any $P \in \mathcal{P}$, the principal T_P-bundle

$$Z_P \to X_P$$

is defined as $Z_P := \Phi^{-1}(P)/\sim$ where, for any facet $Q \subset P$, $Q \in \mathcal{P}$, $\sim$ mods out $\Phi^{-1}(Q)$ by T_Q/T_P, viewed as a subgroup of T_P via the pairing (3.6) on $\mathfrak{t}_P$. Alternately, Z_P may be defined as the T_P-bundle for which $X_{\overline{P}} \to X_P$ is the associated $V_{P^\vee}$-bundle, that is

$$X_{\overline{P}} = Z_P \times_{T_P} V_{P^\vee},$$

where $V_{P^\vee}$ is a T_P-toric manifold with moment polytope $P^\vee$.

EXAMPLE 3.19. For the multiple cut in Figure 3.1 the dual complex is a rectangle and the broken manifold $\mathfrak{X}$ is as in Figure 3.5. Relative submanifolds $X_{P_{ij}}$ and $X_{P_\cap}$ are thickened into toric fibrations $X_{\overline{P}_{ij}}$ and $X_{\overline{P}_\cap}$

$$\mathbb{P}^1 \to X_{\overline{P}_{ij}} \to X_{P_{ij}}, \quad (\mathbb{P}^1)^2 \to X_{\overline{P}_\cap} \to X_{P_\cap}$$

in the broken manifold.

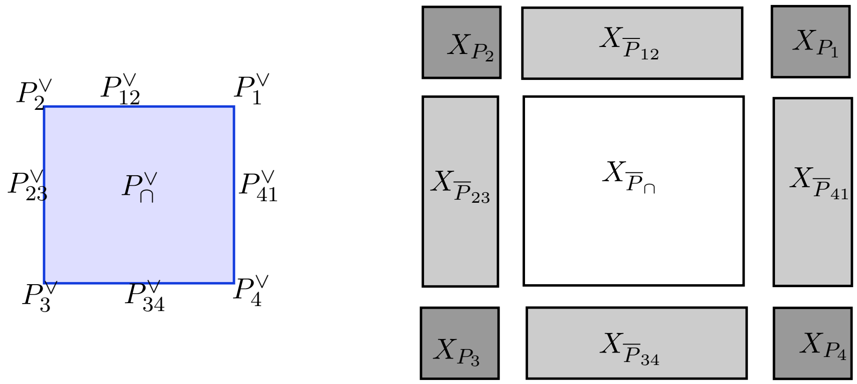

Figure 3.5. Dual complex and broken manifold for the multiple cut in Figure 3.1

.

3.4. Neck-stretched almost complex structures

We define a family of almost complex structures on the manifold X, called neck-stretched almost complex structures. These almost complex structures are 'cylindrical' in the sense of the following Definition:

Definition 3.20. (Cylindrical almost complex structure)

(a) (P-cylinder) For a polytope $P \in \mathcal{P}$ a $T_{P,\mathbb{C}}$-principal bundle

$$Z_{\mathbb{C}} \to M$$

on a manifold M is called a *P-cylinder*. There is a $T_{\mathbb{C}}$-diffeomorphism

$$Z_{\mathbb{C}} \simeq Z \times i\mathfrak{t}_P$$

where $Z \to M$ is a principal T_P-bundle and the $T_{\mathbb{C}}$-action on $Z \times i\mathfrak{t}_P$ is

$$te^{is}(z, s_0) = (tz, s_0 + s), \quad t \in T_P, s \in \mathfrak{t}_P.$$

(b) Let $Z_{\mathbb{C}} \to M$ be a P-cylinder. An almost complex structure J is *P-cylindrical* iff

(i) there exists an almost complex structure J_M on the base manifold M such that the projection

$$\pi : Z_{\mathbb{C}} \to M$$

is almost complex, that is, $d\pi \circ J = J_M \circ d\pi$;

(ii) there exists a connection one-form $\alpha \in \Omega^1(Z, \mathfrak{t}_P)$ on the T_P-bundle $Z \to M$ such that the horizontal sub-bundle

$$H := \ker(\alpha) \subset TZ \subset TZ_{\mathbb{C}} \tag{3.11}$$

is J-invariant;

(iii) on any fiber $\pi^{-1}(m) \subset Z_{\mathbb{C}}$, J is standard in the following sense: For any point $z \in \pi^{-1}(m)$ the map

$$T_{P,\mathbb{C}} \to \pi^{-1}(m), t \mapsto tz$$

is a biholomorphism.

As a result, J is invariant under the $T_{P,\mathbb{C}}$-action on $Z_{\mathbb{C}}$. Denote by

$$\mathcal{J}^{\mathrm{cyl}}(Z_{\mathbb{C}}) = \{ J \in \mathcal{J}(Z_{\mathbb{C}}) \mid (i) - (iii) \}$$

the space of P-cylindrical almost complex structures on $Z_{\mathbb{C}}$. Note that a P-cylindrical almost complex structure J is determined by its projection $d\pi(J)$ to the base M called the *base almost complex structure* and its *associated connection one-form* $\alpha(J)$.

Given a tropical manifold $(X, \mathcal{P}, \Phi)$ we aim to define 'neck-stretched manifolds' where the neck region associated to $P \in \mathcal{P}$ has a P-cylindrical almost complex structure. To achieve this end, a subset of the manifold X with the action of the torus T_P has to be identified with a subset of a P-cylinder. This identification is made via a symplectic cylindrical structure defined below, where a neighborhood of $\Phi^{-1}(P)$ in X possessing a T_P-action is identified to a product $\Phi^{-1}(P) \times P^\vee$. In what follows the identification between $\mathfrak{t}_P$ and $\mathfrak{t}_P^\vee$ from (3.6) is crucial. For the symplectic structure, we view $P^\vee$ as a subset of $\mathfrak{t}_P^\vee$ and the T_P-moment map on $\Phi^{-1}(P) \times P^\vee$ is given by the projection to $P^\vee$. For the P-cylindrical complex structure, we view $P^\vee$ as a subset of $\mathfrak{t}_P$ so that the fibers of the projection

$$\Phi^{-1}(P) \times P^\vee \to \Phi^{-1}(P)/T_P$$

are subsets of $T_{P,\mathbb{C}}$ via

$$T_P \times P^\vee \subset T_P \times \mathfrak{t}_P^\vee \xrightarrow{\simeq} T_P \times \mathfrak{t}_P \xrightarrow{(t,\xi) \mapsto t\exp(i\xi)} T_{P,\mathbb{C}},$$

where the identification $\mathfrak{t}_P^\vee \simeq \mathfrak{t}_P$ in the middle arrow is via the product on $\mathfrak{t}$ from (3.6). Consequently $\Phi^{-1}(P) \times P^\vee$ has a partial $T_{P,\mathbb{C}}$-action, meaning that there is an infinitesimal action

$$\mathfrak{t}_{P,\mathbb{C}} \to \mathrm{Vect}(\Phi^{-1}(P) \times P^\vee) \tag{3.12}$$

whose flows satisfy the axioms for an action of $T_{P,\mathbb{C}}$ wherever they are defined.

Definition 3.21. (Symplectic cylindrical structure on tropical manifolds) Let $(X, \mathcal{P}, \Phi)$ be a symplectic manifold with a tropical Hamiltonian action. A *symplectic cylindrical structure* $\underline{\phi} = (\phi_P)_{P \in \mathcal{P}}$ consists of a T_P-equivariant symplectomorphism ϕ_P

$$\Phi^{-1}(\overline{P}) \xrightarrow{\phi_P} (\Phi^{-1}(P) \times P^\vee, \overline{\omega}), \quad \overline{\omega} := (\omega_X | \Phi^{-1}(P)) + d\langle \alpha_P, \pi_{P^\vee} - c_{P^\vee} \rangle \tag{3.13}$$

for each polytope $P \in \mathcal{P}$. Here

(a) $\alpha_P \in \Omega^1(\Phi^{-1}(P), \mathfrak{t}_P)$ is a T_P-connection one-form;
(b) $\pi_{P^\vee} : \Phi^{-1}(P) \times P^\vee \to P^\vee \subset \mathfrak{t}_P^\vee$ is the projection to the second factor, $c_{P^\vee} \in \mathfrak{t}_P^\vee$ is the constant given by projecting $P \subset \mathfrak{t}^\vee$ to $\mathfrak{t}_P^\vee$, and $\langle \cdot, \cdot \rangle$ is the pairing between $\mathfrak{t}_P$, $\mathfrak{t}_P^\vee$, and so, $\langle \alpha_P, \pi_{P^\vee} - c_{P^\vee} \rangle$ is a one-form on $\Phi^{-1}(P) \times P^\vee$;
(c) the dual polytopes $P^\vee$ are assumed to be small enough that the forms in the right hand side of (3.13) are symplectic;

and for any $P \in \mathcal{P}$ the second component of ϕ_P is the T_P-moment map, that is,

$$\pi_{P^\vee} \circ \phi_P = \Phi,$$

and the maps $\underline{\phi}$ satisfy the following (Patching) condition:

> (Patching) For any pair $Q \subset P$, in the overlap region $\overline{Q} \cap \overline{P}$ the $T_{P,\mathbb{C}}$-action induced by ϕ_P (see the explanation preceding this definition) is the restriction of the $T_{Q,\mathbb{C}}$-action induced by ϕ_Q.

This ends the Definition.

In the above definition the (Patching) condition automatically places the following consistency condition on connection one-forms $(\alpha_P)_{P\in\mathcal{P}}$. For a pair of polytopes $Q \subset P$ the connection α_P on $\Phi^{-1}(P)$ is determined by α_Q in a neighborhood of $\Phi^{-1}(Q)$ in $\Phi^{-1}(P)$ via the consistency condition in the following Lemma.

Lemma 3.22. (Consistency for connection one-forms) *Consider polytopes $Q \subset P$ in $\mathcal{P}$. Given a symplectic structure map ϕ_Q as in (3.13), there is a natural choice of the map $\phi_P|\Phi^{-1}(\overline{P} \cap \overline{Q})$ and the connection one-form $\alpha_P|\Phi^{-1}(\overline{P} \cap \overline{Q})$ for which the (Patching) condition is satisfied. This choice of (ϕ_P, α_P) is unique up to the action of T_P-gauge transformations on the bundle $\Phi^{-1}(P) \to X_P^{\square}$.*

Proof. We first describe the symplectic cylindrical structure map ϕ_P. The inner product on $\mathfrak{t}$ in (3.6) gives the following orthogonal projection

$$Q^\vee \to P^\vee \times \mathfrak{t}_Q^\vee/\mathfrak{t}_P^\vee, \tag{3.14}$$

and the image lies in $P^\vee \times \operatorname{NCone}_Q P$ if we restrict the map to an appropriate neighborhood in $Q^\vee$. Viewing $\operatorname{NCone}_Q P$ as a subset of $Q^\vee$, via the cylindrical structure map ϕ_Q, we obtain an embedding

$$\Phi^{-1}(Q) \times \operatorname{NCone}_Q P \to \Phi^{-1}(P) \tag{3.15}$$

defined on a neighborhood of the origin in $\operatorname{NCone}_Q P$. The map ϕ_P is defined on $\Phi^{-1}(\overline{Q})$ by (3.14) and (3.15). The connection one-form is defined by the condition

$$\ker(\alpha_P(x)) = \ker(\alpha_Q(x)) \oplus (\mathfrak{t}_Q/\mathfrak{t}_P)x, \tag{3.16}$$

for any $x \in \Phi^{-1}(\overline{Q} \cap \overline{P})$ where we view the quotient $\mathfrak{t}_Q/\mathfrak{t}_P$ as a subspace of $\mathfrak{t}_Q$ by the pairing (3.6) on $\mathfrak{t}_P$. $\square$

For the moment, we assume the existence of symplectic cylindrical structures and use them to define neck-stretching. The proof of the existence of symplectic cylindrical structures on tropical manifolds is deferred to Lemma 3.43 at the end of Chapter 3. The cylindrical structure maps $\{\phi_P\}_{P\in\mathcal{P}}$ in (3.13) are fixed throughout the paper.

Remark 3.23. (A decomposition of polytopes) We describe a decomposition of a tropical manifold that is used to define neck-stretched manifolds. Let $(X, \mathcal{P}, \Phi)$ be a tropical manifold and let $B^\vee$ be a dual complex $B^\vee$. For any $P \in \mathcal{P}$ let

$$P^{\blacksquare} := P \backslash (\cup_{Q \subset P} \mathrm{i}_Q(\overline{Q})), \tag{3.17}$$

be the complement of fibered neighborhoods of faces of P, where $\mathrm{i}_Q : \overline{Q} \to \mathfrak{t}^\vee$ is the inclusion from (3.8). Corresponding to any facet $Q \subset P$, $P^\blacksquare$ has a facet $Q^\blacksquare$.

$$\overline{P}^\blacksquare := \pi_P^{-1}(P^\blacksquare) \subset \overline{P} \tag{3.18}$$

be the thickening of $P^\blacksquare$. For a pair $Q \subset P$ with $\mathrm{codim}_P(Q) = 1$, the fibered polytopes $\overline{P}^\blacksquare, \overline{Q}^\blacksquare \subset \mathfrak{t}^\vee$ share a facet, which is isomorphic to $Q^\blacksquare \times P^\vee$, see Figure 3.6. Then the image of Φ has the following cover

$$\mathrm{im}(\Phi) = \cup_{P\in\mathcal{P}} \mathrm{i}_P(\overline{P}^\blacksquare)/\sim, \tag{3.19}$$

where $\sim$ identifies shared facets of polytopes. The partition of $\mathrm{im}(\Phi)$ pulls back to a partition of the manifold (X, ω)

$$(X,\omega) := \left(\bigsqcup_{P\in\mathcal{P}} \Phi^{-1}(\overline{P}^\blacksquare) \right) / \sim \tag{3.20}$$

into manifold with corners, where the identifications in $\sim$ are along the boundaries and are induced by the inclusions $\Phi^{-1}(\overline{P}^\blacksquare) \to X$. Further, the symplectic cylindrical structure map $\underline{\phi} = (\phi_P)_P$ may be used to rewrite the decomposition in (3.20) as

$$(X,\omega) := \left(\bigsqcup_{P\in\mathcal{P}} \Phi^{-1}(P^\blacksquare) \times P^\vee \right) / \sim . \tag{3.21}$$

In (3.21), the equivalence $\sim$ identifies the boundary components

$$\Phi^{-1}(Q^\blacksquare) \times Q^\vee \supset \Phi^{-1}(Q^\blacksquare) \times P^\vee \xrightarrow{\sim} \Phi^{-1}(Q^\blacksquare) \times P^\vee \subset \Phi^{-1}(P^\blacksquare) \times P^\vee \tag{3.22}$$

for all pairs $Q \subset P$, $\mathrm{codim}_P(Q) = 1$.

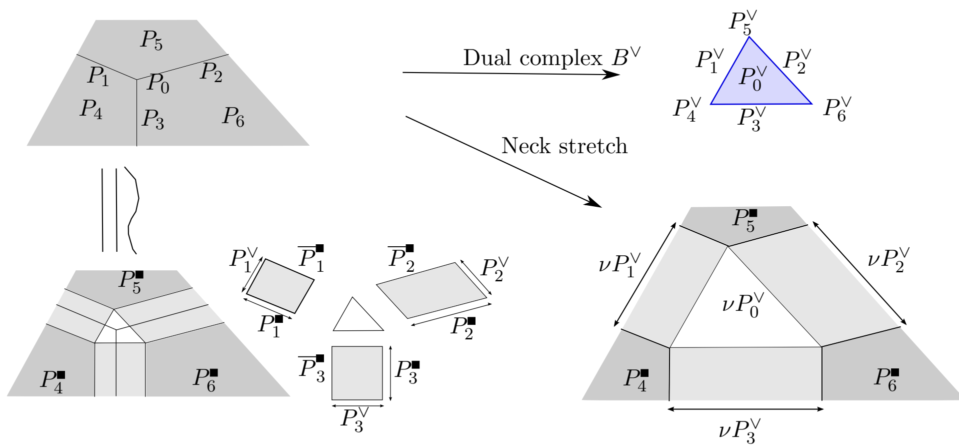

FIGURE 3.6. Stretching

REMARK 3.24. (A warning about Figure 3.6) Figure 3.6 is a schematic representation of neck-stretched manifolds. We recall symplectic broken manifolds are represented by the images of their tropical moment map, for example in Figure 3.4. In contrast, we do not have a moment map on neck-stretched manifolds. A larger polytope is just used to indicate a larger cylinder.

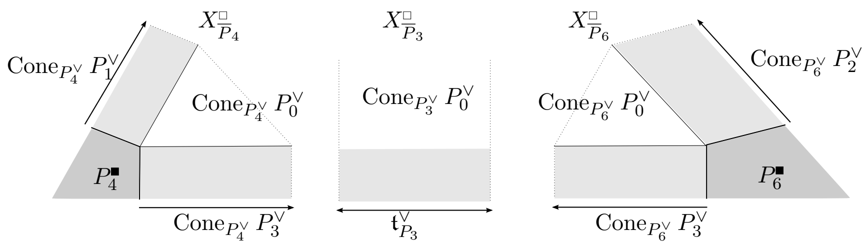

FIGURE 3.7. Some components of the broken manifold in the limit $\nu \to \infty$ of the neck-stretching in Figure 3.6.

DEFINITION 3.25. (Neck-stretched manifolds) Let $(X, \mathcal{P}, \Phi)$ be a tropical manifold with a symplectic cylindrical structure. For any $\nu \in \mathbb{R}_{\geq 1}$, define a *neck-stretched manifold* X^ν as

$$X^\nu = \left(\bigsqcup_{P \in \mathcal{P}} \Phi^{-1}(P^{\blacksquare}) \times \nu P^\vee \right) / \sim, \tag{3.23}$$

where the equivalence relation $\sim$ is exactly as in (3.22), that is, for all pairs $Q \subset P$, $\mathrm{codim}_P(Q) = 1$, the boundary component $\Phi^{-1}(Q^{\blacksquare}) \times \nu P^\vee$ in $\Phi^{-1}(P^{\blacksquare}) \times \nu P^\vee$ is identified to the corresponding boundary component in $\Phi^{-1}(Q^{\blacksquare}) \times Q^\vee$ by the identity map. Thus the neck-stretched manifold X^ν is equipped with

(a) for each P, a P-cylindrical structure on the subset $\Phi^{-1}(P^{\blacksquare}) \times \nu P^\vee \subset X^\nu$, that is, there is a projection

$$\Phi^{-1}(P^{\blacksquare}) \times \nu P^\vee \to \Phi^{-1}(P^{\blacksquare})/T_P \tag{3.24}$$

whose fibers are $T_P \times \nu P^\vee \subset T_{P,\mathbb{C}}$;

(b) and a symplectic form on the base manifold $\Phi^{-1}(P^{\blacksquare})/T_P$, which is a T_P-reduction of the symplectic form ω_X on X.

This ends the definition.

REMARK 3.26. The neck-stretched manifold X^ν is diffeomorphic to X, but the diffeomorphism is not canonical. As a result there is no canonical symplectic form on X^ν.

EXAMPLE 3.27. (Neck-stretched manifolds for a single cut) We describe neck-stretched manifolds in case of a single sut. Let (X, Φ, c) be a tropical manifold with a single cut along $\Phi^{-1}(c)$, and whose polytopes are as in Example 3.10. The symplectic cylindrical structure consists of an S^1-equivariant symplectomorphism defined on a neighborhood

$$\Phi^{-1}([c - \epsilon, c + \epsilon]) \xrightarrow{\phi_{P_0}} (\Phi^{-1}(c) \times [-\epsilon, \epsilon], \omega_{/\!/} + d(t\alpha)), \tag{3.25}$$

where $\omega_{/\!/}$ is the reduced symplectic form on the quotient $\Phi^{-1}(c)/S^1$, $t \in (-\epsilon, \epsilon)$ is the coordinate function, and $\alpha \in \Omega^1(\Phi^{-1}(c))$ is a connection one-form on the S^1-bundle $\Phi^{-1}(c) \to \Phi^{-1}(c)/S^1$. The other symplectic cylindrical structure maps ϕ_{P_-},

ϕ_{P_+} are trivial since P_+, P_- are top-dimensional. The neck-stretched manifold is then

$$X^\nu := \{|\Phi - c| \geq \epsilon\} \cup_\phi (\Phi^{-1}(c) \times [-\epsilon\nu, \epsilon\nu]),$$

where the attachment map ϕ is given by the restriction of ϕ_{P_0} to the boundary $\Phi^{-1}(\{c-\epsilon, c+\epsilon\})$.

Definition 3.28. (Cylindrical almost complex structures on neck-stretched manifolds) Let $(X, \mathcal{P})$ be a tropical Hamiltonian action and let X^ν be a sequence of neck-stretched manifolds. Recall from (3.23) that a neck-stretched manifold has a decomposition

$$X^\nu = \left(\bigsqcup_{P \in \mathcal{P}} \Phi^{-1}(P^\blacksquare) \times \nu P^\vee \right) / \sim,$$

where $\sim$ identifies boundaries of the different components.

(a) An almost complex structure J^ν on the neck-stretched manifold X^ν is *cylindrical* if J^ν is P-cylindrical in the sense of Definition 3.20 in the subset

$$\Phi^{-1}(P^\blacksquare) \times \nu P^{\vee,\circ} \subset X^\nu.$$

We say that (X^ν, J^ν) is a *family of neck-stretched almost complex structures* if there are ν-independent cylindrical almost complex structures J_P on $\Phi^{-1}(P^\blacksquare) \times \mathfrak{t}_P^\vee$ for all P, such that on the P-cylindrical subset $\Phi^{-1}(P^\blacksquare) \times \nu P^{\vee,\circ}$ of X^ν, J^ν is the restriction of J_P.

(b) (Local tamedness) Let (X^ν, J^ν) be a family of neck-stretched almost complex structures. We say that each element J^ν is *locally tamed* if J^1 tames ω_X. (The definition is based on the observation that the neck-stretched manifold X^1 has a canonical diffeomorphism to (X, ω), and this is not so for X^ν, $\nu \neq 1$.)

(c) (Local strong tameness) A cylindrical locally tamed almost complex structure J on X^ν is *locally strongly tamed* if for all $P \in \mathcal{P}$ the connection one-form $\alpha_{P,J}$ underlying the P-cylindrical almost complex structure on $\Phi^{-1}(P^\blacksquare) \times \nu P^{\vee,\circ}$ in Definition 3.20 is the same as the connection one-form underlying the symplectic cylindrical structure in Definition 3.21.

Remark 3.29. The following are some remarks on cylindrical almost complex structures on neck-stretched manifolds.

(a) The definition of local (strong) tameness depends not just on the symplectic form ω_X but also on the symplectic cylindrical structure in Definition 3.21.

(b) In the space of cylindrical almost complex structures the cylindrical coordinate maps, taking values in $\Phi^{-1}(P^\blacksquare) \times \nu P^\vee$, are held fixed, but the connection one-forms $\alpha_{P,J}$, $P \in \mathcal{P}$ underlying a cylindrical almost complex structure J are allowed to vary.

(c) We give an alternate definition of local tamedness. A cylindrical almost complex structure J^ν is *locally tamed* on X^ν if on the P-cylinder $\Phi^{-1}(P^\blacksquare) \times \nu P^\vee$, the base almost complex structure $d\pi_P(J^\nu)$ is tamed by $\omega_{X_P} + \langle d\alpha_P, c\rangle$ for all $c \in P^\vee - c_{P^\vee} \subset \mathfrak{t}_P^\vee$. Here $c_{P^\vee} := \pi_{\mathfrak{t}_P^\vee}(P) \in \mathfrak{t}_P^\vee$ is a constant and α_P is the T_P-connection one-form underlying the symplectic cylindrical structure in Definition 3.21. In other words, local tamedness means that $d\pi_P(J^\nu)$ is tamed by all the symplectic forms that occur in the horizontal complement of the P-cylinder in the symplectic cylindrical structure fixed in Definiton 3.21. Indeed this definition of local tamedness is equivalent to the one in Definition 3.28 (b), since tamedness in the fibers of the P-cylinder is automatic.

The following Lemma is a quantitative version of the statement that local tamedness is a C^0-open property in the space of cylindrical almost complex structures. We use the alternate definition of local tamedness from Remark 3.29 (c).

LEMMA 3.30. *Let J_P^0 be an almost complex structure on $\Phi^{-1}(P^\blacksquare)/T_P$ that is tamed by the form $\omega_{X_P} + \langle d\alpha_P, \tau\rangle$ for all $\tau \in P^\vee \subset \mathfrak{t}_P^\vee$. Then there exist constants ϵ, c such that if $J_P \in \mathcal{J}(\Phi^{-1}(P^\blacksquare)/T_P)$ is such that $\|J_P - J_P^0\|_{C^0} < \epsilon$ then*

(3.26)

$$c^{-1}\omega_{X_P}(v, Jv) \leq (\omega_{X_P} + \langle d\alpha_P, c\rangle)(v, Jv) \leq c\omega_{X_P}(v, Jv) \quad \forall v \in T(\Phi^{-1}(P^\blacksquare)/T_P).$$

PROOF. By the compactness of the spaces $P^\blacksquare/T_P$ and $P^\vee$, and the tamedness of J_P^0 for $\omega_{X_P} + \langle d\alpha_P, \tau\rangle$ for all $\tau \in P^\vee$, we conclude that there are constants $c_0, c_1 > 0$ such that

$$c_0|v|^2 \leq (\omega_{X_p} + \langle d\alpha_P, \tau\rangle)(v, J_P^0 v) \leq c_1|v|^2$$

for all $v \in T(\Phi^{-1}(P^\blacksquare)/T_P$, $\tau \in P^\vee$. For any $\epsilon > 0$ and $J_P \in B_\epsilon(J_P^0)$, there are constants $c_0', c_1' > 0$ such that

$$(c_0 - c_0'\epsilon)|v|^2 \leq (\omega_{X_p} + \langle d\alpha_P, \tau\rangle)(v, J_P^0 v) \leq (c_1 + c_1'\epsilon)|v|^2.$$

The Lemma follows by choosing $\epsilon > 0$ such that $c_0 - c_0'\epsilon > 0$ and $c := \frac{c_1 + c_1'\epsilon}{c_0 - c_0'\epsilon}$. □

3.5. Broken manifold as a degenerate limit

3.5.1. Defining almost complex broken manifolds. The broken manifold is the degenerate limit of neck-stretched almost complex manifolds as we now explain. For a polytope P and $\nu \geq 1$, let

$$X_{\overline{P}}^\nu := \left(\Big(\bigsqcup_{Q \in \mathcal{P}: Q \subseteq P} \Phi^{-1}(P^\blacksquare) \times \nu P^\vee \Big) / \sim \right) \subset X^\nu \tag{3.27}$$

be the subset of X^ν that has a P-cylindrical structure, and thus $X_{\overline{P}}^\nu$ has a partial $T_{P,\mathbb{C}}$-action (see (3.12)) that is free and proper. (Here the equivalence relation $\sim$ is the same as the one in the definition 3.23 of X^ν.) Let

$$X_P^\nu := X_{\overline{P}}^\nu / T_{P,\mathbb{C}} \tag{3.28}$$

be the quotient of $X_{\overline{P}}^\nu$ under the partial $T_{P,\mathbb{C}}$-action. Consequently there are projections

$$X_{\overline{P}}^\nu \xrightarrow{\pi_P'} Z_P^\nu \xrightarrow{\pi_P''} X_P^\nu, \tag{3.29}$$

where $Z_P^\nu \xrightarrow{\pi_P''} X_P^\nu$ is a T_P-bundle, and the fibers of π_P' are subsets of $\mathfrak{t}_P^\vee$.

LEMMA 3.31. *Let (X^ν, J^ν) be a family of neck-stretched almost complex structures (as in Definition 3.28). For any $P \in \mathcal{P}$ and $\nu_0 < \nu_1$ there are natural inclusions*

$$i_P : (X_P^{\nu_0}, J^{\nu_0}) \to (X_P^{\nu_1}, J^{\nu_1}), \quad i_{Z_P} : Z_P^{\nu_0} \to Z_P^{\nu_1}.$$

PROOF. Let $P \in \mathcal{P}$ be a top-dimensional polytope. Then, $X_P^\nu = X_{\overline{P}}^\nu$ has a decomposition as in (3.27). we define inclusions

$$i_P := (\mathrm{Id}, \tau) : \Phi^{-1}(Q^\blacksquare) \times \nu_0 Q^\vee \to \Phi^{-1}(Q^\blacksquare) \times \nu_1 Q^\vee$$

where τ is the restriction of a translation on $\mathfrak{t}_Q^\vee$ (recall that $\nu Q^\vee \subset \mathfrak{t}_Q^\vee$ for all ν) that maps the point $\nu_0 P^\vee \in \nu_0 Q^\vee$ to the point $\nu_1 P^\vee \in \nu_1 Q^\vee$; and the maps glue to yield an inclusion $i_P : X_P^{\nu_0} \to X_P^{\nu_1}$.

Next, consider a face $Q \in \mathcal{P}$ of the top-dimensional polytope P. The inclusion $i_P : X_{\overline{Q}}^{\nu_0} \to X_{\overline{Q}}^{\nu_1}$ descends to $i_{Z_Q} : Z_Q^{\nu_0} \subset Z_Q^{\nu_1}$ and $i_Q : X_Q^{\nu_0} \subset X_Q^{\nu_1}$. We leave it to the reader to check that the inclusion preserves the almost complex structure. □

Next we define cut spaces and the broken manifold as manifolds with cylindrical structures. For both these objects, this is a second definition, the first being Definitions 3.5 and 3.14, where they were defined as symplectic manifolds. See Remark 3.35 for a reconciliation of the two viewpoints.

DEFINITION 3.32. (Cut spaces and broken manifold) Let $(X, \mathcal{P}, \Phi)$ be a tropical manifold and let $\{X^\nu\}_\nu$ be the corresponding family of neck-stretched manifolds.

(a) (Cut space) For a polytope $P \in \mathcal{P}$, the *cut space* $X_P^\square$ is defined as the exhaustion

$$X_P^\square := \cup_\nu X_P^\nu$$

where X_P^ν is the T_P-quotient of the P-cylindrical subset of X^ν, see (3.28).

(b) (T_P-bundle on the cut space X_P) Let

$$Z_P^\square := \cup_\nu Z_P^\nu$$

be the T_P-bundle on $X_P^\square$ which is the exhaustion of the T_P-bundles $Z_P^\nu \to X_P^\nu$ defined in (3.29).

(c) (Broken manifold) The broken manifold $\mathfrak{X}$ is the disjoint union

$$\mathfrak{X}_{\mathcal{P}} \quad \text{or} \quad \mathfrak{X} := \bigsqcup_{P \in \mathcal{P}} X_{\overline{P}}^\square, \tag{3.30}$$

where we use the notation $\mathfrak{X}$ when the polyhedral decomposition $\mathcal{P}$ is obvious from the context. Here,

$$X_{\overline{P}}^\square := Z_P^\square \times \mathfrak{t}_P^\vee,$$

and it is a $T_{P,\mathbb{C}}$-bundle over $X_P^\square$ with projection

$$T_{P,\mathbb{C}} \to X_{\overline{P}}^\square \xrightarrow{\pi_{X_P}} X_P^\square. \tag{3.31}$$

(d) (Cylindrical almost complex structures) An almost complex structure $\mathfrak{J} = (J_{\overline{P}})_{P \in \mathcal{P}}$ on $\mathfrak{X}$ consists of an almost complex structure $J_{\overline{P}}$ on each component $X_{\overline{P}}^\square \subset \mathfrak{X}$. Such a $\mathfrak{J}$ is *cylindrical* if it is the limit of a family $(J^\nu)_\nu$ of neck-stretched cylindrical almost complex structures on X^ν. The almost complex structure $\mathfrak{J}$ is *locally tame* or *locally strongly tamed* if the corresponding property holds for J^ν.

(e) (Cylindrical ends) For a pair of polytopes $Q \subset P$, the *Q-cylindrical end of* $X^\square_P$ is a subset $U_Q(X^\square_P) \subset X^\square_P$ defined as the exhaustion

$$U_Q(X^\square_P) := \cup_\nu U_Q(X^\nu_{\overline{P}})$$

of Q-cylindrical subsets $U_Q(X^\nu_{\overline{P}}) \subset X^\nu_{\overline{P}}$ given by

$$U_Q(X^\nu_{\overline{P}}) := \cup_{R \subseteq Q} \Phi^{-1}(R^\blacksquare) \times \nu R^\vee \subset X^\nu_{\overline{P}}.$$

The Q-cylindrical end in $X^\square_{\overline{P}}$ resp. $Z^\square_{\overline{P}}$ is the lift of $U_Q(X_P)$ by the projection map $\pi''_P : Z^\square_P \to X^\square_P$ resp. $\pi_{X_P} : X^\square_{\overline{P}} \to X^\square_P$, and is denoted by

$$U_Q(X^\square_{\overline{P}}) \subset X^\square_{\overline{P}} \quad \text{resp.} \quad U_Q(Z^\square_P) \subset Z^\square_P.$$

Given a broken manifold with a cylindrical almost structure, any piece $X^\square_{\overline{P}}$ of the broken manifold has a compactification given by an almost complex manifold

- which is diffeomorphic to the broken symplectic manifold $X_{\overline{P}}$,
- and the complement $X_{\overline{P}} \backslash X^\square_{\overline{P}}$ is the union of the relative divisors Y_Q, which are almost complex submanifolds.

REMARK 3.33. The glued manifolds X^ν can be recovered from the broken almost complex manifold equipped with cylindrical coordinates on its ends as

$$X^\nu = \left(\bigcup_{P \in \mathcal{P}, \operatorname{codim}(P)=0} X^\square_{\overline{P}} \right) / \sim_\nu$$

where the equivalence relation $\sim_\nu$ is as follows. For top-dimensional polytopes P_0, P_1, and $Q := P_0 \cap P_1$,

$$X^\square_{P_0} \ni x_0 \sim_\nu x_1 \in X^\square_{P_1} \iff x_0 \in U_{P_0}Q, \;\; x_1 \in U_{P_1}Q, \;\; i_Q^{P_0}(x_0) = e^{\pi(\nu P_0^\vee - \nu P_1^\vee)} i_Q^{P_1}(x_1). \tag{3.32}$$

Here $P_0^\vee - P_1^\vee \in \mathfrak{t}_Q^\vee$ and π is the identification $\mathfrak{t}_Q^\vee \simeq \iota \mathfrak{t}_Q$ from (3.6). This ends the Remark.

EXAMPLE 3.34. (Cylindrical ends in a single cut) We continue Example 3.27 where we described neck-stretched manifolds in case of a single cut. Let (X, Φ, c) be a tropical manifold with a single cut. We describe the P_0-cylindrical end in X_{P_+}, the case of X_{P_-} being analogous. Firstly note that since $\operatorname{codim}(P_+) = 0$, the spaces $X^\square_{P_+}$, $Z^\square_{P_+}$ and $X^\square_{\overline{P}_+}$ are all the same, and are given by

$$X^\square_{P_+} = \{\Phi \geq c + \epsilon\} \cup_\phi ((-\infty, 0] \times Z), \tag{3.33}$$

where

- $Z := Z^\square_{P_0} = \Phi^{-1}(c)$,
- the identification $\Phi^{-1}(c+\epsilon)$ with $\Phi^{-1}(c) \times \{0\}$ is via the map given by the symplectic cylindrical structure, see (3.25).

Note that the normal cone $\mathrm{NCone}_{P_+^\vee} P_0^\vee$ is $(-\infty, 0]$. The subset $(-\infty, 0] \times Z$ in $X_{P_+}^\square$ is the P_0-cylindrical end.

Remark 3.35. (Symplectic versus almost complex broken manifolds) We have defined the broken manifold as a collection of non-compact almost complex manifolds $(X_P^\square, J_P)$, and as compact symplectic manifolds (X_P, ω_{X_P}). The manifold $X_P^\square$ is diffeomorphic to the complement of the relative divisors in X_P, however there is no canonical embedding ϕ that respects the cylindrical structure and for which $\phi^*\omega_{X_P}$ is taming for the cylindrical almost complex structures on $X_P^\square$. In general, one may not have any embedding satisfying these properties. In the complement of the cylindrical ends, there is a natural embedding

$$i_P : X_P^\square \backslash (\cup_{Q \subset P} U_Q(X_P^\square)) \to (X_P, \omega_{X_P}),$$

and $i_P^* \omega_{X_P}$ is a symplectic form on the complement of $X_P^\square \backslash \cup_{Q \subset P} U_Q(X_P^\square)$.

The problem of not having taming embeddings of broken almost complex manifolds into compact symplectic manifolds can be remedied by introducing a slight weakening in the definition of cylindrical almost complex structures on broken manifolds. In particular, if we allow the connection one-forms $(\alpha_P)_{P \in \mathcal{P}}$ and the $\mathfrak{t}$-inner product to be different across the set of cut spaces $\{X_P\}_{P \in \mathcal{P}^0}$, one can construct taming diffeomorphisms of $X_P^\square$ into X_P, see Lemma 5.12. But the cost is that the new kind of almost complex structures on pieces of $\mathfrak{X}$ do not glue to give a neck-stretched almost complex structure in X, that is, they are not *gluable*.

Gluability of almost complex structures on broken manifolds is necessary for the proof of homotopy equivalence of the Fukaya algebras defined on the unbroken manifold (X, L) and the broken manifold $(\mathfrak{X}, L)$. Once this result is proved, one has greater flexibility in choosing almost complex structures on broken manifolds while keeping the compactness and regularity results. In particular, the requirement of gluability can be dropped and one can choose almost complex structures so that there are taming embeddings into compact symplectic manifolds. Domain-dependent almost complex structures that are not gluable are described in Section 10.7. This finishes the Remark.

3.5.2. Coordinates on cylindrical ends. Our next task is to produce identifications between Q-cylindrical ends of different components of a broken manifold. These identifications are used in writing down the matching conditions at nodes of a broken map, as the matching condition compares the evaluation of maps lying in different manifolds (say $X_{\overline{P}}$, $X_{\overline{P}'}$), though neighborhoods of both lifts of the node map to the Q-cylindrical end. The identifications between the cylindrical ends are via certain natural cylindrical 'coordinates' that take values in cones of torus Lie algebras such as $\mathfrak{t}_P^\vee$. We first define the cones.

Definition 3.36. (Cones at faces of polytopes) Let P be a simple polytope in a vector space $\mathfrak{t}$. For a face $Q \subset P$, the *cone* of P at Q is

$$\mathrm{Cone}_Q(P) := \{\lambda(p - q) : p \in P, q \in Q, \lambda \in \mathbb{R}_{\geq 0}\} \subset \mathfrak{t}. \tag{3.34}$$

For any interior point q in Q, let $\mathfrak{t}_Q := T_q Q$ be the tangent space. The *normal cone* of P at Q is

$$\operatorname{NCone}_Q(P) := \operatorname{Cone}_Q(P)/\mathfrak{t}_Q \subset \mathfrak{t}/\mathfrak{t}_Q, \tag{3.35}$$

which is the image of $\operatorname{Cone}_Q(P)$ under the projection $\mathfrak{t} \to \mathfrak{t}/\mathfrak{t}_Q$. See Figure 3.8.

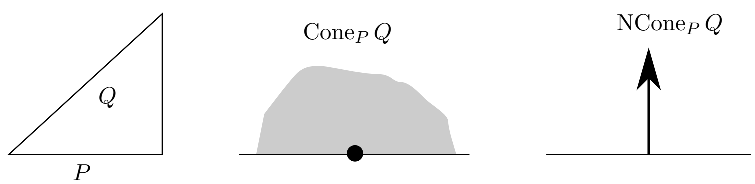

FIGURE 3.8. The cone and normal cone of Q at a face P.

REMARK 3.37. The cones defined above arise naturally in the description of broken manifolds because they are exhaustions of polytopes underlying neck-stretched manifolds. For a pair $Q \subset P$ of polytopes with $\operatorname{codim}(P) = 0$, there are inclusions

$$i_{P,Q} : \nu_0 Q^\vee \to \nu_1 Q^\vee, \quad \forall \nu_0 < \nu_1$$

that are translations and which map the vertex $P^\vee \in \nu_0 Q^\vee$ to $P^\vee \in \nu_1 Q^\vee$. Then $\operatorname{Cone}_{P^\vee} Q^\vee$ is the exhaustion

$$\operatorname{Cone}_{P^\vee} Q^\vee = \cup_\nu \nu Q^\vee \tag{3.36}$$

with respect to the inclusion $i_{P,Q}$. In general for any pair $Q \subset P$ (with $\operatorname{codim}(P) > 0$ possibly) there are inclusions

$$i_{P,Q} : \nu_0 Q^\vee/\mathfrak{t}_P^\vee \to \nu_1 Q^\vee/\mathfrak{t}_P^\vee, \quad \forall \nu_0 < \nu_1$$

which take the point $\nu_0 P^\vee/\mathfrak{t}_P^\vee$ in the domain to the point $\nu_1 P^\vee/\mathfrak{t}_P^\vee$ in the target space. The resulting exhaustion is the normal cone

$$\operatorname{NCone}_{P^\vee} Q^\vee = \cup_\nu \nu Q^\vee/\mathfrak{t}_P^\vee. \tag{3.37}$$

The following lemma gives an identification between the Q-cylindrical end in a component $X_{\overline{P}}^{\square}$ of the broken manifold $\mathfrak{X}$ to the 'standard' Q-cylinder $Z_Q^{\square} \times \mathfrak{t}_Q^\vee$ (which is the same as $X_{\overline{Q}}^{\square}$ by definition).

LEMMA 3.38. (Cylindrical ends on broken manifolds) *For any pair of polytopes $Q \subset P$ in $\mathcal{P}$, there are natural embeddings*

$$\begin{aligned} i_Q^P &: U_Q(X_P^{\square}) \to (Z_Q^{\square}/T_P) \times \operatorname{NCone}_{P^\vee} Q^\vee, \\ i_Q^{Z,P} &: U_Q(Z_P^{\square}) \to Z_Q^{\square} \times \operatorname{NCone}_{P^\vee} Q^\vee, \\ i_Q^{\overline{P}} &: U_Q(X_{\overline{P}}^{\square}) \to Z_Q^{\square} \times (\operatorname{NCone}_{P^\vee} Q^\vee \times \mathfrak{t}_P^\vee) \subset Z_Q^{\square} \times \mathfrak{t}_Q^\vee \simeq X_{\overline{Q}}^{\square}. \end{aligned} \tag{3.38}$$

PROOF. We prove the Lemma for the case $\operatorname{codim}(P) = 0$. In this case, the torus T_P is trivial, $X_P^\square = Z_P^\square = X_{\overline{P}}^\square$, and $\mathrm{NCone}_{P^\vee} Q^\vee = \mathrm{Cone}_{P^\vee} Q^\vee$. The domain of i_Q^P has a decomposition (into sets whose interiors are disjoint)

$$U_Q(X_P) = \cup_{R \subseteq Q} (\Phi^{-1}(R^\blacksquare) \times \mathrm{Cone}_{P^\vee} R^\vee) \tag{3.39}$$

where we use the fact that $\mathrm{Cone}_{P^\vee} R^\vee$ is the limit of the polytopes $\nu R^\vee$, see (3.36). The target space of i_Q^P has a decomposition (into sets whose interiors are disjoint)

$$Z_Q^\square = \Phi^{-1}(Q^\blacksquare) \cup (\cup_{R \subset Q} \Phi^{-1}(R^\blacksquare) \times \mathrm{NCone}_{Q^\vee} R^\vee) \tag{3.40}$$

which is a limit of the decompositions

$$Z_Q^\nu = \Phi^{-1}(Q^\blacksquare) \cup (\cup_{R \subset Q} \Phi^{-1}(R^\blacksquare) \times \nu R^\vee / \mathfrak{t}_Q^\vee),$$

since $\mathrm{NCone}_{Q^\vee} R^\vee$ is the limit of $\nu R^\vee / \mathfrak{t}_Q^\vee$ as $\nu \to \infty$ (see (3.37)). The map i_Q^P maps a subset in the domain decomposition (3.39) to the corresponding subset in the target decomposition (3.40). Here Z_Q^ν is a T_Q-bundle defined in (3.29). For any $R \subseteq Q$, the map

$$\mathrm{i}_Q^P : \Phi^{-1}(R^\blacksquare) \times \mathrm{Cone}_{P^\vee} R^\vee \to (\Phi^{-1}(R^\blacksquare) \times \mathrm{NCone}_{Q^\vee} R^\vee) \times \mathrm{Cone}_{P^\vee} Q^\vee$$

is defined in an obvious way: it is the identity on $\Phi^{-1}(R^\blacksquare)$ and $\mathrm{Cone}_{P^\vee} R^\vee \to \mathrm{NCone}_{Q^\vee} R^\vee \times \mathrm{Cone}_{P^\vee} Q^\vee$ is an orthogonal splitting. We leave to the reader the construction of the maps i_Q^P, $i_Q^{Z,P}$ when $\operatorname{codim}(P) > 0$. Finally we point out that the domain resp. target space of the map $i_Q^{\overline{P}}$ is the product of $\mathfrak{t}_P^\vee$ and the domain resp. target space of $i_Q^{Z,P}$. Therefore the map $i_Q^{\overline{P}}$ is defined in an obvious way once $i_Q^{Z,P}$ is known. □

Neck-stretched manifolds and components of a broken manifold are equipped with a cylindrical metric. We recall that the inner product (3.6) on $\mathfrak{t}$ defines a metric $|\cdot|_{\mathfrak{t}_P}$ resp. $|\cdot|_{\mathfrak{t}_P^\vee}$ on $\mathfrak{t}_P^\vee$ resp. $\mathfrak{t}_P^\vee$ for all $P \in \mathcal{P}$.

DEFINITION 3.39. (Cylindrical metric) A metric g_P on $X_{\overline{P}}^\square \simeq Z_P \times \mathfrak{t}_P^\vee$ is P-*cylindrical* if g_P is a product metric, that is, the product of the linear metric $|\cdot|_{\mathfrak{t}_P^\vee}$ on $\mathfrak{t}_P^\vee$ and a T_P-invariant metric g_{Z_P} on Z_P that satisfies

$$|\xi_{Z_P}|_{g_{Z_P}} = |\xi|_{\mathfrak{t}_P} \quad \xi \in \mathfrak{t}_P$$

where ξ_{Z_P} is the generating vector field as in (3.1). On the multiply-stretched manifolds X^ν, a metric g_ν is *cylindrical* if for any $P \in \mathcal{P}$, g_ν is P-cylindrical in the region $\Phi^{-1}(P^\blacksquare) \times \nu P^\vee$.

3.6. Translations: Relating neck-stretched and broken manifolds

To examine the convergence behavior of maps in neck-stretched manifolds to a limit map in the broken manifold, we need to embed P-cylindrical regions of the neck-stretched manifold into the P-cylindrical component of the broken manifold. Before defining these embeddings we recall some facts about neck-stretched and broken manifolds :

(a) For any polytope $P \in \mathcal{P}$ and any ν, the P-cylindrical subset of X^ν is $X^\nu_{\overline{P}}$ which has a projection

$$X^\nu_{\overline{P}} \xrightarrow{\pi'_{P,\nu}} Z^\nu_P$$

where Z^ν_P is a T_P-bundle over X^ν_P and the fibers of $\pi'_{P,\nu}$ are subsets of $\nu P^\vee$. In particular there is a map

$$\pi^\nu_{P^\vee} : X^\nu_{\overline{P}} \to \nu P^\vee,$$

which, for any $Q \subseteq P$, is defined by orthogonal projection $Q^\vee \to P^\vee$ on the cylindrical ends $\Phi^{-1}(Q^\circ) \times \nu Q^\vee \subset X^\nu_{\overline{P}}$.

(b) The P-cylindrical component $X^\square_{\overline{P}} \subset \mathfrak{X}$ is a product

$$X^\square_{\overline{P}} = Z^\square_P \times \mathfrak{t}^\vee_P$$

where $Z^\square_{\overline{P}} \to X^\square_P$ is a T_P-bundle over $X^\square_P$.

(c) There are natural inclusions $Z^\nu_P \to Z^\square_P$, $X^\nu_P \to X^\square_P$.

Definition 3.40. (P-translation) For any $P \in \mathcal{P}$, $t \in \nu P^\vee$, denote by

$$\mathbb{e}^{-t}_P : X^\nu_{\overline{P}} \to X^\square_{\overline{P}}. \tag{3.41}$$

the lift of the inclusion $Z^\nu_P \to Z^\square_P$ that maps a level set $\{\pi_{\nu P^\vee} = c\} \subset X^\nu_{\overline{P}}$ to $Z^\square_P \times \{c - t\} \subset X^\square_{\overline{P}}$.

For any $t \in \nu P^\vee$ the translation $\mathbb{e}^{-t}_Q$ is the 'same' for all $Q \subseteq P$ because

$$\mathbb{e}^{-t}_P | X^\nu_{\overline{Q}} = \mathbb{e}^{-t}_Q.$$

Indeed, the restriction of $\mathbb{e}^{-t}_P$ to $X^\nu_{\overline{Q}} \subset X^\nu_{\overline{P}}$ maps to the Q-cylindrical end $U_Q(X^\square_{\overline{P}}) \subset X^\square_{\overline{P}}$, which may be viewed as a subset of $X^\square_{\overline{Q}}$. This leads us to view the parameter t in the translation $\mathbb{e}^{-t}$ as an element in the dual complex $B^\vee$ and we get the following notion:

Definition 3.41. (Generalized translation) For any $t \in \nu B^\vee$,

$$\mathbb{e}^{-t} := \mathbb{e}^{-t}_P \quad \text{if } t \in P^{\vee,\circ}.$$

For an element $t \in \nu P^{\vee,\circ}$ the inverse of the translation $\mathbb{e}^{-t}$ is well-defined on a subset $X^\square_{\overline{P},\nu} \subset X^\square_{\overline{P}}$:

$$\mathbb{e}^t := (\mathbb{e}^{-t})^{-1} : X^\square_{\overline{P}} \supset X^\square_{\overline{P},\nu} \to X^\nu_{\overline{P}}. \tag{3.42}$$

The sequence of subsets $X^\square_{\overline{P},\nu}$ exhaust $X^\square_{\overline{P}}$ as $\nu \to \infty$.

Example 3.42. We illustrate translations in a multiple cut using an example with two non-intersecting single cuts. Consider the tropical manifold $(X, \Phi, \mathcal{P})$ where the torus is $T = S^1$, and the polytopes in $\mathcal{P}$ are P_0, P_1, P_2, P_{01}, $P_{12} \subset \mathbb{R}$ shown in Figure 3.9. The dual complex $B^\vee$ is a subset of $\mathbb{R}$. Let the point $P_i^\vee$ be $g_i \in \mathbb{R}$ for $i = 0, 1, 2$. For $\nu > 0$, the neck-stretched manifold is

$$X^\nu := X^\blacksquare_{P_0} \cup ([\nu g_0, \nu g_1] \times Z_0) \cup X^\blacksquare_{P_1} \cup ([\nu g_1, \nu g_2] \times Z_1) \cup X^\blacksquare_{P_2} / \sim,$$

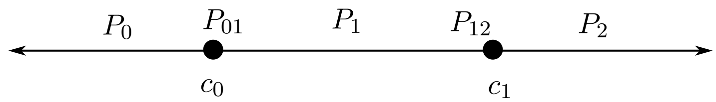

FIGURE 3.9. A polyhedral decomposition $\mathcal{P}$ and its dual complex $B^\vee$.

where $Z_i := \Phi^{-1}(c_i)$, $X^{\blacksquare}_{P_i}$ is X_{P_i} minus a tubular neighbourhood of relative divisors, and $\sim$ identifies the copies of Z_0 and Z_1 on the boundaries. Consider a translation $t \in \nu B^\vee$.

- Suppose $t \in \nu P_1^\vee$. Then $t = \nu g_1$ and $\mathbb{e}^{-t}$ is the embedding

$$([\nu g_0, \nu g_1] \times Z_0) \cup_{Z_0} X^{\blacksquare}_{P_1} \cup_{Z_1} ([\nu g_1, \nu g_2] \times Z_1) \\ \to X^{\square}_{P_1} \simeq ((-\infty, 0] \times Z_0) \cup_{Z_0} X^{\blacksquare}_{P_1} \cup_{Z_1} ([0, \infty) \times Z_1),$$

that is identity on $X^{\blacksquare}_{P_1}$, and on the cylindrical ends it is a translation by νg_1, namely

$$([\nu g_0, \nu g_1] \times Z_0) \xrightarrow{((\cdot - \nu g_1), \mathrm{Id})} (-\infty, 0] \times Z_0, \quad ([\nu g_1, \nu g_2] \times Z_1) \xrightarrow{((\cdot - \nu g_1), \mathrm{Id})} ([0, \infty) \times Z_1).$$

The translation $\mathbb{e}^{-t}$ is defined similarly when $t \in \nu P_0^\vee$ or $\nu P_2^\vee$.

- Suppose $t \in P_{i(i+1)}^\vee$. Then $t \in [\nu g_i, \nu g_{i+1}]$ and

$$\mathbb{e}^{-t} : [\nu g_i, \nu g_{i+1}] \times Z_i \to \mathbb{R} \times Z_i$$

maps $\{c\} \times Z_i$ to $\{c - t\} \times Z_i$.

This ends the example.

3.7. Existence of symplectic cylindrical structures

We prove that tropical Hamiltonian manifolds possess symplectic cylindrical structures, which were used in the definition of neck-stretched almost complex manifolds.

LEMMA 3.43. (Existence of symplectic cylindrical structures) *There exists a symplectic cylindrical structure (see Definition 3.21) for a tropical Hamiltonian manifold $(X, \mathcal{P}, \Phi)$.*

The symplectic cylindrical structure is constructed via an analogous structure on cut spaces, which we define next. We recall that for a pair of polytopes $Q \subset P$ a neighborhood $U_{X_Q} X_P$ of X_Q in X_P is mapped by the tropical moment map Φ to a neighborhood of Q in $\mathrm{Cone}_Q P \subset \mathfrak{t}^\vee$. Thus a model symplectic neighborhood of X_Q in X_P is given by $Z_Q \times \mathrm{Cone}_Q P$ with boundaries modded out by circle actions as in Lerman's construction.

DEFINITION 3.44. (Symplectic structure on cut spaces) A *symplectic structure on cut spaces* consists of a collection of maps $(\phi_P^Q)_{Q \subset P}$ where $P \in \mathcal{P}$ ranges over top-dimensional polytopes. For any pair $Q \subset P$, the map ϕ_P^Q is a T_Q-equivariant symplectomorphic embedding defined on a neighborhood $U_{X_Q} X_P \subset X_P$ of X_Q

$$\phi_P^Q : U_{X_Q} X_P \to (\mathrm{Cone}_Q P \times Z_Q, \omega_{\overline{Q}})/\sim, \quad \omega_{\overline{Q}} := \omega_{X_Q} + d\langle \alpha_Q, \pi_{\mathfrak{t}_Q^\vee} \rangle, \tag{3.43}$$

where

(a) the equivalence $\sim$ mods out the boundary $\mathrm{Cone}_Q R \subset \mathrm{Cone}_Q P$ by the action of $T_R \simeq S^1$, and $R \subset P$ ranges over all facets of P for which $Q \subseteq R$; and thus by Lerman's construction $(\mathrm{Cone}_Q P \times Z_Q, \omega_{\overline{Q}})/\sim$ is a manifold resp. orbifold if $\mathrm{Cone}_Q P$ is Delzant resp. simple;

(b) X_Q is identically mapped by ϕ_P^Q to $Q \times Z_Q/T_Q$;

(c) $\alpha_Q \in \Omega^1(Z_Q, \mathfrak{t}_Q)$ is a connection one-form such that the collection $(\alpha_P)_{P\in\mathcal{P}}$ is consistent in the sense of (3.16), $\pi_{\mathfrak{t}_Q^\vee} : \mathfrak{t}^\vee \to \mathfrak{t}_Q^\vee$ is the projection on $\mathrm{Cone}_Q P \subset \mathfrak{t}^\vee$, and so, $\langle \alpha_Q, \pi_{\mathfrak{t}_Q^\vee}\rangle$ is a one-form on $(Z_Q \times \mathrm{Cone}_P Q)$.

REMARK 3.45. Any symplectic cylindrical structure induces projection maps on neighborhoods of boundary submanifolds

$$\pi_Q^\omega : U_{X_Q} X_P \to X_Q \tag{3.44}$$

for all pairs of polytopes $Q \subset P$.

PROOF OF LEMMA 3.43. We first construct a symplectic cylindrical structure on cut spaces as in Definition 3.44. This data includes a consistent collection of connection one-forms $(\alpha_P)_{P\in\mathcal{P}}$ and symplectomorphisms $(\phi_P^Q)_{Q\subset P}$ in (3.43). For a fixed top-dimensional polytope P, the maps ϕ_P^Q are constructed by induction on the dimension of Q. At every step of the induction the connection one-form α_{Q_0} on $Z_{Q_0} \to X_{Q_0}$ is pre-determined in a neighborhood $U_{X_Q} X_{Q_0}$ of X_Q for strict subsets $Q \subsetneq Q_0$. We may choose any extension of $\alpha_{Q_0}|(\cup_Q U_{X_Q} X_{Q_0})$ to all of X_{Q_0}. The symplectomorphism $\phi_P^{Q_0}$ is also pre-determined in the neighbourhoods $U_{X_Q} X_{Q_0}$, and is extended to all of $U_{X_P} X_{Q_0}$ by the relative symplectic neighborhood theorem (Lemma 3.46).

A symplectic cylindrical structure on X is obtained by gluing the structure maps on cut spaces. Indeed, for a polytope $Q \in \mathcal{P}$, $\mathrm{codim}(Q) > 0$, we can invert the multiple cut operation on the domain and target spaces of the maps

$$\phi_P^Q : U_{X_Q} X_P \to (\mathrm{Cone}_{P^\vee}(Q^\vee) \times Z_Q, \omega_{\overline{Q}})/\sim, \quad P \supset Q, \ \mathrm{codim}(P) = 0, \tag{3.45}$$

to yield the symplectomorphism ϕ_Q of (3.13), because the connection one-form α_Q underlying the symplectic form is the same in each of the target spaces in (3.45). $\square$

LEMMA 3.46. (Relative symplectic neighbourhood theorem) *Let $Y \subset X$ be a compact symplectic submanifold of a symplectic manifold (X, ω). Let $\mathcal{N}$ be a neighborhood of the zero section in the normal bundle $N_X Y$ that is equipped with a symplectic form $\omega_{\mathcal{N}}$. Let $\psi : \mathcal{N} \to \mathcal{U}$ be a diffeomorphism onto a neighborhood $\mathcal{U} \subset X$ of Y that is identity on Y. Further, let $Y \subset S \subset \mathcal{N}$ be a subset satisfying*

$$s \in S \implies ts \in S \ \forall t \in [0,1], \quad \textit{and} \quad (\omega_{\mathcal{N}} - \psi^*\omega)|_{T_s\mathcal{N}} = 0.$$

Then, there is a smaller neighborhood $\mathcal{N}' \subset \mathcal{N}$ of the zero section that contains S and so that ψ can be homotoped to a symplectomorphism $\phi : \mathcal{N}' \to \mathcal{U}$ satisfying $\phi|_S = \psi|_S$.

The proof of the ordinary symplectic neighborhood ([**59**, Lemma 3.14]) can be used to prove the slightly stronger statement of Lemma 3.46.

CHAPTER 4

Broken disks

The goal of this chapter is to define *broken treed holomorphic disks.* These are analogues of what Parker [**68**] calls exploded holomorphic maps. These structures combine the features of treed holomorphic disks and tropical (or broken) maps.

- Treed holomorphic disks consist of surface components that are nodal holomorphic disks or spheres in a symplectic manifold whose boundary lies in a Lagrangian submanifold, with the additional feature that disk nodes may be replaced by tree segments mapping to the Lagrangian submanifold, analogous to pearly trajectories in Biran-Cornea [**9**].
- On the other hand, a broken map is defined on the normalization of a nodal curve, and each of the domain components maps to a different component of the target broken manifold, and the maps satisfy a matching condition at the nodal lifts. The broken map is equipped with the additional data of a tropical graph in the dual complex associated with the degeneration of the symplectic manifold.

In a broken treed holomorphic disk, certain nodes in the domain are 'tropical nodes' which means that the curve components incident on the node map to different target components; and the other nodes, called 'internal nodes', are sphere nodes as seen in Gromov-Witten theory or disk nodes with treed segments as in treed holomorphic disks. Since we assume that the Lagrangian submanifold is disjoint from the relative divisors of the broken manifold, all disk nodes are internal nodes.

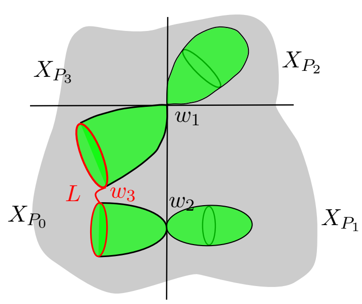

FIGURE 4.1. A broken treed holomorphic disk. The disk node w_3 has been replaced by a treed segment in the Lagrangian L, and w_1, w_2 are tropical nodes.

4.1. Treed disks

The domains of our pseudoholomorphic maps are treed disks, which are analogues of pearly trajectories of Biran-Cornea [9], Cornea-Lalonde [24] and Seidel [80]. A treed disk is a combination of trees, nodal disks, and nodal spheres.

DEFINITION 4.1. (a) (Nodal disks) A nodal disk S is a union

$$S = \left(\bigcup_{\alpha=1,\ldots,d(\circ)} S_{\alpha,\circ} \right) \cup \left(\bigcup_{\beta=1,\ldots,d(\bullet)} S_{\beta,\bullet} \right) / \sim \tag{4.1}$$

of a collection of disk components $S_{\alpha,\circ}$ each biholomorphic to a unit disk $\mathbb{D}^2 \subset \mathbb{C}$, and sphere components $S_{\beta,\bullet}$ each biholomorphic to the projective line $\mathbb{P}^1$, glued together by an equivalence relation $\sim$. The equivalence relation $\sim$ is generated by pairs of interior or boundary *nodal points*

$$w_e = (w_+(e), w_-(e)) \in (\cup_\alpha \partial S_{\alpha,\circ})^2 \cup (\cup_\alpha(\operatorname{int}(S_{\alpha,\circ}) \cup \cup_\beta S_{\beta,\bullet})^2 \tag{4.2}$$

of pairs of boundary points of disks, interior points of disks or spheres, with the property that the boundary ∂S is connected and there are no cycles of components $S_{\alpha_1}, \ldots, S_{\alpha_k} = S_{\alpha_1}$ connected by nodes (so that in particular, each node connects two different disk or sphere components of S). A *marking* of a nodal disk is a collection of boundary and interior points

$$\underline{z}_\circ = (z_{\circ,i} \in \partial S, i = 0, \ldots, d(\circ)), \quad \underline{z}_\bullet = (z_{\bullet,i} \in S \backslash \partial S, i = 1, \ldots, d(\bullet)) \tag{4.3}$$

distinct from the nodes. A marked nodal disk is *stable* if it admits no automorphisms, or equivalently, if for each disk component $S_{\circ,i}$ the sum of the number of special (nodal or marked) boundary points and twice the number of interior special points is at least three, and each sphere component $S_{\bullet,i}$ has at least three special points.

(b) (Combinatorial type of a nodal disk) The *combinatorial type* of a nodal disk S is the tree Γ whose vertices $\operatorname{Vert}(\Gamma)$ correspond to disk or sphere components of S, and whose edges $\operatorname{Edge}(\Gamma)$ correspond to markings or nodes, together with

(i) (Sphere and disk vertices) a partition of the vertex set

$$\operatorname{Vert}(\Gamma) = \operatorname{Vert}_\bullet(\Gamma) \cup \operatorname{Vert}_\circ(\Gamma)$$

into disk vertices $\operatorname{Vert}_\circ(\Gamma)$ and sphere vertices $\operatorname{Vert}_\bullet(\Gamma)$;

(ii) (Interior and boundary edges) a partition of the edge set

$$\operatorname{Edge}(\Gamma) = \operatorname{Edge}_\bullet(\Gamma) \cup \operatorname{Edge}_\circ(\Gamma)$$

into edges $\operatorname{Edge}_\circ(\Gamma)$ of boundary type and edges $\operatorname{Edge}_\bullet(\Gamma)$ of interior type;

(iii) (Leaf and non-leaf edges) a partition of the edge set $\operatorname{Edge}(\Gamma)$

$$\operatorname{Edge}(\Gamma) = \operatorname{Edge}_{-}(\Gamma) \cup \operatorname{Edge}_{\to}(\Gamma)$$

into leaf edges $\operatorname{Edge}_{\to}(\Gamma)$ which correspond to markings, and non-leaf edges $\operatorname{Edge}_{-}(\Gamma)$ which correspond to nodes (and so, leaf edges are incident on a single vertex and non-leaf edges are each incident on two vertices);

(iv) (Ordering of leaves) an ordering of the boundary leaf edges $\mathrm{Edge}_{\circ,\to}(\Gamma)$ and the interior leaf edges $\mathrm{Edge}_{\bullet,\to}(\Gamma)$ given by the corresponding ordering of markings (see (4.3));

(v) (Ribbon structure) a *ribbon structure* on Γ, which is a cyclic ordering $<_v$ on the set of boundary edges (both leaf and non-leaf)

$$\mathrm{Edge}_\circ^v(\Gamma) := \{e \in \mathrm{Edge}_\circ(\Gamma) : v \in e\}$$

incident on each boundary vertex $v \in \mathrm{Vert}_\circ(\Gamma)$ such that the induced cyclic ordering on the set $\mathrm{Edge}_{\circ,\to}(\Gamma)$ of boundary leaves corresponds to the cyclic ordering $z_{\circ,0}, \ldots, z_{\circ,d(\circ)}, z_{\circ,0}$ of boundary markings.

(vi) (Root edge and edge orientations) The edge $e_0 \in \mathrm{Edge}_\circ(\Gamma)$ corresponding to the first boundary marking z_0 is an outgoing edge and is called the *root edge*, all the other boundary markings are incoming edges, and all edges corresponding to nodes are oriented to point towards the root.

(c) (Treed segments) A *treed segment* is obtained from a collection of closed intervals $I_1, \ldots, I_k$, each isomorphic to some subset of $\mathbb{R}$ by

$$I_j \cong [0, \ell(I_j)], \quad (-\infty, 0], \quad \text{or} I_j \cong [a_j, \infty) \quad \text{or} \quad (-\infty, \infty)$$

by gluing along infinite endpoints, e.g.

$$[0,\infty) \cup_\infty (-\infty,\infty) \cup_\infty (-\infty, 0] \tag{4.4}$$

is a treed segment with three components and finite end-points. Each treed segment T has a length $\ell(T) \in [0,\infty]$ and a number of breakings $b(T) \in \mathbb{Z}_{\geq 0}$, with $\ell(T) = \infty$ if $b(T) > 0$. We also consider treed segments with one infinite end such as

$$[0,\infty) \cup_\infty (-\infty,\infty)$$

or both infinite ends such as $(-\infty,\infty) \cup_\infty (-\infty,\infty)$.

(d) (Treed disk) A *treed nodal disk* $C = S \cup T$ is

(i) either obtained from a nodal disk S_0

- by assigning a length $\ell(e) \in [0,\infty]$ to each boundary node $e \in \mathrm{Edge}_{\circ,-}(\Gamma)$, and replacing any boundary node w_e, $e \in \mathrm{Edge}_{\circ,-}(\Gamma)$ with $\ell(e) > 0$ by a treed segment T_e with finite end-points, and the treed segment T_e has length $\ell(e)$ if $\ell(e) < \infty$ or it is a broken segment as in (4.4) if $\ell(e) = \infty$;
- and each boundary marking w_e, $e \in \mathrm{Edge}_{\circ,\to}(\Gamma)$ is replaced by a (possibly broken) treed segment one of whose end-points is infinite, and thus $\ell(e) = \infty$ for all boundary markings;

(ii) or is a treed segment both whose end-points are infinite. The $-\infty$ resp. ∞ end of the segment is regarded as the input resp. output, and therefore, there is one output and one input, $d(\circ) = 1$.

A treed segment T_e, $e \in \mathrm{Edge}_\circ(\Gamma)$, containing a breaking is called a *broken edge*. A non-leaf edge $e \in \mathrm{Edge}_-(\Gamma)$ is broken if and only if $\ell(e) = \infty$.

(e) (Isomorphism of treed disks) An isomorphism of treed disks is a homeomorphism $\phi : C \to C'$ that is a biholomorphism on each sphere or disk component, length-preserving on edges, and preserves the labelling of leaves.

(f) (Combinatorial type of a treed disk) The combinatorial type of a treed disk consists of the combinatorial type Γ of the underlying nodal disk which includes the vertex and edge partitions, ordering of markings, ribbon structure and edge orientations; and in addition a partition

$$\mathrm{Edge}_{\circ,-}(\Gamma) = \mathrm{Edge}^{0}_{\circ,-}(\Gamma) \cup \mathrm{Edge}^{(0,\infty)}_{\circ,-} \cup \mathrm{Edge}^{\infty}_{\circ,-}(\Gamma)$$

of boundary edges corresponding to boundary nodes with zero, finite, and infinite length edges.Note that for a boundary edge e with edge length $\ell(e) \in (0,\infty)$, the length of the treed segment T_e is not part of the combinatorial type.

(g) (Stable treed disk) A treed nodal disk C is *stable* if the underlying disk S is stable, the treed segment at any node $T_e, e \in \mathrm{Edge}_{\circ,-}(\Gamma)$ has at most one breaking, and treed segments at markings $T_e, e \in \mathrm{Edge}_{\to,\circ}(\Gamma)$ are unbroken. A treed disk with no surface component (that is, consisting of a single possibly broken sequence of segments) is not stable.

(h) (Disconnected types) Let Γ be a disjoint union of treed disk types $\Gamma_1, \ldots, \Gamma_k$. A treed curve of type Γ is a collection $u_1, \ldots, u_k$ of treed disks of types $\Gamma_1, \ldots, \Gamma_k$.

For integers $d(\bullet), d(\circ) \geq 0$, denote by $\overline{\mathcal{M}}_{d(\bullet),d(\circ)}$ the moduli space of isomorphism classes of stable treed disks with $d(\circ)$ incoming boundary markings, one outgoing boundary marking and $d(\bullet)$ interior markings. For each combinatorial type Γ, denote by $\mathcal{M}_\Gamma \subset \mathcal{M}_{d(\bullet),d(\circ)}$ the set of isomorphism classes of stable treed disks of type Γ. The moduli space $\mathcal{M}_{d(\bullet),d(\circ)}$ then decomposes as

$$\overline{\mathcal{M}}_{d(\bullet),d(\circ)} = \bigcup_{\Gamma} \mathcal{M}_\Gamma.$$

The top-dimensional cells in $\mathcal{M}_{d(\bullet),d(\circ)}$ correspond to strata where all boundary edges e have finite non-zero length, that is, $\ell(e) \in (0,\infty)$. The dimension of each of the top-dimensional cells is $d(\circ) + 2d(\bullet) - 2$. For the stratum with no finite edges, so containing a single disk, this follows immediately from the fact that each boundary leaf edge resp. interior leaf edge contributes 1 resp. 2 to the dimension, and the automorphism group of the disk is $PSL(2,\mathbb{R})$ which has dimension 3. The dimensions of the other top-dimensional cells can be computed by first computing the corresponding stratum of the moduli spaces of disks without trees, and then adding one for each boundary edge with finite non-zero length. The compactified moduli space $\overline{\mathcal{M}}_{d(\bullet),d(\circ)}$ is a manifold with corners where the codimension k corner stratum consists of curves containing k broken treed edge. We do not give a proof, since for our purposes it is enough to view $\overline{\mathcal{M}}_{d(\bullet),d(\circ)}$ as a cell-complex, with a manifold structure on each of the strata $\mathcal{M}_\Gamma$.

Definition 4.2. (Orientation of moduli of treed disks) We fix an orientation for moduli spaces of treed disks of type Γ, where Γ does not contain boundary edges $e \in \mathrm{Edge}_{\circ,-}(\Gamma)$ of length 0 or ∞.

(a) (Single disk) Let Γ be a treed disk type with a single disk component with ≥ 2 incoming boundary leaves and no other surface component. The moduli space $\mathcal{M}_\Gamma$ is the quotient

$$\mathrm{Conf}_{d(\bullet),d(\circ)}/PSL(2,\mathbb{R}),$$

where

$$\mathrm{Conf}_{d(\bullet),d(\circ)} \subset \{\underline{z} = (\underline{z}^\circ, \underline{z}^\bullet) \in (\partial D)^{d(\circ)+1} \times (\mathrm{int}(D))^{d(\bullet)}\}$$

is the subset of configurations where $\underline{z}$ consists of distinct points and the boundary markings are ordered counter-clockwise. Fixing $z_0^\circ = -1$, $z_1^\circ = 1$, $z_2^\circ = i$ gives a global slice S of the $PSL(2,\mathbb{R})$-action, and $\mathcal{M}_\Gamma$ is oriented via its identification to the slice S which is an open subset of $\mathbb{R}^{d(\circ)-2} \times \mathbb{C}^{d(\bullet)}$. In case Γ has < 2 incoming boundary leaves, we fix the global slice by setting $z_0^\circ = -1$, $z_0^\bullet = 0$.

(b) (Multiple disks) Let Γ be a treed disk type with no sphere components and all whose boundary edges have finite non-zero length. The orientation on $\mathcal{M}_\Gamma$ is chosen by inducting on the number of boundary edges $e \in \mathrm{Edge}_{\circ,-}(\Gamma)$. Suppose the treed disk types Γ, Γ_0, Γ' are related by the following morphisms

$$\Gamma \xleftarrow{\text{Make the edge length } \ell(e) \text{ non-zero}} \Gamma_0 \xrightarrow{\text{Collapse } e} \Gamma'.$$

Here we point out that the edge length $\ell(e)$ in Γ_0 is zero. The moduli space $\mathcal{M}_{\Gamma_0}$ is a codimension one boundary stratum of both $\overline{\mathcal{M}}_\Gamma$ and $\overline{\mathcal{M}}_{\Gamma'}$ via inclusions

$$i_{\Gamma,\Gamma_0} : \mathcal{M}_{\Gamma_0} \to \overline{\mathcal{M}}_\Gamma, \quad i_{\Gamma',\Gamma_0} : \mathcal{M}_{\Gamma_0} \to \overline{\mathcal{M}}_{\Gamma'}.$$

Assuming an orientation on $\mathcal{M}'_\Gamma$, $\mathcal{M}_\Gamma$ is oriented so that the orientations on $\mathcal{M}_{\Gamma_0}$ induced by i_{Γ,Γ_0}, i_{Γ',Γ_0} are the opposite of each other.

(c) (Disks and spheres) Let Γ be a treed disk type, and let Γ' be the type obtained by collapsing interior edges $e \in \mathrm{Edge}_{\bullet,-}(\Gamma)$. Then the orientation on $\mathcal{M}_{\Gamma'}$ induces an orientation on $\mathcal{M}_\Gamma \subset \overline{\mathcal{M}}_{\Gamma'}$ which is equal to the product orientation $\mathcal{M}_{\Gamma_\circ} \times \prod_{v \in \mathrm{Vert}_\bullet(\Gamma)} \mathcal{M}_{\overline{\Gamma}_v}$. Here $\Gamma_\circ \subset \Gamma$ is the subgraph consisting of all the disk vertices and Γ_v is a graph with a single vertex $\{v\}$ and markings corresponding to all edges $e \in \mathrm{Edge}_\bullet(\Gamma)$ incident on v.

Remark 4.3. Let Γ be a treed disk type containing a single edge $e \in \mathrm{Edge}_{\circ,-}(\Gamma)$ of infinite length, and all whose other boundary edges $e' \in \mathrm{Edge}_{\circ,-}(\Gamma)$ have $\ell(e') \in (0,\infty)$. Let Γ' be a treed disk type obtained from Γ by making the edge length of e finite. Then,

$$i_{\Gamma,\Gamma'} : \mathcal{M}_\Gamma \hookrightarrow \overline{\mathcal{M}}_{\Gamma'}$$

is a boundary stratum of codimension one. Suppose Γ_1, Γ_2 are treed disk types obtained by cutting the edge e in Γ (see Definition 6.3 of (Cutting edges) morphism), and suppose the root of Γ is contained in Γ_1. Then the boundary orientation on $\mathcal{M}_\Gamma$ induced by $i_{\Gamma,\Gamma'}$ differs from the product orientation $\mathcal{M}_{\Gamma_1} \times \mathcal{M}_{\Gamma_2}$ by

$$\epsilon(\Gamma_1, \Gamma_2) := (-1)^\circ \tag{4.5}$$

where $\circ$ depends only on the combinatorial types Γ_1, Γ_2, see [**82**, (12.22)].

The moduli spaces admit universal curves, which admit partitions into one and two-dimensional parts. For any combinatorial type Γ, let $\mathcal{U}_\Gamma$ denote the universal treed disk consisting of isomorphism classes of pairs (C, z) where C is a treed disk of type Γ and z is a point in C, possibly on a disk component, sphere component, or one of the edges of the tree. The map

$$\mathcal{U}_\Gamma \to \mathcal{M}_\Gamma, \quad [C, z] \mapsto [C] \tag{4.6}$$

is the universal projection, whose fiber over $[C]$ is a copy of C. The union over types Γ' with $\mathcal{M}_{\Gamma'} \subset \overline{\mathcal{M}}_\Gamma$ is denoted $\overline{\mathcal{U}}_\Gamma$. Denote by

$$\mathcal{S}_\Gamma, \quad \text{resp.} \quad \mathcal{T}_\Gamma$$

the locus of points $[C, z] \in \mathcal{U}_\Gamma$ where z lies on a disk or sphere component resp. an edge of C. Hence $\mathcal{U}_\Gamma = \mathcal{S}_\Gamma \cup \mathcal{T}_\Gamma$, and $\mathcal{S}_\Gamma \cap \mathcal{T}_\Gamma$ is the set of points on the boundary of the disks meeting the edges of the tree. Denote by

$$\overline{\mathcal{S}}_\Gamma, \quad \text{resp.} \quad \overline{\mathcal{T}}_\Gamma$$

be the compactification of $\mathcal{S}_\Gamma$ resp. $\mathcal{T}_\Gamma$ in $\overline{\mathcal{U}}_\Gamma$.

On each stratum, the universal curve admits local trivializations. For a stratum Γ, the spaces

$$\mathcal{S}_\Gamma \to \mathcal{M}_\Gamma, \quad \text{resp.} \quad \mathcal{T}_\Gamma \to \mathcal{M}_\Gamma$$

are smooth fibrations with 2 resp. 1-dimensional fibers and markings are sections

$$\underline{z} = (z_{i,\circ}, z_{j,\bullet}) : \mathcal{M}_\Gamma \to \mathcal{S}_\Gamma, \quad 0 \le i \le d(\circ), 1 \le j \le d(\bullet).$$

We also view breaking points on broken treed segments as sections $\mathcal{M}_\Gamma \to \mathcal{T}_\Gamma$, since on a type Γ, the treed segment T_e corresponding to any edge $e \in \text{Edge}_\circ(\Gamma)$ has a fixed number of breakings. The union $\mathcal{U}_\Gamma = \mathcal{S}_\Gamma \cup \mathcal{T}_\Gamma$ has local trivializations: For a curve $[C] \in \mathcal{M}_\Gamma$, $C = S \cup T$, and a small enough neighborhood $U_C \subset \mathcal{M}_\Gamma$ of $[C]$ there is a homeomorphism

$$\mathcal{U}_\Gamma | U_C \simeq U_C \times C, \tag{4.7}$$

which, when restricted to $\mathcal{S}_\Gamma | U_C$ resp. $\mathcal{T}_\Gamma | U_C$ is a diffeomorphism onto its image $U_C \times S$ resp. $U_C \times T$, and the markings

$$z_{i,\circ}, z_{j,\bullet} | U_C : U_C \to S$$

are constant functions, whose values we denote by $\underline{z_{U_C}} = (z_{i,\circ,U(C)}, z_{j,\bullet,U(C)})$, where $z_{i,\circ,U(C)}, z_{j,\bullet,U(C)} \in S$. The fibers of $\mathcal{S}_\Gamma \to U_C$ possess a conformal structure, and the trivialization in (4.7) induces a map

$$j : U_C \to \mathcal{J}(S) \tag{4.8}$$

such that for $[C_1] \in U_C$, $(C, j([C_1], \underline{z_{U_C}})$ is biholomorphic to $(\mathcal{S}_\Gamma, \underline{z})|_{[C_1]}$.

REMARK 4.4. The structure of a 'smooth fibration with sections' on $\mathcal{S}_\Gamma \to \mathcal{M}_\Gamma$, $\mathcal{T}_\Gamma \to \mathcal{M}_\Gamma$ breaks down on the extension to the compactification $\overline{\mathcal{M}}_\Gamma$. For example, a finite treed segment T_e in the fibers of $\mathcal{T}_\Gamma$ may be transformed to a segment of 'zero length' in the compactification $\overline{\mathcal{T}}_\Gamma$; or two disconnected components in the fibers of $\mathcal{S}_\Gamma$ may connect at a disk node in the compactification $\overline{\mathcal{S}}_\Gamma$.

4.2. Treed pseudoholomorphic disks

Treed pseudoholomorphic disks are maps from treed disks to a symplectic manifold equipped with a Lagrangian submanifold. The symplectic manifold has a tamed almost complex structure and the Lagrangian submanifold has a Morse function on it. On the two-dimensional part of the treed disk the map is pseudoholomorphic and the boundaries of the disks map to the Lagrangian submanifold. On the one-dimensional part of the domain, the map is a gradient flow line of the Morse function on the Lagrangian submanifold, whose length is the same as the length of the tree edge. Later in this chapter we will adapt the definition of treed holomorphic disks in a symplectic manifold to define broken treed holomorphic disks in a broken manifold. Later in the text the almost complex structure, the Morse function, and the metric on the Lagrangian will be given domain-dependent perturbations in order to regularize the moduli spaces of treed holomorphic (broken) disks.

We introduce the necessary notation for defining treed holomorphic disks. Let (X, ω) be a symplectic manifold and $L \subset X$ be a Lagrangian submanifold. Let J be an ω-tame almost complex structure. Let G_L be a Riemannian metric on L and let $F \in C^\infty(L, \mathbb{R})$ be a Morse function such that the pair (F, G_L) is Morse-Smale. The *gradient vector field* is defined by the condition

$$\operatorname{grad}_F \in \operatorname{Vect}(L), \quad \mathrm{d}f(\cdot) = G_L(\operatorname{grad}_F, \cdot).$$

DEFINITION 4.5. A *treed holomorphic disk* with boundary in $L \subset X$ consists of a treed disk $C = S \cup T$ and a continuous map

$$u : C \to X$$

satisfying the following conditions:

(a) The tree components T and the boundary ∂S of the surface components are mapped to the Lagrangian submanifold $L : u(T \cup \partial S) \subset L$;
(b) the map $u|S$ is a pseudoholomorphic map on the surface part: $Jd(u|S) = d(u|S) \circ j$; and
(c) the map $u|T$ is a union of gradient trajectories: $\frac{d}{ds} u|T = -\operatorname{grad}_F(u|T)$ where s is a unit velocity coordinate on T.

A holomorphic treed disk $u : C = S \cup T \to X$ is *stable* if it has finitely many automorphisms $\# \operatorname{Aut}(u) < \infty$, or equivalently

(a) each surface component $S_v \subset S$ on which the map u is constant is stable as a component of a nodal disk S (see Definition 4.1);
(b) each treed segment T_e on which the map u is constant has at most one infinite end, that is, one of the ends of T_e is an attaching point to a sphere or disk $S_v \subset S$.

Note that the case $C \cong \mathbb{R}$ equipped with a non-constant Morse trajectory $u : C \to L$ is allowed under this stability condition and corresponds to the case of a single incoming edge, that is, $d(\circ) = 1$. . The area of a J-holomorphic or sphere disk $u : S \to X$ is the symplectic area

$$\mathrm{Area}(u) = \int_S (u|_S)^* \omega.$$

4.3. Multiply-broken disks

A broken map is a map from a nodal curve to a multiply cut manifold that is discontinuous at nodes. Different components of the nodal curve map to different pieces of the multiply cut manifold, and the lifts of the nodal points satisfy an edge matching condition. The nodal points carry an additional data of intersection multiplicity with relative divisors. This data is packaged into a *tropical structure*, which is part of the combinatorial type of the broken map.

Definition 4.6. (Tropical graph) Let $B^\vee \subset \mathfrak{t}$ be the dual complex for a set of polytopes $\mathcal{P} = \{P \subset \mathfrak{t}^\vee\}$ as in Definition 3.9. A *tropical graph* is a triple $(\Gamma, P, \mathcal{T})$ consisting of

(a) a graph Γ with vertex set $\mathrm{Vert}(\Gamma)$ and edge set $\mathrm{Edge}(\Gamma) \subset \mathrm{Vert}(\Gamma) \times \mathrm{Vert}(\Gamma)$,
(b) polytope assignments

$$P : \mathrm{Vert}(\Gamma) \to \mathcal{P}, \quad P : \mathrm{Edge}(\Gamma) \to \mathcal{P} \tag{4.9}$$

such that for any edge $e = (v_+, v_-)$, $P(e) = P(v_+) \cap P(v_-)$,
(c) edge *slopes*

$$\mathcal{T}(e) \in \mathfrak{t}_{P(e),\mathbb{Z}} \backslash \{0\}, \quad \forall \ e \in \mathrm{Edge}(\Gamma)$$

which is *embeddable* in the following sense: A tropical graph $(\Gamma, P, \mathcal{T})$ is embeddable if it has an *embedding*, which is a collection of *tropical vertex positions*

$$\mathcal{T}(v) \in P(v)^\vee \subset B^\vee, \quad v \in \mathrm{Vert}(\Gamma)$$

satisfying the following *slope condition*

$$\text{(Slope condition)} \quad \mathcal{T}(v_+) - \mathcal{T}(v_-) \in \mathbb{R}_{>0} \mathcal{T}(e).$$

Here we view $\mathcal{T}(e)$ as an element in $\mathfrak{t}^\vee_{P(e)}$ via the identification $\mathfrak{t} \simeq \mathfrak{t}^\vee$ in (3.6). The set of tropical vertex position maps on a tropical graph Γ is denoted

$$\mathcal{W}(\Gamma) = \{(\mathcal{T}(v))_{v \in \mathrm{Vert}(\Gamma)} : \mathcal{T} \text{ is a tropical vertex position map on } \Gamma\}. \tag{4.10}$$

Remark 4.7. A deformation of embeddings for a tropical graph Γ is shown in Figure 4.2 in the dual polytope $B^\vee$ corresponding to two orthogonal cuts (see Figure 1.4 and Figure 1.5). By the definition of tropical graphs, moving the vertices v_3 and v_4 to the dotted positions produces a one-parameter space of embeddings. Thus the space of vertex positions is

$$\mathcal{W}(\Gamma) \simeq (0, 1).$$

This remark is continued in Remark 4.19.

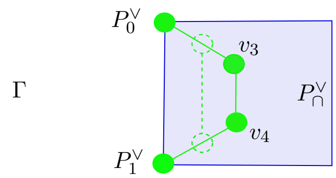

FIGURE 4.2. Moving the vertices v_3 and v_4 to the dotted positions gives a different embedding of the same tropical graph Γ.

REMARK 4.8. In a tropical graph Γ two vertices sharing an edge can not have the same tropical position. Indeed, for any edge $e = (v_+, v_-)$ the edge slope is non-zero and by the slope condition for tropical graphs, the difference in the tropical positions of v_+ and v_- is a positive multiple of the edge slope.

REMARK 4.9. For a tropical graph Γ, the set of vertex positions $\mathcal{W}(\Gamma)$ is a convex set.

DEFINITION 4.10. (Tropical structure on a treed disk type) Let Γ be the combinatorial type of a treed disk (see Definition 4.1(f)). Let $\mathfrak{X}$ be a broken manifold with an underlying polyhedral decomposition $\mathcal{P}$, and a Lagrangian submanifold L contained in the component $X_{P_0} \subset \mathfrak{X}$, $P_0 \subset \mathcal{P}$.

A *tropical structure* on Γ consists of a tropical graph $(\Gamma_{\text{tr}}, P, \mathcal{T})$ and an edge collapse morphism

$$\text{tr} : \Gamma \to \Gamma_{\text{tr}}$$

called the *tropicalization*. Here $P : \text{Vert}(\Gamma_{\text{tr}}) \cup \text{Edge}(\Gamma_{\text{tr}}) \to \mathcal{P}$ is the polytope assignment on the tropical graph (see (4.9)), and $\mathcal{T} = (\mathcal{T}(e) \in \mathfrak{t}_{P(e),\mathbb{Z}})_{e \in \text{Edge}(\Gamma_{\text{tr}})}$ is the set of edge slopes. The edges collapsed by tr are called *internal edges* and uncollapsed edges are called *tropical edges*. The subset of internal resp. tropical edges of Γ is denoted by

$$\text{Edge}_{\text{int}}(\Gamma) \quad \text{resp.} \quad \text{Edge}_{\text{trop}}(\Gamma) \subset \text{Edge}_-(\Gamma).$$

All boundary edges are collapsed. Therefore,

$$\text{Edge}_{-,\circ}(\Gamma) \subset \text{Edge}_{\text{int}}(\Gamma), \quad \text{Edge}_{\text{trop}}(\Gamma) \subset \text{Edge}_{-,\bullet}(\Gamma).$$

The map tr is often suppressed in the notation. The polygon assignment $P \circ \text{tr}$ and edge slope $\mathcal{T} \circ \text{tr}$ maps on Γ are often denoted by P, $\mathcal{T}$.

REMARK 4.11. In a treed disk type Γ equipped with a tropical structure, since the tropicalization map collapses all the boundary edges $e \in \text{Edge}_\circ(\Gamma)$, all the boundary vertices $v \in \text{Vert}_\circ(\Gamma)$ are mapped to a single vertex in Γ_{tr}.

We recall some notation required to define broken maps. The domain of a broken map is a treed disk $C = T \cup S$ (Definition 4.1) where S is the surface part and T is the tree part. For a treed disk C of type Γ equipped with a tropical structure, for every vertex $v \in \text{Vert}(\Gamma)$ we denote by

$$S_v^\circ = S_v \backslash \{w_e : e \in \text{Edge}_{\text{trop}}(\Gamma)\} \tag{4.11}$$

the surface part with cylindrical ends obtained by deleting nodal points corresponding to tropical edges. The target space of a broken map is a broken manifold $\mathfrak{X}$. We recall from (3.30) that a broken manifold with an underlying polyhedral decomposition $\mathcal{P}$ is a disjoint union

$$\mathfrak{X} = \bigsqcup_{P \in \mathcal{P}} X_{\overline{P}}^{\square},$$

and there is a top-dimensional polytope $P_0 \in \mathcal{P}$ such that $X_{\overline{P_0}}^{\square}$ contains a Lagrangian submanifold L. Note that $X_{P_0}^{\square} = X_{\overline{P_0}}^{\square}$ since the torus T_{P_0} is trivial.

DEFINITION 4.12. (Broken map) Let $\mathfrak{X}_{\mathcal{P}}$ be a broken manifold and $L \subset X_{\overline{P_0}}^{\square}$ be a Lagrangian submanifold as above. A *broken map* u to $\mathfrak{X}$ is a datum consisting of

(a) (Domain type and tropical structure) a *domain type* Γ, which is the combinatorial type of a treed disk equipped with a *tropical structure* (Definition 4.10) for which disk vertices map to P_0 :

$$v \in \mathrm{Vert}_{\circ}(\Gamma) \implies P(v) = P_0;$$

(b) (Domain curve) a treed nodal disk $C = S \cup T$ of type Γ;
(c) (Map) a collection of holomorphic maps on punctured surface components

$$u_v : S_v^{\circ} \to X_{\overline{P(v)}}^{\square}, \quad v \in \mathrm{Vert}(\Gamma),$$

(where S_v° is defined in (4.11)) and continuous maps on the treed segments

$$u_e : T_e \to L, \quad e \in \mathrm{Edge}_{\circ}(\Gamma)$$

collectively denoted $u : C \to \mathfrak{X}$, and the restriction

$$u|(\cup_{v \in \mathrm{Vert}_{\circ}(\Gamma)} S_v^{\circ} \cup \cup_{e \in \mathrm{Edge}_{\circ}(\Gamma)} T_e)$$

to the boundary part of C is a treed holomorphic map (in the sense of Definition 4.5) to the target space $(X_{\overline{P_0}}^{\square}, L)$;
(d) (Framing) and for all nodes w_e corresponding to tropical edges $e = (v_+, v_-) \in \mathrm{Edge}_{\mathrm{trop}}(\Gamma)$ a *framing*, which is a linear isomorphism

$$\mathrm{fr}_e : T_{w_+(e)} S_{P(v_+)} \otimes T_{w_-(e)} S_{P(v_-)} \to \mathbb{C}, \tag{4.12}$$

such that the collection $(u_v)_v$ has the following behaviour at punctures, and satisfies a node matching condition stated below.

- (Matching at internal nodes) For an interior internal edge

$$e = (v_+, v_-) \in \mathrm{Edge}_{\bullet}(\Gamma) \cap \mathrm{Edge}_{\mathrm{int}}(\Gamma),$$

the corresponding nodal points $w_{\pm}(e) \in C_{v_{\pm}}$ map to $X_{\overline{P(v_{\pm})}}^{\square}$ and the map u is continuous at the node :

$$u(w_+(e)) = u(w_-(e)) \in X_{\overline{P(v_{\pm})}}^{\square}. \tag{4.13}$$

- (Asymptotic behavior at punctures) For any tropical edge $e = (v_+, v_-)$ with nodal lifts $w_e^{\pm} \in C_{v_{\pm}}$ and a small enough neighborhood $U(w_e^{\pm}) \subset S_{v_{\pm}}$, the punctured neighborhood $U(w_e^\pm)\backslash\{w_e^\pm\}$ is mapped by $u_{v_\pm}$ to the $P(e)$-cylindrical end of $X^\square_{P(v_\pm)}$. Via the cylindrical coordinate map (3.38) which identifies the $P(e)$-cylindrical end to $X^\square_{\overline{P}(e)}$, we may assume that

$$u_{v_\pm}(U(w_e^\pm)\backslash\{w_e^\pm\}) \subset X^\square_{\overline{P}(e)}.$$

Recall that $X^\square_{\overline{P}(e)}$ has a holomorphic $T_{P(e),\mathbb{C}}$-action. At the puncture $w_e^\pm$ the map $u_{v_\pm}$ has the following asymptotic behavior: For any holomorphic coordinate

$$z_\pm : (U_{w_e^\pm}, w_e^\pm) \to (\mathbb{C}, 0) \tag{4.14}$$

the limit

$$x_{e,\pm} := \lim_{z_\pm \to 0} z_\pm^{\pm\mathcal{T}(e)} u_{v_\pm}(z_\pm) \tag{4.15}$$

exists in $X^\square_{\overline{P}(e)}$. The map

$$\mathbb{C}^\times \ni z_\pm \mapsto z_\pm^{\pm\mathcal{T}(e)} x_{e,\pm} \in X^\square_{\overline{P}(e)} \tag{4.16}$$

is called the *vertical cylinder* corresponding to the nodal lift $w_e^\pm$, or that the map $u_{v_\pm}$ is *asymptotic to the vertical cylinder* in (4.16) at the puncture $w_e^\pm$.

- (Matching at tropical nodes) For a tropical edge $e = (v_+, v_-) \in \mathrm{Edge}_{\mathrm{trop}}(\Gamma)$ if the holomorphic coordinates in (4.14) respect the framing

$$dz_+(w_+(e)) \otimes dz_-(w_-(e)) = \mathrm{fr}_e \tag{4.17}$$

then

$$\text{(Matching condition at nodes)} \quad \lim_{z_- \to 0} z_-^{-\mathcal{T}(e)} u_{v_-}(z_-) = \lim_{z_+ \to 0} z_+^{\mathcal{T}(e)} u_{v_+}(z_+). \tag{4.18}$$

The coordinates z_+, z_- are called *matching coordinates* at the node w_e.

A broken map does not have any extraneous components, that is,

(No extraneous components) A broken map cannot have a sphere component $u_v : C_v \to X_{P(v)}$, $v \in \mathrm{Vert}_\bullet(\Gamma)$ with no markings, two nodal points $e_1, e_2 \ni v$, and for which the projection

$$\pi_{P(v)} \circ u_v : C_v^\circ \to X^\square_{P(v)}$$

is a constant map.

This ends the Definition.

Remark 4.13. (Why disallow extraneous components?) An extraneous component in a broken map u can be deleted to yield a broken map with one less component. Indeed, an extraneous component is isomorphic to a trivial cylinder

$$u_v : C_v^\circ \simeq \mathbb{P}^1\backslash\{0,\infty\} \to X^\square_{\overline{P}(v)}, \quad z \mapsto z^\mu x,$$

where $\mu \in \mathfrak{t}_{P,\mathbb{Z}}$ is the slope of both the edges e_1, e_2 incident on v; the vertex v and edges e_1, e_2 can be deleted and replaced by a single edge e of slope μ. Conversely, if extraneous components were allowed in broken maps, any tropical edge

$e \in \mathrm{Edge}_{\mathrm{trop}}(\Gamma)$ can be subdivided to insert an extraneous vertex, and such insertions can be done an arbitrary number of times.

Definition 4.14. (Primitive slope, multiplicity of a tropical edge) In a tropical graph Γ, the slope $\mathcal{T}(e) \in \mathfrak{t}_{P(e),\mathbb{Z}} \backslash \{0\}$ of an edge can be written as a product

$$\mathcal{T}(e) = \mu_e^{\mathrm{trop}} \mathcal{T}(e)_{\mathrm{prim}}$$

of a primitive integer vector $\mathcal{T}(e)_{\mathrm{prim}} \in \mathfrak{t}_{P(e),\mathbb{Z}}$, called the *primitive slope*, and a positive integer $\mu_e^{\mathrm{trop}} \in \mathbb{Z}_{\geq 1}$ which is called the *multiplicity* of the edge e.

Definition 4.15. (Isomorphisms) An *isomorphism* between two broken maps $u : C \to \mathfrak{X}$, $u' : C' \to \mathfrak{X}$ is an isomorphism $\phi : C \to C'$ of treed disks (see 4.1(e)) such that $u = u' \circ \phi$. The group of automorphisms of a map $u : C \to \mathfrak{X}$ is denoted $\mathrm{Aut}(C, u)$.

Definition 4.16. In a broken map u a surface component
$u|S_v : S_v^\circ \to X_{P(v)}^\square$ is *horizontally constant* if its projection to $X_{P(v)}^\square$ is constant.

Definition 4.17. A broken map $u : C \to \mathfrak{X}$ is *stable* if either of the following equivalent conditions hold:

(a) Any surface component $S_v \subset C$ on which the map $u_v : S_v^\circ \to X_{P(v)}^\square$ is horizontally constant is stable as a marked curve, and any tree component $T_e \subset C$ on which $u|T_e$ is constant does not contain an infinite segment $\mathbb{R} \subset T_e$.

(b) The automorphism group $\mathrm{Aut}(C, u)$ (see Definition 4.15) is finite.

Remark 4.18. (Area of a broken map) The *area* of a surface component $u_v : S_v \to X_{P(v)}$ of a broken map is the symplectic area of its horizontal projection $\pi_{P(v)} \circ u_v : S_v \to X_{P(v)}$. That is,

$$\mathrm{Area}(u_v) = \langle (\pi_{P(v)} \circ u_v)_* [S_v], \omega_{X_{P(v)}} \rangle.$$

Note that here we use an extension of the map u_v to the compact curve S_v. The extension exists (after passing to a finite quotient in the orbifold case) by the asymptotic behavior of the map u_v at the punctures of S_v°, see (4.15).

Remark 4.19. (Tropical graph as a combinatorial invariant) The tropical graph of a broken map is a combinatorial invariant consisting of the data of edge slopes $\mathcal{T}(e)$, $e \in \mathrm{Edge}_{\mathrm{trop}}(\Gamma)$ and vertex polytopes $P(v)$, $v \in \mathrm{Vert}(\Gamma)$, since varying the tropical positions of vertices does not produce a new broken map, see Remark 4.7.

The following remarks are easy conclusions of the definition of broken maps, and are intended to help the reader process the definition.

Remark 4.20. (Continuity away from tropical nodes) Suppose the tropical structure on a broken map is given by $\Gamma \xrightarrow{\mathrm{tr}} \Gamma_{\mathrm{tr}}$ where tr is the tropicalization map and Γ_{tr} is a tropical graph. The domain of a broken map breaks up into connected components

$$C \backslash \{w_e : e \in \mathrm{Edge}_{\mathrm{trop}}(\Gamma)\} = \cup_{v \in \mathrm{Vert}(\Gamma_{\mathrm{tr}})} C_v^\circ.$$

For a vertex $v \in \mathrm{Vert}(\Gamma_{\mathrm{tr}})$ in the tropical graph, if $\mathrm{tr}^{-1}(v)$ does not have disk vertices, then C_v is a nodal sphere with punctures, with nodes corresponding to internal edges in $\mathrm{tr}^{-1}(v)$, and punctures corresponding to tropical edges incident on $\mathrm{tr}^{-1}(v)$. If $\mathrm{tr}^{-1}(v)$ has disk components then C_v° is a treed disk with punctures. For any vertex $v \in \mathrm{Vert}(\Gamma_{\mathrm{tr}})$ in the tropical graph, the restriction

$$u|C_v^\circ : C_v^\circ \to X^{\square}_{\overline{P}(v)}$$

is continuous.

Definition 4.21. (Tropical evaluation map) In the notation of Definition 4.12 of broken maps, for a node w_e and coordinates $z_\pm$ in the punctured neighborhoods of the node, the *tropical evaluation* at the nodal lift $w_e^\pm$ is

$$\mathrm{ev}^{\mathcal{T}(e)}_{w_e^\pm} u_{v_\pm} := \lim_{z_\pm \to 0} z_\pm^{\pm\mathcal{T}(e)} u_{v_\pm}(z_\pm) \in X^{\square}_{\overline{P}(e)}. \tag{4.19}$$

Thus the matching condition at a tropical node is that the tropical evaluation maps are equal for both the lifts of the node. The *projected tropical evaluation* at the nodal lift $w_e^\pm$ is

$$\pi^\perp_{\mathcal{T}(e)}(\mathrm{ev}^{\pm\mathcal{T}(e)}_{w_e^\pm} u_{v_\pm}) \in X^{\square}_{\overline{P}(e)}/T_{\mathcal{T}(e),\mathbb{C}}, \tag{4.20}$$

where $T_{\mathcal{T}(e),\mathbb{C}} \subset T_{P(e),\mathbb{C}}$ is the one-torus generated by the slope $\mathcal{T}(e) \in \mathfrak{t}_{P(e),\mathbb{Z}}$ and $\pi^\perp_{\mathcal{T}(e)} : X^{\square}_{\overline{P}(e)} \to X^{\square}_{\overline{P}(e)}/T_{\mathcal{T}(e),\mathbb{C}}$ is the projection to the quotient. The projected evaluation is independent of the choice of coordinate $z_\pm$.

Remark 4.22. (On the projected evaluation map)

(a) (Independence of domain coordinates) The projected tropical evaluation map is independent of the choice of coordinates $z_\pm$ in the neighborhood of the puncture. Indeed, altering the coordinate $z_\pm$ has the effect of multiplying $\mathrm{ev}^{\mathcal{T}(e)}_{w_e^\pm} u_{v_\pm}$ by a factor $a^{\mathcal{T}(e)}$ for some $a \in \mathbb{C}^\times$, and this factor is killed by the projection $\pi^\perp_{\mathcal{T}(e)}$.

(b) (Alternate definition) The projected tropical evaluation map can alternately defined as projection followed by evaluation : the projection $\pi^\perp_{\mathcal{T}(e)}(u_{v_\pm})$ has a bounded image and therefore the projected map extends over the singularity w_e^+, and the evaluation $\pi^\perp_{\mathcal{T}(e)}(u_{v_\pm})(w_e^\pm)$ is equal to the projected tropical evaluation map. This definition was stated in (1.5).

Definition 4.23. (Unframed broken maps) An unframed broken map $u : C \to \mathfrak{X}$ of type Γ is an object consisting of all the data in a broken map except for the framing at tropical nodes, and satisfying all the conditions in Definition 4.12 with the exception that the (Matching at tropical nodes) is replaced by the following condition:

(Unframed matching at tropical nodes) For a node w_e corresponding to $e = (v_+, v_-) \in \mathrm{Edge}_{\mathrm{trop}}(\Gamma)$ and lifts $w_e^\pm \in C_{v_\pm}$, the projected tropical evaluation maps (defined in (4.20)) are equal on w_e^+, w_e^-, that is,

$$\pi^\perp_{\mathcal{T}(e)}(\mathrm{ev}^{\mathcal{T}(e)}_{w_e^+} u_{v_+}) = \pi^\perp_{\mathcal{T}(e)}(\mathrm{ev}^{-\mathcal{T}(e)}_{w_e^-} u_{v_-}) \in X^{\square}_{\overline{P}(e)}/T_{\mathcal{T}(e),\mathbb{C}} \tag{4.21}$$

Remark 4.24. (Node matching for a single cut) In the case of a single cut, the familiar form of the node matching condition from symplectic field theory corresponds to the matching condition (4.21) for unframed maps. Indeed, for any edge e in a tropical graph, $\mathrm{codim}(P(e)) = \dim T_{P(e)} = 1$, and so, $T_{\mathcal{T}(e),\mathbb{C}} = T_{P(e),\mathbb{C}}$; the unframed matching condition (4.21) then translates to the matching of evaluations on $X^{\square}_{P(e)}$, namely, $(\pi_{P(e)} \circ u_{v+})(w_e^+) = (\pi_{P(e)} \circ u_{v-})(w_e^-)$.

Remark 4.25. (Number of framings) Any unframed broken map possesses framings, and the number of framings is given as follows: Let u be an unframed broken map of type Γ. Consider a tropical edge $e \in \mathrm{Edge}_{\mathrm{trop}}(\Gamma)$. Let $n_e \in \mathbb{Z}_{>0}$ be the largest integer such that $\frac{1}{n_e}\mathcal{T}_e \in \mathfrak{t}_{P(e),\mathbb{Z}}$. Then, the trivial cylinder $z \mapsto z^{\mathcal{T}(e)}$ is an n_e-cover. Consequently, if fr_e is a framing for the node w_e, then $e^{2\pi i k/n_e}\mathrm{fr}_e$ is also a framing for $k = 0, 1, \ldots, n_e - 1$. Thus, the total number of possible framings for u is $\Pi_{e \in \mathrm{Edge}_{\bullet,-}(\Gamma)} n_e$.

Remark 4.26. (Removing the singularity at punctures) In the special case that the compactifications of the components of the broken manifold are manifolds (and do not have orbifold singularities), a broken map extends over lifts of nodal points to yield

$$u_v : C_v \to X_{\overline{P}(v)}, \quad v \in \mathrm{Vert}(\Gamma). \tag{4.22}$$

Suppose for a nodal lift $w \in C_v$ corresponding to an edge $e = (v, v')$, $u_v(w)$ lies on the intersection of relative divisors $Y_1, \ldots, Y_k \subset X_{\overline{P}(v)}$. Further, suppose $\nu_1, \ldots, \nu_k \in \mathfrak{t}$ are the primitive outward pointing normal vectors to the facets of $\overline{P}(v)$ corresponding to the divisors $Y_1, \ldots, Y_k$, and that the map u_v intersects the divisor Y_i with multiplicity $\mu_i \in \mathbb{Z}_+$. Then we have the following relation to the slope of the edge:

$$\sum \mu_i \nu_i = \mathcal{T}(e). \tag{4.23}$$

If $X_{\overline{P}(v)}$ has orbifold singularities, that is if the polytope $P(v)$ is simple and not Delzant, the extension (4.22) is defined on a domain curve with orbifold singularities. The intersection multiplicities μ_i may be fractional, though the relation (4.23) still holds. We do not use this point of view in developing the theory. This ends the Remark.

Remark 4.27. (Exponential convergence to the vertical cylinder at tropical nodes) On neighborhoods of tropical nodes, broken maps are asymptotically converge to vertical cylinders. For a tropical node w_e corresponding to $e \in \mathrm{Edge}_{\mathrm{trop}}(\Gamma)$ choose any holomorphic coordinates

$$(s, t) : U(w_e^{\pm}) \backslash \{w_e^{\pm}\} \to \mathbb{R}_{\geq 0} \times S^1 \tag{4.24}$$

in punctured neighborhoods of lifts of nodes. Recall from the definition of a broken map that

$$u_{\mathrm{vert}} : \mathbb{R}_{\geq 0} \times S^1 \to X^{\square}_{\overline{P}(v_{\pm})}, \quad (s,t) \mapsto e^{\mp(s+it)\mathcal{T}(e)} x_e^{\pm}$$

is the vertical cylinder corresponding to the node w_e and coordinates (4.24). Here $x_e^\pm := \lim_{s\to\infty} e^{\pm(s+it)} u_{v_\pm}(s,t)$. The map $u_{v_\pm}$ exponentially converges to the vertical cylinder u_{vert} in the punctured neighborhood $U(w_e^\pm)\backslash\{w_e^\pm\}$, that is, there is a constant $c > 0$ such that

$$d(u(s,t), u_{\text{vert}}(s,t)) \leq ce^{-s}, \quad s \geq s_0,$$

using the cylindrical metric in $X^{\square}_{\overline{P}(e)}$. The proof is left to the reader.

Remark 4.28. (A different view of node matching) In the special case when the multiple cut is a collection of orthogonally intersecting single cuts, and each cut space is a manifold, the matching condition splits into

- a *horizontal matching* condition on the intersection of relative divisors,
- and a *vertical matching* condition involving leading order Taylor coefficients in the directions normal to each of the relative divisors.

We consider a multiple cut given by a tropical moment map $\Phi : X \to \mathbb{R}^k$ with cuts along the hypersurfaces $\Phi_i^{-1}(0)$, $i = 1, \ldots, k$. We denote the relative divisors by $Y_i := \Phi_i^{-1}(0)/S^1$ (though Y_i is a broken manifold). We consider a broken map u with components u_+, u_- in the cut spaces $P_+ := \cap_i\{\Phi_i \geq 0\}$ and $P_- := \cap_i\{\Phi_i \leq 0\}$ respectively, sharing a node w_e. As discussed in Remark 4.26, the map $u_\pm$ extends over the puncture at the nodal point $w_e^\pm$. The intersection multiplicities with the divisors $Y_1, \ldots, Y_k$ are the same for u_+, u_-, and are denoted by $\mu_1, \ldots, \mu_k$. The matching condition consists of

- (Horizontal matching) a horizontal condition which says that
  $$u_+(w_e^+) = u_-(w_e^-) \in X_{P(e)},$$
  where $X_{P(e)} = \Phi^{-1}(0)/(S^1)^k \simeq Y := \cap_i Y_i$;
- (Vertical matching) and a vertical condition which says that for holomorphic coordinates
  $$z_\pm : (U(w_e^\pm), w_e^\pm) \to (\mathbb{C}, 0)$$
  in a neighborhood of the nodal lifts that respect the framing (as in (4.17)), for all i, the $(\mu_i + 1)$-th derivative normal to Y_i is equal for u_{v_+} and u_{v_-}. In the direction normal to Y_i, since derivatives up to order μ_i vanish, the $(\mu_i + 1)$-th normal derivative, or the $(\mu_i + 1)$-jet normal to Y_i, denoted by
  $$\text{ev}_{Y_i}^{\mu_i} u_\pm(w_e^\pm) \in N_\pm Y_i \backslash Y_i$$
  is well-defined (see [**22**, Section 6]), where $N_\pm Y_i$ is the normal bundle of Y_i in $X_{P_\pm}$. The vertical matching condition says that
  $$\text{ev}_{Y_i}^{\mu_i} u_+(w_e^+) = \text{ev}_{Y_i}^{\mu_i} u_-(w_e^-).$$

To see that the horizontal and vertical matching conditions together imply the node matching condition in (4.18), we observe that

- a neighborhood of the zero section of the bundle $\oplus_i(NY_i|Y) \to Y$ has an identification to a neighbourhood of Y in $X_\pm$ that is holomorphic on the fibers;

- and near the node w_e the maps $u_{v_\pm}$ is asymptotically close to the vertical cylinder

$$z_\pm \mapsto \left(\frac{a_i}{(\mu_i+1)!} z_\pm^{\mu_i+1} \right).$$

The node matching in (4.18) follows because both u_{v_+} and u_{v_-} are asymptotic to the 'same' vertical cylinder.

This view of the matching condition is relevant only in case of a collection of orthogonal single cuts. In general cases, given a node, both nodal lifts do not lie on the same set of relative divisors. For example, in the broken map in Figure 2.11 corresponding to a Mikhalkin graph, there are edges of the form

- $e = (v_+, v_-)$, $\mathcal{T}(e) = (1,1)$
- with $P(v_+) = P(v_-) = P_0$ which is the zero-dimensional polytope;
- and one of the lifts lies on a divisor of $X_{\overline{P}_0}$ corresponding to a facet with normal vector $(1,1)$ and another lift lies on the intersection of divisors corresponding to facets with normal vectors $(0,-1)$ and $(-1,0)$.

This ends the Remark.

Example 4.29. We unpack the horizontal and vertical matching conditions from Remark 4.28 in the first broken map in Figure 4.3. The multiple cut consists of two orthogonal cuts shown ion Figure 1.4, and the node w lies between maps $u_0 : C_0 \to X_{P_2}$, $u_1 : C_1 \to X_{P_0}$. The nodes $u_0(w)$, $u_1(w)$ lies in $X_{P_\cap}$. Suppose the leading order Taylor term of u_0 in the direction normal to $X_{P_{12}}$ resp. $X_{P_{23}}$ is $a_+ z_+^{\mu_h^+}$ resp. $b_+ z_+^{\mu_v^+}$, and suppose the leading order Taylor term of u_1 in the direction normal to $X_{P_{30}}$ resp. $X_{P_{01}}$ is $a_- z_-^{\mu_h^-}$ resp. $b_- z_-^{\mu_v^-}$. Then the matching condition says that

- (Matching of intersection multiplicities) $\mu_h^+ = \mu_h^-$, $\mu_v^+ = \mu_v^-$,
- (Horizontal matching) the projections of the nodal evaluation maps to $X_{P_\cap}$ are equal : $\pi_{P_\cap}(u(w_+)) = \pi_{P_\cap}(u(w_-)) \in X_{P_\cap}$,
- (Vertical matching) and the leading order Taylor coefficients are equal: $a_+ = a_-$, $b_+ = b_-$.

Remark 4.30. (Comparison to Ionel's refined matching) In [**50**, p14], Ionel states the edge matching condition using a 'refined evaluation map', which is analogous to the tropical evaluation map defined in (4.19). Suppose a node w_e maps to the intersection of relative divisors $Y_1, \dots, Y_n$ and the intersection multiplicity with each of the divisors is $\mu_1, \dots, \mu_n$. (Here we work with the extension of the broken map u over nodal lifts as in Remark 4.26.) The refined evaluation map $\mathrm{ev}_{\mathrm{ref}}(u, 0)$ for the map u and the nodal lift $0 \in C$ is a point in the weighted projective space

$$\mathrm{ev}_{\mathrm{ref}}(u, 0) = (x_1, \dots, x_n) \in \mathbb{P}_{(\mu_1, \dots, \mu_n)}(\oplus_i NY_i).$$

The point $\mathrm{ev}_{\mathrm{ref}}(u, 0)$ is well-defined since replacing the domain coordinate z by az for some $a \in \mathbb{C}^\times$ has the effect of changing $(x_1, \dots, x_n)$ to $(a^{\mu_1} x_1, \dots, a^{\mu_n} x_n)$. The projected tropical evaluation in (4.21)

$$\pi^\perp_{\mathcal{T}(e)}(\mathrm{ev}^\mu(0)) \in (Z \times \mathbb{R}^n)/T_{\mathcal{T}(e),\mathbb{C}}$$

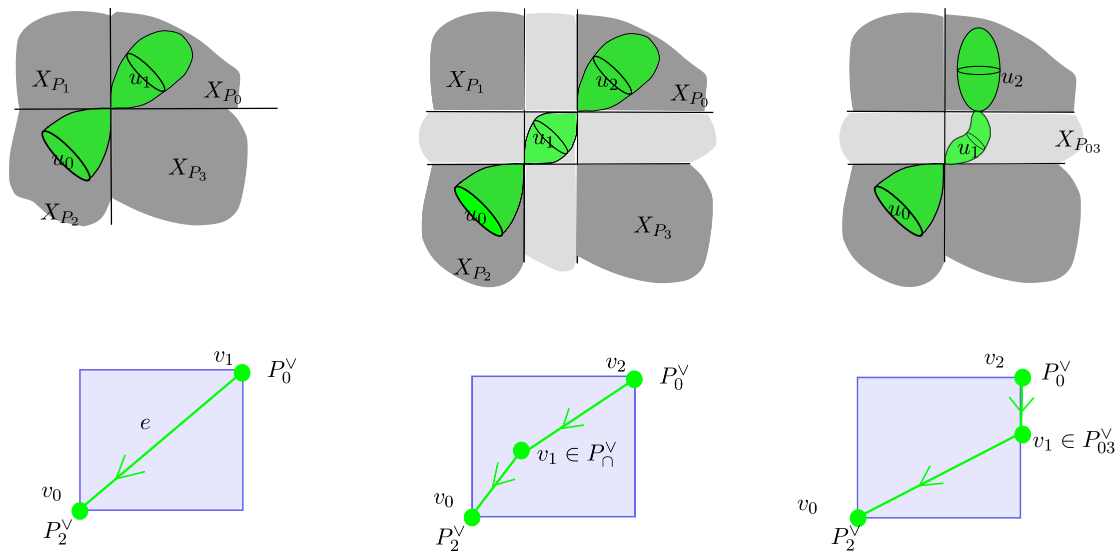

FIGURE 4.3. Broken maps and their dual graphs in the broken manifold of Figure 3.5.

is Ionel's evaluation map written using cylindrical coordinates on the target space, and the slope $\mathcal{T}(e)$ of the edge e is the vector $(\mu_1, \ldots, \mu_n)$. The cylindrical viewpoint appears more natural because in the image of the refined evaluation map none of the coordinates x_i vanish. On the other hand all points in the cylinder $(Z \times \mathbb{R}^n)/T_{\mathcal{T}(e),\mathbb{C}}$ can possibly be in the image of the evaluation map.

REMARK 4.31. (Balancing property) The edge slopes $\mathcal{T}(e)$ for the edges $e \in \operatorname{Edge}(\Gamma)$ coming out of any vertex $v \in \operatorname{Vert}(\Gamma)$ satisfy a *balancing property* in $\mathfrak{t}_{P(v)}$

$$\sum_{e \ni v} \mathcal{T}(e)_{\text{vert}} = c_1((\pi_{P(v)} \circ u_v)^* Z_{P(v)} \to X_{P(v)})); \tag{4.25}$$

Here $(\pi_{P(v)} \circ u_v) : C_v \to X_{P(v)}$ is the projection of the component u_v corresponding to $v \in \operatorname{Vert}(\mathcal{T})$, and $\mathcal{T}(e)_{\text{vert}} \in \mathfrak{t}_{P(v)}$ is the vector sum of intersection multiplicities (see (4.23)) with vertical relative divisors of $X_{\overline{P}(v)} \to X_{P(v)}$ at the node w_e. The projection $\mathcal{T}(e)_{\text{vert}} \in \mathfrak{t}_{P(v)}$ can alternately be defined as the orthogonal projection of $\mathcal{T}(e)$ to $\mathfrak{t}_{P(v)}$ with respect to the inner product (3.6) on $\mathfrak{t}$. Indeed, the map u_v gives a section of the $T_{P(v)}$-principal bundle $(\pi_{P(v)} \circ u_v)^* Z_{P(v)} \to X_{P(v)}$ on the complement of nodal points w_e, $e \ni v$, and the monodromy of the section around each such intersection is determined by $\mathcal{T}(e)_{\text{vert}}$.

In particular, if the horizontal projection $\pi_{P(v)}(u_v)$ is constant (for example, if $X_{P(v)}$ is a point) then one obtains from (4.25) the standard balancing condition

$$\sum_{e \ni v} \mathcal{T}(e)_{\text{vert}} = 0,$$

seen for example, in Mikhalkin's tropical curves (see (2.28)).

REMARK 4.32. (A comparison with symplectic field theory) Holomorphic buildings in symplectic field theory [**11**] are comparable to broken maps on manifolds with a single cut. The target space of a holomorphic building consists of X^+, X^- and $k-1$ copies of the neck piece $Z(\mathbb{P}^1)$ for some $k \geq 1$:

$$\mathfrak{X}[k] := X_+ \cup_Y Z(\mathbb{P}^1) \cup_Y \cdots \cup_Y Z(\mathbb{P}^1) \cup_Y X_-,$$

and any pair of consecutive pieces are identified along a divisor Y. A holomorphic building $u : C \to \mathfrak{X}[k]$ is a continuous map, where nodes map to the divisor Y, and intersection multiplicities $m_{w_\pm(e)}(u_{v_\pm}, Y)$ are equal on both sides.

A holomorphic building differs from a broken map in two ways : A holomorphic building is a continuous map, and the data for a holomorphic building includes an ordering for the neck piece components. In the broken map view, this ordering is not important: With suitable regularity assumptions a broken u map with m components in neck pieces can be glued to give a $2m$-dimensional family of unbroken maps in X^ν for any ν. Any sequence of maps $u_\nu : C_\nu \to X^\nu$ lying in the glued family converges to a broken map $u' : C' \to \mathfrak{X}$ that is related to $u : C \to \mathfrak{X}$ by a tropical symmetry (as in Definition 4.33). In contrast the holomorphic building limits of different sequences in the glued family may not all be the same, since the choice of the sequence $(u_\nu)_\nu$ determines the component of $\mathfrak{X}[k]$ to which a curve component $C_v \subset C$ maps. For a broken map $u = (u_v)$, the ordering of the pieces u_v is not part of the data of the the map. One effect of the differing definitions is the following: Unlike holomorphic buildings, broken maps do not have components u_v that are trivial cylinders in the sense that they map into a fiber of $X_{\overline{P}(v)}$ with only two marked points and so are unstable. In holomorphic buildings, trivial cylinders have to be inserted whenever there is a node $w(e)$ between components C_{v_+}, C_{v_-} that are not in adjacent levels in order to achieve continuity. This ends the Remark.

4.4. Symmetries of broken maps

The definition of equivalence for broken maps involves not only automorphisms of the domains, but the torus actions on the "neck pieces" in the degeneration. These equivalences are encoded in the action of a *tropical symmetry group* defined as follows.

DEFINITION 4.33. (Tropical symmetry) A *tropical symmetry* for a tropical graph Γ is a tuple

$$(\underline{g}, \underline{z}) = ((g_v)_{v \in \mathrm{Vert}(\Gamma)}, (z_e)_{e \in \mathrm{Edge}_{\mathrm{trop}}(\Gamma)}), \quad g_v \in T_{P(v),\mathbb{C}}, \quad z_e \in \mathbb{C}^\times$$

consisting of a translation g_v for each vertex and a change of local coordinate z_e for each tropical edge that satisfies

$$g_{v_+} g_{v_-}^{-1} = z_e^{\mathcal{T}(e)} \quad \forall e = (v_+, v_-) \in \mathrm{Edge}_\bullet(\Gamma), \tag{4.26}$$

where we assume $z_e = 1$ for $e \in \mathrm{Edge}_\bullet(\Gamma) \setminus \mathrm{Edge}_{\mathrm{trop}}(\Gamma)$. A tropical symmetry (g, z) acts on a broken map (u, fr) as

$$u_v \mapsto g_v u_v, \quad \mathrm{fr}_e \mapsto z_e \mathrm{fr}_e.$$

The group of tropical symmetries is denoted

$$T_{\text{trop}}(\Gamma) := \{((g_v)_{v\in \text{Vert}(\Gamma)}, (z_e)_{e\in \text{Edge}_\bullet(\Gamma)})|(4.26)\}. \tag{4.27}$$

The condition (4.26) is a necessary and sufficient condition for the translations $(g_v)_v$ to preserve the matching condition at nodes.

Remark 4.34. (Framing symmetry) There are finitely many tropical symmetries $(\underline{g}, \underline{z}) \in T_{\text{trop}}(\Gamma)$, called *framing symmetries*, for which the action on the unframed map u of type Γ is trivial, that is $g_v = \text{Id}$ for all vertices $v \in \text{Vert}(\Gamma)$. The group of such framing symmetries $(\underline{g}, \underline{z})$ is the product $\prod_{e\in \text{Edge}_\bullet(\Gamma)} \mathbb{Z}_{n_e}$, where n_e is the order of ramification of the broken map at the nodal point w_e, see Remark 4.25.

The next result shows that the identity component of the tropical symmetry group is generated by tropical positions of vertices of the tropical graph.

Lemma 4.35. (Tropical vertex positions generate tropical symmetries)

(a) For a tropical graph Γ, the set of tropical vertex positions $\mathcal{W}(\Gamma)$ (defined in (4.10)) is convex.

(b) (From tropical weights to tropical symmetries) *If a tropical graph Γ has two distinct tropical vertex position maps*

$$(\mathcal{T}_0(v), v \in \text{Vert}(\Gamma)), \quad (\mathcal{T}_1(v), v \in \text{Vert}(\Gamma))$$

then the difference $\mathcal{T}_1 - \mathcal{T}_0$ generates a real-two-dimensional subgroup $\exp((\mathcal{T}_1 - \mathcal{T}_0)(\cdot))$ *of the tropical symmetry group $T_{\text{trop}}(\Gamma)$ (defined in (4.27)).*

(c) The subgroup

$$T_{\text{trop},\mathcal{W}}(\Gamma) := \langle \exp((\mathcal{T}_1 - \mathcal{T}_0)z) | \mathcal{T}_0, \mathcal{T}_1 \in \mathcal{W}(\Gamma), z \in \mathbb{C}\rangle \tag{4.28}$$

generated by tropical vertex position maps is the identity component of $T_{\text{trop}}(\Gamma)$.

Proof. For the first statement, if $\mathcal{T}_0, \mathcal{T}_1 \in \mathcal{W}(\Gamma)$ are vertex position maps for a tropical graph Γ, then for any $t \in [0, 1]$

$$(1 - t)\mathcal{T}_0 + t\mathcal{T}_1 \in \mathcal{W}(\Gamma)$$

is also a vertex position map on Γ. Assume that $\mathcal{T}_0$, $\mathcal{T}_1$ are distinct, and that $l_{e,i} \in \mathbb{R}$ describe the difference between the tropical vertex positions in the sense that

$$\mathcal{T}_i(v_+) - \mathcal{T}_i(v_-) = l_{i,e}\mathcal{T}_e, \quad i = 0, 1, e \in \text{Edge}_\bullet(\Gamma).$$

Then the set of elements

$$g : \mathbb{C} \to T_{\text{trop}}(\Gamma), \quad z \mapsto \begin{cases} e^{(\mathcal{T}_1(v) - \mathcal{T}_0(v))z}, & v \in \text{Vert}(\Gamma) \\ e^{(l_{1,e} - l_{0,e})z}, & e \in \text{Edge}_\bullet(\Gamma) \end{cases}$$

is a non-trivial group of symmetries for broken maps modelled on Γ. In a similar way, we see that the Lie algebra $\mathfrak{t}_{\text{trop}}(\Gamma)$ is equal to the vector space generated by differences of tropical vertex position maps $\mathcal{T}_1 - \mathcal{T}_0$, $\mathcal{T}_0, \mathcal{T}_1 \in \mathcal{W}(\Gamma)$, from where (c) follows. □

Definition 4.36. A tropical graph Γ is *rigid* if its tropical symmetry group $T_{\text{trop}}(\Gamma)$ is finite.

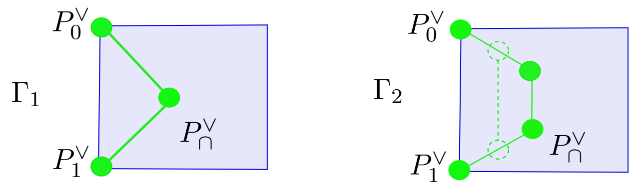

FIGURE 4.4. The dual complex $B^\vee$ for the multiple cut in Figure 1.4 is a rectangle. The tropical graph Γ_1 (left figure) is rigid, but Γ_2 (right figure) is not rigid since the vertices inside the square can be moved to the dotted positions.

LEMMA 4.37. *For any tropical graph Γ, the quotient $T_{\mathrm{trop}}(\Gamma)/T_{\mathrm{trop},\mathcal{W}}(\Gamma)$ is finite. Consequently the group $T_{\mathrm{trop}}(\Gamma)$ has a finite number of connected components.*

PROOF. The quotient $T_{\mathrm{trop}}(\Gamma)/T_{\mathrm{trop},\mathcal{W}}(\Gamma)$ is discrete because by Lemma 4.35 $T_{\mathrm{trop},\mathcal{W}}(\Gamma)$ is the identity component of $T_{\mathrm{trop}}(\Gamma)$. Furthermore, every connected component of $T_{\mathrm{trop}}(\Gamma)$ deformation retracts to a component of the maximal compact subtorus

$$T_{\mathrm{trop}}(\Gamma)^{im} := \{(\underline{g}, \underline{z}) \in T_{\mathrm{trop}}(\Gamma) : g_v \in T_{P(v)}, |z_e| = 1\}.$$

Indeed, an element $(g, z) \in T_{\mathrm{trop}}(\Gamma)$ can be written as

$$g_v = k_v e^{is_v}, \quad k_v \in T_{P(v)}, s_v \in \mathfrak{t}_{P(v)}, \quad z_e = \theta_e e^{\alpha_e}, \quad \theta_e \in S^1, \alpha_e \in \mathbb{R}$$

The tuple $(k, \theta) := ((k_v)_v, (\theta_e)_e)$ is a tropical symmetry element in $T_{\mathrm{trop}}(\Gamma)^{im}$ which is connected to (g, z) via the path

$$[0, 1] \ni \tau \mapsto ((k_v e^{i\tau s_v})_v, (\theta_e e^{\tau\alpha_e})_e) \in T_{\mathrm{trop}}(\Gamma).$$

The quotient $T_{\mathrm{trop}}(\Gamma)/T_{\mathrm{trop},\mathcal{W}}(\Gamma)$ is finite because it is in bijection with the connected components of the compact subgroup $T_{\mathrm{trop}}(\Gamma)^{im}$. □

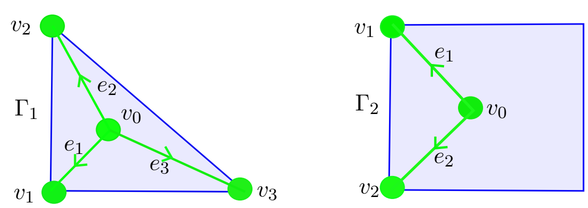

FIGURE 4.5. Rigid tropical graphs in Example 4.38

EXAMPLE 4.38. The tropical graphs in Figure 4.5 are rigid, but have non-trivial tropical symmetry groups. Suppose the tropical graph Γ_1 has edge slopes

$$\mathcal{T}_{e_1} = (-1, -1), \quad \mathcal{T}_{e_2} = (-1, 2), \quad \mathcal{T}_{e_3} = (2, -1).$$

A tropical symmetry $(\underline{g}, \underline{z})$ on Γ_1 satisfies the equations

$$g_1 = g_2 = g_3 = \mathrm{Id}, \quad g_1 g_0^{-1} = z_{e_1}^{(-1,-1)}, \quad g_2 g_0^{-1} = z_{e_2}^{(-1,2)}, \quad g_3 g_0^{-1} = z_{e_3}^{(2,-1)},$$

and is therefore given by

$$g_0 = (\omega, \omega), \quad z_{e_1} = z_{e_2} = z_{e_3} = \omega,$$

where $\omega \in e^{2\pi ik/3}$ is a cube root of unity. Thus

$$|T_{\mathrm{trop}}(\Gamma_1)| = 3. \tag{4.29}$$

This tropical graph is similar to the example studied by Abramovich-Chen-Gross-Siebert [**2**, p51] and Tehrani [**30**, Section 6]. The tropical graph Γ_2 has edge slopes

$$\mathcal{T}_{e_1} = (-1, 1), \quad \mathcal{T}_{e_2} = (-1, -1).$$

By a similar calculation $|T_{\mathrm{trop}}(\Gamma_2)| = 2$.

CHAPTER 5

Stabilizing divisors

We use domain-dependent perturbations of the almost complex structure to regularize the moduli space of pseudoholomorphic disks/spheres. As in Cieliebak-Mohnke [**22**], we use Donaldson divisors to stabilize the domains so that domain-dependent perturbations are possible. We will require the interior markings on the domains of disks/spheres to coincide with the intersections points of the pseudoholomorphic map with the Donaldson divisor. If the divisor has high enough degree, all domain components are stable as curves. Consequently the almost complex structure can be perturbed to attain regularity on all domain components. Of course, one could envision using any of the current perturbation schemes to achieve virtual fundamental chains; the stabilizing divisor approach is merely a convenience.

5.1. Stabilizing divisors in smooth symplectic manifolds

We recall results about stabilizing divisors for unbroken manifolds. We consider divisors of Donaldson-type, any such divisor is Poincaré dual to $k[\omega]$ where $k \gg 0$ is a large integer. To define and construct Donaldson divisors we assume that all the symplectic forms are rational.

Definition 5.1. The symplectic manifold (X, ω) is *rational* if $[\omega] \in H^2(X, \mathbb{Q})$. A *prequantum line bundle* is a line-bundle-with-connection $\tilde{X} \to X$ whose curvature is

$$\operatorname{curv}(\tilde{X} \to X) = (2\pi/i)\omega \in \Omega^2(X, i\mathbb{R}).$$

A Lagrangian $L \subset X$ is *rational* if there exists a prequantum bundle $\tilde{X}$ an integer k and a flat section $s : L \to \tilde{X}^{\otimes k}$ of the restriction $\tilde{X}^{\otimes k}|L$.

Definition 5.2. (a) A *divisor* in X is a symplectic submanifold $D \subset X$ of real codimension $\operatorname{codim}(D) = 2$. A *Donaldson divisor* is a divisor D that is Poincaré dual to $k[\omega]$ for some large integer $k > 0$, and k is called the *degree* of the divisor D.

(b) An almost complex structure $J \in \mathcal{J}(X)$ is *adapted* to D if $J(TD) = TD$. The space of D-adapted almost complex structures is denoted by

$$\mathcal{J}(X, D) := \{J \in \mathcal{J}(X) | J(TD) = TD\}.$$

(c) Let $D \subset X - L$ be a divisor disjoint from the Lagrangian L. For $E > 0$, a D-adapted almost complex structure J_D is E- *stabilized* by the divisor D iff

(i) (No non-constant spheres) D does not contain any J_D-holomorphic sphere $u : \mathbb{P}^1 \to X$ with symplectic area $< E$,

(ii) (Sufficient intersections) any J_D-holomorphic sphere in X with symplectic area $< E$ has at least 3 distinct points of intersection with D, and any J_D-holomorphic disk with symplectic area less than E has at least one intersection with D.

A pair (J, D) consisting of a divisor D and an ω-tamed $J \in \mathcal{J}(X, D)$ is *stabilizing* if J is E-stabilizing for all $E > 0$.

PROPOSITION 5.3. (Existence of a stabilizing pair, [**17**, Section 4], [**22**, Section 8]) *Suppose $\omega \in \Omega^2(X)$ is a rational symplectic form and $L \subset X$ is a compact rational Lagrangian submanifold. Then, there exists*

(a) a divisor $D \subset X - L$ such that L is exact in $X - D$;
(b) and a tamed almost complex structure $J \in \mathcal{J}(X, D)$ such that (J, D) is a stabilizing pair and for every $E > 0$ there is an open neighbourhood

$$\mathcal{J}(X, D; J_0, E) \subset \{J \in \mathcal{J}(X, D) : J|TD = J_0|TD\}$$

of J_0 consisting of E-stabilizing almost complex structures adapted to D.

The proof of the Proposition is outlined below. The construction of a stabilizing pair uses the following notion of a degree bound.

DEFINITION 5.4. A constant $k_* > 0$ is a *degree bound* for a tamed almost complex structure J on (X, ω) if for any J-holomorphic sphere $u : \mathbb{P}^1 \to X$,

$$c_1(u) := \int_{\mathbb{P}^1} u^* c_1(TX) \leq k_* \omega(u).$$

LEMMA 5.5. (Existence of uniform degree bounds) *(*[**22**, Lemma 8.11]*) Let J_0 be a compatible almost complex structure on the compact symplectic manifold (X, ω). For any $0 < \epsilon < 1$ there exists a constant k_* which is a degree bound for all tamed almost complex structures J satisfying $\|J - J_0\|_{C^0(X)} < \epsilon$.*

We reproduce the proof from [**22**] since an adaptation of the technique is used in the case of broken manifolds.

PROOF OF LEMMA 5.5. Let $\gamma \in \Omega^2(X)$ be a closed two-form in the class of the first Chern class $c_1(TX) \in H^2(X)$. Using the norm $|v|^2 := \omega(v, J_0 v)$ on TX, and a C^0 norm on forms, for any $J \in B_\epsilon(J_0)$ and $v \in TX$, we have the bounds

$$\gamma(v, Jv) \leq \|\gamma\|(1 + \|J - J_0\|)|v|^2 \leq \|\gamma\|(1 + \epsilon)|v|^2,$$
$$\omega(v, Jv) \geq \|(\|1 - \|J - J_0\|)|v|^2 \geq (1 - \epsilon)|v|^2.$$

Defining $k_* := \frac{1+\epsilon}{1-\epsilon}$, we have

$$\gamma(v, Jv) \leq k_* \omega(v, Jv).$$

For a K-holomorphic sphere $u : \mathbb{P}^1 \to X$ this implies $c_1(u) \leq k_* \omega(u)$. □

REMARK 5.6. (Using the degree bound for stabilizing) The degree bound on almost complex structures is related to the stabilizing property in Definition 5.2

as follows. By [**22**, Lemma 8.13], if a divisor D is Poincaré dual to $k[\omega]$ and J is adapted to D and has a degree bound k_* satisfying

$$k > 2\max\{k_*, k_* + \tfrac{1}{2}(\dim(X)) - 2\} \tag{5.1}$$

then the expected dimension of moduli space of J-holomorphic spheres that are not stabilized by D is negative. Therefore if J were chosen generically, (J, D) would be a stabilizing pair.

Outline of proof of Proposition 5.3. By the adaptation of Donaldson's construction in Auroux-Gayet-Mohsen [**5**], one can find an approximately holomorphic divisors in $X - L$. By [**17**, Theorem 3.6], if the Lagrangian L is rational, then for any divisor D in $X - L$ produced by Auroux-Gayet-Mohsen [**5**], L is exact in $X - D$. (We refer the reader to [**17**, Example 3.2] for the definition of 'exactness'.) Exactness allows the following relation of area to intersection numbers: If L is exact in $X - D$, then, the intersection number of a disk $u : (C, \partial C) \to (X, L)$ with the divisor D is proportional to the area of the disk: If $[D]^\vee = k[\omega]$ for some $k \in \mathbb{Z}$, then,

$$\# u^{-1}(D) = k \int_C u^*\omega. \tag{5.2}$$

See [**17**, Lemma 3.4]. Consequently for any J that is adapted to D, any non-constant J-holomorphic disk is automatically stabilized since it has at least one marked point.

To construct a stabilizing pair one starts with a preliminary almost complex structure J^{pre} that is ω-compatible. For any $\epsilon > 0$ there exists a constant $k_*(\epsilon, J^{\mathrm{pre}})$ that is a degree bound for any ω-tamed $J \in \mathcal{J}(X)$ satisfying $\|J - J^{\mathrm{pre}}\|_{C^0} \leq \epsilon$. This fact gives enough wiggle room to find a stabilizing pair as follows: For any $\epsilon' > 0$ there is a constant $k_d(\epsilon')$ such that if the degree of the Donaldson divisor D is $\geq k_e(\epsilon')$, then there is a tamed almost complex structure $J_1 \in \mathcal{B}_{\epsilon'}(J^{\mathrm{pre}})$ that is D-adapted. We choose the degree k of the Donaldson divisor to be high enough that it satisfies the bound (5.1) and $k \geq k_d(\epsilon/2)$. Then, there is an open subset of $\mathcal{J}(X, D)$ contained in $B_\epsilon(J^{\mathrm{pre}})$, and for any $J \in \mathcal{J}(X, D) \cap B_\epsilon(J^{\mathrm{pre}})$, k_* is a degree bound in the sense of Definition 5.4. Therefore by Remark 5.6, for a generic J_0 in $\mathcal{J}(X, D) \cap B_\epsilon(J^{\mathrm{pre}})$, the pair (J_0, D) is stabilizing for all spheres. See [**22**, Section 8] for details. □

Remark 5.7. The construction of a stabilizing pair extends to the case when X is a symplectic orbifold with the Lagrangian contained in the smooth locus. Locally if X is the quotient of a symplectic manifold $\overline{X}$ by a finite group Γ, then the divisor D is the quotient of a Γ-invariant divisor in $\overline{X}$. Donaldson's divisor construction is extended to the orbifold case by Gironella-Muñoz-Zhou [**37**]. Since the Lagrangian is smooth the modifications of Auroux-Gayet-Mohsen [**5**] do not interfere with the arguments in Gironella-Muñoz-Zhou [**37**]. The relation between symplectic area and the number of divisor intersections extends to the orbifold case if the Lagrangian does not contain orbifold singularities, the proof from [**17**] carries over verbatim.

5.2. Stabilizing divisors in broken manifolds

In this section, we construct stabilizing divisors in a broken manifold by a modification of Donaldson's construction. A broken divisor $\mathfrak{D}$ in a broken manifold $\mathfrak{X}$ is a Donaldson divisor D_P in each component X_P of $\mathfrak{X}$ that is cylindrical in the ends of X_P. Such submanifolds are constructed as approximately holomorphic submanifolds with respect to a compatible almost complex structure.

Broken manifolds with cylindrical almost complex structures do not have natural embeddings into compact symplectic broken manifolds. We now define a modified version of cylindricity, called '$\omega_{\mathfrak{X}}$-cylindricity' for almost complex structures so that broken symplectic manifolds possess compatible $\omega_{\mathfrak{X}}$-cylindrical almost complex structures, which can be used for the construction of Donaldson divisors. The flipside is that an $\omega_{\mathfrak{X}}$-cylindrical almost complex structure is not *gluable*, that is, it can not be glued on the ends to yield neck-stretched almost complex structures on X^ν. Indeed the $\mathfrak{t}$-inner product underlying an $\omega_{\mathfrak{X}}$-cylindrical almost complex structure is not the same for all the components X_P of the broken manifold $\mathfrak{X}$, and has the following form.

Definition 5.8. ($\mathfrak{X}$-inner product) Let $g^P_Q : \mathfrak{t}_Q \times \mathfrak{t}_Q \to \mathbb{R}$ be a collection of inner products for all $Q \subset P$ and $P \in \mathcal{P}^0$ ranges over top-dimensional polytopes that satisfy

(a) (Restriction) $R \subset Q \implies g^P_R|\mathfrak{t}_Q = g^P_Q$,
(b) (Orthogonality) For $Q \subset P$, $\dim(Q) = 0$, which is the intersection of facets $Q = \cap_{i=1}^n Q_i$ of P, the inner product g^P_Q satisfies the condition that the outward normals to Q_i form are orthogonal.

We re-define symplectic cylindrical structure maps on cut spaces using $\mathfrak{X}$-inner products for the consistency condition.

Definition 5.9. ($\mathfrak{X}$-symplectic cylindrical structure maps) For any top-dimensional polytope P let $(\phi^P_Q)_{Q\subset P}$ be a collection of symplectic cylindrical structure maps

$$U_{X_P}(X_Q) \xrightarrow{\phi_P} ((\mathrm{Cone}_Q P \times Z_Q), \overline{\omega}_Q), \quad \overline{\omega}_Q = \omega_{X_Q} + d\langle \alpha^P_Q, \pi_{\mathfrak{t}^\vee_Q}\rangle \tag{5.3}$$

on neighborhoods $U_{X_Q}(X_P) \subset X_P$ of X_Q, where

- $\alpha^P_Q \in \Omega^1(Z_Q, \mathfrak{t}_Q)$ is a connection one-form on the T_Q bundle $Z_Q \to X_Q$, where the collection $(\alpha^P_Q)_{Q\subset P}$ satisfies the consistency condition (3.16) with respect to the $\mathfrak{X}$-inner product;
- and $\pi_{\mathfrak{t}^\vee_Q} : \mathfrak{t}^\vee \to \mathfrak{t}^\vee_Q$ is the projection on $\mathrm{Cone}_Q P \subset \mathfrak{t}^\vee$, and so, $\langle \alpha_Q, \pi_{\mathfrak{t}^\vee_Q}\rangle$ is a one-form on $(Z_Q \times \mathrm{Cone}_P Q)$.

The map ϕ^P_Q induces a projection map

$$\pi^P_Q : U_{X_Q}(X_P) \to X_Q.$$

Finally we define $\omega_{\mathfrak{X}}$-compatible almost complex structures. In addition to being ω-compatible on each piece, in neighborhoods $U_{X_Q}(X_P)$ of boundary submanifolds these almost complex structures are integrable on the fibers of the map $\pi_{P,Q}$.

DEFINITION 5.10. ($\omega_{\mathfrak{X}}$-cylindrical almost complex structures) Let $\mathfrak{X}$ be a broken symplectic manifold equipped with $\mathfrak{X}$-symplectic cylindrical structure maps as in (5.3) and resultant projection maps

$$\pi_P^Q : U_{X_P}(X_Q) \to X_Q$$

on neighborhoods $U_{X_Q}(X_P) \subset X_P$ of X_Q for $Q \subset P$.

(a) ($\omega_{\mathfrak{X}}$-cylindrical) An $\omega_{\mathfrak{X}}$-*cylindrical* almost complex structure $\mathfrak{J} = (J_P)_{P\in\mathcal{P}}$ on $\mathfrak{X}$ consists of an almost complex structure J_P on each (compactified) cut space X_P, $P \in \mathcal{P}$ satisfying
   (i) (Restriction) for any pair $Q \subset P$ of polytopes, X_Q is a J_P-holomorphic submanifold of X_P, and $J_P|X_Q = J_Q$;
   (ii) (Cylindrical structure) for any pair $Q \subset P$ of polytopes, the fibers of π_P^Q are $\mathfrak{J}$-holomorphic and $\mathfrak{J}$ is integrable on each of the fibers.

(b) ($\omega_{\mathfrak{X}}$-compatibility) An $\omega_{\mathfrak{X}}$-cylindrical almost complex structure $\mathfrak{J} = (J_P)_{P\in\mathcal{P}}$ is $\omega_{\mathfrak{X}}$-*compatible* if for any $P \in \mathcal{P}$, J_P is ω_{X_P}-compatible.

(c) ($\omega_{\mathfrak{X}}$-strong tamedness) An $\omega_{\mathfrak{X}}$-cylindrical almost complex structure $\mathfrak{J} = (J_P)_{P\in\mathcal{P}}$ is $\omega_{\mathfrak{X}}$-*strongly tamed* if for any $P \in \mathcal{P}$, J_P is ω_{X_P}-tame, and additionally for any pair of polytopes $Q \subset P$ and a point $x \in U_{X_Q}(X_P)$, the ω-complement of $\ker(d\pi_Q^P)_x \subset T_x X_P$ is J_P-holomorphic. (Equivalently the connection one-form underlying J_P in the Q-cylindrical end $U_{X_Q}(X_P)$ is equal to α_Q^P occuring in the $\mathfrak{X}$-symplectic cylindrical structure.)

REMARK 5.11. $\omega_{\mathfrak{X}}$-cylindrical almost complex structures can not be glued along the necks because the underlying $\mathfrak{t}$-inner products and connection one-forms vary across cut spaces.

LEMMA 5.12. *An $\omega_{\mathfrak{X}}$-compatible almost complex structure exists.*

PROOF. We construct an $\omega_{\mathfrak{X}}$-compatible almost structure by induction on the face structure of the polytopes in the polyhedral decomposition. Consider a polytope $Q \in \mathcal{P}\backslash\mathcal{P}^0$ that is not top-dimensional, and assume that an $\omega_{\mathfrak{X}}$-compatible

$$\mathfrak{J}| \cup_{R\subset Q, P\in\mathcal{P}^0 : Q\subset P} U_{X_R}(X_P) \tag{5.4}$$

is already defined. First we extend $\mathfrak{J}| \cup_{R\subset Q} U_{X_R}(X_Q)$ to an ω_{X_Q}-compatible almost complex structure over X_Q. Next we extend the definition of $\mathfrak{J}$ to neighborhoods $U_{X_Q}(X_P)$ for any $P \supset Q$, $P \in \mathcal{P}^0$ as follows. Suppose $Q_1, \ldots, Q_q \in \mathcal{P}$ are facets of P such that $\cap_i Q_i = Q$. Identify $T_Q \simeq \prod_{i=1}^q T_{Q_i} \simeq (S^1)^q$. There is a family of T_Q-equivariant symplectic embeddings

$$(\pi_Q^P)^{-1}(x) \xrightarrow{\phi_x} (\mathbb{C}^q, \omega_{std})$$

where $\mathbb{C}^q$ has the standard action of $(S^1)^q$, the embeddings ϕ_x varies smoothly with $x \in Q$, and $(\pi_Q^P)^{-1}(x) \cap X_{Q_i}$ is mapped to the subspace $\{z_i = 0\} \subset \mathbb{C}^q$. The fibers of the projection $\pi_Q^P : U_{X_Q}(X_P) \to X_Q$ are equipped with the complex structure pulled back by ϕ_x. For any $y \in U_{X_Q}(X_P)$, on the ω-complement of $\ker(d\pi_Q^P)_y$ in the tangent space $T_y X_P$, we define $\mathfrak{J}$ to be equal to $\mathfrak{J}_{X_Q}(\pi_Q^P(y))$. The definition of $\mathfrak{J}$

in neighborhoods of X_Q agrees with the existing definition in neighborhoods of X_R for any $R \subset Q$. Indeed, for any $R \subset Q$, and top-dimensional polytope $P \supset Q$, the definition of the $\mathfrak{X}$-inner product implies that $\mathfrak{J}|U_{X_Q}(X_P)$ agrees with $\mathfrak{J}|U_{X_R}(X_P)$ from (5.4) in the overlaps. In the concluding step of the induction, for each top-dimensional polytope $P \in \mathcal{P}^0$, we extend $\mathfrak{J}| \cup_{Q \subset P} U_{X_Q}(X_P)$ to an ω_{X_P}-compatible almost complex structure on X_P. □

DEFINITION 5.13. (Divisor in a broken manifold) Suppose $(\mathfrak{X}, \omega)$ is a broken manifold associated to a collection of polytopes $\mathcal{P}$. Suppose $\mathfrak{X}$ is equipped with a $\mathfrak{X}$-symplectic cylindrical structure (see Definition 5.9) in the neighborhoods of relative divisors. For any pair $Q \subset P$ of polytopes let

$$\pi_Q^P : U_{X_Q}(X_P) \to X_Q.$$

be the projection map on the neighborhood $U_{X_Q}(X_P) \subset X_P$ of X_Q induced by the symplectic cylindrical structure. A *broken divisor* $\mathfrak{D}$ in $\mathfrak{X}$ is a collection of codimension two symplectic submanifolds

$$\mathfrak{D} = \{D_{\overline{P}} \subset X_{\overline{P}}, \quad P \in \mathcal{P}\}$$

such that

(a) $D_{\overline{P}}$ is a lift of a symplectic divisor $D_P \subset X_P$ by the projection map $\pi_P : X_{\overline{P}} \to X_P$, and
(b) for any pair of polytopes $Q \subset P$, D_Q is the intersection $D_P \cap X_Q$.

The broken divisor $\mathfrak{D}$ is *cylindrical*, if for any pair of polytopes $Q \subset P$, $\operatorname{codim}(P) = 0$, the divisor $D_P \subset X_P$ is Q-cylindrical in the neighbourhood of X_Q. That is, in the neighbourhood $U_{X_Q}(X_P) \subset X_P$, $D_P \cap U_{X_Q}(X_P)$ is equal to $(\pi_Q^P)^{-1}(D_Q)$.

As in the unbroken case we consider broken divisors $\mathfrak{D}$ of Donaldson-type, which means that there is a large integer k such that for any P, $D_P \subset X_P$ is Poincaré dual to $k[\omega_{X_P}]$. The divisor is constructed so that for any pseudoholomorphic disk or sphere in any component $X_P \subset \mathfrak{X}$ the number of intersections with $\mathfrak{D}$ is equal to k times the ω_{X_P}-area.

The following is the main result of the section.

PROPOSITION 5.14. *Let $\mathfrak{X}$ be a broken manifold with a $\omega_{\mathfrak{X}}$-compatible almost complex structure $\mathfrak{J}$. For any $\theta > 0$, there is a cylindrical Donaldson-type divisor $\mathfrak{D} \subset \mathfrak{X}$ that is θ-approximately holomorphic.*

We refer the reader to [**26**] for the definition of θ-approximate holomorphicity. Before constructing the divisor, we show that a broken divisor $\mathfrak{D}$ can be isotoped to a cylindrical divisor if $\mathfrak{D}$ intersects relative divisors in $\mathfrak{X}$ ω-orthogonally.

LEMMA 5.15. (Making a broken divisor cylindrical) *Suppose $\mathfrak{X}$ is a broken manifold with a $\omega_{\mathfrak{X}}$-strongly compatible cylindrical almost complex structure J. For any $0 < \theta_0 < 1$, there exists θ_1 such that the following is satisfied. Let $\mathfrak{D} \subset \mathfrak{X}$ be a broken divisor that is θ_1-approximately J-holomorphic and intersects relative divisors ω-orthogonally. Then there is a cylindrical broken divisor $\mathfrak{D}_1$ that is homotopic to $\mathfrak{D}$ and which is θ_0-approximately holomorphic.*

Proof of Lemma 5.15. The proof is an induction on the face structure of the polyhedral decomposition $\mathcal{P}$. For any pair $Q \subset P$, $\operatorname{codim}(P) = 0$ let

$$\psi_Q^P : U_{X_Q} X_P \to (\operatorname{Cone}_{P^\vee}(Q^\vee) \times Z_Q, \omega_{\overline{Q}})/\sim$$

be the symplectic cylindrical structure on a neighborhood $U_{X_Q} X_P \subset X_P$ of boundary submanifold X_Q. Here the symplectic form $\omega_{\overline{Q}}$ is as in (3.43) and the equivalence relation $\sim$ mods out boundaries by circle actions as described in (3.43). Since the divisor $\mathfrak{D}$ intersects boundary submanifolds ω-orthogonally, the relative symplectic neighbourhood theorem 3.46 implies that there is a different symplectic cylindrical structure

$$\phi_Q^P : U_{X_Q} X_P \to (\operatorname{Cone}_{P^\vee}(Q^\vee) \times Z_Q, \omega_{\overline{Q}}), \quad Q \subset P$$

that is adapted to $\mathfrak{D}$. That is, in the neighborhood $U_{X_Q} X_P$, D_P projects to $D_Q \subset X_Q$. The map $(\psi_Q^P)^{-1} \circ \phi_Q^P$ is homotopic to the identity, and is equal to identity on $X_Q \setminus \cup_{R \subset Q} U_{X_R} X_P$. Therefore, there exist truncations of the cylindrical ends $U'_{X_Q} X_P \subset U_{X_Q} X_P$ for all $Q \subset P$ and diffeomorphisms $\phi_P : X_P \to X_P$ for all top-dimensional P,

- that are equal to $(\psi_Q^P)^{-1} \circ \phi_Q^P$ on $U'_{X_Q} X_P \subset U_{X_Q} X_P$;
- that are equal to the identity in the complement of all the cylindrical neighborhoods $U_{X_Q} X_P$; and
- for any face $Q \subset P$, $\phi_P(X_Q) = X_Q$ and for any pair of top-dimensional polytopes P_0, P_1 containing Q, $\phi_{P_0} | X_Q = \phi_{P_1} | X_Q$.

The map ϕ_P is constructed inductively. It is first defined on X_Q where $Q \in \mathcal{P}$ is the least dimensional face of P and then extended to higher dimensional faces. The broken divisor

$$\mathfrak{D}_1 := (D_{1,P} \subset X_P)_{P \in \mathcal{P}, \operatorname{codim}(P) = 0}, \quad D_{1,P} := \phi_P(D_P)$$

is cylindrical. If the truncated neighborhoods $U_{X_Q} X_P$ are small enough, ϕ_P is small in the C^1-norm, and $\mathfrak{D}_1$ is θ_0-approximately holomorphic. Therefore the broken divisor $\mathfrak{D}_1$ satisfies the properties required by the Proposition. □

Remark 5.16. The proof of Lemma 5.15 can be adapted to the case when the components X_P, $P \in \mathcal{P}$ are orbifolds. The maps ψ_Q^P, ψ_Q^P are defined on finite covers $\mathbb{C}^n \to \mathbb{C}^n/\Gamma$ and are equivariant with respect to the action of the finite group Γ.

In the rest of this section, we construct a stabilizing divisor in the cut spaces of a broken manifold that intersects relative divisors ω-orthogonally. The construction of the divisors is via a slight modification of Donaldson's technique [**26**]. We give an outline of Donaldson's construction following Auroux [**4**]. Let (X, ω) be a symplectic manifold with a compatible almost complex structure J. Let $\widehat{X} \to X$ be a Hermitian line-bundle with connection α over X whose curvature two-form $\operatorname{curv}(\alpha)$ satisfies $\operatorname{curv}(\alpha) = (2\pi/i)\omega$; since our symplectic manifolds are rational we may always assume this to be the case after taking a suitable integer multiple of the symplectic form. We will construct approximately holomorphic sections s_k of the line bundles $\widetilde{X}^k$ for large k that are transverse in the sense that $|\overline{\partial} s_k| \ll |\partial s_k|$ on the zero set

$s_k^{-1}(0)$. Then for large enough k the zero set $s_k^{-1}(0)$ is transversely cut out, and is a divisor of X with degree k. To study sections on the bundle $\widetilde{X}^k$, we use the metric

$$g_k := k\omega(\cdot, J\cdot)$$

on X. Under this metric, the effects of the non-integrability of J become negligible as k increases. We define the notions of 'approximate holomorphicity' and 'transversality':

DEFINITION 5.17. (Asymptotically holomorphic sequences of sections) Let (X, ω) be a symplectic manifold with ω-compatible almost complex structure J and a prequantum line bundle $\widehat{X} \to X$. Let $(s_k)_{k\geq 0}$ be a sequence of sections of $\widehat{X}^k \to X$.

(a) The sequence $(s_k)_{k\geq 0}$ is *asymptotically holomorphic* if there exists a constant C and integer k_0 such that for $k \geq k_0$,

$$|s_k| + |\nabla s_k| + |\nabla^2 s_k| \leq C, \quad |\overline{\partial} s_k| + |\nabla \overline{\partial} s_k| \leq Ck^{-1/2}. \tag{5.5}$$

(b) The sequence $(s_k)_{k\geq 0}$ is *uniformly transverse* to 0 if there exists a constant η independent of k, and $k_0 \in \mathbb{Z}$ such that for any $x \in X$ and $k \geq k_0$ with $|s_k(x)| < \eta$, the derivative ∇s_k of s_k is surjective at any point and satisfies $|\nabla s_k(x)| \geq \eta$.

In both definitions the norms of the derivatives ∇s_k are evaluated using the metric $g_k = k\omega(\cdot, J\cdot)$.

Donaldson's construction uses a family of asymptotically holomorphic sections on $\widetilde{X}^k$, that are concentrated at a given point in X :

LEMMA 5.18. (Gaussian sections, [4, Lemma 2]) *Given any point $x \in X$, for all large enough k, there exist asymptotically holomorphic sections s $\sigma_{k,x} : X \to \widetilde{X}^k$ such that $|\sigma_{k,x}| \geq c_0$ at every point of the ball of g_k-radius 1 centered at x, for some universal constant $c_0 > 0$, and such that the sections $\sigma_{k,x}$ have uniform Gaussian decay away from x in the C^2-norm. Here, 'uniform Gaussian decay' means that for any $x, y \in X$, $|\sigma_{k,x}(y)|$, $|\nabla\sigma_{k,x}(y)|_{g_k}$, $|\nabla^2\sigma_{k,x}(y)|_{g_k}$ are all bounded by $P(d_{g_k}(x,y))\exp(-\lambda d_{g_k}(x,y)^2)$, where P is a polynomial, and P, λ are independent of k, x.*

The Gaussian section centered at $x \in X$ is constructed by choosing Darboux coordinates $(z_1, \ldots, z_n) : (U_x, x) \to (\mathbb{C}^n, 0)$ that are approximately holomorphic in the following sense: If $\psi : \mathbb{C}^n \to X$ is the inverse map then the complex structures $\psi^* J$ and i differ by $c|z|$, where c is a uniform constant independent of x, and the derivative $|\nabla(\psi^* J - i)|$ is uniformly bounded. We choose a trivialization of the bundle $\widetilde{X}$ so that the Hermitian connection is $\sum_i (z_i d\overline{z}_i - \overline{z}_i dz_i)$ plus terms of higher order in $z, \overline{z}$. The Gaussian section is then defined as

$$\sigma_{k,x} := \rho_{k,x} \cdot e^{-k|z|^2}$$

on U_x, where $\rho_{k,x}$ is a cut-off function vanishing at a g_k-distance of $k^{1/6}$ from x, and extended by zero outside U_x.

The globalization process in Donaldson's construction uses the following result:

LEMMA 5.19. (Quantitative Sard's theorem, [**26**, Theorem 20]) *Suppose* $0 < \delta < \frac{1}{4}$*, and* $f : B_+ \to \mathbb{C}$ *is a function defined on a ball* $B_+ \subset \mathbb{C}^n$ *of radius* $\frac{11}{10}$ *that satisfies* $\|f\|_{C^1} \leq \eta$*, where* $\eta := \delta \log(\delta^{-1})^{-p}$*. Then, there exists* $w \in \mathbb{C}$*,* $|w| \leq \delta$ *such that* $f - w$ *is* η*-transverse to* 0 *over the interior ball* B *of radius* 1*.*

Lemma 5.19 is used to modify an approximately holomorphic section in order to achieve uniform transversality in the neighborhood of a given point: Given an asymptotically holomorphic section $s_k : X \to \widetilde{X}^k$ and a point $x \in X$, by applying Lemma 5.19 to the section $f_k := s_k/\sigma_{k,x} : B_{g_k}(x, \frac{11}{10}) \to \mathbb{C}$, we obtain a constant $w_k \in \mathbb{C}$ such that $f_k - w_k$ is uniformly transverse to the zero section. Consequently $s_k - w_k\sigma_{k,x}$ is also uniformly transverse to the zero section in a neighborhood of x.

Uniform transversality on the entire manifold is obtained by iteratively adding successively smaller contributions, so that the transversality over previous neighborhoods is not disturbed. In particular, Donaldson's construction chooses a lattice of points $\Lambda_k \subset X$ such that X is covered by the unit balls $B_{g_k}(x, 1)$, $x \in \Lambda_k$; and partitions the lattice into N sets (where N is k-independent)

$$\Lambda_k = I_1^k \cup \cdots \cup I_N^k$$

such that any two points in a set I_j^k are separated by a uniform g_k-distance guaranteeing the following: For any I_j^k, the constants $w_{k,x}$ for the balls centered at points $x \in I_j$ can be chosen simultaneously and independently of each other. If $j_1 < j_2$, the constants $w_{k,x}$ for $x \in I_{j_2}^k$ are chosen to be small enough to not break the transversality in the balls belonging to $I_k^{j_1}$. We start with the zero section, and run the iteration described above. The iteration terminates in N steps, which is k-independent, and the resulting section is a sum

$$\sigma_k := \sum_{x \in \Lambda_k} w_{k,x}\sigma_{k,x}.$$

that is uniformly transverse on X for large k.

REMARK 5.20. In case of a symplectic orbifold (X, ω), we consider the prequantum orbibundle $\widetilde{X}$, and a sequence of sections $s_k : X \to \widetilde{X}^k$ for k such that $\widetilde{X}^k$ is a line bundle on X. Notions of asymptotic holomorphicity and transversality extend naturally to the orbifold case and are defined in Gironella-Muñoz-Zhou [**37**].

PROPOSITION 5.21. (Construction of a broken divisor) *Suppose* $\mathfrak{X}$ *is a broken manifold, for which the symplectic form* ω_{X_P} *on each of the cut spaces* X_P*,* $P \in \mathcal{P}$ *is rational. Given a* $\omega_{\mathfrak{X}}$*-strongly compatible cylindrical almost complex structure* J *on* $\mathfrak{X}$*, there is a sequence of asymptotically holomorphic and uniformly transverse sections* $\{s_{k,P} : X_P \to \widehat{X}_P\}$ *for all* $P \in \mathcal{P}$ *such that*

(a) (Restriction) *for any pair of polytopes* $Q \subset P$*,* $s_{k,P}|X_Q = s_{k,Q}$*,*
(b) (Orthogonality to relative divisors) *and the zero set* $s_{k,P}^{-1}(0)$ *intersects* X_Q ω*-orthogonally.*

PROOF. The sections are constructed by running Donaldson's procedure simultaneously for all the manifolds in the set $\{X_P\}_{P \in \mathcal{P}}$. Our modification of Donaldson's algorithm is limited to choosing appropriate Gaussian sections in order to ensure the conditions (Restriction) and (Orthogonality to relative divisors). The step of achieving global transversality by applying Lemma 5.19 is the same as the original algorithm, and therefore not discussed.

We first describe a set of center points for the Gaussian section for each tensor power of the given line bundle. Given $k \gg 0$, a set of center points $\Lambda_{k,P} \subset X_P$ is defined so that it contains $Ck^{\dim(X_P)}$ number of points, and X_P is covered by the following neighbourhoods:

$$X_P = (\cup_{x \in \Lambda_{k,P}} B_{g_k}(x, r_P)) \cup (\cup_{Q \subset P} B_{g_k}(X_Q, r_Q)),$$

where $(r_P < 1)_{P \in \mathcal{P}}$ is a set of fixed constants satisfying $r_Q > r_P$ if $Q \subset P$. The set of center points also satisfies the condition that for any pair of polytopes $Q \subset P$:

$$\Lambda_{k,Q} \subset \Lambda_{k,P}, \quad \text{and} \quad x \in \Lambda_{k,P} \backslash \Lambda_{k,Q} \implies d_{g_k}(x, X_Q) > r_P.$$

Let Λ_k be the union $\cup_P \Lambda_{k,P}$. The set Λ_k is partitioned into subsets $I_1, \dots, I_N$ where N is independent of k while satisfying the following: For any pair

$$x \in \Lambda_{k,P} \backslash \cup_{Q \subset P} \Lambda_{k,Q}, \quad y \in \cup_{Q \subset P} \Lambda_{k,Q}$$

we have

$$x \in I_\alpha, y \in I_\beta \implies \beta < \alpha.$$

We define Gaussian sections concentrated at the lattice points so that they respect the (Restriction) condition. In particular we require that for any pair of polytopes $Q \subset P$, and $p \in \Lambda_{k,Q}$ the sections $\sigma_{k,p,P} : X_P \to \widetilde{X}_P^k$, $\sigma_{k,p,Q} : X_Q \to \widetilde{X}_Q^k$ satisfy

$$\sigma_{k,p,P} | X_Q = \sigma_{k,p,Q}.$$

To achieve this, for any top-dimensional polytope P, a point $p \in \Lambda_{k,P}$ and for all $Q \subseteq P$ such that $p \in \Lambda_{k,Q}$, Darboux coordinates $z_1^Q, \dots, z_{n(Q)}^Q$ (where $n(Q) := \frac{1}{2}(\dim(X_Q))$ are chosen such that for any $Q \subset P$,

- a subset of the coordinates $(z_1^P, \dots, z_{n(P)}^P)$, say $(z_{n(Q)+1}^P, \dots, z_{n(P)}^P)$ vanishes on X_Q;
- and the remaining coordinates $(z_1^P, \dots, z_{n(Q)}^P)|X_Q$ are a permutation of $(z_1^Q, \dots, z_{n(Q)}^Q)$.

Gaussian sections whose centers are close to relative divisors are multiplied by a factor that vanishes up to order two at the relative divisor. Consider a lattice point p and let P_0 be the smallest polytope that contains p. We know that $d_{g_k}(p, Q) \geq 1$ for all facets $Q \subset P_0$. Let $Q_1, \dots, Q_\alpha \subset P_0$ be facets for which

$$1 \leq d_{g_k}(p, Q_i) \leq k^{-1/6}, \quad \forall k > 0.$$

Since the $\omega_{\mathfrak{X}}$-compatible almost complex structure $\mathfrak{J}_0$ is integrable on the fibers of the map

$$\pi_{Q_i}^{P_0} : U_{X_{Q_i}}(X_{P_0}) \to X_{Q_i},$$

we may choose an approximately holomorphic function z_i in a neighborhood of p that vanishes on X_{Q_i} and $\overline{\partial} z_i(p) = 0$. Indeed, we may view $U_{X_{Q_i}}(X_{P_0}) \to X_{Q_i}$

as the neighborhood of the zero section of a unitary line bundle $L^{P_0}_{Q_i} \to X_{Q_i}$; we may trivialize the bundle in a neighborhood of $\pi^{P_0}_{Q_i}(p)$ and choose z_i to be the fiber coordinate. Further the function z_i satisfies uniform bounds as in (5.5) if the trivialization of $L^{P_0}_{Q_i}$ is chosen so that the Chern connection satisfies uniform bounds. The modified Gaussian section defined as

$$\sigma'_{k,p,P} = \sigma_{k,p,P} \cdot \prod_{i=1}^{\alpha} z_i^2 / z_i(p)^2$$

is asymptotically holomorphic and has a uniform lower bound on $B_{g_k}(p, r_p)$. (The second power ensures that not only does the factor $z_i^2/z_i(p)^2$ vanish along the relative divisor, but the derivative of $z_i^2/z_i(p)^2$ also vanishes, so that the resulting stabilizing divisor will be ω-orthogonal to the relative divisor.) If p is not close to any relative divisors then $\alpha = 0$, and the Gaussian section $\sigma_{k,p,P}$ at p is left unchanged. Such a section $\sigma_{k,p,P}$ vanishes in neighborhoods of relative divisors $X_{Q_1}, \dots, X_{Q_\alpha}$ because of the cutoff function $\rho_{k,p,P}$.

The globalization process consists of finding coefficients $\{w_p \in \mathbb{C}\}_{p \in \Lambda_k}$ such that

$$s_{w,P} := \sum_{p \in \Lambda_{k,P}} w_p \sigma'_{k,p,P}$$

is a uniformly transverse sequence of sections for each P. The proof of globalization carries over from [**26**]. The only new feature is to determine each coefficient w_p in a P-independent way. This can be done by Lemma 5.22 which is a modification of Lemma 5.19.

Finally we show that the zero set of the transverse section satisfies the (Restriction) and (Orthogonality with relative divisors) required by the Proposition. For a pair $Q \subset P$ of polytopes, the section $\sigma'_{w,p,P}$ vanishes near X_Q if $d_k(p, X_Q) > k^{-1/6}$. If for a lattice point $1 < d_k(p, X_Q) < k^{-1/6}$, then both the section $\sigma'_{k,p,P}$ and its derivative normal to X_Q vanishes along X_Q. Therefore, on X_Q, the section $s_{w,P}$ is equal to $s_{w,Q}$, and ω-orthogonality follows from that of the section

$$\sum_{p \in \Lambda_{k,Q}} w_p \sigma'_{k,p,P} = s_{w,Q} e^{|z|^2}.$$

□

LEMMA 5.22. (Quantitative Sard's theorem) *Given a tuple of positive integers $(n_1, \dots, n_k) \in \mathbb{Z}^k$ there is an integer p for which the following is satisfied. Suppose $0 < \delta < \frac{1}{4}$, and $f_i : B_+^{n_i} \to \mathbb{C}$ is a set of k functions on balls $B_+^{n_i} \subset \mathbb{C}^{n_i}$ of radius $\frac{11}{10}$ that satisfy $\|f_i\|_{C^1} \le \eta$, where $\eta := \delta \log(\delta^{-1})^{-p}$. Then, there exists $w \in \mathbb{C}$, $|w| \le \delta$ such that $f_i - w$ is η-transverse to 0 over the interior ball B^{n_i} of radius 1.*

The case $k = 1$ is Theorem 20 of [**26**], and is stated as Lemma 5.19. The proof in the case $k = 1$ is by bounding the size of the image $f(B_+)$ in the range. For a finite k, the volume is multiplied by a constant, and the proof in [**26**] can be replicated by altering the constants.

PROOF OF PROPOSITION 5.14. The proof is a consequence of Lemma 5.15 and Proposition 5.21. □

Remark 5.23. The construction of a broken divisor in Proposition 5.21 extends to the orbifold case. Indeed, in the orbifold adaptation of the Donaldson construction in Gironella-Muñoz-Zhou [**37**], the new features are the choice of an appropriate lattice compatible with the stratification of the orbifold; and adjusting the uniform transversality constants. Both these features are compatible with the modifications we have introduced in the proof of Proposition 5.21.

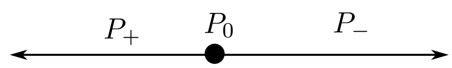

Figure 5.1. The polyhedral decomposition of a single cut.

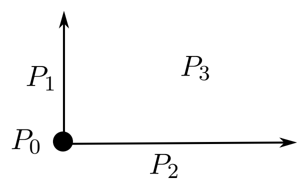

Figure 5.2. Multiple cut case

Remark 5.24. (On the construction of broken stabilizing divisors) This is a technical remark justifying the proof technique of Proposition 5.21 wherein we construct asymptotically holomorphic sequences of sections $\sigma_{k,P}$ simultaneously for all polytopes $P \in \mathcal{P}$ rather than construct one sequence $(\sigma_{k,P})_k$ at a time. We point out that the latter approach, which appears cleaner, is taken in [**18**] for the case of single breaking. Consider a single cut with polyhedral decomposition shown in Figure 5.1. Given a sequence of sections

$$(\sigma_{k,P_0} : X_{P_0} \to \tilde{X}_{P_0})_k,$$

[**18**, Lemma 7.15] constructs an extension of the sequence of sections to

$$(\sigma_{k,P_\pm} : X_{P_\pm} \to \tilde{X}_{P_\pm})_k.$$

The first step in the construction is to consider the extension $\sigma_{k,P_0}e^{-k|x|^2}$ and then turn on the contributions from Gaussians centered away from the divisor $X_{P_0} \subset X_{P_\pm}$ so that the sections become transverse. If we apply this approach to spaces corresponding to the polytopes shown in Figure 5.2, we would construct a sequence in the order σ_{k,P_0}, then σ_{k,P_1}, σ_{k,P_2} and then finally σ_{k,P_3}. We would like the sequence σ_{k,P_3} on X_{P_3} to be an extension of σ_{k,P_0} , σ_{k,P_1} and σ_{k,P_2}. Therefore, we would like to start Donaldson's globalization iteration with a sequence of Gaussian sections that are equal to $\sigma_{k,P_0}e^{-k(|z_1|^2+|z_2|^2)}$ in a neighborhood of X_{P_0}, equal to $\sigma_{k,P_1}e^{-k|z_1|^2}$ in a neighborhood of X_{P_1}, and equal to $\sigma_{k,P_2}e^{-k|z_2|^2}$ in a neighborhood of X_{P_2}. The approach fails because the three sections do not agree on overlaps. Indeed, to define the sequence σ_{k,P_1}, we would have used $\sigma_{k,P_0}e^{-k|z_2|^2}$ as a starting sequence, but then these sections would have been modified when contributions from Gaussians in nearby balls are turned on to achieve transversality.

5.3. Stabilizing pairs in neck-stretched manifolds

In this section we prove the existence of a stabilizing pair $(\mathfrak{J}_0, \mathfrak{D})$ on a broken manifold such that the family (J^ν, D^ν) obtained by gluing the pair $(\mathfrak{J}_0, \mathfrak{D})$ consists of stabilizing pairs on neck-stretched manifolds.

Definition 5.25. (Adapted broken almost complex structure) Given a broken divisor $\mathfrak{D} \subset \mathfrak{X}$, we denote by

$$\mathcal{J}^{\mathrm{cyl}}(\mathfrak{X}, \mathfrak{D}) = \{\mathfrak{J} | \mathfrak{J}(TD_{\overline{P}}) = TD_{\overline{P}} \text{ on } X_{\overline{P}} \ \ \forall P \in \mathcal{P}\}$$

the space of cylindrical almost complex structures that are adapted to $\mathfrak{D}$.

Definition 5.26. (Stabilizing pair in a broken manifold) Let $\mathfrak{D} \subset \mathfrak{X}$ be a broken cylindrical divisor which is disjoint from the Lagrangian $L \subset \mathfrak{X}$. For $E > 0$, an adapted almost complex structure $\mathfrak{J} \in \mathcal{J}^{\mathrm{cyl}}(\mathfrak{X}, \mathfrak{D})$, is *$E$-stabilized* if for every polytope $P \in \mathcal{P}$, the almost complex structure $\mathfrak{J}|X_P$ is E-stabilized in (X_P, ω_{X_P}, D_P). A pair $(\mathfrak{J}, \mathfrak{D})$ is *stabilizing* if $\mathfrak{J}$ is locally strongly tamed and it is E-stabilizing for all $E > 0$.

The existence theorem for stabilizing pairs extends to the broken case.

Proposition 5.27. (Stabilizing pair in a broken manifold) *Suppose $\mathfrak{X}_{\mathcal{P}}$ is a broken manifold, such that on each cut space X_P, $P \in \mathcal{P}$ the symplectic form ω_{X_P} is rational. Suppose $J^{\mathrm{pre}} \in \mathcal{J}(\mathfrak{X})$ is a $\omega_{\mathfrak{X}}$-strongly compatible almost complex structure. Then, there exists a stabilizing pair $(\mathfrak{J}_0, \mathfrak{D})$ in $\mathfrak{X}$ such that*

(a) the family (J^ν, D^ν) obtained by gluing consists of stabilizing pairs for all $\nu \in [1, \infty)$,
(b) and for every $E > 0$ there is a neighbourhood

$$\mathcal{J}^{\mathrm{cyl}}(\mathfrak{X}, \mathfrak{D}; \mathfrak{J}_0, E) \subset \{\mathcal{J} \in \mathcal{J}^{\mathrm{cyl}}(\mathfrak{X}, \mathfrak{D}) : \mathcal{J}|T\mathfrak{D} = \mathfrak{J}_0|T\mathfrak{D}\}$$

of $\mathfrak{J}_0$ consisting of E-stabilizing locally tamed cylindrical almost complex structures adapted to $\mathfrak{D}$.

We recall that to construct a stabilizing pair in the unbroken case, we produced a degree bound for stabilization that was applicable in an ϵ-neighborhood of a fixed almost complex structure J_0. However, in the broken case the preliminary fixed almost complex structure $\mathfrak{J}^{\mathrm{pre}}$ is an $\omega_{\mathfrak{X}}$-compatible almost complex structure, whereas we need to prove a degree bound in stabilization for cylindrical almost complex structures. One cannot expect to find cylindrical almost complex structures arbitrarily close to $\omega_{\mathfrak{X}}$-compatible ones because the underlying connection one-forms and $\mathfrak{t}$-inner product are different for the two types of almost complex structures. Our strategy is to

(a) construct a cylindrical broken divisor $\mathfrak{D}$ using an $\omega_{\mathfrak{X}}$-compatible $\mathfrak{J}^{\mathrm{pre}}$,
(b) find a $\omega_{\mathfrak{X}}$-strongly tamed $\mathfrak{D}$-adapted $\mathfrak{J}_0$ close to $\mathfrak{J}^{\mathrm{pre}}$,
(c) construct a cylindrical $\mathfrak{D}$-adapted $\mathfrak{J}_1 \in \mathfrak{J}^{\mathrm{cyl}}(\mathfrak{X}, \mathfrak{D})$ that has the same base projections as $\mathfrak{J}_0$ on each of the cylindrical ends,

(d) and choose a degree bound for a neighborhood of $\mathfrak{J}^{\mathrm{pre}}$ that is also a degree bound for the neck-stretched almost complex structures J_1^ν obtained by gluing $\mathfrak{J}_1$ along the necks.

Step (c) in the above construction is carried out in the following Lemma:

Lemma 5.28. *Given a $\omega_{\mathfrak{X}}$-strongly tame almost complex structure $\mathfrak{J}_0$ that is cylindrical on the neighborhoods*

$$\underline{U} = (U_{X_Q}(X_P))_{Q\subset P}, \quad U_{X_Q}(X_P) \subset X_P,$$

there exists a cylindrical locally strongly tame almost complex structure $\mathfrak{J}_1$ on $\mathfrak{X}$

(a) whose cylindrical neighborhoods are

$$\underline{U}' \subset \underline{U};$$

(b) for any pair of polytopes $Q \subset P$, there exist projection maps

$$\pi_Q'^P : U_{X_Q}(X_P) \to X_Q \tag{5.6}$$

with $\mathfrak{J}_1$-holomorphic fibers;

(c) for any $Q \in \mathcal{P}\backslash\mathcal{P}^0$, $x \in X_Q\backslash\underline{U}$ and $y \in \pi_Q'^P(x)$, the projection $(d\pi_Q'^P)_y$ is $(\mathfrak{J}_1, \mathfrak{J}_0)$-holomorphic, and for any $P \in \mathcal{P}^0$, $\mathfrak{J}_0 = \mathfrak{J}_1$ on $X_P\backslash\underline{U}$.

Proof of Lemma 5.28. We alter the cylindrical structure maps in the neighborhoods of boundary submanifolds so that the almost complex structures on the pieces of $\mathfrak{X}$ can be glued along the necks. Recall that for any pair of polytopes $Q \subset P$ there is a symplectic embedding of the neighborhood $U_{X_Q}(X_P)$ of X_Q to a model neighborhood, namely

$$\phi_Q^P : U_{X_Q}(X_P) \to (\mathrm{Cone}_Q(P) \times Z_Q, \omega_{X_Q} + d\langle \alpha_Q^P, \pi_{\mathfrak{t}_Q^\vee}\rangle)/\sim \tag{5.7}$$

where the terms are as in (3.43). The image of ϕ_Q^P is

$$U(\mathrm{Cone}_Q P) \times Z_Q,$$

which is a neighborhood of $\{0\} \times Z_Q$. In order to be able to glue the cylindrical ends, we need the Q-connection one-forms α_Q^P to be P-independent for all $P \supset Q$, and also the $\mathfrak{t}$-inner product underlying the cylindrical almost complex structure to be the one fixed in (3.6). Let $(\alpha_P)_{P\in\mathcal{P}}$ be a collection of connection one-forms $\alpha_P \in \Omega^1(Z_P, \mathfrak{t}_P)$ that are consistent in the sense of (3.16) with respect to the standard $\mathfrak{t}$-inner product. We define a family of connection one-forms interpolating between $(\alpha_Q^P)_{Q\subset P}$ and $(\alpha_P)_P$, namely let

$$\beta_Q^P : U(\mathrm{Cone}_Q P) \to \{\tau\alpha_Q^P + (1-\tau)\alpha_Q : \tau \in [0,1]\} \subset \Omega^1(Z_Q, \mathfrak{t}_Q)$$

be a family of connection one-forms that is equal to α_Q in a neighborhood of the origin

$$0 \in U'(\mathrm{Cone}_Q(P)) \subset U(\mathrm{Cone}_Q P),$$

and equal to α_Q^P in a neighborhood of the boundary $\partial U(\mathrm{Cone}_Q P)\backslash(\partial\,\mathrm{Cone}_Q P)$, such that the two-forms

$$\omega_{X_Q} + d\langle \beta_Q^P(\pi_{\mathfrak{t}_Q^\vee}), \pi_{\mathfrak{t}_Q^\vee}\rangle$$

are non-degenerate on $U(\mathrm{Cone}_Q P) \times Z_Q$. The non-degeneracy can be ensured if $\beta_Q^P \neq \alpha_Q^P$ in a small enough neighborhood of $\{0\} \times Z_Q$. There is a symplectomorphism

$$\phi_Q'^P : U_{X_Q}(X_P) \to (\mathrm{Cone}_Q(P) \times Z_Q, \omega_{X_Q} + d\langle \beta_Q^P(\pi_{\mathfrak{t}_Q^\vee}), \pi_{\mathfrak{t}_Q^\vee}\rangle)/\sim \tag{5.8}$$

such that

$$\phi_Q^P = \phi_Q'^P \quad \text{near } \partial U_{X_Q}(X_P). \tag{5.9}$$

Symplectomorphisms $\phi_Q'^P$ satisfying (5.9) indeed exist by a Moser-type argument on $\mathrm{Cone}_Q P \times Z_Q$ equipped with the symplectic forms in (5.7) and (5.8). Let

$$\pi_Q'^P : U_{X_Q}(X_P) \to X_Q$$

be the projection obtained by composing $\phi_Q'^P$ with the projection $Z_Q \to X_Q$. We note that

$$\pi_Q'^P = \pi_Q^P \quad \text{in a neighborhood of } \partial U_{X_Q}(X_P). \tag{5.10}$$

The cylindrical almost complex structure $\mathfrak{J}_1$ required by the lemma can now be defined. On

$$U'_{X_Q}(X_P) := (\phi_Q'^P)^{-1}(U'(\mathrm{Cone}_Q P) \times Z_Q) \tag{5.11}$$

the symplectic form is $\omega_{X_Q} + d\langle \alpha_Q, \pi_{\mathfrak{t}_Q^\vee}\rangle$, and so $U'_{X_Q}(X_P)$ will serve as the 'Q-cylindrical end' for $\mathfrak{J}_1$. We define $\mathfrak{J}_1 \in \mathcal{J}^{\mathrm{cyl}}(\mathfrak{X})$ so that on $U_{X_Q}(X_P)$, the fibers of $\pi_Q'^P$ are $\mathfrak{J}_1$-holomorphic and on the ω-complement, which is equal to $\ker(\beta_Q^P)$, $\mathfrak{J}_1$ is equal to the lift of $\mathfrak{J}_0|(X_Q \backslash \underline{U})$. The definitions of $\mathfrak{J}_1$ on cylindrical ends and their complement patch smoothly on the boundary. Indeed, in a neighborhood of $\partial U_{X_Q}(X_P)$, the projection maps and connection one-forms satisfy

$$\beta_Q^P = \alpha_Q^P, \quad \pi_Q^P = \pi_Q'^P,$$

and thus in this neighborhood $\mathfrak{J}_1|U_{X_Q}(X_P)$ is equal to $\mathfrak{J}_0$. □

Proof of Proposition 5.27. Let $\mathfrak{J}^{\mathrm{pre}}$ be an $\omega_{\mathfrak{X}}$-compatible almost complex structure on $\mathfrak{X}$.

Step 1: We first find a uniform degree bound on a neighborhood of $\mathfrak{J}^{\mathrm{pre}}$ that will also be a degree bound for glued almost complex structures on neck-stretched manifolds. To derive the bound we choose a representative of the first Chern class $[c_1(TX)]$ that is basic on the cylindrical ends. That is, we choose forms

$$\gamma = (\gamma_P)_{P \in \mathcal{P}}, \quad \gamma_P \in \Omega^2(X_P), \quad [\gamma_P] = c_1(TX_P),$$

which are 'basic' in the sense that for any pair $Q \subset P$ of polytopes, on the Q-cylindrical end $U_{X_Q}(X_P) \xrightarrow{\pi_Q^P} X_Q$, we have

$$\gamma_P = (\pi_Q^P)^* \gamma_Q.$$

These forms have the property that they glue along the necks to yield

$$\gamma_\nu \in \Omega^2(X^\nu), \quad \gamma_\nu \in [c_1(TX)],$$

and therefore, can be used to obtain uniform degree bounds for glued almost complex structures on X^ν for all ν. We choose $\epsilon > 0$. For any $P \in \mathcal{P}$, the supremum of the ratio of the first Chern class form and the basic symplectic form

$$k_*(\epsilon, P) := \sup_{\substack{\mathfrak{J}: \|\mathfrak{J}-\mathfrak{J}^{\text{pre}}\|<\epsilon \\ 0 \neq v \in TX_P}} \tfrac{\gamma_P(v, \mathfrak{J}v)}{\omega_{X_P}(v, \mathfrak{J}v)}$$

is finite by Lemma 5.5. Here the norm on $(\mathfrak{J} - \mathfrak{J}^{\text{pre}})$ is the C^0-bound with respect to the metric $\omega_{X_P}(\cdot, \mathfrak{J}^{\text{pre}}\cdot)$ on X_P. We choose k_* satisfying

$$k_* \geq k_*(\epsilon, P) \quad \forall P \in \mathcal{P}^0, \tag{5.12}$$

and one other condition stated below in (5.14).

Step 2 : We find a stabilizing divisor in the broken manifold. We will find a stabilizing divisor of degree

$$k > 2 \max\{k_*, k_* + \tfrac{1}{2}(\dim(X)) - 2\}. \tag{5.13}$$

for reasons explained in Remark 5.6. Let $\theta_0 > 0$ be such that for any θ_0-approximately $\mathfrak{J}^{\text{pre}}$-holomorphic divisor, there is an adapted tamed almost complex structure that is $\epsilon/2$-close to $\mathfrak{J}^{\text{pre}}$. In addition to satisfying (5.13), we may choose k to be large enough that a Donaldson divisor of degree k is θ_0-approximately holomorphic. Thus there is a divisor $\mathfrak{D}$ cylindrical in $\underline{U}$ such that there is a $\mathfrak{D}$-adapted $\omega_{\mathfrak{X}}$-strongly tamed

$$\mathfrak{J}_0 : \quad \text{such that } k_* \text{ is a degree bound for } \mathfrak{J}_0|X_P \text{ for all } P \in \mathcal{P}.$$

In particular, we have the pointwise bound

$$\gamma_P(v, \mathfrak{J}_0 v) \leq k_* \omega_{X_P}(v, \mathfrak{J}_0 v) \quad \forall v \in TX_P, \forall P \in \mathcal{P}.$$

Step 3: We deform $\mathfrak{J}_0$ to make it gluable along the neck and show that the resulting neck-stretched almost complex manifolds continue to satisfy the degree bound. The difficulty in this step is that the deformation required to make an almost complex structure gluable is not C^0-small, and the new almost complex structure may no longer satisfy the degree bound. To overcome the difficulty we use the fact that the representative of the first Chern class is basic, and the almost complex structure is deformed only in the fiber direction. The details are as follows: By Lemma 5.28 there exists $\mathfrak{J}_1 \in \mathcal{J}^{\text{cyl}}(\mathfrak{X})$ that is cylindrical on truncated cylindrical ends

$$\underline{U}' = (U'_{X_Q}(X_P))_{Q \subset P, P \in \mathcal{P}^0}, \quad U'_{X_Q}(X_P) \subset U_{X_Q}(X_P),$$

and the base almost complex structure of $\mathfrak{J}_1$ on $\underline{U}$ is equal to that of $\mathfrak{J}_0$. For any $\nu > 1$ let J_1^ν be the neck-stretched almost complex structure on X^ν obtained by gluing $\mathfrak{J}_1$. We claim that k_* is a degree bound for J_1^ν. Firstly we observe that changing the cylindrical structure has the effect of changing the projection maps π_Q^P to

$$\pi'^P_Q : U_{X_Q}(X_P) \to X_Q$$

for any pair $Q \subset P$, see Lemma 5.28. The π_Q^P-basic forms $(\gamma_P \in \Omega^2(X_P))_{P \in \mathcal{P}}$ representing the first Chern classes $c_1(TX_P)$ may be replaced by cohomologous forms $(\gamma'_P)_P$ that are basic with respect to π'^P_Q, without altering $\gamma_P|(X_P \backslash \underline{U})$ for any

$P \in \mathcal{P}$ (this replacement is possible because of (5.10)). The resulting glued forms are denoted by

$$\gamma'_\nu \in \Omega^2(X^\nu), \quad [\gamma'_\nu] \in c_1(TX).$$

A taming symplectic form on the neck-stretched manifold (X^ν, J^ν) is given by choosing an increasing map $\aleph : \nu B^\vee \to B^\vee$ between dual complexes (defined in Section 7.1), and the resulting form is called $\omega_\aleph$. Let $\pi^\nu_Q : X^\nu_{\underline{Q}} \to X_Q$ denote the projection map on the Q-cylindrical subset $X^\nu_{\underline{Q}}$ of X^ν. For any point $x \in X^\nu_{\underline{Q}}$, on the $\omega_\aleph$-complement of $\ker(d\pi^\nu_Q)_x$ the form $\omega_\aleph$ is

$$\omega_{X_Q} + c d\alpha_Q$$

for some constant $c \in Q^\vee$. To ensure that k_* is a degree bound for J^ν_1, we require that

$$k_* \geq \sup_{0 \neq v \in T(X_Q \setminus \underline{U}), c \in Q^\vee \subset \mathfrak{t}^\vee_Q} \tfrac{\gamma_Q(v, \mathfrak{J}_1 v)}{(\omega_{X_Q} + c d\alpha_Q)(v, \mathfrak{J}_1 v)} \tag{5.14}$$

for all $Q \in \mathcal{P} \backslash \mathcal{P}^0$. In Step 1, we can indeed choose k_* satisfying (5.14) because (5.14) does not depend on subsequent calculations. For a non-zero vector $v \in TX^\nu$ with projection $v_Q := d\pi^\nu_Q(v)$, we have

$$\gamma'_\nu(v, J^\nu_1 v) = \gamma'_Q(v_Q, \mathfrak{J}_1 v_Q) \leq k_*((\omega_{X_P} + c d\alpha^P_Q)(v, \mathfrak{J}_1 v)) \leq k_* \omega_\aleph(v, J^\nu_1 v) \tag{5.15}$$

using (5.14). Therefore, for a J^ν_1-holomorphic map $u : \mathbb{P}^1 \to X^\nu$, $\int_{\mathbb{P}^1} u^* \gamma'_\nu \leq k_* \int_{\mathbb{P}^1} u^* \omega_\aleph$, and consequently k_* is a degree bound for J^ν_1.

One can also obtain a uniform degree bound in a neighborhood of the neck-stretched almost complex structures J^ν_1. Let $k'_* > 0$ be slightly smaller than k_* so that that (5.13) is satisfied when k_* is replaced by k'_*, and there is a constant $\delta > 0$ such that k'_* is a degree bound for any $J^\nu \in \mathcal{J}^{\mathrm{cyl}}(X^\nu)$, $\nu \in [1, \infty)$, obtained by gluing $\mathfrak{J} \in U_{\mathfrak{J}_1}$. Here

$$\begin{aligned} U_{\mathfrak{J}_1} := \{ \mathfrak{J} \in \mathcal{J}^{\mathrm{cyl}}(\mathfrak{X}, \mathfrak{D}) : \mathfrak{J} &\text{ is locally strongly tamed}, \\ |d\pi'^P_Q(\mathfrak{J}) - d\pi'^P_Q(\mathfrak{J}_1)|_{C^0(X_Q \setminus \underline{U}')} &< \delta \ \ \forall Q \subset P, P \in \mathcal{P}^0, \\ |\mathfrak{J} - \mathfrak{J}_1|_{C^0(X_P \setminus \underline{U}')} &< \epsilon. \} \end{aligned}$$

Indeed it is enough to consider the distance between π'^P_Q projections of the almost complex structures to estimate the degree bound, because the representatives γ'_P, $P \in \mathcal{P}$ of the first Chern classes $c_1(TX_P)$ are basic on the cylindrical ends.

Step 4: We finish the proof closely following the method in the unbroken case in Cieliebak-Mohnke [**22**], which is outlined in Proposition 5.3. So far we have shown that there is a neighborhood $U_{\mathfrak{J}_1}$ in the space of locally strongly tamed $\mathfrak{D}$-adapted cylindrical almost complex structures on $\mathfrak{X}$ such that for any $\mathfrak{J} \in U_{\mathfrak{J}_1}$, k_* is a degree bound for the glued family J^ν. By Remark 5.6 generic almost complex structures in the set $U_{\mathfrak{J}_1}$ are stabilizing.

Next, we claim that for any $E > 0$, an open and dense subset of $U_{\mathfrak{J}_1}$ is E-stabilizing for all neck lengths. The subset

$$\mathcal{J}^{\mathrm{reg}, E, \infty} = \{ \mathfrak{J} \in U_{\mathfrak{J}_1} : \mathfrak{J} \text{ is } E\text{-stabilizing on } \mathfrak{X} \}$$

is open and dense: Openness is a consequence of Gromov convergence applied to each of the cut spaces X_P, $P \in \mathcal{P}$ (see the proof of [**22**, Corollary 8.16]), and denseness follows because the E-stabilizing condition is generic. The subset

$$\mathcal{J}^{\mathrm{reg},E} = \{\mathfrak{J} \in \mathcal{J}^{\mathrm{reg},E,\infty} : J^\nu \text{ is } E\text{-stabilizing on } X^\nu \text{ for } \nu \in [1,\infty)\}.$$

is also open in $\mathcal{J}^{\mathrm{reg},E,\infty}$ by Gromov compactness. Openness at the infinite neck length parameter is proved by Lemma 5.29. The subset $\mathcal{J}^{\mathrm{reg},E} \subset \mathcal{J}^{\mathrm{reg},E,\infty}$ is comeager by an application of Sard's theorem : The universal moduli space

$$\mathcal{M}_{\mathcal{J}} := \{u_\nu : \mathbb{P}^1 \to X^\nu : u_\nu \text{ is } J^\nu\text{-holomorphic}, \mathfrak{J} \in \mathcal{J}^{\mathrm{reg},E,\infty}, \nu \in [1,\infty)\}$$

is a Banach manifold and the projection $\pi_{\mathcal{J}} : \mathcal{M}_{\mathcal{J}} \to \mathcal{J}^{\mathrm{reg},E,\infty}$ is a Fredholm map. Therefore the subset $\mathcal{J}^{\mathrm{reg},E}$ of regular values of $\pi_{\mathcal{J}}$ is comeager.

Finally, let $E_k \to \infty$ be any sequence of real numbers with limit infinity. The set of almost complex structures

$$\mathcal{J}^{\mathrm{reg}} = \bigcap_{k=1}^{\infty} \mathcal{J}^{\mathrm{reg},E_k}$$

that is stabilizing for all $\nu \in [1,\infty]$ is the intersection of the set of E_k-stabilizing almost complex structures for all k. The intersection $\mathcal{J}^{\mathrm{reg}}$ is non-empty because of each of the sets in the intersection is open and dense. □

The following Lemma, used above in the proof of Proposition 5.21, is an openness statement for stabilizing almost complex structures at $\nu = \infty$.

LEMMA 5.29. *Suppose $\mathfrak{D} \subset \mathfrak{X}$ is a cylindrical broken divisor, and $\mathfrak{J} \in \mathcal{J}^{\mathrm{cyl}}(\mathfrak{X},\mathfrak{D})$ is a tamed adapted almost complex structure that is E-stabilizing. Suppose the divisors $D^\nu \subset X^\nu$ are obtained by gluing, and the sequence $J^\nu \in \mathcal{J}(X^\nu, D^\nu)$ converges to $\mathfrak{J}$. Then, there exists ν_0 such that J^ν is E-stabilizing for $\nu \geq \nu_0$.*

The proof is given at the end of Section 8.4 because it uses terminology introduced in Chapter 8. The proof of the Lemma however does not use any result proved in Chapter 8.

CHAPTER 6

Coherent perturbations and regularity

In order to obtain the necessary transversality for moduli spaces of treed holomorphic maps, we use the Cieliebak-Mohnke perturbation scheme [**22**] which has been adapted to define Fukaya algebras of Lagrangians by Charest-Woodward [**17**]. The Morse function on the Lagrangian submanifold and the almost complex structure on the broken manifold are allowed to be domain-dependent. Stabilizing divisors from the previous chapter are essential in defining domain-dependent perturbations.

A domain-dependent perturbation is defined as a map from a universal curve to the space of tamed almost complex structures. We recall from Chapter 4 that domain curves are treed disks. For any type Γ of treed disks the moduli space $\mathcal{M}_\Gamma$ of treed disks has a universal curve

$$\mathcal{U}_\Gamma \to \mathcal{M}_\Gamma$$

whose fiber over a point $[C] \in \mathcal{M}_\Gamma$ is isomorphic to C. The perturbation datum is then a map from $\mathcal{U}_\Gamma$ to the space of cylindrical almost complex structures on the broken manifold and Morse functions on the Lagrangian submanifold. Under this perturbation scheme, marked points on the domain curve are the inverse images of the intersection of the holomorphic map with the stabilizing divisor. The stabilizing divisor is Poincaré dual to a large multiple of the symplectic form. This ensures that generically there are no holomorphic spheres in the divisor, and that any non-constant holomorphic sphere has enough divisor intersections to ensure that its domain is stable.

Domain-dependent perturbations are necessary for solving the 'multiple cover problem'. Indeed, for a fixed almost complex structure the compactification of the moduli space of pseudoholomorphic maps may contain nodal maps with components that are multiple covers of maps with negative Chern class. It is not possible to achieve transversality for the moduli space of such maps. Multiple covers can be perturbed away by domain-dependent perturbations. As a toy example, suppose for a domain-independent almost complex structure J_0, there is a simple J_0-holomorphic sphere $S \subset X$ with negative Chern number and one intersection point with the stabilizing divisor. Via a domain-dependent perturbation of J_0, we can ensure that an n-covered sphere ($n > 1$) homologous to $n[S]$ does not occur in the moduli space of perturbed holomorphic maps. Indeed for such a cover the domain would have n marked points labelled $z_1, \ldots, z_n$, and the perturbation is not required to be invariant under automorphisms of the domain curve that permute the marked points. Breaking the permutation symmetry adds a multiplicative factor of $\frac{1}{n!}$ to the curve count (see Remark 6.16).

6.1. Domain-dependent perturbations

We first fix subsets of the universal treed disk outside of which perturbations vanish. Let Γ be a combinatorial type of treed disk and $\overline{\mathcal{U}}_\Gamma = \overline{\mathcal{S}}_\Gamma \cup \overline{\mathcal{T}}_\Gamma$ its universal curve from (4.6). Fix a compact subset

$$\overline{\mathcal{T}}_\Gamma^{\mathrm{cp}} \subset \overline{\mathcal{T}}_\Gamma$$

in the complement of breaking points and disk nodes (that is, disk nodes w_e where the length $\ell(e)$ of the treed segment is zero), and containing in its interior, for every edge $e \in \mathrm{Edge}_\circ(\Gamma)$ and every curve $C \subset \overline{\mathcal{U}}_\Gamma$, at least one point $z \in T_e \subset C$ on any infinite segment. Also fix a compact subset

$$\overline{\mathcal{S}}_\Gamma^{\mathrm{cp}} \subset \overline{\mathcal{S}}_\Gamma - \{w_e \in \overline{\mathcal{S}}_\Gamma,\ e \in \mathrm{Edge}_-(\Gamma)\}$$

disjoint from the boundary and spherical nodes $w(e) \in C, e \in \mathrm{Edge}(\Gamma)$, containing in its interior at least one point $z \in S_v$ on each sphere and disk component $S_v \subset C$ for each fiber $C \subset \overline{\mathcal{U}}_\Gamma$. Furthermore, the complement $\overline{\mathcal{S}}_\Gamma - \overline{\mathcal{S}}_\Gamma^{\mathrm{cp}} \subset \overline{\mathcal{S}}_\Gamma$ is a neighbourhood of the boundary and nodes; these neighborhoods must be chosen compatibly with those already chosen on the boundary for the inductive construction later.

Following Floer [**32**], we use a C^ε-topology on the space of almost complex structures. For a section ξ of a vector bundle $E \to X$, the C^ε-norm is

$$\|\xi\|_{C^\varepsilon} := \textstyle\sum_{k=0}^\infty \varepsilon_k \|\xi\|_{C^k(X,E)}.$$

Here $\varepsilon = (\varepsilon_i)_{i\in\mathbb{N}}$ is a fixed sequence of positive numbers that converges fast enough to 0 as $i \to \infty$. If the convergence is sufficiently rapid, then the space of sections with a bounded norm is a Banach space [**32**, Lemma 5.1] and contains sections supported in arbitrarily small neighbourhoods of X.

Definition 6.1. (a) (Domain-dependent Morse functions) Suppose that Γ is a type of stable treed disk, and $\overline{\mathcal{T}}_\Gamma \subset \overline{\mathcal{U}}_\Gamma$ is the tree part of the universal treed disk. Let

$$(F : L \to \mathbb{R}, G : T^{\otimes 2}L \to \mathbb{R})$$

be a Morse-Smale pair. For an integer $l \geq 0$ a *domain-dependent perturbation* of F of class C^l is a C^l map

$$F_\Gamma : \overline{\mathcal{T}}_\Gamma \times L \to \mathbb{R} \tag{6.1}$$

equal to the given function F away from the compact part:

$$F_\Gamma | (\overline{\mathcal{T}}_\Gamma - \overline{\mathcal{T}}_\Gamma^{\mathrm{cp}}) = \pi_2^* F$$

where π_2 is the projection on the second factor in (6.1). Here F is called the *background Morse function* for the domain-dependent perturbation F_Γ. The set of critical points of the background Morse function is denoted by

$$\mathcal{I}(L) := \mathrm{crit}(F). \tag{6.2}$$

(b) (Domain-dependent almost complex structure) Let $J_0 \in \mathcal{J}^{\mathrm{cyl}}(\mathfrak{X})$ be a locally strongly tamed cylindrical almost complex structure. A *domain-dependent* *almost complex structure* of class C^ε for treed disks of type Γ with *background almost complex structure* J_0 is a map from the two-dimensional part $\overline{\mathcal{S}}_\Gamma$ of the universal curve $\overline{\mathcal{U}}_\Gamma$ to $\mathcal{J}^{\text{cyl}}(X)$ given by

$$J_\Gamma : \overline{\mathcal{S}}_\Gamma \to \mathcal{J}^{\text{cyl}}(\mathfrak{X}) \tag{6.3}$$

equal to the given J_0 away from the compact part:

$$J_\Gamma|(\overline{\mathcal{S}}_\Gamma - \overline{\mathcal{S}}_\Gamma^{\text{cp}}) = J_0, \tag{6.4}$$

and for any fiber $S_\Gamma \subset \overline{\mathcal{S}}_\Gamma$ $J_\Gamma - J_0$ has finite norm in $C^\varepsilon(S_\Gamma \times \mathfrak{X}, \text{End}(T\mathfrak{X}))$.

Definition 6.2. (Perturbation data) A perturbation datum for a type Γ of stable treed disks is a pair $\mathfrak{p}_\Gamma = (F_\Gamma, J_\Gamma)$ consisting of a domain-dependent Morse function F_Γ and a domain-dependent almost complex structure J_Γ.

The following morphisms on the set of combinatorial types of stable treed disks are used to define coherence of perturbations.

Definition 6.3. (Morphisms on treed disk types)

(a) (Cutting edges) There is a (Cutting edges) morphism $\Gamma \to \Gamma'$ between combinatorial types Γ, Γ' of stable treed disks iff Γ is obtained by cutting a boundary edge $e \in \text{Edge}_{\circ,-}$ of Γ' that contains a breaking, see Figure 6.1.
(b) (Collapsing edges) A morphism $\Gamma \to \Gamma'$ is a (Collapsing edges) morphism if Γ' is obtained from Γ by collapsing an interior edge $e \in \text{Edge}_{\bullet,-}(\Gamma)$ or a boundary edge $e \in \text{Edge}_{\circ,-}(\Gamma)$ with length zero, $\ell(e) = 0$.
(c) (Making an edge length finite or non-zero) A morphism $\Gamma \to \Gamma'$ is a (Making an edge length finite or non-zero) morphism if Γ' is obtained from Γ by changing the edge length of a boundary edge $e \in \text{Edge}_{\circ,-}(\Gamma)$ from infinite or zero to finite non-zero.

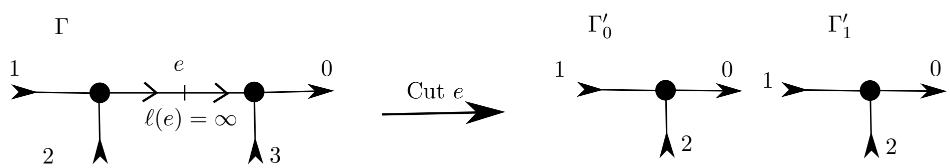

Figure 6.1. Cutting an edge e relabels the boundary and interior markings while preserving their ordering on the pieces Γ'_0, Γ'_1.

Perturbation datum $\mathfrak{p}_\Gamma$ for various treed disk types Γ can not be just chosen independently of each other. Rather for any two treed types Γ, Γ' related by a morphism, there is a coherence condition relating $\mathfrak{p}_\Gamma$, $\mathfrak{p}_{\Gamma'}$.

Definition 6.4. (Morphisms of perturbation data)

(a) (Collapsing edges/Making an edge length finite/non-zero) Let $\Gamma' \neq \Gamma$ be (Collapsing edges/making an edge length finite/non-zero) morphism $\Gamma' \to \Gamma$. For perturbation data $\mathfrak{p}_{\Gamma'}$, $\mathfrak{p}_\Gamma$ there is a morphism $\mathfrak{p}_{\Gamma'} \to \mathfrak{p}_\Gamma$ if $\mathfrak{p}_{\Gamma'}$ is induced by pullback of $\mathfrak{p}_\Gamma$ under the natural inclusion of the universal curve

$$\iota_\Gamma^{\Gamma'} : \overline{\mathcal{U}}_{\Gamma'} \to \overline{\mathcal{U}}_\Gamma.$$

(b) Suppose $\Gamma \to \Gamma'$ is a (Cutting edges) morphism, where an edge $e \in \mathrm{Edge}_{\circ,-}(\Gamma)$ is cut to yield leaf edges $e_+, e_- \in \mathrm{Edge}(\Gamma')$. There is a morphism of perturbation data $\mathfrak{p}_\Gamma \to \mathfrak{p}_{\Gamma'}$ if $\mathfrak{p}_{\Gamma'}$ is obtained by pushing forward $\mathfrak{p}_\Gamma$ under the map

$$\pi_\Gamma^{\Gamma'} : \overline{\mathcal{U}}_\Gamma \to \overline{\mathcal{U}}_{\Gamma'}$$

defined by gluing at the leaf edges $e_\pm$ to form a single non-leaf edge e with $\ell(e) = \infty$. That is, define

$$J_{\Gamma'}(z', x) = J_\Gamma(z, x), \quad \forall z \in (\pi_\Gamma^{\Gamma'})^{-1}(z).$$

The definition for the perturbed Morse datum $F_{\Gamma'}$ is similar.

We are now ready to define coherent collections of perturbation data. These are data that behave well with each type of operation in Definition 6.3.

Definition 6.5. (Coherent families of perturbation data) A collection of perturbation data $\underline{\mathfrak{p}} = (\mathfrak{p}_\Gamma)_\Gamma$ is *coherent* if it is compatible with the morphisms of moduli spaces of different types in the sense that

(a) (Collapsing edges/making an edge length finite/non-zero) if Γ is obtained from Γ' by collapsing an edge or making an edge finite/non-zero, then $\mathfrak{p}_{\Gamma'}$ is the pullback of $\mathfrak{p}_\Gamma$;

(b) (Cutting edges) if Γ is obtained from Γ' by cutting a boundary edge $e \in \mathrm{Edge}_{\circ,-}^{\infty}(\Gamma')$ of infinite length, then $\mathfrak{p}_{\Gamma'}$ is the push-forward of $\mathfrak{p}_\Gamma$. Assuming Γ is the union of types Γ_1, Γ_2, $\mathfrak{p}_\Gamma$ is obtained from $\mathfrak{p}_{\Gamma_1}$ and $\mathfrak{p}_{\Gamma_2}$ as follows: For $k = 1, 2$, let

$$\pi_k : \overline{\mathcal{M}}_\Gamma \cong \overline{\mathcal{M}}_{\Gamma_1} \times \overline{\mathcal{M}}_{\Gamma_2} \to \overline{\mathcal{M}}_{\Gamma_k}$$

denote the projection on the k-th factor, and therefore, $\overline{\mathcal{U}}_\Gamma$ is the union of $\pi_1^*\overline{\mathcal{U}}_{\Gamma_1}$ and $\pi_2^*\overline{\mathcal{U}}_{\Gamma_2}$. We require that $\mathfrak{p}_\Gamma$ is equal to the pullback of $\mathfrak{p}_{\Gamma_k}$ on $\pi_k^*\overline{\mathcal{U}}_{\Gamma_k}$:

$$\mathfrak{p}_\Gamma|\overline{\mathcal{U}}_{\Gamma_k} = \pi_k^*\mathfrak{p}_{\Gamma_k}. \tag{6.5}$$

We also require the perturbation data to satisfy the following locality axiom which ensures that the perturbations on any component only depend on special points on that component, and the length of the treed segments on the boundary of the disk. We first set up some notation: For a type Γ underlying treed disks, and a vertex $v \in \mathrm{Vert}(\Gamma)$, let $\Gamma(v)$ be a graph with a single vertex v and markings

$$\{e \in \mathrm{Edge}(\Gamma) : v \in e\}.$$

corresponding to each edge incident on v. Let $\mathcal{U}_{\Gamma,v} \subset \mathcal{U}_\Gamma$ be a fibration over $\mathcal{M}_\Gamma$ whose fiber over $m \in \mathcal{M}_\Gamma$ consists of the curve component represented by v. Define a map

$$\pi_v : \mathcal{U}_{\Gamma,v} \to \mathcal{U}_{\Gamma(v)} \times ([0,\infty])^{|\mathrm{Edge}_{\circ,-}(\Gamma)|}, \tag{6.6}$$

whose first component $\mathcal{U}_{\Gamma,v} \to \mathcal{U}_{\Gamma(v)}$ is the natural projection map, and the second component is the length function on boundary edges $e \in \mathrm{Edge}_{\circ,-}(\Gamma)$.

(Locality Axiom) The restriction of the perturbation datum $\mathfrak{p}_\Gamma$ to $\mathcal{U}_{\Gamma,v}$ is the pullback via π_v of some datum on $\mathcal{U}_{\Gamma(v)} \times ([0,\infty])^{|\operatorname{Edge}_{\circ,-}(\Gamma)|}$.

This ends the Definition.

Let C be a possibly unstable treed disk of type Γ. The *stabilization* of C is the stable treed disk $\operatorname{st}(C)$ of some type $\operatorname{st}(\Gamma)$ obtained by collapsing unstable surface and tree components. Thus the stabilization $\operatorname{st}(C)$ of any treed disk C is the fiber of a universal treed disk $\mathcal{U}_{\operatorname{st}(\Gamma)}$. Given perturbation datum for the type $\operatorname{st}(\Gamma)$, we obtain a domain-dependent almost complex structure and Morse function for C, still denoted J_Γ, F_Γ, by pull-back under the map $C \to \mathcal{U}_{\operatorname{st}(\Gamma)}$. If Γ does not contain vertices, i.e. if C is a single infinite segment $T_e, e \in \operatorname{Edge}(\Gamma)$, then the perturbation $\mathfrak{p}_\Gamma$ vanishes on C.

Remark 6.6. Coherence conditions for perturbation data are motivated as follows.

(a) If a sequence C_ν of (domain) curves of type Γ converges to a limit curve C of type Γ', we would like the perturbation data on C_ν to converge to the perturbation datum on C. If $\Gamma' \neq \Gamma$ then there is a (Collapsing edges/making an edge length finite/non-zero) morphism $\Gamma' \to \Gamma$, and the convergence of the perturbation datum is ensured by coherence under (Collapsing edges/making an edge length finite/non-zero).

(b) Suppose $\Gamma \to \Gamma'$ is a cutting edges morphism, and suppose the disconnected type Γ' has components Γ'_0, Γ'_1. Once moduli spaces of perturbed holomorphic maps in a symplectic manifold X are defined, we would like to be able to say that the moduli space $\mathcal{M}_\Gamma(X)$ of maps with domain type Γ is a product

$$\mathcal{M}_\Gamma(X) = \mathcal{M}_{\Gamma'_0}(X) \times \mathcal{M}_{\Gamma'_1}(X),$$

The product relation would hold only if the perturbation data is coherent under (Cutting edges).

Remark 6.7. (On the locality axiom) The locality axiom ensures that forgetting a marking $z_{e'}$ on a treed curve C affects the perturbation datum only on the component containing $z_{e'}$. This feature is used in Proposition 8.43. The dependence on the boundary edge lengths is useful for the following reason. Suppose Γ is a combinatorial type of treed disk depicted in Figure 6.1. By cutting an edge $e \in \operatorname{Edge}(\Gamma)$, we obtain two identical types. The cutting edge axiom requires that a coherent perturbation datum $\mathfrak{p}_\Gamma$ for Γ is equal after restriction to the universal curve for the type Γ' on both sides of the edge e. If in the locality axiom, the perturbation is $\mathfrak{p}_\Gamma$ defined by pulling back by the map $\pi_v : \mathcal{U}_{\Gamma,v} \to \mathcal{U}_{\Gamma(v)}$, then the perturbation datum on both surface components will be required to be equal even when the edge length of e is finite. This creates a problem, because in order to obtain transversality in the case $\ell(e) = 0$, we need the perturbation datum on both surface components to be independent of each other.

Remark 6.8. (Perturbations for moduli spaces of spheres) The coherence conditions required to define moduli spaces of spheres are much simpler than those required for treed disks: The only coherence condition is (Collapsing of edges), which translates to the statement that the perturbation datum is continuous on the compactified moduli space $\overline{\mathcal{M}}_n$ of n-marked spheres. An important difference from the disk case is that the perturbation datum on $\overline{\mathcal{M}}_n$ can be defined independently of $\overline{\mathcal{M}}_k$ for any $k < n$.

6.2. Perturbed maps

Given domain-dependent perturbations as in the previous section, the equations defining the moduli space of pseudoholomorphic maps are perturbed as follows.

Definition 6.9. (Perturbed pseudoholomorphic treed disks) Given a coherent perturbation datum $\mathfrak{p} = \{\mathfrak{p}_\Gamma\}_\Gamma$, and $\mathfrak{p}_\Gamma = (J_\Gamma, F_\Gamma)$, a pseudoholomorphic treed disk in X with boundary in L consists of

(a) a treed disk $C = S \cup T$ with stabilized type $\Gamma := \operatorname{st}(C)$,
(b) and continuous maps
$$u : C \to X,$$

such that the following hold:

(a) (Boundary condition) The tree components and the boundary of the surface components are mapped to the Lagrangian submanifold :
$$u(\partial S \cup T) \subset L.$$
(b) (Surface equation) On the surface part S of C the map u is J-pseudoholomorphic for the given domain-dependent almost complex structure: if j denotes the complex structure on S then
$$J_\Gamma(z, u(z))\, \mathrm{d}u_S = \mathrm{d}u_S\, j.$$
(c) (Boundary tree equation) On the boundary tree part $T \subset C$ the map u is a collection of gradient trajectories:
$$\frac{d}{ds} u_T = -\operatorname{grad}_{F_\Gamma(s,u(s))} u_T$$
where s is a coordinate on the segment so that the segment has the given length. Thus for each treed edge $e \in \operatorname{Edge}_{\circ,-}(\Gamma)$ the length of the trajectory $u|_{T_e}$ is equal to $\ell(e)$.

Definition 6.10. (Perturbed pseudoholomorphic broken treed disks) Given a coherent perturbation datum $\mathfrak{p} = \{\mathfrak{p}_\Gamma\}_\Gamma$ and a stable type Γ, a holomorphic broken treed disk in a broken manifold $\mathfrak{X}_\mathcal{P}$ with boundary in $L \subset X_{P_0}$, $P_0 \in \mathcal{P}$, consists of

(a) a treed disk $C = S \cup T$ of type Γ;
(b) a tropical structure $\mathcal{T}$ on Γ;
(c) a collection of maps on surface components
$$u_v : S_v \to X_{P(v)}, \quad v \in \operatorname{Vert}_\circ(\Gamma)$$
that are J_Γ-holomorphic, satisfy the edge matching condition (4.18) at all interior edges $e \in \operatorname{Edge}_\bullet(\mathcal{T})$ and the Lagrangian boundary condition

$u_v((\partial C)_v) \subset L$, and F_Γ-gradient flow lines

$$u_e : T_e \to L \subset X_{P_0}, \quad e \in \mathrm{Edge}_\circ(\Gamma)$$

on the treed parts.

In case Γ is an unstable type, a broken disk is defined similarly except that u_v, u_e are $\mathcal{P}_{\mathrm{st}\,\Gamma}$-holomorphic, where $\mathrm{st}\,\Gamma$ is the stabilization of Γ.

REMARK 6.11. Our domains below are stable on the surface parts because of the additional markings arising from the Donaldson hypersurface. So the content of the last sentence is that on broken Morse trajectories, for segments that are infinite in both directions the gradient flow equation is unperturbed.

We will not introduce notation for the moduli space of broken maps, but instead work directly with the regularized moduli spaces using Donaldson hypersurfaces and domain-dependent almost complex structures.

DEFINITION 6.12. (Perturbations adapted to a stabilizing divisor) Let $k \gg 0$, and $(\mathfrak{J}_0, \mathfrak{D})$ be a stabilizing pair on the broken manifold $\mathfrak{X}$ (as in Definition 5.26), such that $D_P \subset X_P$ is dual to $k[\omega_{X_P}]$ for all polytopes $P \in \mathcal{P}$. Further let $\mathfrak{D}$ be disjoint from the Lagrangian submanifold L. Suppose Γ is a type of treed disk. A perturbation datum $\mathfrak{p}_\Gamma = (J_\Gamma, F_\Gamma)$ on $(\mathfrak{X}, L)$ is adapted to the pair $(\mathfrak{J}_0, \mathfrak{D})$ if

- $\mathfrak{J}_0$ is the background almost complex structure for J_Γ, meaning that each J_Γ is a domain-dependent perturbation of $\mathfrak{J}_0$
- and for any treed curve $C = S \cup T$, and a connected component $S' \subset S$ with $d_\bullet(S')$ interior markings,

$$J_\Gamma(S') \subset \mathcal{J}^{\mathrm{cyl}}(\mathfrak{X}, \mathfrak{D}; \mathfrak{J}_0, \tfrac{1}{k} d_\bullet(S')) \cap \mathcal{U}_{\mathfrak{J}_0}. \tag{6.7}$$

  Here

$$\mathcal{J}^{\mathrm{cyl}}(\mathfrak{X}, \mathfrak{D}; \mathfrak{J}_0, \tfrac{1}{k} d_\bullet(S')) \subset \{\mathfrak{J} \in \mathcal{J}^{\mathrm{cyl}}(\mathfrak{X}, \mathfrak{D}) : \mathfrak{J}|T\mathfrak{D} = \mathfrak{J}_0|T\mathfrak{D}\}$$

  is the neighbourhood of $\mathfrak{J}_0$ consisting of $\frac{1}{k} d_\bullet(S')$-stabilizing almost complex structures (see Proposition 5.27), and $U_{\mathfrak{J}_0} \subset \mathcal{J}^{\mathrm{cyl}}(\mathfrak{X})$ is a neighborhood of $\mathfrak{J}_0$ on which the results on Hofer energy hold (see Lemma 7.23).

The set of perturbation data adapted to $(\mathfrak{J}_0, \mathfrak{D})$ is denoted by $\mathcal{P}_\Gamma(\mathfrak{X}, \mathfrak{J}_0, \mathfrak{D})$.

REMARK 6.13. We explain why we use $\frac{1}{k} d_\bullet(S')$-stabilizing almost complex structures in (6.7). Let $\underline{J} = (J_\Gamma)_\Gamma$ be a collection of coherent domain-dependent almost complex structures satisfying (6.7). Let S_ν be a sequence of nodal curves of type Γ with $d(\bullet)$ markings. Then the limit u of any sequence $u_\nu : S_\nu \to \mathfrak{X}$ of J_Γ-holomorphic broken maps can not have an unstable domain component. Indeed, the area of the maps u_ν is $\frac{d_\bullet}{k}$, and so an unstable component u_v of the u has area $\leq \frac{d_\bullet}{k}$. Therefore u is holomorphic with respect to a domain-independent almost complex structure $\mathfrak{J}_v$ which, by (6.7), is $\frac{d_\bullet}{k}$-stabilizing. This contradicts the existence of a non-constant map u_v.

6.3. Adapted maps

We consider *adapted* holomorphic maps whose domains are equipped with markings that are required to map to a Donaldson divisor.

DEFINITION 6.14. (Adapted stable treed disks) A stable broken treed holomorphic disk $u : (C, \partial C) \to (\mathfrak{X}, L)$ is *adapted* to a divisor $\mathfrak{D}$ iff each interior marking $z_e, e \in \mathrm{Edge}_{\bullet,\to}(\Gamma)$ maps to D under u and each connected component C' of $u^{-1}(\mathfrak{D}) \subset C$ contains an interior marking.

The moduli space of treed holomorphic disks is stratified by combinatorial type.

DEFINITION 6.15. The *combinatorial type of an adapted holomorphic broken treed disk* $u : C \to X$ adapted to a divisor $D \subset X - L$ consists of

(a) the combinatorial type Γ of its domain C,

(b) the tropical structure on Γ, which consists of an assignment of polytopes for vertices, and slopes for tropical edges :

$$\mathrm{Vert}(\Gamma) \ni v \mapsto P(v) \in \mathcal{P}, \quad \mathrm{Edge}_{\mathrm{trop}}(\Gamma) \ni e \mapsto \mathcal{T}(e) \in \mathfrak{t}_{P(e),\mathbb{Z}};$$

(c) a labelling

$$d : \mathrm{Vert}(\Gamma) \to (\cup_{P \in \mathcal{P}} H_2(X_{\overline{P}})) \cup H_2(X_{P_0}, L)$$

that maps each vertex v of Γ to the homology class $d(v) := u_*[(C_v, \partial C_v)]$ of the disk/sphere obtained by extending u over the punctures corresponding to nodal lifts;

(d) a labelling

$$\mu_{\mathfrak{D}} : \mathrm{Edge}_{\bullet,\to}(\Gamma) \to \mathbb{Z}_{>0}$$

that records the order of tangency of the map u to the divisor $\mathfrak{D}$ at markings that do not lie on horizontally constant components (with the convention that a transverse intersection has order 1).

The type is denoted simply as Γ, suppressing $\mathcal{T}, d, \mu_{\mathfrak{D}}$ in the notation, or by Γ_X if we wish to distinguish a type of map Γ_X from a type of treed disk Γ.

We introduce the following notations for moduli spaces. Let $\mathcal{M}^{\mathrm{brok}}(L, \mathfrak{D})$ be the moduli space of isomorphism classes of stable treed broken holomorphic disks in $\mathfrak{X}$ with boundary in L adapted to $\mathfrak{D}$, where isomorphism is modulo reparametrizations of domains and is defined in Definition 4.15. Let

$$\mathcal{M}^{\mathrm{brok}}_{\Gamma}(L, \mathfrak{D}) \subset \mathcal{M}(L, \mathfrak{D})$$

be the *moduli space of broken maps* of combinatorial type Γ modulo the action of domain reparametrizations. We recall that domain reparametrizations are isomorphisms of broken maps as in Definition 4.15. The group of tropical symmetries $T_{\mathrm{trop}}(\Gamma)$ acts naturally on $\mathcal{M}^{\mathrm{brok}}_{\Gamma}(L, \mathfrak{D})$. The quotient

$$\mathcal{M}^{\mathrm{brok}}_{\Gamma,\mathrm{red}}(L, \mathfrak{D}) := \mathcal{M}^{\mathrm{brok}}_{\Gamma}(L, \mathfrak{D}) / T_{\mathrm{trop}}(\Gamma)$$

is the *reduced moduli space of broken maps* of combinatorial type Γ. For $\underline{x} \in \mathcal{I}(L)^{d(\circ)+1}$, let

$$\mathcal{M}^{\mathrm{brok}}_{\Gamma}(L, \mathfrak{D}, \underline{x}) \subset \mathcal{M}^{\mathrm{brok}}_{\Gamma}(L, \mathfrak{D})$$

denote the adapted subset made of holomorphic treed disks of type Γ adapted to D with limits $\underline{x} = (x_0, \ldots, x_{d(\circ)}) \in \mathcal{I}(L)$ along the root and leaves. The union over all types with $d(\circ)$ incoming leaves is denoted

$$\mathcal{M}^{\text{brok}}_{d(\circ)}(L, \mathfrak{D}) = \bigcup_{\Gamma, \underline{x}} \mathcal{M}^{\text{brok}}_{\Gamma}(L, \mathfrak{D}, \underline{x}).$$

The space $\mathcal{M}^{\text{brok}}_{d(\circ)}(L, \mathfrak{D})$ has a natural topology for which convergence is a version of Gromov convergence defined in Chapter 8. In this paper we show that the compactifications of zero and one-dimensional components of the moduli space $\mathcal{M}^{\text{brok}}_{d(\circ)}(L, \mathfrak{D})^{<E}$ are manifolds. For one-dimensional components, the boundary consists of configurations with disk bubbling. For zero and one-dimensional components of $\mathcal{M}^{\text{brok}}_{d(\circ)}(L, \mathfrak{D})$ the symmetry tropical group is finite.

Remark 6.16. (Perturbed maps and curve counts) A perturbed holomorphic map u in a zero-dimensional moduli space makes a contribution of $\frac{1}{d(\bullet)!}$ to curve counts. Here $d(\bullet)$ is the number of interior leaves of u, and hence is the number of intersections of the map u with the stabilizing divisor. To justify this multiplicative factor, let us consider the special case where for a certain type of map, the moduli space can be regularized using a domain-independent perturbation. Then there are $d(\bullet)!$ ways of labelling the divisor intersections as $z_1, \ldots, z_{d(\bullet)}$. Each of these labellings is a 'different' map in the moduli space, and therefore a corrective factor of $\frac{1}{d(\bullet)!}$ needs to be inserted.

6.4. Fredholm theory for broken maps

We introduce a weighted Sobolev space needed for the transversality result. The norm is defined by viewing the domain as having punctures that map to cylindrical ends in the target. With this Sobolev completion, we can enforce higher order tangencies with relative divisors. The Sobolev norm is defined component-wise for a broken map. In this section, we focus on surface components of a broken map. The transversality result uses a standard norm on tree components.

Definition 6.17. (Relative map) A *relative type* is a graph Γ with a single vertex v, a collection $\text{Edge}_{\text{trop}}(\Gamma)$ of edges each with a single end-point, a polytope $P(v) \in \mathcal{P}$, slopes for the edges

$$\mathcal{T}(e) \in \mathfrak{t}_{P(e),\mathbb{Z}}, \quad e \in \text{Edge}_{\text{trop}}(\Gamma),$$

and homology data $d(v) \in H_2(X_{\overline{P}(v)})$. A *relative map* is a map $u : C^\circ \to X^{\square}_{\overline{P}(v)}$ from a punctured domain $C^\circ := C \setminus \{z_e : e \in \text{Edge}_{\text{trop}}(\Gamma)\}$ where C is a disk or sphere, to the manifold $X^{\square}_{\overline{P}(v)}$ with cylindrical ends, such that at a puncture z_e u is asymptotic to a $\mathcal{T}(e)$-cylinder, and the extension $\overline{u}$ of u over the punctures has homology class $d(v) \in H_2(X_{\overline{P}(v)})$.

A relative map must be viewed as a single component of a broken map, so that tropical nodes in the broken map appear as markings on a relative map. We call them *tropical markings*, and they correspond to the one-sided edges $\text{Edge}_{\text{trop}}(\Gamma)$ on the combinatorial type Γ of the relative map.

We introduce the space of smooth relative maps whose Sobolev completion will be given later. Let (C, j) be a connected Riemann surface with a set $\{z_e : e \in \text{Edge}_{\text{trop}}(\Gamma)\}$ of tropical markings. Let

$$C^\circ := C \backslash \{z_e : e \in \text{Edge}_{\text{trop}}(\Gamma)\} \tag{6.8}$$

be the curve with the tropical marked points removed. For the purpose of defining weighted Sobolev norms, we introduce cylindrical coordinates

$$(s_e, t_e) : U_{z_e} \backslash \{z_e\} \to [0, \infty) \times S^1 \tag{6.9}$$

on the neighborhood $U_{z_e} \subset C$ of every puncture z_e. Let

$$\text{Map}_\Gamma(C^\circ, X_{\overline{P}}^\square)$$

be the set of relative maps, which at a puncture $z_e \in C \backslash C^\circ$ corresponding to an edge $e \in \text{Edge}_{\text{trop}}(\Gamma)$ is asymptotically close to a $\mathcal{T}(e)$-cylinder

$$u_{\text{vert},e} : [0, \infty) \times S^1 \to X_{\overline{P}(e)}^\square, \quad (s_e, t_e) \mapsto e^{\mathcal{T}(e)(s_e + it_e)} x_e$$

for some point $x_e \in X_{\overline{P}(e)}^\square$. Here we recall that the punctured neighborhood at z_e in C° maps to a $P(e)$-cylindrical end of $X_{\overline{P}}^\square$ and the $P(e)$-cylindrical end is identified to $X_{\overline{P}(e)}^\square$, see (3.38). An infinitesimal deformation of a $\mathcal{T}(e)$-cylinder is given by a vector $\xi_e \in T_{x_e} X_{\overline{P}(e)}^\square$ and it corresponds to a section

$$\xi_{e,\mathcal{T}(e)} := \frac{d}{d\tau}(e^{\mathcal{T}(e)(s+it)} \exp_{x_e}(\tau \xi_e))|_{\tau=0} \in \Gamma(\mathbb{R}_+ \times S^1, u_{\text{vert},e}^* T X_{\overline{P}(e)}^\square). \tag{6.10}$$

The tangent space of $\text{Map}_\Gamma(C^\circ, X_{\overline{P}}^\square)$ at u consists of sections

$$\xi \in \Gamma(C^\circ, u^* T X_{\overline{P}}^\square)$$

which, for any edge $e \in \text{Edge}_{\text{trop}}(\Gamma)$, has a limit

$$\xi_e := \lim_{z \to z_e} \xi \in T_{x_e} X_{\overline{P}(e)}^\square \tag{6.11}$$

in the sense that ξ is asymptotically close to the section $\xi_{e,\mathcal{T}(e)}$ (6.10) as $z \to z_e$.

REMARK 6.18. The quantity ξ_e defined in (6.11) is the infinitesimal change in the tropical evaluation map $\text{ev}_{z_e}^{\mathcal{T}(e)}(u) \in X_{\overline{P}(e)}^\square$ (see (4.19)) when the broken map u is infinitesimally deformed by ξ. Since the tropical evaluation map depends on the choice of cylindrical coordinates (6.9) in the punctured neighborhood of z_e in the domain curve, the quantity ξ_e depends on infinitesimal changes in the domain cylindrical coordinates. On the other hand, the projected evaluation map $\pi_{\mathcal{T}(e)}^\perp(\text{ev}_{z_e}(u)) \in X_{\overline{P}(e)}^\square / T_{\mathcal{T}(e),\mathbb{C}}$ is independent of the domain cylindrical coordinates, and the same is true of its infinitesimal version $\pi_{\mathcal{T}(e)}^\perp(\xi_e) \in T(X_{\overline{P}(e)}^\square / T_{\mathcal{T}(e),\mathbb{C}})$.

We now define the Sobolev norm. Define a cutoff function

$$\beta \in C^\infty(\mathbb{R}, [0,1]), \quad \begin{cases} \beta(s) = 0 & s \leq 0 \\ \beta(s) = 1 & s \geq 1 \end{cases}. \tag{6.12}$$

Let

$$\kappa : C^\circ \to \mathbb{R} \tag{6.13}$$

be equal to $\beta(s_e)s_e$ near the puncture corresponding to z_e, for any edge e, and 0 outside all the edge neighbourhoods. For a section $\xi : C^\circ \to E$ of a vector bundle E with connection ∇, integers $k \geq 0$, $p > 1$ and constant $\lambda \in (0,1)$ define the $W^{k,p,\lambda}$-norm of ξ as

$$\|\xi\|^p_{W^{k,p,\lambda}} := \sum_{0 \leq i \leq k} \int_{C^\circ} |\nabla^i \xi|^p \exp(\lambda \kappa p) \, \mathrm{dvol}_{C^\circ} .$$

The norm on a section ξ is

$$\|\xi\|^\circ_\Gamma := \textstyle\sum_e |\xi_e| + \|\xi - \sum_e \beta \mathbb{T}^e_u \xi_{e,\mathcal{T}(e)}\|_{W^{1,p,\lambda}}, \tag{6.14}$$

where

$$\xi_e := \lim_{z \to z_e} \xi \in T_{x_e} X^\square_{\overline{P}(e)} \tag{6.15}$$

near the puncture z_e as in (6.15), u is asymptotic to a vertical cylinder $u_{\mathrm{vert},e}$,

$$\mathbb{T}^e_u : u^*_{\mathrm{vert},e} T X^\square_{\overline{P}} \to u^* T X^\square_{\overline{P}}$$

is the parallel transport map along geodesics, and $\xi_{e,\mathcal{T}(e)}$ is a section of $u^*_{\mathrm{vert},e} T X^\square_{\overline{P}}$ as in (6.10).

There are similar Sobolev spaces of maps and one-forms. Let $\mathrm{Map}^{1,p,\lambda}_\Gamma(C^\circ, X^\square_{\overline{P}})$ be the Banach completion of the space of maps $\mathrm{Map}_\Gamma(C^\circ, X^\square_{\overline{P}})$ under the norm (6.14). That is, for a smooth map $u \in \mathrm{Map}_\Gamma(C^\circ, X^\square_{\overline{P}})$ and a section $\xi \in \Gamma(C, u^* T X_{\overline{P}})$ satisfying $\|\xi\|^\circ_\Gamma < \infty$, the map $\exp_u \xi$ belongs to the completion $\mathrm{Map}^{1,p,\lambda}_\Gamma(C^\circ, X^\square_{\overline{P}})$. Let

$$\mathcal{E}^{p,\lambda} \to \mathrm{Map}^{1,p,\lambda}_\Gamma(C^\circ, X^\square_{\overline{P}})$$

be the vector bundle whose fiber

$$\mathcal{E}^{p,\lambda}_u = \Omega^{0,1}(C^\circ, u^* T X^\square_{\overline{P}})_{L^{p,\lambda}}$$

is the space of $(0,1)$-forms with respect to (j, J), where J is a fixed cylindrical almost complex structure. The vertical projection of the linearization of the Cauchy-Riemann operator $\overline{\partial}_{j,J}$ at u

$$D^\circ_u : T_u \, \mathrm{Map}^{1,p,\lambda}_\Gamma \to \mathcal{E}^{p,\lambda}_u$$

is a Fredholm operator by results of Lockhart-McOwen [**56**].

There is a related Fredholm operator on sections with compactified domain, called the *Cauchy-Riemann operator on compactifications*, and denoted by D_u. This operator will be useful for index computations. It is not used for transversality and gluing proofs where we stick to D°_u. We give two descriptions of the operator D_u:

Approach 1 to define D_u: Choose integers

$$k \in \mathbb{Z}_{>0}, \quad p > 1, \quad k > \mu + 2/p$$

where μ is the highest order of tangency $m_{z_e}(u, D_Q)$ to any relative divisor $D_Q \subset X_{\overline{P}}$ over all the edges $e \in \mathrm{Edge}_{\mathrm{trop}}(\Gamma)$. Let $\mathrm{Map}^{k,p}_\Gamma(C, X_{\overline{P}})$ be the set of $W^{k,p}$ maps from

C to $X_{\overline{P}}$ that intersect the relative divisors at marked points with the prescribed order of tangency. Assume the type Γ is such that the image of maps does not intersect orbifold singularities (if any) in $X_{\overline{P}}$. Consider the Banach bundle

$$\mathcal{E}^{k-1,p} \to \mathrm{Map}_\Gamma^{k,p}(C, X_{\overline{P}})$$

whose fiber over any map u is

$$\mathcal{E}_u^{k-1,p} = W_{\Gamma-1}^{k-1,p}(\Omega_{j,J}^{0,1}(C, u^*TX_{\overline{P}})),$$

where the space $W_{\Gamma-1}^{k-1,p}$ consists of sections, which have zeros of order one less than that prescribed by Γ at marked points. More rigorously, a $u^*TX_{\overline{P}}$-valued $(0,1)$-form η is in $W_{\Gamma-1}^{k-1,p}$ if for any marked point z_e and a relative divisor Y with prescribed intersection multiplicity $m_{z_e}(u, Y) \geq 2$, the projection of η to the normal bundle NY has a zero of order $m_{z_e}(u, Y) - 1$. The linearization of the Cauchy-Riemann operator at u

$$D_u : T_u \,\mathrm{Map}_\Gamma^{k,p}(C, X_{\overline{P}}) \to \mathcal{E}_u^{k-1,p}.$$

is a Fredholm operator.

Approach 2 to define D_u : The second approach is more general and is applicable even when the image of the map u contains orbifold singularities in $X_{\overline{P}}$. We extend the bundle $u^*TX_{\overline{P}}^{\square}$ on C° to a bundle $\overline{u^*TX_{\overline{P}}^{\square}}$ on C. To extend the bundle over a tropical marking $z_e \in C \backslash C^\circ$, we choose a toric compactification $X_{\overline{P}'_e}$ of $X_{\overline{P}}^{\square}$ in which $u(z_e)$ maps to a single relative divisor, as opposed to an intersection of a collection of relative divisors. The extension of the bundle $u^*TX_{\overline{P}}^{\square}$ over z_e is defined to be the pullback bundle $u^*TX_{\overline{P}'_e}$. The details are as follows:

Recall from Definition 4.14 that for any tropical marking $z_e \in C \backslash C^\circ$, the slope $\mathcal{T}(e) \in \mathfrak{t}_{P(e),\mathbb{Z}}$ is the product

$$\mathcal{T}(e) = \mu_e^{\mathrm{trop}} \mathcal{T}(e)_{\mathrm{prim}}, \tag{6.16}$$

of the primitive slope $\mathcal{T}(e)_{\mathrm{prim}} \in \mathfrak{t}_{P(e),\mathbb{Z}}$ and a multiplicity $\mu_e^{\mathrm{trop}} \in \mathbb{Z}_{\geq 1}$. Let $\overline{P}'_e$ be the polytope obtained by cutting off the face $P(e)$ in the polytope $\overline{P}$ along a hypersurface whose normal vector is $\mathcal{T}(e)_{prim}$. In other words, the fan of $\overline{P}'_e$ is obtained from the fan of $\overline{P}$ by adding a 1-dimensional cone $\mathbb{R}_{\geq 0}\mathcal{T}(e)_{prim}$. Let $X_{\overline{P}'_e}$ be the compactification of $X_{\overline{P}}^{\square}$ for which $X_{\overline{P}'_e} \backslash X_{\overline{P}}^{\square}$ is a collection of normally intersecting relative divisors, the tropical moment map Φ on $X_{\overline{P}}^{\square}$ extends to a tropical moment map Φ on $X_{\overline{P}'_e}$ with image $\Phi(X_{\overline{P}'_e}) = P'_e$. The cylindrical almost complex structure on $X_{\overline{P}}^{\square}$ extends to $X_{\overline{P}'_e}$ in a way that the relative divisors are pseudoholomorphic. Let $Y_e \subset X_{\overline{P}'_e}$ be the relative divisor corresponding to the cone $\mathbb{R}_{\geq 0}\mathcal{T}(e)_{prim}$. The map $u : C^\circ \to X_{\overline{P}}^{\square}$ extends over the point z_e and $u(z_e)$ maps to Y_e, and the intersection multiplicity with Y_e is μ_e^{trop}. Extend the bundle $u^*TX_{\overline{P}}^{\square}$ over the point z_e by considering a neighborhood $U_{z_e} \subset C$, and gluing the bundle $u^*TX_{\overline{P}'_e}|U_{z_e}$ to $u^*TX_{\overline{P}}^{\square}$. On the intersection, $U_{z_e} \backslash \{z_e\}$ there is a canonical identification between the bundles being glued. The vector bundle $\overline{u^*TX^{\square}_{\overline{P}}}$ on C is obtained by performing such a gluing at each of the points $z_e \in C\backslash C^{\circ}$.

The operator D_u is defined as the Cauchy-Riemann operator

$$D_u : \Omega^0_{\underline{\mu}}(C, \overline{u^*TX^{\square}_{\overline{P}}}) \to \Omega^{0,1}_{\underline{\mu}-1}(C, \overline{u^*TX^{\square}_{\overline{P}}}),$$

where $\underline{\mu} := (\mu_e^{\mathrm{trop}})_{e \in \mathrm{Edge}_{\mathrm{trop}}(\Gamma)}$, and the subscript in the domain and codomain indicate the order of vanishing of the sections at the tropical markings z_e, $e \in \mathrm{Edge}_{\mathrm{trop}}(\Gamma)$ in the direction normal to $T_{u(z_e)}Y_e \subset T_{u(z_e)}X_{\overline{P}'_e}$. We leave it to the reader to check that both approaches give the same operator D_u whenever Approach 1 is applicable, that is, when the map u does not contain orbifold singularities of $X_{\overline{P}}$.

PROPOSITION 6.19. *Let $u \in \mathrm{Map}_{\Gamma}(C, X_{\overline{P}})$ be a relative holomorphic map. There are isomorphisms*

$$\ker(D_u) \simeq \ker(D_u^{\circ}), \quad \mathrm{coker}(D_u) \simeq \mathrm{coker}(D_u^{\circ}).$$

PROOF. Since the proof is by a local computation near each of the special points, it is enough to prove the result for a map whose domain is a single curve component with a single marked point. We may also assume that the marked point is mapped to a single relative divisor $Y \subset X_{\overline{P}}$ with intersection multiplicity $\mu \in \mathbb{Z}_{>0}$, since Approach 2 (in page 132) for the definition of D_u reduces all intersections to the case of a single divisor intersection. We assume that the marked point is $0 \in C$. In this proof, we denote $X := X_{\overline{P}}$ and $X^{\square} := X^{\square}_{\overline{P}}$.

The proof is based on the fact that in a neighbourhood of the divisor the tangent space splits into a vertical and a horizontal subspace. Indeed, because of the cylindrical almost complex structure, there is a neighbourhood $U_Y \subset X$ of Y for which there is a projection $\pi_Y : U_Y \to Y$ with holomorphic fibers. The tangent space splits into a J-holomorphic horizontal and vertical part:

$$TX|U_Y \simeq V \oplus H, \quad H := \pi_Y^* TY, \quad V := \ker(d\pi_Y). \tag{6.17}$$

The operator D_u is defined by the Levi-Civita connection of the cylindrical metric (Definition 3.39). Therefore under the splitting (6.17),

$$D_u = \begin{pmatrix} \overline{\partial} & A \\ 0 & D_u^Y \end{pmatrix} \tag{6.18}$$

in a neighbourhood $U_0 \subset C$ of the marked point. Here $\overline{\partial} : \Gamma(U_0, u^*V) \to \Omega^{0,1}(U_0, u^*V)$ is the standard Cauchy-Riemann operator, $A : \Gamma(U_0, u^*H) \to \Omega^{0,1}(U_0, u^*V)$ is multiplication with a tensor and

$$D_u^Y : \Gamma(U_0, u^*H) \to \Omega^{0,1}(U_0, u^*H)$$

is the lift of the linearized operator D_{u_Y}, where $u_Y := \pi_Y \circ u$.

The correspondence for the kernels follows by decay estimates. For an element $\xi \in W^{k,p}(C, u^*TX)$, we denote the horizontal part and vertical part by

$$\xi_h \in W^{k,p}(U_0, u^*H), \quad \xi_v \in W^{k,p}(U_0, u^*V).$$

in the neighbourhood $U_0 \subset C$ of 0. Denote by $|\cdot|$ resp. $|\cdot|^\circ$ the ordinary metric on TX resp. the cylindrical metric on $X \backslash Y$. Then, we have

$$|\xi_h(z)|^\circ \sim |\xi_h(z)|, \quad |\xi_v(z)|^\circ \sim |\xi_v(z)|/|u(z)|, \quad z \neq 0, \tag{6.19}$$

where the norm on $u(z)$ is a norm on the fiber of the projection $U_Y \to Y$. Since

$$|\xi_h| \leq c, \quad |\xi_v(z)| \leq c|z|^\mu, \quad |u(z)| \sim c|z|^\mu,$$

we conclude that $|\xi_h|^\circ$ and $|\xi_v(z)|^\circ$ are uniformly bounded and have a limit as $z \to 0$. Since ξ is smooth, the convergence is exponential :

$$|\xi(z) - \xi(0)|^\circ, |\nabla \xi_h(z)|^\circ, |\nabla \xi_v(z)|^\circ \sim c|z| \sim ce^{-s}$$

where (s,t) are cylindrical coordinates on $\mathbb{P}^1 \backslash \{0\}$ near the puncture. We conclude

$$\xi \in W^{k,p}_\Gamma(C,X) \implies \xi - \xi(0) \in W^{1,p,\lambda}(C^\circ, u^*TX^\square)$$

for any $0 < \lambda < 1$, which implies $\ker D_u \subset \ker D_u^\circ$. The converse $\ker D_u^\circ \subset \ker D_u$ is proved using removal of singularity, elliptic regularity, and the estimate (6.19).

A similar argument holds for the cokernels. We first prove the inclusion $\operatorname{coker}(D_u) \subset \operatorname{coker}(D_u^\circ)$. Consider $\eta \in \operatorname{coker}(D_u)$. We view $\operatorname{coker}(D_u)$ as a subspace in the $(W^{k-1,p})^\vee$-completion of $\Omega^{1,0}(C, u^*(T^*X))$, which is a space of distributions. (We use $\Omega^{1,0}(u^*(T^*X))$ instead of $\Omega^{0,1}(u^*TX)$ to avoid making choices of metrics on C and X.) By elliptic regularity, the distribution η is represented by a smooth section in the complement of marked points. So we focus attention in a neighbourhood $U_0 \subset C$ of 0. For any $W^{k,p}$-section $\xi : C \to u^*TX$ that is supported in U_0, and is vertical, we have $z^\mu \xi \in W^{k,p}_\Gamma(C, u^*TX)$. So for any such ξ,

$$0 = \int_C (\overline{\partial}(z^\mu \xi), \eta_v).$$

Therefore, $\overline{\partial}(z^\mu \eta_v) = 0$ weakly in U_0 and so, $z^\mu \eta_v$ can be represented by a smooth function. Next, using the split form (6.18), we observe that any section ξ that is supported in U_0 and is horizontal satisfies

$$(A\xi, \eta_v) + (D_u^Y \xi, \eta_h) = 0,$$

and so, $A^*\eta_v + (D_u^Y)^*\xi_h = 0$ weakly in U_0. The tensor A vanishes to order μ at $0 \in C$ in the $|\cdot|$-norm, as a consequence of the transformation relation (6.19) and the fact that A is bounded in the $|\cdot|^\circ$ norm. So, $A^*\eta_v$ is smooth in U_0. By elliptic regularity ξ_h is smooth in U_0. Finally, η is in $\operatorname{coker} D_u^\circ$ because of the following transformations valid in U_0 :

$$|\eta_h(z)|^\circ \sim |\eta_h(z)|, \quad |\eta_v(z)|^\circ \sim |\eta_v(z)| \cdot |u(z)|, \quad z \neq 0. \tag{6.20}$$

Indeed, $|\eta(z)|^\circ$ is bounded, and therefore is in $L^{p^*,-\lambda}$ the dual space of $L^{p,\lambda}$. The reverse inclusion $\operatorname{coker}(D_u^\circ) \subset \operatorname{coker}(D_u)$ follows formally from the inclusion relation

$$W^{k-1,p}_{\mu-1}(\Omega^{0,1}(C, u^*TX)) \to L^{p,\lambda}(\Omega^{0,1}(C^\circ, u^*TX^\square))$$

between the target spaces of D_u and D_u°.

□

Next we relate the Fredholm index of the linearized operator D_u° to the Maslov index of disks.

Definition 6.20. (Adjusted Maslov index) Given a relative map u of type Γ the *adjusted Maslov index* of u is

$$I_{\mathrm{adj}}(\Gamma) := I_{\mathrm{adj}}(u) := \mathrm{ind}(D_u^\circ) + 2|\operatorname{Edge}_{\mathrm{trop}}(\Gamma)|.$$

In the definition of the adjusted Maslov index we add $2|\operatorname{Edge}_{\mathrm{trop}}(\Gamma)|$ because in the definition of D_u° the tropical markings are fixed on the domain curve.

Remark 6.21. (Relation to Maslov index) We give a relation between the adjusted Maslov index and the ordinary Maslov index. We recall that the *boundary Maslov index* $I(E,F)$ is an integer assigned to a pair (E,F) consisting of a complex vector bundle E on a Riemann surface C with boundary, and a totally real subbundle $F \subset E|_{\partial C}$ on the boundary (see [**60**, Appendix C]). In case $\partial C = \emptyset$, the Maslov index $I(E,\emptyset)$ is twice the first Chern number $c_1(E)$.

The adjusted Maslov index is then

$$I_{\mathrm{adj}}(u) = I(\overline{u^*TX^{\square}_{P}}, u^*TL) - \sum_{e\in \operatorname{Edge}_{\mathrm{trop}}(\Gamma)} 2(\mu_e^{\mathrm{trop}} - 1). \tag{6.21}$$

Here $\overline{u^*TX^{\square}_{P(v)}}$ is the extended vector bundle on C defined in Approach 2 in page 132. The second term in the rhs of (6.21) arises because we consider sections that vanish to order μ_e^{trop} at the point z_e, $e \in \operatorname{Edge}_{\mathrm{trop}}(\Gamma)$ in the direction normal to Y_e.

Example 6.22. The Maslov index of a disk or sphere is equal to the adjusted Maslov index if all relative markings are incident on relative divisors (and not relative submanifolds of higher codimension) and have a multiplicity of one. This is an example of a map where the two quantities differ although the edge multiplicity μ_e^{trop} is 1 for all edges $e \in \operatorname{Edge}_{\mathrm{trop}}(\Gamma)$. Suppose $X_{P(v)} = \mathbb{P}^2$ with a Hamiltonian $(S^1)^2$-action, $X^{\square}_{P(v)}$ is the complement of the 3 torus-invariant divisors, $L \subset X_{P(v)}$ is a toric Lagrangian, and $u : (D, \partial D) \to (X_{P(v)}, L)$ is a disk through the point $[1:0:0]$ whose boundary Maslov index is 4. In contrast the adjusted Maslov index $I_{\mathrm{adj}}(u)$ is 2, because the adjusted Maslov index accounts for the constraint that the disk passes through the point $[1:0:0]$ which cuts down the index by 2. The reason for the discrepancy is that the pullback bundle $u^*TX_{P(v)}$ is different from $\overline{u^*TX^{\square}_{P(v)}}$ since a tropical marking maps to the intersection of two divisors.

6.5. The index of a broken map

The index of a broken map is the expected dimension of the moduli space containing the broken map. In this section, we give an expression for the index in terms of the Fredholm index of the linearized Cauchy-Riemann operator defined in the last section. We also prove that collapsing tropical edges in the tropical graph of the broken map does not change the expected dimension of the moduli space.

For a type of broken maps the expected dimension of the moduli space is the same as that of the glued type of unbroken maps. We define ‘glued type’:

DEFINITION 6.23. (Glued type) Let Γ be the combinatorial type of a broken map. Then the *glued type* Γ_{glue} is the type of the unbroken map obtained by collapsing the tropical edges $e \in \mathrm{Edge}_{\mathrm{trop}}(\Gamma)$ in Γ.

NOTATION 6.24. We introduce the following notation for stating the result on expected dimension:

(a) (Maslov index for an unbroken type) The Maslov index $I(\Gamma)$ of an unbroken map type Γ is the sum of Maslov indices of the maps on the surface components, that is,

$$I(\Gamma) := \sum_{v \in \mathrm{Vert}(\Gamma)} I(\Gamma_v).$$

(b) (Adjusted Maslov index for a broken type) The adjusted Maslov index $I(\Gamma)$ of a broken map type Γ is the sum of adjusted Maslov indices (Definition 6.20) of the relative maps on the surface components, that is,

$$I_{\mathrm{adj}}(\Gamma) := \sum_{v \in \mathrm{Vert}(\Gamma)} I_{\mathrm{adj}}(\Gamma_v).$$

(c) (Morse index of a critical point) For a critical point $x \in \mathcal{I}(L)$ of the Morse function of the Lagrangian, the *Morse index* $i_{\mathrm{Morse}}(x)$ is the dimension of its stable manifold in L.

(d) (Morse index of an input/output tuple) Let $\underline{x} \in \mathcal{I}(L)^{d+1}$ be a tuple consisting of an output label x_0 and inputs $x_1, \ldots, x_d$ of a treed holomorphic map. The Morse index of the tuple is the difference

$$i_{\mathrm{Morse}}(\underline{x}) := i_{\mathrm{Morse}}(x_0) - \textstyle\sum_{i=1}^{d(\circ)} i_{\mathrm{Morse}}(x_i)$$

(e) (Index for unbroken and broken maps) Given an unbroken resp. broken map type Γ and an input/output tuple $\underline{x} \in \mathcal{I}(L)^{d(\circ)+1}$, the index of the type Γ with labels $\underline{x}$, denoted by

$$i(\Gamma, \underline{x}) \quad \text{resp.} \quad i^{\mathrm{brok}}(\Gamma, \underline{x}),$$

is the expected dimension of the moduli space $\mathcal{M}_\Gamma(L, \underline{x})$ resp. $\mathcal{M}_\Gamma^{\mathrm{brok}}(L, \underline{x})$.

PROPOSITION 6.25. (Expected dimension) *Suppose* Γ *is the combinatorial type of a broken map with limits* $\underline{x} \in \mathcal{I}(L)^{d(\circ)+1}$ *along the root and leaves.*

(a) (Maslov index sum formula) *The Maslov index of the glued disk type* Γ_{glue} *is*

$$I(\Gamma_{\mathrm{glue}}) = I_{\mathrm{adj}}(\Gamma) - 4|\,\mathrm{Edge}_{\mathrm{trop}}(\Gamma)|. \tag{6.22}$$

(b) The expected dimension of $\mathcal{M}_\Gamma^{\mathrm{brok}}(L, \mathfrak{D}, \underline{x})$ *is given by*

$$i^{\mathrm{brok}}(\Gamma, \underline{x}) := d(\circ) - 2 + i_{\mathrm{Morse}}(\underline{x}) + I(\Gamma_{\mathrm{glue}}) - c(\Gamma), \tag{6.23}$$

where

$$c(\Gamma) := 2|\,\mathrm{Edge}_{\mathrm{int},\bullet}(\Gamma)| + |\,\mathrm{Edge}_{\circ,-}^{0}\,| + |\,\mathrm{Edge}_{\circ,-}^{\infty}\,| + 2\textstyle\sum_{e \in \mathrm{Edge}_{\bullet,\to}(\Gamma)} (\mu_{\mathfrak{D}}(e) - 1)$$

is a factor accounting for internal disk and sphere nodes in Γ *and higher order intersections with the stabilizing divisor, and the other terms are defined in Notation 6.24. Thus, collapsing tropical edges does not affect the index, that is,*

$$i^{\text{brok}}(\Gamma, \underline{x}) = i(\Gamma_{\text{glue}}, \underline{x}).$$

Proof of Proposition 6.25. We first prove the Maslov index sum formula (6.22). Consider a tropical edge $e = (v_+, v_-)$ in Γ whose corresponding node w_e lifts to the points $w_e^+ \in C_{v_+}$, $w_e^- \in C_{v_-}$ in the normalized domain curve $\tilde{C}$. Recall from definition 4.14 that the slope $\mathcal{T}(e) \in \mathfrak{t}_{P(e),\mathbb{Z}}$ is the product

$$\mathcal{T}(e) = \mu_e^{\text{trop}} \mathcal{T}(e)_{\text{prim}}$$

of a primitive integer vector $\mathcal{T}(e)_{\text{prim}} \in \mathfrak{t}_{P(e),\mathbb{Z}}$ and a positive integer μ_e^{trop}. We assume that $u_{v_\pm}(z_e)$ maps to a single relative divisor $Y \subset X_{\overline{P}(v_\pm)}$. Otherwise, we may use a different compactification of $X^{\square}_{\overline{P}(v_\pm)}$ (as in Approach 2 in page 132) to enforce this condition. Because of the cylindrical complex structure on the ends, a neighbourhood of Y in $X_{\overline{P}(v_\pm)}$ can be viewed as a neighbourhood of the zero section of a direct sum of complex line bundles

$$\pi_\pm : L_\pm \oplus H \to Y.$$

In this region, the tangent space splits as the sum of horizontal and vertical subbundles

$$TX_{\overline{P}(v_\pm)} = H \oplus L_\pm, \quad H = \pi_\pm^* TY.$$

Topologically, the maps u_{v_+}, u_{v_-} are obtained by cutting u_{glue} to yield $u^\circ_{v_+}$, $u^\circ_{v_-}$ and adding a capping disk to either side. The horizontal bundle $(u^\circ_{v_\pm})^* H$ extends trivially over the capping disk, but for $(u^\circ_\pm)^* L$, the bundle over the capping disk is glued in with a μ_e^{trop} twist. Thus gluing at the node w_e has the effect of adding the Maslov indices on both sides and subtracting $4\mu_e^{\text{trop}}$. Summing over contributions from all tropical edges $e \in \text{Edge}_{\text{trop}}(\Gamma)$, we get

$$\sum_{v \in \text{Vert}(\Gamma} I(\overline{u^* TX^{\square}_{\overline{P}(v)}}) - \sum_{e \in \text{Edge}_{\text{trop}}(\Gamma)} 4\mu_e^{\text{trop}} = I(u_{\text{glue}}). \tag{6.24}$$

We recall from (6.21) that

$$I_{\text{adj}}(\Gamma_v) = I(\overline{u^* TX^{\square}_{\overline{P}(v)}}) - \sum_{e \in \text{Edge}_{\text{trop}}(\Gamma): v \in e} 2(\mu_e^{\text{trop}} - 1). \tag{6.25}$$

Combining (6.24), (6.25), we obtain the relation (6.22) on Maslov indices claimed by the Proposition.

The expected dimension of $\mathcal{M}^{\text{brok}}_\Gamma(L, \mathfrak{D}, \underline{x})$ can be calculated in a straightforward way. To lighten notation, let us assume that all nodes are tropical, since internal nodes can be accounted for in an obvious way. For the map component u_v, $v \in \mathrm{Vert}(\Gamma)$ the expected dimension of the moduli space of relative maps u_v is

$$i(\Gamma_v) = \begin{cases} (d(\circ)+1) + i(\underline{x}) + I_{\mathrm{adj}}(\Gamma_v) - \mathrm{Aut}(\mathbb{D}^2), & v \in \mathrm{Vert}_\circ(\Gamma), \\ \dim(X) + I_{\mathrm{adj}}(\Gamma_v) - \mathrm{Aut}(\mathbb{P}^1), & v \in \mathrm{Vert}_\bullet(\Gamma). \end{cases}$$

We obtain the dimension formula (6.23) in the Proposition by adding the contributions from all vertices $v \in \mathrm{Vert}(\Gamma)$, subtracting $(\dim(X) - 2)$ for each tropical edge to account for the codimension of the matching condition, and applying the Maslov index sum formula (6.22). □

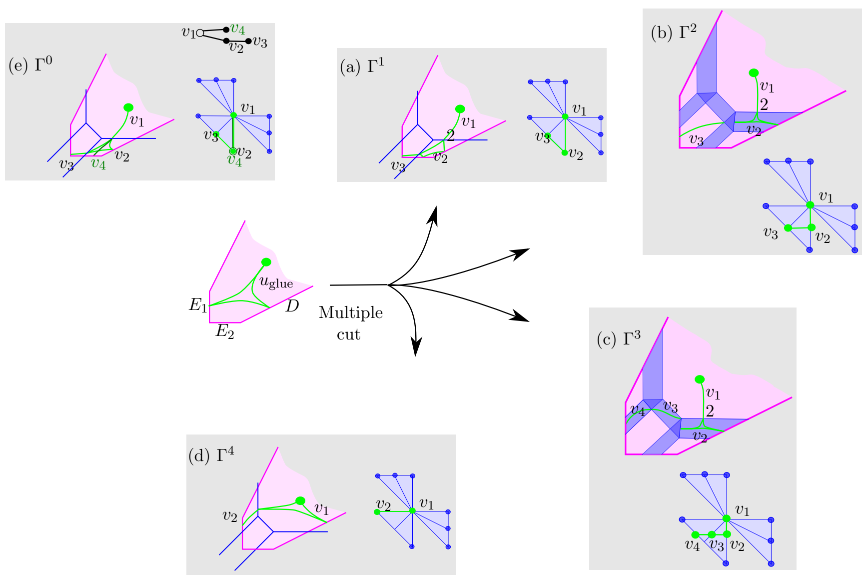

FIGURE 6.2. Broken maps corresponding to the Maslov index four disk class $[u_{\mathrm{glue}}] = 2[\delta_{E_2}] + [E_2] = [\delta_{E_1}] + [\delta_D]$.

EXAMPLE 6.26. (Adjusted Maslov index computation) In this example we consider the multiple cut of the cubic surface from Section 2.1. We list all broken maps whose gluing has homology class $2[\delta_{E_2}] + [E_2]$, and write down the adjusted Maslov indices of all the components. Here $\delta_{E_2} \in H_2(X, L)$ is the class of the disk of Maslov index two with a single intersection with the short divisor E_2. There are five types Γ of broken disks whose gluing has homology class $2[\delta_{E_2}] + [E_2]$ shown in Figure 6.2.

(a) In Γ^1 the edge $e = (v_1, v_2)$ has multiplicity 2. The adjusted Maslov indices are

$$I_{\mathrm{adj}}(v_1) = 2, \quad I_{\mathrm{adj}}(v_2) = 8, \quad I_{\mathrm{adj}}(v_3) = 2.$$

Note that the ordinary Maslov indices are different : $I(v_1) = 4$, $I(v_2) = 10$. The value of I_{adj} is lower for v_1, v_2 because it accounts for the intersection multiplicity of 2 at the node w_e corresponding to the edge $e = (v_1, v_2)$. We point out that the map u_{v_3} is homologous to the (-1)-divisor E_2' which is obtained by cutting the (-2)-divisor E_2, see Figure 2.5 for notation. Since the almost complex structure is a perturbation of the toric almost complex structure, the image of u_{v_3} does not lie on the toric divisor E_2'.

(b) The edge slopes in Γ^2 are

$$\mathcal{T}(v_1, v_2) = (0, -2), \quad \mathcal{T}(v_2, v_3) = (-1, 0).$$

The adjusted Maslov indices of the vertices in Γ^3 are

$$I_{\text{adj}}(v_1) = 2, \quad I_{\text{adj}}(v_2) = 6, \quad I_{\text{adj}}(v_3) = 4.$$

Indeed u_{v_2} is a curve with self-intersection number 2 in $X_{\overline{P}(v_2)}$ which is the second Hirzebruch surface H_2, and therefore, $c_1(u_{v_2}^* TX_{\overline{P}(v_2)}) = 4$. The ordinary Maslov indices are $I(v_1) = 4$, $I(v_2) = 8$. The adjusted Maslov indices $I_{\text{adj}}(v_1)$, $I_{\text{adj}}(v_2)$ are each 2 lower because the node between v_1, v_2 has an intersection multiplicity of 2 with the relative divisor. To calculate $I_{\text{adj}}(v_3)$, we need to choose an alternate almost complex compactification of $X^{\square}_{P(v_3)}$ since u_{v_3} has a nodal point w_e corresponding to the edge $e = (v_2, v_3)$ mapping to an orbifold singularity. We choose the compactification to be a resolution of the A_2-singularity as in Figure 6.3 so that the nodal point w_e lies in a single divisor, from where, we conclude $I_{\text{adj}}(v_3) = 4$.

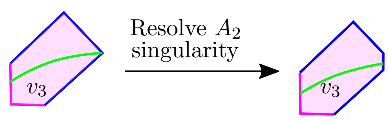

FIGURE 6.3. Orbifold singularities can be eliminated by changing the compactification.

(c) The edge slopes in Γ^3 are

$$\mathcal{T}(v_1, v_2) = (0, -2), \quad \mathcal{T}(v_2, v_3) = (-1, 0), \quad \mathcal{T}(v_3, v_4) = (-1, 0),$$

and the adjusted Maslov indices are

$$I_{\text{adj}}(v_1) = 2, \quad I_{\text{adj}}(v_2) = 6, \quad I_{\text{adj}}(v_3) = 4, \quad I_{\text{adj}}(v_4) = 4.$$

The adjusted Maslov indices of v_3 and v_4 are computed by passing to a toric resolution as in as in Figure 6.3.

(d) In Γ^4 the adjusted Maslov indices are

$$I_{\text{adj}}(v_1) = 4, \quad I_{\text{adj}}(v_2) = 4.$$

(e) In a broken map of type Γ^0, the component u_{v_1} has two distinct intersections with the relative divisor, each of multiplicity 1. For each of the vertices v_i in Γ^0 the Maslov index $I(u_{v_i})$ is equal to the adjusted Maslov index $I_{\text{adj}}(u_{v_i})$, and

$$I_{\text{adj}}(v_1) = 4, \quad I_{\text{adj}}(v_2) = 6, \quad I_{\text{adj}}(v_3) = 2, \quad I_{\text{adj}}(v_4) = 4.$$

For the broken map types $\Gamma^0, \ldots, \Gamma^4$, the Maslov index sum formula (6.22) implies that the Maslov index $I(\Gamma^i_{\text{glue}})$ of the glued disk is 4.

6.6. Transversality

In this section, we show that for a set of comeager domain-dependent perturbations, moduli spaces of broken maps with index at most one are transversely cut out. The perturbation scheme we use here only achieves transversality for certain combinatorial types, as in Cieliebak-Mohnke [**22**]. In particular, the Cieliebak-Mohnke [**22**] perturbation scheme can not achieve transversality on 'crowded components', which are components where the map is constant and which contain more than one marking. The moduli space of such maps cannot be transversally cut out, since the constraint that the first marking in such a component S maps to the stabilizing divisor D together with the fact that $u|S$ is a constant map guarantees that the second marking does as well. In the context of broken maps, the 'constant' condition needs to be replaced by 'horizontally constant'. Recall that a component $u|S_v : S_v \to X_{\overline{P}(v)}$ is horizontally constant if its projection to $X_{P(v)}$ is constant. We define crowdedness for broken maps.

Definition 6.27. The combinatorial type Γ of a broken map $u : C \to \mathfrak{X}$ is *crowded* if there exists a connected subgraph $\Gamma' \subset \Gamma$ such that $u|S_v$ is horizontally constant (see Definition 4.16) on all vertices $v \in \mathrm{Vert}(\Gamma')$ and Γ' contains more than one interior leaf.

Definition 6.28. For a type Γ of treed disks, a perturbation datum $\mathfrak{p}_\Gamma$ is *regular* if any $\mathfrak{p}_\Gamma$-adapted map is regular, and so the moduli space $\mathcal{M}_\Gamma(\mathfrak{X}, L, D)$ is a manifold of expected dimension.

The following theorem on the existence of regular perturbation data for broken maps is the main result of this section. Perturbation data is defined stratawise, and at each step we assume that the data on the 'smaller' strata is fixed, where the ordering on the strata is as follows: For types Γ', Γ of broken maps,

$$\Gamma' < \Gamma \tag{6.26}$$

if Γ is obtained from Γ' by collapsing an edge or making the length of a boundary edge finite or non-zero. The ordering relation helps in describing the boundary of the moduli spaces of treed curves as

$$\overline{\mathcal{M}}_\Gamma \backslash \mathcal{M}_\Gamma = \cup_{\Gamma' < \Gamma} \mathcal{M}_{\Gamma'}.$$

Theorem 6.29. (Transversality) *Let $\mathfrak{X}$ be a broken manifold with a symplectic cylindrical structure as in Definition 3.21. Let $\mathfrak{D} \subset \mathfrak{X}$ be a cylindrical broken divisor and $\mathfrak{J}_0 \in \mathcal{J}^{\mathrm{cyl}}(\mathfrak{X}, \mathfrak{D})$ be a locally strongly tamed cylindrical almost complex structure adapted to $\mathfrak{D}$ such that $(\mathfrak{D}, \mathfrak{J}_0)$ is a stabilizing pair, and let $F_0 : L \to \mathbb{R}$ be a Morse function. Suppose Γ is a type of stable treed disks, and regular perturbation data for types Γ' of stable treed disks with $\Gamma' < \Gamma$ are given. Then there exists a comeager subset*

$$\mathcal{P}^{\mathrm{reg}}_\Gamma(D) \subset \mathcal{P}_\Gamma(D)$$

of regular perturbation data for type Γ with background data $(\mathfrak{J}_0, F_0)$ coherent with the previously chosen perturbation data. For any regular perturbation $\mathfrak{p} \in \mathcal{P}_\Gamma$, and an uncrowded type Γ_X of broken maps whose domain type is Γ (see Definition 6.15) the moduli space $\mathcal{M}_{\Gamma_X}(\mathfrak{X}, L, D)$ is a smooth oriented manifold of expected dimension.

Proof of Theorem 6.29. Transversality, as in the unbroken case, is an application of Sard-Smale as in Cieliebak-Mohnke [**22**] and Charest-Woodward [**18**] on the universal space of maps. The new feature is that in neck pieces of the broken manifold the almost complex structure is fixed in the fiber direction. This does not pose any issues for maps whose horizontal projection is non-constant. Components of the map whose horizontal projection is constant will shown to be automatically transversal.

The moduli space is cut out as a zero set of a section of a Banach bundle which we now describe. We restrict our attention to types of maps for which all intersections with the stabilizing divisor have multiplicity one. Other types are discussed later. We construct the moduli space of broken maps without framing, since the framed version is a finite cover of the unframed one. (See Definition 4.12 (d) of framing and Remark 4.25.) We cover the moduli space of treed disks $\mathcal{M}_\Gamma$ by charts $\cup_i \mathcal{M}^i_\Gamma$, so that on a trivialization of the universal curve $\mathcal{U}^i_\Gamma$, each of the fibers is a fixed treed curve $C = S \cup T$ with fixed special points (see (4.7)). The complex structure on S varies smoothly in the sense that it is given by a map

$$\mathcal{M}^i_\Gamma \to \mathcal{J}(S_\Gamma), \quad m \mapsto j(m).$$

In order to apply Sard-Smale, we pass to maps from the normalized curve of a fixed Sobolev class. The domain of the map is the punctured curve $C^\circ \subset C$ with punctures at tropical nodal points w_e (corresponding to $e \in \mathrm{Edge}_{\mathrm{trop}}(\Gamma)$). Let

$$\tilde{C}^\circ := \bigsqcup_{v \in \mathrm{Vert}(\Gamma)} S^\circ_v \sqcup \bigsqcup_{e \in \mathrm{Vert}(\Gamma)} T_e$$

denote the normalized curve in the sense that nodes in C° (corresponding to internal edges $e \in \mathrm{Edge}_{\mathrm{int}}(\Gamma)$) are lifted to double points in $\tilde{C}^\circ$ and the tree components in C° are detached from the surface components. Choose $p > 2$ and $\lambda \in (0,1)$. Let

$$\mathrm{Map}^{1,p,\lambda}_\Gamma(\tilde{C}^\circ, \mathfrak{X}, L, \mathfrak{D})$$

denote the completion of maps under the weighted Sobolev norm $\|\cdot\|^\circ_\Gamma$ in (6.14) for surface components, and the ordinary $W^{1,p}$-norm for the tree components. That is, an element of this space consists of

(a) a collection of $W^{1,p,\lambda}$-maps

$$u_v : S^\circ_v \to X_{\overline{P}(v)}, \quad u(\partial S_v) \subset L, \quad v \in \mathrm{Vert}(\Gamma)$$

for each vertex v (where S°_v is defined in (4.11)), with the puncture at the node corresponding to any tropical edge $e \in \mathrm{Edge}_{\mathrm{trop}}(\Gamma)$ asymptotic to a $\mathcal{T}(e)$-cylinder;

(b) and $W^{1,p}$ maps

$$u_e : T_e \to L, \quad e \in \mathrm{Edge}(\Gamma)$$

from tree components to L asymptotic to the critical points $\underline{x}$ on the leaves $T_e, e \in \mathrm{Edge}_\circ(\Gamma)$. If $\ell(e)$ is finite and non-zero, that is $e \in \mathrm{Edge}^{(0,\infty)}_\circ(\Gamma)$, then $T_e \simeq [0,1]$ and if $\ell(T_e) = \infty$, each segment in T_e is either $[0,\infty)$, $(-\infty, 0]$ or $\mathbb{R}$.

The metric on $X^{\square}_{\overline{P}(v)}$ is chosen to be cylindrical on the ends, and so that L is totally geodesic.

The perturbation data on the strata lower than Γ glue to give regular perturbation datum in a small neighborhood in the boundary of the moduli space $\overline{\mathcal{M}}_\Gamma$: Any perturbation data for Γ' with $\Gamma' < \Gamma$ induces perturbations for Γ in a neighborhood of $\mathcal{U}_{\Gamma'}$ in $\overline{\mathcal{U}_\Gamma}$ by the gluing construction in (8.3) and a similar gluing construction for the tree parts. Let

$$\mathcal{P}_\Gamma = \{\mathfrak{p}_\Gamma = (J_\Gamma, F_\Gamma)\}$$

be the space of perturbation data defined on $\overline{\mathcal{U}}_\Gamma$ whose distance from the background data $(\mathfrak{J}_0, F_0)$ is bounded in the C^ε-norm, and that agree with the glued perturbation datum on a fixed open neighbourhood $\mathcal{N}_{\Gamma'}$ of $\mathcal{U}_{\Gamma'}$ for the strata $\Gamma' < \Gamma$. For a sufficiently small neighborhood $\mathcal{N}_{\Gamma'}$, any such perturbation $\mathfrak{p}_\Gamma$ is already regular for maps u with domain in $\mathcal{N}_{\Gamma'}$. Indeed, otherwise one could take a sequence u_ν of irregular maps of fixed area with domain converging to $\mathcal{U}_{\Gamma'}$. By Gromov compactness in Chapter 8, after passing to a subsequence we may assume that u_ν converges to a limit u_∞ of type Γ'. By surjectivity of gluing in Section 9.7 and Lemma 9.6, the linearized operator D_{u_ν} is surjective for ν sufficiently large, which is a contradiction. Since the gluing and compactness arguments do not use regularity, circularity of argument is avoided.

We use the C^ε-norm for J_Γ and a C^l-norm for F_Γ where $l > 1$ is a fixed number. For any broken curve C of type Γ we obtain perturbation data on C by identifying it conformally with a fiber of the universal tree disk $\mathcal{U}_\Gamma$. Let

$$\mathcal{B}^i_{p,\lambda,l,\Gamma} := \mathcal{M}^i_\Gamma \times \mathrm{Map}^{1,p,\lambda}_\Gamma(\tilde{C}, \mathfrak{X}, L, \mathfrak{D}) \times \mathcal{P}_\Gamma(\mathfrak{X}, \mathfrak{D}).$$

Let $\mathcal{E}^i = \mathcal{E}^i_{p,\lambda,\Gamma}$ be the Banach bundle over $\mathcal{B}^i_{p,\lambda,l,\Gamma}$ given by

$$(\mathcal{E}^i_{p,\lambda,\Gamma})_{j,u,J} \subset L^{p,\lambda}(\Omega^{0,1}_{j,J}(S, (u|S)^* T\mathfrak{X})) \oplus L^p(\Omega^1(T, (u|T)^* TL)).$$

Here the first summand is the space of $0,1$-forms with respect to $(j(m), J)$. The Cauchy-Riemann and shifted gradient operators applied to the restrictions $u|S$ resp. $u|T$ of u to the two resp. one dimensional parts of $C = S \cup T$ define a C^{l-1} section

(6.27)
$$\overline{\partial}_\Gamma : \mathcal{B}^i_{p,\lambda,l,\Gamma} \to \mathcal{E}^i_{p,\lambda,\Gamma}, \quad (C, u, (J_\Gamma, F_\Gamma)) \mapsto \left(\overline{\partial}_{j(m),J_\Gamma} u|S, \left(\frac{1}{\lambda_e}\frac{d}{ds} + \mathrm{grad}_{F_\Gamma}\right) u|T\right)$$

where s is a local coordinate on the tree components with unit speed, and

$$\lambda_e = \begin{cases} \ell(e), & e \in \mathrm{Edge}^{(0,\infty)}_{\circ,-}(\Gamma), \\ 1, & e \in \mathrm{Edge}^{\infty}_{\circ}(\Gamma). \end{cases}$$

is a factor accounting for the length of treed segments in case the edge length is finite and non-zero. The evaluation maps at lifts of nodal points, markings, and lifts of $S \cap T$ give a smooth map

$$\mathrm{ev}_\Gamma : \mathcal{B}^i_{p,\lambda,l,\Gamma} \to \mathfrak{X}(\Gamma) \tag{6.28}$$

where

$$\mathfrak{X}(\Gamma) = \left(\prod_{e \in \mathrm{Edge}_{\mathrm{trop}}(\Gamma)} (X^{\square}_{\overline{P}(e)}/T_{\mathcal{T}(e),\mathbb{C}})^2 \right) \times \left(\prod_{e=(v_+,v_-) \in \mathrm{Edge}_{\mathrm{int},\bullet}(\Gamma)} (X_{\overline{P}(v_\pm)})^2 \right) \times \left(\prod_{x \in S \cap T} L^2 \right) \times \left(\prod_{e \in \mathrm{Edge}^0_{\circ}(\Gamma)} L^2 \right) \times \left(\prod_{e \in \mathrm{Edge}_{\bullet,\to}(\Gamma)} X^{\square}_{\overline{P}(v(e))} \right). \tag{6.29}$$

The first two factors in (6.29) correspond to lifts of interior nodes, the third term is a lift of boundary nodes w_e with no treed segments (that is, $\ell(e) = 0$), the fourth term is a lift of $S \cap T$, and the last term corresponds to evaluation at an interior marking. All the factors of ev_Γ are standard evaluation maps, except for the first factor which is the projected tropical evaluation map as in (4.20). Let

$$\Delta(\Gamma) \subset \mathfrak{X}(\Gamma)$$

be the submanifold that is the product of diagonals in the first four factors of $\mathfrak{X}(\Gamma)$ in (6.29), and the stabilizing divisor $D_{P(v(e))} \subset X^{\square}_{\overline{P}(v(e))}$ in the last factor. The *local universal moduli space* is

$$\mathcal{M}^{\mathrm{univ},i}_{\Gamma}(L, \mathfrak{D}) = (\overline{\partial}, \mathrm{ev}_\Gamma)^{-1}(\mathcal{B}^i_{p,\lambda,l,\Gamma}, \Delta(\Gamma)), \tag{6.30}$$

where $\mathcal{B}^i_{k,p,l,\Gamma}$ is embedded as the zero section in $\mathcal{E}^i_{p,\lambda,\Gamma}$.

We will next show that this subspace is cut out transversely. We first consider two-dimensional components of $\tilde{C}$ on which the map is not horizontally constant, and show that the linearization of $(\overline{\partial}, \mathrm{ev}_\Gamma)$ is surjective. For components whose target space is not a neck piece, the surjectivity of the differential $D(\overline{\partial}, \mathrm{ev}_\Gamma)$ follows from [**60**, Proposition 3.4.2]. Recall from (3.10) that a neck piece $X^{\square}_{\overline{P}}$ is a fibration

$$V_{P^\vee} \to X^{\square}_{\overline{P}} \to X^{\square}_{P}$$

whose fiber is a T_P-toric manifold $V_{P^\vee}$, and the almost complex structure on $X^{\square}_{\overline{P}}$ is P-cylindrical. For non-neck pieces T_P is trivial. Consider a component $S^\circ_v \simeq \mathbb{P}^1 \backslash \{\text{special points}\}$ that maps to a neck piece $X_{\overline{P}}$, and the horizontal projection $\pi_P \circ u : S^\circ_v \to X^{\square}_P$ is non-constant. The linearized operator

$$D^\circ_{u,J}(\xi, K) = D^\circ_u \xi + \tfrac{1}{2} K D u j.$$

is surjective as follows. Let

$$\eta \in \mathrm{coker}(D^\circ_{u,J}) \subset \Omega^{0,1}(u^* T X^{\square}_{\overline{P}})$$

be a one-form in the cokernel of $D^\circ_{u,J}$. Variations of tamed almost complex structure of cylindrical type are J-antilinear maps

$$K : T X^{\square}_{\overline{P}} \to T X^{\square}_{\overline{P}}$$

that vanish on the vertical sub-bundle and are $T_{P,\mathbb{C}}$-invariant.[1] Since the horizontal part of $D_z u$ is non-zero at some $z \in S^\circ_v$, we may find an infinitesimal variation K of

[1]Perturbing K is equivalent to perturbing the connection one-form α_P of the P-cylindrical almost complex structure, see (3.11).

almost complex structure *of cylindrical type* by choosing $K(z)$ so that $K(z)D_z uj(z)$ is an arbitrary $(j(z), J(z))$-antilinear map from T_zC to $T_{u(z)}X^{\square}_{\overline{P}}$. Choose $K(z)$ so that $K(z)D_z uj(z)$ pairs non-trivially with $\eta(u(z))$ and extend $K(z)$ to an infinitesimal almost complex structure K by a cutoff function on the domain curve. For two-dimensional components that are horizontally constant, by Corollary 6.35 the linearization $d\overline{\partial}$ is surjective, and additionally the evaluation map at a single marked point is surjective. Note that $d\,\mathrm{ev}_\Gamma$ may not be surjective on these components. Finally, for tree components, the linearization of the shifted gradient operator and the evaluation map at the finite end is surjective, since we can perturb the Morse function on the Lagrangian.

From the discussion so far, we conclude that the local moduli space (6.30) is cut out transversely except if there is a connected component in the domain on which the map is horizontally constant, and which contains more than one irreducible surface component. In this exceptional case, it remains to prove that the matching conditions at nodes between two horizontally constant components are cut out transversely. Such nodes are necessarily internal interior nodes, corresponding to edges $e \in \mathrm{Edge}_{\bullet,\mathrm{int}}(\Gamma)$. We consider a maximal connected subgraph $\Gamma' \subset \Gamma$ so that the map is horizontally constant on the vertices of Γ', and Γ' does not have tree components. By uncrowdedness, there is at most one marked point in Γ'. So, it is possible to choose at most one special point (marked point or a lift of a nodal point) on each component of $\tilde{C}_{\Gamma'}$, so that for every nodal point one of its ends is chosen. By Corollary 6.35, for a horizontally constant component with a single marked point z, the linearized map $D(\overline{\partial}, \mathrm{ev}_z)$ is surjective. Since the evaluation map is surjective at each of the chosen lifts, an inverse of the linearized map $D(\overline{\partial}, \mathrm{ev}_\Gamma)$ can be constructed inductively, see [**18**, p63].

For types where the map has higher order intersections with the stabilizing divisor, the universal moduli space is cut out inductively as in [**22**, Lemma 6.5]. Each step of the induction cuts out a moduli space where the tangencies at one of the markings is increased by one. We start out with a moduli space cut out of $W^{k,p,\lambda}$ where $k - \frac{2}{p} > \mu$ and μ is the largest order of tangency with the stabilizing divisor $\mathfrak{D}$ that occur in the type Γ.

By the implicit function theorem, $\mathcal{M}^{\mathrm{univ},i}_\Gamma(L, \mathfrak{D})$ is a smooth Banach manifold, and the forgetful morphism

$$\varphi_i : \mathcal{M}^{\mathrm{univ},i}_\Gamma(L, \mathfrak{D})_{k,p,l} \to \mathcal{P}_\Gamma(L, \mathfrak{D})_l$$

is a smooth Fredholm map. By the Sard-Smale theorem, the set of regular values $\mathcal{P}^{i,\mathrm{reg}}_\Gamma(L, D)$ of φ_i on $\mathcal{M}^{\mathrm{univ},i}_\Gamma(L, \mathfrak{D})_d$ in $\mathcal{P}_\Gamma(L, \mathfrak{D})$ is comeager. Let

$$\mathcal{P}^{\mathrm{reg}}_\Gamma(L, \mathfrak{D}) = \bigcap_i \mathcal{P}^{i,\mathrm{reg}}_\Gamma(L, \mathfrak{D}).$$

A standard argument shows that the set of smooth domain-dependent $\mathcal{P}^{\mathrm{reg}}_\Gamma(L, \mathfrak{D})$ is also comeager. Fix $(J_\Gamma, F_\Gamma) \in \mathcal{P}^{\mathrm{reg}}_\Gamma(L, \mathfrak{D})$. By elliptic regularity, every element of $\mathcal{M}^i_\Gamma(L, \mathfrak{D})$ is smooth. The transition maps for the local trivializations of the universal bundle define smooth maps

$$\mathcal{M}^i_\Gamma(L,\mathfrak{D})|_{\mathcal{M}^i_\Gamma \cap \mathcal{M}^j_\Gamma} \to \mathcal{M}^j_\Gamma(L,\mathfrak{D})_{\mathcal{M}^i_\Gamma \cap \mathcal{M}^j_\Gamma}.$$

This construction equips the space

$$\mathcal{M}_\Gamma(L,\mathfrak{D}) = \cup_i \mathcal{M}^i_\Gamma(L,\mathfrak{D})$$

with a smooth atlas. Since $\mathcal{M}_\Gamma$ is Hausdorff and second-countable, so is $\mathcal{M}_\Gamma(L,\mathfrak{D})$ and it follows that $\mathcal{M}_\Gamma(L,\mathfrak{D})$ has the structure of a smooth manifold. Orientation of the moduli space is discussed in Remark 6.30 below. □

For regular perturbations, the tangent space to the moduli space can be identified with the kernel of an operator called the *linearized operator*. Using notations from the proof of transversality, we write down the linearized operator for future reference. For a regular perturbation $\mathfrak{p}_\Gamma$, consider the map

$$(\overline{\partial}, \mathrm{ev}_\Gamma) : \mathcal{B}^i_{k,p,\Gamma} \to \mathcal{E}^i_{k,p,\Gamma} \times \mathfrak{X}(\Gamma)$$

where ev_Γ is defined in (6.28). Its linearization is denoted

$$D_u : T_{[C],u}\mathcal{B}^i_{k,p,\Gamma} \to (\mathcal{E}^i_{k,p,\Gamma})_{[C],u} \oplus \mathrm{ev}^*_\Gamma T\mathfrak{X}(\Gamma)/T\Delta(\Gamma). \tag{6.31}$$

For a regular broken map u the operator D_u is surjective.

Remark 6.30. (Orientation of moduli spaces of maps) Moduli spaces of broken maps are oriented using a relative spin structure on the Lagrangian as in Fukaya-Oh-Ohta-Ono [**34**] as follows. An orientation for the moduli space of maps consists of a choice of a section of the determinant line bundle of the Fredholm operator D cutting out the moduli space, which is defined as $\det(D) := \Lambda^{top}(\ker(D)) \wedge \Lambda^{top}(\mathrm{coker}(D))$. Let L be an oriented Lagrangian equipped with a relative spin structure (see [**34**, Chapter 44] for the definition). Given a holomorphic disk $u : (\mathbb{D}^2, \partial\mathbb{D}^2) \to (X, L)$ and a complex linear Cauchy-Riemann operator

$$\overline{\partial} : \Gamma(\mathbb{D}^2, \partial\mathbb{D}^2; u^*TX, (\partial u)^*TL) \to \Omega^{0,1}(\mathbb{D}^2, u^*TX),$$

the determinant line bundle can be identified with $\det(T_lL)$ for any point $l \in L$

$$\det(\overline{\partial}) \simeq \det(T_lL),$$

and the identification is canonical upto multiplication by positive scalars $\lambda \in \mathbb{R}_+$, see [**34**, Proposition 44.4] or an explanation in [**19**, Proposition 5.2]. The argument in [**34**, Proposition 44.4] also gives such an identification in case of a nodal disk with interior nodes, since that involves degenerating the bundle u^*TX into a bundle on the disk and a complex vector bundle on the sphere. Since the Fredholm operator for a broken map differs only in the matching condition being defined on a quotient of X by a complex torus, the determinant bundle can be identified to $\det(T_lL)$ in a similar way. Finally the line bundle $\det(\overline{\partial})$ extends to other strata of treed nodal disks as follows. Let Γ be a type of a treed nodal curve with an edge e of finite non-zero length, Γ_0 be the type obtained by making $\ell(e) = 0$ and let Γ_c be obtained from Γ_0 by collapsing e. Then the moduli space of maps $\mathcal{M}_\Gamma(\mathfrak{X})$ can be glued to $\mathcal{M}_{\Gamma_c}(\mathfrak{X})$ along the codimension one boundary $\mathcal{M}_{\Gamma_0}(\mathfrak{X})$. Therefore, if the determinant bundle has an orientation on $\mathcal{M}_{\Gamma_c}(\mathfrak{X})$, it extends to $\mathcal{M}_\Gamma(\mathfrak{X})$.

It remains to consider contributions to the orientation from the infinite treed segments attached at boundary markings. We choose orientations on Morse unstable and stable manifolds

$$o(x) : W^{\pm}(x) \to TW^{\pm}(x), \quad \forall x \in \mathrm{crit}(F)$$

for the Morse function F on the Lagrangian, such that for any $x \in \mathrm{crit}(F)$ the map

$$\det(T_x W^-(x) \oplus T_x W^+(x)) \to \det(T_x L)$$

is orientation preserving. We view a broken treed holomorphic disk as defined by a condition $u(z_i^\circ) \in W^-(x_i)$ on each of the boundary markings $e_i^\circ \in \mathrm{Edge}_{\circ,\to}$. We thus obtain an isomorphism of line bundles

$$\det(D_u) \simeq \det(T\mathcal{M}_\Gamma) \wedge \det(TL) \wedge \det(W^+(x_0)) \wedge (\wedge_{i=1}^{d(\circ)} \det(W^-(x_i))),$$

up to multiplication by positive scalars. The choices made in the previous paragraphs fixes a positive section on the right-hand side, which induces an orientation on the moduli space $\mathcal{M}_\Gamma(\mathfrak{X}, \underline{x})$ for any type Γ of broken maps and a collection of inputs and output $\underline{x}$. The orientation of $\det(T\mathcal{M}_\Gamma)$ is described in Definition 4.2.

Finally, for a broken map type Γ with a broken edge, the boundary orientation induced by (Making an edge length finite) morphism and the product orientation induced by (Cutting an edge) morphism differ by a quantity that depends on the type Γ and the Morse indices of the labels on the end-points of the treed segments, see [**82**, (12.25)]. This ends the Remark.

6.7. The toric case

In this section, we consider the question of regularity for maps in a toric variety with the standard complex structure. The results are also useful for maps in toric fibrations that project to a constant in the base space.

DEFINITION 6.31. A component $X_{\overline{P}} \subset \mathfrak{X}_{\mathcal{P}}$ of a broken manifold is a *tropical toric piece* if

(a) $P \in \mathcal{P}$ is a top-dimensional polytope,
(b) the tropical moment map Φ generates a Hamiltonian T-action which makes X_P a non-singular toric variety with moment polytope P,
(c) all torus-invariant divisors of X_P are relative divisors of the broken manifold $\mathfrak{X}$ (see Definition 3.14 (b)), or in other words, if Q is a facet of the polytope P then $Q \in \mathcal{P}$.

In a tropical toric piece X_P, any Lagrangian torus orbit $L \subset X_P$ is called a *tropical fiber.*

Recall from Definition 6.17 that the domain of a relative map has a single surface component and no tree components. Thus the type of the map prescribes the order of intersection of the map with relative divisors at each of the marked points.

PROPOSITION 6.32. *Let $X := X_P \subset \mathfrak{X}$ be a tropical toric manifold as in Definition 6.31, and let $L \subset X_P$ be a Lagrangian tropical fiber. Let Γ be a type of relative map (see definition 6.17). Any map u of type Γ is regular, and therefore, the moduli space of maps $\mathcal{M}_\Gamma(X)$ is a manifold of expected dimension.*

Furthermore, let Γ *be a type corresponding to maps only containing sphere components. Let* z_e *be a marked point corresponding to an edge* $e \in \mathrm{Edge}_\bullet(\Gamma)$ *that maps to a torus-invariant submanifold* $Y \subset X$*. Then, the projected tropical evaluation map at* z_e

$$\pi^\perp_{\mathcal{T}(e)}(\mathrm{ev}^{\mathcal{T}(e)}_{z_e}) : \mathcal{M}_\Gamma(X) \to T_\mathbb{C}/T_{\mathcal{T}(e),\mathbb{C}}$$

is submersive.

Remark 6.33. The maps considered in Proposition 6.32 can not map to a torus-invariant divisor of X. Indeed since all such divisors are relative divisors of the broken manifold, a relative map can intersect those divisors only at relative marked points.

Remark 6.34. (The orbifold case) It is enough to prove Proposition 6.32 assuming that the target space X_P is a manifold. In case X_P has toric orbifold singularities, we may replace X_P by a toric resolution $\tilde{X}_P$ and the manifold proof carries over. Indeed, the Fredholm properties of the linearized operator are the same in both cases, because either of them is equivalent to the Fredholm operator D°_u defined using a target space $X_P^\square$ with cylindrical ends. In fact both X_P and $\tilde{X}_P$ are complex compactifications of $X_P^\square$.

Proof of Proposition 6.32. As explained in Remark 6.34 above, we may assume that X is a toric manifold corresponding to a Delzant moment polytope $\Delta \subset \mathfrak{t}^\vee$. Each facet $Q_1, \ldots, Q_N$ defines a prime $T_\mathbb{C}$-invariant divisor $X_{Q_1}, \ldots, X_{Q_N}$ in X, whose union is a representative of the first Chern class.

Holomorphic disks to X with boundary on L are regular by Cho-Oh [**20**]. We view the toric variety as a quotient of a vector space $\mathbb{C}^N$ and the Lagrangian torus as a quotient of the standard torus in $\mathbb{C}^N$. Then the disks bounding the torus are Blaschke products, they are defined on the unit disk $\mathbb{D} \subset \mathbb{C}$ and are of the form

$$u : \mathbb{D} \to \mathbb{C}^N, \quad z \mapsto \left(\zeta_i \prod_{j=1}^{d_i} \frac{z - a_{i,j}}{1 - z\overline{a_{i,j}}} \right)_{i=1,\ldots,N}. \tag{6.32}$$

for some constants $a_{i,j}, \zeta_i \in \mathbb{C}, |a_{i,j}| \leq 1, |\zeta_i| = 1$. A Maslov index two Blaschke disk has a single non-vanishing degree d_i equal to 1. Since the intersection points with the divisor are given by roots of the components and can be chosen arbitrarily, the evaluation maps are transverse to any constraints.

Holomorphic spheres meeting the interior of X are also regular by the following argument. As in Delzant [**25**], X can be viewed as a geometric invariant theory quotient $\mathbb{C}^N /\!\!/ G$ where N is the number of prime torus-invariant relative divisors of X, and $G \subset (\mathbb{C}^\times)^N$ is a complex torus whose quotient $(\mathbb{C}^\times)^N/G$ is $T_\mathbb{C}$. Each of the relative divisors $X_{Q_1}, \ldots, X_{Q_N}$ in X lifts to a coordinate hyperplane $\{z_1 = 0\}, \ldots, \{z_N = 0\}$ in $\mathbb{C}^N$. Consider a holomorphic sphere $u : \mathbb{P}^1 \to X$, that is not contained in any toric divisor. The vector bundle $u^*TX \oplus \underline{\mathfrak{g}}$ on $\mathbb{P}^1$ is a sum of line bundles

$$u^*TX \oplus \underline{\mathfrak{g}} = \bigoplus_{i=1}^N u^*\mathcal{O}(X_{Q_i})$$

where $\underline{\mathfrak{g}} := \mathfrak{g} \times \mathbb{P}^1$ is the trivial bundle. The degree $\deg(u^*\mathcal{O}(X_{Q_i}))$ of the line bundle $u^*\mathcal{O}(X_{Q_i})$ is given by the intersection of u with X_{Q_i}. Hence each of the degrees $\deg(u^*\mathcal{O}(X_{Q_i}))$ is non-negative. As a result, the operator

$$\overline{\partial} : \Gamma(\mathbb{P}^1, \oplus_i u^*\mathcal{O}(X_{Q_i})) \to \Omega^{0,1}(\mathbb{P}^1, \oplus_i u^*\mathcal{O}(X_{Q_i}))$$

is onto. Consequently the cohomology group

$$H^{0,1}(\mathbb{P}^1, \oplus_i u^*\mathcal{O}(X_{Q_i})) \simeq H^1(\mathbb{P}^1, \oplus_i u^*\mathcal{O}(X_{Q_i}))$$

vanishes. Consider the long exact sequence in Čech cohomology, corresponding to the short exact sequence of sheaves

$$0 \to \underline{\mathfrak{g}} \to \oplus_i u^*\mathcal{O}(X_{Q_i}) \to u^*TX \to 0.$$

Vanishing of the zeroth resp. first cohomology of the first resp. second terms implies that the first cohomology of the third term

$$H^{0,1}(\mathbb{P}^1, u^*TX) \simeq H^1(\mathbb{P}^1, u^*TX)$$

also vanishes. Therefore the sphere u is regular in X.

Next, we will show that the moduli spaces of spheres or disks with prescribed tangencies at relative marked points are cut out transversely. The Maslov index $I(u)$ of a holomorphic disk or sphere $u : C \to X$ (not contained in a relative divisor) is twice the sum of its intersection multiplicities $(u.X_Q)$ with all relative divisors X_Q. In fact, the moduli space containing u is parametrized by the set of intersection points $z \in C$ with the toric divisors. For maps with a higher order tangency with a toric divisor X_Q, some subsets of intersection points of u with X_Q coincide. Since the zeros of sections of positive line bundles may be chosen arbitrarily in genus zero, the locus of such maps is cut out transversely from the space of all maps. Therefore the moduli space of relative maps of any type Γ is regular.

The linearization of the evaluation map at relative marked points is surjective because of the torus-equivariance of the evaluation map. Recall that if the marked point z_e maps to an intersection of relative divisors $Y = \cap_i Y_i$, the target space for a higher order evaluation map is the normal bundle NY. The normal bundle can be identified with a neighbourhood of Y in a standard way. With this identification,

$$\mathrm{ev}_{z_e}^{\mathcal{T}(e)}(tu) = t\,\mathrm{ev}_{z_e}^{\mathcal{T}(e)}(u), \quad \forall t \in T_{\mathbb{C}}.$$

A $T_{\mathbb{C}}$-orbit is an open set in X, and therefore $\mathrm{ev}_{z_e}^{\mathcal{T}(e)}(u)$ is a regular value of the map $\mathrm{ev}_{z_e}^{\mathcal{T}(e)}$. □

Proposition 6.32 implies an analogous result for toric fibrations, for which all torus-invariant divisors are relative divisors. This result was used to prove transversality for horizontally constant components $u : \mathbb{P}^1 \to X_{\overline{P}}$ in a broken map. The target space $X_{\overline{P}}$ is a fibration $X_{\overline{P}} \to X_P$, whose fibers are smooth symplectic toric varieties with the action of a compact torus T_P. The torus action on the fibers is the restriction of a T_P-action on the total space $X_{\overline{P}}$. The torus-invariant divisors of $X_{\overline{P}}$ are precisely the vertical divisors of $X_{\overline{P}}$, and each of these divisors is a relative divisor (as in Definition 3.14 (b)). Indeed every facet of a polytope $P^\vee$ in the dual complex $B^\vee$ is a polytope $P_0^\vee$ for some $P_0 \in \mathcal{P}$.

Corollary 6.35. *Suppose* $P \in \mathcal{P}$ *is a polytope and the toric fibration* $X_{\overline{P}} \to X_P$ *has an almost complex structure that is standard on the fibers. Let* Γ *be a type of relative map that has a single sphere component* $S_v, v \in \mathrm{Vert}(\Gamma)$ *mapping to* $X_{\overline{P}}$*, which projects to a constant* $\pi_P \circ u_v : S_v \to X_P$*, and has no tree components. Any map* u *of type* Γ *is regular, and therefore, the moduli space of maps* $\mathcal{M}_\Gamma(X_{\overline{P}})$ *is a manifold of expected dimension. For a marking* z_e *corresponding to any edge* $e \in \mathrm{Edge}_\bullet(\Gamma)$ *that maps to a torus-invariant submanifold* $Y \subset X_{\overline{P}}$*, the projected tropical evaluation map*

$$\pi^\perp_{\mathcal{T}(e)} \circ \mathrm{ev}^{\mathcal{T}(e)}_{z_e} : \mathcal{M}_\Gamma(X) \to T_{\mathbb{C}}/T_{\mathcal{T}(e),\mathbb{C}} \oplus TX_P$$

is submersive.

CHAPTER 7

Hofer energy and exponential decay

Compactness for sequences of broken pseudoholomorphic maps requires bounds on area. Neck-stretched manifolds do not have a naturally defined taming symplectic form. However, since neck-stretched manifolds are diffeomorphic to a tropical Hamiltonian manifold $(X, \omega, \mathcal{P}, \Phi)$, they are equipped with a cohomology class $[\omega]$. Hofer energy of a pseudoholomorphic curve is defined as a supremum of 'symplectic areas' where the symplectic form ranges over a family of taming forms in the class $[\omega]$. Each symplectic form in the family is given by a map of complexes

$$\aleph : B_J \to B_\omega$$

between the *J-complex* underlying the cylindrical almost complex structure and the *ω-complex* underlying the symplectic form on the tropical manifold. For example, in a neck-stretched manifold X^ν the ω-complex is the dual complex $B^\vee$ and the J-complex is $\nu B^\vee$. Such a map $\aleph$ between complexes gives a map between manifolds

$$\psi_\aleph : (X^\nu, J^\nu) \to (X, \omega)$$

which is taming (that is, J^ν is $\psi_\aleph^* \omega$-tame) if $\aleph$ satisfies certain conditions. In case of a single cut, $\psi_\aleph$ is taming if $\aleph : [\frac{-\nu}{2}, \frac{\nu}{2}] \to [0, 1]$ is an increasing diffeomorphism. This notion of energy as a supremum over a family of symplectic areas was originally defined by Hofer in the context of symplectic field theory [49]. The taming condition is more complicated for a multiple cut and leads us to define a class of maps between complexes, called *squashing maps*. A squashing map is a continuous piecewise smooth map between complexes, which on any piece, is the composition of a translation, dilation and an orthogonal projection. For a pseudoholomorphic map $u : C \to X^\nu$ in a neck-stretched manifold (or a broken manifold), we define the *Hofer energy* as

$$E_{\mathrm{Hof}}(u) := \sup_{\aleph \text{ is a squashing map}} \int_C u^*(\psi_\aleph^* \omega).$$

Of course, if C is a closed curve, or if C is a disk and u maps the boundary ∂C to a Lagrangian submanifold, the value of the integral is independent of $\aleph$ since the cohomology class of $\psi_\aleph^* \omega$ is independent of $\aleph$.

The main result of this chapter is that a punctured pseudoholomorphic curve with finite Hofer energy has a removable singularity :

PROPOSITION 7.1. (Removal of singularities) *Suppose $u : \mathbb{D}^2 \backslash \{0\} \to X_P^\square$ is a perturbed J-holomorphic curve with respect to the domain-dependent almost complex structure $J : \mathbb{D}^2 \to U_{\mathfrak{J}_0}$*

$$E_{\mathrm{Hof}}(u) < \infty, \quad \|du\|_{L^\infty(\mathrm{Cyl})} < \infty. \tag{7.1}$$

Then, u extends to a holomorphic map $u : \mathbb{D}^2 \to X_{\overline{P}}$ possibly, in the orbifold case, after passing to a finite cover.

7.1. Symplectic forms on neck-stretched manifolds

In this section we define a family of symplectic forms on neck-stretched manifolds that tame cylindrical almost complex structures that are locally strongly tame. For a pseudoholomorphic disk $u : (C, \partial C) \to (X^\nu, L)$ in a neck stretched manifold X^ν whose boundary maps to a Lagrangian submanifold $L \subset X^\nu$, the symplectic area

$$\mathrm{Area}(u) := \langle u_*[C], [\omega] \rangle,$$

is well-defined since a neck-stretched manifold X^ν is diffeomorphic (not canonically) to a tropical manifold (X, ω). Our task is to construct a family of symplectic forms in the class $[\omega] \in H^2(X^\nu)$.

We recall some details from the definition of neck-stretched manifolds from Chapter 3 to justify why the ω-complex and J-complex are $B^\vee$ and $\nu B^\vee$ respectively. We recall from (3.13) that on a tropical manifold $(X, \omega, \mathcal{P}, \Phi)$ there is a symplectic cylindrical structure on neck regions, that is, there is a T_P-equivariant symplectomorphism

$$(X, \omega) \supset \Phi^{-1}(\overline{P}) \xrightarrow{\phi_P} (\Phi^{-1}(P) \times P^\vee, \overline{\omega}_P), \qquad \overline{\omega}_P := (\omega_X | \Phi^{-1}(P)) + d\langle \alpha_P, \pi_{P^\vee} \rangle, \tag{7.2}$$

for each polytope $P \in \mathcal{P}$, where $\overline{P} \subset \mathfrak{t}^\vee$ is a tubular neighborhood of the polytope $P \subset \mathfrak{t}^\vee$ with projection (see Figure 3.3)

$$\pi_P : \overline{P} \to P,$$

$\alpha_P \in \Omega^1(\Phi^{-1}(P), \mathfrak{t}_P)$ is a T_P-connection one-form. The projection $\pi_{P^\vee} : \Phi^{-1}(\overline{P}) \to P^\vee \subset \mathfrak{t}_P^\vee$ of the moment map Φ to $\mathfrak{t}_P^\vee$ is a moment map for the T_P-action. Let $P^\blacksquare \subset P$ be the complement of a neighborhood of faces of P, namely

$$P^\blacksquare := P \backslash (\cup_{Q \subset P} \overline{Q}), \tag{7.3}$$

and let

$$\overline{P^\blacksquare} := \pi_P^{-1}(P^\blacksquare) \subset \overline{P} \tag{7.4}$$

be the thickening of $P^\blacksquare$. For a pair $Q \subset P$ with $\mathrm{codim}_P(Q) = 1$, the fibered polytopes $\overline{P^\blacksquare}, \overline{Q^\blacksquare} \subset \mathfrak{t}^\vee$ share a facet, which is isomorphic to $Q^\blacksquare \times P^\vee$. Then the image of Φ has a cover

$$\mathrm{im}(\Phi) = \cup_{P \in \mathcal{P}} \mathrm{i}_P(\overline{P^\blacksquare}) / \sim, \tag{7.5}$$

where $\sim$ identifies shared facets of polytopes, see Figure 7.1. The partition of $\mathrm{im}(\Phi)$ pulls back to a partition of the manifold X

$$(X, \omega) \simeq \left(\bigsqcup_{P \in \mathcal{P}} \Phi^{-1}(\overline{P^\blacksquare}) \right) / \sim \tag{7.6}$$

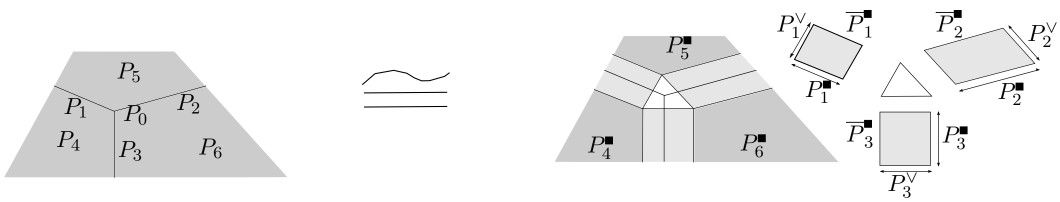

FIGURE 7.1. Decomposition of moment polytope induced by thickenings of polytopes in $\mathcal{P}$.

into manifolds with corners, where the identifications in $\sim$ are along the boundaries and are induced by the inclusions $\Phi^{-1}(\overline{P}^{\blacksquare}) \to X$. Further, the symplectic cylindrical structure map $\underline{\phi} = (\phi_P)_P$ may be used to rewrite the decomposition in (7.6) as

$$(X, \omega) \simeq \left(\bigsqcup_{P \in \mathcal{P}} \Phi^{-1}(P^{\blacksquare}) \times P^{\vee} \right) / \sim . \tag{7.7}$$

(See the derivation of (3.21) for more details.) We have thus shown the following:

LEMMA 7.2. (ω-complex) *There is a continuous projection map*

$$\pi^{\omega}_{B^{\vee}} : (X, \omega) \to B^{\vee} \tag{7.8}$$

such that for any polytope $P \in \mathcal{P}$, $\pi_{B^{\vee}}$ is a T_P-moment map for the T_P-action on $\pi^{-1}_{B^{\vee}}(P^{\vee,\circ})$.

PROOF. The map $\pi^{\omega}_{B^{\vee}}$ is defined by projecting $\Phi^{-1}(P^{\blacksquare}) \times P^{\vee} \subset X$ to $P^{\vee}$ for each $P \in \mathcal{P}$. $\square$

Thus $P^{\vee}$ is the ω-polytope on $\Phi^{-1}(P^{\blacksquare}) \times P^{\vee} \subset X$, and the union $B^{\vee} = \cup_{P \in \mathcal{P}} P^{\vee}$ is the *ω-complex* for X. As an almost complex manifold, the neck-stretched manifold X^{ν} is the union

$$X^{\nu} = \left(\bigsqcup_{P \in \mathcal{P}} \Phi^{-1}(P^{\blacksquare}) \times \nu P^{\vee} \right) / \sim \tag{7.9}$$

where the equivalence $\sim$ from (3.23) identifies boundary components. Analogous to Lemma 7.2, neck-stretched manifolds project to the scaled dual complex:

LEMMA 7.3. (J-complex) *For any $\nu \geq 1$ there is a continuous projection map*

$$\pi_{\nu B^{\vee}} : X^{\nu} \to \nu B^{\vee} \tag{7.10}$$

such that for any polytope $P \in \mathcal{P}$, $(\pi_{\nu B^{\vee}})^{-1}(P^{\vee,\circ})$ is a P-cylinder $\Phi^{-1}(P^{\blacksquare}) \times \nu P^{\vee,\circ}$.

Thus the J-polytope for the P-cylindrical subset of X^{ν} is $\nu P^{\vee}$, and the *J-complex* of X^{ν} is $\nu B^{\vee}$.

Neck-stretched almost complex manifolds are mapped to compact symplectic manifolds via maps of complexes between the J-complex of the neck-stretched manifold and the ω-complex of the symplectic manifold:

Definition 7.4. For any $\nu \geq 1$ a map

$$\aleph : \nu B^\vee \to B^\vee$$

is a *map of complexes* if it is a collection of maps of polytopes, that is, for any $P \in \mathcal{P}$, $\aleph(\nu P^\vee) \subseteq P^\vee$.

A map $\aleph$ of complexes induces a map of manifolds

$$\psi^\nu_\aleph : X^\nu \to (X, \omega), \tag{7.11}$$

where, for any $P \in \mathcal{P}^\vee$, $\psi_\aleph$ maps $\nu P^\vee \times \Phi^{-1}(P^\blacksquare)$ to $P^\vee \times \Phi^{-1}(P^\blacksquare)$ by $(\aleph, \mathrm{Id})$. We leave it to the reader to verify that the maps on the subsets patch to yield a continuous map.

The map of complexes must satisfy an (Increasing condition) stated in the Lemma below. In case of a single cut, the increasing condition is equivalent to the condition that the map $\aleph : [\frac{-\nu}{2}, \frac{\nu}{2}] \to [-\frac{1}{2}, \frac{1}{2}]$ is increasing, which is precisely the class of maps $\aleph$ used to define Hofer energy in [**11**].

Lemma 7.5. (Symplectic forms from increasing maps) *Let J^ν be a locally strongly tamed cylindrical almost complex structure on X^ν. Let $\aleph : \nu B^\vee \to B^\vee$ be a map of complexes such that for any polytope $P \in \mathcal{P}$, $\aleph|\nu P^\vee$ is a diffeomorphism onto $P^\vee$ and*

$$\text{(Increasing condition)} \quad \langle (d\aleph|\nu P^\vee)_x(\xi), \xi \rangle > 0, \quad \forall x \in \nu P^\vee,\ \xi \in \mathfrak{t}^\vee_P. \tag{7.12}$$

Then $\psi^\nu_\aleph : (X^\nu, J^\nu) \to (X, \omega)$ is a diffeomorphism and $(\psi^\nu_\aleph)^*\omega$ tames J^ν.

Proof. We recall from Definition 3.28 that local strong tamedness implies that the fibers of the projection map

$$\pi_P : \Phi^{-1}(P^\blacksquare) \times \nu P^\vee \to \Phi^{-1}(P^\blacksquare)/T_P \subset X^\square_P$$

are J^ν-holomorphic, and the horizontal tangent sub-bundle $\ker(\alpha_P) \subset T\Phi^{-1}(P^\blacksquare)$ is J^ν-invariant. On $\Phi^{-1}(P^\blacksquare) \times \nu P^\vee$ the form $(\psi^\nu_{P,\aleph})^*\omega$ is equal to $\omega_{X_P} + d\langle \aleph, \alpha_P \rangle$. Indeed, $\aleph(\nu P^\vee) \subseteq P^\vee \subset \mathfrak{t}^\vee_P$, and therefore $\langle \aleph | \nu P^\vee, \alpha_P \rangle$ is a one-form. On the vertical sub-bundle $\ker(d\pi_P)$, the form $(\psi^\nu_{P,\aleph})^*\omega$ is equal to $\langle d(\aleph|\nu P^\vee), \alpha_P \rangle$, which is taming because of the increasing condition (7.12) :

$$\langle d(\aleph|\nu P^\vee), \alpha_P \rangle (v, J^\nu v) > 0, \quad v \in \ker(d\pi_P).$$

On the horizontal sub-bundle $\ker(\alpha_P) \subset T\Phi^{-1}(P^\blacksquare)$, the form is $\omega_{X_P} + \aleph_P d\alpha_P$ which tames J^ν, since J^ν is locally strongly tamed. □

Increasing maps produce taming forms only if the almost complex structure is locally strongly tamed, and not if the almost complex structure is locally tamed. However, we cannot attain transversality if the domain-dependent complex structures are constrained to be locally strongly tamed. Indeed the connection one-forms underlying the almost complex structures need to be perturbed in order to attain transversality of the linearized $\overline{\partial}$ operator. Therefore, we need the following stronger condition on maps of complexes to produce taming symplectic forms for locally tamed cylindrical almost complex structures on neck-stretched manifolds.

Definition 7.6. (Directionally increasing) A map of complexes $\aleph : \nu B^\vee \to B^\vee$ is *directionally increasing* resp. *strictly directionally increasing* if for any polytope $P \in \mathcal{P}$ and a point $x \in \nu P^\vee$, the derivative $D(\aleph|\nu P^\vee)_x : \mathfrak{t}_P^\vee \to \mathfrak{t}_P^\vee$ is diagonalizable with a set of orthogonal eigenvectors (with respect to the $\mathfrak{t}$-inner product (3.6)), and the eigenvalues $n_1, \dots, n_k$ lie in the interval $[0,1]$ resp. $(0,1]$.

Definition 7.7. An almost complex structure is *weakly taming* if $\omega(v, Jv) \geq 0$ for all tangent vectors v.

Lemma 7.8. *Suppose $\mathfrak{J}_0 \in \mathcal{J}^{\mathrm{cyl}}(\mathfrak{X})$ is a locally strongly tamed cylindrical almost complex structure. There is a C^0-small neighbourhood $U_{\mathfrak{J}_0} \subset \mathcal{J}^{\mathrm{cyl}}(\mathfrak{X})$ of $\mathfrak{J}_0$ such that the following is satisfied for any $\mathfrak{J}_1 \in U_{\mathfrak{J}_0}$ and any $\nu \in [1, \infty)$. Let $\aleph : \nu B^\vee \to B^\vee$ be a map of complexes, and let $\psi^\nu_\aleph : (X^\nu, J^\nu) \to (X, \omega)$ be the resulting map of manifolds.*

(a) *If $\aleph$ is directionally increasing, then the form $(\psi^\nu_\aleph)^*\omega$ weakly tames J^ν_1, where $J^\nu_1 \in \mathcal{J}^{\mathrm{cyl}}(X^\nu)$ is obtained by gluing $\mathfrak{J}_1$ at cylindrical ends.*
(b) *If $\aleph$ is strictly directionally increasing, then the form $(\psi^\nu_\aleph)^*\omega$ tames J^ν_1.*

Proof. We prove the first statement for directionally increasing maps, since the second one is similar. We consider a polytope $P \in \mathcal{P}$, and prove the (Weakly taming) property in the P-cylindrical region

$$Z_{\mathbb{C}} := \Phi^{-1}(P^\blacksquare) \times \nu P^\vee \subset X^\nu.$$

We denote $Z := \Phi^{-1}(P^\blacksquare)$. We recall that for a cylindrical almost complex structure $\mathfrak{J}$, the fibers of the projection

$$T_P \times \nu P^\vee \to Z_{\mathbb{C}} \xrightarrow{\pi_P} Z/T_P \subset X_P^\square$$

are J^ν-holomorphic. Further on $Z_{\mathbb{C}}$, J^ν is determined by its projection to $X_P^\square$, denoted $D\pi_P(J)$, and the associated connection one-form $\alpha_{P,\mathfrak{J}} \in \Omega^1(Z, \mathfrak{t}_P)$ defined by the condition that the horizontal complement of $\ker(D\pi_P)$ given by

$$\ker(\alpha_{P,\mathfrak{J}_1}) \subset TZ \subset TZ_{\mathbb{C}}$$

is J^ν-invariant (also see (3.11)).

We first consider perturbations to the horizontal projection of the almost complex structure $\mathfrak{J}_0$. The horizontal part of the two-form $\psi_\aleph^*\omega$ is $\omega_{X_P} + \aleph \mathrm{d}\alpha_{\mathfrak{J}_0}$, which is a symplectic form since it occurs as a horizontal component of the symplectic form in (7.2); and it tames $D\pi_P(\mathfrak{J}_0)$ since $\mathfrak{J}_0$ is locally tamed (see Remark 3.29 (c)). Since $\aleph$ takes values in a compact set $P^\vee$, and tameness on a symplectic manifold is a C^0-open condition, we conclude the following : There is a constant ϵ_H such that if

$$\|D\pi_P(\mathfrak{J}_1) - D\pi_p(\mathfrak{J}_0)\|_{C^0} < \epsilon_H \tag{7.13}$$

on (Z/T_P) for some $\mathfrak{J}_1 \in \mathcal{J}^{\mathrm{cyl}}(\mathfrak{X})$, then $D\pi_P(\mathfrak{J}_1)$ is tamed by $\omega_{X_P} + \aleph \mathrm{d}\alpha_{\mathfrak{J}_0}$, see Lemma 3.30 for a proof.

Next, we study the effect of changing the connection associated to the almost complex structure. Suppose $\mathfrak{J}_1 \in \mathcal{J}^{\mathrm{cyl}}$ is a cylindrical almost complex structure whose horizontal projection $J_P := D\pi_P(\mathfrak{J}_1)$ is C^0-close to $D\pi_P(\mathfrak{J}_0)$ as in (7.13).

Let $\mathfrak{J}_{10} \in \mathcal{J}^{\text{cyl}}(Z_{\mathbb{C}})$ be a locally strongly tamed almost complex structure whose horizontal projection is the same as that of $\mathfrak{J}_1$, that is, $D\pi_P(\mathfrak{J}_{10}) = J_P$. Denote the P-connection forms by $\alpha_0 := \alpha_{\mathfrak{J}_{10}} = \alpha_{\mathfrak{J}_0}, \alpha_1 := \alpha_{\mathfrak{J}_1}$. The difference

$$A := \alpha_1 - \alpha_0$$

descends to a $\mathfrak{t}_P$-valued one-form on X_P. For a vector $(v,t) \in TZ_{\mathbb{C}}$, where $v \in TZ$, $t \in \mathfrak{t}_P^\vee$, the difference between $\mathfrak{J}_1$ and $\mathfrak{J}_{10}$ is

$$(\mathfrak{J}_1 - \mathfrak{J}_{10})(v,t) = -(A(J_P v_P))_Z + J_F(Av_P)_Z, \quad v_P := D\pi_P(v),$$

where J_F is the complex structure on the fibers of π_P, and for any $\xi \in \mathfrak{t}$, $\xi_Z \in \text{Vect}(Z)$ is the vector field $\xi_Z(z) := \frac{d}{dt}(z\exp(-t\xi))|_{t=0}$ for any $z \in Z$. We write

$$(\psi_\aleph^*\omega)((v,t),\mathfrak{J}_1(v,t)) = (\psi_\aleph^*\omega)((v,t),\mathfrak{J}_{10}(v,t)) + (\psi_\aleph^*\omega)((v,t),(\mathfrak{J}_1 - \mathfrak{J}_{10})(v,t)). \tag{7.14}$$

By (7.13), there is a constant $c > 0$ such that the first term in the right-hand side of (7.14) is bounded on the P-region as

$$(\psi_\aleph^*\omega)((v,t),\mathfrak{J}_{10}(v,t)) \geq c|v_P|^2 + \langle D\aleph_x(\alpha_0(v)),\alpha_0(v)\rangle + \langle D\aleph_x(t),t\rangle. \tag{7.15}$$

In the second term in the right-hand side of (7.14), the difference $(\mathfrak{J}_1 - \mathfrak{J}_{10})(v,t)$ is in the fiber direction. The form $\psi_\aleph^*\omega$ in the fiber is $D\aleph \wedge \alpha_0$ and therefore,

$$(\psi_\aleph^*\omega)((v,t),(\mathfrak{J}_1 - \mathfrak{J}_{10})(v,t)) = \langle \alpha_0(v), D\aleph_x(A(v_P))\rangle + \langle D\aleph_x(t), A(J_P v_P)\rangle \tag{7.16}$$

By the directionally increasing property, $D\aleph_x$ is diagonalizable with orthogonal eigen-vectors and eigenvalues $n_1,\dots,n_k \in [0,1]$. For any element $\xi \in \mathfrak{t}_P$ resp. $\mathfrak{t}_P^\vee$, we denote by $\xi_i \in \mathfrak{t}_P$ resp. $\mathfrak{t}_P^\vee$ the projection of ξ to the i-th eigenspace. We write

$$D\aleph(\xi) = \textstyle\sum_{i=1}^k n_i\xi_i.$$

Using this eigen-decomposition, and the equations (7.14), (7.15) and (7.16), we get

$$\begin{aligned}(\psi_\aleph^*\omega)((v,t),\mathfrak{J}_1(v,t)) &\geq c|v_P|^2 + \textstyle\sum_{i=1}^k n_i(|\alpha_0(v)_i|^2 + |t_i|^2 + \alpha_0(v)_i A(v_P)_i + t_i A(J_P v_P)_i)\\ &\geq \textstyle\sum_{i=1}^k n_i(\frac{c}{k}|v_P|^2 + |\alpha_0(v)_i|^2 + |t_i|^2 + \alpha_0(v)_i A(v_P)_i + t_i A(J_P v_P)_i).\end{aligned}$$

The last two terms are bounded as

$$\alpha_0(v)_i A(v_P)_i \geq -\tfrac{1}{2}(|A|^2|v_P^2| + |\alpha_0(v)_i|^2), \quad t_i A(J_P v_P)_i \geq -\tfrac{1}{2}(|A|^2|v_P^2| + |t_i|^2)$$

where $|A| := \|A\|_{C^0}$. Therefore, $(\psi_\aleph^*\omega)((v,t),\mathfrak{J}_1(v,t))$ is non-negative if $|A|^2 \leq \frac{c}{2k}$, leading to the proof of the weak tamedness of J_1^ν. □

Remark 7.9. For future use, we point out that the following is a consequence of the last paragraph in the proof of Lemma 7.8: In the notation used in the proof, if $|A|^2 \leq \frac{c}{2k}$, then

$$(\psi_\aleph^*\omega)((v,t),\mathfrak{J}_1(v,t)) \geq \tfrac{1}{2}(\psi_\aleph^*\omega)((v,t),\mathfrak{J}_{10}(v,t)). \tag{7.17}$$

7.2. Squashing maps

In the last section we saw 'directionally increasing maps' between complexes of polytopes induce two-forms that are weakly taming for cylindrical almost complex structures in a neighborhood of a locally strongly tamed almost complex structure. In this section we construct a large class of directionally increasing maps, called *squashing maps*. Squashing maps are continuous and piecewise smooth; and on each subset where the map is smooth, it is a composition of an orthogonal projection, a translation and a dilation. In the next section Hofer energy for a multiple cut is defined as the supremum over the weak symplectic areas induced by squashing maps.

The building blocks for squashing maps are 'closest point' maps to polytopes.

Definition 7.10. (Closest point map) Let V be a vector space with an inner product. Let $Q \subset V$ be a compact polytope. The *closest point* map

$$\aleph : V \to Q$$

sends any point $v \in V$ to its nearest point in Q (which is unique by the convexity of Q).

The closest point map is, piecewise, an orthogonal projection composed with a translation. See Figure 7.2 for an example. We consider rational polytopes, and so,

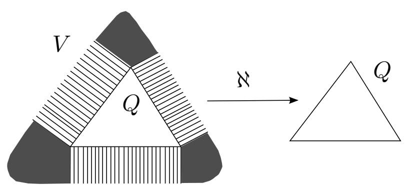

Figure 7.2. The closest point map $\aleph : V \to Q$. Each of the solidly shaded regions is mapped to a vertex of Q, ruled regions are mapped to the sides of Q by contracting each ruling to a point, and the blank region is mapped isometrically to the interior of Q.

the closest-point maps are rational orthogonal projections, as defined next.

Definition 7.11. (Rational projection) Let $V \simeq \mathbb{R}^n$ be a vector space equipped with a dense rational lattice $V_{\mathbb{Q}} \simeq \mathbb{Q}^n$. A *rational subspace* of V is a linear subspace $W \subset V$ in which $W \cap V_{\mathbb{Q}}$ is dense. A linear projection map $\pi : V \to W$ is *rational* if $\ker \pi$ is a rational subspace.

Lemma 7.12. *Let V be a vector space with an inner product and a dense rational lattice $V_{\mathbb{Q}} \subset V$. Let $Q \subset V$ be a compact polytope whose edge slopes are rational. Let $\aleph : V \to Q$ be the closest point map to the compact polytope $Q \subset V$. Then there is a decomposition $V = \cup_{i=1}^k P_i$ such that $\aleph | P_i$ is a rational orthogonal projection composed with a translation.*

Proof. We note that for a codimension k face R of Q and a point $x \in R^\circ$ (here $R^\circ \subset R$ is the complement of the facets of R), the inverse image $\aleph^{-1}(x)$ is contained in an k-dimensional plane through x orthogonal to R. Indeed, for a fixed $y \in \aleph^{-1}(x)$, x is a local minimum in R for the distance from y. Thus $\pi^{-1}(R^\circ)$ is isomorphic to a product $\aleph^{-1}(x) \times R^\circ$ of orthogonal polytopes, on which the map $\aleph$ is an orthogonal projection. The projection is rational since the edge slopes of the polytope are rational. □

Next we define the closest point map to a 'parted polytope'. A parted polytope is obtained by applying a polyhedral decomposition to a compact polytope Q, and then translating the top-dimensional cut spaces away from each other, such that the relative motion for a pair of polytopes with a shared facet is in a direction orthogonal to the facet.

Definition 7.13. (Parted polytope) Let V be a vector space with an inner product. A parted polytope is constructed from the following data:

(a) (Polytope) Let $Q \subset V$ be a compact polytope.
(b) (Polyhedral decomposition) Let $\mathcal{Q}$ be a polyhedral decomposition of Q,

$$Q = \cup_{R \in \mathcal{Q}} R,$$

into polytopes $R \subset V$ such that
- all the edge slopes of the polytopes $R \in \mathcal{Q}$ are rational,
- if Q_0, $Q_1 \in \mathcal{Q}$ have non-empty intersection, then $Q_0 \cap Q_1$ is a face of both Q_0, Q_1 and $Q_0 \cap Q_1 \in \mathcal{Q}$,
- the set of top-dimensional polytopes in $\mathcal{Q}$ is denoted by $\mathcal{Q}^0$, and any lower dimensional polytope in $\mathcal{Q}$ is the intersection of a collection of top-dimensional polytopes. (This decomposition is similar to the polyhedral decomposition of $\mathfrak{t}^\vee$ in Definition 1.1 with the difference that here the polytopes in $\mathcal{Q}$ need not be simplicial.)

(c) (Translations) For each top-dimensional polytope $R \in \mathcal{R}^0$, let $\tau_R \in V$ be a translation such that if $R_+, R_- \in \mathcal{R}^0$ are top-dimensional polytopes sharing a facet $R_0 \in \mathcal{R}$ given by

$$R_0 \simeq \{x \in V : \langle x, \nu_{R_0} \rangle = c_{R_0}\}, \quad \text{for some} \quad \nu_{R_0} \in V, c_{R_0} \in \mathbb{R},$$

with $R_\pm \subset \{\pm(\langle x, \nu_{R_0} \rangle - c_0) \geq 0\}$, then

$$\text{(Orthogonality)} \quad \tau_{R_+} - \tau_{R_-} \in \mathbb{R}_{\geq 0} \nu_{R_0}.$$

Given a polytope Q, a decomposition $\mathcal{Q}$, and translations $(\tau_R)_{R \in \mathcal{R}^0}$ as above, the *parted polytope* $Q(\mathcal{Q}, \tau)$ is the image of the embedding

$$\bigsqcup_{R \in \mathcal{Q}^0} R \to V, \quad R \ni x \mapsto x + \tau_R.$$

The parted polytope $Q(\mathcal{Q}, \tau)$ has a natural projection

$$i_Q : Q(\mathcal{Q}, \tau) \to Q \tag{7.18}$$

that maps $\tau_R + R$ to $R \subset Q$ via translation by $-\tau_R$.

Remark 7.14. The images of the interiors of two distinct top-dimensional polytopes $Q_0, Q_1 \in \mathcal{Q}$ are disjoint in the parted polytope $Q(\mathcal{Q}, \tau)$ as a consequence of the (Orthogonality) condition.

REMARK 7.15. (A dual complex determines orthogonal translations) A dual complex $B_{\mathcal{Q}}^{\vee}$ of the polyhedral decomposition $\mathcal{Q}$ corresponds to a translation satisfying orthogonality: For any top-dimensional polytope $R \in \mathcal{Q}$, define $\tau_R := R^{\vee} \in V$. For any pair $Q_+, Q_- \in \mathcal{Q}^0$ sharing a facet Q_0, the translations τ_{R_+}, τ_{R_-} satisfy the (Orthogonality) condition because the 1-polytope $Q_0^{\vee}$, which connects the points $Q_+^{\vee}, Q_-^{\vee}$ is orthogonal to Q_0. For example, in the parted polytope in Figure 7.3, the dual complex $B_{\mathcal{Q}}^{\vee}$ is the dotted triangle.

DEFINITION 7.16. (Closest point in a parted polytope) Let $Q(\mathcal{Q}, \tau)$ be a parted polytope. The *closest point* map to $Q(\mathcal{Q}, \tau)$ is

$$\aleph = i_Q \circ \aleph_{\text{pre}} : V \to Q$$

where $\aleph_{\text{pre}} : V \to Q(\mathcal{Q}, \tau)$ maps any point $x \in V$ to its closest point in $Q(\mathcal{Q}, \tau)$. We note that the map $\aleph_{\text{pre}}$ is not uniquely defined, but the composition $\aleph$ is uniquely defined. See Figure 7.3 for an example of $\aleph$.

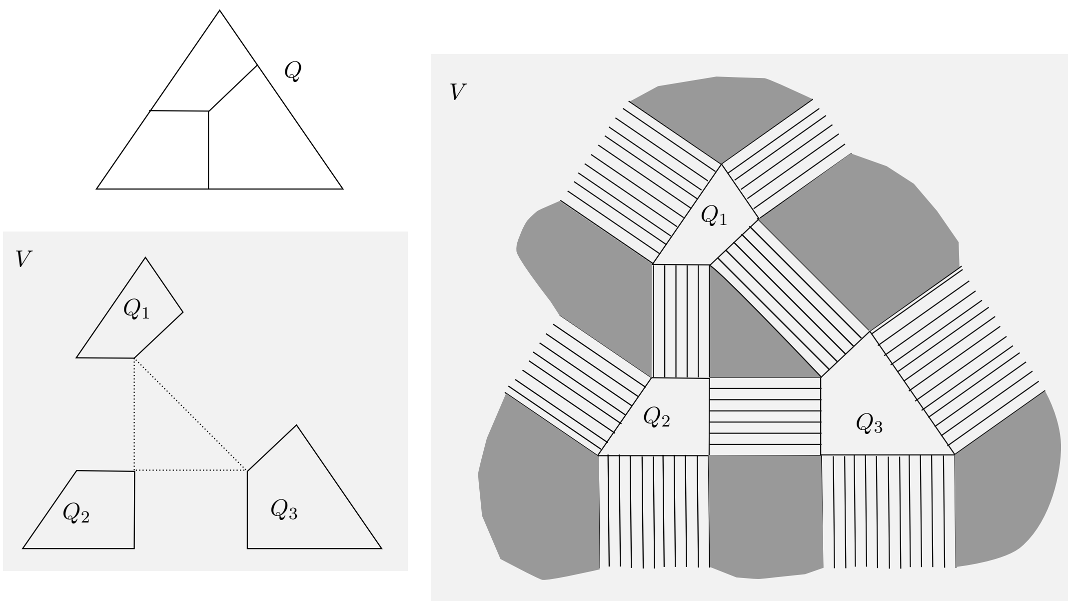

FIGURE 7.3. Top left: A polytope Q with a polyhedral decomposition $\mathcal{Q}$. Bottom left: A parted polytope $Q(\mathcal{Q}, \tau)$ in V. Right: 'The closest point in a parted polytope' map $\aleph : V \to Q$. Solidly shaded regions are mapped to points, ruled regions are mapped to lines by contracting each ruling to a point, and blank regions are mapped isometrically.

LEMMA 7.17. *For a parted polytope $Q(\mathcal{Q}, \tau)$, the closest point map $\aleph : V \to Q$ is well-defined and continuous. Further, then there is a decomposition $V = \cup_{i=1}^{k} P_i$ such that $\aleph|P_i$ is a rational orthogonal projection composed with a translation.*

Proof. To show well-definedness and continuity of $\aleph$, it is enough to focus on points in V which are equidistant from multiple components of the parted polytope. Let $x \in V$, and let $Q_1, \dots, Q_k \in \mathcal{Q}^0$ be top-dimensional polytopes such that $y_i \in \tau_i + Q_i$ is the closest point to x in $\tau_i + Q_i$, and the distances $d(x, y_i + \tau_i)$, $i = 1, \dots, k$ are all the same. The orthogonality condition implies that the $i_Q(y_1), \dots, i_Q(y_k)$ are the same points in Q, and continuity of $\aleph$ follows. The second claim about orthogonal projections follows from Lemma 7.12. □

Definition 7.18. (Squashing maps) Let V be a vector space with an inner product, and let $P, Q \subset V$ be top-dimensional polytopes with rational edge slopes. A *squashing map* $\aleph : P \to Q$ is a composition

$$\aleph := \aleph_0^P \circ \delta_t,$$

of a map $\delta_t : P \to \frac{1}{t}P$ that scales by a factor of $t \geq 1$ and

$$\aleph_0^P : \tfrac{1}{t}P \to Q, \quad \aleph_0^P := \aleph_0|(\tfrac{1}{t}P),$$

where $\aleph_0 : V \to Q$ is the closest point map to some parted polytope $Q(\mathcal{Q}, \tau)$ of Q.

Definition 7.19. (Unpartitioned squashing map) We say that a squashing function $\aleph : P \to Q$ between polytopes P, Q is *unpartitioned* if the underlying partition of the target space Q has a single element, which is Q itself.

7.3. Multi-directional Hofer energy

Hofer energy for a multiple cut is defined as the supremum over the weak symplectic areas induced by squashing maps.

7.3.1. Hofer energy for neck-stretched manifolds. We recall that for a neck-stretched manifold X^ν, the ω-complex is the dual complex $B^\vee$ and the J-complex is $\nu B^\vee$ (see Lemmas 7.2 and 7.3).

Definition 7.20. (Hofer energy for neck-stretched manifolds) Let $\{X^\nu\}_\nu$ be a family of neck-stretched manifolds. For any $\nu \geq 1$, the *Hofer energy* of a map $u : C \to X^\nu$ is

$$E_{\mathrm{Hof}}(u) = \sup_{\aleph : \nu B^\vee \to B^\vee} \int_C (\psi_\aleph \circ u)^* \omega,$$

where the supremum is over all maps of complexes $\aleph : \nu B^\vee \to B^\vee$ for which $\aleph | \nu P^\vee : \nu P^\vee \to P^\vee$ is a squashing map for each polytope $P \in \mathcal{P}$; and $\psi_\aleph : X^\nu \to (X, \omega)$ is the map induced by $\aleph$ as in (7.11). For any squashing map $\aleph$, the form $\psi_\aleph^* \omega$ is called a *squashed area form*. See Figures 7.4 and 7.5 for examples of squashing maps for neck-stretched manifolds.

Remark 7.21. (On squashing maps for neck-stretched manifolds)

(a) (Degeneracy of a squashed area form) For a squashing map $\aleph : \nu B^\vee \to B^\vee$, let $S_0 \subset \nu B^\vee$ be the subset where $\aleph$ is an isometry. Then, $\psi_\aleph^* \omega$ is a non-degenerate form on the subsets

$$\Phi^{-1}(Q^\blacksquare) \times (S_0 \cap \nu Q^\vee) \subset X^\nu, \quad Q \in \mathcal{P}.$$

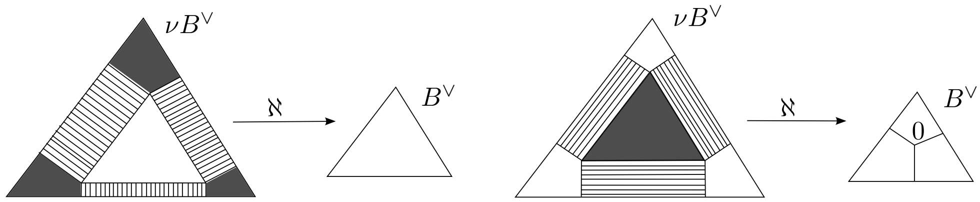

FIGURE 7.4. Examples of squashing maps $\aleph : \nu B^\vee \to B^\vee$. In both examples, solidly shaded regions are mapped to points, ruled regions are mapped to lines by contracting each ruling to a point, and blank regions are mapped isometrically.

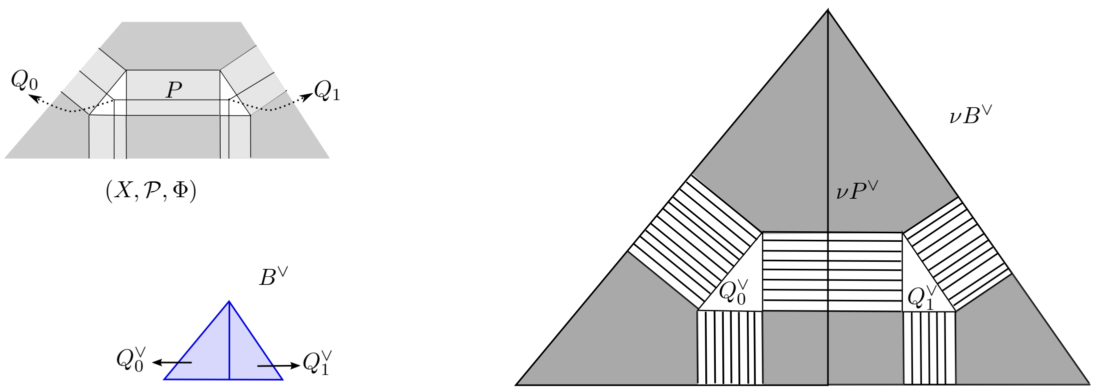

FIGURE 7.5. A squashing map $\aleph : \nu B^\vee \to B^\vee$ where there are two top-dimensional dual polytopes $Q_0^\vee$, $Q_1^\vee$ in $B^\vee$. The maps $\aleph|\nu Q_0^\vee$, $\aleph|\nu Q_1^\vee$ are equal on $\nu P^\vee$.

Suppose on a subset $S_{\mathfrak{k}} \subset \nu Q^{\vee,\circ}$, $Q \in \mathcal{P}$, $\aleph$ is a projection with kernel $\mathfrak{k}^\vee \subset \mathfrak{t}_Q^\vee$, then on the subset

$$\Phi^{-1}(Q^\blacksquare) \times (S_{\mathfrak{k}} \cap \nu Q^{\vee,\circ}) \subset X^\nu$$

the squashed form $\psi_\aleph^* \omega$ is a lift of a non-degenerate form on the quotient $(\Phi^{-1}(Q^\blacksquare) \times (S_{\mathfrak{k}} \cap \nu Q^{\vee,\circ}))/T_{K,\mathbb{C}}$.

(b) (Squashing property on complexes) A map of complexes $\aleph : \nu B^\vee \to B^\vee$ is a squashing map on each polytope $\nu P^\vee$, $P \in \mathcal{P}$ if and only if it is a squashing map on top-dimensional dual polytopes $\nu Q^\vee$ (that is $\dim(Q) = 0$). Indeed, for any face $R^\vee \subset Q^\vee$ the parted polytope corresponding to $\aleph : \nu R^\vee \to R^\vee$ can be read off from the parted polytope of $Q^\vee$ by restricting the embedding $Q^\vee(\mathcal{Q}, \tau) \hookrightarrow \mathfrak{t}^\vee$ to the faces corresponding to $R^\vee$ and projecting the image to $\mathfrak{t}_R^\vee$.

(c) (Surjectivity of squashing maps) A squashing map $\aleph : \nu B^\vee \to B^\vee$ for a neck-stretched manifold is surjective by the following reason: It is enough to prove surjectivity on each top-dimensional polytope $Q^\vee$ in the dual complex. For any facet $R^\vee \subset Q^\vee$, since $\aleph(\nu R^\vee) \subset R^\vee$, the inverse image $\aleph^{-1}(Q^{\vee,\circ}) \subset \mathfrak{t}^\vee$ is on the same side of $\nu R^\vee$ as $Q^\vee$. Since this is true for all facets $R^\vee \subset Q^\vee$, we conclude that the inverse image $\aleph^{-1}(Q^{\vee,\circ})$ of the interior of $Q^\vee$ is contained in $\nu Q^\vee \subset \mathfrak{t}^\vee$, and as a result $\aleph : \nu Q^\vee \to Q^\vee$ is surjective.

REMARK 7.22. For a compact curve C with boundary, and a map $u : (C, \partial C) \to (X^\nu, L)$, the Hofer energy is equal to the area of u :

$$E_{\mathrm{Hof}}(u) = \langle u_*[C], [\omega]\rangle.$$

The next result shows that the squashed area $(\psi_\aleph \circ u)^*\omega$ is pointwise non-negative if the map u is pseudoholomorphic with respect to an almost complex structure that is close to a locally strongly tamed almost complex structure.

LEMMA 7.23. (Monotonicity of Hofer energy) *Suppose $\mathfrak{J}_0 \in \mathcal{J}^{\mathrm{cyl}}(\mathfrak{X})$ is a locally strongly tamed cylindrical almost complex structure. There is a C^0-small neighbourhood $U_{\mathfrak{J}_0} \subset \mathcal{J}^{\mathrm{cyl}}(\mathfrak{X})$ of $\mathfrak{J}_0$ such that the following is satisfied for any $\mathfrak{J}_1 \in U_{\mathfrak{J}_0}$ and any ν :*

(a) (Weakly taming) *For any squashing map $\aleph$,*

$$\forall v \in TX^\nu \quad \psi_\aleph^*\omega(v, J_1^\nu v) \geq 0,$$

where $J_1^\nu \in \mathcal{J}^{\mathrm{cyl}}(X^\nu)$ is obtained by gluing $\mathfrak{J}_1$ at cylindrical ends.

(b) (Monotonic) *Let $u : C \to X^\nu$ be a perturbed J^ν-holomorphic map for a domain dependent almost complex structure $J^\nu : C \to U_{\mathfrak{J}_0}$. For any open subset $\Omega \subset C$,*

$$E_{\mathrm{Hof}}(u, \Omega) \leq E_{\mathrm{Hof}}(u, C).$$

PROOF. The neighborhood $U_{\mathfrak{J}_0}$ is determined by Lemma 7.8. Indeed, for the neighborhood $U_{\mathfrak{J}_0}$ of locally tamed cylindrical almost complex structures, Lemma 7.8 proves the weakly taming property for weak symplectic forms given by maps $\psi_\aleph$ induced by directionally increasing maps $\aleph$ of complexes; and by Lemma 7.17 a squashing map is a directionally increasing map. The (Monotonic) property follows from (Weakly Taming). □

7.3.2. Hofer energy on a broken manifold. Hofer energy on broken manifolds is defined in a similar way as that of neck-stretched manifolds. The only new feature is that the ω-complex and J-complex are different.

We describe the ω-complexes for pieces of the broken manifold. For a polytope $P \in \mathcal{P}$, the ω-complex of the cut space X_P is the subset

$$\mathbb{B}_P^\vee := \pi_{B^\vee}^\omega(X_P^\square, \omega) \subset B^\vee,$$

where $\pi_{B^\vee}^\omega : (X, \omega) \to B^\vee$ is the projection to the dual complex from (7.8) and we recall that the symplectic cut space $X_P^\square = \Phi^{-1}(P)$ is a subset of (X, ω). The complex $\mathbb{B}_P^\vee$ may alternately be defined as

$$\mathbb{B}_P^\vee := \bigcup_{Q \in \mathcal{P}, \dim(Q) = 0, Q \in P} \mathrm{i}_Q^{-1}(P) \subset B^\vee, \tag{7.19}$$

where $\mathrm{i}_Q : Q^\vee \to \mathrm{im}(\Phi) \subset \mathfrak{t}^\vee$ is the embedding from (3.8). See Figure 7.6. The space $\mathbb{B}^\vee_P$ inherits the structure of a complex from $B^\vee$. It is a union of polytopes

$$\mathbb{B}^\vee_P = \cup_{Q \in \mathcal{P}: P \subseteq Q} \mathbb{B}^\vee_{P,Q}, \quad \mathbb{B}^\vee_{P,Q} := \mathbb{B}^\vee_P \cap Q^\vee, \tag{7.20}$$

and for any pair $R \subset Q \subseteq P$, $\mathbb{B}^\vee_{P,Q}$ is identified to a face of $\mathbb{B}^\vee_{P,R}$. Note that $\dim(\mathbb{B}^\vee_P) = \dim(P)$, and $\dim(\mathbb{B}^\vee_{P,Q}) = \dim(P) - \dim(Q)$.

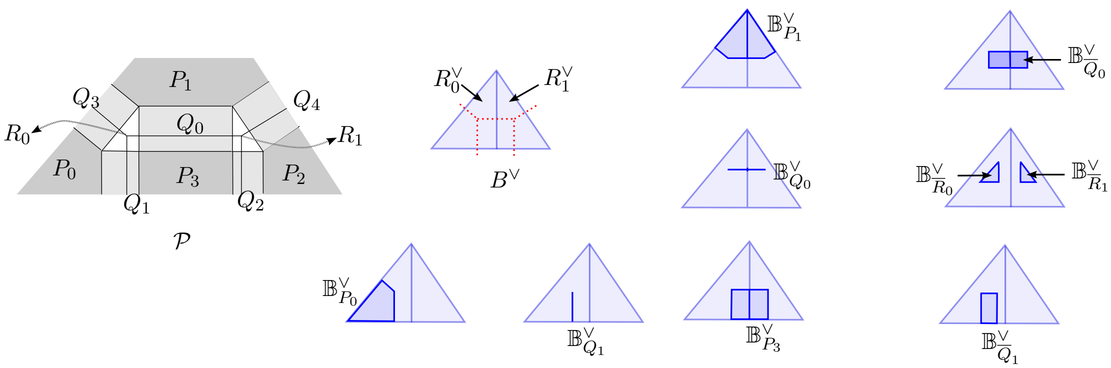

FIGURE 7.6. For a polyhedral decomposition $\mathcal{P}$ of a tropical manifold, $\mathbb{B}^\vee_P$ resp. $\mathbb{B}^\vee_{\overline{P}}$ is the ω-complex of the cut space X_P resp. broken manifold $X_{\overline{P}}$ for $P \in \mathcal{P}$.

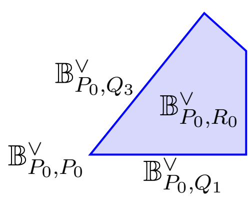

FIGURE 7.7. Polytopes in the complex $\mathbb{B}^\vee_{P_0}$ from Figure 7.6.

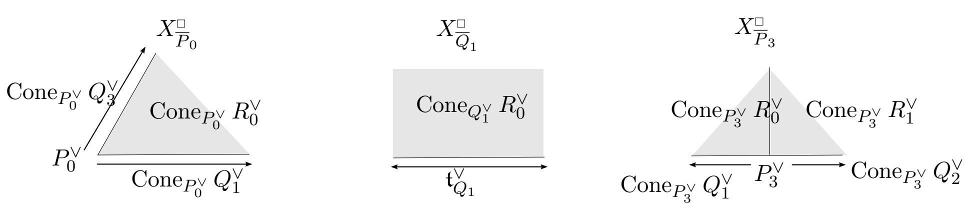

FIGURE 7.8. Some J-complexes for the polyhedral decomposition $\mathcal{P}$ in Figure 7.6

.

We also associate an ω-complex to thickened complexes $\overline{P}$ in order to define Hofer energy for maps in the broken manifold $X_{\overline{P}}$. We define $\mathbb{B}^\vee_{\overline{P}}$ as a thickening of $\mathbb{B}^\vee_P$:

$$\mathbb{B}^\vee_{\overline{P}} := \mathbb{B}^\vee_P \times \delta P^\vee,$$

where $\delta > 0$ is a small constant such that there is an embedding $\mathbb{B}^\vee_{\overline{P}} \to B^\vee$ whose image is a tubular neighborhood of $\mathbb{B}^\vee_P \subset B^\vee$. The space $\mathbb{B}^\vee_{\overline{P}}$ inherits the structure of a complex from $B^\vee$, and thus, consists of polytopes

$$\mathbb{B}^\vee_{\overline{P},Q} := \mathbb{B}^\vee_{P,Q} \times \delta P^\vee \quad \forall Q \subseteq P.$$

The choice of the ω-complexes $\mathbb{B}^\vee_P$, $\mathbb{B}^\vee_{\overline{P}}$ is justified by the fact that the complement of relative divisors in a symplectic broken manifold or cut space can be re-constructed from the complexes as

$$\begin{gathered}\left(\cup_{Q\subseteq P}\Phi^{-1}(Q^\blacksquare) \times \mathbb{B}^\vee_{\overline{P},Q}\right) / \sim = \Phi^{-1}(P^\circ) = (X_{\overline{P}}, \omega_{X_{\overline{P}}}) - \text{relative divisors},\\ \left(\cup_{Q\subseteq P}\Phi^{-1}(Q^\blacksquare)/T_P \times \mathbb{B}^\vee_{P,Q}\right) / \sim = \Phi^{-1}(P^\circ)/T_P = (X_P, \omega_{X_P}) - \text{relative divisors},\end{gathered} \tag{7.21}$$

where $\sim$ is an identification on boundaries, which is a restriction of the equivalence relation in (3.21). The decomposition in (7.21) is a consequence of the decomposition of (X, ω) in (7.7).

Next, we describe the J-complexes. For any $P \in \mathcal{P}$, the J-complex for the broken manifold $X^\square_{\overline{P}}$ is

$$\mathrm{Cone}_{P^\vee} B^\vee := (\bigcup_{Q\subseteq P} \mathrm{Cone}_{P^\vee} Q^\vee)/\sim,$$

where, for any pair $Q_0 \subset Q_1$, the equivalence relation $\sim$ identifies $\mathrm{Cone}_{P^\vee} Q_1^\vee$ to a face of $\mathrm{Cone}_{P^\vee} Q_0^\vee$. The J-complex for the cut space $X^\square_P$ is the corresponding normal cone

$$\mathrm{NCone}_{P^\vee} B^\vee := (\bigcup_{Q\subseteq P} \mathrm{NCone}_{P^\vee} Q^\vee)/\sim,$$

and thus $\mathrm{Cone}_{P^\vee} B^\vee$ is a product of orthogonal spaces

$$\mathrm{Cone}_{P^\vee} B^\vee = \mathrm{NCone}_{P^\vee} B^\vee \times \mathfrak{t}^\vee_P.$$

The almost complex broken manifold and cut spaces can be reconstructed from the J-complexes

$$\begin{aligned} X^\square_{\overline{P}} &= \left(\cup_{Q\subseteq P}\Phi^{-1}(Q^\blacksquare) \times \mathrm{Cone}_{P^\vee} Q^\vee\right) / \sim,\\ X^\square_P &= \left(\cup_{Q\subseteq P}\Phi^{-1}(Q^\blacksquare)/T_P \times \mathrm{NCone}_{P^\vee} Q^\vee\right) / \sim,\end{aligned} \tag{7.22}$$

where, for any facet $Q \subset R$, $\sim$ identifies the boundary component

$$\Phi^{-1}(Q^\blacksquare) \times \mathrm{Cone}_{P^\vee}(Q^\vee) \subset \Phi^{-1}(Q^\blacksquare) \times \mathrm{Cone}_{P^\vee}(Q^\vee)$$

with the boundary component

$$\Phi^{-1}(Q^\blacksquare) \times \mathrm{Cone}_{P^\vee}(Q^\vee) \subset \Phi^{-1}(R^\blacksquare) \times \mathrm{Cone}_{P^\vee}(R^\vee)$$

by the identity map. Here we use the viewpoint that broken manifolds are degenerate limits of neck-stretching, and therefore the decomposition in (7.22) is the limit of the J-decomposition in (7.9). As a consequence of the decomposition in (7.22) there are projection maps

$$\pi_{\mathrm{Cone}_{P^\vee} B^\vee} : X^\square_{\overline{P}} \to \mathrm{Cone}_{P^\vee} B^\vee, \quad \pi_{\mathrm{NCone}_{P^\vee} B^\vee} : X^\square_P \to \mathrm{NCone}_{P^\vee} B^\vee \tag{7.23}$$

for all polytopes $P \in \mathcal{P}$.

DEFINITION 7.24. Let $\mathfrak{X}_{\mathcal{P}}$ be a broken manifold and let $P \in \mathcal{P}$ be a polytope.

(a) (Squashing map for cut spaces) A *squashing map for the cut space* $X_P^{\square}$ is a map of complexes
$$\aleph : \mathrm{NCone}_{P^\vee} B^\vee \to \mathbb{B}_P^\vee$$
for which $\aleph(\mathrm{NCone}_{P^\vee} Q^\vee) \subset \mathbb{B}_{P,Q}^\vee$ for any polytope $Q \subseteq P$, and $\aleph : \mathrm{Cone}_{P^\vee} Q^\vee \to \mathbb{B}_{P,Q}^\vee$ is a squashing map of polytopes.

(b) (Squashing map for a broken manifold) A *squashing map* for the component $X_{\overline{P}} \subset \mathfrak{X}_{\mathcal{P}}$ is a map
$$\aleph = (\aleph^P, \aleph^{P^\vee}) : \mathrm{Cone}_{P^\vee} B^\vee \to \mathbb{B}_{\overline{P}}^\vee$$
where $\aleph^P : \mathrm{NCone}_{P^\vee} B^\vee \to \mathbb{B}_P^\vee$ is a squashing map for the cut space $X_P^{\square}$, and $\aleph^{P^\vee} : \mathfrak{t}_P^\vee \to \delta P^\vee$ is a squashing map of polytopes. Note that $\aleph$ is itself a squashing map of complexes.

See Figure 7.9 for examples.

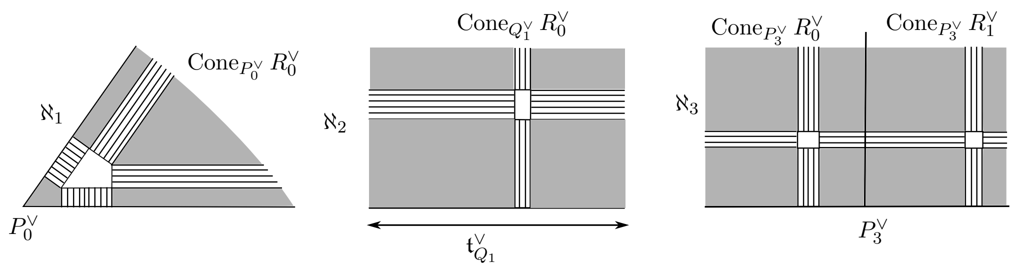

FIGURE 7.9. Squashing maps for a broken manifold $\aleph_1 : \mathrm{Cone}_{P^0} R_0^\vee \to \mathbb{B}_{P_0}^\vee$, $\aleph_2 : \mathrm{Cone}_{Q_1^\vee} R_0^\vee \to \mathbb{B}_{Q_1}^\vee$, $\aleph_3 : Cone_{P_3^\vee} B^\vee \to \mathbb{B}_{P3}^\vee$ with ω-complexes and J-complex from Figures 7.6, 7.8

.

Via the decompositions (7.21), (7.22) of a cut space into a union of fibrations over polytopes, a squashing map $\aleph$ for the cut space $X_P^{\square}$ induces a map of manifolds
$$\psi_\aleph : X_P^{\square} \to (X_P, \omega_{X_P}),$$
that is, piecewise, a smooth submersion, and whose image is the complement of the relative divisors in X_P. Similarly for a component $X_{\overline{P}}^{\square}$ of a broken manifold $\mathfrak{X}$ a squashing map $\aleph$ induces a map of manifolds
$$\psi_\aleph : X_{\overline{P}}^{\square} \to (X_{\overline{P}}, \omega_{X_{\overline{P}}}).$$

DEFINITION 7.25. (Hofer energy for a broken manifold) Let $P \in \mathcal{P}$ be a polytope. The *Hofer energy* of a map $u : C \to X_{\overline{P}}^{\square}$ is
$$E_{\mathrm{Hof}}(u) = \sup_{\aleph} \int_C (\psi_\aleph \circ u)^* \omega_{X_{\overline{P}}},$$

where the supremum is over all squashing maps $\aleph$ for $X^{\square}_{\overline{P}}$ (as in Definition 7.24). The Hofer energy of a holomorphic map $u : C \to X^{\square}_{P}$ is analogously defined.

The proof of the following monotonicity result is the same as the proof in the neck-stretched case (Lemma 7.23).

LEMMA 7.26. (Monotonicity of Hofer energy for broken manifolds) *Suppose* $\mathfrak{J}_0 \in \mathcal{J}^{\rm cyl}(\mathfrak{X})$ *is locally strongly tamed. Then there is a neighborhood* $U_{\mathfrak{J}_0} \subset \mathcal{J}^{\rm cyl}$ *of* $\mathfrak{J}_0$ *such that for any* $\mathfrak{J} \in U_{\mathfrak{J}_0}$, $P \in \mathcal{P}$ *and any map*

$$\psi_{\aleph} : (X^{\square}_{\overline{P}}, \mathfrak{J}) \to (X_{\overline{P}}, \omega_{X_{\overline{P}}})$$

induced by a squashing map $\aleph : \mathrm{Cone}_{P^\vee} B^\vee \to B^\vee_P$, $\mathfrak{J}|X^{\square}_{\overline{P}}$ *satisfies* $\psi_{\aleph}^* \omega_{X_{\overline{P}}}(v, \mathfrak{J}v) \geq 0$ *for all* $v \in TX^{\square}_{\overline{P}}$.

PROPOSITION 7.27. (Limit of maps and Hofer energy) *Let* $\Omega \subset \mathbb{C}$ *be an open domain, and let* $u_\nu : \Omega \to X^\nu$ *be such that there is a polytope* $P \in \mathcal{P}$ *and a sequence of translations* $t_\nu \in \nu P^\vee$ *such that* $d(t_\nu, \nu P_0^\vee) \to \infty$ *for all* $P_0 \supset P$ *and the sequence of translated maps converges uniformly to a limit* $u : \Omega \to X^{\square}_{\overline{P}}$. *Then,*

$$E_{\rm Hof}(u) \leq \lim_{\nu} E_{\rm Hof}(u_\nu).$$

PROOF. We first consider the case that the polytope $P \in \mathcal{P}$ is top-dimensional in $\mathfrak{t}^\vee$ and there is only one point polytope $Q \in \mathcal{P}$ that is a vertex of P. The top-dimensionality of P implies that $\nu P^\vee$ is a point, and $t_\nu = \nu P^\vee$. The inclusion of the P-cylindrical manifold

$$\mathbb{e}^{-t_\nu} : X^{\nu}_{\overline{P}} \to X^{\square}_{\overline{P}}$$

is induced by a map

$$i_\nu : \nu B^\vee \to \mathrm{Cone}_{P^\vee} B^\vee$$

that is a translation in $\mathfrak{t}^\vee$ and sends the point $\nu P^\vee$ to the vertex $P^\vee$ of the cone. The images of i_ν exhaust the target space as $\nu \to \infty$.

The Proposition is proved by constructing a sequence of squashing maps for neck-stretched manifolds that converge to a given squashing map for the broken manifold $X_{\overline{P}}$. In particular, the proof of the Proposition is a consequence of the following Claim:

CLAIM 7.28. Given a squashing map $\aleph : \mathrm{Cone}_{P^\vee} B^\vee \to \mathbb{B}^\vee_P$, there is a sequence of squashing maps $\aleph_\nu : \nu B^\vee \to B^\vee$ such that there is a sequence of subsets $S_\nu \subset \nu B^\vee$ that exhaust $\mathrm{Cone}_{P^\vee}(Q^\vee)$ and for which $\aleph_\nu | S_\nu = \aleph \circ i_\nu$.

PROOF. To prove the Claim, we extend the polyhedral decomposition underlying $\aleph$ to one underlying the sequence $\aleph_\nu$, and define translations for the top-dimensional polytopes in a way that the $\aleph$-polytopes stay fixed, and the non-$\aleph$ polytopes are at a distance proportional to ν. Consequently in the limit $\nu \to \infty$ the non-$\aleph$ polytopes 'disappear into infinity' and do not contribute to the limit map $\aleph$. See Figure 7.10. The details are as follows.

First we describe the polyhedral decomposition underlying the sequence $\aleph_\nu$. Recall that a squashing map $\aleph$ is given by a partition $\mathcal{Q}_P$ of $\mathbb{B}^\vee_P$ and a translation

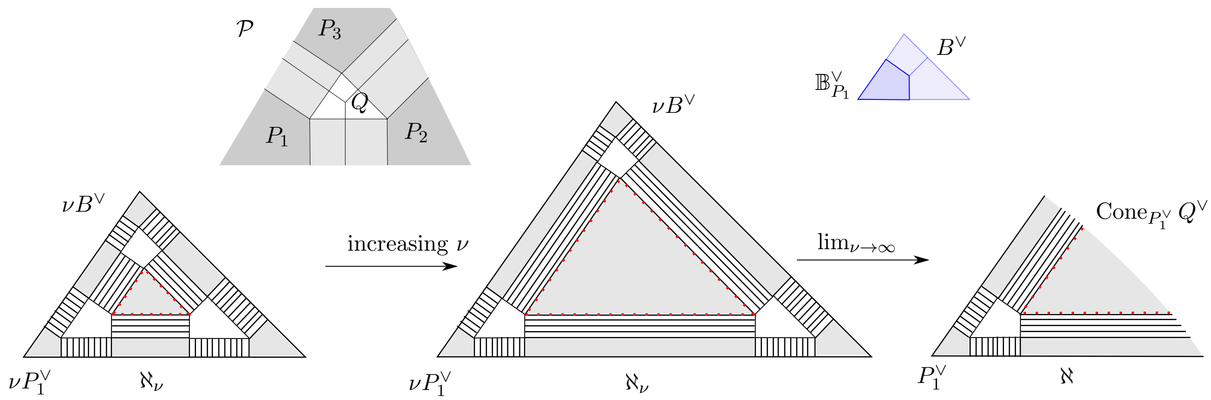

FIGURE 7.10. The squashing maps $\aleph_\nu : \nu B^\vee \to B^\vee$ converge to $\aleph : \mathrm{Cone}_{P_1^\vee} Q^\vee \to \mathbb{B}^\vee_{P_1}$ as $\nu \to \infty$.

$\tau_R \in \mathfrak{t}^\vee$ for each top-dimensional polytope $R \in \mathcal{Q}^0_P$. Let $\mathcal{Q}_{\rm pr}$ be a polyhedral decomposition of $B^\vee$, which we call the *primary decomposition*, whose top-dimensional polytopes are

$$Q_R := i_Q^{-1}(R) \subset B^\vee, \quad R \in \mathcal{P}^0.$$

The squashing maps $\aleph_\nu$ are defined via a polyhedral decomposition $\mathcal{Q}$ of $B^\vee$ that is a refinement of $\mathcal{Q}_{\rm pr}$ (that is, any top-dimensional polytope of $\mathcal{Q}$ is a subset of a polytope in $\mathcal{Q}^0_{\rm pr}$), and $\mathcal{Q}|\mathbb{B}^\vee_P = \mathcal{Q}_P$. See Figure 7.11. We remark that in the special case that $\mathcal{Q}_P$ is a trivial decomposition (that is, $\mathcal{Q}_P = \{\mathbb{B}^\vee_P\}$), $\mathcal{Q}$ can be taken to be equal to $\mathcal{Q}_{\rm pr}$.

The translations corresponding to the squashing maps $\aleph_\nu$ extend the translations underlying $\aleph$ as follows: Suppose $\aleph$ is given by translations μ_R, $R \in \mathcal{Q}^0_P$. For any ν, and a polytope $R \in \mathcal{Q}^0$, the translation τ^ν_R is defined so that

$$\tau^\nu_R := \mu_R, \quad R \in \mathcal{Q}^0_P; \tag{7.24}$$

and for a pair $R_0, R_1 \in \mathcal{Q}^0$ that are contained in distinct top-dimensional polytopes R'_0, R'_1 of the primary decomposition $\mathcal{Q}^0_{\rm pr}$,

$$|\tau^\nu_{R_0} - \tau^\nu_{R_1}| \text{ grows linearly with } \nu. \tag{7.25}$$

The existence of a sequence of translations satisfying (7.24), (7.25) is seen by constructing the dual polytopes underlying the pairs $(\mathcal{Q}, \tau^\nu)$ (see Remark 7.15) as follows: In the special case that $\mathcal{Q}_P$ is a trivial decomposition, and consequently $\mathcal{Q} = \mathcal{Q}_{\rm pr}$, the dual complex $B^\vee_{\mathcal{Q}}$ underlying $\mathcal{Q}$ is taken to be $c\nu Q^\vee$, where $c > 0$ is a constant chosen so that the parted polytope corresponding to $(\mathcal{Q}, \tau)$ is contained in $\nu B^\vee$. In general, the dual complex $B^\vee_{\mathcal{Q},\nu}$ underlying $(\mathcal{Q}, \tau^\nu)$ is constructed as follows: Let

$$\psi : \mathrm{Vert}(B^\vee_{\mathcal{Q}}) \to \mathfrak{t}^\vee, \quad \psi_{\rm pr} : \mathrm{Vert}(B^\vee_{\mathcal{Q}_{\rm pr}}) \to \mathfrak{t}^\vee$$

denote the positions of the point polytopes $B^\vee_{\mathcal{Q}}$, $B^\vee_{\mathcal{Q}_{\rm pr}}$. (Recall that $B^\vee_{\mathcal{Q}}$, $B^\vee_{\mathcal{Q},\rm pr}$ are embedded in $\mathfrak{t}^\vee$.) There is a map of vertices $\kappa : \mathrm{Vert}(B^\vee_{\mathcal{Q}}) \to \mathrm{Vert}(B^\vee_{\mathcal{Q}_{\rm pr}})$ that sends

$\mathfrak{p}^\vee \in B^\vee_{\mathcal{Q}}$ corresponding to $\mathfrak{p} \in \mathcal{Q}^0$ to the vertex $\mathfrak{p}_1^\vee \in B^\vee_{\mathcal{Q}}$, where $\mathfrak{p}_1 \in \mathcal{Q}^0_{\mathrm{pr}}$ is the top-dimensional polytope containing $\mathfrak{p}$. We observe that any edge $e = (v_+, v_-)$ in $B^\vee_{\mathcal{Q}}$ that is not collapsed in $B^\vee_{\mathcal{Q}_{\mathrm{pr}}}$ has the same slope in both the complexes, and therefore, the vector $\psi(v_+) - \psi(v_-) \in \mathfrak{t}^\vee$ is a positive multiple of $\psi_{\mathrm{pr}} \circ \kappa(v_+) - \psi_{\mathrm{pr}} \circ \kappa(v_-)$. Define the vertex position ψ_ν of the dual complex $B^\vee_{\mathcal{Q},\nu}$ underlying $(\mathcal{Q}, \tau^\nu)$ as

$$\mathrm{Vert}(B^\vee_{\mathcal{Q},\nu}) = \mathrm{Vert}(B^\vee_{\mathcal{Q}}) \xrightarrow{\psi_\nu} \mathfrak{t}^\vee, \quad v \mapsto \psi(v) + \nu(\psi_{\mathrm{pr}} \circ \kappa)(v).$$

With this definition, the slopes of edges $e \in B^\vee_{\mathcal{Q},\nu}$, $\dim(e) = 1$ is the same as the slopes in $B^\vee_{\mathcal{Q},\nu}$. The length of any edge collapsed by κ is the same in $B^\vee_{\mathcal{Q},\nu}$ and $B^\vee_{\mathcal{Q}}$; and for any edge not collapsed by κ the length grows linearly in ν. See Figure 7.11. This finishes the proof of Claim 7.28. □

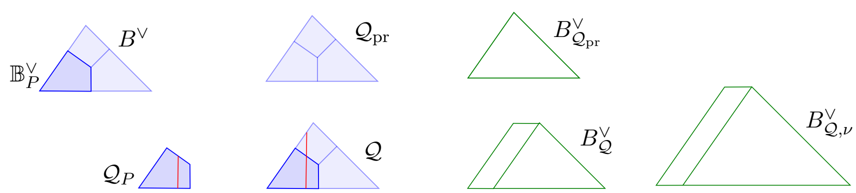

FIGURE 7.11. The polyhedral decomposition $\mathcal{Q}$ of $B^\vee$ extends the decomposition $\mathcal{Q}_{\mathrm{pr}}$, and is equal to $\mathcal{Q}_P$ on $\mathbb{B}^\vee_P \subset B^\vee$. The dual complexes $B^\vee_{\mathcal{Q}}$, $B^\vee_{\mathcal{Q}_{\mathrm{pr}}}$ corresponding to the decompositions $\mathcal{Q}$, $\mathcal{Q}_{\mathrm{pr}}$. The dual complex $B^\vee_{\mathcal{Q},\nu}$ stretches the edges of $\mathcal{B}^\vee_{\mathcal{Q}}$ that occur in $B^\vee_{\mathcal{Q}_{\mathrm{pr}}}$ leaving the others unchanged.

In case P is not a top-dimensional polytope, the proof of the Proposition is the same with the following minor differences:

- The map $i_\nu : \nu B^\vee \to \mathrm{Cone}_{P^\vee} B^\vee$ is a translation in $\mathfrak{t}^\vee$ that is defined by the condition that $\nu P^\vee$ is mapped to $\mathfrak{t}^\vee_P \subset \mathrm{Cone}_{P^\vee} B^\vee$ by translation by t_ν. In fact the translation map $\mathbb{e}^{-t_\nu} : X^\nu_{\overline{P}} \to X^\square_{\overline{P}}$ is induced by the map i_ν of complexes. The condition $d(t_\nu, P_0^\vee) \to \infty$ for all $P_0 \supset P$ ensures that the images $i_\nu(\nu B^\vee)$ exhaust $\mathrm{Cone}_{P^\vee} B^\vee$.
- In the proof of Claim 7.28, the primary decomposition of $B^\vee$ is defined so that $\mathbb{B}^\vee_{\overline{P}}$ is one of the top-dimensional polytopes produced by the decomposition. Unlike the case of $\mathrm{codim}(P) = 0$, there is no natural choice, but any choice of primary decomposition works for the proof. See Figure 7.12 for example.

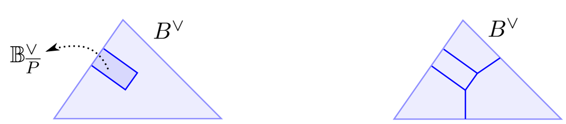

FIGURE 7.12. A primary decomposition of $B^\vee$ (right) that has $\mathbb{B}^\vee_{\overline{P}}$ as a top-dimensional polytope.

Finally in the case that there are more than one point polytopes in $\mathcal{P}$ that are vertices of P, the same proof carries over with heavier notation. The ω-complex $\mathbb{B}_P^\vee$ and J-complex $\mathrm{Cone}_{P^\vee} B^\vee$ each has more than one top-dimensional polytope, and the same construction of $\aleph_\nu$ may be applied to each of the top-dimensional polytopes. This finishes the proof of Proposition 7.27. □

Proposition 7.29. (Hofer energy and quotients) *Let $P \in \mathcal{P}$ be a polytope with* $\mathrm{codim}(P) > 0$*. Let $u : C \to X_{\overline{P}}^\square$ be a perturbed holomorphic map with respect to a domain-dependent almost complex structure $J : C \to U_{\mathfrak{J}_0}$ for which Hofer energy is monotonic (see Lemma 7.26). Let $\pi_P : X_{\overline{P}}^\square \to X_P^\square$ be the quotient under the action of $T_{P,\mathbb{C}}$. Then,*

$$E_{\mathrm{Hof}}(\pi_P \circ u) \leq E_{\mathrm{Hof}}(u).$$

Proof. We recall that a squashing map for $X_{\overline{P}}^\square$ is a product map

$$\aleph = (\aleph^P, \aleph^{P^\vee}) : \mathrm{Cone}_{P^\vee} B^\vee \to \mathbb{B}_{\overline{P}}^\vee$$

where $\aleph^P$: $\mathrm{NCone}_{P^\vee} B^\vee \to \mathbb{B}_P^\vee$ is a squashing map for the cut space $X_P^\square$, and $\aleph^{P^\vee} : \mathfrak{t}_P^\vee \to \delta P^\vee$ is a squashing map of polytopes. We assume $\delta P^\vee$ is embedded in $\mathfrak{t}_P^\vee$ so that the image of $\pi_{\mathfrak{t}_P^\vee}(P)$ is mapped to the origin. For any map $u : C \to X_{\overline{P}}^\square$,

$$(\pi_P \circ u)^*(\psi_{\aleph^P}^* \omega_{X_P}) = u^* \psi_\aleph^* \omega_X$$

if and only if

$$\aleph = (\aleph^P, 0), \quad 0 \in \delta P^\vee$$

on the image of $(\pi_{\mathrm{Cone}_{P^\vee} B^\vee} \circ u)$, where

$$\pi_{\mathrm{Cone}_{P^\vee} B^\vee} : X_{\overline{P}}^\square \to \mathrm{Cone}_{P^\vee} B^\vee$$

is the projection from (7.23).

Claim 7.30. There is a sequence of squashing maps $\aleph_\nu : \mathfrak{t}_P^\vee \to \delta P^\vee$ for which $\mathfrak{t}_P^\vee$ is exhausted by the sequence of sets

$$U_\nu := \mathrm{interior}(\aleph_\nu^{-1}(0)).$$

Proof. Choose a polyhedral decomposition $\mathcal{Q}$ of $\delta P^\vee$ that contains a single point polytope R located at $0 \in \delta P^\vee$, and let $B_{\mathcal{Q}}^\vee$ be a dual complex of the decomposition. Define $\aleph_\nu$ by a sequence of embeddings of the parted polytope $\delta P^\vee(\mathcal{Q})$ into $\mathfrak{t}_P^\vee$ such that the underlying dual complex is $\nu B_{\mathcal{Q}}^\vee$, see Remark 7.15. □

The Lemma follows from the Claim : For any squashing map $\aleph^P$ of $X_P^\square$, define a sequence of squashing maps $\overline{\aleph}_\nu := (\aleph^P, \aleph_\nu)$ of $X_{\overline{P}}^\square$, where $\aleph_\nu$ is the sequence of maps defined by Claim 7.30. We have

$$E_{\mathrm{Hof}}(u) \geq \int_{C_\nu} (\psi_{\overline{\aleph}_\nu} \circ u)^* \omega_X = \int_{C_\nu} (\psi_{\aleph_\nu} \circ \pi_P \circ u)^* \omega_{X_P},$$

where $C_\nu := (\pi_{\mathrm{Cone}_{P^\vee} B^\vee} \circ u)^{-1}(U_\nu)$. Since the sets C_ν exhaust the domain C, we conclude that for any squashing map $\aleph_P$ of $X_P^\square$,

$$\int_C (\psi_{\aleph_\nu} \circ \pi_P \circ u)^* \omega_{X_P} \leq E_{\mathrm{Hof}}(u),$$

and consequently $E_{\mathrm{Hof}}(\pi_P \circ u) \leq E_{\mathrm{Hof}}(u)$. □

7.4. Basic area forms

In this section we define a sequence of two-forms called 'basic area forms' on neck-stretched manifolds that resemble symplectic forms in the complement of necks. On the neck regions, these area forms are basic. Such a family of area forms is useful in the proof of Gromov convergence, as it gives us a taming symplectic form on any compact subset of cut spaces.

The basic area forms are given by a family of squashing maps, called 'basic squashing maps', which we first describe informally. For any top-dimensional polytope $Q^\vee \subset B^\vee$ (corresponding to a point polytope $Q \in \mathcal{P}$), the map $\aleph_\nu^{\mathrm{bas}} : \nu Q^\vee \to Q^\vee$ is given by a polyhedral decomposition of $Q^\vee$ which looks like the decomposition $\mathcal{P}$ in a neighborhood of Q. The top-dimensional polytopes in this decomposition are $\mathbb{B}_{P,Q}^\vee \subset Q^\vee$ corresponding to every $P \in \mathcal{P}^0$ (see (7.20) and Figures 7.6, 7.7). To define the squashing map $\aleph_\nu^{\mathrm{bas}}$, we also need an embedding $\mathbb{B}_{P,Q}^\vee \hookrightarrow \nu Q^\vee$ for every $P \in \mathcal{P}$, which is defined so that the vertex $P^\vee \in \mathbb{B}_{P,Q}^\vee$ is mapped to $\nu P^\vee \in \nu Q^\vee$. The squashing map in the right side of Figure 7.4 is basic. Figure 7.13 is another example of a basic squashing map.

Definition 7.31. (Basic area form for neck-stretched manifolds) For any $\nu \geq 0$, the *basic squashing map* is an undilated squashing map

$$\aleph_\nu^{\mathrm{bas}} : \nu B^\vee \to B^\vee$$

- given by a polyhedral decomposition $\mathcal{Q}$ of $B^\vee$, where for every top-dimensional polytope $Q^\vee \subset B^\vee$ (corresponding to a point polytope $Q \in \mathcal{P}$) the top-dimensional polytopes in the decomposition $\mathcal{Q}|Q^\vee$ are

$$\{\mathbb{B}_{P,Q}^\vee \simeq i_Q^{-1}(P) \cap Q^\vee : P \in \mathcal{P}^0, Q \in P\},$$

  where $\mathbb{B}_{P,Q}^\vee$ is defined in (7.20).
- To define $\aleph_\nu^{\mathrm{bas}}$ the translation of $\mathbb{B}_{P,Q}^\vee \subset Q^\vee$ is defined so that

$$(\aleph_\nu^{\mathrm{bas}})^{-1}(P^\vee) = \{\nu P^\vee\} \quad \forall P \in \mathcal{P}^0. \tag{7.26}$$

The map $\aleph_\nu^{\mathrm{bas}}$ defines a *basic area form* as

$$\omega_\nu^{\mathrm{bas}} := \phi^*_{\aleph_\nu^{\mathrm{bas}}} \omega \in \Omega^2(X^\nu).$$

See Figure 7.13 for an example of a squashing map underlying a basic area form.

Definition 7.32. (Dilated basic area form) Let $\aleph_\nu^{\mathrm{bas}} : \nu B^\vee \to B^\vee$ be the family of basic squashing maps. For all $1 \leq t \ll \nu$ *dilated basic squashing maps* $\aleph_{\nu,t}^{\mathrm{bas}}$ are defined as

$$\aleph_{\nu,t}^{\mathrm{bas}} := \aleph_{\nu/t}^{\mathrm{bas}} \circ \delta_t : \nu B^\vee \to B^\vee,$$

where the map $\delta_t : \nu B^\vee \to \frac{\nu}{t} B^\vee$ dilates by a factor of t. The maps define *dilated basic area forms* as

$$\omega_{\nu,t}^{\mathrm{bas}} := \phi^*_{\aleph_{\nu,t}^{\mathrm{bas}}} \omega \in \Omega^2(X^\nu).$$

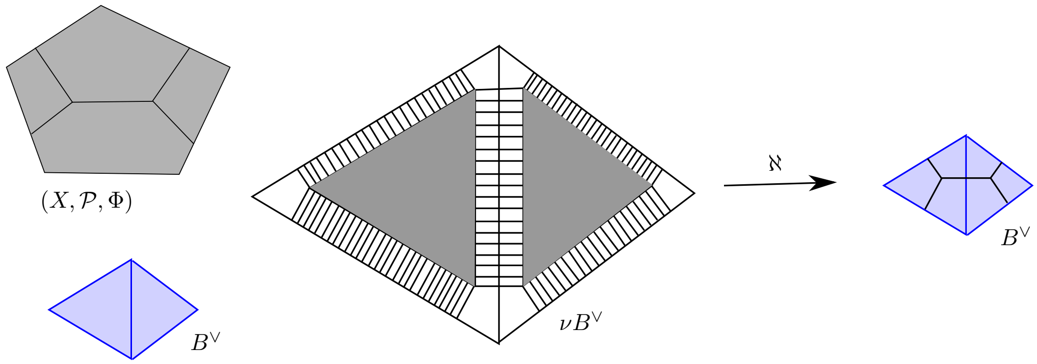

FIGURE 7.13. Example of a squashing map underlying a basic area form.

Basic area forms on neck-stretched manifolds converge to basic area forms on cut spaces which are defined as follows.

DEFINITION 7.33. (Basic area on cut spaces, broken manifolds) Let $P \in \mathcal{P}$. The *basic squashing map*

$$\aleph_P^{\mathrm{bas}} : \mathrm{NCone}_{P^\vee} B^\vee \to \mathbb{B}_P^\vee$$

on the cut space $X_P^\square$ is an unpartitioned squashing map (as in Definition 7.19) defined by the condition

$$(\aleph_P^{\mathrm{bas}})^{-1}(\mathbb{B}_{P,P}^\omega) = P^\vee. \tag{7.27}$$

Here we recall that $\mathbb{B}_{P,P}^\omega$ is a vertex of the polytope $\mathbb{B}_{P,Q}^\omega$ (see (7.20)) occuring in the complex $\mathbb{B}_P^\vee$. The *basic area form* on the cut space $X_P^\square$ is defined as

$$\omega_P^{\mathrm{bas}} := \psi^*_{\aleph_P^{\mathrm{bas}}} \omega_{X_P} \in \Omega^2(X_P^\square).$$

Dilated basic area forms $\omega_{P,t}^{\mathrm{bas}} \in \Omega^2(X_P^\square)$ are defined analogously by dilating the squashing map $\aleph_P^{\mathrm{bas}}$ by a factor of t.

PROPOSITION 7.34. (Basic area forms on neck stretched manifolds) *The basic area forms $\omega_\nu^{\mathrm{bas}}$ and $\omega_{\nu,t}^{\mathrm{bas}}$ satisfy the following properties:*

(a) (Boundedness) *For any ν, t, and a map $u : C \to X^\nu$,*

$$\omega_{\nu,t}^{\mathrm{bas}}(u) \le E_{\mathrm{Hof}}(u).$$

(b) (Cohomology) *The form $\omega_{\nu,t}^{\mathrm{bas}}$ is cohomologous to $\omega \in \Omega^2(X)$ and is weakly taming for $J \in U_{\mathfrak{J}_0} \subset \mathcal{J}^{\mathrm{cyl}}(\mathfrak{X})$ from Lemma 7.23.*

(c) (Convergence to horizontal forms) *Suppose for a polytope $P \in \mathcal{P}$, $t_\nu \in \nu P^\vee$ is a sequence of translations $t_\nu \in \nu P^\vee$ satisfying $d(t_\nu, \nu P_0^\vee) \to \infty$ for any $P_0 \supsetneq P$. Then*

$$(\mathbb{e}^{t_\nu})^* \omega_\nu^{\mathrm{bas}} \quad \textit{converges to} \quad \pi_P^* \omega_P^{\mathrm{bas}}, \tag{7.28}$$

where the embedding $\mathbb{e}^{t_\nu} : X_P^\square \to X_P^\nu$ is defined in (3.42).

(d) (Non-degeneracy on compact sets) *For any $P \in \mathcal{P}$, $t \geq 1$ there is an open set $U_{P,t} \subset X_P^{\square}$ on which $\omega_{P,t}^{\text{bas}}$ is non-degenerate, and the family $\{U_{P,t}\}_t$ exhausts the cut space $X_P^{\square}$ as $t \to \infty$.*

Proof. Parts (a) and (b) are true for any squashed area form. To prove part (c) consider a polytope $P \in \mathcal{P}$. Since we may prove the convergence (7.28) on one cylindrical region at a time, for the sake of simplifying notation, we assume that there is only one point polytope $Q \in \mathcal{P}$. We recall that $\mathbb{e}^{t_\nu}$ is defined as the inverse of $\mathbb{e}^{-t_\nu}$, and the translation map $\mathbb{e}^{-t_\nu} : X_{\overline{P}}^{\nu} \to X_{\overline{P}}^{\square}$ is defined is induced by the map

$$i_\nu : \nu B^\vee \to \text{Cone}_{P^\vee} B^\vee$$

which is a translation in $\mathfrak{t}^\vee$ defined by the condition that $\nu P^\vee$ is mapped to $\mathfrak{t}_P^\vee \subset \text{Cone}_{P^\vee} B^\vee$ by translation by t_ν. The condition $d(t_\nu, P_0^\vee) \to \infty$ for all $P_0 \supset P$ ensures that the images $i_\nu(\nu B^\vee)$ exhaust $\text{Cone}_{P^\vee} B^\vee$. The convergence (7.28) follows from the following Claim :

Claim 7.35. The sequence of maps $\aleph_\nu^{\text{bas}} \circ i_\nu^{-1}$ converges to a lift of $\pi_P^* \aleph_P^{\text{bas}}$ where $\pi_P : \text{Cone}_{P^\vee} B^\vee \to \text{NCone}_{P^\vee} B^\vee$ is the projection map.

Proof of Claim. For any ν, let $S_{P,\nu} \subset \nu B^\vee$ denote the subset on which $\aleph_\nu^{\text{bas}}$ is invariant under $\mathfrak{t}_P^\vee$ translation of the domain. The condition $d(t_\nu, P_0^\vee) \to \infty$ for all $P_0 \supset P$ together with the definition of basic area forms ensures that $i_\nu(S_{P,\nu})$ exhausts $\text{Cone}_{P^\vee} B^\vee$. The map i_ν, $\aleph_\nu^{\text{bas}}$ descend to

$$\bar{i}_\nu : S_{P,\nu}/\mathfrak{t}_P^\vee \to \text{NCone}_{P^\vee} B^\vee, \quad \overline{\aleph}_\nu^{\text{bas}} : S_{P,\nu}/\mathfrak{t}_P^\vee \to \mathbb{B}_P^\vee.$$

The defining condition of basic squashing maps (7.26), (7.27) ensures that the maps $\aleph_\nu^{\text{bas}} \circ \bar{i}_\nu^{-1}$ converges to $\aleph_P^{\text{bas}}$, implying the statement in the Claim. □

Part (d) follows from the fact that the sequence of sets

$$S_{P,t} \subset \text{NCone}_{P^\vee} B^\vee$$

on which $\aleph_{p,t}^{\text{bas}}$ is a dilation (and not a projection) exhausts $\text{NCone}_{P^\vee} B^\vee$ as $t \to \infty$. □

Definition 7.36. (Basic area of a region of a map) Let Ω be a complex curve possibly with boundary, and let $u : \Omega \to X^\nu$ be a perturbed holomorphic map. The *basic area* of the map is

$$\text{Area}^{\text{bas}}(u, \Omega) := \int_\Omega u^* \overline{\omega}_\nu^{\text{bas}},$$

where $\omega_\nu \in \Omega^2(X^\nu)$ is the basic area form on X^ν. If a map u has closed domain Ω, or it maps the boundary of the domain $\partial\Omega$ to the Lagrangian L, then $\text{Area}(u) = \text{Area}^{\text{bas}}(u)$. For a perturbed holomorphic map $u : \Omega \to X^\nu$ as in Definition 7.36, the *dilated basic area* is

$$\text{Area}_t^{\text{bas}}(u, \Omega) := \int_\Omega u^* \overline{\omega}_{\nu,t}^{\text{bas}}.$$

7.5. Removal of singularities

In this section, we prove a removal of singularities result for punctured holomorphic maps in a piece of a broken manifold. The standard approach to prove this result by embedding the broken almost complex manifold into a compact symplectic manifold fails, because the almost complex manifold $X_P^\square$ can not be embedded into the symplectic cut space X_P^ω via increasing maps.[1] Consequently a Hofer energy bound in $X_P^\square$ does not translate to a bound on symplectic area in X_P^ω. Therefore we take a different approach where we prove that any punctured holomorphic map with finite Hofer energy lies in some Q-cylindrical region and the image of the projection to $X_Q^\square$ lies in a compact set in the complement of relative divisors. Therefore the removal of singularities result applies on the projected map, and consequently on the original map.

For the removal of singularities result we consider maps on a punctured disk $B_1 \backslash \{0\}$ which is holomorphically identified with

$$\mathrm{Cyl} := \mathbb{R}_{\geq 0} \times S^1.$$

For any $l \geq 0$, we refer to a truncated semi-infinite cylinder by

$$\mathrm{Cyl}(l) := [l, \infty) \times S^1.$$

Proposition 7.1. (Removal of singularities) *Suppose $u : [0,\infty) \times S^1 \to X_{\overline{P}}^\square$ is a perturbed J-holomorphic curve with respect to the domain-dependent almost complex structure $J : B_1 \to U_{\mathfrak{J}_0}$ (holomorphically identifying $\mathrm{Cyl} \simeq B_1 \backslash \{0\}$) with*

$$E_{\mathrm{Hof}}(u) < \infty, \quad \|\mathrm{d}u\|_{L^\infty(\mathrm{Cyl})} < \infty. \tag{7.29}$$

Then, u extends to a holomorphic map $u : B_1 \to X_{\overline{P}}$ possibly, in the orbifold case, after passing to a finite cover.

Monotonicity for pseudoholomorphic maps is the main technical tool in the proof of Proposition 7.1. We state the monotonicity result (see for example [**90**, Proposition 3.12]).

Proposition 7.37. (Monotonicity) *Let (X, ω) be a compact symplectic manifold, and let $U_{\mathcal{J}}$ be a C^0-neighborhood on the space of tamed almost complex structures. There exist constants $c, r_0 > 0$ such that for any $x \in X$, $0 < r \leq r_0$, a Riemann surface C with boundary ∂C and a pseudoholomorphic map $u : C \to X$ with respect to a domain-dependent almost complex structure $J : C \to U_{\mathcal{J}}$ whose image contains x and $u(\partial C) \subset \partial B(x, r)$,*

$$\int_C u^* \omega \geq c r^2.$$

We first prove the removal of singularities result (Proposition 7.1) in the case of a single cut to serve as a warm-up for the more complicated proof in the case of multiple cuts.

[1] For any vertex $Q \in \mathcal{P}$ of P, an increasing map may not exist in the Q-cylindrical corner of $X_P^\square$ because the fixed $\mathfrak{t}$-inner product from (3.6) is not equal to the natural $\mathfrak{t}$-inner product at the Q-corner. The natural inner product is the one for which the edges $P_1 \subset P$ emanating from Q form an orthogonal basis.

Proof of Proposition 7.1 in the case of a single cut. We set up some notation first. We consider a single cut with polyhedral decomposition

$$\mathcal{P} = \{P_- := (-\infty, c], P_0 := \{c\}, P_+ := [c, \infty)\},$$

and dual complex $B^\vee \cong [\frac{-\delta}{2}, \frac{\delta}{2}]$.

We consider a map $u : \mathrm{Cyl} \to X_{P_+}$ satisfying the hypothesis of Proposition 7.1, the case of a map in X_{P_-} being similar. We recall there is a projection to the J-complex

$$\pi_{B_J} : X^{\square}_{P_+} \to \mathrm{Cone}_{P_+^\vee} \simeq (-\infty, 0]$$

We choose any unpartitioned squashing map (as in Definition 7.19)

$$\mathrm{Cone}_{P_+^\vee} \simeq (-\infty, 0] \xrightarrow{\aleph} [\tfrac{-\delta}{2}, \tfrac{\delta}{2}] \simeq B^\vee.$$

Such a map is a translation on an interval $[\tau, \tau + \delta] \subset \mathbb{R}_-$ and a locally constant map on the complement. Denote

$$X_{B^{\vee,\circ}} := \pi_{B_J}^{-1}(\overline{\aleph^{-1}(B^{\vee,\circ})}) \simeq \pi_{B_J}^{-1}([\tau, \tau + \delta]).$$

Definition 7.38. (Crossing) A connected component $C \subset \mathrm{Cyl}$ of $u^{-1}(X_{B^{\vee,\circ}})$ is called a *crossing* if $u(C)$ intersects both boundary components of $X_{B^{\vee,\circ}}$

We claim that there is a lower bound on the squashed area of crossings. That is, there is a constant $c > 0$ such that for any crossing C

$$\int_C u^*(\psi_\aleph^* \omega) \geq c. \tag{7.30}$$

Choose a small constant ϵ and a closed interval $U \subset B^{\vee,\circ}$ such that

$$\tilde{U} := \overline{B_\epsilon(U)} \subset B^{\vee,\circ} = (\tfrac{-\delta}{2}, \tfrac{\delta}{2}],$$

and define

$$X_{\tilde{U}} := \pi_{B_J}^{-1}(\aleph^{-1}(\tilde{U})), \quad X_U := \pi_{B_J}^{-1}(\aleph^{-1}(U)) \subset X^{\square}_{P_+}.$$

The squashed form $\psi_\aleph^* \omega$ is a symplectic form on $X_{\tilde{U}}$ and $B_\epsilon(X_U) \subset X_{\tilde{U}}$. For a crossing C, $u(C)$ intersects X_U. Therefore by the monotonicity theorem applied to the manifold $(X_{\tilde{U}}, \psi_\aleph^* \omega)$, the lower bound (7.30) on the squashed area follows.

Next we claim that one of the following two possibilities occur :

(a) Either the image $u(\mathrm{Cyl})$ is contained in a compact subset of $X^{\square}_{P_+}$,
(b) or after truncating the domain cylinder by a finite amount, the image of u is contained in $\pi_{B_J}^{-1}(-\infty, \tau]$, which is equal to $(-\infty, \tau] \times Z_{P_0}$, where $Z_{P_0} \to X_{P_0}$ is a principal S^1-bundle.

Indeed, after passing to a truncation of the domain cylinder, we may assume that the squashed area of the map is small enough to ensure that there are no compact crossings. If there is a non-compact crossing $C \subset \mathrm{Cyl}$, the image of u is contained in a compact set. Indeed, C intersects the image $u(\{\ell\} \times S^1)$ for all $\ell \geq \ell_0$ for some ℓ_0, and therefore the image $u(\mathrm{Cyl}(\ell_0))$ is contained within a radius $2\pi \|du\|_{L^\infty}$ of $X_{\tilde{U}}$, which is a compact set.

In case the image of u is contained in a compact subset of X_{P_+}, the result follows from the removal of singularities result for compact symplectic manifolds. Indeed by Proposition 7.34 (d), for any compact subset K of X_{P_+} there is a dilated basic form ω_t^{bas} that is symplectic on K, and $\int_{\mathrm{Cyl}} u^*\omega_t^{\mathrm{bas}} < E_{\mathrm{Hof}}(u)$.

Next we consider the case that the image of u is in the semi-infinite cylinder $Z_{P_0} \times (-\infty, \tau]$. The $(\omega_{X_{P_0}} - \frac{\epsilon}{2} d\alpha_{P_0})$-area of the projection $u_{P_0} := \pi_{P_0} \circ u$ is bounded by $E_{\mathrm{Hof}}(u)$ because if we define the squashing map $\aleph_0$ so that

$$\aleph_0 = \tfrac{-\epsilon}{2} \quad \text{on} \quad (-\infty, \tau],$$

then

$$\int_{\mathrm{Cyl}} u_{P_0}^*(\omega_{X_{P_0}} - \tfrac{\epsilon}{2} d\alpha_{P_0}) = \int_{\mathrm{Cyl}} u^*\omega_{\aleph_0} \leq E_{\mathrm{Hof}}(u).$$

By the removal of singularities result for compact symplectic manifolds, we conclude that the projected map u_{P_0} extends holomorphically to

$$u_{P_0} : B_1 \to X_{P_0}.$$

For the rest of the proof we view u as mapping to $X^{\square}_{\overline{P}_0} \simeq \mathbb{Z}_{P_0} \times \mathbb{R}$, since the P_0-cylindrical end is embedded in $X^{\square}_{\overline{P}_0}$. Consider a holomorphic trivialization of the pullback bundle $u_{P_0}^* X^{\square}_{\overline{P}_0} \to B_1$. The projection of u to the fiber, denoted by

$$u_v : B_1 \backslash \{0\} \to S^1 \times \mathbb{R}, \tag{7.31}$$

is holomorphic, and the $\mathbb{R}$-coordinate has an upper bound. Therefore u_v extends over 0 to a holomorphic map in $\mathbb{P}^1$. Consequently, u extends over 0 in the compactified space $X_{\overline{P}}$.

Next we consider the case of a map $u : \mathrm{Cyl} \to X^{\square}_{\overline{P}_0}$. As in the previous paragraph, the singularity at ∞ can be removed for the projected map $\pi_{P_0} \circ u : \mathrm{Cyl} \to X_{P_0}$. To prove that the singularity can be removed for the vertical component u_v (as in (7.31)), it is enough to show that the $\mathbb{R}$-component has either an upper or a lower bound so that essential singularities can be ruled out. This bound is a consequence of a lower bound on the squashed area $\psi^*\aleph$ for crossings. The details are exactly as in the case of $X_{P_\pm}$ and are therefore omitted. □

The proof of removal of singularities in the case of a single cut proceeded by showing that the map u either lies in a compact set or in the P_0-cylinder, and we proved this fact by ruling out ‘crossings’ using an area bound. To prove the result in the case of multiple cuts, we analyze a variety of crossings, each corresponding to a P-cylindrical region, for some $P \in \mathcal{P}$. In the proof in the single cut case, if the image of the map were in a P_0-cylindrical subset, we needed to show that the $\mathfrak{t}^\vee_{P_0}$-coordinate (which is an $\mathbb{R}$-coordinate) either has an upper or lower bound, or in other words, it lies in a cone. Analogously in case of a multiple cut, if the map lies in a P-cylindrical subset, we will show that the $\mathfrak{t}^\vee_P$-coordinate on the image of the map lies in a proper cone in the sense defined below:

Definition 7.39. A cone C in an affine space V is a *proper cone* if it has a non-empty interior and it is the intersection of n half-spaces where $n := \dim(V)$.

Most of the technical work in the proof of removal of singularities for multiple cuts is carried out in the following Lemma.

LEMMA 7.40. *For a holomorphic map*

$$u : \mathrm{Cyl} \to X_P^{\square}, \quad \textit{resp.} \quad u : \mathrm{Cyl} \to X_{\overline{P}}^{\square}$$

with finite Hofer energy and $\|du\|_{L^\infty} < \infty$*, there is a constant* ℓ_0 *and a polytope* $Q \in \mathcal{P}$, $Q \subseteq P$ *such that*

(a) *the image* $u(\mathrm{Cyl}(\ell_0))$ *is contained in the* Q*-cylindrical end of* $X_P^{\square}$ *resp.* $X_{\overline{P}}^{\square}$*,*
(b) *the image of* $\pi_{X_Q} \circ u : \mathrm{Cyl}(\ell_0) \to X_Q^{\square}$ *is contained in a compact subset of* $X_Q^{\square}$*,*
(c) $\pi_{Q^\vee} \circ u$ *is contained in a proper cone of* $\mathfrak{t}_Q^\vee/\mathfrak{t}_P^\vee$ *resp.* $\mathfrak{t}_Q^\vee$*.*

REMARK 7.41. If we apply Lemma 7.40 to a cut space $X_P^{\square}$, and it turns out that $Q = P$, then the image of u is contained in a compact set since the map π_{X_P} on $X_P^{\square}$ is the identity map. Further, in this case, the conclusion (c) is vacuous because $\mathfrak{t}_Q^\vee/\mathfrak{t}_P^\vee$ is a point.

PROOF OF LEMMA 7.40. We first prove the result for the case that P is a point, and therefore $P^\vee \subset B^\vee$ is top-dimensional. We may assume that the target space is $X_{\overline{P}}^{\square}$, since the other case is that of $X_P^{\square}$ which is a compact symplectic manifold. The Lemma then requires us to prove that the $\mathfrak{t}_P^\vee$-coordinate on the image of u is contained in a proper cone. We denote the complex and symplectic polytopes as B_J and B_ω, and therefore in the current case

$$B_J = \mathfrak{t}_P^\vee, \quad B_\omega = P^\vee.$$

We recall that there is a projection map

$$\pi_{B,J} : X_{\overline{P}}^{\square} \to B_J.$$

Fix any squashing map

$$\aleph : B_J \to B_\omega.$$

The set of 'special faces' of B_ω plays a role in the proof. In the current case of $P = \textit{point}$, all the faces of $B_\omega \simeq P^\vee$ are special faces. We say that for any face R of B_ω the interior R° of R is the complement of the special faces of R.

STEP 1: *We choose a collection neighborhoods in* B_ω*, each corresponding to a special face of* B_ω*, which cover* B_ω*, and such that for each of the neighborhoods the* $\aleph$*-inverse image is contained in a proper cone.*

We choose small constants $\epsilon_{\mathfrak{q}} > 0$ for all $\mathfrak{q}$; and closed subsets

$$\mathfrak{q}^{\blacksquare} \Subset \tilde{\mathfrak{q}}^{\blacksquare} \Subset \mathfrak{q} \tag{7.32}$$

such that $\mathfrak{q}^{\blacksquare}$ resp. $\tilde{\mathfrak{q}}^{\blacksquare}$ is contained in the interior of $\tilde{\mathfrak{q}}^{\blacksquare}$ resp. $\mathfrak{q}$, and if $\dim(\mathfrak{q}) = 0$ then $\mathfrak{q}^{\blacksquare} := \tilde{\mathfrak{q}}^{\blacksquare} = \mathfrak{q}$; and closed tubular neighborhoods

$$U_{\mathfrak{q}} = \mathfrak{q}^{\blacksquare} \times I_{\mathfrak{q}}, \quad \tilde{U}_{\mathfrak{q}} = \tilde{\mathfrak{q}}^{\blacksquare} \times I_{\mathfrak{q}} \subset B_\omega, \tag{7.33}$$

where $I_{\mathfrak{q}}$ is a closed neighborhood of the origin in the normal cone $\mathrm{NCone}_{\mathfrak{q}} B_\omega$ such that

(a) $B_\omega = \cup_{\mathfrak{q}:\text{special face of } B_\omega} U_{\mathfrak{q}}$,
(b) $U_{\mathfrak{q}_1} \cap U_{\mathfrak{q}_2} \neq \emptyset$ iff $\mathfrak{q}_1 \subseteq \mathfrak{q}_2$ or $\mathfrak{q}_2 \subseteq \mathfrak{q}_1$,

(c) the vertical boundary of $U_{\mathfrak{q}}$ resp. $\tilde{U}_{\mathfrak{q}}$, that is,

$$\mathfrak{q}^{\blacksquare} \times (\partial I_{\mathfrak{q}} \backslash \partial B_{\omega}), \quad \text{resp.} \quad \tilde{\mathfrak{q}}^{\blacksquare} \times (\partial I_{\mathfrak{q}} \backslash \partial B_{\omega}),$$

is contained in $\cup_{\mathfrak{q} \subset \mathfrak{r}} U_{\mathfrak{r}}$.

Note that for any special face $\mathfrak{q}$, the inverse image $\aleph^{-1}(\tilde{U}_{\mathfrak{q}})$ is contained in a proper cone of $\mathfrak{t}_P^{\vee}$. See Figure 7.14. For any special face $\mathfrak{q}$ of B_{ω}, we define

$$X_{\mathfrak{q}^{\blacksquare}} := \pi_{B_J}^{-1} \aleph^{-1}(\mathfrak{q}^{\blacksquare}), \quad X_{U_{\mathfrak{q}}} := \pi_{B_J}^{-1} \aleph^{-1}(U_{\mathfrak{q}}), \quad \text{etc.}$$

Viewing $X_{\tilde{U}_{\mathfrak{q}}}$ as a fibration over $X_{\tilde{\mathfrak{q}}^{\blacksquare}}$, we define the vertical boundary as

$$\partial_{\text{ver}} X_{\tilde{U}_{\mathfrak{q}}} = \pi_{B_J}^{-1} \aleph^{-1}(\mathfrak{q}^{\blacksquare} \times (\partial I_{\mathfrak{q}} \backslash \partial B_{\omega})) \tag{7.34}$$

and the horizontal boundary as

$$\partial_{\text{hor}} X_{\tilde{U}_{\mathfrak{q}}} = \pi_{B_J}^{-1} \aleph^{-1}(\partial \tilde{\mathfrak{q}}^{\blacksquare} \times I_{\mathfrak{q}}). \tag{7.35}$$

Example 7.42. Consider the complex in Figure 7.14.

(a) For the top-dimensional polytope $\mathfrak{q} = Q^{\vee}$, $I_{\mathfrak{q}} = point$. So, $\mathfrak{q}^{\blacksquare} = U_{\mathfrak{q}}$ and $\tilde{\mathfrak{q}}^{\blacksquare} = \tilde{U}_{\mathfrak{q}}$ are contained in the interior of $Q^{\vee}$ and have the same dimension as $Q^{\vee}$. The set $X_{\tilde{U}_{\mathfrak{q}}}$ only has a horizontal boundary and no vertical boundary.

(b) For $\mathfrak{q} = \mathfrak{q}_1$ (and symmetrically $\mathfrak{q}_2$), $I_{\mathfrak{q}}$ is an interval, and $X_{U_{\mathfrak{q}}}$ is a manifold with corners. The boundary component intersecting $X_{U_{Q^{\vee}}}$ is the vertical boundary, and the boundary component which only intersects $X_{U_{\mathfrak{p}}}$ is the horizontal boundary.

(c) For the point polytope $\mathfrak{q} = \mathfrak{p}$, $I_{\mathfrak{p}}$ is a top-dimensional cone. The manifold $X_{U_{\mathfrak{q}}}$ only has a vertical boundary, and no horizontal boundary.

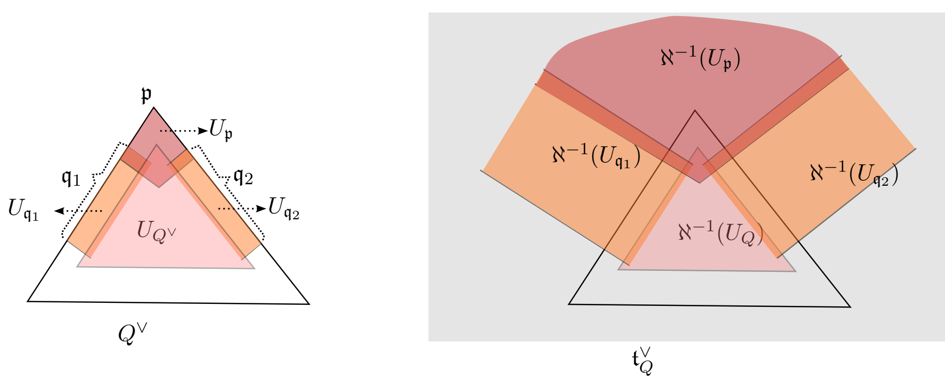

Figure 7.14. Left : Neighborhoods of interiors of faces of $Q^{\vee}$ that cover $Q^{\vee}$. Right : The inverse images $\aleph^{-1}(U_Q)$, $\aleph^{-1}(U_{\mathfrak{q}_1})$, $\aleph^{-1}(U_{\mathfrak{q}_2})$, $\aleph^{-1}(U_{\mathfrak{p}_1})$ are cones in $\mathfrak{t}_Q^{\vee}$.

Definition 7.43. ($\mathfrak{q}$-crossing) For a special face $\mathfrak{q}$ of B_{ω} with $\dim(\mathfrak{q}) \geq 1$, we say that a connected component $C \subset \text{Cyl}$ of $u^{-1}(\overline{X_{\tilde{U}_{\mathfrak{q}}}})$ is a $\mathfrak{q}$-*crossing* if $u(C)$ intersects $X_{U_{\mathfrak{q}}}$, and $u(\partial C)$ lies on the horizontal boundary $\partial_{\text{hor}} X_{\tilde{U}_{\mathfrak{q}}}$.

STEP 2 : *For every special face $\mathfrak{q}$ of B_ω there is a constant $\epsilon_{\mathfrak{q}}$ such that for any $\mathfrak{q}$-crossing $C \subset$ Cyl,*

$$\int_C u^*(\psi_\aleph^*\omega) \geq \epsilon_{\mathfrak{q}}. \tag{7.36}$$

Consequently, after truncating the domain cylinder by a finite amount there are no compact crossings.

As a warm-up we consider the easiest case of $\mathfrak{q} = B_\omega$. Suppose $C \subset$ Cyl is a B_ω-crossing. The lower bound for area is obtained by applying the monotonicity theorem (Proposition 7.37) to the space $X_{\tilde{U}_{\mathfrak{q}}}$, on which the squashed area form $\psi_\aleph^*\omega$ is a symplectic form. There is a constant $\epsilon > 0$ such that $B_\epsilon(X_{U_{\mathfrak{q}}})$ is contained in $X_{\tilde{U}_{\mathfrak{q}}}$. Indeed, since $\mathfrak{q} = B_\omega$ is top-dimensional, $I_{\mathfrak{q}}$ is a point, and $U_{\mathfrak{q}} = \mathfrak{q}^{\blacksquare}$, $\tilde{U}_{\mathfrak{q}} = \tilde{\mathfrak{q}}^{\blacksquare}$; further $U_{\mathfrak{q}}$ lying in the interior of $\tilde{U}_{\mathfrak{q}}$ together with a uniform bound on the derivative $d(\aleph \circ \pi_{B_J})$ for all squashing maps $\aleph$ implies that ϵ can be chosen uniformly for all $\aleph$. Therefore by the monotonicity theorem there is a constant $\epsilon_{\mathfrak{q}} > 0$ such that for any J-holomorphic map $u : C \to X_{\tilde{U}_{\mathfrak{q}}}$, with $J \in U(\mathfrak{J}_0)$, $u(\partial C) \subset \partial X_{\tilde{U}_{\mathfrak{q}}}$, and with $u(C) \cap X_{U_{\mathfrak{q}}} \neq \emptyset$, there is a lower bound on squashed area given by $\int_C u^*(\psi_\aleph^*\omega) \geq \epsilon_{\mathfrak{q}}$.

Next we consider a proper special face $\mathfrak{q}$ of B_ω and prove the area lower bound on $\mathfrak{q}$-crossings. We observe that $\psi_\aleph^*\omega$ is not a symplectic form on $X_{\tilde{\mathfrak{q}}^{\blacksquare}}$, since the form vanishes along $K_{\mathfrak{q},\mathbb{C}}$-orbits, where $K_{\mathfrak{q},\mathbb{C}} \subset T_{\mathbb{C}}$ is the subtorus generated by $\mathfrak{k}_{\mathfrak{q},\mathbb{C}} = \mathfrak{k}_{\mathfrak{q}} \otimes \mathbb{C} \subset \mathfrak{t} \otimes \mathbb{C}$, and $\mathfrak{k}_{\mathfrak{q}} \subset \mathfrak{t}$ is the orthogonal complement of $\mathfrak{t}_{\mathfrak{q}}$. However, $\psi_\aleph^*\omega$ descends to a symplectic form on the quotient $X_{\tilde{\mathfrak{q}}^{\blacksquare}}/K_{\mathfrak{q},\mathbb{C}}$, which we denote by $\omega_{/\!\!/ K_{\mathfrak{q}}}$. In fact, by the definition of $\psi_\aleph$, there is a symplectomorphism

$$(X_{\tilde{\mathfrak{q}}^{\blacksquare}}/K_{\mathfrak{q},\mathbb{C}}, \omega_{/\!\!/ K_{\mathfrak{q}}}) \simeq \Phi^{-1}(Q)/K_{\mathfrak{q}} \times \tilde{\mathfrak{q}}^{\blacksquare}. \tag{7.37}$$

We also note that any cylindrical almost complex structure J on $X_Q^{\square}$ descends to the quotient

$$\pi_{K_{\mathfrak{q},\mathbb{C}}} : X_{\tilde{U}_{\mathfrak{q}}} \to X_{\tilde{U}_{\mathfrak{q}}}/K_{\mathfrak{q},\mathbb{C}} \simeq X_{\tilde{\mathfrak{q}}^{\blacksquare}}/K_{\mathfrak{q},\mathbb{C}},$$

and therefore we may apply the monotonicity theorem to the projection $\pi_{K_{\mathfrak{q},\mathbb{C}}} \circ (u|C)$ lying in the symplectic manifold $(X_{\tilde{\mathfrak{q}}^{\blacksquare}}/K_{\mathfrak{q},\mathbb{C}}, \omega_{/\!\!/ K_{\mathfrak{q}}})$. We obtain a constant $\epsilon'_{\mathfrak{q}}$, that is independent of u, and satisfies

$$\int_C (\pi_{K_{\mathfrak{q},\mathbb{C}}} \circ u)^* \omega_{/\!\!/ K_{\mathfrak{q}}} \geq \epsilon'_{\mathfrak{q}}. \tag{7.38}$$

It remains to relate the left-hand side= of (7.38) with $\int_C u^*(\psi_\aleph^*\omega)$. As discussed earlier in (7.37), the forms $\pi^*_{K_{\mathfrak{q},\mathbb{C}}}\omega_{/\!\!/ K_{\mathfrak{q}}}$ and $\psi_\aleph^*\omega$ are equal on $X_{\tilde{\mathfrak{q}}^{\blacksquare}}$. However, on the complement $X_{\tilde{U}_{\mathfrak{q}}} \backslash X_{\tilde{\mathfrak{q}}^{\blacksquare}}$, $\psi_\aleph$ is a diffeomorphism

$$X_{\tilde{U}_{\mathfrak{q}}} \backslash X_{\tilde{\mathfrak{q}}^{\blacksquare}} \xrightarrow{\psi_\aleph} \Phi^{-1}(Q) \times (\tilde{U}_{\mathfrak{q}} \backslash \tilde{\mathfrak{q}}^{\blacksquare}),$$

and so, the form $\psi_\aleph^*\omega$ is symplectic on $X_{\tilde{U}_{\mathfrak{q}}} \backslash X_{\tilde{\mathfrak{q}}^{\blacksquare}}$. For any $x \in X_{\tilde{U}_{\mathfrak{q}}} \backslash X_{\tilde{\mathfrak{q}}^{\blacksquare}}$, denote by $\omega_{x,\mathrm{hor}}$ the restriction of $\psi_\aleph^*\omega$ to the complement of $\ker D\pi_{K_{\mathfrak{q},\mathbb{C}}}$. For a locally strongly tamed J_0, we have for any $v \in T_x X_{\tilde{U}_{\mathfrak{q}}}$,

$$(\psi_{\aleph}^* \omega)(v, J_0 v) \geq \omega_{x,\mathrm{hor}}(D\pi_{K_{\mathfrak{q},\mathbb{C}}}(v), D\pi_{K_{\mathfrak{q},\mathbb{C}}}(J_0 v)), \tag{7.39}$$

and by Lemma 3.30 there is a constant c (independent of J, x) such that

$$c\omega_{x,\mathrm{hor}}(D\pi_{K_{\mathfrak{q},\mathbb{C}}}(v), D\pi_{K_{\mathfrak{q},\mathbb{C}}}(J_0 v)) \geq \omega_{/\!\!/ K_{\mathfrak{q}}}(D\pi_{K_{\mathfrak{q},\mathbb{C}}}(v), D\pi_{K_{\mathfrak{q},\mathbb{C}}}(J_0 v)), \tag{7.40}$$

and combining (7.39), (7.40) we get

$$c(\psi_{\aleph}^* \omega)(v, J_0 v) \geq \omega_{/\!\!/ K_{\mathfrak{q}}}(D\pi_{K_{\mathfrak{q},\mathbb{C}}}(v), D\pi_{K_{\mathfrak{q},\mathbb{C}}}(J_0 v)). \tag{7.41}$$

For any cylindrical $J \in U_{\mathfrak{J}}$ from Lemma 7.8, let J_0 be a cylindrical almost complex structure on $X_Q^{\square}$ that is strongly tamed and whose horizontal projection to X_Q is same as that of J. Then the calculation in the last paragraph of the proof of Lemma 7.8 (in particular, the estimate (7.17)) implies that there is a constant c (independent of J, x) such that

$$c(\psi_{\aleph}^* \omega)(v, Jv) \geq (\psi_{\aleph}^* \omega)(v, J_0 v) \tag{7.42}$$

Combining (7.41), (7.42) and using the fact that u is holomorphic with respect to a domain-dependent almost complex structure taking values in $U_{\mathfrak{J}}$, we conclude that there is a constant $c_{\mathfrak{q}}$ (determined by the sets $\mathfrak{q}$, $U_{\mathfrak{q}}$ etc.) such that

$$\int_C (\pi_{K_{\mathfrak{q},\mathbb{C}}} \circ u)^* \omega_{/\!\!/ K_{\mathfrak{q}}} \leq c \int_C u^* (\psi_{\aleph})^* \omega. \tag{7.43}$$

The lower bound (7.36) on squashed area follows from (7.43) and the lower bound on $\int_C (\pi_{K_{\mathfrak{q},\mathbb{C}}} \circ u)^* \omega_{/\!\!/ K_{\mathfrak{q}}}$ from (7.38).

Step 3: *We finish the proof for the case P = point by proving the following Claim.*

Claim 7.44. There is a constant ℓ_0 and a special face $\mathfrak{q} \subset B_\omega$ such that the image $u(\{\ell\} \times S^1)$ intersects $X_{\tilde{U}_{\mathfrak{q}}}$ for all $\ell \geq \ell_0$.

The Claim implies the Lemma for the following reason : The inverse image $\aleph^{-1}(\tilde{U}_{\mathfrak{q}})$ is contained in a proper cone $\mathcal{C} \subset \mathfrak{t}_Q^\vee$. By the boundedness of the derivative $\|du\|_{L^\infty}$, we conclude that the image $\pi_{\mathfrak{t}_Q^\vee} \circ u(\mathrm{Cyl})$ is contained within a bounded distance of $\pi_{\mathfrak{t}_Q^\vee} \circ u(C)$. Therefore, $\pi_{\mathfrak{t}_Q^\vee} \circ u(\mathrm{Cyl})$ is contained in a proper cone that is obtained by translating the cone $\mathbb{C}$.

Proof of Claim. By truncating the cylinder by a finite amount we may assume the squashed area of the map is small enough to ensure that there are no compact crossings. If there is a non-compact $\mathfrak{q}$-crossing C then for all large enough ℓ, the crossing intersects the circle $\{\ell\} \times S^1$, and the Claim follows. Therefore we proceed, assuming that the map u does not have any crossings. The Claim is a consequence of the following statement.

Inductive statement : Suppose for a special face $\mathfrak{q} \subseteq B_\omega$ with $\dim(\mathfrak{q}) \geq 1$, the image $u(\mathrm{Cyl})$ does not intersect $X_{U_{\mathfrak{p}}}$ for any special face $\mathfrak{p} \supset \mathfrak{q}$ of B_ω. Then there is a constant ℓ_0 such that

Case 1: either the image $u(\{\ell\} \times S^1)$ intersects $X_{\tilde{U}_\mathfrak{q}}$ for all $\ell \geq \ell_0$.
Case 2: or the image of $u(\mathrm{Cyl}(\ell_0))$ does not intersect $X_{U_\mathfrak{q}}$.

Indeed applying the statement inductively on the special faces of B_ω, starting with the full face B_ω, lets us conclude that the Claim holds for some $\mathfrak{q}$, $\dim(\mathfrak{q}) \geq 1$, or for some ℓ, the image $u(\mathrm{Cyl}(\ell))$ is contained in $\cup_{\mathfrak{q}:\dim(\mathfrak{q})=0} X_{\tilde{U}_\mathfrak{q}}$. But since the sets $X_{\tilde{U}_\mathfrak{q}}$, $\dim(\mathfrak{q}) = 0$, are disjoint from each other, the image $u(\mathrm{Cyl}(\ell))$ is contained in a unique $X_{\tilde{U}_\mathfrak{q}}$ with $\dim(\mathfrak{q}) = 0$, which implies the Claim.

It remains to prove the Inductive statement. Assuming the contrapositive of Case 1, we will prove Case 2. The contrapositive of Case 1 implies that there is a sequence of points $z_i \to \infty$ in the domain cylinder with $u(z_i) \in X_{U_\mathfrak{q}}$. Since u does not have any crossings, the connected component of C_i of $u^{-1}(X_{\tilde{U}_\mathfrak{q}})$ containing z_i intersects the boundary $\partial\mathrm{Cyl} = \{0\} \times S^1$. Indeed, the image $u(\partial C_i)$ of the boundary does not map to the vertical boundary $\tilde{\mathfrak{q}}^\blacksquare \times (\partial I_\mathfrak{q} \backslash \partial B_\omega)$ of $\tilde{U}_\mathfrak{q}$ since the image $u(\mathrm{Cyl})$ does not intersect $X_{U_\mathfrak{p}}$ for any special face $\mathfrak{p} \supset \mathfrak{q}$. Since $z_i \to \infty$, for any $\ell \geq 0$, $u(\{\ell\} \times S^1)$ intersects $X_{\tilde{U}_\mathfrak{q}}$. This finishes the proof of the Claim, and the Lemma in case $P = \mathit{point}$. □

STEP 4 : *Proof of Lemma for a cut space* $X_P^\square$.
The new feature in the proof for cut spaces is that the special faces of the ω-polytope do not coincide with the faces of the ω-polytope. In order to list the special faces, we first list the faces of B_ω. We recall that the ω and J-polytopes are

$$B_\omega = \mathbb{B}_P^\vee, \quad B_J = \mathrm{NCone}_{P^\vee} B^\vee.$$

The faces of B_ω are given by the bijective correspondence

$$\{(Q_0, Q_1) : Q_0, Q_1 \in \mathcal{P}, Q_1 \subseteq Q_0 \subseteq P\} \to \mathrm{Faces}(\mathbb{B}_P^\vee),$$
$$(Q_0, Q_1) \mapsto F_{Q_0,Q_1} := i_R^{-1}(Q_0) \cap Q_1^\vee,$$

where $R \subset Q_0$ is a point polytope in $\mathcal{P}$. If there are more than one point polytopes in Q_0 the space $i_R^{-1}(Q_0) \subset B^\vee$ is independent of the choice of $R \supset Q_0$.

The special faces of $\mathbb{B}_P^\vee$ are

$$\mathcal{F}_Q := \cup_{R \subseteq Q} F_{Q,R}, \quad Q \subseteq P, Q \in \mathcal{P}. \tag{7.44}$$

See Figure 7.15. Note that $\mathcal{F}_Q$ is a subcomplex of $\mathbb{B}_P^\vee$ of dimension $\dim(Q)$. In particular, $\mathcal{F}_Q$ contains as many top-dimensional polytopes as the number of point polytopes in Q. For a pair $R \subseteq Q$ of polytopes in $\mathcal{P}$, $\mathcal{F}_R \subseteq \mathcal{F}_Q$, and we say that $\mathcal{F}_R$ is a special face of $\mathcal{F}_Q$. The interior $\mathcal{F}_Q^\circ$ is the complement of the proper special faces of F_Q. Let

$$\mathcal{F}_Q^\square \subset \tilde{\mathcal{F}}_Q^\square \subset \mathcal{F}_Q^\circ$$

be neighborhoods in the interiors of special faces (as in (7.32)), and let

$$U_{\mathcal{F}_Q}, \tilde{U}_{\mathcal{F}_Q}$$

be thickenings of $\mathcal{F}_Q^\square, \tilde{\mathcal{F}}_Q^\square$ as in (7.33).

Choose a squashing map $\aleph : B_J \to B_\omega$ such that

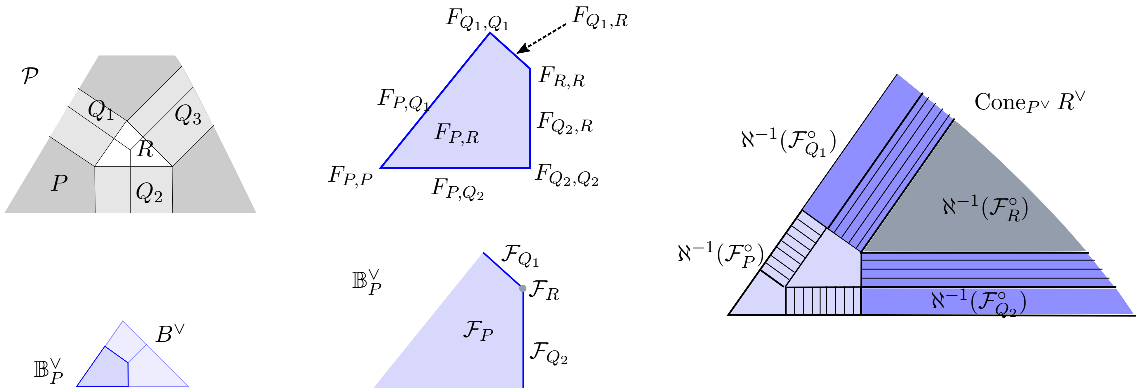

FIGURE 7.15. Left : A polyhedral decomposition $\mathcal{P}$ on (X, ω), dual complex $B^\vee$, and the ω-polytope $\mathbb{B}_P^\vee$ for X_P. Middle top : Faces of $\mathbb{B}_P^\vee$. Middle bottom : Special faces of $\mathbb{B}_P^\vee$. Right : Inverse images of special under a squashing map $\aleph : \mathrm{Cone}_{P^\vee} R^\vee \to \mathbb{B}_P^\vee$.

- for any vertex $Q \in \mathcal{P}$ of P, $\mathbb{B}_{P,Q}^\vee \subset \mathbb{B}_P^\vee$ is unpartitioned in the polyhedral decomposition underlying $\aleph$,
- and the translation $\tau_Q : \mathbb{B}_{P,Q}^\vee \to \mathrm{NCone}_{P^\vee} Q^\vee$ is such that

$$d(\tau_Q(\mathbb{B}_{P,Q}^\vee), \mathrm{NCone}_{P^\vee} R^\vee) > 2\pi \|du\|_{L^\infty}. \tag{7.45}$$

We describe some features of the special faces. On the inverse image $\aleph^{-1}(\mathcal{F}_Q^\circ)$ of the interior of a special face $\mathcal{F}_Q$, the squashing map $\aleph$ is equal to orthogonal projection to $\mathfrak{t}_Q^\vee$. As a consequence the following properties hold on the special faces and their thickenings. (See Figure 7.15.) Let S_Q denote one of the spaces: $S_Q := \mathcal{F}_Q^\circ, \mathcal{F}_Q^\square, \tilde{\mathcal{F}}_Q^\square, U_{\mathcal{F}_Q},$ or $\tilde{U}_{\mathcal{F}_Q}$.

- The inverse image $\pi_{B_J}^{-1} \aleph^{-1}(S_Q)$ is contained in a Q-cylindrical end of $X_P^\square$.
- The projection $\aleph^{-1}(S_Q)/\mathfrak{t}_Q^\vee$ is contained in a compact subset of $NCone_{P^\vee} B^\vee$.

The proofs of both Step 2 and Step 3 carry over to the case of cut spaces, and we conclude that there exists ℓ_0 such that for all $\ell \geq \ell_0$, the image $u(\{\ell\} \times S^1)$ intersects $X_{\tilde{U}_{\mathcal{F}_Q}}$ for some $Q \subseteq P$, $Q \in \mathcal{P}$. We claim that $u(\mathrm{Cyl}(\ell_0))$ lies in the Q-cylindrical end of $X_P^\square$. Indeed $u(\mathrm{Cyl}(\ell_0))$ lies within a $(2\pi\|du\|_{L^\infty})$-radius of $X_{\tilde{U}_{\mathcal{F}_Q}}$, which is part of the Q-cylindrical end by (7.45), since the Q-cylindrical end is

$$\aleph_J^{-1}(\mathrm{NCone}_{P^\vee} B^\vee \backslash \cup_{R \subseteq P : Q \not\subseteq R} \mathrm{NCone}_{P^\vee} R^\vee).$$

Next, the projection from the Q-cylindrical end to $X_Q^\square$ corresponds to the orthogonal projection

$$\pi_{\mathfrak{t}_Q^\vee} : \mathrm{NCone}_{P^\vee} B^\vee \to \mathfrak{t}_Q^\vee / \mathfrak{t}_P^\vee.$$

The image of the projection $\pi_{\mathfrak{t}_Q^\vee}$ of $\aleph^{-1}(tU_{\mathcal{F}_Q})$ is compact, therefore $\pi_{X_Q} \circ u(\mathrm{Cyl}(\ell_0))$ is contained in a compact subset of $X_Q^\square$. Finally the $\mathfrak{t}_Q^\vee/\mathfrak{t}_P^\vee$-coordinate on the image of u is contained in a translate of the cone $\mathrm{NCone}_{P^\vee} Q^\vee$, which is a proper cone.

STEP 5 : *Proof of Lemma for a broken manifold* $X^{\square}_{\mathcal{P}}$.
The only new feature of the proof is the set of special faces of the ω-polytope. The ω-polytope is a product

$$B_\omega = \mathbb{B}^{\vee}_{\mathcal{P}} := \mathbb{B}^{\vee}_{P} \times \delta P^{\vee}.$$

Any special face in $\mathbb{B}^{\vee}_{\mathcal{P}}$ is a product $\mathcal{F}_Q \times \delta R^{\vee}$ of a special face $\mathcal{F}_Q$ of $\mathbb{B}^{\vee}_{P}$ for some $Q \subseteq P$ (see (7.44)); and a special face of $\delta P^{\vee}$ which is $\delta R^{\vee}$ for some $R \supset P$, $R \in \mathcal{P}$. The proof is the same as the case when $\dim(P) = 0$. Finally if the truncation of the cylinder is mapped by $\aleph \circ \pi_{B_J} \circ u$ to a neighborhood of $(\mathcal{F}_Q \times \delta R^{\vee})^{\circ}$, then we show that the image of u is contained in the Q-cylindrical end of $X^{\square}_{\mathcal{P}}$. The other conclusions of the Lemma follow as in Step 3 and Step 4. □

PROOF OF PROPOSITION 7.1 FOR MULTIPLE CUTS. Applying Lemma 7.40 to the map u, we obtain a polytope Q such that, after truncating the domain cylinder by a finite amount, the image of u is contained in the Q-cylindrical end of the target manifold.

We first show that the horizontal projection

$$u_Q := \pi_Q \circ u : \mathrm{Cyl} \to X^{\square}_{Q}$$

has a removable singularity at infinity. Since the image of u_Q lies in a compact subset of $X^{\square}_{Q}$, there is a basic dilated form $\omega^{\mathrm{bas}}_{Q,t} \in \Omega^2(X^{\square}_{Q})$ (Proposition 7.34 (d)) which is non-degenerate on the image of u_Q; and $\omega^{\mathrm{bas}}_{Q,t}$ is taming, since it is constructed using a squashing map. By Proposition 7.29, $\int_{\mathrm{Cyl}} u_Q^* \omega^{\mathrm{bas}}_{Q,t} < E_{\mathrm{Hof}}(u) < \infty$. Therefore the removal of singularity result for compact symplectic manifolds implies that u_Q extends holomorphically to $u_Q : B_1 \to X^{\square}_{Q}$.

We finish the proof by showing that the singularity is removable for the vertical component of the map also. The Q-cylindrical end of $X^{\square}_{P}$ resp. $X^{\square}_{\mathcal{P}}$ may be identified with $X^{\square}_{Q}/T_P$ resp. $X^{\square}_{Q}$, and we view u as mapping to these latter spaces. We consider a holomorphic trivialization of the pullback bundle $u_Q^* X^{\square}_{Q}/T_P$ resp. $u_Q^* X^{\square}_{Q}$ over B_1, and denote the vertical component of the map u by

$$u_v : \mathrm{Cyl} \to T_{Q,\mathbb{C}}/T_{P,\mathbb{C}} \quad \text{resp.} \quad T_{Q,\mathbb{C}}.$$

Identifying $T_{Q,\mathbb{C}} \simeq T_Q \times \mathfrak{t}_Q$ via the Cartan map, and observing that $\pi_{\mathfrak{t}_Q} \circ u_v$ is contained in a proper cone of $\mathfrak{t}_Q/\mathfrak{t}_P$ resp. $\mathfrak{t}_Q$, we conclude that u_v extends over the singularity to a holomorphic map (possibly after passing to a finite quotient) whose target space is a toric orbifold. The target toric orbifold is chosen based on the compactifications of the broken manifold and cut spaces. This concludes the proof of Proposition 7.1. □

The following is a version of the removal of singularities result for punctures on the boundary. This result does not involve any technical difficulties arising from neck-stretching because the Lagrangian submanifold does not intersect the neck regions.

PROPOSITION 7.45. (Removal of singularities on the boundary) *Let*

$$u : (\mathbb{R}_+ \times [0,1], \mathbb{R}_+ \times \{0,1\} \to (X_{P_0}, L)$$

be a perturbed pseudoholomorphic map and

$$E_{\mathrm{Hof}}(u) < \infty, \quad \|\mathrm{d}u\|_{L^\infty(\mathrm{Cyl})} < \infty.$$

Then, u extends to a holomorphic map

$$u : (B_1 \cap \mathbb{H}, B_1 \cap \mathbb{R}) \to (X_{\overline{P}}, L).$$

PROOF. By the uniform bound on the derivative, the image of u is contained in a compact neighborhood of the Lagrangian L. Then there is a dilated basic form $\omega \in \Omega^2(X^{\square}_{\overline{P}})$ that is symplectic on the closure of im(u) and which tames the almost complex structure. Therefore, the Proposition follows from the removal of singularity theorem for compact symplectic manifolds. □

7.6. Hofer energy for Gromov compactness of broken maps

In Section 7.3.2 we gave a definition of Hofer energy which will be used in the proof of Gromov compactness of maps in neck-stretched manifolds. For the proof of convergence of a sequence of broken maps (in the upcoming Section 8.5), we define a different version of Hofer energy, called P-Hofer energy, by choosing a different target space for the squashing map beween complexes.

The reason for having a different definition for P-Hofer energy is so that the following analog of Proposition 7.27, holds with the same proof as that of Proposition 7.27.

PROPOSITION 7.46. (Hofer energy and limits of maps, broken version) *Let $Q \subseteq P$ be a pair of polytopes in $\mathcal{P}$. Let $u_\nu : C \to X^{\square}_{\overline{P}}$ be a sequence of maps, and $t_\nu \in \mathrm{Cone}_{P^\vee} Q^\vee$ a sequence of translations such that*

$$d(t_\nu, \mathrm{Cone}_{P^\vee} R^\vee) \to \infty \quad \forall R \in \mathcal{P} : Q \subset R \subseteq P,$$

the sequence of translated maps $\mathbb{e}^{-t_\nu} u_\nu : C \to X^{\square}_{\overline{Q}}$ converges uniformly to u. Then,

$$E_{\mathrm{Hof},P}(u) \le \liminf_\nu E_{\mathrm{Hof},P}(u_\nu).$$

The proof of Proposition 7.27 (which is the neck-stretched version of Proposition 7.46) was based on the fact that the ω-complex of $X_{\overline{P}}$ is a subset of the ω-complex of X. To imitate the proof of Proposition 7.27 in the broken case, we require that for any $Q \subseteq P$, the ω-complex used to define P-Hofer energy for maps in $X^{\square}_{\overline{Q}}$ is a subset of $\mathbb{B}^\vee_{\overline{P}}$ which is the ω complex of $X_{\overline{P}}$.

For a pair of polytopes $Q \subseteq P$ we now define the ω-complex used for defining P-Hofer energy for maps in $X^{\square}_{\overline{Q}}$. P-Hofer energy is defined by embedding the almost complex manifold $X^{\square}_{\overline{Q}}$ into the symplectic broken manifold $(X_{\overline{P}}, \omega_{X_{\overline{P}}})$. We recall that the tropical moment polytope of $X_{\overline{P}}$ is $\overline{P}$ which is a fibered polytope over P, that is, there is a fibration $\pi_P : \overline{P} \to P$. (See Definition 3.12.) Suppose $\pi_P^{-1}(Q) \subset \overline{P}$ is the intersection of facets $\overline{Q}_1, \dots, \overline{Q}_k$, where

$$\overline{Q}_i = \{x \in \mathfrak{t}^\vee : \langle x, v_i \rangle = \epsilon_i\} \cap \overline{P} \quad \text{for some } v_i \in \mathfrak{t},\ \epsilon_i \in \mathbb{R},$$

and $\overline{P} \subset \{\langle x, v_i\rangle \geq \epsilon_i\}$ for all i. Then the ω-complex for defining P-Hofer energy for maps in $X^{\square}_{\overline{Q}}$ is

$$\mathbb{B}^{P,\vee}_{Q} := \cap_i \{\langle x, v_i\rangle \leq \epsilon_i + \delta_i\} \cap \overline{P},$$

where $\delta_1, \ldots, \delta_k > 0$ are small constants. Thus $\mathbb{B}^{P,\vee}_{Q}$ is a fibered polytope over Q. For a map $u : C \to X^{\square}_{\overline{Q}}$, the *$P$-Hofer energy* is

$$E_{\mathrm{Hof},P}(u) := \sup_{\aleph} \int_C \psi^*_{\aleph} \omega_{\overline{X}_P},$$

where the supremum is over squashing maps $\aleph : \mathrm{Cone}_{Q^\vee} B^\vee \to \mathbb{B}^{P,\vee}_{Q}$. Note that if $Q = P$ then $\mathbb{B}^{P,\vee}_{P} = \mathbb{B}^{\vee}_{P}$, and therefore, $E_{\mathrm{Hof}}(u : C \to X^{\square}_{\overline{P}}) = E_{P,\mathrm{Hof}}(u : C \to X^{\square}_{\overline{P}})$.

With this definition of P-Hofer energy, the proof of Proposition 7.46 is the same as the proof of Proposition 7.27. Next we will see that other features of Hofer energy – such as it being a topological quantity for maps on closed curves, and the finiteness of Hofer energy implying removability of singularities at punctures – also hold in the broken case.

Recall from Remark 7.22 that the Hofer energy of a map $u : (C, \partial C) \to (X^\nu, L)$ on a neck-stretched manifold is equal to the symplectic area. The following is the analogous result for broken maps. The Hofer energy for maps on closed curves is the sum of basic area and a weighted sum of the intersection multiplicities.

PROPOSITION 7.47. *For any polytope $P \in \mathcal{P}$, there is a constant c such that for a perturbed holomorphic map $u : C \to X_{\overline{P}}$ on a closed curve C without boundary that intersects boundary divisors at the points $z_1, \ldots, z_{d(\bullet)}$ with multiplicities $\mathcal{T}(z_i) \in \mathfrak{t}_{\mathbb{Z}}$,*

$$E_{P,\mathrm{Hof}}(u) \leq \int_C (\pi_P \circ u)^* \omega_{X_P} + c \textstyle\sum_{i=1}^{d(\bullet)} |\mathcal{T}(z_i)|.$$

If $u(z_i)$ maps to an intersection of relative divisors, then $|\mathcal{T}(z_i)|$ is the sum of the intersection multiplicities of the map with each of the divisors at z_i.

PROOF. We remark that for a map in $X_{\overline{P}}$, the P-Hofer energy is equal to the ordinary Hofer energy, and both are defined by pulling back the symplectic form on the broken symplectic manifold $(X_{\overline{P}}, \omega_{X_{\overline{P}}})$. For a closed domain C, the Hofer energy of a map $u : C \to X_{\overline{P}}$ is a topological quantity, namely,

$$E_{\mathrm{Hof}}(u) = \int_C (\phi_{\aleph} \circ u)^* \omega_{X_{\overline{P}}} = \langle u_*[C], \omega_{X_{\overline{P}}}\rangle$$

for any squashing map $\aleph : \mathrm{Cone}_{P^\vee} B^\vee \to \mathbb{B}^\vee_P$ and the induced weakly taming map $\phi_{\aleph}$. Now suppose the moment polytope for the fiber of $\pi_P : X_{\overline{P}} \to X_P$ is

$$P^\vee = \{x \in \mathfrak{t}^\vee_P : \langle \mu_i, x\rangle \leq c_i, i = 1, \ldots, N\}$$

where $c_i \in \mathbb{R}$ and $\mu_i \in \mathfrak{t}_{P,\mathbb{Z}}$, $i = 1, \ldots, N$ are primitive outward pointing normals to the facets of $P^\vee$. Suppose the base symplectic form ω_{X_P} on X_P is given by reduction at the origin $0 \in \mathfrak{t}^\vee_P$. Then,

$$\langle u_*[C], \omega_{X_{\overline{P}}}\rangle = \langle (\pi_P \circ u)_*[C], \omega_{X_P}\rangle + 2\pi \textstyle\sum_{i=1}^{d(\bullet)} \sum_j c_j \cdot m(z_i, D_j), \tag{7.46}$$

where $m(z_i, D_j)$ is the intersection multiplicity at z_i with the toric divisor $D_j = \{\langle \mu_j, x\rangle = c_j\}$. Since the origin $0 \in \mathfrak{t}_P^\vee$ is in the interior of $P^\vee$, the constant c_j is positive for all j. The Proposition follows from (7.46). □

Remark 7.48. The removal of singularities result continues to hold (with the same proof) when E_{Hof} is replaced by $E_{\mathrm{Hof},P}$.

CHAPTER 8

Gromov compactness

We prove two convergence results for maps in broken manifolds in this chapter. The first theorem concerns the limit of holomorphic maps to the neck-stretched manifolds X^ν in the limit $\nu \to \infty$. The second theorem is about the convergence behavior of broken maps. Both theorems require only an area bound on the sequence of holomorphic maps.

We sketch the definition of Gromov convergence, with a more precise and detailed definition postponed to Section 8.1. As in symplectic field theory, convergence in neck-stretched manifolds is modulo target rescalings. Rescalings are analogous to 'translations' defined in Section 3.5. Given $\nu \in \mathbb{R}_+$, the set of translations is parametrized by the dual complex $\nu B^\vee$. A translation $t \in \nu P^\vee \subset B^\vee$ is an embedding

$$\mathbb{e}^{-t} : X^\nu_P \to X^\square_P \subset \mathfrak{X}$$

of the P-cylindrical subset X^ν_P of the neck-stretched manifold X^ν. Gromov convergence of a sequence of maps $u_\nu : C_\nu \to X^\nu$ on neck-stretched manifolds to a broken map $u : C \to \mathfrak{X}$ encapsulates the following conditions:

(a) (Domain convergence) The sequence of treed disks C_ν converges to a treed disk C;

(b) (Convergence of maps) for each component C_v of C, there is a sequence of translations $t_\nu(v) \in \nu P(v)^\vee$ such that the translated maps $\mathbb{e}^{-t_\nu} u_\nu$ converges to a map u_v;

(c) (Thin cylinder convergence) for a tropical node w in the domain C of u corresponding to an edge $e = (v_+, v_-)$, on the midpoints

$$m_\nu := (0,0) \in [-l_\nu/2, l_\nu/2] \times \mathbb{R}/\mathbb{Z} \subset C_\nu$$

of the sequence of annuli in C_ν converging to the node w, the translated evaluations $\mathbb{e}^{\frac{1}{2}(t_\nu(v_+)+t_\nu(v_-))} u_\nu(m_\nu)$ converge to the tropical evaluation of u at the nodal point w.

Translation sequences yield tropical positions of vertices in the tropical graph of the limit map, justifying the representation in Figure 1.8. For any ν

$$\mathcal{T}(v) := \tfrac{t_\nu(v)}{\nu} \in P(v)^\vee, \quad v \in \mathrm{Vert}(\Gamma)$$

is a vertex position map for the combinatorial type Γ of the limit map. For a Gromov-converging sequence of maps, area is automatically preserved in the limit, since the number of domain marked points is preserved in the limit, and this number is a constant multiple of area. We remark that (Thin cylinder convergence) is not part of the definition of convergence of stable maps in smooth manifolds, since it can be deduced from convergence of maps. However, in case of neck-stretching, (Thin cylinder convergence) does not follow from (Convergence of maps). Moreover (Thin cylinder convergence) is used in the proof of surjectivity of gluing in Chapter 9.

In the statement of Gromov convergence for maps on neck-stretched manifolds, the perturbations on neck-stretched manifolds are obtained by gluing a perturbation datum on the broken manifold. A cylindrical divisor $\mathfrak{D}$ in a broken manifold $\mathfrak{X}$ can be glued to give a family of divisors D^ν in neck-stretched manifolds X^ν. There is a natural correspondence of tamed cylindrical divisor-adapted almost complex structures

$$\rho_\nu : \mathcal{J}^{\mathrm{cyl}}(\mathfrak{X}, \mathfrak{D}) \to \mathcal{J}^{\mathrm{cyl}}(X^\nu, D^\nu). \tag{8.1}$$

Definition 8.1. (Breaking perturbation datum) Let $\mathfrak{p}$ be a perturbation datum for the broken manifold $\mathfrak{X}$ adapted to a cylindrical divisor $\mathfrak{D}$. A *breaking perturbation datum* is a family of perturbation data $\{\mathfrak{p}_\nu\}_{\nu \in [\nu_0, \infty]}$ for the neck-stretched manifolds (X^ν, D^ν) satisfying the following property: there exists a constant $\nu(\Gamma)$ so that if $\nu > \nu(\Gamma)$ then $\mathfrak{p}_{\nu,\Gamma} := \rho_\nu(\mathfrak{p}_\Gamma)$.

The statement of the convergence theorem does not require perturbations be regular, but it does require perturbations to be adapted to a stabilizing pair (D^ν, J^ν) on X^ν, which is obtained by gluing a stabilizing pair $(\mathfrak{D}, \mathfrak{J})$ on the broken manifold (constructed in Section 5.3). We recall that by definition of a stabilizing pair (D^ν, J^ν), any non-constant J^ν-holomorphic sphere in X^ν intersects D^ν transversely at atleast 3 points.

Theorem 8.2. (Gromov convergence for breaking maps) *Suppose* $(\mathfrak{J}_0, \mathfrak{D})$ *is a stabilizing pair for the broken manifold* $\mathfrak{X}$*, and* $\mathfrak{p}^\infty$ *is a perturbation datum on* $\mathfrak{X}$ *adapted to* $(\mathfrak{J}_0, \mathfrak{D})$*. Suppose* $\{\mathfrak{p}^\nu\}_\nu$ *is a sequence of breaking perturbation data (as in Definition 8.1) on neck-stretched manifolds* $\{X^\nu\}_\nu$*. Let*

$$u_\nu : C_\nu \to X^\nu$$

be a sequence of treed $\mathfrak{p}^\nu$*-adapted disks with uniformly bounded area* $\mathrm{Area}(u_\nu)$*. There is a subsequence of* $\{u_\nu\}_\nu$ *that Gromov converges to a* $\mathfrak{p}^\infty$*-adapted stable broken disk* $u : C \to \mathfrak{X}_{\mathcal{P}}$ *modelled on a tropical graph* Γ*. The limit* u *is unique up to domain reparametrizations and the action of the identity component of the tropical symmetry group* $T_{\mathrm{trop},\mathcal{W}}(\Gamma)$ *(see (4.28)).*

The second result of the Chapter concerns convergence of a sequence of broken maps with uniformly bounded area. After passing to a subsequence the maps u_ν have the same type, say Γ with tropical structure $\mathcal{T}$. The limit u typically has additional domain components. The tropical structure $\mathcal{T}'$ of the limit is related to $\mathcal{T}$ by a *tropical edge collapse morphism*, which is a morphism of graphs $\mathcal{T}' \xrightarrow{\kappa} \mathcal{T}$ that collapses a subset of edges of $\mathcal{T}'$, and the slopes of uncollapsed edges are the same in $\mathcal{T}$ and $\mathcal{T}'$. See Figure 8.1 for an example.

Theorem 8.3. (Gromov convergence for broken maps) *Suppose* $(\mathfrak{J}_0, \mathfrak{D})$ *is a stabilizing pair for the broken manifold* $\mathfrak{X}_{\mathcal{P}}$*, and* $\mathfrak{p}$ *is a perturbation datum on* $\mathfrak{X}_{\mathcal{P}}$

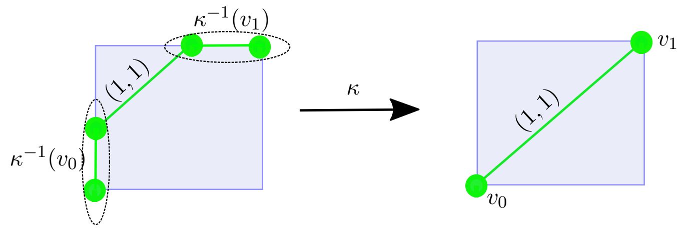

FIGURE 8.1. Tropical edge collapse morphism between tropical graphs.

adapted to $(\mathfrak{J}_0, \mathfrak{D})$. *Suppose* $u_\nu : C_\nu \to \mathfrak{X}_{\mathcal{P}}$ *is a sequence of adapted broken* $\mathfrak{p}$-*holomorphic disks with uniformly bounded area.*

After passing to a subsequence the type of the broken disks u_ν *is* ν-*independent, and is equal to, say,* Γ. *The sequence* u_ν *Gromov converges to a broken adapted* $\mathfrak{p}$-*disk* $u : C \to \mathfrak{X}_{\mathcal{P}}$ *of type* Γ' *for which there is a tropical edge-collapse morphism* $\Gamma' \to \Gamma$. *The limit* u *is unique up to domain reparametrizations and the action of the identity component of the tropical symmetry group* $T_{\mathrm{trop},\mathcal{W}}(\Gamma)$ *(see (4.28)).*

If the edge-collapse $\Gamma' \to \Gamma$ *is non-trivial (in the sense of Definition 8.26) then* $\dim_{\mathbb{C}}(T_{\mathrm{trop}}(\Gamma')) > \dim_{\mathbb{C}}(T_{\mathrm{trop}}(\Gamma))$.

REMARK 8.4. In theorems 8.2 and 8.3, the limit broken map is uniquely determined if it is rigid. Indeed for a rigid tropical graph Γ the tropical symmetry group is finite.

8.1. Gromov convergence

In this section, we define the notion of convergence in the compactness results, Theorem 8.2 and Theorem 8.3. We first fix identifications between domain curves that are close to each other in the compactified moduli space of stable curves. This is done via a choice of an exponential map in the neighborhood of nodes. Different choices lead to the same notion of convergence. We start by recalling a construction of deformation of a stable nodal curve. We describe the construction for rational stable nodal curves (without boundary). The extensions to treed disks is explained later in Remark 8.9. Let Γ be the combinatorial type of a rational stable nodal curve with $d(\bullet)$ markings. The moduli space of curves $\mathcal{M}_\Gamma$ is a submanifold of the compactified moduli space $\mathcal{M}_{d(\bullet)}$ whose tubular neighborhood can be described as follows. For a stable curve S in $\mathcal{M}_\Gamma$, let $\tilde{S}$ be the normalization of S at the nodal points. For any edge $e \in \mathrm{Edge}_{\bullet,-}(\Gamma)$, we fix a map

$$\exp^S_{w_\pm(e)} : (U(T_{w_\pm(e)}\tilde{S}), 0) \to (U_{w_\pm(e)}(\tilde{S}), w_\pm(e)), \qquad U(T_{w_\pm(e)}\tilde{S}) \subset T_{w_\pm(e)}\tilde{S}, \quad U_{w_\pm(e)}(\tilde{S}) \subset \tilde{S} \tag{8.2}$$

that biholomorphically maps a neighborhood $U(T_{w_\pm(e)}\tilde{S})$ of the origin in the tangent space onto a neighborhood $U_{w_\pm(e)}(\tilde{S})$ of the lift of the node w_e, satisfies $d\exp^S_{w_\pm(e)}(0) = \mathrm{Id}$, and varies smoothly with S. The family of maps $\{\exp_{w_\pm(e)} : e \in \mathrm{Edge}_{\bullet,-}(\Gamma)\}$ is fixed for the rest of this paper. Whenever we choose complex coordinates

$$z_\pm : (U_{w_\pm(e)}(\tilde{S}), w_\pm(e)) \to (\mathbb{C}, 0)$$

in neighborhoods of the node we assume that they are compatible with $\exp^S_{w_\pm(e)}$. That is, $z_\pm$ is the composition of $\exp_{w^\pm}$ with a linear map on the tangent space. On a neighborhood

$$U_{\mathcal{M}_\Gamma} \subset \mathcal{M}_{d(\bullet)}$$

of $\mathcal{M}_\Gamma$, there is a projection map

$$\pi_\Gamma : U_{\mathcal{M}_\Gamma} \to \mathcal{M}_\Gamma.$$

such that curves in a fiber $\pi_\Gamma^{-1}(S)$ are obtained by gluing the interior nodes of S as follows. (See, for example, Siebert [**84**, Proposition 2.4] in the closed case; the open case is similar.)

Definition 8.5. Given a nodal sphere S of type Γ, a *collection of gluing parameters* is a tuple

$$\delta = (\delta_e)_{e \in \mathrm{Edge}_-(\Gamma)} : \delta_e \in T_{w_+(e)}\tilde{S} \otimes T_{w_-(e)}\tilde{S}.$$

The *curve corresponding to a gluing parameter* δ is defined by

$$S^\delta := (S \backslash \cup_e U'_{w_\pm(e)}(\tilde{S}))/\sim, \qquad z_+ \sim z_- \Leftrightarrow (\exp^S_{w_+(e)})^{-1}(z_+) \otimes (\exp^S_{w_-(e)})^{-1}(z_-) = \delta_e. \tag{8.3}$$

Here $U'_{w_\pm(e)}(\tilde{S}) \subset \tilde{S}$ is a neighborhood of $w_\pm(e)$ such that the boundary $\partial U'_{w_\pm(e)}(\tilde{S})$ is identified to $\partial U_{w_e^\mp}(\tilde{S})$ (with reversed orientation) by the equivalence relation $\sim$. Thus, for any node w_e, the relation $\sim$ identifies the pair of annuli

$$U_{w_+(e)}(\tilde{S}) \backslash U'_{w_+(e)}(\tilde{S}) \xrightarrow{\sim} U_{w_-(e)}(\tilde{S}) \backslash U'_{w_-(e)}(\tilde{S}). \tag{8.4}$$

The resulting annulus in S^ν is called $\mathrm{Neck}_e(S^\nu)$, and

$$\mathrm{Neck}(S^\nu) := \cup_e \mathrm{Neck}_e(S^\nu) \subset S^\nu. \tag{8.5}$$

The gluing construction maps a neighborhood of zero in the space of gluing parameters to a neighbourhood of S in the fiber $\pi_\Gamma^{-1}(S)$. Curves in the neighborhood $U_{\mathcal{M}_\Gamma}$ of $\mathcal{M}_\Gamma$ possess a *gluing parameter* function for every smoothened node

$$\delta = (\delta_e)_e : U_{\mathcal{M}_\Gamma} \to \prod_{e \in \mathrm{Edge}_{\bullet,-}(\Gamma)} T_{w_+(e)}\tilde{S} \otimes T_{w_-(e)}\tilde{S}, \quad S^\delta \mapsto \delta.$$

Remark 8.6. At a node w_e if we choose a framing $\mathrm{fr} : T_{w_+(e)}\tilde{S} \otimes T_{w_-(e)}\tilde{S} \to \mathbb{C}$ as in (4.12) or if we choose complex coordinates

$$z_+ : (\tilde{S}, w_+(e)) \to (\mathbb{C}, 0), \quad z_- : (\tilde{S}, w_-(e)) \to (\mathbb{C}, 0),$$

then the gluing parameter δ_e is a complex number.

Remark 8.7. (Identifications between nearby curves) Consider rational stable curves S, S' with $[S] \in \mathcal{M}_\Gamma$ and $[S'] \in \mathcal{M}_{d(\bullet)}$ in the neighborhood of S. The complement of nodes S° can be identified to subsets of the curve S' as follows. First suppose that S' is obtained by gluing S. That is, $S' \in \pi_\Gamma^{-1}(S)$. For a vertex

$v \in \mathrm{Vert}(\Gamma)$, let $S'(v) \subset S'$ be the subset corresponding to the component $S_v \subset S$ including the necks $\mathrm{Neck}_e(S')$, $v \in e$. The subsets $S'(v)$ cover S'. By the gluing construction there are natural inclusions

$$i_{S_v,S'} : S'(v) \to S_v \subset S. \tag{8.6}$$

Next suppose S' is not in $\pi_\Gamma^{-1}(S)$. There is a neighborhood $U_{\Gamma,S} \subset \mathcal{M}_\Gamma$ of S on which the universal curve $\mathcal{U}_\Gamma$ can be trivialized, so that the variation of the complex structure is given by a map

$$U_{\Gamma,S} \to \mathcal{J}(\mathcal{U}_\Gamma), \quad m \mapsto j(m) \tag{8.7}$$

for which $j(m)$ is m-independent in the neighborhood of special points. The trivialization of the universal curve $\mathcal{U}_\Gamma$ gives a diffeomorphism $\phi : \pi_\Gamma(S') \to S$. The maps $i_{S_v,S'}$ are defined as compositions

$$i_{S_v,S'} := \phi \circ i_{\pi_\Gamma(S')_v,S'}.$$

If a sequence S_ν of curves converges to S, then the images $i_{S_v,S_\nu}(S_\nu(v))$ exhaust the complement of nodes S_v° as $\nu \to \infty$.

Remark 8.8. (Uniqueness of identifications) The identifications between regions of nearby curves are unique in the following sense. Suppose a sequence $[S_\nu] \in \overline{\mathcal{M}}_{d(\bullet),d(\circ)}$ converges to a curve $[S]$. For any node $e = (v_+, v_-)$ in S, the identification in the neck region

$$i_{S_{v_\pm},S^\nu} : \mathrm{Neck}_e(S^\nu) \to S_{v_\pm} \tag{8.8}$$

is uniquely determined by the choice of the exponential map (8.2). On the complement of the neck region let

$$\phi_\nu, \phi'_\nu : S^\nu \setminus \mathrm{Neck}(S^\nu) \to S$$

be two possible identifications given by trivializations of the universal curve. Then ϕ_ν, ϕ'_ν have the same image, and the maps $\phi'_\nu \circ \phi_\nu^{-1}$ converge to the identity uniformly in all derivatives.

Remark 8.9. (Identifications of nearby treed disks) The identifications between nearby nodal spheres extend to treed disks. Let C be a stable treed disk of type Γ with surface part S and treed part T.

(a) (Glued treed disk) For a disk node $e \in \mathrm{Edge}_\circ^0(\Gamma)$ of zero length, $\ell(e) = 0$, a *gluing parameter* at the node $w_e \in C$ is an element

$$\delta_e \in (T_{w_e^+}\partial S \otimes T_{w_e^-}\partial S)_{\leq 0} \subset T_{w_e^+}S \otimes T_{w_e^-}S$$

where the orientation on $T_{w_e^+}\partial S \otimes T_{w_e^-}\partial S$ is derived from the orientation on the boundaries of components of S. Given gluing parameters

$$\delta := (\delta_e)_{e \in \mathrm{Edge}_{\bullet,-}(\Gamma) \cup \mathrm{Edge}_\circ^0(\Gamma)}$$

for all nodes in the surface component S, a glued surface

$$S^\delta$$

is defined exactly as in (8.3). The glued treed curve is $C^\delta := S^\delta \cup T$, that is, the treed part in C^δ is the same as that of C.

(b) (Domain identifications) Let C be a treed curve of type Γ, and let Γ' be a treed disk type obtained by collapsing a subset of interior edges $e \in \mathrm{Edge}_{\bullet,-}(\Gamma)$ and zero length boundary edges $e \in \mathrm{Edge}^0_{\circ,-}(\Gamma)$. Let $C' = S' \cup T'$ be a treed disk of type Γ' that is close to C. For any surface component $S_v \subset C$ corresponding to a vertex $v \in \mathrm{Vert}(\Gamma)$, there is a subset $S'(v) \subset S'$ and an identification

$$i_{S_v,S'} : S'(v) \to S_v$$

exactly as in (8.6).

Next we identify punctured neighborhoods of interior nodes to annuli in nearby curves. These identifications are needed for stating the (Thin cylinder convergence) for tropical nodes in the definition Gromov convergence. We make these identifications only for interior nodes since all tropical nodes are interior nodes.

Definition 8.10. (Annuli converging to a node) Let C_ν be a sequence of treed disks converging to a limit curve C for which the arguments $\frac{\delta_e(C_\nu)}{|\delta_e(C_\nu)|}$ of the gluing parameters converge for all interior edges $e \in \mathrm{Edge}_{\bullet,-}(\Gamma)$. Let $w \in C$ be an interior node corresponding to the edge $e = (v_+, v_-) \in \mathrm{Edge}_{\bullet,-}(\Gamma)$.

(a) A sequence of *annuli* $A_\nu \subset C_\nu$ *converge to a node* w in C if there are $U_{w_\pm} \subset S_{v_\pm}$ of the lifts w_+, w_- of w such that

$$A_\nu = i^{-1}_{S_{v_+},C_\nu}(U_{w_+}) \cap i^{-1}_{S_{v_-},C_\nu}(U_{w_-}).$$

We say that the sequence of annuli A_ν is obtained by *gluing the node* w.

(b) (Centered annuli converging to a node) Suppose the node $w \in C$ is equipped with complex coordinates

$$z_\pm : (U_{w_\pm}, w_\pm) \to (\mathbb{C}, 0)$$

on neighborhoods $U_{w_\pm} \subset C_{v_\pm}$ of its lifts w_+, w_-. A *sequence of centered annuli converging to the node* w is a sequence of parametrized annuli

$$A_\nu := [-l'_\nu/2, l'_\nu/2] \times \mathbb{R}/2\pi\mathbb{Z} \hookrightarrow C_\nu$$

for which there exists $\epsilon > 0$ such that

$$A_\nu = i^{-1}_{S_{v_+},C_\nu}(\{|z_+| \le \epsilon\}) \cap i^{-1}_{S_{v_-},C_\nu}(\{|z_-| \le \epsilon\}),$$

and the map

$$A_\nu \xrightarrow{i_{S_{v_\pm},C_\nu}} C_{v_\pm} \xrightarrow{z_\pm} \mathbb{C} \quad \text{is equal to} \quad (s,t) \mapsto \exp(\mp(s+it) - (l_\nu + i\theta_\nu)/2).$$

Here $e^{-(l_\nu + i\theta_\nu)} \in \mathbb{C}^\times$ is the neck length parameter for the curve C_ν resulting from the choice of coordinates $z_\pm$.

Informally we may restate the above definition as follows: A sequence of annuli converging to a node w is obtained by gluing neighborhoods U_{w_+}, U_{w_-} of the nodal lifts w_+, w_-. Given holomorphic coordinates in a neighborhood of the nodal w_+, w_-, a sequence of centered annuli converging to a node w is given by gluing neighborhoods of w_+, w_- that have the same radius.

The following definition describes rescalings of the target spaces required in the statement of convergence of components of the broken map. We recall from Definition 3.41 that a translation is a way of identifying a subset of a neck-stretched manifold X^ν to a subset of the broken manifold $\mathfrak{X}$, and the set of translations is parametrized by $\nu B^\vee$ where $B^\vee$ is the dual complex of the neck-stretching. A translation sequence for a tropical graph consists of a translation sequence corresponding to each of the vertices of the tropical graph that are compatible with the edge slopes of the tropical graph.

Definition 8.11. (Translation sequence for a tropical graph) Suppose Γ is a tropical graph. A Γ-*translation sequence* consists of a collection of sequences

$$\{t_\nu(v) \in \nu B^\vee\}_\nu, \quad v \in \mathrm{Vert}(\Gamma)$$

such that the following conditions hold :

(a) (Polytope) For each vertex v

$$t_\nu(v) \in \nu P(v)^\vee \subset \nu B^\vee, \tag{8.9}$$

and for any polytope $P_0 \in \mathcal{P}$, $P_0 \supset P(v)$,

$$d_{B^\vee}(t_\nu(v), \nu P_0^\vee) \to \infty \quad \text{as} \quad \nu \to \infty.$$

(b) (Slope) For any node e between vertices v_+, v_-, there is a sequence $l_\nu(e) \to \infty$ such that

$$t_\nu(v_+) - t_\nu(v_+) = \mathcal{T}(e) l_\nu. \tag{8.10}$$

Definition 8.12. (Gromov convergence, multiple cuts) Let $C_\nu = S_\nu \cup T_\nu$ be a sequence of stable treed disks of type Γ_0. A sequence of holomorphic maps $u_\nu : (C_\nu, \underline{z}_\nu) \to X^\nu$ converges to a broken map $u : C = (S \cup T) \to \mathfrak{X}_{\mathcal{P}}$ of type Γ and framing fr if the following are satisfied.

(a) (Convergence of domains) The sequence of treed disks C_ν converges to C and for any tropical node w_e of the limit map u the arguments $\frac{\delta_e(C_\nu)}{|\delta_e(C_\nu)|}$ of the gluing parameters converge to a limit. Using the fixed holomorphic exponential map (8.2), let $S_\nu(v) \subset S_\nu$ be the subset corresponding to a vertex $v \in \mathrm{Vert}(\Gamma)$, and let

$$i_{v,\nu} := i_{S_v, C_\nu} : S_\nu(v) \to S_v, \quad S_\nu(v) \subset S_\nu,$$

be embeddings from (8.8) whose images $i_{v,\nu}(S_\nu(v))$ exhaust the complement of nodes S_v° as $\nu \to \infty$.

(b) (Convergence of maps) There is a Γ-translation sequence $\{t_\nu(v)\}_{v,\nu}$ such that for any irreducible surface component $S_v \subset S$, the sequence of maps

$$S_v^\circ \supset i_{v,\nu}(S_\nu(v)) \xrightarrow{\mathbb{e}^{-t_\nu(v)}(u_\nu \circ i_{v,\nu}^{-1})} X^\square_{\overline{P(v)}}$$

converges in $C^\infty_{\mathrm{loc}}(S_v^\circ)$ to $u_v : S_v^\circ \to X^\square_{\overline{P(v)}}$. The map $\mathbb{e}^{-t_\nu(v)} : X^\nu_{P(v)} \to X_{\overline{P(v)}}$ is defined in (3.41). For each boundary edge e in Γ_0, the map $u_\nu | T_{e,\nu}$ on the treed segment converges to a (possibly broken) treed segment in u.

(c) (Thin cylinder convergence) For a node w in C corresponding to a tropical edge $e = (v_+, v_-) \in \mathrm{Edge}_{\mathrm{trop}}(\Gamma)$, let

$$z_\pm : (U_{w_\pm}, w_\pm) \to (\mathbb{C}, 0)$$

be matching coordinates (see Definition 4.12 following (4.18)) on neighborhoods $U_{w_\pm} \subset S_{v_\pm}$ of the nodal point which respect the framing fr_e. Let

$$A(l_\nu) := [-l_\nu/2, l_\nu/2] \times S^1 \subset S_\nu$$

be a sequence of centered annuli converging to the node w, see Definition 8.10. Then the sequence

$$x_\nu := \mathbb{e}^{-\frac{1}{2}(t_\nu(v_+)+t_\nu(v_-))} u_\nu(0,0) \in X^{\square}_{\overline{P(e)}}$$

converges to a limit x_0, and the components $u_{v_\pm}$ of the broken map are asymptotically close to

$$z_\pm \mapsto z_\pm^{\mathcal{T}(e)} x_0.$$

The uniqueness of the limit, when it exists, is explained in Theorem 8.2.

8.2. Horizontal convergence

We discuss a notion of compactness for sequence of points in neck-stretched manifolds, that is useful in the proof of convergence for breaking maps.

Definition 8.13. (Horizontal convergence) A sequence of points $x_\nu \in X^\nu$ *horizontally converges* to a point $x \in X^\square_P$ for a polytope $P \in \mathcal{P}$ if

- $x_\nu \in X^\nu_{\overline{P}}$ for all ν, and the sequence $\pi^\nu_P(x_\nu)$ converges to x, where $\pi^\nu_P : X^\nu_{\overline{P}} \to X^\square_P$ is the projection map $X^\nu_{\overline{P}} \to X^\nu_P$ from (3.28) composed with the inclusion $X^\nu_P \to X^\square_P$,
- and for any subsequence of $\{x_\nu\}_\nu$, the above condition is not satisfied for any polytope $P_0 \supset P$.

We present some results on horizontal convergence of sequences:

Lemma 8.14. *For any sequence $x_\nu \in X^\nu$, there exists a subsequence of $\{x_{\nu_k}\}_k$ that converges horizontally. There is a unique polytope $P \in \mathcal{P}$ for which the subsequence $\{x_{\nu_k}\}_k$ converges horizontally in X_P.*

Proof. Recall from (7.10) that there is a projection to the dual polytope $\pi_{\nu B^\vee} : X^\nu \to \nu B^\vee$ for all ν. There is a polytope $P \in \mathcal{P}$ such that, after passing to a subsequence, $(x_\nu)_\nu$ satisfies

$$\sup_\nu d(\pi_{\nu B^\vee}(x_\nu), \nu P^\vee) < \infty, \quad d(P_0^\vee, \pi_{\nu B^\vee}(x_\nu)) \to \infty \;\; \forall P_0 \supset P. \tag{8.11}$$

It is then clear that a subsequence of $(x_\nu)_\nu$ and the polytope P satisfy the conclusions of the Lemma. □

The following Lemma relates horizontal convergence to a property of translation sequences.

LEMMA 8.15. *Suppose* $x_\nu \in X^\nu$ *is a sequence of points, and* $t_\nu \in \nu P^\vee$ *is a sequence of translations such that* $\mathbb{e}^{-t_\nu} x_\nu$ *converges in* $X_P^\square$*. Then, the following are equivalent:*

(a) The sequence $x_\nu \in X^\nu$ *converges horizontally in* X_P*.*
(b) For any $P_0 \supset P$ *(that is,* P *is a proper face of* P_0*),* $d(t_\nu, \nu P_0^\vee) \to \infty$*.*

PROOF. Firstly we observe that the convergence of the translated sequence $\mathbb{e}^{-t_\nu} x_\nu$ implies that $d(\pi_{\nu B^\vee}(x_\nu), \nu P^\vee)$ is uniformly bounded, and that $(\pi_P(x_\nu))_\nu$ converges to a limit in $X_P^\square$. The condition in (b) then implies that a $(x_\nu)_\nu$ horizontally converges in X_P. The converse is similar and is left to the reader. □

The next result may be seen as the horizontal convergence version of the Arzela-Ascoli theorem.

LEMMA 8.16. *Suppose* C *is a connected curve and* $u_\nu : C \to X^\nu$ *is a sequence of differentiable maps satisfying* $\sup_\nu \|\mathrm{d}u_\nu\|_{L^\infty} < \infty$*. There exists a subsequence of maps* $\{u_{\nu_k}\}_k$ *and a polytope* $P \in \mathcal{P}$ *such that for all* $z \in C$ *the sequence* $u_{\nu_k}(z)$ *converges horizontally in* X_P*. For the subsequence* $\{u_{\nu_k}\}_k$*, the polytope* P *is unique.*

PROOF. First assume that C is compact. The proof of Lemma 8.14 carries over. Indeed, by the uniform bound on $\sup_\nu \|\mathrm{d}u_\nu\|_{L^\infty}$, the sequence $d(\pi_{\nu B^\vee}(u_\nu(z_0)), P^\vee)$ is uniformly bounded for a fixed point $z_0 \in C$ iff the sequence $d(\pi_{\nu B^\vee}(u_\nu(z)), P^\vee)$ is uniformly bounded for any $z \in C$. The result also holds for non-compact curves since they are exhausted by a sequence of compact curves. □

8.3. Breaking annuli

The next proposition governs the behavior of annuli with small base area and bounded Hofer energy. For any $L > 0$, we denote by

$$A(L) := [-\tfrac{L}{2}, \tfrac{L}{2}] \times S^1 \cong \{z \in \mathbb{C} \mid |z| \in [e^{-L/2}, e^{L/2}]\}$$

the cylinder with length L equipped with the product metric.

PROPOSITION 8.17. (Breaking annulus lemma) *Let* $\mathfrak{J}_0$ *be a cylindrical almost complex structure on* $\mathfrak{X}$ *for which Hofer energy is monotonic (see Lemma 7.23), and suppose that for any* ν *the almost complex structures* J_0^ν *is obtained by gluing* $\mathfrak{J}_0$ *on the necks. There are constants* $0 < \rho < 1$*,* $c > 0$ *such that the following hold. For a sequence* $l_\nu \to \infty$*, let*

$$u_\nu : A(l_\nu) \to X^\nu$$

be a sequence of J_0^ν*-holomorphic maps satisfying*

(1) $\sup_\nu E_{\mathrm{Hof}}(u_\nu) < \infty$, $\sup_{z \in \mathbb{R}\times S^1, \nu} |\mathrm{d}u_\nu(z)| < \infty$*, where the domain annuli* $A(l_\nu) := [-\frac{l_\nu}{2}, \frac{l_\nu}{2}] \times S^1$ *are equipped with the product metric.*
(2) For all $t \geq 1$*,*

$$\lim_{L\to\infty} \lim_{\nu\to\infty} \mathrm{Area}_t^{\mathrm{bas}}(u_\nu, A(l_\nu - L)) = 0. \tag{8.12}$$

(3) There exist polytopes P_+, P_- *such that the sequence of maps* $u_\nu(\cdot \pm \frac{l_\nu}{2})$ *converges horizontally in* $X_{P_\pm}$*.*

(4) There exists a sequence of translations $t_\nu^\pm \in \nu P_\pm^\vee$ *such that the sequence* $\mathrm{e}^{-t_\nu} u_\nu(\cdot \mp \frac{l_\nu}{2})$ *converges in* C^∞_{loc} *to*

$$u_\pm : \mathbb{R}_\mp \times S^1 \to X^\square_{P_\pm},$$

and the map extends holomorphically over $\mp\infty$*, possibly after passing to a finite cover in the orbifold case.*

Then, there exists $\mu \in \mathfrak{t}_{P,\mathbb{Z}}$*,* $P_\cap := P_+ \cap P_-$ *for which the following hold.*

(a) The sequence $t_\nu^+ - t_\nu^- - \mu l_\nu \in \mathfrak{t}^\vee_{P_\cap}$ *is uniformly bounded.*

(b) (Horizontal matching) *The points* $\pi_{P_+}(u_+(-\infty))$*,* $\pi_{P_-}(u_-(+\infty))$ *lie in* $X^\square_{P_\cap} \subset X_{P_\pm}$ *and* $\pi_{P_+}(u_+(-\infty)) = \pi_{P_-}(u_-(+\infty))$.

(c) (Asymptotic decay) *Let* ξ_ν *be a section defined by the relation*

$$u_\nu = \exp_{u_{\nu,\mathrm{triv}}} \xi_\nu, \quad u_{\nu,\mathrm{triv}}(s,t) := e^{\mu(s+it)} u_\nu(0,0).$$

There exists $l \geq 0$ *and a subsequence such that for* $k = 0, 1$,

$$|D^k \xi_\nu(s,t)| \leq c(e^{\rho(s-\frac{l_\nu}{2})} + e^{\rho(-s-\frac{l_\nu}{2})}), \quad \forall s \in \left[-\tfrac{l_\nu}{2} + l, \tfrac{l_\nu}{2} - l\right]. \tag{8.13}$$

Proof of Proposition 8.17. The Proposition is proved using the annulus lemma on compact symplectic manifolds.

Step 1 : *Proof of horizontal matching.*
Suppose $x_0 := \pi_{P_-}(u_-(\infty)) \in X^\square_P$. Then $P \subseteq P_-$.

First we will show that, after truncating the domain cylinders by a ν-independent amount, the π_P-projections of the images are contained in a small neighborhood of x_0. Choose t such that the dilated basic area form $\omega^{\mathrm{bas}}_{P,t} \in \Omega^2(X^\square_P)$ (defined in Proposition 7.34) is a symplectic form in a neighborhood $U_t \subset X_P$ of x_0. Since $P \subseteq P_-$ the convergence of $\mathrm{e}^{-t_\nu^-} u_\nu$ to u_- implies that the projection $\pi_P(u_\nu)$ converges in C^∞_{loc} to $\pi_P(u_-)$. Consider a constant $\epsilon > 0$ such that $\overline{B}_\epsilon(x_0) \subset U_t$. Choose a sequence of points $z_\nu = s_\nu + it_\nu \in A(l_\nu)$ such that

$$x_\nu := \pi_P(u_\nu(z_\nu)) \in B_{\epsilon/2}(x_0), \quad |\pm \tfrac{l_\nu}{2} - s_\nu| \to \infty.$$

We apply the monotonicity theorem to $\pi_P \circ u_\nu$ on the ball $B(x_\nu, \frac{\epsilon}{2})$ and obtain a constant $\rho > 0$ such that

$$\int_{u_\nu^{-1}(B_{\epsilon/2}(x_\nu))} (\pi_P \circ u_\nu)^* \omega_{P,t} > \rho.$$

Away from the ends of the cylinder the basic area goes to zero as $\nu \to \infty$ by (8.18), and so, there exists L_0 and ν_0 such that

$$\mathrm{Area}^t_{\mathrm{bas}}(u_\nu, A(l_\nu - L_0)) \leq \rho.$$

Consequently

$$\pi_P u_\nu(A(l_\nu - L_0)) \subset B_{\epsilon/2}(x_\nu) \subset B_\epsilon(x_0) \subset X^\square_P \tag{8.14}$$

for $\nu \geq \nu_0$.

We claim $P \subseteq P_+$. Indeed by (8.14) a subsequence of $u_\nu(\cdot + \frac{l_\nu}{2})$ horizontally converges in a polytope $P' \supseteq P$. By the hypothesis of the Proposition and the uniqueness of P' (by Lemma 8.14) we conclude that $P' = P_+$.

Next we prove (Horizontal Matching) using the fact that any neighborhood of x_0 contains the images of the cylinders truncated by a ν-independent amount. Since $P \subseteq P_+$ the convergence of the translated maps $\mathbb{e}^{-t_\nu^+} u_\nu(\cdot + \frac{l_\nu}{2})$ implies that the sequence $\pi_P(u_\nu(\cdot + \frac{l_\nu}{2}))$ converges to $\pi_P(u_+)$. By (8.14) we conclude

$$\pi_P(u_+(-\infty)) \in \overline{B}_{\epsilon/2}(z_\nu) \subset B_\epsilon(x_0).$$

Since ϵ can be chosen to be arbitrarily small,

$$\pi_P(u_+(-\infty)) = \pi_P(u_-(+\infty)).$$

This finishes the proof of horizontal matching modulo the proof of the fact that $P = P_+ \cap P_-$. We have shown that $P \subseteq P_+ \cap P_-$ and the proof of equality is postponed to the last step of the Proof.

Step 2 : *Determining the edge slope μ.*
In this step we read off the slope μ of the edge from the topology of the cylinders u_ν and show that it is equal to the slopes of the limit maps at the nodal point. We have shown that there exists L_0 such that the images $\pi_P(u_\nu(A(l_\nu - L_0)))$ lie in a neighborhood $B_\epsilon(x_0) \subset X_P^\square$. Consider a trivialization of $X_{\overline{P}}^\square \to X_P^\square$

$$B_\epsilon(x_0) \times T_{P,\mathbb{C}} \simeq \pi_P^{-1}(B_\epsilon(x_0)) \subset X_{\overline{P}}^\square.$$

Viewing the target space of u_ν as the product $B_\epsilon \times T_{P,\mathbb{C}}$, the homotopy class $(u_\nu)_*[A(l_\nu)] \in \pi_1(T_P)$ corresponds to an element $\mu_\nu \in \mathfrak{t}_{P,\mathbb{Z}}$. Let $\mu \in \mathfrak{t}_{P,\mathbb{Z}}$ be defined so that u_- is asymptotically close to a trivial cylinder of slope μ at ∞. Therefore $\mu_\nu = \mu$ for large ν. Similarly u_+ is also asymptotically close to trivial cylinder of slope μ at $-\infty$.

Step 3: *The sequence of twisted maps*

$$\overline{u}_\nu : A(l_\nu) \to X_{\overline{P}}^\square, \quad (s,t) \mapsto e^{-\mu(s+\frac{l_\nu}{2}+it)}(\mathbb{e}^{-t_\nu^-} u_\nu)$$

converges to a pair of disks connected at an interior point.

Because the derivatives of the maps $(\overline{u}_\nu)_\nu$ are uniformly bounded, the sequence has a Gromov limit consisting of two disks connected by a path of spheres. Each of the components is obtained by a sequence of rescalings of the cylinder of the form $(s,t) \mapsto (s + s_\nu, t)$ for some sequence $s_\nu \in \mathbb{R}$.

We show that one of the components in this Gromov-Floer limit of $\overline{u}_\nu$ is a disk $\overline{u}_-$ which is a twisted version of u_-. The convergence of u_ν to u_- implies that the sequence $\overline{u}_\nu(\cdot - \frac{l_\nu}{2})$ converges to $\overline{u}_- := e^{-\mu(s+it)}u_-$. The image of $\overline{u}_-$ is compact in $X_{\overline{P}}^\square$ since u_- is asymptotically close to the μ-cylinder $(s,t) \mapsto e^{\mu(s+it)}$. Since $\pi_P \circ u_-$ extends over ∞, the same is also true for $\overline{u}_-$. We denote $\overline{x}_0 := \overline{u}_-(\infty)$.

Next, we describe the component of the limit map attached to the disk $\overline{u}_-$ at the nodal point ∞ by a soft rescaling argument. Choose a taming symplectic form $\omega_{\overline{P}}$ defined in a neighborhood $U'_{\overline{x}_0} \subset X_{\overline{P}}^\square$ of $\overline{x}_0$. The form $\omega_{\overline{P}}$ can, for example, be taken to be a dilated basic area form as in Proposition 7.34 but the particular properties of basic forms are not relevant to us here. Let $U_{\overline{x}_0} \Subset U'_{\overline{x}_0}$ be a smaller open neighborhood whose closure is contained in $U'_{\overline{x}_0}$, and let the constants c, $\hbar$,

μ be from the annulus lemma on compact manifolds (Proposition 8.18) applied to $(\overline{U}_{\overline{x}_0}, \omega_{\overline{P}}, J_0)$. Since $\overline{u}_\nu(\cdot - \frac{l_\nu}{2})$ converges to $\overline{u}_-$, there exists a constant r_0 such that for any $l \geq 0$ there exists $\nu_0(l)$ such that

$$\overline{u}_\nu([\tfrac{-l_\nu}{2} + r_0, \tfrac{-l_\nu}{2} + r_0 + l] \times S^1) \subset U_{\overline{x}_0} \quad \text{for} \quad \nu \geq \nu_0(l).$$

Let

$$e_0 := \omega_{\overline{P}}(\overline{u}_-, [r_0, \infty) \times S^1) + \tfrac{\hbar}{2}.$$

Define

$$r_\nu := \sup\{r : \overline{u}_\nu([\tfrac{-l_\nu}{2} + r_0, r] \times S^1) \subset U_{\overline{x}_0},\ \omega_{\overline{P}}(u_\nu, [\tfrac{-l_\nu}{2} + r_0, r] \times S^1) \leq e_0\}. \tag{8.15}$$

The definition implies $\frac{l_\nu}{2} + r_\nu \to \infty$. Consider the rescaled maps

$$\overline{u}_\nu^+(s,t) := \overline{u}_\nu(\cdot + r_\nu) : [\tfrac{-l_\nu}{2} - r_\nu, \tfrac{l_\nu}{2} - r_\nu] \times S^1 \to X_{\overline{P}}^{\square}.$$

Since $|d\overline{u}_\nu^+|$ is uniformly bounded and $\overline{u}_\nu^+(\{0\} \times S^1)$ lies in a compact set of $X_{\overline{P}}^{\square}$ we conclude that a subsequence of $\{\overline{u}_\nu^+\}_\nu$ converges in C^∞_{loc} to a limit $\overline{u}_+ : (-\infty, L_1) \times S^1 \to X_{\overline{P}}^{\square}$. Here

$$L_1 := \lim_\nu (\tfrac{l_\nu}{2} - r_\nu) \tag{8.16}$$

could possibly be ∞, and by the definition of r_ν in (8.15),

$$\omega_{\overline{P}}(\overline{u}_+, (-\infty, L_1) \times S^1) \leq \tfrac{\hbar}{2}$$

By removal of singularities on the compact symplectic manifold $(\overline{U}_{\overline{x}_0}, \omega_{\overline{P}}, J_0)$, the map $\overline{u}_+$ extends over $-\infty$. The images of the components $\overline{u}_-$ and $\overline{u}_+$ connect at the nodal point. Indeed, by applying Proposition 8.19 to the sequence of maps $\overline{u}_\nu(\cdot - \frac{(r_0+r_\nu)}{2})$ on the cylinders $A(r_\nu - r_0)$ with target space $(\overline{U}_{\overline{x}_0}, \omega_{\overline{P}})$ we conclude $\overline{u}_-(\infty) = \overline{u}_+(-\infty)$.

So far we have shown that the sequence of twisted maps $\overline{u}_\nu$ contains a disk $\overline{u}_-$ in its Gromov limit, and there is a non-constant component $\overline{u}_+$ that is attached to $\overline{u}_-$ at an interior node. We claim that the component $\overline{u}_+$ is a punctured disk and not a punctured sphere. To prove the Claim we assume that $\overline{u}_+$ is a punctured sphere, and arrive at a contradiction by showing that it is a constant map. Our assumption implies that the domain of $\overline{u}_+$ is the cylinder $\mathbb{R} \times S^1$, and we showed in the previous paragraph that $\overline{u}_+$ extends over $-\infty$ with $\overline{u}_+(-\infty) \in X_{\overline{P}}^{\square}$. Next, we show that the projected map $\pi_P \circ \overline{u}_+$ is constant. The π_P-projections of u_ν and $\overline{u}_\nu$ are the same, and hence,

$$\omega_{P,t}^{\text{bas}}(\pi_P \circ \overline{u}_+) \leq \limsup_\nu \omega_{P,t}^{\text{bas}}(u_\nu) \leq E_{\text{Hof}}(u_\nu) < \infty.$$

By (8.14) the image of $\pi_P \circ \overline{u}_+$ lies in a neighborhood of $X_P^{\square}$ where $\omega_{P,t}^{\text{bas}}$ is a symplectic form. Therefore $\pi_P \circ \overline{u}_+$ extends over $+\infty$ and is a holomorphic sphere. The projected map $\pi_P(\overline{u}_+)$ is constant because by (8.14) $\pi_P(\overline{u}_+) \subset B_\epsilon(x_0)$. Indeed a contractible neighborhood in a symplectic manifold can not contain a symplectic sphere. The map $\overline{u}_+ : \{-\infty\} \cup (\mathbb{R} \times S^1) \to X_{\overline{P}}^{\square}$ is also constant, because it lies in a single fiber of the $T_{P,\mathbb{C}}$-bundle $X_{\overline{P}}^{\square} \to X_P^{\square}$, and has only a single pole at $+\infty$. Therefore we conclude that L_1 is finite and $\overline{u}_+$ is a disk.

We have shown that the limit of the twisted maps is a pair of disks $(\overline{u}_-, \overline{u}_+)$. Since $L_1 = \lim_\nu(\frac{l_\nu}{2} - r_\nu)$ is finite, after truncating the domain cylinders by a ν-independent amount, we may replace r_ν by $\frac{l_\nu}{2}$. The limit $\overline{u}_+$ will be altered by a domain reparametrization and $\overline{u}_+(-\infty) = \overline{u}_-(\infty)$ continues to hold.

Part (a) of the Proposition is now proved as follows. We have shown that the maps

$$\overline{u}_\nu^+(s,t) := e^{-\mu(s+it)}(\mathbb{e}^{-t_\nu^- - \mu l_\nu} u_\nu(s + \tfrac{l_\nu}{2}, t))$$

and $\mathbb{e}^{-t_\nu^+} u_\nu(\cdot + \frac{l_\nu}{2})$ converge on $\mathbb{R}_{\geq 0} \times S^1$. At the point $(s,t) = (0,0)$ the sequences $\overline{u}_\nu^+(s,t)$ and $\mathbb{e}^{-t_\nu^+} u_\nu(\cdot + \frac{l_\nu}{2})$ differ by a translation by $e^{t_\nu^+ - t_\nu^- - \mu l_\nu}$. Since both sequences of points converge we conclude that the limit

$$\delta := \lim_\nu(-\mu l_\nu - t_\nu^- + t_\nu^+)$$

exists (which proves (a)) and that

$$\overline{u}_+(s,t) = e^\delta e^{-\mu(s+it)} u_+(s,t).$$

Step 4 : *Proof of the decay estimate.*
We have shown that after truncating the domain cylinders by a ν-independent amount, the sequence of the twisted maps $\overline{u}_\nu$ converge to a pair of disks $(\overline{u}_-, \overline{u}_+)$, and the images of the maps lie in a compact set $\overline{U}_{\overline{x}_0}$ with a taming symplectic form $\omega_{\overline{P}}$. The sequence of domains can truncated again by a finite amount to ensure that $\omega_{\overline{P}}(\overline{u}_\nu) < \hbar$, since the $\omega_{\overline{P}}$-area on the sequence of cylinders converges to the $\omega_{\overline{P}}$-area of the pair of disks $(\overline{u}_-, \overline{u}_+)$. We apply the annulus lemma for compact manifolds (Proposition 8.18) to the maps

$$\overline{u}_\nu : A(l_\nu - L) \to (\overline{U}_{\overline{x}_0}, \omega_{\overline{P}}) \subset X_{\overline{P}}^\square.$$

The decay estimate for the twisted maps $\overline{u}_\nu$ implies the asymptotic decay estimate (8.13) for u_ν required by the Proposition.

Step 5 : $P = P_+ \cap P_-$.
The polytope $P \in \mathcal{P}$ was chosen so that the nodal point on u_- is on $X_P^\square : (\pi_{P_-} u_-)(\infty) \in X_P^\square \subset X_{P_-}$. Therefore at the nodal point ∞, u_- intersects all the divisors X_Q of X_{P_-} which contain X_P, that is, $P \subseteq Q \subseteq P_-$. Consequently the slope μ lies in $\mathfrak{t}_P^\vee \setminus \cup_{Q \supset P} \mathfrak{t}_Q^\vee$. On the other hand since $t_\nu^+ - t_\nu^- - \nu l_\nu$ is uniformly bounded, and $t_\nu^\pm \in \nu P_\pm^\vee \subset \mathfrak{t}_{P_\pm}^\vee$, we conclude $\mu \in \mathfrak{t}_{P_\cap}^\vee$, where $P_\cap := P_+ \cap P_-$. Therefore $P_\cap = P$. $\square$

The proof of the breaking annulus lemma was based on the following results (Propositions 8.18 and 8.19) on compact symplectic manifolds.

Proposition 8.18. (Annulus lemma on compact manifolds) *Suppose (X, ω) is a compact symplectic manifold with a tamed almost complex structure J. There exists constants $0 < \rho < 1$, $\hbar > 0$, $c > 0$ such that the following holds for any J-holomorphic map $u : A(L) \to X$ with $E(u) \leq \hbar$. For $x = u(0,0)$, there is a map*

$$\xi : A(L-1) \to T_x X \quad \textit{such that} \quad u = \exp_x \xi$$

on $A(L-1)$ *and for* $k = 0, 1$,

$$|D^k\xi(s,t)| \leq c(e^{\rho(s-L)} + e^{\rho(-s-L)}), \quad \forall s \in [-L+1, L-1]. \tag{8.17}$$

The constants ρ, $\hbar$, c *depend continuously on* J *with respect to the* C^2*-topology.*

Proof. The proposition is a consequence of the annulus lemma, see [**60**, Lemma 4.7.3] and elliptic regularity for holomorphic maps. □

Proposition 8.19. (Convergence of long cylinders) *Suppose* (X, ω) *is a compact symplectic manifold with a tamed almost complex structure* J*. Let* $u_\nu : A(l_\nu) \to X$ *be a sequence of holomorphic cylinders satisfying* $\omega(u_\nu) \leq \hbar$*, where* $\hbar$ *is as in Proposition 8.18. After passing to a subsequence,* $u_\nu(\cdot \pm \frac{l_\nu}{2})$ *converges in* C^∞_{loc} *to*

$$u_\pm : \mathbb{R}_\mp \times S^1 \to X,$$

the map $u_\pm$ *extends holomorphically over* $\mp\infty$*, and* $u_-(\infty) = u_+(-\infty)$*. Further,*

$$\lim_\nu \text{Area}_\omega(u_\nu) = \text{Area}_\omega(u_+) + \text{Area}_\omega(u_-).$$

Proof. This Proposition is proved as part of the 'bubbles connect' result in [**60**, Proposition 4.7.1]. □

We give an annulus lemma for sequences of holomorphic strips in X^ν whose boundaries lie on the Lagrangian submanifold L. The result is simpler than the breaking annulus lemma since the strips converge to a nodal disk in the complement of the relative divisors, and not a tropical node. The result is stated for maps whose domains are strips, defined as

$$A_\circ(L) := [-\tfrac{L}{2}, \tfrac{L}{2}] \times [0,1]$$

for any $L > 0$. The result is a boundary version of the 'bubbles connect' result in McDuff-Salamon [**60**, Proposition 4.7.1]. But since the manifold X_{P_0} has cylindrical ends, we need to use a Hofer energy bound.

Proposition 8.20. (Annulus lemma with boundary) *Suppose the almost complex structures* $\mathfrak{J}_0 \in \mathcal{J}^{\text{cyl}}(\mathfrak{X})$ *and* $J_0^\nu \in \mathcal{J}^{\text{cyl}}(X^\nu)$ *are as in the statement of Proposition 8.17. There are constants* $0 < \rho < 1$*,* $c > 0$ *such that the following hold. For a sequence* $l_\nu \to \infty$*, let*

$$u_\nu : A_\circ(l_\nu) \to X^\nu, \quad u_\nu(\partial A_\nu) \subset L$$

be a sequence of J_0^ν*-holomorphic strips satisfying the following.*

(1) $\sup_\nu E_{\text{Hof}}(u_\nu) < \infty$, $\sup_{z \in \mathbb{R}\times[0,1],\nu} |du_\nu(z)| < \infty$.

(2) For all $t \geq 1$,

$$\lim_{L\to\infty} \lim_{\nu\to\infty} \text{Area}_t^{\text{bas}}(u_\nu, A(l_\nu - L)) = 0. \tag{8.18}$$

(3) The sequence of maps $u_\nu(\cdot \pm \frac{l_\nu}{2})$ *converges in* C^∞_{loc} *to a limit*

$$u_\pm : \mathbb{R}_\mp \times [0,1] \to X^\square_{P_0},$$

and the map extends holomorphically over $\mp\infty$.

Then, $u_+(-\infty) = u_-(+\infty) \in L$.

The proof carries over verbatim from Step 1 of the proof of Proposition 8.17.

8.4. Proof of convergence for breaking maps

Proof of Theorem 8.2. Step 1 : *Domain components for the limit map.* The sequence of domains (C_ν, z_ν) converges to a stable treed nodal curve (C, z) modelled on a tree Γ. The perturbation maps $\mathfrak{p}_\nu$ converge to a perturbation datum $\mathfrak{p}_\infty = (J_\infty, F_\infty)$ defined on C. In the next few steps we will show that there are no additional domain components in the limit map.

In the rest of the proof we focus on the convergence of maps on irreducible surface components. The proof of convergence of gradient trajectories on treed components does not involve any technical difficulties arising from multiple cutting, and are therefore left to the reader.

Step 2 : *Boundedness of derivatives.*
In this step we show a bound on the derivatives of the sequence of maps by ruling out bubble trees with unstable domains in the limit. In particular, we will show that for any $v \in \mathrm{Vert}(\Gamma)$, the derivative of

$$u_{v,\nu} := u_\nu \circ i_{v,\nu}^{-1} : C_v^\circ \to X^\nu$$

is uniformly bounded for all ν. The norm on the derivative is with respect to the cylindrical metric on X^ν. On the domain we use a metric on C° that is cylindrical (or strip-like) in the punctured neighborhood of nodal points and on the neck regions in C_ν.

Assuming, for the sake of contradiction, that the derivatives are not uniformly bounded, we will construct a sequence of rescaled maps. After passing to a subsequence, there exists a sequence of points $z_\nu \in C_v^\circ$ and a point $z_\infty \in C_v$ for some $v \in \mathrm{Vert}(\Gamma)$ such that

$$z_\nu \to z_\infty, \quad |\mathrm{d}u_{v,\nu}(z_\nu)| \to \infty.$$

We first carry out the proof assuming that z_∞ is not a nodal point and does not lie on the boundary ∂C_v, and thus $z_\infty \in C_v^\circ$. We apply Hofer's lemma 8.21 to the function $|\mathrm{d}u_{v,\nu}|$, $x = z_\nu$ and the constant $\delta = |\mathrm{d}u_\nu(z_\nu)|^{-1/2}$.

Lemma 8.21. (Hofer's Lemma, [**60**, Lemma 4.6.4]) *Suppose (X, d) is a metric space, $f : X \to \mathbb{R}_{\geq 0}$ is a continuous function, and $x \in X$, $\delta > 0$ are such that the ball $B_{2\delta}(x)$ is complete. Then there exists a positive constant $\epsilon \leq \delta$ and a point $\zeta \in B_{2\epsilon}(x)$ such that*

$$\sup_{z \in B_\epsilon} f(z) \leq 2f(\zeta), \quad \epsilon f(\zeta) \leq \delta f(x).$$

We obtain another sequence $\zeta_\nu \in C_v^\circ$ converging to z_∞, and a sequence of constants $\epsilon_\nu \to 0$ such that

$$c_\nu := |\mathrm{d}u_{v,\nu}(\zeta_\nu)| \to \infty, \quad \sup_{z \in B_{\epsilon_\nu}} |\mathrm{d}u_\nu(z)| \leq 2c_\nu, \quad c_\nu \epsilon_\nu \to \infty.$$

The rescaled maps are

$$\tilde{u}_\nu := u_{v,\nu}((\cdot - \zeta_\nu)/c_\nu) : B_{\epsilon_\nu c_\nu} \to X^\nu,$$

and $|d\tilde{u}_\nu| \leq 2$, $|d\tilde{u}_\nu(0)| = 1$.

Each of the rescaled maps converges to a limiting map with domain the affine line as follows. By Lemma 8.16 there is a polytope $P \in \mathcal{P}$ such that, a subsequence of $\tilde{u}_\nu$ horizontally converges in X_P. We may view $\tilde{u}_\nu$ as mapping to the P-cylinder, since for any translation $t_\nu \in \nu P^\vee \subset \mathfrak{t}_P^\vee$ there is an embedding

$$\mathbb{e}^{-t_\nu} : X_{\overline{P}}^{\nu} \to X_{\overline{P}}^{\square},$$

see (3.41). We choose a sequence of translations $t_\nu \in t_P^\vee$ so that $\pi_{\mathfrak{t}_P^\vee} \circ (\mathbb{e}^{-t_\nu}\tilde{u}_\nu)(0)$ is ν-independent. The horizontal convergence of $\tilde{u}_\nu$ and the uniform bound on the derivatives of $\tilde{u}_\nu$ imply that for any compact set $K \subset \mathbb{C}$, the images $\mathbb{e}^{-t_\nu}\tilde{u}_\nu(K)$ are contained in a compact set of $X_{\overline{P}}^{\square}$. By the Arzela-Ascoli theorem, the sequence $\mathbb{e}^{-t_\nu}\tilde{u}_\nu$ converges in C^∞_{loc} to a non-constant J_{z_∞}-holomorphic limit $\tilde{u} : \mathbb{C} \to X_{\overline{P}}^{\square}$.

The limit map in the previous paragraph extends to a holomorphic map from the projective line by removal of singularities. Hofer energy $E_{\mathrm{Hof}}(u_\nu, C_\nu)$ is equal to area (Remark 7.22) and therefore is uniformly bounded for all ν. By monotonicity of Hofer energy (Lemma 7.23) $E_{\mathrm{Hof}}(\tilde{u}_\nu, B_{\epsilon_\nu c_\nu})$ is uniformly bounded. The quantity $E_{\mathrm{Hof}}(\tilde{u}, \mathbb{C})$ is finite because E_{Hof} is preserved in the neck-stretching limit. Therefore by Proposition 7.1 the singularity at ∞ can be removed to produce an extension $\tilde{u} : \mathbb{P}^1 \to X_{\overline{P}}$. (The extension is defined on the weighted projective line $\mathbb{P}(1, n)$ for some n if the singular point ∞ maps to an orbifold singularity of $X_{\overline{P}}$.)

Finally, we arrive at a contradiction by showing that the limit map is constant. The domain reparametrizations $\phi_{v,\nu}$ were derived from the stable map compactification, and so, there is at most a single marked point $z_{i,\nu}$ that is contained in each of the regions $B_{\epsilon_\nu}(\zeta_\nu)$. Therefore, the projection $\pi_P \circ \tilde{u} : \mathbb{C} \to X_P^{\square}$ either lies in the stabilizing divisor D_P, or it has at most one intersection with the divisor D_P. Since $(\mathfrak{J}_0, \mathfrak{D})$ is a stabilizing pair for $\mathfrak{X}$, and the perturbation $\mathfrak{p}$ is adapted to $(\mathfrak{J}_0, \mathfrak{D})$ neither possibilities can happen if $\pi_P \circ \tilde{u}$ is non-constant. Therefore, $\pi_P \circ \tilde{u}$ is constant, and so the image of $\tilde{u}$ lies in a single toric fiber $V_{P^\vee}$. The image $\tilde{u}(\mathbb{C})$ does not intersect torus-invariant divisors of $V_{P^\vee}$, and therefore $\tilde{u}$ is a constant map.

We now consider the case that the sequence of points with increasing derivatives converges to an interior nodal point w_e. The sequence z_ν lies on the neck region

$$A_\nu := [\tfrac{-l_\nu}{2}, \tfrac{l_\nu}{2}] \times S^1 \subset C_\nu$$

obtained by gluing the node w_e, and $l_\nu \to \infty$ as $\nu \to \infty$. Since z_ν converges to the node w_e, there is a constant $\epsilon > 0$ such that the ball $B_\epsilon(z_\nu)$ is contained in A_ν for all ν. The earlier proof using Hofer's Lemma and a rescaling can now be applied to the maps $u_\nu | B_\epsilon(z_\nu)$, because there are no marked points in A_ν. We conclude that $|\mathrm{d}u_{v,\nu}|$ is uniformly bounded for all ν.

Finally if the sequence of points with increasing derivatives converges to a point on the boundary of the domain, the same steps can be repeated to show the existence of a non-constant disk bubble with unstable domain, leading to a contradiction. The removal of singularities in this case is proved by Proposition 7.45.

STEP 3 : *Determining stable components of the limit map.*
In Step 2 we showed that the derivatives of the maps $u_{v,\nu}$ are uniformly bounded, where v is any vertex of Γ. Therefore, by Lemma 8.16 the maps $u_{v,\nu}$ converge

horizontally in some polytope $P(v) \in \mathcal{P}$, and the image of $u_{v,\nu}$ is contained in $X^{\nu}_{P(v)}$. For any choice of translation $t_\nu(v) \in \nu P^\vee \subset \mathfrak{t}^\vee_P$, there is an embedding $\mathbb{e}^{-t_\nu(v)} : X^{\nu}_{P(v)} \to X^{\square}_{\overline{P}(v)}$ into the $P(v)$-cylinder, see (3.41). Choose a smooth point $z_v \in C_v$ in each curve component. Fix the translation $t_\nu(v) \in \nu P(v)^\vee$ so that $\pi_{\mathfrak{t}^\vee_{P(v)}}(\mathbb{e}^{-t_\nu(v)} u_{v,\nu}(z_v))$ is ν-independent. Then, for any compact set in the complement of nodal points $K \subset C^\circ_v$, the images $(\mathbb{e}^{-t_\nu(v)} u_{v,\nu})(K)$ are contained in a uniformly bounded region of $X^{\square}_{\overline{P}(v)}$. Since the derivatives of $\mathbb{e}^{-t_\nu(v)} u_{v,\nu}$ are locally uniformly bounded, the sequence $\mathbb{e}^{-t_\nu(v)} u_{v,\nu}$ converges in C^∞_{loc} to a limit $u_v : C^\circ_v \to X^{\square}_{\overline{P}(v)}$. By arguments as in Step 2, $E_{\text{Hof}}(u_v, C^\circ_v)$ is bounded. Therefore, by removal of singularity (Proposition 7.1),

- u_v has a removable singularity at the node w,
- or in a punctured neighborhood of the node w, u_v is asymptotic to a vertical cylinder, and this latter case can only happen if w is an interior point of the domain.

We obtain a J_∞-holomorphic map $u_v : C^\circ_v \to X^{\square}_{\overline{P}(v)}$ that is adapted to the stabilizing divisor $D_{\overline{P}(v)} \subset \mathfrak{D}$.

For later use, we remark that the convergence continues to hold if the sequence $t_\nu(v)$ is replaced by a sequence $t_\nu(v)'$ for which $\sup_\nu |t_\nu(v)' - t_\nu(v)| < \infty$. In that case, after passing to a subsequence, the limit u_v would be replaced by $e^{-t} u_v$, where $t := \lim_\nu (t_\nu(v)' - t_\nu(v))$.

Step 4 : *No area is lost at nodes.*
Next we show that for any node w in the limit curve C and a sequence of annuli (resp. strips) in the domain curves C_ν that are obtained by gluing neighborhoods $U_+, U_- \subset S \subset C$ of the nodal lifts w_+, w_-, the area limit on the annuli is equal to the area of the limit map on $U_+ \cup U_-$. We first consider the case that w is an interior node. Suppose the annuli

$$A(l_\nu) = [-\tfrac{l_\nu}{2}, \tfrac{l_\nu}{2}] \times S^1 \subset C_\nu \tag{8.19}$$

obtained by gluing the node w in C. First, we pass to a subsequence so that the limit

$$\text{Area}^{\text{bas}}(w) := \lim_{l\to\infty} \lim_{\nu\to\infty} \text{Area}^{\text{bas}}(u_\nu, [-\tfrac{l_\nu}{2} + l, \tfrac{l_\nu}{2} - l] \times S^1) \tag{8.20}$$

exists. It is enough to prove that $\text{Area}^{\text{bas}}(w)$ vanishes. Suppose for the sake of contradiction that $\text{Area}^{\text{bas}}(w) > 0$. We recall that by Step 3, the sequence of maps $\mathbb{e}^{-t_\nu(v_-)} u_\nu(\cdot + l_\nu/2)$ converges in C^∞_{loc} to $u_{v_-} : \mathbb{R}_+ \times S^1 \to X_{\overline{P}(v_-)}$. We now find a limit unstable bubble attached to C_{v_-} at ∞. The reparametrization sequence $r_\nu \in [0, l_\nu]$ and a subsequence of $\{u_\nu\}_\nu$ are chosen so that

$$\text{Area}^{\text{bas}}(u_\nu, [\tfrac{-l_\nu}{2}, \tfrac{-l_\nu}{2} + r_\nu] \times S^1) = \text{Area}^{\text{bas}}(u_{v_-}, \mathbb{R}_+ \times S^1) + m_w,$$

where

$$m_w := \tfrac{1}{2} \min\{\text{Area}^{\text{bas}}(w), \hbar\}.$$

If $\mathrm{Area}^{\mathrm{bas}}(w) > 0$, then $l_\nu - r_\nu \to \infty$. A sequence of rescaled maps is defined as

$$\tilde{u}_\nu := u_\nu(\cdot - \tfrac{l_\nu}{2} + r_\nu).$$

We will arrive at a contradiction by showing that the rescaled sequence of maps has a limit that is horizontally non-constant. Since the derivatives of the maps u_ν are uniformly bounded, arguments as in Steps 2 and 3 imply that

(a) there is a polytope $P_w \in \mathcal{P}$ such that the sequence $\tilde{u}_\nu$ converges horizontally in X_{P_w},
(b) and there is a sequence of translations $t_{w,\nu} \in \nu P_w^\vee$ such that the sequence $\mathbb{e}^{-t_{w,\nu}}\tilde{u}_\nu$ converges to a $\mathfrak{J}_0$-holomorphic limit $u_w : \mathbb{R} \times S^1 \to X^{\square}_{\overline{P}_w}$, which extends holomorphically to $u_w : \mathbb{P}^1 \to X_{\overline{P}_w}$.

By Lemma 8.15, the translations $t_{w,\nu}$ diverge from faces of $\nu P_w^\vee$, that is,

$$d(\nu P_0^\vee, t_{w,\nu}) \to \infty \quad \forall P_0 \supset P_w.$$

Therefore, by Proposition 7.34 (c), the maps $\tilde{u}_\nu$ map to a region where the basic area form $\omega_\nu^{\mathrm{bas}}$ is $\mathfrak{t}_{P_w}^\vee$-translation invariant. Therefore, for any compact set K,

$$\lim_\nu \mathrm{Area}^{\mathrm{bas}}(\tilde{u}_\nu, K) = \mathrm{Area}^{\mathrm{bas}}(u_w, K).$$

We conclude that

$$\mathrm{Area}^{\mathrm{bas}}(u_w, \mathbb{R}_- \times S^1) = m_w > 0.$$

Since there are no marked points in the neck region $A(l_\nu)$, the sphere $\pi_{P_w} \circ u_w$ either maps to the stabilizing divisor D_{P_w} or intersects the divisor at a maximum of two points denoted $+\infty, -\infty \in C_w$. Since $(\mathfrak{J}_0, D_{P_w})$ is a stabilizing pair on X_{P_w}, both possibilities are ruled out. We conclude that $\mathrm{Area}^{\mathrm{bas}}(w) = 0$ for all nodes.

Next we consider the case that w is a boundary node in C with no treed segment. That is, for the edge $e \in \mathrm{Edge}_{\circ,-}(\Gamma)$ corresponding to w, the length $\ell(e)$ is zero. The same argument as above works with annuli replaced by strips

$$A_\circ(l_\nu) := [-\tfrac{l_\nu}{2}, \tfrac{l_\nu}{2}] \times [0,1] \subset C_\nu$$

which are obtained by gluing at the boundary node w. Assuming that

$$\mathrm{Area}^{\mathrm{bas}}(w) := \lim_{l\to\infty} \lim_{\nu\to\infty} \mathrm{Area}^{\mathrm{bas}}(u_\nu, [-\tfrac{l_\nu}{2} + l, \tfrac{l_\nu}{2} - l] \times [0,1])$$

is positive, we find a limit map $u_w : \mathbb{R} \times [0,1] \to X_{P_0}$ that has positive area and whose boundary $\mathbb{R} \times \{0,1\}$ maps to the Lagrangian submanifold, and the image of u_w does not have any intersections with the stablizing divisor. The singularities at $+\infty, -\infty$ are removable by Proposition 7.45. Thus we have a non-constant $\mathfrak{J}_0$-holomorphic disk in (X_{P_0}, L) with no intersections with the stabilizing divisor $\mathfrak{D}$. Since $(\mathfrak{J}_0, \mathfrak{D})$ is a stabilizing pair in X_{P_0}, we have arrived at a contradiction. We conclude that $\mathrm{Area}^{\mathrm{bas}}(w) = 0$ for boundary nodes in C.

For later use, we observe that the same arguments apply for any dilated basic area form (see Definition 7.36). Thus we have proved for any $t \geq 1$, and an interior node w in C,

$$\mathrm{Area}_t^{\mathrm{bas}}(w) := \lim_{l\to\infty} \lim_{\nu\to\infty} \mathrm{Area}_t^{\mathrm{bas}}(u_\nu, [-\tfrac{l_\nu}{2} + l, \tfrac{l_\nu}{2} - l] \times S^1) = 0, \tag{8.21}$$

and for a boundary node w in C,

$$\text{Area}_t^{\text{bas}}(w) := \lim_{l\to\infty}\lim_{\nu\to\infty}\text{Area}_t^{\text{bas}}(u_\nu, [-\tfrac{l_\nu}{2}+l, \tfrac{l_\nu}{2}-l]\times[0,1]) = 0. \tag{8.22}$$

REMARK 8.22. (Unstable domain components in the compactness result) In the compactness theorem, unstable domain components are ruled out because the limit is adapted with respect to a stabilizing divisor. In contexts where there is no stabilizing divisor, the proof can be modified using the technique of [**11**], wherein additional sequences of marked points are added to the domain curves to ensure that the limit domain curve is stable. The additional marked points are added at sequences of points where the derivative of the map blows up with respect to a hyperbolic metric on the domain. Our proof then entirely carries over.

STEP 5 : *Determining the tropical structure and constructing translation sequences.* We have so far determined the domain treed disk of the limit, and the limit map u_v corresponding to each vertex. We know the polytope $P(v) \in \mathcal{P}$ for each component $v \in \text{Vert}(\Gamma)$ of the limit. Further for every $v \in \text{Vert}(\Gamma)$ the map u_v is the limit of $\mathbb{e}^{-t_\nu(v)}u_\nu$ where $t_\nu(v) \in \nu P^\vee$ is a sequence of translations.

First we observe that all boundary edges are internal edges, and are collapsed by the tropicalization morphism. Disk components $v \in \text{Vert}_\circ(\Gamma)$ map to a single vertex $v_\circ$ in the tropical graph and the translation sequence $t_\nu(v_\circ)$ is the point $P_0^\vee$.

Edge slopes of the tropical graph are determined using the breaking annulus lemma. For any node w_e of C corresponding to an interior node $e = (v_+, v_-) \in \text{Edge}_{\bullet,-}(\Gamma)$, choose a complex coordinate

$$z_e^\pm : (U_\pm, w_e) \to (\mathbb{C}, 0) \tag{8.23}$$

in a neighborhood $U_\pm \subset C_{v_\pm}$ of the lift $w_e^\pm$ of the node, and let $A_{e,\nu} := A(l_\nu(e)) \subset C_\nu$ be a sequence of centered annuli converging to the neighborhoods U_+, U_- of the nodal lifts w_e^+, w_e^- (as in Definition 8.10). The dilated basic area decays to zero in the middle of A_e, see (8.21). Therefore the breaking annulus lemma (Proposition 8.17) is applicable on the sequence of annuli $u_\nu|A_{e,\nu}$, and we obtain a slope $\mathcal{T}(e) \in \mathfrak{t}_{P(e),\mathbb{Z}}$ corresponding to the edge e. The breaking annulus lemma implies that

$$\sup_\nu\{|t_\nu(v_+) - t_\nu(v_-) - \mathcal{T}(e)l_\nu(e)|\} < \infty, \tag{8.24}$$

which is the (Approximate slope) condition in Definition 8.23.

The tropicalization of Γ is determined as follows. An edge $e \in \text{Edge}_{\bullet,-}(\Gamma)$ is an internal (non-tropical) edge exactly if $\mathcal{T}(e)$. In that case, the sequence is $|t_\nu(v_+) - t_\nu(v_-)|$ is uniformly bounded for all ν. Therefore we can replace $\mathfrak{t}_\nu(v_-)$ by $t_\nu(v_+)$ (or the other way around), and the convergence of Step 2 still holds. By performing such replacements for all internal edges, we may assume $t_\nu(v_+) = t_\nu(v_-)$ for all $e = (v_+, v_-) \in \text{Edge}_{\text{int}}(\Gamma)$. Therefore t_ν descends to the tropical graph Γ_{tr} and is an approximate Γ_{tr}-translation sequence.

Next we obtain translation sequences and determine tropical vertex positions. Since $\{t_\nu(v) : v \in \text{Vert}(\Gamma_{\text{tr}})\}$ is an approximate Γ_{tr}-translation sequence, by Lemma

8.25, there is a Γ_{tr}-translation sequence $(t_\nu(v)')_{v,\nu}$ such that

$$\sup_{\nu,v} |t_\nu(v)' - t_\nu(v)| < \infty. \tag{8.25}$$

The convergence of maps in Step 2 continues to hold if t_ν is replaced by $t_\nu(v)'$. Any element of the translation sequence gives a tropical vertex position map. Indeed, since

$$t'_\nu = (t'_\nu(v))_{v\in \mathrm{Vert}(\Gamma_{\mathrm{tr}})}, \quad t'_\nu(v) \in \nu P(v)^\vee$$

satisfies the (Slope) condition (8.10),

$$\mathrm{Vert}(\Gamma_{\mathrm{tr}}) \ni v \mapsto \tfrac{t'_\nu(v)}{\nu}$$

is a tropical vertex position map.

STEP 6: *Finishing the proof of convergence.*
To finish the proof of convergence it remains to show that the collection of limit maps satisfy matching conditions at nodes, and that (Thin cylinder convergence) is satisfied at tropical nodes. We first consider disk nodes. For a disk node $w_e \in C$ corresponding to an edge $e = (v_+, v_-) \in \mathrm{Edge}_\circ(\Gamma)$ with length $\ell(e) = 0$, by (8.22) in Step 4, the dilated basic area decays to zero on the strips $A_{e,\nu} \subset C_\nu$ converging to the node w_e. Therefore the boundary version of the annulus lemma (Proposition 8.20) is applicable on $u_\nu | C_\nu$ and we conclude that node matching holds, that is, $u_{v_+}(w_e^+) = u_{v_-}(w_e^-) \in L$.

Next we determine matching coordinates for tropical nodes and prove node matching using the breaking annulus Lemma. Matching coordinates at a tropical node w_e, $e \in \mathrm{Edge}_{\mathrm{trop}}(\Gamma)$ are obtained by 'correcting' the holomorphic coordinate chosen in (8.23). Multiplying a constant to the coordinates z_e^+, z_e^- has the effect of adding a constant to the sequence of neck length parameters $l_\nu(e) + i\theta_\nu(e)$. We recall that as a consequence of (8.24) and (8.25) the Γ-translation sequence t'_ν satisfies

$$\sup_\nu \{ |t'_\nu(v_+) - t'_\nu(v_-) - \mathcal{T}(e) l_\nu(e)| \} < \infty.$$

Since t'_ν satisfies the (Slope) condition (8.10), there is a sequence $l'_\nu(e) \to \infty$ for every $e \in \mathrm{Edge}_{\bullet,-}(\Gamma)$ such that

$$t'_\nu(v_+) - t'_\nu(v_-) = \mathcal{T}(e) l'_\nu(e), \quad \text{and} \quad \sup_\nu |l_\nu(e) - l'_\nu(e)| < \infty.$$

Therefore, we can adjust the coordinates z_e^+, z_e^- by scalar multiplication so that a subsequence of neck length parameters $l_\nu(e) + i\theta_\nu(e)$ satisfies

$$\lim_\nu (t'_\nu(v_+) - t'_\nu(v_-) - \mathcal{T}(e) l_\nu(e)) = 0, \quad \lim_\nu \theta_\nu(e) = 0. \tag{8.26}$$

We will show that (z_e^+, z_e^-) are matching coordinates at the node w_e for the broken map u. For interior non-tropical nodes, we leave (z_e^+, z_e^-) from (8.23) unchanged, and the following discussion is valid. To simplify calculations, we use logarithmic coordinates in the neighborhood of the node

$$C_\pm \supset U_\pm \backslash \{w_\pm(e)\} \xrightarrow{\pm \ln z_e^\pm} \mathbb{R}_\mp \times S^1.$$

The annulus $A(l_\nu(e))$ is then identified to the limit curve by translations

$$A(l_\nu(e)) \to U_\pm, \quad (s,t) \mapsto s + it \mp \tfrac{1}{2}(l_\nu + i\theta_\nu).$$

By Step 2,

$$\mathbb{e}^{-t'_\nu(v_\pm)} u_\nu(s + it \pm \tfrac{1}{2}(l_\nu + i\theta_\nu)) \to u_{v_\pm} \quad \text{in } C^\infty_{\text{loc}}(U_\pm \backslash \{w_\pm(e)\}). \tag{8.27}$$

We apply the breaking annulus lemma on the maps u_ν on the annuli $A(l_\nu(e))$. The resulting decay estimate together with (8.27), (8.26) implies the convergence of the sequence

$$\mathbb{e}^{-\frac{1}{2}(t'_\nu(v_+)+t'_\nu(v_-))} u_\nu(0,0) \to x_0 \quad \text{in } X^\square_{\overline{P(e)}}.$$

It follows that $u_{v_\pm}$ are asymptotically close to the cylinder $(s,t) \mapsto e^{\mathcal{T}(e)(s+it)}x_0$. We conclude that z_e^+, z_e^- are matching coordinates at w_e and (Thin cylinder convergence) is satisfied. Note that we have shown node matching for all interior edges, both tropical and internal, in a unified way. In case of an internal node w_e, $\mathcal{T}(e) = 0$, and we have shown that $u_{v_\pm}(\mp\infty) = x_0 \in X^\square_{\overline{P(v_\pm)}}$.

Step 7 : *Uniqueness of the limit.*
The limit of the domain curves is unique up to reparametrization, because the limit is a stable curve. The identifications between subsets of C^ν to the limit curve are unique in the following sense (see Remark 8.8) : The neck regions in C^ν are parametrized in a unique way, and the difference between any two choices of identifications of the complement of the neck in C^ν to C converge uniformly to identity as $\nu \to \infty$. Let Γ be the combinatorial type of the limit curve C. For every vertex $v \in \text{Vert}(\Gamma)$, the polytope $P(v)$ in the limit map is uniquely determined as follows: Suppose there is a translation sequence $\{t'_\nu(v) \in P'(v)^\vee : v \in \text{Vert}(\Gamma)\}_\nu$ for which Gromov convergence holds. Then the property

$$d_{B^\vee}(t_\nu(v), P_0^\vee) \to \infty \quad \forall P_0 \in \mathcal{P}, P_0 \supset P'(v)$$

of a translation sequence implies, by Lemma 8.15, that the maps $u_{\nu,v}$ horizontally converge in $P'(v)$. However, $P'(v) = P(v)$ because the polytope of horizontal convergence is unique by Lemma 8.16.

Translation sequences are well-determined up to uniformly bounded perturbations as follows : Suppose t_ν, t'_ν are two distinct translation sequences, such that the sequence $\mathbb{e}^{-t_\nu} u_\nu$ resp. $\mathbb{e}^{-t'_\nu} u_\nu$ converges to a broken map u resp. u' . Then for all vertices $v \in \text{Vert}(\Gamma)$, there is a uniform bound

$$\sup_\nu |t_\nu(v) - t'_\nu(v)| < \infty,$$

because both the sequences $\mathbb{e}^{-t_\nu(v)} u_\nu$, $\mathbb{e}^{-t'_\nu(v)} u_\nu$ converge pointwise in C_v°. After passing to a subsequence, we may assume that there exists a limit

$$t(v) := \lim_\nu t_\nu(v) - t'_\nu(v).$$

Then, for each vertex v, $u_v = e^{t(v)}u'_v$. Since u_v and u'_v satisfy matching conditions at nodes we conclude that t is an element of $T_{\mathrm{trop},\mathcal{W}}(\Gamma)$, which is the identity component of $T_{\mathrm{trop}}(\Gamma)$. We have thus shown that the limit is unique up to the action of $T_{\mathrm{trop},\mathcal{W}}(\Gamma)$. □

Definition 8.23. (Approximate translation sequence) Suppose Γ is a tropical graph. An *approximate* Γ*-translation sequence* consists of sequences $\{t_\nu(v) \in \nu P(v)^\vee\}_\nu$ for each $v \in \mathrm{Vert}(\Gamma)$ such that

- (Approximate Slope) For any edge $e = (v_+, v_-) \in \mathrm{Edge}_-(\Gamma)$, there exists a sequence $l_\nu(e) \to \infty$ such that
$$\sup_\nu (t_\nu(v_+) - t_\nu(v_+) - \mathcal{T}(e) l_\nu) < \infty.$$

The differences appearing in the (Approximate Slope) condition will be referred to later using the following notation:

Definition 8.24. On a tropical graph Γ define the *discrepancy* function on any edge $e = (v_+, v_-) \in \mathrm{Edge}_{\bullet,-}(\Gamma)$ as

$$\mathrm{Diff}_e : \oplus_{\mathrm{Vert}(\Gamma)} \mathfrak{t}^\vee_{P(v)} \to \mathfrak{t}^\vee / \langle \mathcal{T}(e) \rangle, \quad (t_v)_{v \in \mathrm{Vert}(\Gamma)} \mapsto (t_{v_+} - t_{v_-}) \mod \mathcal{T}(e).$$

Lemma 8.25. (From approximate to exact translation sequences) *Suppose Γ is a graph with tropical structure $\mathcal{T}$ and t_ν is an approximate Γ-translation sequence. Then, after passing to a subsequence, there is a Γ-translation sequence $\bar{t}_\nu$ such that $|\bar{t}_\nu(v) - t_\nu(v)|$ is uniformly bounded for all ν, $v \in \mathrm{Vert}(\Gamma)$.*

Proof. The (Approximate slope) condition says that the sequences of discrepancies $(\mathrm{Diff}_e(t_\nu))_\nu$ are uniformly bounded. Via uniformly bounded adjustments to t_ν, we aim to make this quantity vanish for all edges.

We give an algorithm that transforms t_ν into a bounded sequence $t^k_\nu \in \oplus_{\mathrm{Vert}(\Gamma)} \mathfrak{t}^\vee_{P(v)}$, and will prove later that $t_\nu - t^k_\nu$ is an exact translation sequence. The algorithm is as follows:

Step 1: *Relativisation.*
In this step, we replace t_ν by
$$t^0_\nu(v) := t_\nu(v) - \nu \lim_\nu (t_\nu(v)/\nu) \in \mathfrak{t}^\vee_{P(v)}.$$
The limit in the right-hand side exists after passing to a subsequence because the original translation sequences t_ν lie in $\nu B^\vee$ and $B^\vee$ is compact. For any $v \in \mathrm{Vert}(\Gamma)$, the discrepancies across edges are preserved :
$$\mathrm{Diff}_e(t_\nu) = \mathrm{Diff}_e(t^0_\nu). \tag{8.28}$$

Step 2: *Subtracting fastest growing sequences.*
By a sequence of further transformations, we will change t^0_ν to a bounded sequence $t^k_\nu \in \mathfrak{t}^\vee$. At each step, the sequence t^i_ν is replaced by t^{i+1}_ν defined as follows. Choose a vertex $v_0 \in \mathrm{Vert}(\Gamma)$ for which the rate of increase of the sequence $|t^i_\nu(v_0)|$ is the maximum. That is, for all $v \in \mathrm{Vert}(\Gamma)$, $\lim_\nu |t^i_\nu(v)|/|t^i_\nu(v_0)|$ is finite. Such

a vertex can indeed be chosen, because after passing to a subsequence, the limit $\lim_\nu |t^i_\nu(v_i)|/|t^i_\nu(v_j)|$ exists in $[0,\infty]$ for any pair of vertices. Now, define

$$r^i_\nu := |t^i_\nu(v_0)|,$$

and

$$t^{i+1}_\nu(v) := t^i_\nu(v) - r^i_\nu \lim_\nu \frac{t^i_\nu(v)}{r^i_\nu} \in \mathfrak{t}^\vee_{P(v)}. \tag{8.29}$$

We stop the iteration when the sequence $t^i_\nu(v)$ corresponding to every vertex is bounded, and suppose the final sequence is t^k_ν.

The process terminates in a finite number of steps. Indeed, notice that $t^{i+1}_\nu(v_0) = 0$ for all ν. The number of vertices $v \in \operatorname{Vert}(\Gamma)$ for which $t^{i+1}_\nu(v)$ vanishes is at least one more than the number of vertices v for which $t^i_\nu(v)$ vanishes.

The iterations of the algorithm preserve the discrepancies across edges: For all tropical edges $e \in \operatorname{Edge}_{\text{trop}}(\Gamma)$

$$\operatorname{Diff}_e(t^{i+1}_\nu) = \operatorname{Diff}_e(t^i_\nu). \tag{8.30}$$

Indeed, (8.29) implies

$$\operatorname{Diff}_e(t^{i+1}_\nu) = \operatorname{Diff}_e(t^i_\nu) - r^i_\nu \lim_\nu \frac{\operatorname{Diff}_e(t^i_\nu)}{r^i_\nu},$$

and the second term in the right-hand-side vanishes because $\operatorname{Diff}_e(t^i_\nu)$ is uniformly bounded and $r^i_\nu \to \infty$ as $\nu \to \infty$.

We claim that $t_\nu - t^k_\nu$ is an exact translation sequence. For all vertices v the (Polytope) condition (8.9) $(t_\nu - t^k_\nu)(v) \in P(v)^\vee$ is satisfied because $t_\nu(v) \in P(v)^\vee$ and $t^k_\nu(v) \in \mathfrak{t}^\vee_{P(v)} \simeq TP(v)^\vee$. The (Slope) condition is satisfied because $\operatorname{Diff}_e(t_\nu) - \operatorname{Diff}_e(t^k_\nu) = 0$ by (8.28) and (8.30). □

At this point we give the proof of Lemma 5.29 that was used in the construction of stabilizing divisors for broken manifolds. The proof is by a hard rescaling argument similar to the one used in the proof of Gromov convergence for breaking maps.

Proof of Lemma 5.29. Suppose $u_\nu : \mathbb{P}^1 \to X^\nu$ is a sequence of non-constant J^ν-holomorphic spheres with area $\leq E$ that are not stabilized. That is, either the images are contained in the divisor D^ν or they have ≤ 2 distinct points of intersection with the divisor. We will show that there is a unstabilized sphere in X_P for some polytope $P \in \mathcal{P}$.

First consider the situation that the derivatives of u_ν are uniformly bounded. Then, the sequence u_ν converges horizontally in some polytope P. As in Step 3 of the proof of Theorem 8.2, there is a sequence of translations $t_\nu \in \nu P^\vee$ such that a subsequence of $\mathbb{e}^{-t_\nu} u_\nu$ uniformly converges to a $J_{\overline{P}}$-holomorphic map $u : \mathbb{P}^1 \to X^\square_{\overline{P}}$, that is unstabilized. The basic area forms $u^*_\nu \omega^{\text{bas}}_\nu$ converge to $u^* \omega^{\text{bas}}_P$, and therefore the area of $\pi_P(u)$ is positive and $\leq E$.

If the derivatives on u_ν are not uniformly bounded, we produce a sphere by hard rescaling. By following the procedure in Step 2 of the proof of Theorem 8.2, we obtain a rescaled sequence of J^ν-holomorphic maps $v_\nu : B_{r_\nu} \to X^\nu$ on balls B_{r_ν}

that exhaust $\mathbb{C}$, a polytope P, and a sequence of translations $t_\nu \in \nu P^\vee$ such that $\mathbb{e}^{-t_\nu} u_\nu$ converges in C^∞_{loc} to a non-constant map $v : \mathbb{C} \to X^{\square}_{\overline{P}}$. As in the proof of Theorem 8.2, the map v extends to $v : \mathbb{P}^1 \to X_{\overline{P}}$. The projection $\pi_P \circ v : \mathbb{P}^1 \to X_P$ is non-constant : otherwise the image of v is in a fiber $V_{P^\vee}$ which is a toric variety, and there is only one point $\infty \in \mathbb{P}^1$ that maps to toric divisors $V_{Q^\vee}, Q \supset P$ of $V_{P^\vee}$, and therefore, v is constant. Since the basic area forms $v_\nu^* \omega_\nu^{\text{bas}}$ for v_ν converge to the basic area form $v^* \omega_P^{\text{bas}}$ for v, so $\text{Area}(v) \le E$. The sphere v is not stabilized and the Lemma follows. □

8.5. Convergence for broken maps

In this section we prove Theorem 8.3 on Gromov compactness for broken maps. The limit map may have additional components because of bubbling and consequently the tropical graph of the limit map may have additional vertices. The tropical graph of the limit map is related to the tropical graph of the maps in the sequence by a 'tropical edge collapse' relation defined below. In this section we show that such bubbling happens only in families whose dimension is at least two, and so does not occur in the zero-dimensional moduli spaces we use to define the Fukaya algebra.

Definition 8.26. (Collapsing edges tropically) A *tropical edge collapse* is a morphism of tropical graphs $\Gamma' \xrightarrow{\kappa} \Gamma$ that collapses a subset of edges $\text{Edge}(\Gamma') \setminus \text{Edge}(\Gamma)$ in Γ' inducing a surjective map on the vertex sets

$$\kappa : \text{Vert}(\Gamma') \to \text{Vert}(\Gamma),$$

and satisfies the following conditions:

(a) for any vertex $v \in \text{Vert}(\Gamma')$, $P(v) \subseteq P(\kappa(v))$; and
(b) the edge slope is unchanged for uncollapsed edges, i.e. if $\mathcal{T}$, $\mathcal{T}'$ are the edge slope functions for Γ, Γ', then $\mathcal{T}(\kappa(e)) = \mathcal{T}'(e)$ for any uncollapsed edge $e \in \text{Edge}(\Gamma')$.

Since the edge slope function $\mathcal{T}'$ extends $\mathcal{T}$, we often use the same notation for both. A tropical edge collapse $\Gamma' \xrightarrow{\kappa} \Gamma$ is *trivial* if no edge is collapsed, and $P_{\Gamma'}(v) = P_\Gamma(\kappa(v))$ for all v.

Let Γ_1, Γ_2 be combinatorial types of treed disks that are equipped with a tropical structure given by tropicalization maps

$$\text{tr}_1 : \Gamma_1 \to \Gamma_{1,\text{tr}}, \quad \text{tr}_2 : \Gamma_2 \to \Gamma_{2,\text{tr}}.$$

An edge collapse morphism $\Gamma_1 \to \Gamma_2$ is a tropical edge collapse if it is a lift of a tropical edge collapse map $\Gamma_{1,\text{tr}} \to \Gamma_{2,\text{tr}}$ between the tropical graphs. This finishes the definition.

Example 8.27. In Figure 4.4, collapsing the middle edge in Γ_2 gives a tropical edge collapse morphism $\kappa : \Gamma_2 \to \Gamma_1$. See Figure 8.1 for another example of a tropical edge collapse morphism.

We next define relative vertex position maps for a tropical edge collapse morphism. Analogous to the case of Gromov convergence for breaking maps, relative vertex positions for the limit of a sequence of broken maps will be read off from relative translation sequences in the convergence. A relative translation sequence for a converging sequence of broken maps 'goes to infinity' (see Definition 8.33), and relative vertex positions of the limit map are obtained by scaling any relative translation is the sequence. Therefore we think of a relative translation as an infinitesimal version of a relative vertex position, which we make precise below in Definition 8.28 and Remark 8.30.

Definition 8.28. (Relative vertex positions, relative translations) Suppose $\kappa : \Gamma' \to \Gamma$ is a tropical edge-collapse morphism.

(a) (Relative vertex position) A *relative vertex position* $\mathcal{T}_{\Gamma',\Gamma}$ is the difference between a vertex position of Γ' and a vertex position of Γ, and the space of relative vertex positions is

$$\mathcal{W}(\Gamma',\Gamma) := \{(\mathcal{T}'(v) - \mathcal{T}(\kappa v) \in \mathrm{Cone}(\kappa, v) \subset \mathfrak{t}^\vee)_{v\in \mathrm{Vert}(\Gamma')} : \\ \mathcal{T}' \in \mathcal{W}(\Gamma'), \mathcal{T} \in \mathcal{W}(\Gamma)\},$$

where

$$\mathrm{Cone}(\kappa, v) := \mathrm{Cone}_{P(\kappa(v))^\vee}(P(v)^\vee) \\ := \{\alpha(t - t_0) \in \mathfrak{t}^\vee : t \in P(v)^\vee, t_0 \in P(\kappa v)^\vee, \alpha \in \mathbb{R}_{\geq 0}\} \tag{8.31}$$

is the cone in the polytope $P(v)^\vee$ based at points in $P(\kappa(v))^\vee$ (same as the Definition in (3.34)).

(b) (Relative translation) A (Γ',Γ)-*translation* or a *relative translation* is an element in the $\mathbb{R}_+$-span of the space of (Γ',Γ)-vertex positions. That is, a relative translation is the sum of $\mathbb{R}_+$-multiples of relative vertex positions. The space of relative translations is denoted by

$$w(\Gamma',\Gamma) := \mathbb{R}_+\langle \mathcal{W}(\Gamma',\Gamma)\rangle.$$

Example 8.29. For the tropical edge collapse $\kappa : \Gamma' \to \Gamma$ in both Figure 4.4 and Figure 8.1, the space $\mathcal{W}(\Gamma',\Gamma)$ of relative vertex positions is (linearly) isomorphic to the interval $[0, 1)$, where 0 corresponds to the vertex position where the edges collapsed by κ have length 0, and 1 corresponds to some vertex v lying in a lower dimensional polytope $Q^\vee \subset P^\vee(v)$. Note that the latter is not a legitimate vertex position. In both cases the space $w(\Gamma',\Gamma)$ of relative translations is the cone $\mathbb{R}_{\geq 0} \subset \mathbb{R}$.

Remark 8.30. A relative translation can alternately be defined as an infinitesimal relative vertex position, that is for any vertex position $\mathcal{T}_\Gamma$ of Γ, and a (Γ',Γ)-translation $\mathcal{T}_{(\Gamma',\Gamma)}$, there exists $t_0 > 0$ such that $\mathcal{T}_\Gamma + t\mathcal{T}_{(\Gamma',\Gamma)}$ is a Γ'-vertex position for all $t \in [0, t_0)$. This is a consequence of the convexity of the space $\mathcal{W}(\Gamma')$ of vertex positions of Γ', and since $\mathcal{W}(\Gamma) \subset \mathcal{W}(\Gamma')$.

Relative translations can also be characterized in terms of slope relations as in the following Lemma, whose proof is left to the reader.

LEMMA 8.31. (Slope condition on relative translations) *Suppose* $\kappa : \Gamma' \to \Gamma$ *is a tropical edge-collapse morphism. A tuple* $(\mathcal{T}_{\Gamma',\Gamma}(v) \in \operatorname{Cone}(\kappa, v))_{v \in \operatorname{Vert}(\Gamma')}$ *is a* (Γ', Γ)*-translation if and only if it satisfies*

$$\mathcal{T}_{\Gamma',\Gamma}(v_+) - \mathcal{T}_{\Gamma',\Gamma}(v_-) \in \begin{cases} \mathbb{R}_{\geq 0}\mathcal{T}(e), & e \notin \operatorname{Edge}(\Gamma), \\ \mathbb{R}\mathcal{T}(e), & e \in \operatorname{Edge}(\Gamma) \end{cases} \tag{8.32}$$

for all edges $e = (v_+, v_-)$ *in* Γ'.

Relative translations (or relative vertex maps) are ways of moving vertices in the tropical graph without changing the edge slope. Therefore they are generators of the tropical symmetry group.

LEMMA 8.32. (Relative translations generate tropical symmetry) *Suppose there is a non-zero relative translation* $\mathcal{T}$ *for the tropical edge collapse* $\Gamma' \to \Gamma$. *Then, a broken map modelled on* Γ' *has non-trivial tropical symmetry group* $T_{\text{trop}}(\Gamma')$ *(Definition 4.33).*

PROOF. Given a non-zero relative translation $(\mathcal{T}(v))_{v \in \operatorname{Vert}(\Gamma'}$,

$$\mathbb{C} \mapsto T_{\text{trop}}(\Gamma'), \quad z \mapsto (e^{z\mathcal{T}(v)})_{v \in \operatorname{Vert}(\Gamma')}, \tag{8.33}$$

is a non-trivial one-parameter subgroup in $T_{\text{trop}}(\Gamma')$. Here we view $\mathcal{T}(v)$ as lying in $\mathfrak{t}$. The slope relations for the relative translation imply corresponding slope relations on $(e^{z\mathcal{T}(v)})_v$, see Lemma 4.35. □

Components of a relative translation correspond to maps of broken manifolds that rescale the coordinates on cylindrical ends. For a pair $Q \subseteq P$ of polytopes, we recall that (3.38) gives an embedding

$$i_Q^{\overline{P}} : U_Q(X_{\overline{P}}^{\square}) \to X_{\overline{Q}}^{\square}. \tag{8.34}$$

where $U_Q(X_{\overline{P}}^{\square}) \subset X_{\overline{P}}^{\square}$ is the Q-cylindrical end of $X_{\overline{P}}^{\square}$. If $Q = P$, $U_P(X_{\overline{P}}^{\square}) = X_{\overline{P}}^{\square}$ and the map in (8.34) is the identity map. A relative translation $t \in \operatorname{Cone}_{P^\vee} Q^\vee$ gives an embedding

$$e^{-t} : U_Q(X_{\overline{P}}^{\square}) \to X_{\overline{Q}}^{\square} \tag{8.35}$$

defined as $i_Q^{\overline{P}}$ in (8.34) composed with a translation by $-t$ in the $\mathfrak{t}_Q$-coordinate in $X_{\overline{Q}}^{\square} \simeq Z_Q^{\square} \times \mathfrak{t}_Q^\vee$.

Relative translations appearing in the convergence of broken maps 'go to infinity' in the following sense:

DEFINITION 8.33. (Relative translation sequence going to infinity) Suppose $\kappa : \Gamma' \to \Gamma$ is a tropical edge collapse. A (Γ', Γ)-translation sequence

$$t_\nu(v) \in \operatorname{Cone}(\kappa, v) = \operatorname{Cone}_{P(\kappa v)^\vee}(P(v)^\vee), \quad v \in \operatorname{Vert}(\Gamma')$$

(where $\operatorname{Cone}(\kappa, v)$ is as defined in (8.31)) *goes to infinity* if

(a) (Polytope) For any vertex $v \in \operatorname{Vert}(\Gamma')$ and a polytope $Q \in \mathcal{P}$ such that $P(v) \subset Q \subseteq P(\kappa(v))$,

$$d(t_\nu(v), \operatorname{Cone}_{P(\kappa v)^\vee}(Q^\vee)) \to \infty.$$

(b) (Slope for collapsed edge) For any edge $e \in \mathrm{Edge}(\Gamma_{\kappa^{-1}(v)})$ connecting vertices v_+, v_-, the sequence l_ν defined as

$$t_\nu(v_+) - t_\nu(v_-) = \mathcal{T}(e) l_\nu$$

goes to infinity, that is, $l_\nu \to \infty$. (Note that the existence of l_ν follows from the Slope condition on relative translations from Lemma 8.31.)

Definition 8.34. (Gromov convergence for broken maps) Suppose Γ', Γ are combinatorial types of stable treed disks that are equipped with a tropical structure. Suppose $\Gamma' \xrightarrow{\kappa} \Gamma$ is a tropical edge collapse morphism which induces a vertex map $\kappa : \mathrm{Vert}(\Gamma') \to \mathrm{Vert}(\Gamma)$. A sequence of broken maps $u_\nu : C_\nu \to \mathfrak{X}_{\mathcal{P}}$ of type Γ converges to a limit broken map $u : C \to \mathfrak{X}_{\mathcal{P}}$ of type Γ' if the following conditions are satisfied.

(a) (Convergence of domains) The sequence of treed disks C_ν converge to the treed disk C and for any tropical node w_e, $e \in \mathrm{Edge}_{\mathrm{trop}}(\Gamma')$ that is collapsed by κ, the arguments $\frac{\delta_e(C_\nu)}{|\delta_e(C_\nu)|}$ of the gluing parameters converge to a limit. Let $S_\nu(v) \subset (S_\nu)_{\kappa(v)} \subset C_\nu$ be the subset corresponding to a vertex $v \in \mathrm{Vert}(\Gamma')$, and let

$$i_{v,\nu} := i_{S_v, S_{\nu,\kappa(v)}} : S_\nu(v) \to S_v, \quad S_\nu(v) \subset S_\nu,$$

be embeddings from (8.8) whose images $i_{v,\nu}(S_\nu(v))$ exhaust S_v° as $\nu \to \infty$. Here

$$S_v^\circ := S_v \backslash \{w_e : v \in e, \;\; e \in \mathrm{Edge}\_(\Gamma') \text{ is collapsed by } \kappa\}.$$

(b) (Convergence of maps) There is a (Γ', Γ)-translation sequence $(t_\nu(v))_{v \in \mathrm{Vert}(\Gamma')}$ going to infinity in the sense of Definition 8.33, such that for any vertex $v \in \mathrm{Vert}(\Gamma')$, the sequence of maps

$$S_v^\circ \supset i_{v,\nu}(S_\nu(v)) \xrightarrow{e^{-t_\nu(v)}(u_\nu \circ i_{v,\nu}^{-1})} X^{\square}_{\overline{P}(v)}$$

converges in $C^\infty_{\mathrm{loc}}(S_v^\circ)$ to $u_v : S_v^\circ \to X^{\square}_{\overline{P}(v)}$. The map $e^{-t_\nu(v)} : U_{P(v)}(X_{\overline{P}(\kappa v)}) \to X^{\square}_{\overline{P}(v)}$ is defined on the $P(v)$-cylindrical end $U_{P(v)}(X^{\square}_{\overline{P}(\kappa v)}) \subset X^{\square}_{\overline{P}(\kappa v)}$, see (8.35). For each boundary edge $e \in \mathrm{Edge}_\circ(\Gamma)$, the maps $u_\nu | T_{e,\nu}$ on the treed segment converge to a (possibly broken) treed segment in u.

(c) (Thin cylinder convergence) holds for all tropical nodes w_e corresponding to edges $e \in \mathrm{Edge}_{\mathrm{trop}}(\Gamma')$ that are collapsed by κ (that is e is not an edge in Γ).

Remark 8.35. An informal description of the above definition is that for each vertex $v \in \mathrm{Vert}(\Gamma)$ the sequence of smooth maps u_v with markings corresponding to nodal lifts $w_e \in C_v$ converges to a limit map $u|(\cup_{v' \in \kappa^{-1}(v)} C_{v'})$.

Proposition 8.36. (Finite number of types of broken maps) *For any $E > 0$, $d(\circ) \geq 1$ there are a finite number of tropical graphs Γ that are types of broken maps $u : C \to \mathfrak{X}$ of area at most E and $d(\circ)$ boundary leaves.*

PROOF. Consider a broken map $u : C \to X$ of type Γ and area at most E_0, whose edge slopes are given by $\mathcal{T} : \mathrm{Edge}(\Gamma) \to \mathfrak{t}_{\mathbb{Z}}$.

STEP 1: *Uniform bound on the number of vertices.*
The number of interior markings in Γ is bounded by kE_0, where k is the degree of the stabilizing divisor (as in Definition 5.2). Since u is an adapted map, all the surface components $S_v, v \in \mathrm{Vert}(\Gamma)$ in its domain C are stable. Given a uniform bound on the number of interior and boundary markings on the type of domain curve, we obtain a uniform bound on the number irreducible surface components.

STEP 2: *Uniform bound on the sum of vertical components of edge slopes.*
Let $u_v : C \to X_{\overline{P}}$ be a component of the broken map corresponding to a vertex $v \in \mathrm{Vert}(\Gamma)$, and $P := P(v)$. By the balancing property (4.25) the sum of the edge slopes projected to $\mathfrak{t}_P^\vee$ is

$$\sum_{e \ni v} \pi_{\mathfrak{t}_P^\vee}(\mathcal{T}(e)) = c_1((\pi_P \circ u_v)^* Z_{P(v)} \to X_P))$$

The right-hand side, which is the pairing of $(\pi_P \circ u)_*[C]$ and the Chern class $c_1(Z_P \to X_P)$ has an E_0-dependent bound. Indeed, for any $\epsilon \in (0,1)$, there is a constant k_0 such that for any domain-dependent almost complex structure $J \in B_\epsilon(J_{X_P})_{C^0}$ and a J-holomorphic sphere $u : \mathbb{P}^1 \to X_P$,

$$\int_{\mathbb{P}^1} u^* c_1(Z_P \to X_P) \leq k_0 \int_{\mathbb{P}^1} u^* \omega_{X_P} \leq k_0 E_0.$$

This estimate is similar to the one in Lemma 5.5 and the proof is the same – by choosing a two-form on X_P representing the Chern class and bounding it pointwise by ω_{X_P}.

STEP 3: *Uniform bound on the horizontal components of edge slopes.*
For a vertex v and an incident edge e, the horizontal component of the slope $\mathcal{T}(e)$ is the sum of intersection multiplicities at the node w_e with horizontal relative divisors of X_P:

$$\pi^\perp_{\mathfrak{t}_P^\vee}(\mathcal{T}(e)) = \textstyle\sum_{Q \subset P} m_{w_e}(u_{v,P}, X_Q)\nu_Q, \quad u_{v,P} := \pi_P \circ u_v : S_v \to X_P,$$

where the sum ranges over facets $Q \in \mathcal{P}$ of P, and ν_Q is the normal vector of the facet $Q \subset P$. For any relative divisor $X_Q \subset X_P$, the sum $\sum_{e \ni v} m_{w_e}(u_v, X_Q)$ is bounded by $c\omega_{X_P}(u_{v,P})$ for a uniform constant $c(X_P, X_Q)$. The proof is similar to the vertical case, by expressing the intersection number with any divisor as an integral of a two-form.

STEP 4: *Finishing the proof.*
We will show that the tropical edge slopes $\mathcal{T}(e)$ of Γ are uniformly bounded in $\mathfrak{t}^\vee$ for the set of all broken maps with area at most E_0. So far we have shown that (Step 3) the horizontal components of the edge slopes are uniformly bounded, and the vector sum of the vertical slope components of the edges incident on any vertex are uniformly bounded. From here, we conclude that for a vertex v and an edge e_0

incident on v:

$$\exists c(E) : |\mathcal{T}(e_0)| \leq \sum_{e \ni v, e \neq e_0} |\mathcal{T}(e)| + c(E) \tag{8.36}$$

Recall that Γ is a tree, any edge $e \in \operatorname{Edge}_\bullet(\Gamma)$ is oriented so that it points away from the root vertex. The slope of any incoming edge can be bounded by the slope of outgoing edges by (8.36). Applying (8.36) iteratively, we conclude that for any edge e in Γ

$$|\mathcal{T}(e)| \leq c(E)|\operatorname{Vert}(\Gamma)|.$$

where the constant $c(E)$ is the same as the one in (8.36). The Proposition now follows from the bound on the number of vertices in Step 1. □

The following definition is used in the proof of compactness for broken maps.

Definition 8.37. (Approximate (Γ', Γ)-translation sequence) Suppose the tropical graph Γ is obtained by collapsing edges in Γ' and the induced map on the vertex set is $\kappa : \operatorname{Vert}(\Gamma') \to \operatorname{Vert}(\Gamma)$. Then, an *approximate* (Γ', Γ)-*translation sequence* consists of a sequence

$$t_\nu(v) \in \operatorname{Cone}(\kappa, v),$$

for every $v \in \operatorname{Vert}(\Gamma')$ satisfying the (Polytope) and (Slope for collapsed edges) conditions in Definition 8.33 and the following weakened version of the (Slope for uncollapsed edges) condition.

- (Approximate slope for uncollapsed edges) For an edge e of Γ' that is not collapsed in Γ,
$$t_\nu(v_+) - t_\nu(v_-) \mod \mathcal{T}(e)$$
is a bounded sequence in $\mathfrak{t}^\vee/\mathbb{R}\mathcal{T}(e)$. (Recall that $\operatorname{Cone}(\kappa, v_\pm)$ as embedded in $\mathfrak{t}^\vee$ with vertex $\mathcal{T}(v_\pm)$ mapped to the origin in $\mathfrak{t}^\vee$, and identify $\mathfrak{t}$ to $\mathfrak{t}^\vee$ via a pairing (3.6).)

An approximate (Γ', Γ)-translation sequence can be adjusted by a uniformly bounded amount to produce an actual (Γ', Γ)-translation sequence.

Lemma 8.38. (From an approximate to an exact (Γ', Γ)-translation sequence) *Let $\kappa : \Gamma' \to \Gamma$ be a tropical edge collapse, and let $\{t_\nu\}_\nu$ be an approximate (Γ', Γ)-translation sequence. There exists a (Γ', Γ)-translation sequence $\{\bar{t}_\nu\}_\nu$ such that*

$$\sup_\nu |\bar{t}_\nu(v) - t_\nu(v)| < \infty$$

for all $v \in \operatorname{Vert}(\Gamma')$.

Proof. The proof is by replicating the iteration in Step 2 of the proof of Lemma 8.25. At the start, we set $t_\nu^0 := t_\nu$. At the $(i+1)$-th step, we construct t_ν^{i+1} as follows. As in the proof of Lemma 8.25, there exists a vertex $v_0 \in \operatorname{Vert}(\Gamma')$ such that the sequence $|t_\nu^i(v)|$ has the fastest growth rate. That is, for all $v \in \operatorname{Vert}(\Gamma')$, $\lim_\nu |t_\nu^i(v)|/|t_\nu^i(v_0)|$ is finite. Define

$$t_\nu^{i+1}(v) := t_\nu^i(v) - |t_\nu^i(v_0)| \lim_\nu \frac{t_\nu^i(v)}{|t_\nu^i(v_0)|}.$$

For the sequences $\{t_\nu^{i+1}(v)\}_\nu$, $v \in \mathrm{Vert}(\Gamma')$, the quantity

$$\pi^\perp_{\mathcal{T}(e)}(t_\nu^{i+1}(v_+) - t_\nu^{i+1}(v_-)), \quad e = (v_+, v_-) \in \mathrm{Edge}(\Gamma')$$

vanishes for collapsed edges, and is uniformly bounded for uncollapsed edges. Further for any vertex v, $t_\nu^{i+1}(v) \in \mathfrak{t}^\vee_{P(v)}$. After, say, k steps, the sequence $|t_\nu^k(v)|$ is uniformly bounded for all vertices v. Then, $\bar{t} := t - t^k$ is an exact (Γ', Γ)-translation sequence. □

The notion of horizontal convergence extends in a natural way to broken manifolds, though it is easier to state in this case.

Definition 8.39. (Horizontal convergence in a broken manifold) Let $P \in \mathcal{P}$ be a polytope. A sequence of points $x_\nu \in X_{\overline{P}}$ horizontally converges in $Q \subseteq P$ if the sequence $\pi_P(x_\nu) \in X_P$ converges to a point $x \in X_Q \subseteq X_P$, and x is not contained in a submanifold $X_{Q'}$ for any $Q' \subset Q$.

Analogs of Lemma 8.15 and 8.16 hold for horizontal convergence in broken manifolds.

Proof of Theorem 8.3. After passing to a subsequence, the tropical graph Γ underlying the maps u_ν is ν-independent. Indeed by Proposition 8.36, there is a finite number of tropical graphs that underlie broken maps with area $< E$. By Proposition 7.47 the Hofer energy for a map $u : (C, \partial C) \to (\mathfrak{X}, L)$ can be read off from its area and its intersection data with relative divisors, which is provided by the slopes $\mathcal{T}(e)$ of the edges of the tropical graph Γ. Consequently the Hofer energy $E_{\mathrm{Hof}}(u_\nu)$ of the maps u_ν is uniformly bounded.

We first find the limit map at each of the vertices of the tropical graph Γ. For each vertex $v \in \mathrm{Vert}(\Gamma)$, we assume that the domain curve $C_{\nu,v}$ for $u_{\nu,v}$ has marked points corresponding to lifts of nodes $e \in \mathrm{Edge}_-(\Gamma)$, $v \in e$, in addition to the marked points corresponding to leaves $e \in \mathrm{Edge}_\to(\Gamma)$. We apply the proof of the convergence for breaking maps to the sequence of maps $u_{\nu,v}$. The conclusion is that there is a limit map u_v modelled on a tropical graph Γ_v, and the convergence is via a translation sequence $t_\nu(v')$, $v' \in \Gamma_v$. By connecting the graphs $\{\Gamma_v\}_{v \in \mathrm{Vert}(\Gamma)}$ using the edges of Γ, we obtain a tropical graph Γ'. Indeed Γ' possesses a tropical vertex position map, defined by pulling back the tropical vertex position map on Γ by κ.

Next we prove that the limit map u satisfies matching conditions on the edges of Γ. The matching conditions are proved by ensuring that the relative translation sequences corresponding to each vertex of Γ, when put together, form a (Γ', Γ)-translation sequence. As a first step we show that the translation sequence t_ν in the previous paragraph is an approximate (Γ', Γ)-translation sequence. In particular, we need to prove that the (Approximate slope for uncollapsed edges) is satisfied. Consider an edge e of Γ, that is incident on v_+, $v_- \in \mathrm{Vert}(\Gamma')$, and the nodal point corresponding to e is the pair $w^e_\pm$ on $C_{v_\pm}$. The edge matching condition for broken maps (4.21) implies

$$\pi^\perp_{\mathcal{T}(e)}(u_\nu(w^e_{\nu,+})) = \pi^\perp_{\mathcal{T}(e)}(u_\nu(w^e_{\nu,-})). \tag{8.37}$$

On a sequence of converging relative maps, the projected tropical evaluations of the relative marked points converge, i.e.

$$\pi^{\perp}_{\mathcal{T}(e)}(e^{-t_\nu(v)}u_\nu(z_{i,\nu})) \to \pi^{\perp}_{\mathcal{T}(e)}(u(z)) \quad \text{in } X^{\square}_{\overline{P(e)}}/T_{\mathcal{T}(e),\mathbb{C}}.$$

This convergence is a consequence of the convergence of the marked points $z_{i,\nu}$ on the domain curve and the convergence of maps $e^{-t_\nu(v)}u_\nu$. Therefore, for the edge $e = (v_+, v_-)$,

$$d(\pi^{\perp}_{\mathcal{T}(e)}(e^{-t_\nu(v_+)}u_\nu(w^e_{\nu,+}), \pi^{\perp}_{\mathcal{T}(e)}(e^{-t_\nu(v_-)}u_\nu(w^e_{\nu,-})) \to d(\pi^{\perp}_{\mathcal{T}(e)}(u(w^e_+)), \pi^{\perp}_{\mathcal{T}(e)}(u(w^e_-))),$$

where d is the cylindrical metric distance $X^{\square}_{\overline{P(e)}}/T_{\mathcal{T}(e),\mathbb{C}}$. Using (8.37) the sequence

$$\pi^{\perp}_{\mathcal{T}(e)}(t_\nu(v_+) - t_\nu(v_-)) \in \mathfrak{t}_{\mathbb{C}}/\mathfrak{t}_{\mathcal{T}(e),\mathbb{C}} \tag{8.38}$$

is bounded, implying that $\{t_\nu(v) : v \in \text{Vert}(\Gamma')\}$ is an approximate (Γ', Γ)-translation sequence. An approximate (Γ', Γ)-translation sequence can be adjusted by a uniformly bounded amount to produce an actual (Γ', Γ)-translation sequence by Lemma 8.38. After making such an adjustment the quantity in (8.38) vanishes. Then by (8.37) the limit map u satisfies the edge matching condition (4.21) for edges in Γ, and u is therefore a broken map.

The proof of uniqueness is similar to the case of breaking maps: The domain curve is uniquely determined, and any two (Γ', Γ)-translation sequences t_ν, t'_ν differ by a uniformly bounded amount : $\sup_\nu |t_\nu(v) - t'_\nu(v)| < \infty$. After passing to a subsequence, the limit

$$t_\infty(v) := \lim_\nu (t_\nu(v) - t'_\nu(v))$$

exists. The tuple $(t_\infty(v))_v$ is an unsigned version of a (Γ', Γ)-translation since $t_\infty(v) \in \text{Cone}(\kappa, v)$ and for any edge $e = (v_+, v_-) \in \text{Edge}_{\text{trop}}(\Gamma')$,

$$t_\infty(v_+) - t_\infty(v_-) \in \mathbb{R}\mathcal{T}(e).$$

Therefore we obtain a tropical symmetry transformation $e^{-t_\infty} \in T_{\text{trop}}(\Gamma')$ that relates the limit maps $\lim_\nu e^{-t_\nu}u_\nu$, $\lim_\nu e^{-t'_\nu}u_\nu$. Therefore the limit is unique up to the action of the identity component $T_{\text{trop},\mathcal{W}}(\Gamma')$ of the tropical symmetry group.

It remains to prove the last statement of the Theorem. Suppose the tropical edge collapse $\kappa : \Gamma' \to \Gamma$ is non-trivial. Then either a tropical edge e_0 of Γ' is collapsed by κ, or there is a vertex v_0 of Γ' for which $P(v_0) \subsetneq P(\kappa(v_0))$. In both cases, for large enough ν the translation $(t_\nu(v))_{v\in\text{Vert}(\Gamma')}$ (viewed as a relative vertex position map) generates a subgroup T_{t_ν} of $T_{\text{trop}}(\Gamma')$ that is not contained in $T_{\text{trop}}(\Gamma)$. Indeed, in the first case of a collapsed tropical edge $e_0 = (v_+, v_-)$

$$t_\nu(v_+) - t_\nu(v_-) \in \mathbb{R}_+\mathcal{T}(e_0)$$

and in the second case where $P(v_0) \subsetneq P(\kappa(v_0))$ we have

$$t_\nu(v_0) \notin P(\kappa(v_0))^\vee.$$

Therefore $\dim_{\mathbb{C}}(T_{\text{trop}}(\Gamma')) > \dim_{\mathbb{C}}(T_{\text{trop}}(\Gamma))$. □

8.6. Boundaries of rigid strata

In the moduli space of maps, a strata is 'rigid' if it does not deform to another type; that is, it does not occur in the boundary of a compactification of a different stratum. To count isolated maps under generic perturbations one must restrict attention to rigid strata. In this section we formally define rigid strata of broken and unbroken maps. We will show that the codimension one boundary of rigid strata consists of maps that have a boundary edge that is either broken or has length zero. Later in Section 9.8, we will see that the strata containing a zero length edge is in fact a 'fake boundary strata' as it is the boundary of two different rigid strata with opposite induced boundary orientation. Thus the 'true boundary strata' are those that have a broken boundary edge.

Definition 8.40. (Rigid types of unbroken maps) The combinatorial type Γ of a (unbroken) treed holomorphic map is *rigid* if the only internal edges $e \in \mathrm{Edge}_{\mathrm{int}}(\Gamma)$ are boundary edges $e \in \mathrm{Edge}_{\circ}(\Gamma)$ of finite non-zero length $\ell(e) \in (0, \infty)$, and intersection multiplicity $m_{w(e)}(u, D_P)$ of u of type Γ at $e \cap S$ with the stabilizing divisor $\mathfrak{D} = (D_P, P \in \mathcal{P})$ is 1.

Definition 8.41. (Rigid types of broken maps) The combinatorial type Γ of a broken treed holomorphic map is *rigid* if the tropical graph $\mathcal{T}(\Gamma)$ is rigid, and the only internal edges $e \in \mathrm{Edge}_{\mathrm{int}}(\Gamma)$ are boundary edges $e \in \mathrm{Edge}_{\circ}(\Gamma)$ of finite non-zero length $\ell(e) \in (0, \infty)$, and intersection multiplicity $m_{w(e)}(u, D_P)$ of u of type Γ at $e \cap S$ with the stabilizing divisor $\mathfrak{D} = (D_P, P \in \mathcal{P})$ is 1.

Remark 8.42. The rigidity for types of broken maps includes the additional condition of the rigidity of the tropical graph. We chose this definition because any non-rigid tropical graph Γ is related to a rigid tropical graph Γ_0 by a tropical edge collapse morphism, and therefore tropical symmetry orbits of broken maps modelled on Γ occur in the compactification of the moduli space of maps modelled on Γ_0. Another justification for this definition is that moduli spaces of broken maps with a non-rigid tropical graph have a positive (even) dimensional tropical symmetry group, and therefore these types of broken maps do not contribute to counts of isolated curves.

Next, we address the question of which strata of broken maps occur in the compactification of a one-dimensional moduli space $\mathcal{M}_{\Gamma}^{\mathrm{brok}}(\underline{x})$ where Γ is a rigid type. The following result proves that the compactification does not contain broken maps of crowded type. Once this is established, we prove in Proposition 8.44 that the only strata that occur in the codimension one boundary are those that contain either a broken edge, or an edge with length zero.

Proposition 8.43. (No crowded strata in the compactification) *Let $\underline{\mathfrak{p}}$ be a regular perturbation datum. Let Γ be a type of an uncrowded rigid broken map and let $\underline{x} \in (\mathcal{I}(L))^{d(\circ)+1}$ be a tuple of limit points of boundary leaves such that the expected dimension of the moduli space $\mathcal{M}_{\Gamma}(L, \underline{\mathfrak{p}}, \underline{x})$ is at most 1. Then, for any curve u in the compactification $\overline{\mathcal{M}}_{\Gamma}(L, \underline{\mathfrak{p}}, \underline{x})$, there are no horizontally constant components that contain interior markings.*

Proof. First we point out that a horizontally constant component containing an interior marking can not be a disk component. Indeed the Lagrangian L is contained in a top-dimensional cut space $X^{\square}_{P_0}$ and the torus $T_{P_0,\mathbb{C}}$ is trivial. Therefore, a horizontally constant component $u_v : S_v \to X^{\square}_{P_0}$ is in fact constant and maps to a point in L. Since L is disjoint from the stabilizing divisor there are no interior markings on S_v.

Next we rule out horizontally constant spherical components that contain markings. Suppose the map $u \in \mathcal{M}_\Gamma(L, \underline{\mathfrak{p}}, \underline{x})$ is horizontally constant on a component S_v, for a vertex v of Γ, and suppose S_v has an interior marking z_e. The restriction $u|S_v$ cannot have more than two nodes – otherwise moving the marking z_e on the component gives a two-dimensional family of adapted regular maps. By stability there are exactly two nodes on S_v. By rigidity of the type Γ all edges e corresponding to interior nodes w_e are tropical edges. Since $u|_{S_v}$ is horizontally constant, the balancing property (4.25) implies that the edges $e_1, e_2 \in \operatorname{Edge}(\Gamma)$ corresponding to both nodes have the same slope $\mathcal{T}(e_1) = \mathcal{T}(e_2)$. This means the type Γ has a non-trivial tropical symmetry group $T_{\mathrm{trop}}(\Gamma)$, contradicting the rigidity of Γ.

Next, we consider a map in a boundary of the moduli space. Suppose u_ν is a sequence in $\mathcal{M}_\Gamma(L, \underline{\mathfrak{p}}, \underline{x})$ that converges to a limit $u : C \to \mathfrak{X}$ of type Γ'. First we prove the proposition assuming that the limit u is uncrowded. Then u is an adapted regular map. Since the index of u_ν is at most 1, standard arguments (see the proof of Proposition 8.44, for example) imply that the type Γ of u_ν is obtained from the limit type Γ' by the morphism (Making an edge length finite/non-zero). In particular the tropical graph associated to Γ' is the same as that of Γ. By the argument in the previous paragraph, there are no markings in horizontally constant components of the limit map u.

Next, consider the case that the limit u has a crowded component. Forgetting all but one leaf e meeting each of the crowded components $S_v, v \in \operatorname{Vert}^{\mathrm{crowded}}(\Gamma)$ yields an adapted map of u' of type Γ_s. If a crowded component $S_v \subset S$ becomes unstable after forgetting all but one of its leaves $e, e \cap S_v \neq \emptyset$ it is collapsed, and in Γ_s the remaining leaf e is assigned a multiplicity of $\mu(e)$ plus the number of forgotten leaves $e' \neq e$. The limit u' is $\mathfrak{p}_{\Gamma_s}$-adapted because of the (Locality axiom). Indeed, forgetting markings changes the type of the limit curve, but it does not affect the perturbation datum $\mathfrak{p}_\Gamma$ on the other curve components on which the map is horizontally non-constant. Therefore, u' is regular. If no component is collapsed in $\Gamma' \to \Gamma_s$, then, the expected dimension of the type Γ_s is that same as that of Γ'. In this case u is an uncrowded map with a marking in a horizontally constant component. This possibility has been ruled out in the last paragraph. If a component is collapsed, the expected dimension of Γ_s is at least two lower than Γ and therefore, the map u does not exist. $\square$

Proposition 8.44. (Boundary strata) *Suppose Γ is a rigid type for a broken map and $\underline{x} \in (\mathcal{I}(L))^{d(\circ)}$ is a tuple of limit points of boundary leaves such that the expected dimension $i(\Gamma, \underline{x})$ is ≤ 1. The moduli space $\mathcal{M}_\Gamma(\mathfrak{X}_\mathcal{P}, L, \underline{x})$ admits a compactification $\overline{\mathcal{M}}_\Gamma(\mathfrak{X}_\mathcal{P}, L, \underline{x})$. The boundary strata $\overline{\mathcal{M}}_\Gamma(\mathfrak{X}_\mathcal{P}, L, \underline{x}) \backslash \mathcal{M}_\Gamma(\mathfrak{X}_\mathcal{P}, L, \underline{x})$ correspond to* *types* Γ' *with a single broken boundary trajectory* $e \in \operatorname{Edge}_\circ(\Gamma'), \ell(e) = \infty$ *or a single boundary edge with length zero* $\ell(e) = 0$.

PROOF. For a sequence of maps $u_\nu : C \to \mathfrak{X}$ in the moduli space $\mathcal{M}_\Gamma(\mathfrak{X}_{\mathcal{P}}, L)$, a subsequence converges to a limit $u : C \to \mathfrak{X}$ by Theorem 8.3. By Proposition 8.43 u is uncrowded and therefore, u is regular and adapted. Interior nodes w_e corresponding to non-tropical edges $e \notin \operatorname{Edge}_{\text{trop}}(\Gamma)$ are ruled out in u for dimension reasons using the index relation (6.23). We next claim that the tropical type of u is Γ. If not it is of type Γ' and there is a non-trivial tropical edge collapse morphism $\Gamma' \to \Gamma$. The last statement in Theorem 8.3 implies that $\dim(T_{\text{trop}}(\Gamma')) \geq 2$. By Proposition 6.25 the expected dimension of the moduli space of broken maps is unaffected by collapsing tropical edges. Therefore, the index of u is at most 1. But the positive dimensionality of the tropical symmetry group $T_{\text{trop}}(\Gamma')$, and the fact that the action of the tropical symmetry group does not have infinitesimal stabilizers implies that $i^{\text{brok}}(u) \geq 2$, leading to a contradiction. We conclude that the tropical type of u is the same as that of the sequence u_ν. The only other phenomenon which occurs in the limit of $\{u_\nu\}_\nu$ is the formation of a boundary node $w \in C$ corresponding to an edge e of length $\ell(e)$ zero, or the length of a boundary edge $\ell(e)$ going to zero or infinity. □

CHAPTER 9

Gluing

In this chapter, we show that a rigid broken map can be glued at nodes to produce a family of unbroken map, one in each neck-stretched manifold. A fixed perturbation datum $\mathfrak{p}$ on a broken manifold $\mathfrak{X}$ can be glued in a natural way to produce a perturbation datum $\mathfrak{p}^\nu$ for X^ν which is equal to $\mathfrak{p}$ away from the neck regions. With respect to these perturbation data, we construct a bijection between rigid broken map and rigid maps in neck-stretched manifolds X^ν with sufficiently large neck lengths.

THEOREM 9.1. (Gluing) *Suppose that $u^0 : C \to \mathfrak{X}$ is a regular broken disk of index zero. There exists $\nu_0 > 0$ such that if $\nu \geq \nu_0$ there exists a family of unbroken disks $u_\nu : C^\nu \to X^\nu$ of index zero, with the property that $\lim_{\nu \to \infty}[u_\nu] = [u^0]$. For any area bound $E > 0$ there exists ν_0 such that for $\nu \geq \nu_0$ the correspondence $[u] \mapsto [u_\nu]$ defines an orientation-preserving bijection between the rigid moduli spaces $\mathcal{M}^{<E}(X^\nu, L)_0$ and $\mathcal{M}^{<E}_{\text{brok}}(\mathfrak{X}, L)_0$ for $\nu \geq \nu_0$.*

REMARK 9.2. The gluing operation is defined on broken maps, and not on tropical symmetry orbits of broken maps. In fact if the tropical symmetry group is non-trivial (it is necessarily finite for rigid broken maps), gluing different elements in the tropical symmetry orbit produces distinct sequences of unbroken maps. The convergence result also distinguishes between rigid maps in the same tropical symmetry orbit. Indeed, in Theorem 8.2, if the limit map is rigid then the limit is uniquely determined up to domain reparametrization.

Similar to other gluing theorems in pseudoholomorphic curves, the proof of Theorem 9.1 is an application of a quantitative version of the implicit function theorem for Banach manifolds. The steps are: construction of an approximation solution called the pre-glued map; construction of an approximate inverse to the linearized operator; quadratic estimates; application of the contraction mapping principle, and surjectivity of the gluing construction. Through the proof of the gluing theorem, the notation c denotes a ν-independent constant whose value may be different in every occurrence. The fact that gluing preserves the orientation of rigid moduli spaces is discussed in Remark 9.6.

9.1. The approximate solution

A pre-glued family for a broken map is constructed using tropical vertex positions of the broken map. A rigid map is modelled on a rigid tropical graph Γ. We recall that a rigid tropical graph has a unique tropical position map $\{\mathcal{T}(v) : v \in \text{Vert}(\Gamma)\}$

on its vertices. These positions determine the neck lengths for the approximate solution as follows. For any edge $e = (v_+, v_-)$ of Γ, there exists $l_e > 0$ such that

$$\mathcal{T}(v_+) - \mathcal{T}(v_-) = l_e \mathcal{T}(e), \tag{9.1}$$

where $\mathcal{T}(e) \in \mathfrak{t}_{\mathbb{Z}}$ is the slope of the edge.

The domain of any map in the glued family is a treed disk obtained by replacing tropical nodes in C with necks and leaving the tree part in C unchanged: For any $\nu \in \mathbb{R}_+$ the curve C_ν has a neck of length νl_e in place of a tropical node w_e corresponding to $e \in \mathrm{Edge}_{\mathrm{trop}}(\Gamma)$ in C. Denote by

$$\Gamma_{\mathrm{glue}} \tag{9.2}$$

the type of the glued curve. Recall that the lift of a tropical node w_e in C has matching coordinates (see (Matching at tropical nodes) in Definition 4.12) in the neighbourhoods U_e^+, U_e^- of the lifts $w_+(e)$, $w_-(e)$, which are denoted by

$$S_{v_\pm} \supset (U_e^\pm, w_\pm(e)) \xrightarrow{z_e^\pm} (\mathbb{C}, 0). \tag{9.3}$$

The coordinates can be chosen to be compositions of the complex exponential map (8.2) and linear functions from tangent spaces to $\mathbb{C}$. The glued curve C_ν is obtained from C by deleting a small disk in $U_e^\pm$ for every edge e in Γ, and gluing the remainder of the neighbourhoods $U_e^\pm$ using the identification $z_e^+ \sim e^{-\nu l_e} z_e^-$, and leaving the tree part T unchanged. So the surface and tree parts of C_ν are

$$C_\nu = S_\nu \cup T.$$

For future use in the proof we point out that the punctured neighbourhoods $U_e^\pm \backslash \{w_\pm(e)\}$ have matching logarithmic coordinates

$$(s_e, t_e) : U_e^- \backslash \{w_-(e)\} \to [0, \infty) \times S^1, \quad (s_e, t_e) : U_e^+ \backslash \{w_+(e)\} \to (-\infty, 0] \times S^1$$

given by $(s_e, t_e) := \pm \ln(z_e^\pm)$.

Translated maps can be glued at the nodal points to yield a sequence of approximate solutions for the holomorphic curve equation in neck-stretched manifolds. We recall that for a rigid tropical graph, the translation sequences are unique and are given by

$$t_\nu : \mathrm{Vert}(\Gamma) \to \nu B^\vee, \quad v \mapsto \nu \mathcal{T}(v). \tag{9.4}$$

Translation sequences give identifications between broken manifolds and neck-stretched manifolds. For any ν, $v \in \mathrm{Vert}(\Gamma)$ the map

$$\mathbb{e}^{t_\nu(v)} : X^\square_{\overline{P}(v)} \to X^\nu_{\overline{P}} \subset X^\nu$$

is defined in (3.42). The complement of the neck region in S^ν is a disjoint union

$$S^\nu \backslash \mathrm{Neck}(S^\nu) = \cup_{v \in \mathrm{Vert}(\Gamma)} S^\nu(v).$$

On $S^\nu(v)$, the approximate solution is defined as the translation of the map u_v^0, that is,

$$u_\nu^{\mathrm{pre}} := \mathbb{e}^{t_\nu(v)} u_v^0 \quad \text{on } S^\nu(v). \tag{9.5}$$

On the treed part $T \subset C^\nu$, we have

$$u_\nu^{\mathrm{pre}} := u_v^0 \quad \text{on } T \subset C^\nu.$$

Next, we define u_ν^{pre} on the neck region $\mathrm{Neck}_e(S^\nu)$ corresponding to a tropical edge $e = (v_+, v_-) \in \mathrm{Edge}_{\mathrm{trop}}(\Gamma)$. The matching condition on u^0 at the tropical node w_e implies that there is a point $x_e \in X^{\square}_{\overline{P}(e)}$ such that the map $u_{v_\pm}$ is asymptotically close to the trivial cylinder

$$u_e^{\mathrm{vert}} : \mathbb{R}_\pm \times S^1 \to X^{\square}_{\overline{P}(e)}, \quad (s,t) \mapsto e^{\mathcal{T}(e)(s+it)} x_e \tag{9.6}$$

near the nodal lift $w_e^\pm \in C_v^\pm$. That is, if we define $\zeta_e^\pm \in T_{u_e^{\mathrm{vert}}} X^{\square}_{\overline{P}(e)}$ by the condition

$$u_{v_\pm}^0 = \exp_{u_e^{\mathrm{vert}}} \zeta_e^\pm, \tag{9.7}$$

then

$$\|D^k \zeta_e^\pm(s,t)\| \leq c e^{-|s|} \tag{9.8}$$

for any $k \geq 0$. In (9.7) we note that the target space of the map $u_{v_\pm}^0$ resp. $\exp_{u_e^{\mathrm{vert}}} \zeta_e^\pm$ is $X^{\square}_{\overline{P}(v_\pm)}$ resp. $X^{\square}_{\overline{P}(e)}$, but the relation (9.7) is well-defined because near the node the image of $u_{v_\pm}$ lies in the $P(e)$-cylindrical end of $X^{\square}_{\overline{P}(v_\pm)}$ which is canonically identified to $X^{\square}_{\overline{P}(e)}$. We define a cylinder in $\mathrm{Neck}_e(C_\nu)$ corresponding to u_e^{vert} as

$$u_{e,\nu}^{\mathrm{vert}} := \mathbb{e}^{\frac{1}{2}(t_\nu(v_+)+t_\nu(v_-))} u_e^{\mathrm{vert}} : [\tfrac{-\nu l_e}{2}, \tfrac{\nu l_e}{2}] \times S^1 \to X^\nu.$$

The approximate solution on $\mathrm{Neck}_e(C_\nu)$ is the trivial cylinder $u_{e,\nu}^{\mathrm{vert}}$ corrected by a section obtained by patching ζ_e^+, ζ_e^- defined as

$$\begin{aligned} u_\nu^{\mathrm{pre}}(s,t) &:= \exp_{u_{e,\nu}^{\mathrm{vert}}}(\zeta_e^\nu(s,t)) : [\tfrac{-\nu l_e}{2}, \tfrac{\nu l_e}{2}] \times S^1 \to X^\nu, \\ \zeta_e^\nu(s,t) &:= \beta(-s)\zeta_e^-(s + \tfrac{\nu l_e}{2}, t) + \beta(s)\zeta_e^+(s - \tfrac{\nu l_e}{2}, t). \end{aligned} \tag{9.9}$$

where

$$\beta \in C^\infty(\mathbb{R}, [0,1]), \quad \begin{cases} \beta(s) = 0 & s \leq 0 \\ \beta(s) = 1 & s \geq 1 \end{cases} \tag{9.10}$$

is the cutoff function from (6.12). In other words, in the cylinder in C_ν corresponding to an edge e, one translates the domain on both ends by an amount $\frac{\nu l_e}{2}$, and then patches the map together using the cutoff function and geodesic exponentiation. The definition (9.5) of u_ν^{pre} in the complement of the neck, and its definition (9.9) on the neck together give a smooth map $u_\nu^{\mathrm{pre}} : C_\nu \to X^\nu$. Indeed using the relation $t_\nu(v_+) - t_\nu(v_-) = \nu l_e$, we conclude that the expression in (9.9) is equal to $\mathbb{e}^{t_\nu(v_\pm)} u_{v_\pm}^0$ in the neighborhood of the boundary component $\frac{\pm \nu l_e}{2} \times S^1$.

REMARK 9.3. The number $\nu \in \mathbb{R}_+$ is called the *map gluing parameter* to differentiate it from the parameter used for gluing nodes in a curve. In standard gluing results the map gluing parameter ν is taken to be the neck length in the domain C_ν of the glued map u^ν. In our setting, the map gluing parameter is the neck length ν in the target space X^ν. The domain neck lengths, approximately equal to νl_e, are allowed to vary in the Picard iteration argument. Note that resolving nodes w_e corresponding to tropical edges e does not increase the index of the map. As a result gluing an index zero broken map u produces an isolated map u^ν in each X^ν.

9.2. Fredholm theory for glued maps

We define a map between suitable Banach spaces whose zeroes describe pseudoholomorphic curves close to the approximate solution. Pseudoholomorphic maps are zeroes of the section

(9.11)
$$\mathcal{F}_\nu := (\overline{\partial}, \mathrm{ev}) : \mathcal{M}_{\Gamma_{\mathrm{glue}}} \times \mathrm{Map}(C_\nu, X^\nu) \to \Omega^{0,1}(S_\nu, u^*TX) \oplus \Omega^1(T, u^*TL) \oplus X(\Gamma_{\mathrm{glue}})$$

defined in (6.27) and consisting of $\overline{\partial}$ operators on surfaces, shift gradient operators on treed segments, and various evaluation maps. In this section we define metrics/norms in each of the spaces. We recall the operator $\mathcal{F}_\nu$ later in the section.

We describe a metric on the moduli space of treed disks. Recall that the domain of u is C, which is a curve of type Γ with surface part (S, j_S) and tree part T, and recall from (9.2) that the type Γ_{glue} of the glued curve is given by collapsing all the interior edges in Γ. In a neighborhood $U_{\mathcal{M}_\Gamma} \subset \overline{\mathcal{M}}_{\Gamma_{\mathrm{glue}}}$ of $\mathcal{M}_\Gamma$, there is a projection map

$$\pi_\Gamma : U_{\mathcal{M}_\Gamma} \to \mathcal{M}_\Gamma$$

such that any curve $C' \in U_{\mathcal{M}_\Gamma}$ is obtained by gluing at tropical nodes of the curve $\pi_\Gamma(C')$ using the coordinates (9.3). We remark that the domain curves C_ν of the approximate solutions which were constructed by gluing the tropical nodes in C lie in the neighborhood $U_{\mathcal{M}_\Gamma}$, and $\pi_\Gamma(C_\nu) = C$. The subset $U_{\mathcal{M}_\Gamma} \cap \mathcal{M}_{\Gamma_{\mathrm{glue}}}$ close to the boundary strata $\mathcal{M}_\Gamma$ is equipped with a metric

$$g_{\Gamma_{\mathrm{glue}}} : T(\mathcal{M}_{\Gamma_{\mathrm{glue}}} \cap U_{\mathcal{M}_\Gamma})^{\otimes 2} \to \mathbb{R} \tag{9.12}$$

that is cylindrical in the fibers of π_Γ. That is, each fiber of π_Γ is isometric to a product of cylinders $\Pi_{e \in \mathrm{Edge}_{\mathrm{trop}}(\Gamma)}(\mathbb{R} \times S^1)$ parametrized by gluing parameters (s_e, t_e).

In the neighborhood of each of the glued curves, we give a description of the complex structures on the surface components. As a preliminary step we describe the complex structures on curves close to the domain C of u. Let the complex curve (S, j_S) be the surface part C. A trivialization

$$\mathcal{S}_\Gamma | U_{\Gamma,C} \simeq S \times U_{\Gamma,C} \tag{9.13}$$

of the universal curve $\mathcal{S}_\Gamma \to \mathcal{M}_\Gamma$ in a neighborhood $U_{\Gamma,C} \subset \mathcal{M}_\Gamma$ of C yields a family of complex structures (as in (4.8))

$$U_{\Gamma,C} \to \mathcal{J}(S), \quad m \mapsto j_\Gamma(m).$$

We write j_Γ as a sum

$$j_\Gamma(m) = j_S + \Delta j_\Gamma(m)$$

and assume that the trivialization in (9.13) is chosen so that $\Delta j_\Gamma(m) = 0$ in neighborhoods of special points. Next, we consider a glued treed disk $C_\nu \in \mathcal{M}_{\Gamma_{\mathrm{glue}}}$ whose surface part $(S_\nu, j_{S_\nu}) \subset C_\nu$ is obtained by gluing the tropical nodes in $S \subset C$. In a neighborhood $U_{C_\nu} \subset \mathcal{M}_{\Gamma_{\mathrm{glue}}}$, we choose a trivialization of the universal curve $\mathcal{S}_{\Gamma_{\mathrm{glue}}} \to \mathcal{M}_{\Gamma_{\mathrm{glue}}}$ so that the induced family of complex structures

$$U_{C_\nu} \to \mathcal{J}(S_\nu), \quad m \mapsto j_\nu(m), \tag{9.14}$$

satisfies the following : The function j_ν is a sum

$$j_\nu = j_{S_\nu} + \Delta j_\bullet(m) + \Delta j^\nu_{neck}(m)$$

where

- j_{S_ν} is the complex structure on the glued curve S_ν and is m-independent,
- the function $m \mapsto \Delta j_\bullet(m)$ is supported in the complement of the neck regions of S_ν and is equal to $\Delta j_\Gamma(\pi_\Gamma(m))$,
- and $\Delta j^\nu_{\mathrm{neck}}(m)$ is supported in the neck regions of S_ν, and the support is contained in a uniformly bounded neighborhood the boundary of the neck, that is, there is a ν-independent constant L such that

$$\operatorname{supp}(\Delta j^\nu_{neck}) \subset \bigcup_{e\in \mathrm{Edge}_{\mathrm{trop}}(\Gamma)} \{s : \tfrac{\nu l_e}{2} - L \le |s| \le \tfrac{\nu l_e}{2}\} \times S^1 \subset A(l_\nu) \subset S_\nu. \tag{9.15}$$

  Further there is a ν-independent constant c such that on the neck region corresponding to any edge $e \in \mathrm{Edge}_{\mathrm{trop}}(\Gamma)$ and for any ν

$$\|\Delta j^\nu_{neck}(m)\|_{C^1} \approx c|\delta_e(m) - \nu l_e|,$$

  where $\delta_e(m) \in \mathbb{R}_+ \times S^1$ is the gluing parameter used at the node e to smoothen the node w_e in the curve $\pi_\Gamma(m)$.

Such a choice of the trivialization of the universal curve ensures that there is a ν-independent constant c such that for any two curves represented by m_1, m_2

$$|\Delta j^\nu(m_2) - \Delta j^\nu(m_1)| \le c d_{g_{\Gamma_{\mathrm{glue}}}}(m_1, m_2). \tag{9.16}$$

These uniform estimates are used in the proof of the quadratic estimate in Section 9.5.

The second space in the domain of (9.11) is a space of $W^{1,p}_{\mathrm{loc}}$maps

$$\mathrm{Map}(C_\nu, X^\nu)_{1,p} \subset \mathrm{Map}(S_\nu, X^\nu)_{1,p} \times \mathrm{Map}(T, X^\nu)_{1,p}$$

defined by requiring that the maps from S_ν and T agree on the intersection $S_\nu \cap T$. We recall that T is the disjoint union of treed segments $\cup_e T_e$ corresponding to boundary edges $e \in \mathrm{Edge}_\circ(\Gamma)$ with positive lengths; and $T_e = [0,1]$ if the segment is finite, and $\mathbb{R}_{\ge 0}$, $\mathbb{R}_{\le 0}$ or $\mathbb{R}$ if $\ell(e) = \infty$. The tangent space of $\mathrm{Map}(C_\nu, X^\nu)_{1,p}$ at a map $u : C_\nu \to X^\nu$ is the space of sections

$$\Omega^0(C_\nu, u^*TX) = \Omega^0(S_\nu, (u|S)^*TX^\nu) \oplus \Omega^0(T, (u|T)^*TL). \tag{9.17}$$

As in Abouzaid [**1**, 5.38] the first summand in (9.17) is equipped with a weighted Sobolev norm based on the decomposition of the section into a part constant on the neck and the difference on the neck corresponding to each edge $e \in \mathrm{Edge}_{\mathrm{trop}}(\Gamma)$ described as follows. Denote by

$$(s_e, t_e) \in [-\nu l_e/2, \nu l_e/2] \times S^1$$

the coordinates on the neck region created by the gluing at the node corresponding to the edge $e \in \mathrm{Edge}_{\mathrm{trop}}(\Gamma)$. Let

$$\lambda \in (0,1)$$

be a *Sobolev weight*. Define a *Sobolev weight function*

$$\kappa_\nu : C_\nu \to [0,\infty), \quad \kappa_\nu := \textstyle\sum_{e \in \mathrm{Edge}_{\mathrm{trop}}(\Gamma)} \beta(\nu l_e/2 - |s_e|)(\nu l_e/2 - |s_e|). \tag{9.18}$$

Here, the function $\beta(\nu l_e/2 - |s_e|)$ is extended by zero outside the neck region corresponding to e. As $\nu \to \infty$, κ_ν converges to the Sobolev weight function κ defined on the punctured curve $C - \{w_e : e \in \mathrm{Edge}_{\mathrm{trop}}(\Gamma)\}$ in (6.13). Given a section

$$\xi = (\xi_S, \xi_T) \in \Omega^0(C_\nu, u^*TX)$$

define

$$\begin{aligned} \|\xi\|_{1,p,\lambda} &:= \|\xi_S\|_{1,p,\lambda} + \|\xi_T\|_{1,p}^p \\ \|\xi_S\|_{1,p,\lambda} &:= \Big(\textstyle\sum_e \|\xi_{S,e}(0,0)\|^p + \displaystyle\int_{C_\nu} (\|\nabla \xi_S\|^p \\ &\quad + \|\xi_S - \textstyle\sum_e \beta(\nu l_e/2 - |s_e|)\mathbb{T}_s^u \xi_{S,e}(0,0)\|^p) \exp(\kappa_\nu \lambda p) \mathrm{d}\,\mathrm{Vol}_{C_\nu} \Big)^{1/p} \end{aligned} \tag{9.19}$$

where $\xi_{S,e}$ is the restriction of ξ_S to the neck region $[\frac{-\nu l_e}{2}, \frac{\nu l_e}{2}] \times S^1$ corresponding to the edge e, and $\mathbb{T}_s^u$ is parallel transport from $u(0,t)$ to $u(s,t)$ along $u(s',t)$, $s' \in [0,s]$. Let $\Omega^0(C_\nu, u^*TX)_{1,p,\lambda}$ be the Sobolev completion of $W^{1,p}_{\mathrm{loc}}$ sections with finite norm (9.19); these are sections whose difference from a covariant-constant section on the neck has an exponential decay behavior governed by the Sobolev constant λ.

The target space of (9.11) is a space of $(0,1)$-forms, which we equip with a weighted L^p norm. For a $(0,1)$-form $\eta \in \Omega^{0,1}(S_\nu, u^*TX^\nu)$ define

$$\|\eta\|_{0,p,\lambda} = \left(\int_{S_\nu} \|\eta\|^p \exp(\kappa_\nu \lambda p) \mathrm{d}\,\mathrm{Vol}_{S_\nu} \right)^{1/p}.$$

The implicit function theorem is applied on the $\overline{\partial}$ map pulled back by an exponential map. Pointwise geodesic exponentiation defines a map (using Sobolev multiplication estimates)

$$\exp_{u_\nu^{\mathrm{pre}}} : \Omega^0(C_\nu, (u_\nu^{\mathrm{pre}})^* TX^\nu)_{1,p,\lambda} \to \mathrm{Map}_{1,p}(C_\nu, X^\nu) \tag{9.20}$$

where $\mathrm{Map}_{1,p}(C_\nu, X^\nu)$ denotes maps of class $W^{\mathrm{loc}}_{1,p}$ from C_ν to X^ν. We define

(9.21)

$$\begin{gathered} \overline{\partial} : U_{C_\nu} \times \Omega^0(C_\nu, (u_\nu^{\mathrm{pre}})^*TX^\nu)_{1,p,\lambda} \to \Omega^{0,1}(S_\nu, (u_\nu^{\mathrm{pre}})^*TX^\nu)_{0,p} \oplus \Omega^1(T, (u_\nu^{\mathrm{pre}})^*TL), \\ (m,\xi) \mapsto \mathbb{T}^{-1}_{u_\nu^{pre},\xi} \left(\overline{\partial}_{j(m),J}(\exp_{u_\nu^{\mathrm{pre}}} \xi)|S_\nu, \left(\frac{1}{\lambda_e} \frac{d}{ds} + \mathrm{grad}_F \right) (\exp_{u_\nu^{\mathrm{pre}}} \xi)|T \right), \end{gathered}$$

where

$$\begin{aligned} \mathbb{T}_{u_\nu^{\mathrm{pre}},\xi} : \ &\Omega^{0,1}(S_\nu, (u_\nu^{\mathrm{pre}})^*TX^\nu)_{0,p,\lambda} \oplus \Omega^1(T, (u_\nu^{\mathrm{pre}})^*TL)_{0,p} \\ &\to \Omega^{0,1}(S_\nu, (\exp_{u_\nu^{\mathrm{pre}}}(\xi))^*TX^\nu)_{0,p,\lambda} \oplus \Omega^1(T, (\exp_{u_\nu^{\mathrm{pre}}}(\xi))^*TL)_{0,p} \end{aligned}$$

is the parallel transport defined using an almost-complex connection, and we recall that $U_{C_\nu} \subset \mathcal{M}_{\Gamma_{\text{glue}}}$is a neighborhood of C_ν; and

$$\lambda_e := \begin{cases} \ell_e(m), & e \in \text{Edge}_{\circ,-}^{(0,\infty)}(\Gamma), \\ 1, & e \in \text{Edge}_{\circ}^{\infty}(\Gamma). \end{cases}$$

In order to construct local models for moduli of adapted tree disks, we require that the treed disks C_ν have a collection of interior leaves $e_1, \ldots, e_{d(\bullet)}$ and

$$(\exp_{u_\nu^{\text{pre}}}(\xi))(e_i) \in D, \quad i = 1, \ldots, n.$$

Additionally we require matching conditions at boundary nodes and lifts of $S_\nu \cap T_\nu$. Using notation from the proof of transversality (Theorem 6.29), these constraints may be incorporated into $\mathcal{F}_\nu$ to produce a map

$$\begin{aligned} \mathcal{F}_\nu : U_{C_\nu} \oplus \Omega^0(S_\nu, (u_\nu^{\text{pre}})^* TX^\nu)_{1,p,\lambda} \oplus \Omega^0(T, (u_\nu^{\text{pre}})^* TL)_{1,p} \\ \to \Omega^{0,1}(C_\nu, (u_\nu^{\text{pre}})^* TX^\nu)_{0,p,\lambda} \oplus \Omega^1(T, (u_\nu^{\text{pre}})^* TL)_{0,p} \oplus TX(\Gamma_{\text{glue}})/\Delta(\Gamma_{\text{glue}}). \end{aligned}$$

whose zeroes correspond to *adapted* pseudoholomorphic maps near the pre-glued map u_ν^{pre}.

9.3. Error estimate

We estimate the failure of the approximate solution to be an exact solution in the Banach norms of the previous section. To derive the estimate we split the treed disk C_ν into neck regions corresponding to tropical nodes in C, namely

$$\text{Neck}_e(S_\nu) := \{(s_e, t_e) \in [-\tfrac{\nu l_e}{2}, \tfrac{\nu l_e}{2}] \times S^1\} \subset S_\nu, \quad \forall e \in \text{Edge}_{\text{trop}}(\Gamma);$$

and its complement

$$C_\nu^\bullet := C_\nu \setminus \cup_e \text{Neck}_e(S_\nu), \quad \text{and} \quad C_\nu^\bullet = S_\nu^\bullet \cup T.$$

The one-form $\mathcal{F}_\nu(0)$ has contributions from the difference between the perturbation data $P(C)$ and $P(C_\nu)$ on the complement of the neck regions, and from the cutoff function and the difference between J_u and $J_{u_\nu^{\text{pre}}}$ on the neck regions :

$$\begin{aligned} \|\mathcal{F}_\nu(0)\|_{L^{p,\lambda}(S_\nu)} = \|\overline{\partial}_{J_{u_\nu^{\text{pre}}}} u^\nu_{\text{pre}}\|_{L^{p,\lambda}(S_\nu^\bullet)} + \sum_{e \in \text{Edge}_\circ(\Gamma)} \|\tfrac{1}{\lambda_e} + (\text{grad}_F)_{u_\nu^{\text{pre}}}\|_{L^p(T_e)} \\ + \textstyle\sum_e \|\overline{\partial} \exp_{e^{(s+it)\mathcal{T}(e)} x_e}(\beta(-s_e)\zeta_e^-(s_e + \nu l_e/2, t_e) \\ + \beta(s_e)\zeta_e^+(s_e - \nu l_e/2, t_e))\|_{L^{p,\lambda}(\text{Neck}_e(S_\nu))}. \end{aligned} \tag{9.22}$$

The first two terms may not vanish because the perturbation is domain dependent: in the complement of the neck regions, the map u_ν^{pre} is $P(C)$-holomorphic but not $P(C_\nu)$-holomorphic. For any metric $d_{\overline{\mathcal{M}}}$ on the compactified moduli space $\overline{\mathcal{M}}_{\Gamma_{\text{glue}}}$ of treed disks, the distance between the domain curves is bounded as

$$d_{\overline{\mathcal{M}}}(C_\nu, C) \le c \max_{e \in \text{Edge}_{\text{trop}}(\Gamma)} \exp(-\nu l_e).$$

Therefore, the distance between the domain-dependent perturbations has a similar bound. On the complement of the necks u_ν^{pre} is $J(C)$-holomorphic, so

$$\|\overline{\partial}_{J_{u_\nu^{\text{pre}}}} u^\nu_{\text{pre}}\|_{L^{p,\lambda}(S_\nu^\bullet)} \le c\|J(C) - J(C_\nu)\|_{L^\infty} \le c \max_{e\in \text{Edge}_{\text{trop}}(\Gamma)} \exp(-\nu l_e). \tag{9.23}$$

The last term in the right-hand side of (9.22) is equal to

$$\begin{aligned}\textstyle\sum_e \|(D\exp_{e^{(s+it)\mathcal{T}(e)}x_e}(\mathrm{d}\beta(-s_e)\zeta_e^-(s_e+\nu l_e/2,t_e) + \mathrm{d}\beta(s_e)\zeta_e^+(s_e-\nu l_e/2,t_e)) +\\ (\beta(-s_e)\mathrm{d}\zeta_e^-(s_e+\nu l_e/2,t_e) + \beta(s_e)\mathrm{d}\zeta_e^+(s_e-\nu l_e/2,t_e)))^{0,1}\|_{L^{p,\lambda}(\text{Neck}_e(S_\nu))}.\end{aligned}$$

On the neck regions, the almost complex structure is domain-independent since it is equal to the background almost complex structure (see (6.4)). Holomorphicity of u implies that the terms are non-zero only in the support of $d\beta_e$ which is contained in the interval

$$[-1,1]\times S^1 \subset [-\tfrac{\nu l_e}{2}, \tfrac{\nu l_e}{2}]\times S^1 = \text{Neck}_e(S_\nu).$$

Both $\zeta_e^\pm$ and its derivative decay at the rate of e^{-s_e} on the cylindrical end $\mathbb{R}_+\times S^1$ in S°, see (9.8). As a result both the difference between J_u and $J_{u_\nu^{\text{pre}}}$, and the terms containing $d\beta$ are bounded by $ce^{-l_e\nu/2}$ where c is a constant independent of ν. The Sobolev weight function (9.18) has a multiplicative factor of $e^{\lambda l_e \nu/2}$, and therefore,

$$\|\mathcal{F}_\nu(0)\| \le c\textstyle\sum_{e\in\text{Edge}_{\text{trop}}(\Gamma)} e^{-(1-\lambda)l_e\nu/2}, \tag{9.24}$$

with c a constant independent of ν. (See Abouzaid [**1**, 5.10]).

9.4. Uniform right inverse

We construct a uniformly bounded right inverse for the linearized operator of the approximate solution from the given right inverses of the pieces of the broken map. We denote the linearized operator by

$$D_{u_\nu^{\text{pre}}} := D\mathcal{F}_\nu([C_\nu],0).$$

Given an element

$$\eta = (\eta_S,\eta_T) \in \Omega^{0,1}(S_\nu, (u_\nu^{\text{pre}}|S_\nu)^*TX^\nu)_{L^{p,\lambda}} \oplus \Omega^1(T,(u_\nu^{\text{pre}}|T)^*TL)_{L^p}$$

in the target space of $D_{u_\nu^{\text{pre}}}$ one obtains an element in the target space of the linearized operator D_{u^0} of the broken map

$$\tilde\eta = (\eta_v)_{v\in\text{Vert}(\Gamma)} \oplus \eta_T, \quad \eta_v \in \bigoplus_{v\in\text{Vert}(\Gamma)} \Omega^{0,1}(S_v^\circ, (u_v^0)^*TX^\circ_{\overline{P}(v)})$$

as follows. The element $\tilde\eta$ is equal to η in the tree components $T\subset C$ and in the complement of the neck region on the surface components $S_v \subset C$. On the neck region for an edge $e=(v_+,v_-)\in\text{Edge}_{\text{trop}}(\Gamma)$, $\tilde\eta$ is defined by restricting $\eta|\,\text{Neck}_e(S_\nu)$ to half the neck and extending by zero :

$$\eta_{v,+} : (-\infty,0]\times S^1 \to X^\square_{\overline{P}(e)}, \quad (s,t)\mapsto \begin{cases}\mathbb{T}_e^{+,\nu}\eta(s+\frac{\nu l_e}{2},t), & s\ge \frac{-\nu l_e}{2}\\ 0, & s<\frac{-\nu l_e}{2},\end{cases}$$

and

$$\eta_{v,-} : [0,\infty) \times S^1 \to X^{\square}_{\overline{P}(e)}, \quad (s,t) \mapsto \begin{cases} \mathbb{T}_e^{-,\nu}\eta(s - \frac{\nu l_e}{2}, t), & s \le \frac{\nu l_e}{2} \\ 0, & s > \frac{\nu l_e}{2}, \end{cases},$$

where

$$\mathbb{T}_e^{\pm,\nu} : \Gamma((u_\nu^{\mathrm{pre}})^* T X^{\square}_{\overline{P}(e)}) \to \Gamma((u_v^{\pm})^* T X^{\square}_{\overline{P}(e)}) \tag{9.25}$$

is parallel transport along the path

$$\exp_{e^{(s+it)\mathcal{T}(e)} x_e}(\rho \zeta_e^\nu(s \pm \tfrac{\nu l_e}{2}, t) + (1-\rho)\zeta_e^{\pm}(s,t)), \quad \rho \in [0,1],$$

ζ_ν^e is defined in (9.9) and $\zeta_e^{\pm}$ is defined in (9.7). Since the broken map u^0 is regular and isolated, its linearized operator is bijective. We recall that the linearized operator is a map of Banach spaces (see (6.31))

$$D_{u^0} : T_m\mathcal{M}_\Gamma \times \mathrm{Map}(C,\mathfrak{X})_{1,p,\lambda} \to \Omega^{0,1}(S,(u^0|S)^*T\mathfrak{X}) \oplus \Omega^1(T,(u^0|T)^*TL) \oplus \mathrm{ev}_\Gamma^* T\mathfrak{X}/T\Delta.$$

Bijectivity of D_{u^0} implies there is an inverse (m,ξ) for the element $(\tilde{\eta}, 0 \in \mathrm{ev}_\Gamma^* T\mathfrak{X}/T\Delta)$. We write $\xi = ((\xi_v)_{v\in \mathrm{Vert}(\Gamma)}, \xi_T)$. The vanishing of the last term in $D_{u^0}(m,\xi)$ means that ξ satisfies matching conditions at tropical and boundary nodes, and the interior markings z_i satisfy the divisor constraint :

$$\xi(z_i) \in T_{u_\pm(z_i)}\mathfrak{D}.$$

The matching at tropical nodes implies that for any tropical edge $e = (v_+, v_-)$, the limit of ξ_{v_+}, ξ_{v_-} at the cylindrical end e is equal :

$$\xi_{v_+,e} = \xi_{v_-,e} =: \xi_e \in T_{x_e} X^{\square}_{\overline{P}(e)}.$$

We now define the approximate inverse by patching the inverse of the linearization of the broken map. We denote the approximate inverse of $D_{u_\nu^{\mathrm{pre}}}$ by Q^ν. On the complement of the neck regions in C_ν we define

$$Q^\nu(\eta) := \xi \quad \text{on } C_\nu \backslash \cup_e \mathrm{Neck}_e(S_\nu);$$

and on the neck region corresponding to a tropical edge $e \in \mathrm{Edge}_{\mathrm{trop}}(\Gamma)$ we patch the solutions $\xi_{v_\pm}$ together using a cutoff function:

$$\begin{aligned} Q^\nu \eta := {} & \beta\left(-s + \tfrac{1}{4}\nu l_e\right)((\mathbb{T}_e^{-,\nu})^{-1}\xi_{v_-}(s + \nu l_e/2) - \mathbb{T}_e^\nu \xi_e) \\ & + \beta\left(s + \tfrac{1}{4}\nu l_e\right)((\mathbb{T}_e^{+,\nu})^{-1}\xi_{v_+}(s - \nu l_e/2)) - \mathbb{T}_e^\nu\xi_e) + \mathbb{T}_e^\nu \xi_e \\ & \in \Omega^0(C_\nu, (u_\nu^{\mathrm{pre}})^*TX)_{1,p,\lambda}, \end{aligned} \tag{9.26}$$

where $\mathbb{T}_e^{\pm,\nu}$ is defined in (9.25) and $\mathbb{T}_e^\nu$ is the parallel transport from $\Gamma((u_{\nu,e}^{\mathrm{vert}})^* T X^{\square}_{\overline{P}(e)})$ to $\Gamma((u_\nu^{\mathrm{pre}})^* T X^{\square}_{\overline{P}(e)})$.

The approximate inverse Q^ν is uniformly bounded for all ν: It follows easily from the construction of Q^ν that the operations $\eta \mapsto \tilde{\eta}$, $\tilde{\eta} \mapsto \xi$ are uniformly bounded operators for all ν, and that on the complement of the neck regions there is a uniform bound on the norm of the operator $\xi \mapsto (Q_\nu \eta)|C_\nu^\bullet$. On the neck $\mathrm{Neck}_e(S_\nu)$ corresponding to a tropical edge $e \in \mathrm{Edge}_{\mathrm{trop}}(\Gamma)$ there is a uniform bound

$$\|((Q^\nu\eta) - \mathbb{T}_e\xi_e)e^{\kappa_\nu \lambda}\|_{L^p(\mathrm{Neck}_e(S_\nu))} \le c\|\xi\|_{W^{1,p,\lambda}}.$$

However, we need a κ_ν-weighted Sobolev norm bound on $((Q^\nu \eta) - \mathbb{T}_e(Q^\nu\eta)(0,0)_e)$, which follows from the observation

$$|Q_\nu(0,0)_e - \xi_e| \leq |\xi_{v_+}(\tfrac{-\nu l_e}{2},0)_e - \xi_e| + |\xi_{v_-}(\tfrac{\nu l_e}{2},0)_e - \xi_e| \leq ce^{-\lambda\nu l_e/2}\|\eta\|_{L^{p,\lambda}},$$

which implies that

$$\|(\mathbb{T}_e Q^\nu(0,0)_e - \mathbb{T}_e\xi_e)e^{\kappa_\nu\lambda}\| \leq c\|\eta\|_{L^{p,\lambda}}.$$

We have thus shown that there is a ν-independent constant c such that

$$\|Q^\nu\| \leq c.$$

Next, we give an error estimate for the approximate inverse. We need to bound the quantity $D_{u^\nu_{\text{pre}}}Q^\nu\eta - \eta$. On the complement of the neck regions of C_ν (including treed segments), this quantity is bounded by the difference in the domain-dependent perturbation P on (C,j) and (C_ν, j^ν). Similar to (9.23) in the estimate of $\mathcal{F}^\nu(0)$, we have the bound

$$\|D_{u^\nu_{\text{pre}}}Q^\nu\eta - \eta\|_{L^{p,\lambda}(C^\bullet_\nu)} \leq c \max_{e\in \text{Edge}_{\text{trop}}(\Gamma)} \exp(-\nu l_e). \tag{9.27}$$

Next we bound $\|D_{u^\nu_{\text{pre}}}Q^\nu\eta - \eta\|_{L^{p,\lambda}}$ on $\text{Neck}_e(S_\nu)$ corresponding to an edge $e = (v_+, v_-) \in \text{Edge}_{\text{trop}}(\Gamma)$. We often drop the parallel transport notation in the analysis noting that both the transport map and its derivative contribute smooth multiplicative factors that decay as $ce^{\nu l_e/2 - |s|}$ on the neck. We recall that $\text{Neck}_e(S_\nu) \simeq \{(s,t) \in [\frac{-\nu l_e}{2}, \frac{\nu l_e}{2}] \times S^1\}$, and analyze the half $\{s \geq 0\}$, the other half being symmetric. On $\{s \geq 0\}$, we have

$$\eta = \eta_{v_+} = D_{u^0}\xi_{v_+}, \quad Q^\nu\eta = \xi_{v_+} + \beta(-s + \tfrac{\nu l_e}{4})(\xi_{v_-} - \mathbb{T}_e\xi_e),$$

and therefore,

$$D_{u^\nu_{\text{pre}}}Q^\nu\eta - \eta = D_{u^\nu_{\text{pre}}}(\beta(-s + \tfrac{\nu l_e}{4})(\xi_{v_-} - \mathbb{T}_e\xi_e)) + (D_{u^0_{v_+}} - D_{u^\nu_{\text{pre}}})\xi_{v_+}. \tag{9.28}$$

We have

$$\|(D_{u^0_{v_+}} - D_{u^\nu_{\text{pre}}})\xi_{v_+}\|_{L^{p,\lambda}(S_\nu)} \leq ce^{-\nu l_e/2}\|\xi_{v_+}\|_{W^{1,p,\lambda}(S^\circ_{v_+})} \leq ce^{-\nu l_e/2}\|\eta\|_{L^{p,\lambda}(C_\nu)} \tag{9.29}$$

since $d_{X^\nu}(u^0_{v_+}(s - \frac{\nu l_e}{2}, t), u^\nu_{\text{pre}}(s,t)) \leq ce^{-\nu l_e/2}$. Since $D_{u^0_{v_-}}\xi_{v_-} = 0$ on $\{s \geq 0\}$, by a similar bound to (9.29) on $(D_{u^0_{v_-}} - D_{u^\nu_{\text{pre}}})\xi_{v_-}$, we have

$$\|D_{u^\nu_{\text{pre}}}\xi_{v_-}\|_{L^{p,\lambda}(\{s\geq 0\})} \leq ce^{-\nu l_e/2}\|\xi_{v_-}\|_{W^{1,p,\lambda}(S^\circ_{v_-})} \leq ce^{-\nu l_e/2}\|\eta\|_{L^{p,\lambda}(C_\nu)}. \tag{9.30}$$

From (9.28), (9.29), (9.30) we conclude

$$\|D_{u^\nu_{\text{pre}}}Q^\nu\eta - \eta\|_{L^{p,\lambda}} \leq \|\beta(D_{u^\nu_{\text{pre}}}(\mathbb{T}_e\xi_e))\| + \|d\beta(\xi_{v_-} - \mathbb{T}_e\xi_e)\| + ce^{-\nu l_e/2}\|\eta\|. \tag{9.31}$$

To bound the first term in the right-hand side of (9.31) we observe that $\mathbb{T}_e\xi_e$ is the parallel transport of a covariant constant section on the trivial cylinder u^{ver}_e, and therefore,

$$\|D_{u^\nu_{\text{pre}}}(\mathbb{T}_e\xi_e)\| = \|(D_{u^\nu_{\text{pre}}} - D_{u^{\text{ver}}_e})(\mathbb{T}_e\xi_e)\| \leq ce^{-\nu l_e/2}|\xi_e| \leq ce^{-\nu l_e/2}\|\eta\|_{L^{p,\lambda}}.$$

It remains to bound the second term which is supported in the unit interval $[\frac{\nu l_e}{4}, \frac{\nu l_e}{4} + 1] \times S^1$ in the neck $[\frac{-\nu l_e}{2}, \frac{\nu l_e}{2}] \times S^1$. In this interval, the Sobolev weight in the curve S_{v_-} and the glued curve S_ν differ by $\frac{\nu l_e}{2}$:

$$\kappa_\nu(s,t) = \kappa_{S^\circ_{v_-}}(s + \nu l_e/2) - \nu l_e/2, \quad \mp s \geq \nu l_e/2.$$

Therefore there is a ν-independent constant c such that

$$\|\mathrm{d}\beta(s - \nu l_e/4)\xi_{v_-}\|_{L^{p,\lambda}(S_\nu)} \leq ce^{-\lambda\nu l_e/2}\|\xi_{v_-}\|_{W^{1,p,\lambda}(S^\circ_v)} \leq ce^{-\lambda\nu l_e/2}\|\eta\|_{L^{p,\lambda}(S_\nu)}. \tag{9.32}$$

From (9.27), (9.32), one obtains an estimate as in Fukaya-Oh-Ohta-Ono [**34**, 7.1.32], Abouzaid [**1**, Lemma 5.13]: For some constant $c > 0$, for any ν

$$\|D_{u_\nu^{\mathrm{pre}}} Q^\nu - \mathrm{Id}\,\| < c \min_{e \in \mathrm{Edge}_{\mathrm{trop}}(\Gamma)} (\exp(-\lambda\nu l_e/2), \exp(-(1-\lambda)\nu l_e/2)). \tag{9.33}$$

It follows that for ν sufficiently large an actual inverse may be obtained from the Taylor series formula

$$D^{-1}_{u_\nu^{\mathrm{pre}}} = Q^\nu (D_{u_\nu^{\mathrm{pre}}} Q^\nu)^{-1} = Q^\nu \sum_{k \geq 0} (I - Q^\nu D_{u_\nu^{\mathrm{pre}}})^k.$$

The approximate inverse Q^ν is uniformly bounded for all ν. For large enough ν, (9.33) implies that $\|D_{u_\nu^{\mathrm{pre}}} Q^\nu - \mathrm{Id}\| \leq \frac{1}{2}$, and so,

$$\|D^{-1}_{u_\nu^{\mathrm{pre}}}\| \leq 2\|Q^\nu\| \leq c. \tag{9.34}$$

9.5. Uniform quadratic estimate

We obtain a uniform quadratic estimate for the non-linear terms in the map cutting out the moduli space locally. We will prove that there exists a constant c such that for all ν

$$\|D_{(m_1,\xi_1)}\mathcal{F}_\nu(m_2,\xi_2) - D_{u_\nu^{\mathrm{pre}}}(m_2,\xi_2)\| \leq c\|(m_1,\xi_1)\|_{1,p,\lambda}\|(m_2,\xi_2)\|_{1,p,\lambda}. \tag{9.35}$$

We prove the quadratic estimate for the $\overline{\partial}$ term on surface components. The other terms in the operator $\mathcal{F}_\nu$ are left to the reader as the proof is similar. As a preliminary step we prove a quadratic estimate on a simpler operator. Define

$$\mathcal{G} : \Omega^0(S_\nu, (u_\nu^{\mathrm{pre}})^* TX^\nu)_{W^{1,p,\lambda}} \to \Omega^1(S_\nu, (u_\nu^{\mathrm{pre}})^* TX^\nu)_{L^{p,\lambda}}, \quad \xi \mapsto \mathbb{T}_\xi^{-1} d(\exp_{u_\nu^{\mathrm{pre}}} \xi).$$

Claim 9.4. There is a ν-independent constant c such that

$$\|D_{\xi_1}\mathcal{G}(\xi_2) - D_{u_\nu^{\mathrm{pre}}}\mathcal{G}(\xi_2)\|_{0,p,\lambda} \leq c\|\xi_1\|_{1,p,\lambda}\|\xi_2\|_{1,p,\lambda}$$

if $\xi_{1 1,p,\lambda}$ is small enough.

Proof of Claim. For $x \in X^\nu$, $\xi_1, \xi_2 \in T_x X^\nu$, define

$$\overline{\mathbb{T}}^x_{-\xi_1} : T_x X \to T_x X, \quad \xi_2 \mapsto \mathbb{T}_{\xi_1}^{-1} \frac{d}{\mathrm{d}\tau} \exp_x(\xi_1 + \tau\xi_2)|_{\tau=0}.$$

It extends to a map on sections. For a map $u : S_\nu \to X_\nu$, and $\xi_1 \in \Gamma(S_\nu, u^* TX_\nu)$ define

$$\overline{\mathbb{T}}^u_{-\xi_1} : \Gamma(S_\nu, u^* TX_\nu) \to \Gamma(S_\nu, u^* TX_\nu), \quad \xi_2 \mapsto (z \mapsto \overline{\mathbb{T}}^{u(z)}_{-\xi_1(z)}(\xi_2(z)).$$

We have

$$D_{\xi_1}\mathcal{G}(\xi_2) - D_{u_\nu^{\mathrm{pre}}}\mathcal{G}(\xi_2) = \tfrac{d}{d\tau}\mathbb{T}^{-1}_{\xi_1+\tau\xi_2}\mathrm{d}_z(\exp_{u_\nu^{\mathrm{pre}}}(\xi_1+\tau\xi_2)) - \tfrac{d}{d\tau}\mathbb{T}^{-1}_{\tau\xi_2}\mathrm{d}_z(\exp_{u_\nu^{\mathrm{pre}}}(\tau\xi_2))$$

$$= \nabla_z(\overline{\mathbb{T}}^u_{-\xi_1}\xi_2 - \xi_2), \tag{9.36}$$

where in the second line ∇_z is differentiation along the domain curve S_ν. In the second equality we switch the order of differentiation between $\frac{d}{d\tau}$ and d_z. For any $z \in S_\nu$ we have a uniform pointwise estimate

$$\begin{aligned}|D_{\xi_1}\mathcal{G}(\xi_2) - D_{u_\nu^{\mathrm{pre}}}\mathcal{G}(\xi_2)| &\le |\nabla_z(\overline{\mathbb{T}}^{u(z)}_{-\xi_1(z)} - \mathrm{Id})|\cdot|\xi_2(z)| + |\overline{\mathbb{T}}^{u_\nu^{\mathrm{pre}}(z)}_{-\xi_1(z)} - \mathrm{Id}|\cdot|(\nabla_z\xi_2)(z)|\\ &\le c((|du_\nu^{\mathrm{pre}}(z)|\cdot|\xi_1(z)| + |\nabla_z\xi_1(z)|)\cdot|\xi_2(z)| + |\xi_1(z)|\cdot|\nabla_z\xi_2(z)|)\end{aligned}$$

where the constant c is ν-independent. Indeed, such a uniform constant exists because the complement of the neck regions is compact and identical for all ν; and the neck regions have a cylindrical metric, and only the lengths of the cylinders vary with ν. Then,

$$\begin{aligned}\|D_{\xi_1}\mathcal{G}(\xi_2) - D_{u_\nu^{\mathrm{pre}}}\mathcal{G}(\xi_2)\| &\le c(\|\mathrm{d}u_\nu^{\mathrm{pre}}\|_{L^\infty}\cdot\|\xi_1\|_{L^\infty}\cdot\|\xi_2\|_{L^{p,\lambda}} + \|\nabla_z\xi_1(z)\|_{L^{p,\lambda}}\cdot\|\xi_2\|_{L^\infty}\\ &\quad + \|\xi_1\|_{L^\infty}\cdot\|\nabla_z\xi_2(z)\|_{L^{p,\lambda}}) \le c\|\xi_1\|_{1,p,\lambda}\|\xi_2\|_{1,p,\lambda},\end{aligned}$$

where for the last inequality, we use the fact that $\|\mathrm{d}u_\nu^{\mathrm{pre}}\|_{L^\infty}$ is uniformly bounded for all ν by construction, and the following bound from Sobolev embedding : For any section $\xi \in W^{1,p,\lambda}(S_\nu)$ there is a ν-independent constant c such that $\|\xi\|_{L^\infty} \le c\|\xi\|_{W^{1,p,\lambda}}$. This proves the Claim. □

We obtain the quadratic estimate (9.35) for $\mathcal{F}_\nu$ by adapting the proof of the above claim. We note that, compared to $\mathcal{G}$, the operator $\mathcal{F}_\nu$ additionally consists of a projection to $(0,1)$-forms :

$$\mathcal{F}_\nu : (m,\xi) \mapsto \mathbb{T}_\xi^{-1}\pi^{0,1}_{j(m),J(\xi)}\mathrm{d}(\exp_{u_\nu^{\mathrm{pre}}}\xi),$$

where we abbreviate $J_{\exp_{u_\nu^{\mathrm{pre}}}\xi}$ as $J(\xi)$. The analog of (9.36) is

$$\begin{aligned}&D_{(m_1,\xi_1)}\mathcal{F}_\nu(\xi_2) - D_{u_\nu^{\mathrm{pre}}}\mathcal{F}_\nu(\xi_2)\\ &\quad= \tfrac{d}{d\tau}\mathbb{T}^{-1}_{\xi_1+\tau\xi_2}\pi^{0,1}_{j_\nu(m_1+\tau m_2),J(\xi_1+\tau\xi_2)}\mathrm{d}_z(\exp_{u_\nu^{\mathrm{pre}}}(\xi_1+\tau\xi_2))|_{\tau=0}\\ &\qquad - \tfrac{d}{d\tau}\mathbb{T}^{-1}_{\tau\xi_2}\pi^{0,1}_{j_\nu(\tau m_2),J(\tau\xi_2)}\mathrm{d}_z(\exp_{u_\nu^{\mathrm{pre}}}(\tau\xi_2))|_{\tau=0}\\ &\quad= \pi^{0,1}_{j_\nu(m_1),J(\xi_1)}\nabla_z(\overline{\mathbb{T}}^u_{-\xi_1}\xi_2) + (\nabla_z(\pi^{0,1}_{j_\nu(m_1),J(\xi_1)}))(\overline{\mathbb{T}}^u_{-\xi_1}\xi_2)\\ &\qquad - \pi^{0,1}_{j(C_\nu),J(0)}(\nabla_z\xi_2) - (\nabla_z\pi^{0,1}_{j(C_\nu),J(0)})\xi_2,\end{aligned} \tag{9.37}$$

where τm_2 resp. $m_1 + \tau m_2$, $\tau \in [0,1]$ is a path in a neighborhood U_{C_ν} of C_ν in the moduli space $\mathcal{M}_{\mathrm{glue}}$. The estimate (9.35) can be obtained from the above expression in a similar way to the proof of Claim 9.4. We point out that the metric g_{C_ν} on U_{C_ν} is cylindrical on the non-compact ends of $\mathcal{M}_{\mathrm{glue}}$, and with respect to this metric $\|j_\nu(m_1) - j_\nu(m_2)\|_{C^1} \le c d_{g_{\Gamma_{\mathrm{glue}}}}(m_1,m_2)$, see (9.16).

9.6. Picard iteration

We apply the implicit function theorem to obtain an exact solution. We recall a version of the Picard lemma [**60**, Proposition A.3.4].

LEMMA 9.5. (Picard lemma) *Let X and Y be Banach spaces, $U \subset X$ be an open set containing 0, and $f : U \to Y$ be a smooth map. Suppose $df(0)$ is invertible with inverse $Q : Y \to X$. Suppose c and $\epsilon > 0$ are constants such that $\|Q\| \leq c$, $B_\epsilon \subset U$, and*

$$\|df(x) - df(0)\| \leq \frac{1}{2c} \quad \forall x \in B_\epsilon(0).$$

Suppose $f(0) \leq \frac{\epsilon}{4c}$. Then, there is a unique point $x_0 \in B_\epsilon$ satisfying $f(x_0) = 0$.

Picard's Lemma and the estimates (9.24), (9.34), (9.35) imply the existence of a solution $(m(\nu), \xi(\nu))$ to the equation

$$\mathcal{F}^\nu(m(\nu), \xi(\nu)) = 0$$

for each ν. The map

$$u_\nu := \exp_{u_\nu^{\mathrm{pre}}}(\xi(\nu))$$

is a $(j(m(\nu)), J^\nu)$-holomorphic map to X^ν. Additionally, there is a ν-independent constant $\epsilon > 0$ such that $(m(\nu), \xi(\nu))$ is the unique zero of $\mathcal{F}^\nu$ in an ϵ-neighbourhood of $((C_\nu, j^\nu), u_\nu^{\mathrm{pre}})$ with respect to the $g_{\Gamma_{\mathrm{glue}}}$-norm on $m(\nu)$ and the weighted Sobolev norm $W^{1,p,\lambda}$ on $\xi(\nu)$.

9.7. Surjectivity of gluing

We show that the gluing construction gives a bijection. Note that any family $[u'_\nu : C'_\nu \to X^\nu]$ converges to a broken map $u : C \to \mathfrak{X}$ by Theorem 8.2. To prove the bijection we must show that any such family of maps is in the image of the gluing construction. Since the implicit function theorem used to construct the gluing gives a unique solution in a neighbourhood, it suffices to show that the maps u'_ν are close, in the Sobolev norm used for the gluing construction, to the approximate solution u_ν^{pre} defined by (9.9).

We first show that the domain curves of the converging sequence of maps are close enough to the domains of the approximate solution with respect to the cylindrical metric $g_{\Gamma_{\mathrm{glue}}}$ from (9.12). In the definition of Gromov convergence, the convergence of domains implies that $C'_\nu \to C$ in the compactified moduli space $\overline{\mathcal{M}}_{\Gamma_{\mathrm{glue}}}$, which implies

$$\pi_\Gamma(C'_\nu) \to C \quad \text{in } \mathcal{M}_\Gamma. \tag{9.38}$$

We additionally need to prove that the distance $d_{g_{\Gamma_{\mathrm{glue}}}}(C'_\nu, C_\nu) \to 0$ where the metric $g_{\Gamma_{\mathrm{glue}}}$ is cylindrical in the non-compact ends of $\overline{\mathcal{M}}_{\Gamma_{\mathrm{glue}}}$. By assumption the limit map u does not have any tropical symmetry. Therefore, the translation sequence t_ν is uniquely determined by the tropical graph of u and coincides with the translations used for gluing. By (Thin cylinder convergence), for any tropical edge $e = (v_+, v_-) \in \mathrm{Edge}_{\mathrm{trop}}(\Gamma)$, the gluing parameters $l'_\nu(e) + i\theta'_\nu(e)$ of the curves C'_ν satisfy

$$\lim_{\nu\to\infty} \theta'_\nu(e) = 0, \quad \lim_{\nu\to\infty} (t_\nu(v_+) - t_\nu(v_-) - \mathcal{T}(e) l'_\nu(e)) = 0$$

The gluing parameter of C_ν at the edge e is νl_e, which satisfies the relation

$$t_\nu(v_+) - t_\nu(v_-) = \mathcal{T}(e)\nu l_e.$$

Therefore, $l'_\nu(e) - \nu l_e \to 0$, and $d_{g_{\Gamma_{\text{glue}}}}(C_\nu, C'_\nu) \to 0$ as $\nu \to \infty$. In addition there is a biholomorphism

$$\phi_\nu : (C_\nu, j_\nu([C'_\nu])) \to C'_\nu, \tag{9.39}$$

where j_ν is defined in (9.14).

Next, we show that the maps in the converging sequence are close enough to the approximate solutions. Via the identification (9.39), we view u'_ν as a map on C_ν. We need to bound the section $\xi'_\nu \in \Omega^0(C_\nu, (u_\nu^{pre})^* TX^\nu)$ defined by the equation $u'_\nu = \exp_{u_\nu^{pre}} \xi'_\nu$ in the weighted Sobolev norm (9.19). Consider the neck region in C_ν corresponding to an edge e with coordinates

$$(s_e, t_e) \in [-\tfrac{\nu l_e}{2}, \tfrac{\nu l_e}{2}] \times \mathbb{R}/2\pi\mathbb{Z}.$$

Denote the midpoint of the neck as

$$0_e := \{(s_e, t_e) = (0,0)\} \in C_\nu.$$

In the neck region, the maps u_ν^{pre} and u'_ν are equal to a vertical cylinder perturbed by a quantity that decays exponentially in the middle of the cylinder. The vertical cylinder is determined by $u_\nu^{pre}(0_e)$ resp. $u'_\nu(0_e)$. The sequence $u_\nu^{pre}(0_e)$ converges to x_e because of the asymptotic decay of the sections $\zeta^\pm$. The sequence $u'_\nu(0_e)$ converges to x_e by (Thin cylinder convergence). Indeed, since the complex structure $\Delta j^\nu(C'_\nu)$ is standard on a truncation $[-\frac{\nu l_e}{2} - L, \frac{\nu l_e}{2} - L] \times S^1$ of the neck (see (9.15)), and the mid point of the cylinder is preserved by the biholomorphism ϕ in (9.39), (Thin cylinder convergence) is applicable with the coordinates (s_e, t_e). On the neck region, the section ξ'_ν and its derivatives decay exponentially :

$$|D^k \xi'_\nu(s_e, t_e)| \le c e^{-(\nu l_e/2 - |s_e|)}, \quad k \in \{0,1\}.$$

This inequality follows from the decay of the terms $\zeta_\nu^\pm$ in the definition of u_ν^{pre}, and the breaking annulus lemma applied to u'_ν. Consequently $\|\xi'_\nu\|_{W^{1,p,\lambda}}$ can be made small enough by taking a large ν and shrinking the neck by a fixed amount: that is, we decrease the cylinder length to $\nu l_e - C$ where C is a constant independent of ν. Next we consider the complement of the neck regions. Here, the sequences u'_ν and u_ν^{pre} uniformly converge to u. So, by taking ν large enough the maps u_ν^{pre} and u'_ν are $W^{1,p}$-close enough in the complement of the neck regions.

Remark 9.6. (Gluing preserves orientations) From the definition of orientation of moduli spaces in Remark 6.30 it is easily seen that gluing preserves orientation in the special case that the Cauchy-Riemann operator on each surface component of the broken map u^0 is complex linear. It remains to show that the interpolation from the linearized operator D_{u^0} resp. D_{u_ν} to a complex linear operator $\overline{\partial}_{u^0}$ resp. $\overline{\partial}_{u_\nu}$ changes the orientation signs of the broken maps and the glued maps by the same amount.

We recall from [**60**, Appendix 2] that for a Fredholm operator of index 0 the trivialization of the determinant line bundle along a path of operators may be interpreted as a crossing number. For a surjective Fredholm operator $D : X \to Y$ of index 0 between Banach spaces X, Y, the determinant bundle is $\det(D) := \mathbb{R}$. Let $D_t : X \to Y$, $t \in [0, 1]$ be a path of Fredholm operators of index 0 whose end-points D_0, D_1 are surjective. For a generic path there is a finite number of non-surjective operators in the path and each of them is a simple regular crossing. A crossing at t is *regular* if $\dot{D}_t(\ker(D_t))$ is transverse to $\mathrm{im}(D_t)$, and *simple* if $\dim(\ker(D_t)) = 1$. The orientation signs at D_0 and D_1 differ by the number of such crossings in the path, see [**60**, Proposition A.2.4].

To prove that gluing preserves orientation we need to construct a family of paths of operators of Fredholm index zero, namely

$$D_t^\nu : \mathcal{B}_\nu \to \mathcal{E}_\nu, \quad \nu \in [\nu_0, \infty], t \in [0, 1], \quad D_0^\infty = D_{u^0}, D_0^\nu = D_{u_\nu},$$

each of whose starting point D_0^ν is the linearized operator of the glued map u_ν if $\nu < \infty$ and the broken map u^0 if $\nu = \infty$; and the end point D_1^ν is a regular complex linear Cauchy-Riemann operator, and the path D_t^ν has simple regular crossings for all ν, and each crossing D_t^∞ corresponds to a family of crossings $D_{t_\nu}^\nu$ with t_ν converging to t. We leave this construction to the reader. This finishes the Remark.

9.8. Tubular neighbourhoods and true boundary

In Section 8.6 we proved a convergence result that the boundaries of one-dimensional moduli spaces of rigid broken maps consist of maps that either have a length zero boundary edge, or a broken boundary edge. Theorem 9.7 is the converse gluing result, and it shows that given a broken map u containing a boundary edge that is broken or has zero length, the map u indeed occurs as the codimension one boundary of the expected moduli spaces. The convergence and gluing results, together, allow us to identify the true boundary of one-dimensional moduli spaces. A stratum containing a zero length edge is actually a 'fake boundary stratum' as it is the boundary of two different rigid strata with opposite induced boundary orientations. Thus the 'true boundary strata' are those that have a broken boundary edge, see Remark 9.8.

THEOREM 9.7. *Let $\underline{\mathfrak{p}} = (\mathfrak{p}_\Gamma)_\Gamma$ be a coherent regular perturbation datum for all types Γ. Suppose Γ is a type of broken treed disks and $\underline{x} \in \mathcal{I}(L)^{n+1}$ is a set of limits for boundary leaves such that $i(\Gamma, \underline{x}) = 1$.*

- *(a)* (Tubular neighbourhoods) *If a type Γ_{glue} of broken maps is obtained from a type Γ by collapsing an edge $e \in \mathrm{Edge}_{\circ,-}(\Gamma')$ with $\ell(e) = 0$, or by making an edge length finite/non-zero, then the stratum $\mathcal{M}_\Gamma^{\mathrm{brok}}(L, \underline{x})$ is a codimension one boundary of in $\overline{\mathcal{M}}_{\Gamma_{\mathrm{glue}}}^{\mathrm{brok}}(L, \underline{x})$ and has a tubular neighborhood in it; and*
- *(b)* (Orientations) *the orientations are compatible with the morphisms* (Cutting an edge) *and* (Collapsing an edge/Making an edge finite/non-zero) *in the following sense:*

(i) *Suppose* Γ, Γ_0, Γ_c *are types of broken maps related by the morphisms*

$$\Gamma \xrightarrow{\text{Make } \ell(e) \text{ zero}} \Gamma_0 \xrightarrow{\text{Collapse } e} \Gamma_c,$$

and $i^{\mathrm{brok}}(\Gamma, \underline{x}) = i^{\mathrm{brok}}(\Gamma_c, \underline{x}) = 1$, $i^{\mathrm{brok}}(\Gamma_0, \underline{x}) = 0$. *Then the boundary orientation induced by* $\mathcal{M}^{\mathrm{brok}}_{\Gamma}(L, \underline{x})$ *on* $\mathcal{M}^{\mathrm{brok}}_{\Gamma_0}(L, \underline{x})$ *is the opposite of the boundary orientation induced by* $\mathcal{M}^{\mathrm{brok}}_{\Gamma_c}(L, \underline{x})$ *on* $\mathcal{M}^{\mathrm{brok}}_{\Gamma_0}(L, \underline{x})$.

(ii) *Suppose* Γ_f, Γ, Γ_1, Γ_2 *are types of broken maps and* $\underline{x}_1 \in \mathcal{I}(L)^{d_\circ(\Gamma_1)+1}$, $\underline{x}_2 \in \mathcal{I}(L)^{d_\circ(\Gamma_2)+1}$ *are labels such that there are morphisms*

$$(\Gamma_f, \underline{x}) \xleftarrow{\text{Make } \ell(e) \text{ finite}} (\Gamma, \underline{x}) \xrightarrow{\text{Cut } e} (\Gamma_1, \underline{x}_1) \times (\Gamma_2, \underline{x}_2),$$

and $i^{\mathrm{brok}}(\Gamma_f, \underline{x}) = 1$, $i^{\mathrm{brok}}(\Gamma, \underline{x}) = i^{\mathrm{brok}}(\Gamma_1, \underline{x}_1) = i^{\mathrm{brok}}(\Gamma_2, \underline{x}_2) = 0$. *Then there is an isomorphism*

$$\mathcal{M}^{\mathrm{brok}}_{\Gamma}(L, \underline{x}) \simeq \mathcal{M}^{\mathrm{brok}}_{\Gamma_1}(L, \underline{x}_1) \times \mathcal{M}^{\mathrm{brok}}_{\Gamma_1}(L, \underline{x}_2),$$

and the boundary orientation on $\mathcal{M}^{\mathrm{brok}}_{\Gamma}(L, \underline{x})$ *induced by* $\mathcal{M}^{\mathrm{brok}}_{\Gamma_f}(L, \underline{x})$ *is related to the product orientation by a sign* $(-1)^{\circ}$ *that depends only on the domain type* Γ *and the labels* $\underline{x}$.

Outline of proof. The tubular neighborhood is constructed by gluing and there are different cases depending on the morphism:

Case 1 : *Collapsing a boundary edge* e *of zero length.*
The proof is by gluing a boundary node, which is non-tropical since the Lagrangian is disjoint from relative divisors. The gluing proof is on the same lines as the gluing of tropical nodes in Theorem 9.1, so we point out the differences, referring the reader to [**34**, Chapter 7.1] for the full proof.

- Firstly we choose strip-like coordinates on a neighborhood of the boundary node w_e and we use weighted Sobolev spaces on these strips, see Definition 7.1.3 in [**34**]. A difference from our Fredholm set-up for broken maps in Section 6.4 is that we do not need cylindrical coordinates in the target space since there are no relative divisors.
- The next difference is that the map gluing parameter ν refers to the neck length parameter in the glued curve (see Remark 9.3), and therefore in the domain of the operator $\mathcal{F}_\nu$ in (9.11) the component $\mathcal{M}_{\Gamma_{\mathrm{glue}}}$ is replaced by a subspace where the neck length corresponding to the node w_e is fixed to be ν. The proof of 'surjectivity of gluing' also simplifies, since it is enough to show convergence of the domain curves away from the neck as in (9.38) and the arguments following (9.38) in that paragraph are not necessary.

Cases 2 and 3 do not involve gluing of surface components.

Case 2 : *Making a boundary edge length* $\ell(e)$ *non-zero.*
The proof structure of Theorem 9.1 can be used in this case also. However, the map gluing parameter ν corresponds to the edge length $\ell(e)$ in the glued curve and goes to zero in the limit. The domain curve for the glued map $u_\nu : C_\nu \to \mathfrak{X}$ is allowed to vary among the set of curves of type Γ_{glue} with $\ell(e) = \nu$. The approximate solution $u^{\mathrm{pre}}_\nu : C_\nu \to \mathfrak{X}$ is defined to be equal to u, and the map u^{pre}_ν is defined to be constant

on the treed segment T_e, so the error $\|D\mathcal{F}_\nu(u_\nu^{\text{pre}})\|$ is $\leq c\nu$. The rest of the steps are analogous to the proof of Theorem 9.1.

CASE 3 : *Making a boundary edge length $\ell(e)$ finite.*
Similar to Case 2, the map gluing parameter ν is equal to the length $\ell(e)$ of the edge e, but unlike Case 2, it goes to infinity in the limit. The domain curve for the glued map $u_\nu : C_\nu \to \mathfrak{X}$ is allowed to vary among the set Γ_{glue} curves with $\ell(e) = \nu$. The approximate solution $u_\nu^{\text{pre}} : C_\nu \to \mathfrak{X}$ is defined to be equal to u on all components except the treed segment T_e, where it is defined by pre-gluing the Morse trajectory (see (2.66) in Schwarz [**78**]) at the edge e in the map u. The error estimate is $\leq ce^{-\nu}$ and is derived in a similar way to Section 9.3. The rest of the steps are analogous to the proof of Theorem 9.1.

For the orientations result, part (bi) follows from the definition of orientations in Remark 6.30 and the sign computation for part (bii) is carried out in Seidel's book [**82**, (12.25)]. □

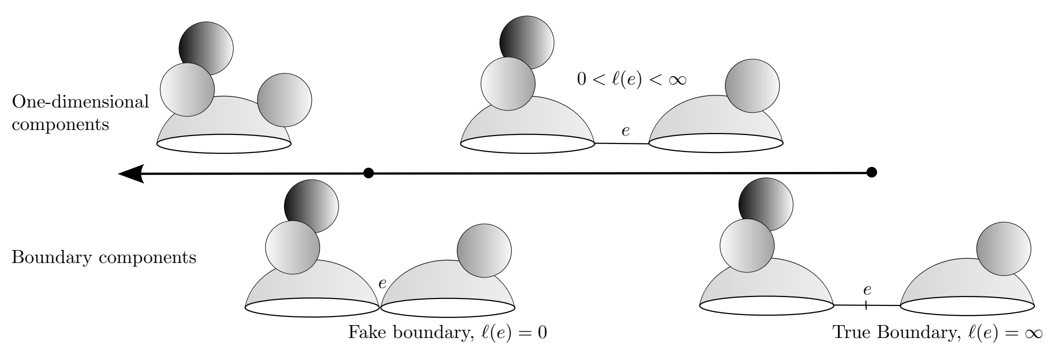

FIGURE 9.1. True and fake boundary strata of a one-dimensional component of the moduli space of treed holomorphic disks. The sphere components lie in different pieces of the tropical manifold.

REMARK 9.8. (True and fake boundary strata) There are two types of strata that occur as the codimension one boundary of a one-dimensional moduli space – one with a boundary edge of length zero, and the second with a boundary edge containing a breaking. The first is a *fake* boundary, and the second one is a *true* boundary as we explain: If a type Γ_0 of broken maps has a boundary edge of length zero, we may make the edge length $\ell(e)$ non-zero or we may collapse the edge (by disk gluing) to produce the following types

$$\Gamma \xrightarrow{\text{Make } \ell(e) \text{ zero}} \Gamma_0 \xrightarrow{\text{Collapse } e} \Gamma_c.$$

Let $\underline{x}$ be a tuple of input and output such that $i^{\text{brok}}(\Gamma, \underline{x}) = 0$. By Theorem 9.7, the moduli space $\mathcal{M}_{\Gamma_0}^{\text{brok}}(\underline{x})$ is the boundary of the one-dimensional moduli spaces $\mathcal{M}_{\Gamma_c}^{\text{brok}}(\underline{x})$ and $\mathcal{M}_{\Gamma}^{\text{brok}}(\underline{x})$ with opposite orientations. So, $\mathcal{M}_{\Gamma_0}^{\text{brok}}(\underline{x})$ does not represent a component in the topological boundary of the compactified moduli space

$$\bigcup_{\Gamma:\Gamma \text{ is rigid, } i^{\text{brok}}(\Gamma, \underline{x})=1} \overline{\mathcal{M}}_{\Gamma}^{\text{brok}}(L, \underline{x}).$$

This is the fake boundary in Figure 9.1. The only (true) boundary components of one-dimensional strata thus consist of maps with a single broken Morse trajectory, see Figure 9.1.

CHAPTER 10

Broken Fukaya algebras

In this chapter we describe A_∞ -algebra structures defined by counting treed holomorphic disks on broken and unbroken manifolds, and show that they are equivalent up to A_∞ -homotopy.

10.1. A_∞ algebras

The set of treed holomorphic disks has the structure of an A_∞ -algebra. A_∞ -algebras were introduced by Stasheff [**85**] in order to capture algebraic structures on the space of cochains on loop spaces. We follow the sign convention in Seidel [**81**]. A $\mathbb{Z}_2$-*graded* A_∞ *algebra* consists of a $\mathbb{Z}_2$-graded vector space A together with for each $d \geq 0$ a multilinear degree zero *composition map*

$$m^d : \ A^{\otimes d} \to A[2-d]$$

satisfying the A_∞ -*associativity equations* [**81**, (2.1)]

$$0 = \sum_{j,k \geq 0, j+k \leq d} (-1)^{j + \sum_{i=1}^j |a_i|} m^{d-k+1}(a_1, \ldots, a_j, m^k(a_{j+1}, \ldots, a_{j+k}), a_{j+k+1}, \ldots, a_d) \tag{10.1}$$

for any $d \geq 0$ and any tuple of homogeneous elements $a_1, \ldots, a_d$ with degrees $|a_1|, \ldots, |a_d| \in \mathbb{Z}_2$. The notation $[2-d]$ indicates a degree shift of $2-d$, so that the degree of $m^d(a_1 + \cdots + a_d)$ is $\sum_i |a_i| + 2 - d \in \mathbb{Z}_2$. The signs appearing in (10.1) are the *shifted Koszul signs*, that is, the Koszul signs for the shifted grading in which the structure maps have degree one as in Kontsevich-Soibelman [**53**]. One of the first of these associativity relations is

$$m^2(m^0, a) - (-1)^{|a|} m^2(a, m^0) + m^1(m^1(a)) = 0, \quad \forall a \in A. \tag{10.2}$$

A *strict unit* for A is an element $1_A \in A$ such that

$$m^2(1_A, a) = a = (-1)^{|a|} m^2(a, 1_A), \quad m^d(\ldots, 1_A, \ldots) = 0, \forall d \neq 2. \tag{10.3}$$

A *strictly unital* A_∞ algebra is an A_∞ algebra equipped with a strict unit.

The element $m^0(1) \in A$ (where $1 \in \mathbb{R}$ is the unit) is called the *curvature* of the algebra. The A_∞ algebra A is *flat* if the curvature vanishes. The *cohomology* of a flat A_∞ algebra A is defined by

$$H(m^1) = \frac{\ker(m^1)}{\operatorname{im}(m^1)}.$$

The algebra structure on $H(m^1)$ is given by

$$[a_1 a_2] = (-1)^{|a_1|}[m^2(a_1, a_2)]. \tag{10.4}$$

The A_∞-algebras in our work are 'curved', that is $m^0(1)$ does not vanish. Cohomology can be defined for a curved strictly unital A_∞ -algebra if the curvature is a multiple of the unit : $m^0(1) \in \Lambda 1_A$, because in this case, $(m^1)^2 = 0$ by (10.2).

More generally, the cohomology exists for any solution to the projective Maurer-Cartan equation [**34**]. The projective Maurer-Cartan equation for $b \in A$ is

$$m^0(1) + m^1(b) + m^2(b, b) + \ldots \in \Lambda 1_A. \tag{10.5}$$

A solution $b \in A$ of odd degree to the equation (10.5) is called a *weakly bounding cochain* and the set of all the odd solutions is denoted $MC(A)$. Given a weakly bounding cochain $b \in MC(A)$, we define a *deformed composition map*

$$m_b^n(a_1, \ldots, a_n) = \sum_{i_1,\ldots,i_{n+1}} m^{n+i_1+\ldots+i_{n+1}}(\underbrace{b, \ldots, b}_{i_1}, a_1, \underbrace{b, \ldots, b}_{i_2}, a_2, b, \ldots, b, a_n, \underbrace{b, \ldots, b}_{i_{n+1}}) \tag{10.6}$$

over all possible combinations of insertions of the element $b \in A^+$ between (and before and after) the elements $a_1, \ldots, a_n$. The maps m_b^n define an A_∞ structure on A if b has odd degree. By the weakly bounding cochain condition $m_b^0(1) \in \Lambda 1_A$, and the A_∞ relations imply

$$(m_b^1)^2(a) = m_b^2(m_b^0(1), a) - m_b^2(a, m_b^0(1)) = 0.$$

Consequently the cohomology

$$H(m_b^1) = \ker(m_b^1)/\operatorname{im}(m_b^1)$$

is well-defined. The function

$$W : MC(A) \to \Lambda, \quad b \mapsto W(b),$$

where $m_b^0 = W(b)1_A$, is called the *potential* of the curved A_∞ algebra A.

One also has homotopy notions of algebra morphisms. Let A_0, A_1 be A_∞ algebras. An A_∞ *morphism* $\mathcal{F}$ from A_0 to A_1 consists of a sequence of linear maps

$$\mathcal{F}^d : \ A_0^{\otimes d} \to A_1[1-d], \quad d \geq 0$$

such that the following holds:

$$\sum_{i+j \leq d} (-1)^{i+\sum_{j=1}^{i}|a_j|} \mathcal{F}^{d-j+1}(a_1, \ldots, a_i, m_{A_0}^j(a_{i+1}, \ldots, a_{i+j}), a_{i+j+1}, \ldots, a_d) = \sum_{i_1+\ldots+i_m=d} m_{A_1}^m(\mathcal{F}^{i_1}(a_1, \ldots, a_{i_1}), \ldots, \mathcal{F}^{i_m}(a_{i_1+\ldots+i_{m-1}+1}, \ldots, a_d)) \tag{10.7}$$

where the first sum is over integers i, j with $i + j \leq d$, the second is over partitions $d = i_1 + \ldots + i_m$. The A_∞ -morphisms we consider are *convergent* in the sense that $\mathcal{F}^0 \in$ The first relation in the family of relations in (10.7) is

$$\mathcal{F}^1(m^0_{A_0}) = m^0_{A_1} + m^1_{A_1}(\mathcal{F}^0) + m^2_{A_1}(\mathcal{F}^0, \mathcal{F}^0) + \ldots$$

An A_∞ morphism $\mathcal{F}$ is *unital* if and only if

$$\mathcal{F}^1(1_{A_0}) = 1_{A_1}, \quad \mathcal{F}^k(a_1, \ldots, a_i, 1_{A_0}, a_{i+2}, \ldots, a_k) = 0 \tag{10.8}$$

for every $k \geq 2$ and every $0 \leq i \leq k-1$, where 1_{A_0} resp. 1_{A_1} is the strict unit in A_0 resp. A_1.

10.2. Composition maps

In this section, we describe A_∞ algebras whose composition maps are given by counts of treed holomorphic maps, both in an ordinary symplectic manifold and in a broken manifold. The boundary of the disks map to a Lagrangian, which in the broken case, is contained in the complement of relative divisors

The *Fukaya algebra* is an A_∞ algebra whose structure coefficients are defined by counts of pseudoholomorphic treed disks as follows. Let q be a formal variable and Λ the *universal Novikov field* of formal sums with rational coefficients

$$\Lambda = \left\{ \sum_{i \in \mathbb{Z}_{\geq 0}} c_i q^{\alpha_i} \;\middle|\; c_i \in \mathbb{C},\ \alpha_i \in \mathbb{R}, \quad \alpha_i \to \infty \right\}.$$

Denote by $\Lambda_{\geq 0}$ resp. $\Lambda_{>0}$ the subalgebra with only non-negative resp. positive exponents. Denote by

$$\Lambda^\times = (\mathbb{C} - \{0\}) + \Lambda_{>0} \subset \Lambda_{\geq 0}$$

the subgroup of formal power series with invertible leading coefficient and non-negative exponents. An A_∞ -algebra A with coefficients in Λ is *convergent* if $\mu^0_A \in \Lambda_{>0}$. This property ensures that the expression μ^b_0 in the Maurer-Cartan equation is well-defined for any $b \in A$ with positive q-valuation.

Lagrangians will be equipped with additional data called *brane structures*. We assume Lagrangians are compact, connected and oriented. A brane structure consists of

- a relative spin structure (see [**34**, Chapter 44] for the definition),
- and a local system, which is an element

$$\mathrm{Hol}_L \in \mathcal{R}(L) = \mathrm{Hom}(\pi_1(L), \Lambda^\times).$$

When working with multiple Lagrangian submanifolds, one needs a grading on each of the Lagrangians, see Seidel [**79**]. We do not use gradings in this paper.

For a Lagrangian L with a brane structure define the space of Floer cochains over the Novikov ring

$$CF^{\mathrm{geom}}(L) := \bigoplus_{d \in \mathbb{Z}_2} CF^d(L), \quad CF^d(L) := \bigoplus_{x \in \mathcal{I}_d(L)} \Lambda_{\geq 0}\langle x \rangle$$

where $\mathcal{I}_d(L)$ is the set of index d critical points of the Morse function $F : L \to \mathbb{R}$, see Definition 4.5. The composition maps are as follows.

DEFINITION 10.1. (Composition maps) For admissible perturbation data $(\mathfrak{p}_\Gamma)_\Gamma$ on the manifold X define

$$m_{d(\circ)} : (CF^{\text{geom}}(L))^{\otimes d(\circ)} \to CF^{\text{geom}}(L)$$

on generators by

$$m_{d(\circ)}(x_1, \ldots, x_{d(\circ)}) = \sum_{x_0, u \in \mathcal{M}_\Gamma(X, L, \mathfrak{p}_\Gamma, \underline{x})_0} w(u) x_0 \tag{10.9}$$

where the sum is over all types Γ of maps with simple intersections with the stabilizing divisor and with $d(\circ)$ incoming boundary edges, and

$$w(u) := (-1)^{\heartsuit} (d_\bullet(\Gamma)!)^{-1} \operatorname{Hol}([\partial u]) \epsilon(u) q^{A(u)}. \tag{10.10}$$

The terms in (10.10) are as below :

(a) $\heartsuit = \sum_{i=1}^{d(\circ)} i|x_i|$,
(b) $\operatorname{Hol}_L([\partial u]) \in \Lambda^\times$ is the evaluation of the local system $\operatorname{Hol}_L \in \mathcal{R}(L)$ on the homotopy class of loops $[\partial u] \in \pi_1(L)$ defined by going around the boundary of each disk component in the treed disk once,
(c) $d_\bullet(\Gamma)$ the number of interior markings on the map u,
(d) $\epsilon(u) \in \{\pm 1\}$ is the orientation sign, see Remark 6.30.

The A_∞ relation follows from the description of the true boundary of the moduli space:

THEOREM 10.2. (A_∞ algebra for a Lagrangian in a symplectic manifold) *For any admissible perturbation system $\underline{\mathfrak{p}} = (\mathfrak{p}_\Gamma)_\Gamma$ on the manifold (X, ω), the maps $(m_{d(\circ)})_{d(\circ) \geq 0}$ on $CF^{\text{geom}}(L)$ satisfy the axioms of a convergent (possibly curved) A_∞ -algebra $CF^{\text{geom}}(L)$.*

PROOF. We prove the A_∞-associativity relations by counting the ends of one-dimensional moduli spaces of treed holomorphic disks. We perform the count for a fixed number of interior and boundary markings, and fixed limits on the input and output treed segments, and then sum over all choices. Consider integers $d(\circ), d(\bullet) \geq 0$ and a tuple $\underline{x} \in (\mathcal{I}(L))^{d(\circ)+1}$ of inputs and outputs. We denote the one-dimensional component of the moduli space of treed disks of rigid type with $d(\circ)$ boundary inputs and $d(\bullet)$ interior markings by

$$\mathcal{M}_{d(\bullet), d(\circ)}(L, \underline{x}) := \cup_{\Gamma : i(\Gamma, \underline{x}) = 1} \mathcal{M}_\Gamma(L, \underline{x}),$$

where Γ ranges over rigid types containing $d(\bullet)$ interior markings and $d(\circ)$ boundary inputs. By the unbroken version of Proposition 8.44 and Remark 9.8 (see [**18**, Remark 4.22]), the true boundary of $\mathcal{M}_{d(\bullet), d(\circ)}(L, \underline{x})$ consists of disks with a single broken edge, that is, an $e \in \operatorname{Edge}_{\circ,-}$ with $\ell(e) = \infty$. The boundary points of one-dimensional moduli spaces occur in pairs, and therefore,

$$\sum_{u \in \mathcal{M}_\Gamma(L, \underline{x})} \epsilon_\partial(u) (d(\bullet)!)^{-1} = 0, \tag{10.11}$$

where, in the summation, Γ ranges over all types in the true boundary of $\mathcal{M}_{d(\bullet),d(\circ)}(L,\underline{x})$, and the sign $\epsilon_\partial(u) \in \{\pm 1\}$ is given by the induced orientation on the boundary of $\mathcal{M}_{d(\bullet),d(\circ)}(L,\underline{x})$. Consider one such type Γ, and suppose Γ_+, Γ_- are the treed disk types obtained by cutting the broken edge, each containing $d_\pm(\bullet) := d_\bullet(\Gamma_\pm)$ interior markings. For any critical point $x \in \mathcal{I}(L)$, denote by $(\underline{x},x)_+$, $(\underline{x},x)_-$ the input-output labelling on Γ_+, Γ_- where both ends of the broken edge are labelled x. There is a bijection

$$\mathcal{M}_\Gamma(L,\underline{x}) \simeq \bigcup_{x\in\mathcal{I}(L)} \mathcal{M}_{\Gamma_+}(L,(\underline{x},x)_+))_0 \times \mathcal{M}_{\Gamma_-}(L,(\underline{x},x)_-))_0. \tag{10.12}$$

However, any map in the right-hand side corresponds to $\binom{d(\bullet)}{d_+(\bullet)}$ maps in the true boundary of $\mathcal{M}_{d(\bullet),d(\circ)}(L,\underline{x})$ as follows: There are $\binom{d(\bullet)}{d_+(\bullet)}$ bijections

$$f : \{1,\ldots,d_+(\bullet)\} \sqcup \{1,\ldots,d_-(\bullet)\} \to \{1,\ldots,d(\bullet)\}$$

whose restrictions to $\{1,\ldots,d_+(\bullet)\}$, $\{1,\ldots,d_-(\bullet)\}$ are order-preserving. Given a pair in (u_+,u_-) in the right-hand side of (10.12), each bijection f corresponds to a map u in the true boundary of $\mathcal{M}_{d(\bullet),d(\circ)}(L,\underline{x})$, where the i-th interior marking in $u_\pm$ is labelled $f(i)$ in u. Therefore the expression in (10.11) is equal to

$$0 = \sum_{(\Gamma_+,\Gamma_-),x} \left\{ \left(\sum_{u_+\in\mathcal{M}_{\Gamma_+}(L,(\underline{x},x)_+)} \epsilon(u_+)(d_\bullet(\Gamma_+)!)^{-1} \right) \cdot \left(\sum_{u_-\in\mathcal{M}_{\Gamma_-}(L,(\underline{x},x)_-)} \epsilon(u_-)(d_\bullet(\Gamma_-)!)^{-1} \right) \frac{\epsilon_\partial(u)}{\epsilon(u_+)\epsilon(u_-)} \right\}, \tag{10.13}$$

where x ranges over critical points in $\mathcal{I}(L)$ and (Γ_+,Γ_-) range over all pairs of types obtained by cutting an edge in a treed disk type occurring in the boundary of $\mathcal{M}_{d(\bullet),d(\circ)}(L,\underline{x})$. The sign contribution $\frac{\epsilon_\partial(u)}{\epsilon(u_+)\epsilon(u_-)}$ is equal to the difference in the orientation of the moduli space of treed disks with a broken edge when viewed as a boundary of a larger stratum, and when viewed as a product, and is equal to the shifted Koszul sign in the A_∞-associativity relation (10.1) by [**82**, (12.25)]. Adding the equations (10.13) corresponding to all $d(\bullet) \geq 0$, we obtain the associativity relation on the maps $(m_d)_{d\geq 0}$ on $CF^{\text{geom}}(L)$ corresponding to $d = d(\circ)$. □

Composition maps for the broken Fukaya algebra are defined analogously by counts of broken disks. Assume that $L \subset \mathfrak{X}$ is a Lagrangian submanifold that is contained in a single piece of $\mathfrak{X}$ and does not intersect relative divisors. The Floer cochains and the brane structure on L are defined as in the unbroken case :

$$CF^{\text{geom}}_{\text{brok}}(L) := \bigoplus_{d\in\mathbb{Z}_2} CF^d_{\text{brok}}(L), \quad CF^d_{\text{brok}}(L) := \bigoplus_{x\in\mathcal{I}_d(L)} \Lambda_{\geq 0}\langle x\rangle.$$

DEFINITION 10.3. (Composition maps for the broken Fukaya algebra) For admissible perturbation data $(\mathfrak{p}_\Gamma)_\Gamma$ define

$$m^{\text{brok}}_{d(\circ)} : (CF^{\text{geom}}_{\text{brok}}(L))^{\otimes d(\circ)} \to CF^{\text{geom}}_{\text{brok}}(L)$$

on generators by

$$m^{\text{brok}}_{d(\circ)}(x_1, \ldots, x_{d(\circ)}) = \sum_{x_0, u \in \mathcal{M}^{\text{brok}}_{\Gamma}(\mathfrak{X}, L, D, \underline{x})_0} w(u) x_0 \tag{10.14}$$

where the combinatorial type Γ of the broken map u ranges over all rigid types (see Definition 8.41) with $d(\circ)$ boundary inputs, and $w(u)$ is as in (10.10). The orientation sign $\epsilon(u) \in \{\pm 1\}$, which is factor in $w(u)$, is determined as in Remark 6.30.

THEOREM 10.4. (A_∞ algebra for a Lagrangian in a broken manifold) *For any admissible perturbation system* $\underline{\mathfrak{p}} = (\mathfrak{p}_\Gamma)_\Gamma$ *on the broken manifold* $\mathfrak{X}$*, the maps* $(m^{\text{brok}}_{d(\circ)})_{d(\circ) \geq 0}$ *on* $CF^{\text{geom}}_{\text{brok}}(L)$ *satisfy the axioms of a convergent (possibly curved)* A_∞ *-algebra* $CF^{\text{geom}}_{\text{brok}}(L)$*.*

The A_∞ relation follows from the description of the true boundary of the moduli space, see Proposition 8.44. The combinatorial factors arising from the distribution of interior markings are accounted exactly as in the unbroken case in Theorem 10.2.

10.3. Homotopy units

In the Fukaya algebra constructed in the previous section, a homotopy unit construction can be applied to produce a strictly unital A_∞ -algebra. Recall that the Morse function $F : L \to \mathbb{R}$ used in the construction of $CF^{\text{geom}}(L)$ is assumed to have a unique maximum point denoted $x^{\blacktriangledown} \in \text{crit}(F)$. In an idealized situation where domain-dependent perturbations are not required, $\langle x^{\blacktriangledown} \rangle$ is a strict unit for $CF^{\text{geom}}(L)$. This is because a boundary marked point mapping to the unstable locus of $x^{\blacktriangledown}$ is an empty constraint, and such a marking can be forgotten without affecting the disk. In our setting, marked points can not be forgotten because domain-dependent perturbations depend on them. The homotopy unit construction is a way of enhancing the Fukaya algebra so that the perturbation system admits forgetful maps.

We outline the idea of the homotopy units construction: For the Fukaya algebra to have a unit, we would like the perturbations on the Morse function to be domain-independent on the treed segments asymptotic to the maximum point $x^{\blacktriangledown} \in \text{crit}(F)$. However, we can not impose such a condition on the perturbation, because the perturbation depends only on the 'domain' which does not include the information about which critical point a Morse trajectory asymptotes to. One can not include the critical point label into the domain data, because Morse trajectories may break in the limit, and the label at the breaking point is not known beforehand. Therefore, we add a new kind of boundary marking to the domain, which is 'forgettable', and will serve as a unit. The treed segment at a forgettable leaf is required to asymptote to the maximum of the Morse function F, but is labelled $x^{\triangledown}$ (to distinguish it from $x^{\blacktriangledown}$). We also add 'weighted boundary markings', labelled $x^{\blacktriangledown}$, to the domain which give us a way of homotoping between treed segments with domain-dependent perturbations asymptoting to $x^{\blacktriangledown}$ and those with domain-independent perturbations asymptoting to $x^{\triangledown}$. The data of the domain now includes the information about whether a leaf is forgettable, weighted or unforgettable. This new kind of domain is called a *weighted treed disk* and is defined below after the statement of the main Theorem.

THEOREM 10.5. (Homotopy unit construction) *Suppose $\mathfrak{p}$ is a coherent perturbation datum for treed holomorphic disks, and suppose $CF^{\text{geom}}(L) := CF^{\text{geom}}(L, \mathfrak{p})$ is the A_∞-algebra whose composition maps count $\mathfrak{p}$-adapted disks. Then there exists a strictly unital A_∞ structure on the vector space*

$$CF(L) := CF^{\text{geom}}(L) \oplus \Lambda x^{{\color{gray}\blacktriangledown}}[1] \oplus \Lambda x^{\triangledown}, \tag{10.15}$$

with gradings

$$|x^{\triangledown}| = 0, \quad |x^{{\color{gray}\blacktriangledown}}| = -1,$$

whose composition maps count $\tilde{\mathfrak{p}}$-adapted weighted treed holomorphic disks (see Definition 10.7), where $\tilde{\mathfrak{p}}$ is an extension of the perturbation datum $\mathfrak{p}$ to weighted disks; and in the resulting A_∞ -algebra,

- *(a) $x^{\triangledown}$ is a strict unit,*
- *(b) $CF^{\text{geom}}(L) \subset CF(L)$ is a A_∞ sub-algebra,*
- *(c) and*

$$m_1(x^{{\color{gray}\blacktriangledown}}) = x^{\blacktriangledown} - x^{\triangledown} \mod \Lambda_{>0}.$$

The theorem is proved later in the section after defining weighted treed holomorphic disks.

The condition that $x^{\triangledown}$ is a strict unit determines all A_∞ structure maps involving occurrences of $x^{\triangledown}$. In the following geometric construction of a homotopy unit, the axioms are designed keeping this fact in mind.

DEFINITION 10.6. (a) (Weightings) A *weighting* of a treed disk $C = S \cup T$ of type Γ, with $S \neq \emptyset$, consists of a partition of the boundary semi-infinite edges

$$\text{Edge}^{\blacktriangledown}(\Gamma) \sqcup \text{Edge}^{{\color{gray}\blacktriangledown}}(\Gamma) \sqcup \text{Edge}^{\triangledown}(\Gamma) = \text{Edge}_{\circ,\to}(\Gamma)$$

into *unforgettable* resp. *weighted* resp. *forgettable*, and a *weight* on semi-infinite edges $\rho : \text{Edge}_{\circ,\to}(\Gamma) \to [0, \infty]$ satisfying

$$\rho(e) \in \begin{cases} \{0\} & e \in \text{Edge}^{\blacktriangledown}(\Gamma) \\ [0, \infty] & e \in \text{Edge}^{{\color{gray}\blacktriangledown}}(\Gamma) \\ \{\infty\} & e \in \text{Edge}^{\triangledown}(\Gamma). \end{cases}$$

The weighting ρ satisfies the following axiom:

(Outgoing edge axiom) A disk output $e_0 \in \text{Edge}_\circ(\Gamma)$ can be weighted only if the disk has exactly one weighted input $e_1 \in \text{Edge}^{{\color{gray}\blacktriangledown}}(\Gamma)$, all the other inputs $e_i \in \text{Edge}(\Gamma), i \neq 1$ are forgettable, and there are no interior leaves, $\text{Edge}_\bullet(\Gamma) = \emptyset$. In this case, the output e_1 has the same weight $\rho(e_1) = \rho(e_0)$ as the weighted input e_0. A disk output e_0 can be forgettable only if all the inputs are forgettable, and there are no interior leaves. In all the other cases, the output of a disk is unforgettable.

In the exceptional case that the treed disk C is an infinite tree segment and does not have surface components, the only possible labels are

$$\color{gray}{\blacktriangledown} \color{black} \to \triangledown, \quad \color{gray}{\blacktriangledown} \color{black} \to \blacktriangledown, \quad \text{or} \quad \blacktriangledown \to \blacktriangledown.$$

In the first two cases, the input has weight $\rho(e)$ equal to ∞ resp. 0.

(b) (Stability) A weighted treed disk $C = S \cup T$ with $S \neq \emptyset$ is *stable* if C is stable as a treed disk. In case $S = \emptyset$ and C is an infinite segment, then C is stable iff the labels are ${\color{gray}\blacktriangledown} \to \triangledown$ or ${\color{gray}\blacktriangledown} \to \blacktriangledown$.

(c) (Isomorphism) Two weighted treed disks C and C' are isomorphic if there is an isomorphism of treed disks $\phi : C \to C'$, the edge labels are identical, and the following is true.

 (i) If the output edge is not weighted, then the weights on the inputs of C_1 and $\phi(C_1)$ are equal;

 (ii) if the output edge e_0 is weighted, then the weights on the inputs are equal up to scalar multiplication, i.e.

$$\exists \lambda \in (0, \infty) : \ \forall e \in \operatorname{Edge}_{\circ,\to}(C) \backslash \{e_0\} \quad \rho(e) = \lambda \rho'(\phi(e)). \tag{10.16}$$

Consequently, if the output edge is weighted, since there is exactly one incoming edge by the (Outgoing edge axiom), the weights do not matter.

(d) (Combinatorial type) The *type* of a weighted treed disk is given by the type of the treed disk, and the labels $\{\triangledown, {\color{gray}\blacktriangledown}, \blacktriangledown\}$ at the inputs and outputs, and whether the weight at any vertex is zero, infinite or neither. Thus the combinatorial type of a weighted treed disk includes a partition of weighted edges

$$\operatorname{Edge}^{\color{gray}\blacktriangledown}(\Gamma) = \operatorname{Edge}_0^{\color{gray}\blacktriangledown}(\Gamma) \cup \operatorname{Edge}_{(0,\infty)}^{\color{gray}\blacktriangledown}(\Gamma) \cup \operatorname{Edge}_\infty^{\color{gray}\blacktriangledown}(\Gamma)$$

into edges of weight 0, non-zero finite, and infinity.

The moduli space $\mathcal{M}_\Gamma$ of weighted treed disks can be identified with

$$\begin{cases} \mathcal{M}_{\Gamma'} \times [0, \infty]^{|\operatorname{Edge}^{\color{gray}\blacktriangledown}(\Gamma)|}, & \text{if the output label is not } {\color{gray}\blacktriangledown} \\ \mathcal{M}_{\Gamma'}, & \text{if the output edge is } {\color{gray}\blacktriangledown}, \end{cases}$$

where Γ' is the type of treed disk obtained by forgetting the weighting. If the type Γ is ${\color{gray}\blacktriangledown} \to \triangledown$ resp. ${\color{gray}\blacktriangledown} \to \blacktriangledown$, then $\mathcal{M}_\Gamma$ is a point.

The (Cutting edges) morphism has some additional features for weighted treed disks. Given a type Γ of a weighted treed disk, suppose Γ_+, Γ_- (here Γ_+ contains the root of Γ) are the treed disk types produced by cutting an edge $e \in \operatorname{Edge}_{\circ,-}(\Gamma)$ in Γ, and $e_\pm \in \operatorname{Edge}_{\circ,\to}(\Gamma_\pm)$ be the pair of new edges created by the cutting. The label ($\blacktriangledown$, ${\color{gray}\blacktriangledown}$ or $\triangledown$) and the weight at e_+ are the same as that of e_-, and the label and weight at e_- is determined by the (Outgoing edges axiom) applied to Γ_-. In this case there is a new type of (Cutting edges) morphism, wherein we cut a weighted incoming edge of weight 0 or ∞ (although it is not a broken edge):

(Cutting a weighted input edge) Suppose $e \in \operatorname{Edge}^{\color{gray}\blacktriangledown}(\Gamma)$ is an input, and $\rho(e) = 0$ resp. ∞. Cutting e produces two types : Γ_- is an infinite segment ${\color{gray}\blacktriangledown} \to \blacktriangledown$ resp. ${\color{gray}\blacktriangledown} \to \triangledown$, and Γ_+ is Γ with e as an unforgettable resp. forgettable edge, see Figure 10.1.

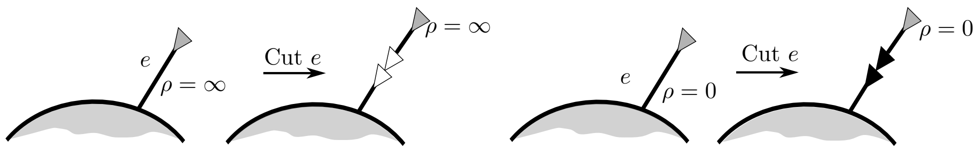

FIGURE 10.1. Cutting a weighted input edge

Perturbation datum defined on the moduli space of weighted treed disks is coherent with respect to (Collapsing of edges, making an edge length finite or non-zero), (Cutting of edges) and the (Locality axiom). The perturbation vanishes on infinite segments ${\color{gray}\blacktriangledown} \to \blacktriangledown$ and ${\color{gray}\blacktriangledown} \to \triangledown$. Additionally the following coherence conditions are satisfied:

(a) (Making an edge weight finite/non-zero) If the weighted treed disk type Γ is obtained from Γ' by making the weight of an edge $e \in \operatorname{Edge}^{\color{gray}\blacktriangledown}(\Gamma')$ finite or non-zero, then $\mathfrak{p}_{\Gamma'}$ is the pullback of $\mathfrak{p}_\Gamma$ by the inclusion of the universal moduli space $\overline{\mathcal{U}}_{\Gamma'} \to \overline{\mathcal{U}}_\Gamma$.

(b) (Forgetting edges) Suppose e is an input edge in a weighted treed type Γ that is either forgettable or weighted with infinite weight, and Γ' is the type obtained by forgetting e. Then $\mathfrak{p}_\Gamma$ is the pullback of $\mathfrak{p}_{\Gamma'}$.

DEFINITION 10.7. (Adapted weighted treed holomorphic disks) Let $\underline{\mathfrak{p}} = (\mathfrak{p}_\Gamma)_\Gamma$ be a coherent perturbation datum for weighted treed disks. An $\mathfrak{p}_\Gamma$-*adapted weighted treed holomorphic disk* is a map $u : C \to X$ that is adapted in the sense of treed holomorphic disks, and additionally satisfies the following:

(Label axiom) A treed input or output segment labelled ${\color{gray}\blacktriangledown}$ resp. $\triangledown$ asymptotes to the maximum point $x^{\blacktriangledown} \in \operatorname{crit}(F)$.

An adapted weighted holomorphic disk $u : C \to X$ is *stable* if for any component of the domain C_v, either the map u_v is non-constant, or the domain C_v is stable in the sense of weighted treed disks. The new feature of stability for weighted holomorphic disks is that a stable map u may be constant on an infinite tree segment labelled ${\color{gray}\blacktriangledown} \to \blacktriangledown$ or ${\color{gray}\blacktriangledown} \to \triangledown$. The *combinatorial type* of a weighted treed holomorphic disk consists of the type of the domain weighted treed disk, together with the combinatorial type data of the treed holomorphic map, namely homology classes of the maps on surface components, and intersection multiplicities with the stabilizing divisor at the markings. The type Γ of a weighted treed holomorphic disk is *rigid* if all the edges $e \in \operatorname{Edge}_-(\Gamma)$ are boundary edges with finite non-zero length, and in case of weighted inputs or output, the weight $\rho(e)$ is finite and non-zero. This ends the Definition.

The expanded set of labels on the ends of treed segments is denoted by

$$\hat{\mathcal{I}}(L) := \mathcal{I}(L) \cup \{x^{\color{gray}\blacktriangledown}, x^{\blacktriangledown}\}, \tag{10.17}$$

where we recall from (6.2) that $\mathcal{I}(L)$ is the set of critical points of the Morse function on the Lagrangian L.

Proposition 10.8. (Transversality for weighted treed holomorphic disks) *Given a regular coherent perturbation datum* $\mathfrak{p}$ *for treed holomorphic disks,* $\mathfrak{p}$ *extends to a regular perturbation datum* $\tilde{\mathfrak{p}}$ *on weighted treed holomorphic disks such that the following holds: For an uncrowded type* Γ *of weighted holomorphic maps, and a prescribed tuple of inputs* $\underline{x} := (x_1, \ldots, x_{d(\circ)}) \in \hat{\mathcal{I}}(L)^{d(\circ)}$ *and an output* $x_0 \in \hat{\mathcal{I}}(L)$ *respecting the (Label axiom) and for which* $i(\Gamma, \underline{x}) \leq 1$*, the moduli space*

$$\mathcal{M}_\Gamma(L, \underline{\mathfrak{p}}, \underline{x})$$

of weighted $\underline{\mathfrak{p}}$*-adapted treed disks with limits* $\underline{x}$ *is a smooth manifold of expected dimension. The moduli space* $\mathcal{M}_\Gamma(L, \underline{\mathfrak{p}}, \underline{x})$ *is compact if* $i(\Gamma, \underline{x}) = 0$*, and if* $i(\Gamma, \underline{x}) = 1$*,* $\mathcal{M}_\Gamma(L, \underline{\mathfrak{p}}, \underline{x})$ *has a compactification whose boundary consists of configurations with a boundary node with length* 0 *or* ∞*, or a weighted edge with weight* 0 *or* ∞*.*

Proof of Proposition 10.8. We extend the perturbation datum $\underline{\mathfrak{p}} = \{\mathfrak{p}_\Gamma\}_\Gamma$ on domain curves that only contain unforgettable leaves to those that have other kinds of leaves as well. We first consider domain types whose output is unforgettable. We consider the case when the domain type Γ has a single weighted/forgettable input leaf e, since the other cases follow inductively. Let Γ_0 resp. Γ_∞ be the domain type obtained by making the weight of e in Γ zero resp. infinity. By the (Cutting a weighted edge) morphism, on the subset $\{\rho(e) = 0\}$ resp. $\{\rho(e) = \infty\}$ of $\mathcal{M}_\Gamma$, the perturbation $\tilde{\mathfrak{p}}_\Gamma$ is given by $\tilde{\mathfrak{p}}_{\Gamma_0}$ resp. $\tilde{\mathfrak{p}}_{\Gamma_\infty}$. The perturbation $\tilde{\mathfrak{p}}_{\Gamma_0}$ is equal to $\mathfrak{p}_{\Gamma_0}$. The perturbation $\tilde{\mathfrak{p}}_{\Gamma_\infty}$ is defined via the (Forgetting edges) axiom, which means that the perturbation is independent of forgettable treed segment on the domain curve. By standard arguments, the perturbation $\tilde{\mathfrak{p}}_\Gamma$ can be extended over $\{\rho(e) \in (0, \infty)\}$ while satisfying regularity.

For domain curves whose output is weighted or forgettable, we define the perturbation to be domain-independent. Indeed, on such domains maps are constant and lie in the maximum point in $\mathrm{crit}(F) \subset L$; and strata with only unforgettable leaves and root do not occur in the compactification of strata where the output is forgettable/weighted.

Next we prove the statement on compactification. We analyze the cases where a sequence of weighted maps of type Γ that are non-constant on either the surface or tree part converges to a limit that has a breaking on a weighted/forgettable leaf e, and the map is non-constant on both segments incident at the breaking. The other cases are covered either by the compactification result for treed holomorphic maps without weights; or, if the maps in the sequence are constant, by straightforward formal arguments.

If e is weighted, for dimension reasons, the infinite segment $(-\infty, \infty)$ in the broken segment maps to the maximum point in L, and therefore, the map is constant on this segment. If e is forgettable, by the (Forgetting edges) axiom, the only case to be considered is when $i(\Gamma, \underline{x}) = 1$ and Γ is of the form shown in Figure 10.2. The map is constant on the surface, and the Morse index of x_0, $x_1 \in \mathrm{crit}(F)$ differ by 1. We may assume the domain-dependent Morse perturbation on the unforgettable edges is small enough that the ends of the moduli space are given by the breaking of the unforgettable input or output edge as in Figure 10.3 below. This rules out the possibility of breaking on the forgettable edge. □

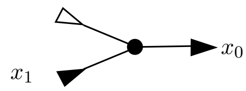

FIGURE 10.2. A type of map

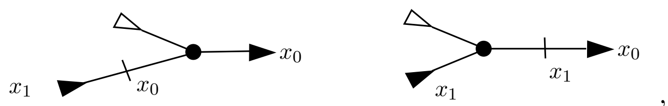

FIGURE 10.3. Breaking an edge

For an admissible perturbation datum $\underline{\mathfrak{p}} = (\mathfrak{p}_\Gamma)_\Gamma$ for weighted treed types define composition maps on the generators of the Fukaya algebra

$$CF(L, \underline{\mathfrak{p}}) := CF^{\text{geom}}(L, \underline{\mathfrak{p}}) \oplus \Lambda\langle x^{\blacktriangledown}\rangle \oplus \Lambda\langle x^{\triangledown}\rangle$$

as

$$m_{d(\circ)}(x_1, \ldots, x_{d(\circ)}) = \sum_{x_0, u \in \overline{\mathcal{M}}_\Gamma(L, \underline{\mathfrak{p}}, \underline{x})_0} w(u) x_0 \tag{10.18}$$

where

$$w(u) = (-1)^{\heartsuit} (d(\bullet)!)^{-1} y(u) q^{A(u)} \epsilon(u), \quad \heartsuit = \sum_{i=1}^{d(\circ)} i|x_i|,$$

and the sum ranges over all rigid types Γ, and the weight for any weighted input is finite non-zero in the type Γ.

We now prove the main Theorem of this section.

PROOF OF THEOREM 10.5. For a one-dimensional moduli space of weighted maps of rigid type, the true boundary strata contain one of the following configurations : a configuration with a weight 0 or ∞ at a weighted input which is equivalent to the broken configuration in Figure 10.1, a boundary node with a broken segment, or a broken Morse trajectory. These configurations exactly correspond to the terms in the A_∞ -associativity relations, and so $CF(L)$ is an A_∞ -algebra. The geometric part $CF^{\text{geom}}(L)$ is a sub-algebra because if the inputs to a treed disk are unforgettable, the output is also unforgettable.

The element $x^{\triangledown}$ is a strict unit for the following reasons. For $d(\circ) > 2$, we have

$$m_{d(\circ)}(\ldots, x^{\triangledown}, \ldots) = 0$$

because the input $x^{\triangledown}$ is an empty constraint, and can be forgotten because the perturbation satisfies the (Forgetting edges) axiom. The term $m_1(x^{\triangledown})$ is also zero : any disk that is counted has interior markings, and therefore, placing the marked point $x^{\triangledown}$ adds one to the dimension, and therefore the moduli space is not zero-dimensional. Finally, by the same argument, $m_2(x^{\triangledown}, y)$ and $m_2(y, x^{\triangledown})$ do not count any disk with interior markings. The only contributions are from constant disks. We conclude that both terms are equal to $\pm y$, for any generator y. $\square$

Remark 10.9. (Leading order term in the first composition map) Constant trajectories $\blacktriangledown \to \triangledown$, $\blacktriangledown \to \blacktriangledown$ contribute to the first composition map $m_1(x^{\blacktriangledown})$, and the choice of orientation [**18**, Remark 4.23] implies

(10.19)
$$m_1(x^{\blacktriangledown}) \ = \ x^{\triangledown} - x^{\blacktriangledown} + \sum_{x_0,[u]\in\overline{\mathcal{M}}_\Gamma(L,D,x^{\blacktriangledown},x_0)_0,E(u)>0} (-1)^{\heartsuit}(d(\bullet)!)^{-1}q^{E(u)}\epsilon(u)y(u)x_0.$$

This formula is similar to that in Fukaya-Oh-Ohta-Ono [**34**, (3.3.5.2)].

10.4. Quilted disks

Morphisms between Fukaya algebras are defined by counts of quilted holomorphic disks. These are ordinary holomorphic disks with perturbed almost complex structures, whose domain disks have an additional structure of a 'quilting'.

10.4.1. Quilted disks.

Definition 10.10. (Quilted disks)

(a) A *quilted disk* is a disk $S \simeq \mathbb{D}^2 \subset \mathbb{C}$ with boundary markings $z_0^\circ, \dots, z_{d_\circ}^\circ \in \partial S$, interior markings $z_0^\bullet, \dots, z_{d_\circ}^\bullet \in S^\circ$ and a biholomorphism

$$\phi : (S \backslash \{z_0\}, \partial S) \xrightarrow{\phi} (\mathbb{H}, \partial\mathbb{H}),$$

which we call an *affine structure*. Two affine structures ϕ_0, ϕ_1 are equivalent if there exists $\xi \in \mathbb{R}$ such that $\phi_1(z) = \phi_0(z) + \xi$. Two quilted disks $(S_0, \underline{z}^0, \phi_0)$, $(S_1, \underline{z}^1, \phi_1)$ are *isomorphic* if there exists $\xi \in \mathbb{R}$ such that, defining $\tau(z) := z + \xi$, the biholomorphism $\phi_1^{-1} \circ \tau \circ \phi_0 : S_0 \to S_1$ maps each marking in S_0 to the corresponding marking in S_1.

(b) (Treed quilted disk) A *treed quilted disk* is a treed nodal disk C, a subset of whose disk components are quilted disks and have an affine structure. The subset of quilted disks satisfies the following:
- A path in C from any boundary (non-root) leaf z_e, $e \in \mathrm{Edge}_{\circ,\to} \backslash \{0\}$ to the root z_0 intersects exactly one quilted disk.
- (Balancing condition) The sum of lengths of treed segments $\mp\ell(e)$ on the path $\gamma \subset C$ connecting any two quilted disks, with sign $\pm$ at e determined by whether the orientation of the path agrees with the orientation on the tree, is zero.

(c) (Combinatorial types) The combinatorial type Γ of a quilted treed disk is the combinatorial type of the treed disk (without the quilting datum) together with the subset

$$\mathrm{Vert}_\circ^{\mathrm{col}}(\Gamma) \subseteq \mathrm{Vert}_\circ(\Gamma)$$

of vertices corresponding to quilted disks, which are called *colored vertices*.

(d) (Stability) A treed quilted disk $C = S \cup T$ is stable if the automorphism group of any surface component $S_v, v \in \mathrm{Vert}(\Gamma)$ is trivial, and there are no treed segments both whose ends are infinite. This means a quilted disk is stable if it has at least two special points.

Definition 10.11. (Distance from the seam function) Let C be a treed quilted disk. For any point $z \in C$, let γ_z be a non self-intersecting path in C that connects z to a point in the quilted disk of C that is closest to z. Define the *distance from the seam* as the sum of lengths of boundary edges lying on γ_z, that is,

$$d(z) := \pm \sum_{e \in \mathrm{Edge}_{\circ,-}(\Gamma): T_e \subset \gamma_z} \ell(e) \in [-\infty, \infty], \tag{10.20}$$

where the sign is $+$ resp. $-$ if z is above resp. below the quilted disk components (that is, further from resp. closer to the root than the quilted disk components). Note that d is constant on the connected surface components of C, or in other words, d is constant on the connected components of $C \setminus \cup_{e \in \mathrm{Edge}_\circ, \ell(e)>0} T_e$. The distance function d is zero on a quilted disk component and the spheres attached to them.

The moduli space of stable quilted disks with interior and boundary markings is a compact cell complex. As the interior and boundary markings go to infinity, they bubble off onto either quilted disks or unquilted disks or spheres. The case of combined boundary and interior markings is a straight-forward generalization of Ma'u-Woodward [**58**]. In the case with no interior markings, the moduli spaces are cell complexes called the 'multiplihedra' introduced by Stasheff [**85**].

Let $\overline{\mathcal{M}}^q_{d(\bullet),d(\circ)}$ denote the moduli space of stable marked quilted treed disks with $d(\circ)$ boundary leaves and $d(\bullet)$ interior leaves. See Figure 10.4 for a picture of $\overline{\mathcal{M}}^q_{0,2}$. The quilted disks $S_v \subset S, v \in \mathrm{Vert}^{\mathrm{col}}(\Gamma)$ are those with two shadings; while the ordinary disks $S_v, v \notin \mathrm{Vert}^{\mathrm{col}}(\Gamma)$ have either light or dark shading depending on whether the distance from the seam function is positive or negative. The hashes on the line segments T_e indicate breakings.

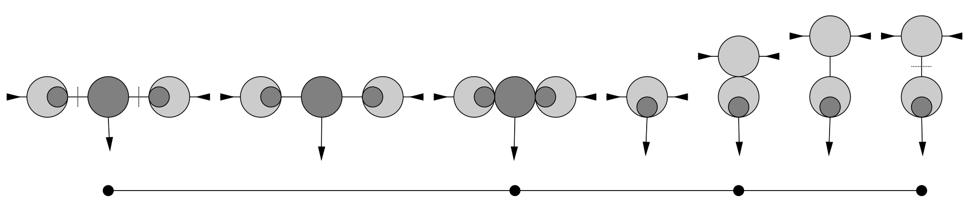

Figure 10.4. Moduli space $\overline{\mathcal{M}}^q_{0,2}$ of stable quilted treed disks

Remark 10.12. (A quilting circle) A quilted disk may alternately defined as a marked complex disk $(S, \underline{z})$ together with a *seam*, which is a circle $Q \subset S$ tangent to the 0-th boundary marking z_0, see Ma'u-Woodward [**58**, Definition 4.1]. A choice of a seam is equivalent to an affine structure, by taking the seam to be $Q = \{\mathrm{Im}(z) = 1\}$. This viewpoint explains our representation of quilted disks in Figures 10.4 and 10.5. An interior point in a quilted disk component may lie above, below or on the seam;

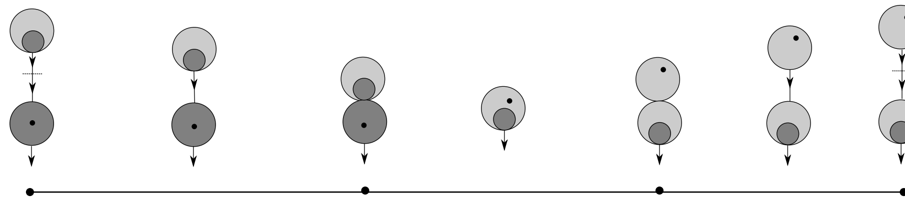

FIGURE 10.5. Moduli space $\overline{\mathcal{M}}^q_{1,0}$ of stable quilted treed disks

the relative position of interior points with respect to the seam does not affect the combinatorial type of the treed quilted disk.

We use the quilting circle viewpoint to justify the compactification of an open stratum in $\mathcal{M}^q_{1,0}$. Let Γ be a type of quilted disk with a single interior marking, a boundary output and no boundary inputs, and no boundary nodes (fourth figure from the left in Figure 10.5). The stratum $\mathcal{M}^q_\Gamma \subset \mathcal{M}^q_{1,0}$ is an open interval $(0,\infty)$ parametrizing disks with markings

$$z^\circ_{0,\nu} = \infty, \quad z^\bullet_{0,\nu} = i\nu, \quad \nu \in (0,\infty), \tag{10.21}$$

and quilting circle $\{\mathrm{Im}(z) = 1\}$, where the disk is identified to $\mathbb{H} \cup \{\infty\}$. In the limit $\nu \to \infty$ resp. $\nu \to 0$, we obtain the strata to the left resp. right of Γ in Figure 10.5. In particular, in the limit $\nu \to \infty$, the limit of the markings in (10.21) gives the quilted disk in the limit, and the unquilted dark disk is given by the reparametrization $z \mapsto \frac{z}{\nu}$, so that the markings in the sequence are $z^\circ_{0,\nu} = \infty$, $z^\bullet_{0,\nu} = i$, and the quilting circle is $\{\mathrm{Im}(z) = \frac{1}{\nu}\}$, giving rise to a limit disk which is fully in the dark region.

As in the unquilted case, the top-dimensional cells in $\mathcal{M}^q_{d(\bullet),d(\circ)}$ consist of strata $\mathcal{M}_\Gamma$ which do not have any interior edges, that is, $\mathrm{Edge}_{\bullet,-}(\Gamma) = \emptyset$, and all boundary edges $e \in \mathrm{Edge}_{\circ,-}(\Gamma)$ have finite non-zero length, and the dimension of these strata is equal to

$$d(\circ) + 2d(\bullet) - 1.$$

The true boundary of $\mathcal{M}^q_{d(\bullet),d(\circ)}$ consists of strata $\mathcal{M}_\Gamma$ where

- Γ either has a single broken edge $e \in \mathrm{Edge}_\circ(\Gamma)$,
- or Γ has a collection of broken edges, such that the types $\Gamma_0, \ldots, \Gamma_k$ obtained by disconnecting Γ at the breakings consist of an unquilted disk type Γ_0, all whose vertices have negative distance from the seam in Γ, and a collection $\Gamma_1, \ldots, \Gamma_k$ of quilted disk types.

For example, the combinatorial types in Figure 10.6 occur in the codimension one stratum of $\mathcal{M}^q_{0,4}$, $\mathcal{M}^q_{0,5}$.

10.4.2. Morphisms of types. Morphisms of graphs (Cutting an edge, collapsing edges, making edge lengths finite or non-zero) induce morphisms of moduli spaces of stable quilted treed disks as in the unquilted case. The new feature is that (Cutting an edge) is done such that one of the pieces is quilted and the other unquilted. This implies that output edges of quilted disks are cut simultaneously, and therefore the output has a disconnected type.

For example, in Figure 10.6, in the left picture, one can cut the e at the breaking to obtain an unquilted disk with positive distance from the seam, and a quilted disk. In the picture to the right, the edges e_1 and e_2 get cut simultaneously to yield an unquilted disk with negative distance from the seam and a disconnected type consisting of two quilted disks.

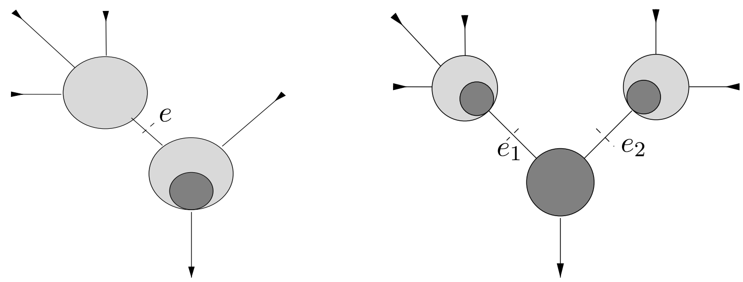

FIGURE 10.6. A quilted treed disk with edges of infinite length.

In a similar vein, the morphisms (Collapsing an edge) and (Making an edge length finite or non-zero) may involve several edges instead of a single one. For example, in the quilted disk with three boundary edges of length zero shown in Figure 10.7, there is no (Collapsing an edge) or (Making an edge length finite)

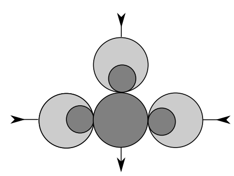

FIGURE 10.7. A quilted disk with length zero edges.

morphism for one of the edges alone; these operations can only be performed for all three edges simultaneously in order to respect the (Balancing condition) in Definition 10.10 (b).

For any combinatorial type Γ of quilted disk there is a *universal quilted treed disk* $\overline{\mathcal{U}}_\Gamma \to \overline{\mathcal{M}}_\Gamma$ which is a cell complex whose fiber over $[C]$ is isomorphic to C. The universal disk splits into surface and tree parts $\overline{\mathcal{U}}_\Gamma = \overline{\mathcal{S}}_\Gamma \cup \overline{\mathcal{T}}_{\circ,\Gamma} \cup \overline{\mathcal{T}}_{\bullet,\Gamma}$, where the last two sets are the boundary and interior parts of the tree respectively

10.4.3. Weightings. Weights can be added to the inputs and output of quilted treed disks as in the case of treed disks. We suppose there is a partition of the boundary markings

$$\mathrm{Edge}^{\blacktriangledown}(T) \sqcup \mathrm{Edge}^{\triangledown}(T) \sqcup \mathrm{Edge}^{\blacktriangledown}(T) = \mathrm{Edge}_{\circ,\to}(T)$$

into *weighted* resp. *forgettable* resp. *unforgettable* edges as in the unquilted case. The outgoing edge axiom is the same as in the unquilted case. In the quilted case, the trees in Figure 10.8 are stable. Isomorphism of weighted quilted disks is the same as the unquilted case, and therefore, the moduli space with a single weighted leaf and no markings is a point.

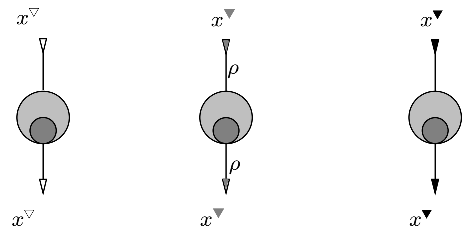

FIGURE 10.8. Unmarked stable treed quilted disks

10.4.4. Orientations. Orientations of the moduli space of quilted treed disks are defined as follows. Each main stratum of $\overline{\mathcal{M}}^q_{d(\bullet),d(\circ)}$ can be oriented using the isomorphism of the stratum made of quilted treed disks having a single disk with $\mathbb{R}$ times $\overline{\mathcal{M}}_{d(\bullet),d(\circ)}$, the extra factor corresponding to the quilting parameter. The boundary of the moduli space is naturally isomorphic to a union of moduli spaces:

$$\partial\overline{\mathcal{M}}^q_{d(\bullet),d(\circ)} \cong \bigcup_{\substack{i,j \\ I_1\cup I_2=[d(\bullet)]}} \left(\overline{\mathcal{M}}^q_{|I_1|,d(\circ)-i+1} \times \overline{\mathcal{M}}_{|I_2|,i}\right) \cup \bigcup_{\substack{r\geq 1, I_1\cup\ldots I_r=[d(\bullet)],\\ m_1+\cdots+m_r=d(\circ)}} \left(\overline{\mathcal{M}}_{|I_0|,r} \times \prod_{j=1}^r \overline{\mathcal{M}}^q_{|I_j|,m_j}\right). \tag{10.22}$$

Here $[d(\bullet)] := \{1,\ldots,d(\bullet)\}$. In the first union j is the index of the attaching leaf in the quilted tree and so ranges from 1 to $n-i+1$, and the union ranges over all partitions of the set $[d(\bullet)]$ into I_1, I_2. The second union ranges over partitions of $[d(\bullet)]$ into $I_0,\ldots,I_r$. By construction, for the facet of the first type, the sign of the inclusions of boundary strata are the same as that for the corresponding inclusion of boundary facets of the moduli space $\overline{\mathcal{M}}_{d(\bullet),d(\circ)}$ of unquilted disks, that is, $(-1)^{i(n-i-j)+j}$. For facets of the second type, the gluing map is

$$(0,\infty)\times\mathcal{M}_{r,m_0}\times\bigoplus_{j=1}^r \mathcal{M}^q_{|I_j|,m_j} \to \mathcal{M}^q_{d(\bullet),d(\circ)}$$

given for boundary markings by

$$(\delta, x_1,\ldots,x_r,(x_{1,j}=0,x_{2,j},\ldots,x_{m_j,j})_{j=1}^r) \mapsto (x_1, x_1+\delta^{-1}x_{2,1},\ldots,x_1+\delta^{-1}x_{m_1,1},\ldots,x_r,x_r+\delta^{-1}x_{2,r},\ldots,x_r+\delta^{-1}x_{m_r,r}). \tag{10.23}$$

This map views the markings as lying in the affine half plane $\mathbb{H}\subset\mathbb{C}$; an interior $x\in I_j$ is mapped to $x_j+\delta^{-1}x$; and the map is well-defined for δ that is large enough to ensure that the ordering of the boundary markings is preserved. This map changes orientations by $\sum_{j=1}^r (r-j)(m_j-1)$; in case of non-trivial weightings, m_j should be replaced by the number of incoming markings plus non-trivial but finite weightings on the j-th component.

10.5. Quilted pseudoholomorphic disks

The Fukaya algebra of a Lagrangian submanifold is independent of the choice of perturbation data up to homotopy equivalence. We outline the proof in the case that the degree of the stabilizing divisors is fixed. One considers two systems of perturbations and extends them to a set of perturbations for the moduli space of quilted treed disks. A morphism is defined between the two A_∞ -algebras by counts of quilted treed holomorphic disks.

Proposition 10.13. *Suppose $\underline{\mathfrak{p}}^0$, $\underline{\mathfrak{p}}^1$ are regular perturbation data that are defined using stabilizing pairs (J^0, D^0) and (J^1, D^1), which are connected by a path of stabilizing pairs $\{(J^t, D^t)\}_{t\in[0,1]}$. There exists a coherent perturbation datum $\underline{\mathfrak{p}}^{01}$ which induces a unital A_∞ -morphism (as in (10.8))*

$$\phi_{01} : CF(L, \underline{\mathfrak{p}}^0) \to CF(L, \mathfrak{p}^1)$$

which is a homotopy equivalence and such that $\phi_{01}^0(1)$ has positive q-valuation, that is $\phi_{01}^0(1) \in \Lambda_{>0}\langle \hat{\mathcal{I}}(L)\rangle$.

Remark 10.14. (Convergent A_∞ morphisms) An A_∞-morphism $\mathcal{F} : A_0 \to A_1$ between A_∞ algebras with Novikov coefficients is said to be *convergent* if $\mathcal{F}^0$ has positive q-valuation. Proposition 10.13 thus shows that the A_∞ morphism ϕ is convergent. A convergent A_∞ morphism $\mathcal{F} : A_0 \to A_1$, induces a well-defined map $MC(\mathcal{F}) : MC(A_0) \to MC(A_1)$ on the space of Maurer-Cartan solutions, see [**18**, Lemma 5.2].

Proposition 10.13 is used in Section 10.6 to show that the Fukaya algebra of a Lagrangian in a neck-stretched manifold is independent of the neck length parameter up to homotopy equivalence. We present a partial proof of Proposition 10.13 later in the section after defining quilted holomorphic disks. We refer the reader to [**18**] for the definition and proof of homotopy equivalence. Our motivation in presenting a partial proof is that the proof technique is used again in Chapter 11. In this section, we only consider unbroken manifolds. A broken version of quilted holomorphic disks will be encountered in Chapter 11.

Definition 10.15. Given perturbation data

$$\underline{\mathfrak{p}}^0 = (J^0_\Gamma, F^0_\Gamma)_\Gamma, \quad \underline{\mathfrak{p}}^1 = (J^1_\Gamma, F^1_\Gamma)_\Gamma$$

on unquilted treed disks with respect to stabilizing divisors D^0 and D^1 that have the same degree, a *perturbation morphism* $\underline{\mathfrak{p}}^{01}$ from $\underline{\mathfrak{p}}^0$ to $\underline{\mathfrak{p}}^1$ for the type Γ of a quilted treed disk consists of

(a) a domain-dependent Morse function

$$F^{01}_\Gamma : \overline{\mathcal{T}}_{\circ,\Gamma} \to \mathbb{R}$$

is equal to the domain-independent Morse function F^0 resp. F^1 on the neighbourhood $\overline{\mathcal{T}}_{\circ,\Gamma} - \overline{\mathcal{T}}^{\mathrm{cp}}_{\circ,\Gamma}$ of the endpoints for which the distance to the seam d is $-\infty$ resp. ∞, where $F^k : L \to \mathbb{R}$ is the background Morse function for F^k_Γ for $k = 0, 1$;

(b) and a domain-dependent almost complex structure

$$J_\Gamma^{01} : \overline{\mathcal{S}}_\Gamma \to \{J \in \mathcal{J}(X) : J \text{ is } \omega\text{-tamed}\}$$

with the property that on the curve associated to any point $m \in \overline{\mathcal{M}}_\Gamma$, J_Γ^{01} is constant on $\overline{\mathcal{S}}_{\Gamma,m} - \overline{\mathcal{S}}^{\text{cp}}_{\Gamma,m}$, where the compact set $\overline{\mathcal{S}}^{\text{cp}}$ is as defined in Section 6.1. Further, denoting by Γ_0 resp. $\Gamma_1 \subset \Gamma$ the sub-tree where the distance from the seam d is $-\infty$ resp. ∞, J_Γ^{01} is equal to the complex structures $J^0_{\Gamma_0}$ resp. $J^1_{\Gamma_1}$ on the (unquilted) treed disk components of type Γ_0, Γ_1.

A collection of perturbation morphisms $\underline{\mathfrak{p}} = (\mathfrak{p}_\Gamma)_\Gamma$ defined on moduli spaces of quilted disks is *coherent* if it is compatible with the (Cutting of edges), (Making an edge length/weight finite or non-zero) and (Forgetting edges) morphisms on weighted quilted disk types, and satisfies the (Locality axiom) from Definition 6.5.

Definition 10.16. (Holomorphic quilted treed disk) Let $\underline{\mathfrak{p}} = (\mathfrak{p}_\Gamma)_\Gamma$ be a coherent perturbation morphism. Let

$$\delta : [-\infty, \infty] \to [0, 1] \tag{10.24}$$

be an increasing diffeomorphism, which is assumed to be fixed. A holomorphic quilted treed disk is a map $u : C \to X$ where C is a quilted disk of type Γ, and u is $\mathfrak{p}_\Gamma$-holomorphic in the sense of ordinary perturbed holomorphic disks (see Definition 6.10), and *adapted* to a family of divisors $\{D^t\}_{t\in[0,1]}$ in the sense that

- each the interior marking z_i maps to $D^{\delta\circ d(z_i)}$, where d is the distance-to-seam function and δ is from (10.24),
- and for each $t \in [0, 1]$, each component of $u^{-1}(D^t) \cap (\delta \circ d)^{-1}(t)$ contains a marking.

If the domain quilted treed disk C is unstable, we obtain a stable quilted treed disk C' by collapsing unstable surface and tree components, and we denote the collapsing map by $f : C \to C'$. Pulling back by f we obtain a datum on C, still denoted $(J_\Gamma^{01}, F_\Gamma^{01})$. A quilted holomorphic treed disk $u : C \to X$ is *stable* if every (surface or tree) component of C on which u is non-constant is stable in the sense of weighted quilted disks. The *type* of a quilted holomorphic disk consists of the combinatorial type of the domain quilted treed disk, and the tangency and homology data of the map components as in the type of a treed holomorphic disk, see Definition 6.15.

Remark 10.17. (Symplectic area and the number of markings) For a quilted treed disk $C = S \cup T$, on any surface component $S_v \subset S$, the function $\delta \circ d|S_v$ is constant and therefore a quilted holomorphic disk $u : C \to X$ is adapted to the divisor $D^{\delta\circ d(S_v)}$ on the component S_v. Further the divisor $D^{\delta\circ d(S_v)}$ is holomorphic with respect to the background almost complex structure $J^{\delta\circ d(S_v)}$, and therefore all intersections between the map $u|S_v$ and the divisor $D^{\delta\circ d(S_v)}$ are positive. The intersection number of u with D^δ is $k\omega[u]$, where the divisor D_0 (and D_1) is Poincaré dual to $k\omega$.

Remark 10.18. (Quilted holomorphic disks for a path of perturbations)

(a) Often a perturbation morphism is constructed using a generic path of perturbations $\mathfrak{p}^t_\Gamma = (J^t_\Gamma, F_\Gamma)$, $t \in [0, 1]$, whose Morse datum F_Γ is t-independent.

Let $\tilde{\Gamma}$ be a quilted disk type, for which forgetting the quilting yields the disk type Γ. One may define a perturbation morphism $\mathfrak{p}^{01}_{\tilde{\Gamma}}$ connecting $\mathfrak{p}^0_\Gamma$ and $\mathfrak{p}^1_\Gamma$ by setting the domain-dependent data to be

$$J^{01}_{\tilde{\Gamma}}(z) := J^{\delta\circ d(z)}_{\Gamma}(z), \quad F^{01}_{\tilde{\Gamma}}(z) := F_\Gamma(z) \tag{10.25}$$

for any $z \in \mathcal{S}_{\tilde{\Gamma}}$ where d is the 'distance from the seam' function from (10.20), and $\delta : [-\infty, \infty] \to [0, 1]$ is a fixed increasing diffeomorphism from (10.24). Thus on any surface component of a $J^{01}_{\tilde{\Gamma}}$-holomorphic quilted disk, the map is J^t_Γ-holomorphic for some $t \in [0, 1]$.

(b) (A path of regular perturbations) In some special cases, a one-dimensional component of quilted disks is made up of a family of unquilted $\mathfrak{p}_t$-holomorphic disks for $t \in [0, 1]$. Such a special case arises when the perturbation $\mathfrak{p}^t_\Gamma = (J^t_\Gamma, F_\Gamma)$ is regular for a disk homology class $\beta \in H_2(X, L)$ and input/output tuple $\underline{x} = (x_0, \ldots, x_d)$ for all t. Then the zero dimensional component of the moduli space of perturbed holomorphic maps $\mathcal{M}_\beta(\mathfrak{p}^t, \underline{x})_0$ is bijective to $\mathcal{M}_\beta(\mathfrak{p}^0, \underline{x})_0$ for any $t \in [0, 1]$. Further, for a perturbation morphism as in (10.25), a path of moduli spaces $\cup_{t\in[0,1]}\mathcal{M}_\Gamma(\mathfrak{p}^t, \underline{x})^{\leq E_0}_0$ of $\mathfrak{p}^t$-holomorphic disks gives a one-dimensional component $(u_t)_{t\in[0,1]}$ of quilted $\mathfrak{p}^{01}$-holomorphic disks where
   - if $\delta^{-1}(t) > 0$ (where $\delta : [-\infty, \infty] \to [0, 1]$ is a fixed increasing diffeomorphism), u_t is of the form in Figure 10.9 (a), u_t is constant on quilted components, J_t-holomorphic on the dark shaded disk, and boundary edges have length $\tau := \delta^{-1}(t)$;
   - if $\delta^{-1}(t) < 0$, u_t is of the form in Figure 10.9 (c), u_t is constant on quilted components and J_t-holomorphic on the light shaded disk, and boundary edges have length $\tau := -\delta^{-1}(t)$;
   - if $\delta^{-1}(t) = 0$, u_t is of the form in Figure 10.9 (b), on the quilted components u_t is J_t-holomorphic.

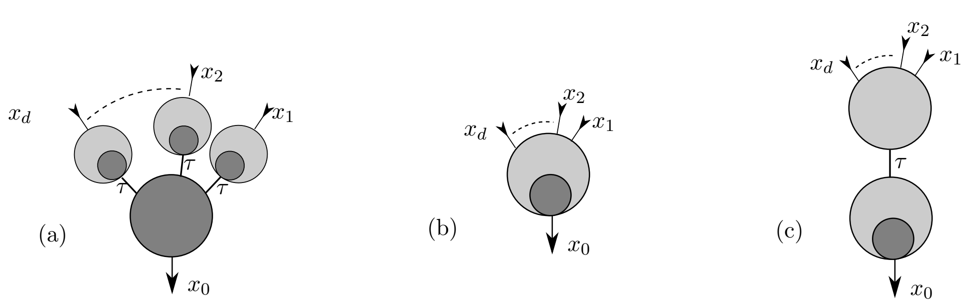

FIGURE 10.9. Types of quilted holomorphic disks occuring in the one-dimensional moduli space described in Remark 10.18 (b). In (a) and (c) the map is constant on the quilted component.

The analysis of the special case shows that in an idealized setting moduli spaces of quilted holomorphic disks interpolate between moduli spaces of unquilted holomorphic disks. But in general cases, one has to account for disk bubbling in the moduli spaces, leading to a variety of configurations occuring in the codimension one boundary of moduli spaces of quilted holomorphic disks, see Remark 10.20.

For any combinatorial type Γ of quilted holomorphic disks we denote by $\overline{\mathcal{M}}^q_\Gamma(L,D)$ the compactified moduli space of equivalence classes of adapted quilted holomorphic treed disks. The moduli space of quilted disks breaks into components depending on the limits along the root and leaf edges. Denote by $\mathcal{M}^q_\Gamma(L,D,\underline{x}) \subset \overline{\mathcal{M}}^q_\Gamma(L,D)$ the moduli space of isomorphism classes of stable adapted holomorphic quilted treed disks with boundary in L and limits $\underline{x}$ along the root and leaf edges, where $\underline{x} = (x_0,\ldots,x_{d(\circ)}) \in \widehat{\mathcal{I}}(L)$. For a rigid type Γ of quilted holomorphic disks, that is, where all edges $e \in \operatorname{Edge}_\circ(\Gamma)$ have finite non-zero length, and an input/output tuple $\underline{x}$, the expected dimension of the moduli space $\mathcal{M}^q_\Gamma(L,D,\underline{x})$ is

$$i^q(\Gamma,\underline{x}) = i(\Gamma',\underline{x}) + 1, \tag{10.26}$$

where i is the index for types of treed holomorphic disks (see (6.23)), and Γ' is the unquilted type obtained by forgetting the quilting in Γ. The extra dimension for the moduli space of quilted disks arises from the quilting datum. In case n quilted disk components, there are $n-1$ equations in the (Balancing condition), leading to an extra dimension of 1 compared to the unquilted case.

For a comeager subset of perturbation morphisms extending those chosen for unquilted disks, the uncrowded moduli spaces of expected dimension at most one are smooth and of expected dimension. For sequential compactness, it suffices to consider a sequence $u_\nu : C_\nu \to X$ of quilted treed disks of fixed combinatorial type $\Gamma = \Gamma_\nu$ for all ν. Coherence of the perturbation morphism implies the existence of a stable limit $u : C \to X$ which we claim is adapted. If a sequence of markings $z_{i,\nu} \in C_\nu$ converges to $z_i \in C$, then, $u(z_i) \in D^{\delta\circ d(z_i)}$. Indeed, since the distance from the seam $d(z_{i,\nu})$ converges to $d(z_i)$, we have $D^{\delta\circ d(z_i)}$ is the limit of the divisors $D^{\delta\circ d(z_{i,\nu})}$. For types of index at most one, each component of $u^{-1}(D^{\delta\circ d(z_i)})$ is a limit of a unique component of $u_\nu^{-1}(D^{\delta\circ d(z_{i,\nu})})$, otherwise the intersection degree would be more than one which is a codimension two condition. Therefore, every marking in C is a transverse divisor intersection. There are no other divisor intersections because the intersection number with $D^{\delta\circ d}$ is preserved in the limit for topological reasons.

The moduli space of quilted broken disks then has the same transversality and compactness property as in the unquilted case, by similar arguments.

Given a regular, stabilized and coherent perturbation morphism $\underline{\mathfrak{p}}^{01}$ from $\underline{\mathfrak{p}}^0$ to $\underline{\mathfrak{p}}^1$, define an A_∞-morphism $\phi = (\phi_d)_{d\geq 0} : CF(L,\mathfrak{p}^0) \to CF(L,\mathfrak{p}^1)$ as

$$\begin{aligned} \phi_d : CF(L;\underline{\mathfrak{p}}^0)^{\otimes d} &\to CF(L;\underline{\mathfrak{p}}^1) \\ (x_1,\ldots,x_d) &\mapsto \sum_{x_0,u\in\mathcal{M}_\Gamma(L,D,x_0,\ldots,x_d)_0} (-1)^{\heartsuit} w(u) x_0 \end{aligned} \tag{10.27}$$

where the weight $w(u)$ is given by

$$w(u) = \epsilon([u])(d_\bullet(\Gamma)!)^{-1} q^{E([u])} \operatorname{Hol}_L(u) x_0 \tag{10.28}$$

the sum is over uncrowded strata Γ of weighted treed quilted holomorphic disks whose boundary edges $e \in \text{Edge}_{\circ,-}(\Gamma)$ have finite non-zero length $\ell(e) \in (0, \infty)$ and whose input and output labels are compatible with $(x_0, \dots, x_d)$ in terms of the (Label axiom) in Definition 10.7, and $\epsilon([u]) = \pm 1$ is the orientation sign.

Remark 10.19. (Lowest area terms) For any $x \in \text{crit}(F^0) \cup \{x^{\blacktriangledown}, x^{\triangledown}\}$, the element $\phi^1(x)$ contains zero area terms coming from the count of a quilted treed disk with no interior marking, that is, a treed disk with only one disk that is quilted and mapped to a point. The domain is one of those in Figure 10.8. If x is $x^{\blacktriangledown}$ resp. $x^{\triangledown}$ there is one such configuration whose output is weighted resp. forgettable. In the latter case, it will be the only term with a forgettable output.

Remark 10.20. (Codimension one boundary strata) The codimension one strata are of several possible types: either there is one (or a collection of) edge(s) e of length $\ell(e)$ infinity, there is one (or a collection of) edge(s) e of length $\ell(e)$ zero, or equivalently, boundary nodes, or there is an edge e with zero or infinite weight $\rho(e)$. The case of an edge of zero or infinite weighting is equivalent to breaking off a constant trajectory, and so may be ignored. In the case of edges of infinite length(s), then either Γ is

(a) (Breaking off an uncolored tree) a pair $\Gamma_1 \sqcup \Gamma_2$ consisting of a colored tree Γ_1 and an uncolored tree Γ_2 as in the left side of Figure 10.6; necessarily the breaking must be a leaf of Γ_1; or
(b) (Breaking off colored trees) a collection consisting of an uncolored tree Γ_0 containing the root and a collection $\Gamma_1, \dots, \Gamma_r$ of colored trees attached to each of its r leaves as in the right side of Figure 10.6. Such a stratum $\mathcal{M}_\Gamma$ is codimension one because of the (Balanced Condition) which implies that if the length of any edge between e_0 to e_i is infinite for some i then the path from e_0 to e_i for any i has the same property.

In the case of a zero length(s), one obtains a fake boundary component with normal bundle $\mathbb{R}$, corresponding to either deforming the edge(s) to have non-zero length or deforming the node(s). This ends the Remark.

Proof of Proposition 10.13. The true boundary strata of one-dimensional moduli spaces of quilted holomorphic disks are those described in Remark 10.20 and correspond to the terms in the axiom for A_∞ morphisms (10.7). The signs are similar to those in [**18**] and omitted. The assertion on the strict units is a consequence of the existence of forgetful maps for infinite values of the weights. By assumption the $\phi^{d(\circ)}$ products involving $x^{\triangledown}$ as inputs involve counts of quilted treed disks using perturbation that are independent of the disk boundary incidence points of the first leaf marked $x^{\triangledown}$ asymptotic to $x_M \in X$. Since forgetting that semi-infinite edge gives a configuration of negative expected dimension, if non-constant, the only configurations contributing to these terms must be the constant maps. Hence

$$\phi^1(x_M^{\triangledown}) = x_M^{\triangledown}, \quad \phi^{d(\circ)}(\dots, x_M^{\triangledown}, \dots) = 0, n \geq 2.$$

In other words, the only regular quilted trajectories from the maximum, considered as $x_M^{\triangledown}$, being regular are the ones reaching the other maximum that do not have interior markings (i.e. non-constant disks). The proof of homotopy equivalence is via twice-quilted disks, and we refer to [**18**] for details. ☐

10.6. Homotopy equivalence: unbroken to broken

In this section, we show that the Fukaya algebra of the broken manifold $\mathfrak{X}$ is homotopy equivalent to the unbroken Fukaya algebra on the manifold X.

A rough outline of the proof is as follows. An A_∞ -morphism from an unbroken Fukaya algebra $CF(X^\nu, L)$ to a broken one $CF^{\text{brok}}(\mathfrak{X}, L)$ is defined as a limit of A_∞ -morphisms $\phi^n : CF(X^\nu, L) \to CF(X^{\nu+n}, L)$ as $n \to \infty$. The A_∞ -morphism ϕ^n is defined via a count of quilted holomorphic disks in (X, L). The existence of the limit relies on the bijection between moduli spaces of broken and unbroken disks obtained from the compactness and gluing theorems (Theorem 8.2 and Theorem 9.1).

We use a breaking perturbation datum on neck-stretched manifolds, so that the bijection between moduli spaces of broken and unbroken disks holds. We recall from Definition 8.1 that given a perturbation datum $\mathfrak{p}^\infty$ on the broken manifold $\mathfrak{X}$, by gluing on the neck we obtain a perturbation datum $\rho_\nu(\mathfrak{p}^\infty)$ on any X^ν. By Proposition 5.21 there is a stabilizing pair $(\mathfrak{J}, \mathfrak{D})$ consisting of a cylindrical almost complex structure $\mathfrak{J}$ on $\mathfrak{X}$ and a $\mathfrak{J}$-holomorphic cylindrical stabilizing divisor $\mathfrak{D}$, for which the glued family (J^ν, D^ν) is a stabilizing pair on the neck-stretched manifold X^ν for all ν. For the perturbation datum $\mathfrak{p}^\infty$ we use $\mathfrak{D}$ as the stabilizing divisor and $\mathfrak{J}$ as the background almost complex structure.

PROPOSITION 10.21. *Let $\mathfrak{p}^\infty$ be a regular perturbation datum on $\mathfrak{X}$. For any $E_0 > 0$, there exists $\nu_0(E_0)$ such that the following holds.*

(a) There is a bijection between the zero-dimensional moduli spaces of rigid broken maps and rigid unbroken maps:

$$\mathcal{M}^{\text{brok}}(\mathfrak{X}, L, \underline{\mathfrak{p}}^\infty)_0^{<E_0} \simeq \mathcal{M}(X^\nu, L, \rho_\nu(\mathfrak{p}^\infty))_0^{<E_0}, \quad \nu \geq \nu_0.$$

(b) For any $\nu \in \mathbb{Z}_{>0}$, there exists a regular perturbation datum $\mathfrak{p}^\nu$, and a perturbation morphism $\mathfrak{p}^{\nu,\nu+1}$ extending $\mathfrak{p}^\nu$ and $\mathfrak{p}^{\nu+1}$ such that for all $E_0 > 0$ and $\nu \geq \nu_0(E_0)$, the A_∞ morphism induced by $\mathfrak{p}^{\nu,\nu+1}$

$$\phi_\nu : CF(X^\nu, L, \underline{\mathfrak{p}}^\nu) \to CF(X^{\nu+1}, L, \underline{\mathfrak{p}}^{\nu+1})$$

is identity modulo q^{E_0}. Similarly there is an A_∞ morphism

$$\psi_\nu : CF(X^{\nu+1}, L, \underline{\mathfrak{p}}^{\nu+1}) \to CF(X^\nu, L, \underline{\mathfrak{p}}^\nu)$$

that is identity modulo q^{E_0}.

PROOF. The proof of bijection of moduli spaces is a consequence of the compactness and gluing theorems – Theorem 8.2 and Theorem 9.1. By the gluing Theorem 9.1, any regular $\mathfrak{p}^\infty$-disk $u : C \to \mathfrak{X}$ can be glued to yield regular disks $u_\nu : C_\nu \to X^\nu$. Conversely, any sequence $(u_\nu)_\nu$ of maps with area $\leq E_0$ converges to a broken disk u_∞. By surjectivity of gluing in Theorem 9.1, for large enough ν, u_ν is contained in the image of the gluing map for u_∞. Since the moduli space $\mathcal{M}(\mathfrak{X}, L)_0^{\leq E_0}$ has a finite number of points, the constant $\nu_0(E_0)$ can be chosen as the minimum obtained from gluing each of the broken maps.

To prove part (b), we construct regular perturbations and perturbation morphisms on neck-stretched manifolds. A regular perturbation datum is constructed by extending the breaking perturbation datum. For any $E_0 > 0$ and $\nu > \nu_0(E_0)$, the perturbation datum is defined as

$$\mathfrak{p}^{\nu}_{\Gamma} := \rho_\nu(\mathfrak{p}^{\infty}_{\Gamma}) = (J^{\nu}_{\Gamma}, F^{\infty}_{\Gamma}) \tag{10.29}$$

if $E(\Gamma) \leq E_0$. For the other strata, a regular perturbation $\mathfrak{p}^{\nu}_{\Gamma}$ can be chosen using the transversality result Theorem 6.29 for all integers ν.

For strata with low enough area, perturbation morphisms are constructed using the breaking perturbation data. For $\nu \in \mathbb{Z}_+$, Γ satisfying $\nu \geq \nu_0(E(\Gamma))$, the perturbation morphism is defined using the recipe in Remark 10.18 as

$$\mathfrak{p}^{\nu,\nu+1}_{\Gamma} = (J^{\nu,\nu+1}_{\Gamma}, F^{\infty}_{\Gamma'})$$

where $F^{\infty}_{\Gamma'}$ is part of the perturbation data $\underline{\mathfrak{p}}^{\infty} = (\underline{J}^{\infty}, \underline{F}^{\infty})$ for broken disks, Γ' is the type of broken disk obtained by forgetting the quilting, and for any $z \in \mathcal{S}_\Gamma$,

$$J^{\nu,\nu+1}_{\Gamma}(z) := J^{\nu+(\delta\circ d(z))}_{\Gamma'}(z),$$

where d is the distance to seams function from (10.20), $\delta : [-\infty, \infty] \to [0, 1]$ is a fixed increasing function from (10.24), and $J^{\nu+\delta\circ d(z)}_{\Gamma'}$ is the neck-stretched domain-dependent almost complex structure from (10.29). The perturbation morphisms are extended to higher strata so that they are coherent and regular.

Finally we will show that this perturbation morphism satisfies (b). The A_∞-morphism ϕ^ν is defined by counting $\mathfrak{p}^{\nu,\nu+1}$-adapted quilted holomorphic disks. First, we point out that this count includes quilted holomorphic disks with a single input and output with the same label (as in Figure 10.8) and a constant map on the surface and tree components, which gives an identity term in ϕ^ν_1. We claim that for any E_0 there is a constant $\nu_0(E_0)$ such that for $\nu \geq \nu_0(E_0)$ there are no other $\mathfrak{p}^{\nu,\nu+1}$-adapted quilted disks with area $\leq E_0$. We prove this statement by assuming the contrapositive, which means that, there is a sequence u_ν of quilted $\mathfrak{p}^{\nu,\nu+1}$-disks with bounded area. Any surface component of u_ν is $J^{\nu'}$-holomorphic for some $\nu' \in [\nu, \nu+1]$. Therefore, after passing to a subsequence, the surface part of the map u_ν converge to a limit J^∞-holomorphic broken map u. For large enough ν, the Maslov index $I(u_\nu)$ is ν-independent, and therefore the number of disk inputs is bounded. After passing to a subsequence, we can assume that the disk input/output tuple $\underline{x}$ and the type Γ is the same for all the quilted disks u_ν. Since the A_∞-morphism counts quilted holomorphic disks with index 0, we may assume that the index of the quilted disks is zero: $i^q(\Gamma, \underline{x}) = 0$. The sequence u_ν converges to a $\mathfrak{p}^\infty$-unquilted broken disk u of type Γ_τ and input/output labels given by $\underline{x}$. Here Γ_τ is the type of a broken treed holomorphic map for which collapsing the tropical edges yields a type Γ' of a treed holomorphic map, which is the same as the type obtained by forgetting the quilting in Γ. The unquilted map u is stable in all cases except the one contributing to the identity term, that is, when there is a single input and output, and the map is constant on the surface and tree part of the domain. In all other cases, the index of the unquilted disk is one less (see (10.26)), and we have $i(\Gamma', \underline{x}) = -1$, and since collapsing tropical edges does not affect the dimension, we have $i(\Gamma_\tau, \underline{x}) = -1$. The existence of such a disk u is a contradiction, since $\mathfrak{p}^\infty$ is a regular perturbation. We have thus shown that if $\nu \geq \nu_0(E_0)$,

$$\phi_1^\nu = \mathrm{Id} \mod q^{E_0}, \quad \phi_k^\nu \mod q^{E_0} = 0, \ \ k \neq 1.$$

The A_∞ morphism ψ_ν is defined similarly. □

PROPOSITION 10.22. *For any ν_0 let $\underline{\mathfrak{p}}_{\nu_0}$ be the perturbation defined in Proposition 10.21. Then, there are strictly convergent unital A_∞ morphisms*

$$\phi : CF(X^{\nu_0}, L, \underline{\mathfrak{p}}^{\nu_0}) \to CF_{\mathrm{brok}}(L, \underline{\mathfrak{p}}^\infty), \quad \psi : CF_{\mathrm{brok}}(L, \underline{\mathfrak{p}}^\infty) \to CF(X^{\nu_0}, L, \underline{\mathfrak{p}}^{\nu_0})$$

such that $\psi \circ \phi$ and $\phi \circ \psi$ are A_∞ -homotopy equivalent.

PROOF. The A_∞ morphisms ϕ and ψ required by the Proposition are defined to be the limits of the A_∞ morphisms from Proposition 10.21 (b), namely,

$$\phi_\nu : CF(X^\nu, L, \underline{\mathfrak{p}}^\nu) \to CF(X^{\nu+1}, L, \underline{\mathfrak{p}}^{\nu+1}), \quad \psi_\nu : CF(X^{\nu+1}, L, \underline{\mathfrak{p}}^{\nu+1}) \to CF(X^\nu, L, \underline{\mathfrak{p}}^\nu).$$

For any E_0, if $\nu \geq \nu'(E_0)$, the A_∞ morphisms ϕ_ν, ψ_ν are equal to Id modulo q^{E_0}. Therefore, there exist limits of the successive compositions: letting

$$\phi_n := \phi_{\nu_0} \circ \phi_{\nu_0+1} \circ \ldots \circ \phi_{\nu_0+n} : CF(X^{\nu_0}, L, \underline{\mathfrak{p}}^{\nu_0}) \to CF(X^{\nu_0+n+1}, L, \underline{\mathfrak{p}}^{\nu_0+n+1})$$

for any choice of ν_0, the limit

$$\phi = \lim_{n\to\infty} \phi_n : CF(X^{\nu_0}, L, \underline{\mathfrak{p}}^{\nu_0}) \to \lim_{n\to\infty} CF_{\mathrm{brok}}(L, \underline{\mathfrak{p}}^\infty)$$

exists and similarly for

$$\psi = \lim_{n\to\infty} \psi_n, \quad \psi_n := \psi_\nu \circ \psi_{\nu+1} \circ \ldots \circ \psi_{\nu+n}.$$

Since the composition of strictly unital morphisms is strictly unital, the composition ψ is strictly unital mod terms divisible by q^E for any E, hence strictly unital. The limit map ψ resp. ϕ is an A_∞ morphism whose domain resp. target is $CF_{\mathrm{brok}}(L)$ because the composition maps in $CF(X^\nu, L, \mathfrak{p}^\nu)$ converge to $CF_{\mathrm{brok}}(L, \underline{\mathfrak{p}}^\infty)$. The composition maps of $CF(X^\nu, L, \mathfrak{p}^\nu)$ converge to those of $CF_{\mathrm{brok}}(L, \underline{\mathfrak{p}}^\infty)$ by the bijection of the moduli spaces of disks in Proposition 10.21 (a). Indeed the bijection preserves the orientation $\epsilon(u)$, area $A(u)$, and holonomy $\mathrm{Hol}([\partial u])$ since gluing preserves these quantities.

We claim that ϕ and ψ are homotopy equivalences. Let h_k, g_k denote the homotopies satisfying

$$\phi_k \circ \psi_k - \mathrm{Id} = m_1(h_k), \quad \psi_k \circ \phi_k - \mathrm{Id} = m_1(g_k),$$

from the homotopies relating $\phi_\nu \circ \psi_\nu$ and $\psi_\nu \circ \phi_\nu$ to the identities, see [**18**, Theorem 5.10]. In particular, h_{k+1}, g_{k+1} differ from h_k, g_k by expressions involving counting *twice-quilted* breaking disks. For any $E > 0$, for ν sufficiently large all terms in $h_{k+1} - h_k$ are divisible by q^E. It follows that the infinite composition

$$h = \lim_{k\to\infty} h_k, \quad g = \lim_{k\to\infty} g_k$$

exists and gives a homotopy equivalence between $\phi \circ \psi$ resp. $\psi \circ \phi$ and the identities. □

10.7. Decoupling : Split perturbations

So far we have used perturbations on broken manifolds that are gluable on the neck regions. This requirement was needed for proving the A_∞-homotopy equivalence between broken and unbroken Fukaya algebras. Now, with that task accomplished, we allow perturbations of almost complex structures that are independently defined in different components of the broken manifold. In doing so, we drop the requirement that the perturbation is coherent under collapsing tropical edges. We rigorously define this new kind of perturbation, called *split perturbation*, and show that the broken Fukaya algebras it produces are homotopy equivalent to the ones we have defined so far. A split perturbation depends not just on the domain curve but also on an accompanying tropical structure on it, recorded by the *base tropical graph*. Maps adapted to split perturbations may have additional components than those in the base tropical graph Γ. However, for generic perturbations, for maps of index ≤ 1 the tropical graph is the same as the base tropical graph.

DEFINITION 10.23. (Curve with base type) A *curve with base type* consists of

(a) a stable treed disk C of type $\tilde{\Gamma}$,
(b) a *base tropical graph* Γ,
(c) and an edge collapse morphism $\kappa : \tilde{\Gamma} \to \Gamma$ that necessarily collapses all disk edges $e \in \mathrm{Edge}_\circ(\tilde{\Gamma})$ and treed segments T_e, $e \in \mathrm{Edge}_\circ(\tilde{\Gamma})$, and possibly some interior edges. (The map κ is an edge collapse of graphs, and not of tropical graphs, because $\tilde{\Gamma}$ does not have a tropical structure.)

The *type* of such a curve consists of the datum $\tilde{\Gamma} \xrightarrow{\kappa} \Gamma$. The moduli space of stable treed curves of type $\tilde{\Gamma}$ with base type Γ is denoted by $\mathcal{M}_{\tilde{\Gamma},\Gamma}$, or simply $\mathcal{M}_{\tilde{\Gamma}}$ when the context allows it. The topology on the space of based curves is the standard topology on the space of stable curves with the additional axiom that the base tropical graph is preserved in the limit. This ends the Definition.

DEFINITION 10.24. (Based graph morphisms) The following morphisms are defined on types of curves with a base :

(a) A based curve type $\tilde{\Gamma}_1 \xrightarrow{\kappa_1} \Gamma$ is obtained by *(Collapsing edges)* resp. *(Making an edge length/weight finite/non-zero)* in the based curve type $\tilde{\Gamma}_0 \xrightarrow{\kappa_0} \Gamma$ if there is a (Collapsing edges) resp. (Making an edge length/weight finite/non-zero) morphism $\tilde{\Gamma}_0 \xrightarrow{\tilde{\kappa}} \tilde{\Gamma}_1$ of treed curve types (see Definition 6.5) and $\tilde{\kappa} \circ \kappa_0 = \kappa_1$. Note that the (Collapsing edges) morphism does not collapse an edge e in $\tilde{\Gamma}_0$ that is present in the base graph Γ.
(b) A type with base $\tilde{\Gamma} \xrightarrow{\kappa} \Gamma$ is obtained by *(Cutting an edge)* in $\tilde{\Gamma}' \xrightarrow{\kappa'} \Gamma'$ if
    - $\tilde{\Gamma}$ is obtained by cutting an $e = (v_+, v_-) \in \mathrm{Edge}_{\circ,-}(\tilde{\Gamma}')$ containing a breaking,
    - Γ' is obtained by identifying the vertices $\kappa(v_+)$, $\kappa(v_-)$ in Γ.

DEFINITION 10.25. (Split perturbation datum) Let Γ be a tropical graph. A *coherent split perturbation datum* is a collection

$$\underline{\mathfrak{p}} = (\mathfrak{p}_{(\tilde{\Gamma},\Gamma)})_{(\tilde{\Gamma},\Gamma)},$$

consisting of a perturbation datum

$$\mathfrak{p}_{(\tilde{\Gamma},\Gamma)} = (J_{(\tilde{\Gamma},\Gamma)}, F_{(\tilde{\Gamma},\Gamma)}), \quad J_{\tilde{\Gamma},\Gamma} : \mathcal{S}_{\tilde{\Gamma}} \to \mathcal{J}^{\mathrm{cyl}}(\mathfrak{X}), \quad F_{\tilde{\Gamma},\Gamma} : \mathcal{T}_{\tilde{\Gamma}} \to C^{\infty}(L, \mathbb{R})$$

for each based curve type $(\tilde{\Gamma}, \Gamma)$, that is coherent under the based graph morphisms of (Collapsing of edges), (Making an edge length or weight finite/non-zero) and (Cutting edges) from Definition 10.24, and the (Locality axiom) from Definition 6.5.

Remark 10.26. An important difference between ordinary perturbations (from Definition 6.5) on $\mathfrak{X}$ and the split perturbations in the above Definition 10.25 is that the split perturbation datum $\mathfrak{p}_{\tilde{\Gamma},\Gamma}$ is not required to be coherent under collapsing edges that are present in the base tropical graph Γ. Indeed, for a based graph $\tilde{\Gamma} \xrightarrow{\kappa} \Gamma$, the (Collapsing edges) morphism is only applicable on edges $e \in \mathrm{Edge}(\tilde{\Gamma}) \setminus \mathrm{Edge}(\Gamma)$.

Remark 10.27. (Background J for split perturbations) For a split perturbation datum $\mathfrak{p}_{\tilde{\Gamma},\Gamma}$, the stabilizing divisor and background almost complex structure can be chosen to be different for curve components corresponding to different vertices of Γ. Coherence can be ensured by the following scheme:

Definition 10.28. For a based curve type $(\tilde{\Gamma}, \Gamma)$, the *root vertex* $v_\circ$ of the base tropical graph is the vertex $v_\circ \in \mathrm{Vert}(\Gamma)$ containing the disk components of $\tilde{\Gamma}$, that is, $\kappa(v) = v_\circ$ for all $v \in \mathrm{Vert}_\circ(\tilde{\Gamma})$. The type $(\tilde{\Gamma}, \Gamma)$ is *atomic* if $v_\circ$ has at most one incident edge in the tropical graph Γ.

Stabilizing pairs underlying a coherent split perturbation datum can be chosen by making the following choices independently:

- a stabilizing pair on (D_0, J_0) on $X^{\square}_{P_0}$,
- and a stabilizing pair (D_v, J_v) on $X^{\square}_{P(v)}$ for all $v \in \mathrm{Vert}(\Gamma) \setminus \{v_\circ\}$, where Γ ranges over all atomic tropical graphs. The pair (D_0, J_0) is used for the root vertex in all graphs.

Here, we note that two tropical graphs Γ_0, Γ_1 are isomorphic if there is an isomorphism of graphs $\phi : \Gamma_0 \to \Gamma_1$ that preserves polytope assignments on vertices, edge slopes, and the root vertex. Two isomorphic tropical graphs are assigned the same stabilizing pairs on their vertices.

The above choice of stabilizing pairs on based curves with an atomic base tropical graph extends to a coherent choice of stabilizing pairs for all components of all based curves. Indeed, given a based curve type $(\tilde{\Gamma}, \Gamma)$ where the root vertex in the tropical graph Γ has k incident edges, there are atomic tropical graphs $\Gamma_1, \ldots, \Gamma_k$ such that Γ is obtained by identifying the root vertices in $\Gamma_1, \ldots, \Gamma_k$; and the stabilizing pairs assignments on $\Gamma_1, \ldots, \Gamma_k$ yield stabilizing pairs on components of based curves of type $(\tilde{\Gamma}, \Gamma)$. This finishes Remark 10.27.

Remark 10.29. (Cylindrical structures for split perturbations) The cylindrical almost complex structures used to define based perturbations need not be gluable. Therefore, we may choose the cylindrical almost complex structure J_P on $X^{\square}_P$ so that there is a taming embedding of $(X^{\square}_P, J_P)$ into the symplectic broken manifold (X_P, ω_{X_P}).

A coherent split perturbation datum $\underline{\mathfrak{p}}$ lets us define adapted broken maps with base.

DEFINITION 10.30. (Adapted broken maps with base) Let $\underline{\mathfrak{p}}$ be a coherent split perturbation datum.

(a) A $\underline{\mathfrak{p}}$-*adapted broken map with base* consists of
- a based type $\kappa : \tilde{\Gamma} \to \Gamma$,
- and a broken map u whose domain is of type $\tilde{\Gamma}$ and u is $\mathfrak{p}_{\tilde{\Gamma},\Gamma}$-adapted in the ordinary sense (as in Definitions 6.10, 6.14). Further if $\tilde{\Gamma}_{\mathrm{tr}}$ is the tropical graph of u then κ factors through $\tilde{\Gamma}_{\mathrm{tr}}$ as
$$\kappa = \kappa_1 \circ \mathrm{tr}, \quad \tilde{\Gamma} \xrightarrow{\mathrm{tr}} \tilde{\Gamma}_{\mathrm{tr}} \xrightarrow{\kappa_1} \Gamma,$$
where κ_1 is a tropical edge collapse.

(b) (Type) The *type* of an adapted broken map with base consists of the tropical graph $\tilde{\Gamma}$, the tropical edge collapse map $\tilde{\Gamma} \to \Gamma$, and the homology and tangency data for the map u (as in the type of a broken map, see Definition 6.15). Whenever it is possible, the base tropical type Γ is suppressed in the notation.

(c) (Tropical symmetry) For a type $\tilde{\Gamma}$ of a broken map with base, the tropical symmetry group $T_{\mathrm{trop}}(\tilde{\Gamma})$ is the symmetry group obtained by viewing $\tilde{\Gamma}$ as a type of broken map (by Definition 4.33) and forgetting the base type Γ.

(d) (Rigidity) The type $\tilde{\Gamma} \to \Gamma$ of a broken map with base is *rigid* if $\tilde{\Gamma}$ is rigid as a type of broken map. Thus for a rigid type, the base tropical graph Γ is rigid, the tropical graph $\tilde{\Gamma}_{\mathrm{tr}}$ of the map is equal to Γ, and the morphism $\tilde{\Gamma} \to \Gamma$ only collapses treed components and boundary nodes.

Split perturbations can be used to define broken Fukaya algebras. Given a coherent split perturbation datum $\underline{\mathfrak{p}}$, we define a Fukaya algebra as the graded vector space
$$CF_{\mathrm{brok}}(\mathfrak{X}, L, \underline{\mathfrak{p}}) := CF^{\mathrm{geom}}(L) \oplus \Lambda x^{\blacktriangledown}[1] \oplus \Lambda x^{\triangledown}$$
with composition maps
$$m_d^{\mathrm{def}}(x_1, \ldots, x_d) = \sum_{x_0, u \in \mathcal{M}_{\tilde{\Gamma},\Gamma}(\underline{\mathfrak{p}}, \underline{\eta}, \underline{x})_0} w_s(u) x_0, \tag{10.30}$$
where u ranges over all rigid types $(\tilde{\Gamma}, \Gamma)$ of broken maps with base that have d inputs (see Definition 10.30 (d) for rigidity), and
$$w_s(u) := (-1)^{\heartsuit} (d_{\bullet}(\Gamma)!)^{-1} \operatorname{Hol}([\partial u]) \epsilon(u) q^{A(u)}, \tag{10.31}$$
where the symbols in (10.31) are as in (10.10), and in particular the orientation is computed in the same way as for broken map without base, see Remark 6.30. As in the case of broken maps (without base), for a one-dimensional component of the moduli space, the configurations with a boundary edge of length zero constitute a fake boundary, whereas those with a broken boundary edge constitute the true boundary of the moduli space. Setting the counts of the maps occuring in the

true boundary to zero yields the A_∞-associativity relations, and we conclude that $CF(L, \underline{\mathfrak{p}})$ is an A_∞-algebra.

Broken Fukaya algebras defined by counts of broken maps with base are independent of the choice of perturbation up to homotopy equivalence. The proof is the same as the case of broken Fukaya algebras without base. If we assume that the stabilizing divisors have the same degree, the proof is by constructing an A_∞-morphism by counts of quilted broken holomorphic disks with base. Compared to quilted broken holomorphic disks without base, the only new feature is that the domains additionally have a base tropical graph, and perturbations are not coherent under collapsing the edges in the base tropical graph. The extension is straightforward because quilting phenomena take place on disk components, and the base tropical graph collapses disk components, and therefore the two features do not interact with each other. We thus obtain the following result, which is an extension of Proposition 10.13.

PROPOSITION 10.31. (Independence of perturbations, split version) *Suppose* $\underline{\mathfrak{p}}^0$, $\underline{\mathfrak{p}}^1$ *are regular split perturbation data that are defined using stabilizing pairs* (J^0, D^0) *and* (J^1, D^1)*, which are connected by a path of stabilizing pairs* $\{(J^t, D^t)\}_{t\in[0,1]}$. *There exists a coherent perturbation datum* $\underline{\mathfrak{p}}^{01}$ *which induces a convergent unital* A_∞ *-morphism (as in* (10.8)*)*

$$\phi_{01} : CF(L, \underline{\mathfrak{p}}^0) \to CF(L, \underline{\mathfrak{p}}^1)$$

which is a homotopy equivalence.

In fact, Fukaya algebras defined using ordinary perturbation data are A_∞ -homotopy equivalent to those defined using split perturbation data:

PROPOSITION 10.32. *Let* $\underline{\mathfrak{p}}$ *be a regular coherent perturbation datum for the broken manifold* $(\mathfrak{X}, L)$ *and let* $\underline{\mathfrak{p}}' = (\underline{\mathfrak{p}}'_\Gamma)_\Gamma$ *be a collection of regular coherent split perturbation data for all rigid tropical graphs* Γ. *Then the broken Fukaya algebras* $CF(\underline{\mathfrak{p}})$, $CF(\underline{\mathfrak{p}}')$ *are homotopy equivalent.*

PROOF. The result follows from the homotopy equivalence of any two broken Fukaya algebras defined by split perturbations, and the fact that a non-split perturbation may be regarded as a split perturbation in the following way: given a perturbation $\underline{\mathfrak{p}}$ without base, we define a split perturbation $\underline{\mathfrak{p}}'$ by forgetting the base tropical graph. The resulting curve counts defining the composition maps are the same for $\underline{\mathfrak{p}}$ and $\underline{\mathfrak{p}}'$. Indeed, a $\underline{\mathfrak{p}}$-adapted map of type Γ corresponds to a unique $\underline{\mathfrak{p}}'$-adapted map of type $\Gamma \xrightarrow{\mathrm{tr}} \Gamma_{\mathrm{tr}}$ where tr is the tropicalization map; the uniqueness follows from the fact that since Γ is rigid, there is no other tropical graph that can be obtained by edge collapses. $\square$

CHAPTER 11

Tropical Fukaya algebras

In this chapter we introduce a degeneration of broken maps into split maps. In split maps, there is no edge matching condition on a subset of edges, called *split edges*. A homotopy equivalence between the moduli space of broken maps and that of split maps is constructed by deforming the edge matching condition at the split edges. The maps with a deformed edge matching condition on split edges are called *deformed maps*. We refer the reader to Section 1.5 for an introduction to split maps, and to Figure 1.14 for an example.

11.1. Deformed maps

A deformed map is a version of a broken map where the edge matching condition for split edges is replaced by a deformed matching condition. Let $(X, \Phi, \mathcal{P})$ be a symplectic manifold with a tropical Hamiltonian action as in Definition 3.3.

Definition 11.1. (Split edges) Let $\mathcal{P}_s \subset \mathcal{P}$ be the set of polytopes P for which

(a) the tropical moment map Φ generates a T/T_P-action on X_P and this action makes X_P a toric manifold,
(b) and any torus-invariant divisor D of X_P is a relative divisor. That is, there is a polytope $Q \subset P$ such that $X_Q = D$.

For a tropical graph Γ, $e \in \mathrm{Edge}(\Gamma)$ is a *split edge* if $P(e) \in \mathcal{P}_s$. The set of split edges of a tropical graph is denoted by

$$\mathrm{Edge}_s(\Gamma) \subset \mathrm{Edge}(\Gamma).$$

By definition, a split edge e satisfies the property that the space $X_{\overline{P(e)}}$ is a toric variety that is a compactification of $T_{\mathbb{C}}$, and the same is true of the space $X_{\overline{Q}}$ where $Q \subset P(e)$ is a face of the polytope $P(e)$.

Example 11.2. An edge whose polytope $P(e) \in \mathcal{P}$ is zero dimensional is a split edge. The example of more interest to us will be a top-dimensional component $X_{P_0} \subset \mathfrak{X}$ which is a toric variety, and all whose torus-invariant divisors are relative divisors. That is, for any torus-invariant divisor $D \subset X_{P_0}$, there is a polytope $Q \subset P$ in $\mathcal{P}$ such that $D = X_Q$.

Definition 11.3. A *deformation parameter* for a tropical graph Γ is an element

$$\eta = (\eta_e)_{e \in \mathrm{Edge}_s(\Gamma)} \in \bigoplus_{e \in \mathrm{Edge}_s(\Gamma)} \mathfrak{t}/\mathfrak{t}_{\mathcal{T}(e)} =: \mathfrak{t}_\Gamma, \tag{11.1}$$

where $\mathfrak{t}_{\mathcal{T}(e)} \subset \mathfrak{t}$ is the linear span $\langle \mathcal{T}(e) \rangle$ of the slope $\mathcal{T}(e) \in \mathfrak{t}_{P(e),\mathbb{Z}}$ of the edge e.

Deformed maps are defined on based curves. We recall from Definition 10.23 that a based curve is a treed disk of type $\tilde{\Gamma}$ together with a base tropical graph Γ and an edge collapse relation $\tilde{\Gamma} \to \Gamma$. Thus the domain contains the additional information of a polytope $P(v)$ for each component. The set of split edges is a subset of the edges in the base tropical graph Γ.

DEFINITION 11.4. (Deformed map) Let η be a deformation parameter for a tropical graph Γ. A η-deformed map of type Γ consists of a

(a) a treed curve C of type Γ,
(b) a framing $\mathrm{fr}_e : T_{w_+(e)}\tilde{C} \otimes T_{w_-(e)}\tilde{C} \to \mathbb{C}$ for each tropical edge $e \in \mathrm{Edge}_{\mathrm{trop}}(\Gamma)$,
(c) and maps

$$u_v : C_v \to X_{\overline{P}(v)}, \quad v \in \mathrm{Vert}(\Gamma)$$

satisfying edge matching conditions as in Definition 4.12 of broken maps for all edges $e \in \mathrm{Edge}(\Gamma) \setminus \mathrm{Edge}_s(\Gamma)$. On split edges the maps $(u_v)_v$ satisfy a deformed matching condition described as follows: For a node $w = (w_+, w_-)$ corresponding to a split edge $e = (v_+, v_-) \in \mathrm{Edge}_s(\Gamma)$, and coordinates on the neighborhoods of $w_\pm$ that respect the framing

$$x_+ = e^{i\eta_e} x_-, \quad \text{where } x_\pm := \mathrm{ev}^{\mathcal{T}(e)}\, u_{v_\pm}(w_\pm(e)) \in X^{\square}_{\overline{P}(e)}.$$

Here $\mathcal{T}(e) \in \mathfrak{t}_{P(e),\mathbb{Z}}$, $\mathrm{ev}^{\mathcal{T}(e)}$ is the tropical evaluation map (see (4.19)), and $\eta_e \in \mathfrak{t}/\mathfrak{t}_{\mathcal{T}(e)}$ is the deformation parameter for the edge e. (We identify $\mathfrak{t}/\mathfrak{t}_{\mathcal{T}(e)}$ to a subspace of $\mathfrak{t}$ using a fixed inner product.)

REMARK 11.5. We take the deformation parameter to be in the quotient $\mathfrak{t}/\mathfrak{t}_{\mathcal{T}(e)}$ (and not in $\mathfrak{t}$) because altering it in the $\mathcal{T}(e)$-direction affects only the framing, not the map. For example, consider a node where the edge matching condition is satisfied for a broken map, and so, $x_+ = x_-$. Changing the domain trivialization has the effect of changing the matching to $x_+ = e^{\mathcal{T}(e)z} x_-$ for some $z \in \mathbb{C}$.

11.2. Moduli spaces of deformed maps

The deformation datum for a deformed map is domain-dependent and satisfies coherence conditions.

DEFINITION 11.6. (Deformation datum) A coherent deformation datum $\underline{\eta} = (\eta_{\tilde{\Gamma},\Gamma})_{(\tilde{\Gamma},\Gamma)}$ consists of a continuous map

$$\eta_{\tilde{\Gamma},\Gamma} : \mathcal{M}_{\tilde{\Gamma}} \to \mathfrak{t}_\Gamma \simeq \bigoplus_{e \in \mathrm{Edge}_s(\Gamma)} \mathfrak{t}/\mathfrak{t}_{\mathcal{T}(e)}$$

for all types $(\tilde{\Gamma}, \Gamma)$ of treed disks with base, that are coherent under morphisms (Cutting edges), (Collapsing edges), (Making an edge length/weight finite/non-zero), (Forgetting edges) for based curve types, and the following (Marking independence) axiom: For any type $\tilde{\Gamma}$, the map $\eta_{\tilde{\Gamma},\Gamma} | \mathcal{M}_{\tilde{\Gamma}}$ factors as

$$\text{(11.2)} \quad \text{(Marking independence)} \quad \eta_{\tilde{\Gamma},\Gamma} = \eta'_{\tilde{\Gamma}} \circ f_{\tilde{\Gamma}}, \quad \mathcal{M}_{\tilde{\Gamma}} \xrightarrow{f_{\tilde{\Gamma}}} [0,\infty]^{\mathrm{Edge}_{\circ,-}(\tilde{\Gamma})} \xrightarrow{\eta'_{\tilde{\Gamma}}} \mathfrak{t}_\Gamma,$$

where $f_{\tilde{\Gamma}}([C])$ is equal to the edge lengths $\ell(e)$ of treed components at the boundary nodes of the treed curve C.

REMARK 11.7. If one were defining moduli space of deformed spheres (instead of deformed disks), then the deformation datum can be taken to be constant on the moduli space. In other words, one could just assign a constant deformation parameter η_e to each split edge e. A domain-dependent deformation datum is needed in order to achieve coherence with respect to the (Cutting an edge) morphism, and this requirement is absent for the moduli space of spheres.

A perturbation datum for deformed maps $\underline{\mathfrak{p}} = (\mathfrak{p}_{\tilde{\Gamma}})_{\tilde{\Gamma}}$ is a split perturbation datum as in Definition 10.25, and consists of maps

$$\mathfrak{p}_{\tilde{\Gamma}} = (J_{\tilde{\Gamma}}, F_{\tilde{\Gamma}}), \quad J_{\tilde{\Gamma}} : \mathcal{S}_{\tilde{\Gamma}} \to \mathcal{J}^{\mathrm{cyl}}(\mathfrak{X}), \ \ F_{\tilde{\Gamma}} : \mathcal{T}_{\tilde{\Gamma}} \to C^\infty(L, \mathbb{R})$$

for all types $\tilde{\Gamma}$ of curves with base, with coherence conditions corresponding to morphisms of stable treed curves with base.

DEFINITION 11.8. (Adapted deformed maps) Let $\tilde{\Gamma} \to \Gamma$ be a type of based curve. Let $\mathfrak{p}_{\tilde{\Gamma}}$ be a perturbation datum and $\eta_{\tilde{\Gamma}}$ be a deformation datum. An *adapted deformed map* is the same as an adapted broken map $u : C \to \mathfrak{X}$ with base (see Definition 10.30) with the difference that on any split edge $e \in \mathrm{Edge}_s(\Gamma)$ the edge matching condition is replaced by the condition that the matching condition at the node w_e corresponding to e is $\eta_{\tilde{\Gamma}}([C], e)$-deformed. The combinatorial type, tropical symmetry group, and the rigidity condition on adapted deformed maps are the same as the corresponding properties on adapted broken maps with base.

For a type $\tilde{\Gamma}$ of adapted deformed maps, let

$$\mathcal{M}^{\mathrm{def}}_{\tilde{\Gamma}}(\underline{\mathfrak{p}}, \underline{\eta})$$

denote the moduli space of adapted deformed maps with deformation parameter $\underline{\eta}$ modulo domain reparametrizations.

REMARK 11.9. Compactified moduli spaces of adapted deformed maps of different base tropical types do not intersect. Indeed, a stratum $\mathcal{M}_{\mathrm{def}, \tilde{\Gamma}_0}$ of deformed maps is contained in the compactification $\overline{\mathcal{M}}_{\mathrm{def}, \tilde{\Gamma}}$ only if $\tilde{\Gamma}_0$ is obtained by (Collapsing edges) in $\tilde{\Gamma}$ or by (Making an edge length/weight finite/non-zero) in $\tilde{\Gamma}$, and both these morphisms preserve the base tropical graph.

The following is an analogue of the theorems for broken (undeformed) maps, and the proofs are analogous.

PROPOSITION 11.10. *Given a coherent deformation datum $\underline{\eta}$, there is a comeager set $\mathcal{P}^{\mathrm{reg}}$ of coherent regular perturbations for which the following hold.*

(a) *For any uncrowded type $\tilde{\Gamma}$ of adapted deformed maps, regular perturbation $\mathfrak{p}_{\tilde{\Gamma}} \in \mathcal{P}^{\mathrm{reg}}$, and for any disk output and inputs $x_0, \ldots, x_{d(\circ)} \in \hat{\mathcal{I}}(L)$ such that $i(\tilde{\Gamma}, \underline{x}) \leq 1$, the moduli space $\mathcal{M}_{\mathrm{def}, \tilde{\Gamma}}(\underline{\mathfrak{p}}, \underline{\eta}, \underline{x})$ is a manifold of expected dimension.*

(b) *Any one-dimensional component of the moduli space of rigid deformed broken treed disks* $\mathcal{M}_{\mathrm{def},\tilde{\Gamma}}(\underline{\mathfrak{p}},\underline{\eta},\underline{x})$ *admits a compactification as a topological manifold with boundary. The true boundary is equal to the union of zero-dimensional strata whose domain treed disks* $u : C \to \mathfrak{X}_{\mathcal{P}}$ *have a boundary edge* $e \in \mathrm{Edge}_{\circ}(\Gamma)$ *that is broken.*

(c) *For any* $E > 0$*, there are finitely many zero and one-dimensional components of the moduli space of rigid broken treed disks with area* $\leq E$*.*

We define a deformed version of the Fukaya algebra by counting deformed broken maps:

Definition 11.11. (Deformed Fukaya algebra) Let $\underline{\eta}$ be a coherent deformation datum, and let $\underline{\mathfrak{p}} = (\mathfrak{p}_{\tilde{\Gamma}})_{\tilde{\Gamma}}$ be a coherent regular perturbation datum for all types $\tilde{\Gamma}$ of based treed disks whose base tropical graph Γ is rigid. The *deformed Fukaya algebra* is the graded vector space

$$CF_{\mathrm{def}}(L,\underline{\eta}) := CF^{\mathrm{geom}}(L) \oplus \Lambda x^{\blacktriangledown}[1] \oplus \Lambda x^{\triangledown}$$

equipped with composition maps

$$m^{\mathrm{def}}_{d_\circ}(x_1,\ldots,x_{d_\circ}) = \sum_{x_0, u \in \mathcal{M}_{\mathrm{def},\tilde{\Gamma}}(\underline{\mathfrak{p}},\underline{\eta},\underline{x})_0} w_s(u) x_0, \tag{11.3}$$

where the type $\tilde{\Gamma}$ of u ranges over all rigid types with $d_\circ$ inputs (see Definition 10.30 (d) for rigidity), and

$$w_s(u) := (-1)^{\heartsuit}(d_{\bullet}(\Gamma)!)^{-1}(s(\tilde{\Gamma})!)^{-1}\,\mathrm{Hol}([\partial u])\epsilon(u)q^{A(u)}, \tag{11.4}$$

where $s(\tilde{\Gamma})$ is the number of split edges in the type $\tilde{\Gamma}$ and the other symbols in (11.4) are as in (10.10).

As in the case of broken maps, for a one-dimensional component of the moduli space, the configurations with a boundary edge of length zero constitute a fake boundary, whereas those with a broken boundary edge constitute the true boundary of the moduli space. The A_∞-axioms for the Fukaya algebra $CF_{\mathrm{def}}(L,\underline{\eta})$ follow from counts of the true boundary points of one-dimensional components of the moduli space of deformed maps, and we obtain the following result.

Proposition 11.12. *The composition maps in* (11.3) *satisfy the* A_∞ *axioms, and the Fukaya algebra* $CF_{\mathrm{def}}(L,\underline{\eta})$ *is strictly unital.*

The Fukaya algebras are independent of deformation parameters $\underline{\eta}$ up to homotopy equivalence. The proof is by constructing an A_∞ morphism by counts of a version of quilted disks which we now define. Perturbation morphisms contain all the data as in the unbroken case as in Definition 10.15, and additionally include a path of deformation data.

Definition 11.13. (Deformed quilted holomorphic disks)

(a) (Perturbation morphism) Let $\underline{\eta}_0$, $\underline{\eta}_1$ be deformation data, and let $\underline{\mathfrak{p}}_k$ be a regular perturbation datum for $\underline{\eta}_k$, $k = 0, 1$. A *perturbation morphism* $(\underline{\mathfrak{p}}^{01},\underline{\eta}^{01})$ consists of

(i) a path of deformation data $\eta^{01} := \{\underline{\eta}_t\}_{t\in[0,1]}$ connecting $\underline{\eta}_0$ and $\underline{\eta}_1$,

(ii) a perturbation morphism $\mathfrak{p}^{01}$ as in the unbroken case connecting $\underline{\mathfrak{p}}_0$ and $\underline{\mathfrak{p}}_1$,

(b) (Adapted deformed quilted disk) Suppose $(\underline{\mathfrak{p}}^{01}, \underline{\eta}^{01})$ is a perturbation morphism for quilted deformed maps. A map $u : C \to \mathfrak{X}_{\mathcal{P}}$ is an *adapted deformed quilted disk* if

(i) it is $\mathfrak{p}^{01}$-adapted in the sense of a holomorphic quilted broken treed disk,

(ii) and the deformation parameter at the node w_e corresponding to a split edge $e \in \mathrm{Edge}_s(\Gamma)$ is $\eta_{\delta(d(w_e)),\tilde{\Gamma}}(C, e)$. Here we note that w_e is a point in the universal curve $\overline{\mathcal{U}}_{\tilde{\Gamma}}$, $d(w_e) \in [-\infty, \infty]$ is the distance from the seam (see (10.20)) and $\delta : [-\infty, \infty] \to [0, 1]$ is a fixed increasing diffeomorphism from (10.24).

The following Proposition is analogous to Proposition 10.13, and the proof of that Theorem carries over.

PROPOSITION 11.14. *Given a path of deformation data $\{\eta_t\}_{t\in[0,1]}$, and regular perturbation data $\underline{\mathfrak{p}}^0$, $\underline{\mathfrak{p}}^1$ for the end-points, there is a comeager set of regular perturbation morphisms extending $\underline{\mathfrak{p}}^0$, $\underline{\mathfrak{p}}^1$. Any such perturbation morphism induces a convergent unital A_∞ morphism*

$$\phi = (\phi_r)_{r\geq 0} : CF_{\mathrm{def}}(L, \underline{\mathfrak{p}}^0, \underline{\eta}_0) \to CF_{\mathrm{def}}(L, \underline{\mathfrak{p}}^1, \underline{\eta}_1)$$

defined by counts of quilted deformed disks. The A_∞ morphism ϕ is a A_∞ homotopy equivalence and has zero-th composition map $\phi_0(1)$ with positive q-valuation, that is $\phi^0_{01}(1) \in \Lambda_{>0}\langle\hat{\mathcal{I}}(L)\rangle$.

11.3. Split tropical graphs

Split maps arise as limits of deformed maps as the deformation parameters fo to infinity. A split map is a version of a broken map where there is no matching condition along split edges. The tropical properties of deformed maps, such as the tropical graph and the symmetry group, look exactly like that of broken maps; this is explained by the fact that deformation of an edge matching condition is a small scale phenomenon that is not seen by the tropical graph which only detects large scale behavior. As the deformation parameters go to infinity, split maps, the limit objects, have a new kind of tropical graph called *split tropical graph*, which is introduced in this section. Split tropical graphs do not satisfy the slope condition at split edges, and they satisfy a 'cone condition' which ensures that the tropical symmetry group is as large as possible. In particular, in the rigid case, the dimension of the tropical symmetry group is equal to the codimension of the matching conditions at split edges.

We first define a notion of genericity and a notion of an increasing sequence in a vector space. Both these notions are used to describe criteria under which rational polyhedral cones in a vector space are top-dimensional, and will be used to describe the cone condition for split tropical graphs.

DEFINITION 11.15. Let $V \simeq \mathbb{R}^n$ be a vector space equipped with a dense rational lattice $V_{\mathbb{Q}} \simeq \mathbb{Q}^n$.

(a) (Rational subspace) A *rational subspace* of V is a linear subspace $W \subset V$ in which $W \cap V_{\mathbb{Q}}$ is dense.
(b) (Generic vector) A vector $\eta \in V$ is *generic* if it is not contained in any proper rational subspace of V.
(c) (Rational linear map) Suppose W is a real vector space with a dense rational lattice $W_{\mathbb{Q}}$. A linear map $f : V \to W$ is *rational* if $f(V_{\mathbb{Q}}) \subseteq W_{\mathbb{Q}}$. The subspaces $\ker(f)$ and $\mathrm{im}(f)$ are rational subspaces.

REMARK 11.16. The notion of genericity is applicable to the Lie algebra $\mathfrak{t}$ of the torus T since it has an integral lattice $\mathfrak{t}_{\mathbb{Z}} \subset \mathfrak{t}$ and consequently a rational lattice. For a slope $\mathcal{T}(e) \in \mathfrak{t}_{\mathbb{Z}}$ of a tropical edge, the projection

$$\pi^{\perp}_{\mathcal{T}(e)} : \mathfrak{t} \to \mathfrak{t}/\mathfrak{t}_{\mathcal{T}(e)}$$

projects the rational lattice $\mathfrak{t}_{\mathbb{Q}} \subset \mathfrak{t}$ to a dense subset $(\mathfrak{t}/\mathfrak{t}_{\mathcal{T}(e)})_{\mathbb{Q}} \subset \mathfrak{t}/\mathfrak{t}_{\mathcal{T}(e)}$, which happens to be the rational lattice obtained by viewing the quotient $\mathfrak{t}/\mathfrak{t}_{\mathcal{T}(e)}$ as the Lie algebra of the torus $T/T_{\mathcal{T}(e)}$. Thus the projection $\pi^{\perp}_{\mathcal{T}(e)}$ maps generic elements of $\mathfrak{t}$ to generic elements in the quotient.

DEFINITION 11.17. (Rational polyhedral cone) Let V be a vector space with a dense rational subspace $V_{\mathbb{Q}} \subset V$. A *rational polyhedral cone* $C \subset V$ is a finite intersection of rational half planes, that is

$$C = \cap_{i=1}^{k} H_i \quad \text{where} \quad H_i = \{x \in V : \langle x, \nu_i \rangle \geq 0\}, \;\; \nu_i \in V_{\mathbb{Q}}^{\vee}.$$

Equivalently, $C \subset V$ is a rational polyhedral cone if an only if there are rational elements $u_1, \dots, u_m \in V_{\mathbb{Q}}$ such that

$$C = \left\{ \sum_i t_i u_i : t_i \geq 0 \right\}.$$

LEMMA 11.18. *Suppose $C \subset V$ is a rational polyhedral cone containing a generic vector. Then C is top-dimensional.*

PROOF. The unsigned cone $C_{\pm} := \mathbb{R}\langle C \rangle \subset V$ defined as the $\mathbb{R}$-span of C is a rational subspace of V that contains a generic vector. Therefore $C_{\pm} = V$ and C is top-dimensional. □

The 'increasing condition' defined next is another characterization of top-dimensional cones.

DEFINITION 11.19. (Increasing cone) For any $1 \leq i \leq n$, define a face

$$F_i := \{x_{i+1} = \cdots = x_n = 0\} \subset (\mathbb{R}_{\geq 0})^n.$$

A rational polyhedral cone $C \subset (\mathbb{R}_{\geq 0})^n$ is an *increasing cone* if for any i, $C \cap F_i$ is an i-dimensional cone in F_i.

The increasing condition on a cone implies that it is top-dimensional in $\mathbb{R}^n$. The increasing condition can be equivalently stated in terms of increasing sequences of tuples.

DEFINITION 11.20. (Increasing sequences of tuples) A sequence of tuples

$$(x_{1,\nu}, \ldots, x_{n,\nu}) \in \mathbb{R}^n_+$$

is *increasing* if $x_{i,\nu} \to \infty$ for all i as $\nu \to \infty$, and

$$\lim_{\nu} \frac{x_{i,\nu}}{x_{i-1,\nu}} = 0$$

for $i = 2, \ldots, n$.

LEMMA 11.21. *The following are equivalent for a rational polyhedral cone $C \subset (\mathbb{R}_{\geq 0})^n$.*

(a) C is an increasing cone.
(b) (Weak sequence) *There exists an increasing sequence of tuples in C.*
(c) (Strong sequence) *The tail of any increasing sequence of tuples lies in C.*

PROOF. The increasing cone condition can be characterized inductively : $C \subset (\mathbb{R}^n_+)$ is an increasing cone if and only if $e_1 := (1, 0, \ldots, 0) \in C$ and the *normal cone*

$$N_{e_1} C := \{(v_2, \ldots, v_n) \in \mathbb{R}^{n-1}_+ : \exists t_0 > 0 : (1, tv_2, \ldots, tv_n) \in C \quad \forall 0 \leq t \leq t_0\}$$

is an increasing cone in $(\mathbb{R}_+)^{n-1}$. Conditions (b) and (c) can also be characterized in the same way. That is, a cone $C \subset (\mathbb{R}^n_+)$ satisfies the (Weak sequence) resp. (Strong sequence) condition if and only if $(1, 0, \ldots, 0) \in \overline{C}$ and the normal cone $N_{e_1} C$ satisfies the (Weak sequence) resp. (Strong sequence) condition. The proof of the lemma then follows by induction. □

The property of being increasing and containing a generic vector can be combined to detect top-dimensionality of cones:

LEMMA 11.22. *Let $V = \sum_{i=1}^k V_i$ be a direct sum of vector spaces, each of which has a dense rational lattice $V_{i,\mathbb{Q}} \subset V_i$. For each i, let $\eta_i \in V_i$ be a generic element, Suppose $C \subset V$ be a rational polyhedral cone that contains a sequence*

$$\eta_\nu := \sum_{i=1}^k c_{\nu,i} \eta_i$$

where $c_\nu = (c_{\nu,i})_i \in \mathbb{R}^n_+$ is an increasing sequence. Then C is top-dimensional in V.

PROOF. We show C is top-dimensional by proving that the $\mathbb{R}$-span of C, denoted by $C_\pm \subset V$ is, in fact, equal to V. First, we observe that $\eta_1 \in C_\pm$, since, by the increasing condition, η_1 is the limit of the sequence $\frac{\eta_\nu}{c_{\nu,1}} \in C$. Then, since $V_1 \cap C_\pm$ is a rational subspace of V_1 and it contains the generic element η_1, we conclude $V_1 = V_1 \cap C_\pm$. Therefore $V_1 \subset C_\pm$. The proof can now be finished by induction: The projection of C to $\oplus_{i=2}^n V_i$ is a rational polyhedral cone and it contains the sequence $\sum_{i=2}^k c_{\nu,i} \eta_i$. □

With the properties of genericity and increasing cones at hand, we return to the task of defining split tropical graphs. We first describe a preliminary version called a 'quasi-split tropical graph' and then describe the genericity condition under which such a graph is a split tropical graph.

DEFINITION 11.23. (Quasi-split tropical graph) A *quasi-split tropical graph* $\tilde{\Gamma}$ consists of

(a) a tropical graph Γ, called a *base tropical graph* with a collection of split edges $\mathrm{Edge}_s(\Gamma) \subseteq \mathrm{Edge}(\Gamma)$ as in Definition 11.1,
(b) a graph $\tilde{\Gamma}$ with an edge collapse morphism $\kappa : \tilde{\Gamma} \to \Gamma$,
(c) an ordering $\prec_{\tilde{\Gamma}}$ on the split edges $\mathrm{Edge}_s(\Gamma)$,
(d) and a tropical structure on each connected component of the graph $\tilde{\Gamma} \setminus \mathrm{Edge}_s(\Gamma)$ so that the restricted map

$$\kappa : \tilde{\Gamma} \setminus \mathrm{Edge}_s(\Gamma) \to \Gamma \setminus \mathrm{Edge}_s(\Gamma)$$

is a tropical edge collapse. Consequently for any vertex $v \in \mathrm{Vert}(\tilde{\Gamma})$, $P_{\tilde{\Gamma}}(v) \subset P_{\Gamma}(v)$.

REMARK 11.24. A quasi-split tropical graph can alternately be defined as the combinatorial type $\kappa : \tilde{\Gamma} \to \Gamma$ of a curve with base (see Definition 10.23) together with a tropical structure on the connected components of $\tilde{\Gamma} \setminus \mathrm{Edge}_s(\Gamma)$ so that κ is a tropical edge collapse on $\tilde{\Gamma} \setminus \mathrm{Edge}_s(\Gamma)$.

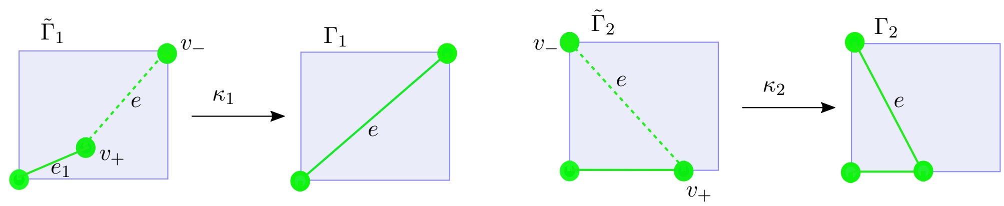

FIGURE 11.1. Quasi-split tropical graphs $(\tilde{\Gamma}_1, \Gamma_1)$, $(\tilde{\Gamma}_2, \Gamma_2)$. In both cases e is the only split edge, and the slope condition is not satisfied for e in $\tilde{\Gamma}_i$.

DEFINITION 11.25. (Relative vertex positions, relative translations for a quasi-split graph) A *relative vertex position map* for a quasi-split tropical graph $\tilde{\Gamma} \xrightarrow{\kappa} \Gamma$ is the difference between a vertex position of $\tilde{\Gamma}$ and a vertex position of Γ, and the space of relative vertex positions is

$$\mathcal{W}(\tilde{\Gamma}, \Gamma) := \{(\tilde{\mathcal{T}}(v) - \mathcal{T}(\kappa v) \in \mathrm{Cone}(\kappa, v) \subset \mathfrak{t}^{\vee}_{P_{\tilde{\Gamma}}(v)})_{v \in \mathrm{Vert}(\Gamma')} : \tilde{\mathcal{T}} \in \mathcal{W}(\tilde{\Gamma}), \mathcal{T} \in \mathcal{W}(\Gamma)\}.$$

A $(\tilde{\Gamma}, \Gamma)$-*translation* or a *relative translation* is an element in the $\mathbb{R}_+$-span of the space of relative vertex positions. The space of relative translations is denoted by

$$w(\tilde{\Gamma}, \Gamma) := \mathbb{R}_+ \langle \mathcal{W}(\tilde{\Gamma}, \Gamma) \rangle \simeq \mathrm{Cone}_{\mathcal{W}(\Gamma)} \mathcal{W}(\tilde{\Gamma}). \tag{11.5}$$

As in the case of relative tropical graphs, relative translations of a quasi-split tropical graph can be characterized by an edge slope condition as in the Lemma below. The proof of the Lemma is left to the reader.

LEMMA 11.26. (Slope condition on relative translations) *Let* $\tilde{\Gamma} \xrightarrow{\kappa} \Gamma$ *be a quasi-split tropical graph. A tuple* $(\mathcal{T}(v) \in \operatorname{Cone}(\kappa, v))_{v \in \operatorname{Vert}(\tilde{\Gamma})}$ *is a relative translation, that is* $\mathcal{T} \in w(\tilde{\Gamma}, \Gamma)$*, if and only if*

$$\text{(Slope)} \quad \mathcal{T}(v_+) - \mathcal{T}(v_-) \in \begin{cases} \mathbb{R}_{\geq 0}\mathcal{T}(e), & e \in \operatorname{Edge}(\tilde{\Gamma}) \setminus \operatorname{Edge}(\Gamma), \\ \mathbb{R}\mathcal{T}(e), & e \in \operatorname{Edge}(\Gamma) \setminus \operatorname{Edge}_s(\Gamma). \end{cases} \tag{11.6}$$

There is no (Slope) condition on the split edges $e \in \operatorname{Edge}_s(\Gamma)$*.*

REMARK 11.27. (Relative translation cone) For a quasi-split tropical graph $\tilde{\Gamma} \xrightarrow{\kappa} \Gamma$, the set

$$w(\tilde{\Gamma}, \Gamma) \subset \oplus_{v \in \operatorname{Vert}(\tilde{\Gamma})} \mathfrak{t}^{\vee}_{P_{\tilde{\Gamma}}(v)}$$

of relative translations is a rational polyhedral cone. Indeed, the set $w(\tilde{\Gamma}, \Gamma)$ is cut out of $\oplus_{v \in \operatorname{Vert}(\tilde{\Gamma})} \mathfrak{t}^{\vee}_{P_{\tilde{\Gamma}}(v)}$ by the (Slope) condition in (11.6), and for each of the edges, the slope condition can be written as a set of linear inequalities with rational coefficients, each cutting out a half-plane. We call $w(\tilde{\Gamma}, \Gamma)$ the *relative translation cone*.

EXAMPLE 11.28. The relative translation cones for the quasi-split tropical graphs in Figure 11.1 are

$$w(\tilde{\Gamma}_1, \Gamma_1) = \mathbb{R}_{\geq 0}, \quad w(\tilde{\Gamma}_2, \Gamma_2) = \mathbb{R}.$$

Indeed, in $\tilde{\Gamma}_2$, the vertex v_+ is free to move to both sides of its position in Γ_2. On the other hand in $\tilde{\Gamma}_1$, the edge e_1 is collapsed by κ_1, the vertex $\kappa_1(v_+)$ in Γ_1 lies at the vertex of the dual complex, and in $\tilde{\Gamma}_1$, v_+ is free to move along a line of slope $\mathcal{T}(e_1)$, and therefore the set of relative positions of v_+ in $\tilde{\Gamma}_1$ is $\mathbb{R}_{\geq 0}$.

DEFINITION 11.29. (Discrepancy cone) Let $\tilde{\Gamma} \to \Gamma$ be a quasi-split tropical graph. Define

$$\begin{aligned} \operatorname{Diff} := (\operatorname{Diff}_e)_{e \in \operatorname{Edge}_s(\Gamma)} : w(\tilde{\Gamma}, \Gamma) \to \mathfrak{t}_\Gamma \simeq \bigoplus_{e \in \operatorname{Edge}_s(\Gamma)} \mathfrak{t}/\langle \mathcal{T}(e) \rangle, \\ \mathcal{T} \mapsto (\pi^{\perp}_{\mathcal{T}(e)}(\mathcal{T}(v_+) - \mathcal{T}(v_-)))_{e=(v_+,v_-) \in \operatorname{Edge}_s(\Gamma)} \end{aligned} \tag{11.7}$$

as the amount by which a relative translation $\mathcal{T}$ fails to satisfy the (Slope) condition at split edges $e \in \operatorname{Edge}_s(\Gamma)$. The *discrepancy cone* for a quasi-split tropical graph is the image

$$\operatorname{Disc}(\tilde{\Gamma}, \Gamma) := \operatorname{Diff}(w(\tilde{\Gamma}, \Gamma)) \subset \bigoplus_{e \in \operatorname{Edge}_s(\Gamma)} \mathfrak{t}/\langle \mathcal{T}(e) \rangle. \tag{11.8}$$

REMARK 11.30. The set $\operatorname{Disc}(\tilde{\Gamma}, \Gamma)$ is in fact a rational polyhedral cone, because $w(\tilde{\Gamma}, \Gamma)$ is a rational polyhedral cone, see Remark 11.27.

Definition 11.31. (Split tropical graph) Let $\eta_0 \in \mathfrak{t}^\vee$ be a generic element, which we call the *cone direction*. A *split tropical graph* $\tilde{\Gamma}$ with cone direction η_0 is a quasi-split tropical graph $\tilde{\Gamma}$ (as in Definition 11.23) whose discrepancy cone $\operatorname{Disc}(\tilde{\Gamma}, \Gamma)$ (see (11.8)) satisfies the following *cone condition*: There is an increasing cone $C \subset (\mathbb{R}_{\geq 0})^n$ such that for all $(c_1, \ldots, c_n) \in C$,

$$(c_i \pi^{\perp}_{\mathcal{T}(e_i)}(\eta_0))_i \in \operatorname{Disc}(\tilde{\Gamma}, \Gamma). \tag{11.9}$$

The cone condition in Definition 11.31 can alternately be expressed in terms of increasing sequences of tuples. Firstly, we adapt the definition of increasing sequence of tuples (Definition 11.20) to split edges using the ordering on the set of split edges.

Definition 11.32. (Increasing sequence of tuples for split edges) A sequence of tuples

$$(c_\nu(e))_{e \in \operatorname{Edge}_s(\Gamma)} \in (\mathbb{R}_+)^{\operatorname{Edge}_s(\Gamma)}$$

is *increasing* if $c_\nu(e) \to \infty$ for all split edges $e \in \operatorname{Edge}_s(\Gamma)$ and for a pair of split edges $e_i \prec e_j, e_i, e_j \in \operatorname{Edge}_s(\Gamma)$ (by the ordering in Definition 11.23(c))

$$\lim_{\nu} \tfrac{c_\nu(e_j)}{c_\nu(e_i)} = 0. \tag{11.10}$$

Lemma 11.33. *Let* $\tilde{\Gamma} \to \Gamma$ *be a quasi-split tropical graph with cone direction* $\eta_0 \in \mathfrak{t}^\vee$. *The following are equivalent for the discrepancy cone* $\operatorname{Disc}(\tilde{\Gamma}, \Gamma)$*:*

(a) $\operatorname{Disc}(\tilde{\Gamma}, \Gamma)$ *satisfies the (Cone condition).*
(b) (Strong sequential cone condition) *For any increasing sequence of tuples* $(c_\nu(e))_{e \in \operatorname{Edge}_s(\Gamma)} \in (\mathbb{R}_+)^{\operatorname{Edge}_s(\Gamma)}$*, there exists* ν_0 *such that*

$$(c_\nu(e) \pi^{\perp}_{\mathcal{T}(e)}(\eta_0))_e \in \operatorname{Disc}(\tilde{\Gamma}, \Gamma) \quad \forall \nu \geq \nu_0. \tag{11.11}$$

(c) (Weak sequential cone condition) *There exists an increasing sequence of tuples* $(c_\nu(e))_{e \in \operatorname{Edge}_s(\Gamma)} \in (\mathbb{R}_+)^{\operatorname{Edge}_s(\Gamma)}$ *such that*

$$(c_\nu(e) \pi^{\perp}_{\mathcal{T}(e)}(\eta_0))_e \in \operatorname{Disc}(\tilde{\Gamma}, \Gamma) \quad \forall \nu. \tag{11.12}$$

The Lemma is a consequence of the corresponding result (Lemma 11.21) on cones in $(\mathbb{R}_{\geq 0})^n$.

The main result of the section is that the cone condition for a split tropical graph implies that the discrepancy cone is top-dimensional:

Proposition 11.34. (Dimension of discrepancy cone) *Suppose* $\tilde{\Gamma} \to \Gamma$ *is a split tropical graph, with cone direction* η_0. *Then,*

$$\dim(\operatorname{Disc}(\tilde{\Gamma}, \Gamma)) = |\operatorname{Edge}_s(\Gamma)|(\dim(\mathfrak{t}) - 1). \tag{11.13}$$

Proof. Proposition 11.34 follows from Lemma 11.22 and the fact that the discrepancy cone $\operatorname{Disc}(\tilde{\Gamma}, \Gamma)$ is a rational polyhedral cone that satisfies the (Cone condition). □

11.4. Split maps

A split map is modelled on a split tropical graph, and it is a version of a broken map where there is no matching condition along split edges. The definition of a split tropical graph requires us to choose a generic cone direction $\eta_0 \in \mathfrak{t}^\vee$ as in Definition 11.31. Therefore, the definition of a split map also requires the choice of a vector η_0, although it is not explicitly mentioned in this Section.

DEFINITION 11.35. (Split map) Given a split tropical graph $\tilde{\Gamma} \to \Gamma$, a *split map* is a collection of maps

$$u : C \to \mathfrak{X} = (u_v : C_v \to X_{\overline{P}_{\tilde{\Gamma}}(v)})_{v \in \mathrm{Vert}(\tilde{\Gamma})}$$

whose domain C is a treed disk of type $\tilde{\Gamma}$, and framing isomorphisms for all *non-split* tropical edges

$$\mathrm{fr} : T_{w_+(e)}\tilde{C} \otimes T_{w_-(e)}\tilde{C} \to \mathbb{C}, \quad e \in \mathrm{Edge}_{\mathrm{trop}}(\tilde{\Gamma}) \setminus \mathrm{Edge}_s(\Gamma),$$

such that the edge-matching condition (as in Definition 4.12) is satisfied for edges $e \in \mathrm{Edge}(\tilde{\Gamma}) \setminus \mathrm{Edge}_s(\Gamma)$. At the split edges $e \in \mathrm{Edge}_s(\Gamma)$, the edge matching condition is dropped.

The moduli spaces decompose into strata indexed by combinatorial types, as before. The *type* of a split map is given by the split tropical graph $\tilde{\Gamma} \to \Gamma$ and the tangency and homology data on $\tilde{\Gamma}$. Thus the type of a split map is the same as the type of a broken map with base in Definition 10.30 (b) with the only difference that the tropical graph $\tilde{\Gamma}$ does not satisfy the (Slope) condition (8.32) on the split edges. The type of a split map is denoted by $\tilde{\Gamma}$ when the base tropical type Γ is clear from the context.

We wish to count rigid split maps to define the tropical Fukaya algebra.

DEFINITION 11.36. (Rigidity for split maps) The type $\tilde{\Gamma} \to \Gamma$ of a split map is *rigid* if

(a) Γ is a rigid tropical graph,
(b) in $\tilde{\Gamma}$ the only non-tropical edges $e \in \mathrm{Edge}(\tilde{\Gamma}) \setminus \mathrm{Edge}_{\mathrm{trop}}(\tilde{\Gamma})$ are boundary edges $e \in \mathrm{Edge}_\circ(\tilde{\Gamma})$ with finite non-zero length $\ell(e) \in (0, \infty)$,
(c) for all interior markings z_e, $e \in \mathrm{Edge}_{\bullet,\to}(\tilde{\Gamma})$, the intersection multiplicity with the stabilizing divisor is 1,
(d) and the cone of tropical vertex positions $\mathcal{W}(\tilde{\Gamma}, \Gamma)$ has dimension

$$\dim(\mathcal{W}(\tilde{\Gamma}, \Gamma)) = |\mathrm{Edge}_s(\Gamma)|(\dim(\mathfrak{t}) - 1). \tag{11.14}$$

The definition of tropical symmetry in the case of split maps is adjusted to reflect the lack of matching conditions at split edges.

DEFINITION 11.37. A *tropical symmetry* on a split map with graph $\tilde{\Gamma} \to \Gamma$ consists of a translation $g_v \in T_{P(v),\mathbb{C}}$ for each vertex $v \in \mathrm{Vert}(\tilde{\Gamma})$, and a framing translation $z_e \in \mathbb{C}^\times$ for each non-split tropical edge $e \in \mathrm{Edge}_{\mathrm{trop}}(\tilde{\Gamma}) \setminus \mathrm{Edge}_s(\Gamma)$ that satisfies

$$g(v_+)g(v_-)^{-1} = z_e^{\mathcal{T}(e)}, \quad \forall e = (v_+, v_-) \in \mathrm{Edge}_\bullet(\tilde{\Gamma}) \setminus \mathrm{Edge}_s(\Gamma), \tag{11.15}$$

where we assume $z_e = 1$ for non-tropical edges $e \in \operatorname{Edge}_-(\tilde{\Gamma}) \backslash \operatorname{Edge}_{\text{trop}}(\tilde{\Gamma})$. Denote the group of tropical symmetries as

$$T_{\text{trop}}(\tilde{\Gamma}, \Gamma) = \{((g_v)_v, (z_e)_e)\}. \tag{11.16}$$

For a split tropical graph $\tilde{\Gamma} \to \Gamma$, the space of relative vertex positions $\mathcal{W}(\tilde{\Gamma}, \Gamma)$ generates a subgroup of tropical symmetries (see (8.33)).

Remark 11.38. (Dimension of the symmetry group) For a split tropical graph $\tilde{\Gamma}$

$$\dim_{\mathbb{C}}(T_{\text{trop}}(\tilde{\Gamma})) \geq |\operatorname{Edge}_s(\Gamma)|(\dim(\mathfrak{t}) - 1). \tag{11.17}$$

Indeed the set of relative translations $w(\tilde{\Gamma}, \Gamma)$ generates a subgroup of $T_{\text{trop}}(\tilde{\Gamma}, \Gamma)$, and by (11.13) $w(\tilde{\Gamma}, \Gamma)$ is a cone of dimension at least $|\operatorname{Edge}_s(\Gamma)|(\dim(\mathfrak{t}) - 1)$. If $\tilde{\Gamma}$ is rigid, then (11.17) is an equality.

Remark 11.39. (Splitting of the tropical symmetry group) For a split tropical graph $\tilde{\Gamma} \to \Gamma$, the group $T_{\text{trop}}(\tilde{\Gamma}, \Gamma)$ is a product

$$T_{\text{trop}}(\tilde{\Gamma}, \Gamma) = T_{\text{trop}}(\tilde{\Gamma}_1, \Gamma_1) \times \cdots \times T_{\text{trop}}(\tilde{\Gamma}_s, \Gamma_s) \tag{11.18}$$

of tropical symmetry groups of the connected subgraphs

$$\tilde{\Gamma}_1, \ldots, \tilde{\Gamma}_s \subset \tilde{\Gamma} \backslash \operatorname{Edge}_s(\Gamma), \quad \Gamma_i := \kappa(\tilde{\Gamma}_i) \subset \Gamma.$$

Indeed by definition there are no matching conditions (11.15) on split edges.

To relate deformed maps to split maps, we will need a variation of split maps where there is a matching condition on split nodes also. Such maps should be thought of as lying on a slice of the $T_{\text{trop}}(\tilde{\Gamma}, \Gamma)$-action.

Definition 11.40. (a) (Framed split map) A *framed split map* is a split map u together with an additional datum of framings on split edges

$$\operatorname{fr}_e : T_{w_+(e)}\tilde{C} \otimes T_{w_-(e)}\tilde{C} \to \mathbb{C}, \quad e \in \operatorname{Edge}_s(\Gamma)$$

such that the matching condition is satisfied at the split edges.

(b) (Tropical symmetry for framed split maps) A *tropical symmetry* for a framed split map with graph $\tilde{\Gamma} \to \Gamma$ is a tuple

$$((g_v \in T_{P(v),\mathbb{C}})_{v \in \operatorname{Vert}(\Gamma)}, (z_e \in \mathbb{C}^\times)_{e \in \operatorname{Edge}_{\text{trop}}(\tilde{\Gamma})})$$

satisfying a matching condition on all interior edges :

$$g(v_+)g(v_-)^{-1} = z_e^{\mathcal{T}(e)} \quad \forall e \in \operatorname{Edge}_\bullet(\tilde{\Gamma}). \tag{11.19}$$

The group of tropical symmetries for framed split maps is denoted by

$$T_{\text{trop,fr}}(\tilde{\Gamma}, \Gamma) = \{((g_v)_v, (z_e)_e) \mid (11.19)\}. \tag{11.20}$$

(c) (Multiplicity of a split tropical graph) For a rigid split tropical graph $\tilde{\Gamma} \to \Gamma$, the group $T_{\text{trop,fr}}(\tilde{\Gamma})$ is finite, and

$$\operatorname{mult}(\tilde{\Gamma}) := |T_{\text{trop,fr}}(\tilde{\Gamma})| \tag{11.21}$$

is called the multiplicity of $\tilde{\Gamma}$. See Example 11.52.

REMARK 11.41. A framed split map is not a broken map, because the underlying tropical graph $\tilde{\Gamma}$ does not satisfy the (Slope) condition on split edges.

11.5. Tropical Fukaya algebras

We define moduli spaces of split maps and use them to define composition maps for tropical Fukaya algebras. To define the moduli spaces, we fix a generic cone direction $\eta_0 \in \mathfrak{t}^\vee$. A perturbation datum for split maps $\underline{\mathfrak{p}} = (\mathfrak{p}_{\tilde{\Gamma}})_{\tilde{\Gamma}}$ is a split perturbation datum as in Definition 10.25, and consists of maps

$$\mathfrak{p}_{\tilde{\Gamma}} = (J_{\tilde{\Gamma}}, F_{\tilde{\Gamma}}), \quad J_{\tilde{\Gamma}} : \mathcal{S}_{\tilde{\Gamma}} \to \mathcal{J}^{\mathrm{cyl}}(\mathfrak{X}), \;\; F_{\tilde{\Gamma}} : \mathcal{T}_{\tilde{\Gamma}} \to C^\infty(L, \mathbb{R})$$

for all types $\tilde{\Gamma}$ of curves with base, with coherence conditions corresponding to morphisms of stable treed curves with base. For each such type $\tilde{\Gamma}$, $J_{\tilde{\Gamma}}$ is a domain-dependent almost complex structure on the broken manifold, and $F_{\tilde{\Gamma}}$ is a domain-dependent Morse function on the Lagrangian L. This datum is the same as the one used for based maps in Section 10.7 and deformed maps in Section 11.2.

DEFINITION 11.42. (Moduli spaces of split maps) Let $\eta_0 \in \mathfrak{t}^\vee$ be a generic cone direction. The *moduli space of split maps* of type $\tilde{\Gamma}$ with cone direction η_0 modulo the action of domain reparametrizations is denoted

$$\mathcal{M}^{\mathrm{split}}_{\tilde{\Gamma}}(L, \mathfrak{p}_{\tilde{\Gamma}}, \eta_0).$$

Quotienting by the action of the tropical symmetry group defines the *reduced moduli space*

$$\mathcal{M}^{\mathrm{split}}_{\tilde{\Gamma},\mathrm{red}}(L, \eta_0) := \mathcal{M}^{\mathrm{split}}_{\tilde{\Gamma}}(L, \mathfrak{p}_{\tilde{\Gamma}}, \eta_0)/T_{\mathrm{trop}}(\tilde{\Gamma}, \Gamma).$$

The moduli space of framed split maps of type $\tilde{\Gamma}$ modulo the action of domain reparametrizations is denoted

$$\mathcal{M}^{\mathrm{split}}_{\mathrm{fr},\tilde{\Gamma}}(L, \mathfrak{p}_{\tilde{\Gamma}}, \eta_0).$$

The quotient $\mathcal{M}^{\mathrm{split}}_{\mathrm{fr},\tilde{\Gamma}}/T_{\mathrm{trop,fr}}(\tilde{\Gamma})$ is equal to the reduced moduli space $\mathcal{M}^{\mathrm{split}}_{\tilde{\Gamma},\mathrm{red}}(L, \eta_0)$. (Split maps and framed split maps are defined in Definitions 11.35, 11.40 respectively.) This ends the definition.

REMARK 11.43. If the split tropical graph $\tilde{\Gamma}$ is rigid then $T_{\mathrm{trop,fr}}(\tilde{\Gamma})$ is a finite group, and $\mathcal{M}^{\mathrm{split}}_{\mathrm{fr},\tilde{\Gamma}}(L, \eta_0)$ is a finite cover of the reduced moduli space.

REMARK 11.44. The expected dimension of strata of split maps can be computed as follows. Let $(\tilde{\Gamma}, \Gamma)$ be a type of split map. Let $\hat{\Gamma}$ be the broken map type obtained by collapsing all the tropical edges of $\tilde{\Gamma}$ that do not occur in Γ, that is we collapse the edges

$$e \in \mathrm{Edge}_{\mathrm{trop}}(\tilde{\Gamma}) \setminus \mathrm{Edge}(\Gamma).$$

Thus the tropical graph of $\hat{\Gamma}$ is equal to Γ, and the boundary edges of $\hat{\Gamma}$ are the same as that of $\tilde{\Gamma}$. Then the moduli space of split maps $\mathcal{M}^{\mathrm{split}}_{\tilde{\Gamma}}(L, \eta)$ with boundary end points $\underline{x} \in (\mathcal{I}(L))^{d(\circ)}$ has expected dimension

$$i^{\mathrm{split}}(\tilde{\Gamma}, \underline{x}) := i^{\mathrm{brok}}(\hat{\Gamma}, \underline{x}) + 2|\,\mathrm{Edge}_s(\Gamma)|(\dim(\mathfrak{t}) - 1).$$

Here i^{brok} is the index function for types of broken maps defined in (6.23). Indeed dropping the node matching condition (4.18) adds $2(\dim(\mathfrak{t})-1)$ dimensions for each split edge, and gluing along tropical edges e doesn't change dimension as proved in Proposition 6.25. For a rigid split map, the tropical symmetry group $T_{\mathrm{trop}}(\tilde{\Gamma}, \Gamma)$ has dimension $2|\operatorname{Edge}_s(\Gamma)|(\dim(\mathfrak{t})-1)$. The expected dimension of the quotiented moduli space is then

$$i^{\mathrm{split}}_{\mathrm{red}}(\tilde{\Gamma}, \underline{x}) := i^{\mathrm{split}}(\tilde{\Gamma}, \underline{x}) - 2|\operatorname{Edge}_s(\Gamma)|(\dim(\mathfrak{t})-1) = i^{\mathrm{brok}}(\hat{\Gamma}, \underline{x}).$$

This ends the Remark.

Proposition 11.45. (Moduli spaces of split maps) *Let $\eta_0 \in \mathfrak{t}^\vee$ be a generic element. Let $(\tilde{\Gamma}, \Gamma)$ be a type of an uncrowded split map, and suppose regular perturbation data for types $(\tilde{\Gamma}', \Gamma)$ of based treed disks with $\tilde{\Gamma}' < \tilde{\Gamma}$ is given (where the ordering on types is as in (6.26)). Then there is a comeager subset $\mathcal{P}^{reg}_{\tilde{\Gamma}} \subset \mathcal{P}_{\tilde{\Gamma}}$ of regular perturbations for split maps of type $\tilde{\Gamma}$ coherent with the previously chosen data such that if $\mathfrak{p}_{\tilde{\Gamma}} \in \mathcal{P}^{reg}_{\tilde{\Gamma}}$ then the following holds. Let $\underline{x} \in \mathcal{I}(L)^{d(\circ)}$ be boundary end points for which the index $i^{\mathrm{split}}_{\mathrm{red}}((\tilde{\Gamma}, \Gamma), \underline{x})$ is ≤ 1. Then*

(a) (Transversality) *the moduli space of split maps $\mathcal{M}^{\mathrm{split}}_{\tilde{\Gamma}}(\mathfrak{p}_{\tilde{\Gamma}}, L, \eta_0)$, framed split maps $\widetilde{\mathcal{M}}^{\mathrm{split}}_{\mathrm{fr},\tilde{\Gamma}}(\mathfrak{p}_{\tilde{\Gamma}}, L, \eta_0)$ and the reduced moduli space $\mathcal{M}^{\mathrm{split}}_{\tilde{\Gamma},\mathrm{red}}(\mathfrak{p}_{\tilde{\Gamma}}, L, \eta_0)$ are manifolds of expected dimension.*

(b) (Compactness) *The moduli space $\mathcal{M}^{\mathrm{split}}_{\tilde{\Gamma},\mathrm{red}}(\mathfrak{p}_{\tilde{\Gamma}}, L, \eta_0)$ is compact if $i^{\mathrm{split}}_{\mathrm{red}}(\tilde{\Gamma} \to \Gamma, \underline{x}) = 0$. If this index is 1, then the compactification $\overline{\mathcal{M}}^{\mathrm{split}}_{\tilde{\Gamma},\mathrm{red}}(\mathfrak{p}_{\tilde{\Gamma}}, L, \eta_0)$ consists of codimension one boundary points which contain a boundary edge $e \in \operatorname{Edge}_\circ(\tilde{\Gamma})$ with $\ell(e) = 0$ or $\ell(e) = \infty$.*

(c) (Tubular neighborhoods) *Let $\tilde{\Gamma}$ be a type of split map with a single edge e that is broken or has length $\ell(e)$ zero, and let the input/output tuple $\underline{x}$ be such that $i^{\mathrm{split}}_{\mathrm{red}}(\tilde{\Gamma}, \underline{x}) = 0$. Then $\mathcal{M}^{\mathrm{split}}_{\tilde{\Gamma},\mathrm{red}}(L, \eta_0)$ has a one-dimensional tubular neighborhood in any adjoining one-dimensional strata of split maps with the expected orientations. (The set of adjoining strata and the expected orientations are as in Theorem 9.7.)*

(d) (True boundary) *The true boundary of the oriented topological manifold*

$$\bigcup_{\tilde{\Gamma}:\tilde{\Gamma} \textit{ is rigid},\ i^{\mathrm{split}}_{\mathrm{red}}(\tilde{\Gamma}, \underline{x})=0} \overline{\mathcal{M}}^{\mathrm{split}}_{\tilde{\Gamma},\mathrm{red}}(L, \underline{x}).$$

is isomorphic to the reduced moduli space of split maps with a single broken edge, with the orientation sign adjusted by a factor $(-1)^\circ$ which depends only on the type $\tilde{\Gamma}$, $\underline{x}$ (and is the same as the corresponding sign for broken maps in Theorem 9.7 (bii)).

Proof of Proposition 11.45. The proof of the first statement is by a Sard-Smale argument similar to the case of broken maps. The only variation is that there is no matching condition (11.19) on the split edges.

We prove the compactness result for the moduli space of framed split maps since that is a finite cover over the reduced moduli space. The compactness statement for framed split maps can be proved in exactly the same way as the corresponding proof for broken maps. In fact, given a sequence of framed split maps u_ν of the same type, the limit is obtained by translations $\{t_\nu(v)\}_{v\in \mathrm{Vert}(\tilde{\Gamma})}$ that satisfy the slope condition (8.32) for all edges, including split edges. Consequently, the limit is a framed split map u. The codimension one boundary strata are as required in the Proposition via the arguments in the proof of Proposition 8.44. The finiteness of the tropical symmetry group used in the proof of Proposition 8.44 is replaced by the finiteness of the framed symmetry group $T_{\mathrm{trop,fr}}(\tilde{\Gamma})$ of the split tropical graph $\tilde{\Gamma}$.

The proofs of parts (c) and (d) are the same as the proofs of the corresponding results (Theorem 9.7 and Remark 9.8) for broken maps. □

PROPOSITION 11.46. *Given $E > 0$, there are finitely many types of split maps that have area $\leq E$.*

The proof is identical to the corresponding result Proposition 8.36 for broken maps, and is omitted.

Next we define the *tropical Fukaya algebra* whose composition maps are given by counts of rigid split maps. In order to obtain homotopy equivalence with other versions of Fukaya algebras, the definition must use framed split maps. We instead use a count of split maps in the reduced moduli space weighted by the size of the discrete tropical symmetry group, since the latter is a finite quotient of the former.

DEFINITION 11.47. Let $\eta_0 \in \mathfrak{t}^\vee$ be a generic cone direction, and let $\underline{\mathfrak{p}}$ be a regular perturbation datum for all based curve types. The *tropical Fukaya algebra* is the graded vector space

$$CF_{\mathrm{trop}}(L, \eta_0) := CF^{\mathrm{geom}}(L) \oplus \Lambda x^{\blacktriangledown}[1] \oplus \Lambda x^{\triangledown}$$

equipped with composition maps

$$m^{\mathrm{trop}}_{d(\circ)}(x_1, \ldots, x_{d(\circ)}) = \sum_{x_0, u \in \mathcal{M}_{\tilde{\Gamma},\mathrm{split,red}}(\mathfrak{X}, L, D, \underline{x})_0} \mathrm{mult}(\tilde{\Gamma}) w_s(u) x_0 \tag{11.22}$$

where

$$w_s(u) := (-1)^{\heartsuit} (d_\bullet(\Gamma)!)^{-1} (s(\tilde{\Gamma})!)^{-1} \, \mathrm{Hol}([\partial u]) \epsilon(u) q^{A(u)}, \tag{11.23}$$

where $s(\tilde{\Gamma})$ is the number of split edges in the type $\tilde{\Gamma}$ and the other symbols in (11.23) are as in (10.10). The orientation sign for split maps is defined in the same way as broken maps, see Remark 6.30. (Indeed, dropping a matching condition, which is a complex condition does not affect the determinant line bundle.) In (11.22), the combinatorial type $\tilde{\Gamma}$ of the split map u ranges over all rigid types with $d(\circ)$ inputs (see Definition 11.36 for rigidity), and $\mathrm{mult}(\tilde{\Gamma}) := |T_{\mathrm{trop,fr}}(\tilde{\Gamma})|$ is the multiplicity of the split map from (11.21).

REMARK 11.48. Any one-dimensional moduli space of split maps has true and fake boundary strata as in Figure 11.2, which look very similar to their analogue for broken maps shown in Figure 9.1.

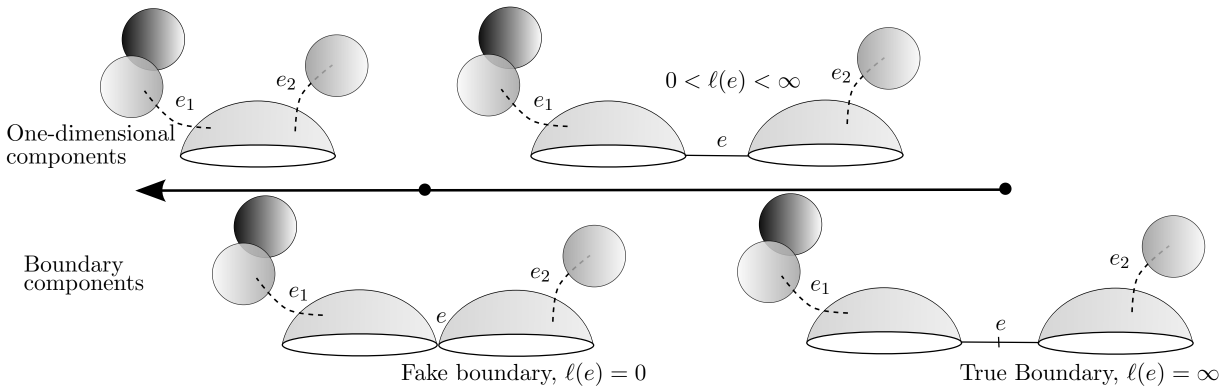

FIGURE 11.2. True and fake boundary strata of a one-dimensional component of the moduli space of split maps with split edges e_1, e_2. The sphere components lie in different pieces of the tropical manifold.

The composition maps $(m^{\text{trop}}_{d(\circ)})_{d(\circ)\geq 0}$ satisfy the A_∞ -axioms. Indeed, as in the broken and unbroken cases, the A_∞ -axioms are proved by a count of end-points of one-dimensional moduli spaces of rigid split maps; the boundary components are described by Proposition 11.45 (d); and the combinatorial factors arising from the distribution of interior markings are accounted exactly as in the unbroken case in Theorem 10.2. Therefore, $CF_{\text{trop}}(L, \underline{\mathfrak{p}}, \eta_0)$ is a A_∞ -algebra.

REMARK 11.49. The composition map m^{trop}_d can be expressed as a sum of products. The sum is over rigid types $\tilde{\Gamma}$ of split maps, and the product is over the connected components $\tilde{\Gamma}_1, \ldots, \tilde{\Gamma}_{s(\tilde{\Gamma})}$ of the graph $\tilde{\Gamma}\setminus \text{Edge}_s(\Gamma)$. Such a decomposition is possible because split maps do not have any edge matching conditions. Therefore, for any type $\tilde{\Gamma}$ of a rigid tropical map, the reduced moduli space $\mathcal{M}^{\text{split}}_{\tilde{\Gamma},\text{red}}(L, \underline{\eta})$ is a product of reduced moduli spaces :

$$\mathcal{M}^{\text{split}}_{\tilde{\Gamma},\text{red}}(L, \underline{\eta}) = \prod_{i=1}^{s(\tilde{\Gamma})} \mathcal{M}_{\tilde{\Gamma}_i,\text{red}}.$$

Here, $\mathcal{M}_{\tilde{\Gamma}_i}$ is the quotient

$$\mathcal{M}_{\tilde{\Gamma}_i,\text{red}} := \mathcal{M}_{\tilde{\Gamma}_i}/T_{\text{trop}}(\tilde{\Gamma}_i, \Gamma_i), \tag{11.24}$$

of the moduli space $\mathcal{M}_{\tilde{\Gamma}_i}$ of maps modelled on the subgraph $\tilde{\Gamma}_i \subset \tilde{\Gamma}$ by the component of the tropical symmetry arising from $\tilde{\Gamma}_i$ as in (11.18). Thus the composition map decomposes as

$$m^d_{\text{split}} = \sum_{\tilde{\Gamma}: d_\circ(\tilde{\Gamma})=d} m^d_{\text{split},\tilde{\Gamma}}, \quad m^d_{\text{split},\tilde{\Gamma}} = \frac{(-1)^{\heartsuit}\,\text{mult}(\tilde{\Gamma})}{d_\bullet(\tilde{\Gamma})!s(\tilde{\Gamma})!} \prod_{i=1}^{s(\tilde{\Gamma})} m_{\tilde{\Gamma}_i}, \tag{11.25}$$

where

- the sum ranges over all rigid types $\tilde{\Gamma}$ of split maps with d inputs
- and

$$m_{\tilde{\Gamma}_i} := \sum_{u \in \mathcal{M}_{\tilde{\Gamma}_i, \mathrm{red}}(L, \eta_0)_0} \mathrm{Hol}([\partial u]) \epsilon(u) q^{A(u)}$$

is a weighted count of rigid elements in $\mathcal{M}_{\tilde{\Gamma}_i,\mathrm{red}}$, where y, ϵ are as in (10.10).

We remark that in the product decomposition (11.24) of moduli spaces, exactly one of the components $\tilde{\Gamma}_0$ has boundary on the Lagrangian. For all other components $\tilde{\Gamma}_i$, $i \neq 0$, the domain components are spheres. Therefore, the number m_{Γ_i} is invariant under choices of perturbations, and may be viewed as a relative Gromov-Witten invariant. Finally, we point out that the decomposition in (11.24) is not a full splitting since any component $\tilde{\Gamma}_i$ may contain tropical nodes satisfying matching conditions.

11.6. Examples of split graphs and maps

We give examples of split tropical graphs and split maps to illustrate some of the definitions and results from the previous sections. We see the (Cone condition) in action and demonstrate the splitting of the tropical symmetry group for split tropical graphs. In most of our examples the multiple cut $\mathcal{P}$ consists of two orthogonally intersecting single cuts, so that the dual complex is a square as in Figure 11.3.

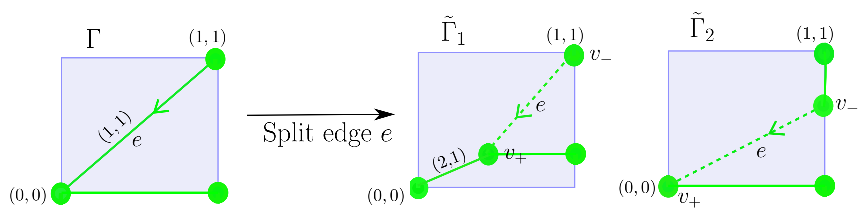

FIGURE 11.3. Split graphs $(\tilde{\Gamma}_1, \Gamma)$, $(\tilde{\Gamma}_2, \Gamma)$ with split edge e and a cone direction η_0 satisfying $\langle \eta_0, (1,-1) \rangle > 0$.

EXAMPLE 11.50. In this example we show that the quasi-split graphs $(\tilde{\Gamma}_1, \Gamma)$, $(\tilde{\Gamma}_2, \Gamma)$ in Figure 11.3 are both split graphs with the same cone direction, but they have different splittings of the tropical symmetry group. The torus T in this case is $(S^1)^2$. Since $\mathfrak{t}/\mathbb{R}\mathcal{T}(e)$ is one-dimensional, any $\eta_0 \in \mathfrak{t}^\vee$ satisfying $\langle \eta_0, (1,-1) \rangle \neq 0$ is (trivially) a generic cone direction. There are two non-equivalent choices of cone direction. The graphs in Figure 11.3 satisfy the cone condition with the cone direction η_0 satisfying $\langle \eta_0, (1,-1) \rangle > 0$.

(a) A relative vertex position $\mathcal{T} \in \mathcal{W}(\tilde{\Gamma}_1, \Gamma)$ satisfies

$$\mathcal{T}(v_+) = (2,1)t \quad \text{for some } t \geq 0, \qquad \mathcal{T}(v_-) = (0,0).$$

The discrepancy cone $\mathrm{Disc}(\tilde{\Gamma}_1, \Gamma)$ is

$$\mathbb{R}_{\geq 0} \pi^\perp_{\mathcal{T}(e)} \{ (\mathcal{T}(v_+) - \mathcal{T}(v_-)) : \mathcal{T} \in \mathcal{W}(\tilde{\Gamma}_1, \Gamma) \} = \mathbb{R}_{\geq 0}(\pi^\perp_{\mathcal{T}(e)} \eta_0).$$

For the connected component $\tilde{\Gamma}_1^{\pm} \subset \tilde{\Gamma}_1 \backslash \{e\}$ containing $v_{\pm}$, the identity component of the tropical symmetry group (see (11.18)) is

$$T_{\text{trop}}(\tilde{\Gamma}_1^+) = \exp(\mathbb{C}(2,1)), \quad T_{\text{trop}}(\tilde{\Gamma}_1^-) = \{\text{Id}\}.$$

(b) A relative vertex position $\mathcal{T} \in \mathcal{W}(\tilde{\Gamma}_2, \Gamma)$ satisfies

$$\mathcal{T}(v_+) = (0,0), \quad \mathcal{T}(v_-) = -(1,0)t \ \text{ for some } t \geq 0, .$$

and the discrepancy cone is

$$\mathbb{R}_{\geq 0}(\pi^{\perp}_{\mathcal{T}(e)}\{(\mathcal{T}(v_+) - \mathcal{T}(v_-)) : \mathcal{T} \in \mathcal{W}(\tilde{\Gamma}_2, \Gamma)\}) = \mathbb{R}_{\geq 0}(\pi^{\perp}_{\mathcal{T}(e)}\eta_0).$$

In $\tilde{\Gamma}_2$, the tropical symmetry group splits as

$$T_{\text{trop}}(\tilde{\Gamma}_2^+) = \{\text{Id}\}, \quad T_{\text{trop}}(\tilde{\Gamma}_2^-) = \exp(\mathbb{C}(1,0)).$$

Both split tropical graphs $\tilde{\Gamma}_1$, $\tilde{\Gamma}_2$ are rigid.

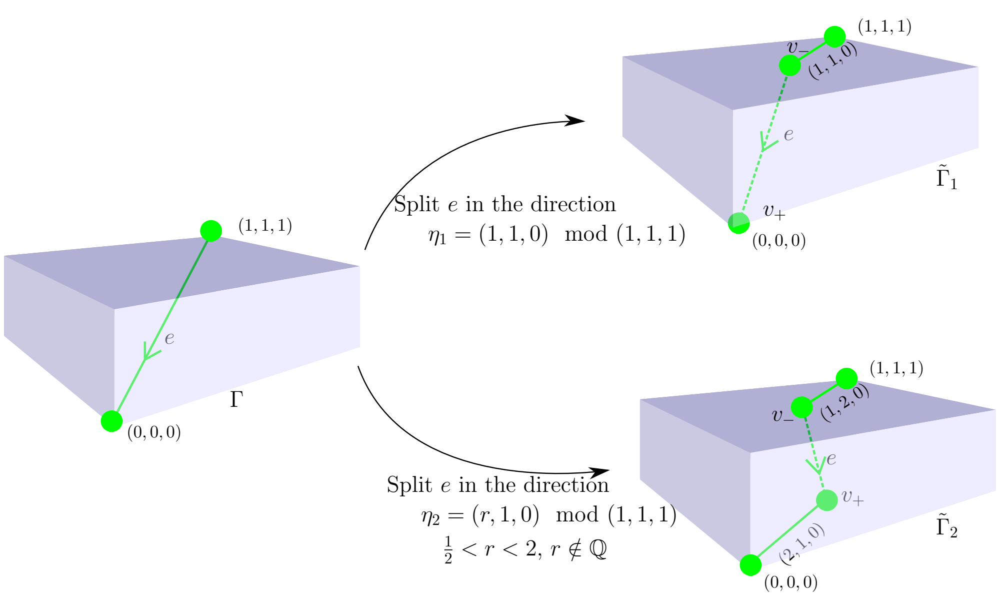

FIGURE 11.4. $(\tilde{\Gamma}_1, \Gamma)$ is not a split graph. $(\tilde{\Gamma}_2, \Gamma)$ is a split graph since the cone direction η_2 is generic.

EXAMPLE 11.51. In this example, we will see that the genericity of the cone direction is necessary to ensure that the discrepancy cone is top-dimensional. This example is also a case where the splitting of the tropical symmetry group contains two non-trivial factors. Consider a multiple cut $\mathcal{P}$ consisting of three orthogonal single cuts, so the structure torus is $T = (S^1)^3$, and the dual complex $B^{\vee} \subset \mathfrak{t}^{\vee}$ is a cube as in Figure 11.4. The quasi-split graphs $(\tilde{\Gamma}_1, \Gamma)$, $(\tilde{\Gamma}_2, \Gamma)$ in Figure 11.4 have a single split edge e and different cone directions.

(a) The quasi-split graph $(\tilde{\Gamma}_1, \Gamma)$ satisfies the (Cone condition) for the cone direction $\eta_1 := (1,1,0) \mod (1,1,1)$, which is not a generic direction in $\mathbb{R}^3/\langle(1,1,1)\rangle$. For any relative vertex position $\mathcal{T} \in \mathcal{W}(\tilde{\Gamma}_1)$,

$$\pi^{\perp}_{\mathcal{T}(e)}(\mathcal{T}(v_+)) = 0; \quad \pi^{\perp}_{\mathcal{T}(e)}(\mathcal{T}(v_-)) \in \{-(1,1,0)t_2 \mod (1,1,1) : t_2 \geq 0\}.$$

and the discrepancy cone is

$$\mathrm{Disc}(\tilde{\Gamma}_1, \Gamma) = \mathbb{R}_{\geq 0}\pi^{\perp}_{\mathcal{T}(e)}(\mathcal{T}(v_+) - \mathcal{T}(v_-)) = \{(1,1,0)t \mod (1,1,1) : t \geq 0\}.$$

The discrepancy cone contains the cone direction $(1,1,0)$, but is not top-dimensional in $\mathbb{R}^3/\langle(1,1,1)\rangle$. Consequently $\tilde{\Gamma}_1$ is not a split graph.

(b) The quasi-split tropical graph $(\tilde{\Gamma}_2, \Gamma)$ satisfies the (Cone condition) for a generic cone direction

$$\eta_2 := (r,1,0),\ r \notin \mathbb{Q},\ \tfrac{1}{2} < r < 2$$

and therefore $\tilde{\Gamma}_2$ is a split tropical graph. For any relative vertex position $\mathcal{T} \in \mathcal{W}(\tilde{\Gamma}_2)$,

$$\pi^{\perp}_{\mathcal{T}(e)}(\mathcal{T}(v_+)) \in \{(2,1,0)t_1 \mod (1,1,1) : t_1 \geq 0\};$$
$$\pi^{\perp}_{\mathcal{T}(e)}(\mathcal{T}(v_-)) \in \{-(1,2,0)t_2 \mod (1,1,1) : t_2 \geq 0\},$$

and the discrepancy cone is

$$\mathrm{Disc}(\tilde{\Gamma}_2, \Gamma) = \{(2,1,0)t_1 + (1,2,0)t_2 : t_1, t_2 \geq 0\},$$

which contains the cone direction $(r,1,0)$ and is top-dimensional in $\mathbb{R}^3/\langle(1,1,1)\rangle$. Let $\tilde{\Gamma}_\pm \subset \tilde{\Gamma}_2 \backslash \{e\}$ be the connected component containing $v_\pm$. The group of tropical symmetries is a product

$$T_{\mathbb{C}}/T_{\mathcal{T}(e),\mathbb{C}} = T_{\mathrm{trop}}(\tilde{\Gamma}^+) \times T_{\mathrm{trop}}(\tilde{\Gamma}^-),$$

where

$$T_{\mathrm{trop}}(\tilde{\Gamma}^+) = \exp(\mathbb{C}(2,1,0)), \quad T_{\mathrm{trop}}(\tilde{\Gamma}^-) = \exp(\mathbb{C}(1,2,0)).$$

Example 11.52. (Framed tropical symmetry group) The split tropical graph $\tilde{\Gamma}_2 \to \Gamma$ in Figure 11.4 is rigid, and consequently the framed tropical symmetry group $T_{\mathrm{trop,fr}}(\tilde{\Gamma}_2)$ is finite. However, this group is non-trivial. An element $(g,z) \in T_{\mathrm{trop,fr}}(\tilde{\Gamma}_2)$ satisfies the equations

$$g_{v_0} = g_{v_1} = \mathrm{Id}, \quad g_{v_0}g_{v_+}^{-1} = z_{e_+}^{(2,1,0)}, \quad g_{v_+}g_{v_-}^{-1} = z_e^{(1,1,1)}, \quad g_{v_-}g_{v_1}^{-1} = z_{e_-}^{(1,2,0)}.$$

There are 3 elements in $T_{\mathrm{trop,fr}}(\tilde{\Gamma}_2)$ given by $z_{e_+} = z_{e_-} = \omega$ where ω is a cube root of unity.

Example 11.53. A split tropical graph need not always have additional vertices compared to the base tropical graph. In some cases, dropping the slope condition (1.6) on split edges (and analogously the matching condition on the corresponding split nodes $w(e)$) produces a top-dimensional discrepancy cone $\mathrm{Disc}(\tilde{\Gamma})$. For example in Figure 11.5 dropping the slope condition at edge e in the graph Γ_1 produces one dimension of symmetry in the tropical graph $\tilde{\Gamma}_1$. That is, the graph $\tilde{\Gamma}_1$, which is Γ_1

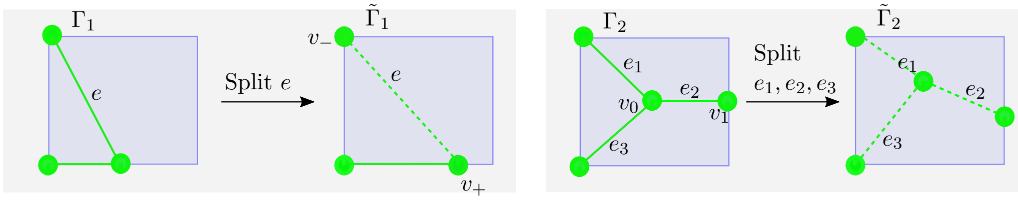

FIGURE 11.5. In these examples dropping the slope condition on split edges automatically produces split tropical graphs.

with the slope condition at e forgotten, is a split tropical graph with discrepancy cone $w(\tilde{\Gamma}_1) = \mathbb{R}$. Similarly dropping the slope condition on edges e_1, e_2, e_3 in Γ_2 in Figure 11.5 produces three dimensions of symmetry in $\tilde{\Gamma}_2$. Indeed in $\tilde{\Gamma}_2$, the vertex v_0 is free to move in two dimensions, and the vertex v_1 is independently free to move in one dimension. Therefore, $\tilde{\Gamma}_2$ is a split tropical graph with discrepancy cone $w(\tilde{\Gamma}_2) = \mathbb{R}^3$. In general, dropping the slope condition on split edges does not guarantee a top-dimensional discrepancy cone, such as in the graph Γ in Figure 11.3 and in Figure 11.6 below.

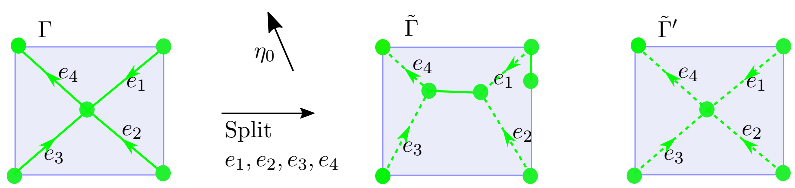

FIGURE 11.6. $(\tilde{\Gamma}, \Gamma)$ is a split graph with cone direction η_0 and split edges ordered as $e_1, \ldots, e_4$. $(\tilde{\Gamma}', \Gamma)$ is not a split tropical graph.

EXAMPLE 11.54. In this example we show that the increasing cone condition for split tropical graphs is necessary, and it can not be replaced by the condition that the cone direction is contained in the individual discrepancy cones of each of the split edges. For example, the quasi-split graph $(\tilde{\Gamma}', \Gamma)$ in Figure 11.6 satisfies the latter condition, namely, for each split edge $e = (v_+, v_-) \in \mathrm{Edge}_s(\tilde{\Gamma}')$,

$$\pi^{\perp}_{\mathcal{T}(e)}(\eta_0) \in \mathbb{R}_+\langle\{\mathcal{T}(v_+) - \mathcal{T}(v_-) : \mathcal{T} \in \mathcal{W}(\tilde{\Gamma}', \Gamma)\}\rangle,$$

but the discrepancy cone is not top-dimensional. In fact the discrepancy cone of $\tilde{\Gamma}'$ is two-dimensional since the vertex v_0 can only move in two dimensions and the other vertices are fixed. However, $(\tilde{\Gamma}, \Gamma)$ in Figure 11.6 satisfies the (Cone condition), and therefore, is a split tropical graph. We continue analyzing this graph in the next example.

EXAMPLE 11.55. The increasing cone condition for split graphs is equivalent to saying that we may split one edge at a time (following the ordering $\prec$), and at each step the cone direction η_0 is in the discrepancy cone. More precisely, given a tropical graph Γ with a set of ordered split edges $e_1, \ldots, e_n$, the increasing cone condition is equivalent to the existence of a sequence of split tropical graphs $(\tilde{\Gamma}_1, \Gamma), \ldots, (\tilde{\Gamma}_n, \Gamma)$ such that for any k the split edges of $(\tilde{\Gamma}_k, \Gamma)$ are $e_1, \ldots, e_k$, and for a small enough relative vertex position $\mathcal{T} \in \mathcal{W}(\tilde{\Gamma}_{k-1}, \Gamma)$,

$$\pi^{\perp}_{\mathcal{T}(e_k)}(\eta_0) \in \mathbb{R}_+\{\mathrm{Diff}_{e_k}(\mathcal{T}') : \mathcal{T}' \in \mathcal{W}(\tilde{\Gamma}_k, \Gamma),$$
$$\mathrm{Diff}_e(\mathcal{T}') = \mathrm{Diff}_e(\mathcal{T}) \text{ for } e = e_1, \ldots, e_{k-1}\}.$$

For the split tropical graph $(\tilde{\Gamma}, \Gamma)$ in Figure 11.6 the intermediate split graphs are given in Figure 11.7. Splitting each of the last two edges e_3, e_4 has the effect of 'forgetting the edge slope', and is therefore not shown separately in Figure 11.7.

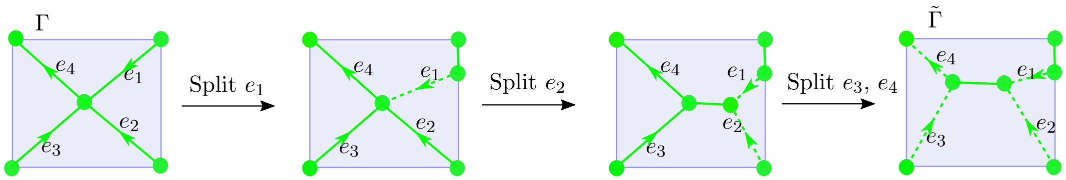

FIGURE 11.7. Splitting edges one at a time yields the increasing cone condition in the final graph.

EXAMPLE 11.56. This is an example of a split map in a multiple cut of the second Hirzebruch surface $X := H_2$ shown in Figure 11.8. We start by describing the moduli space of unbroken disks under consideration. Consider disks $u : (\mathbb{D}, \partial\mathbb{D}) \to (X, L)$ with no input leaves and a single output

- in the homology class $\delta_{F_2} + \delta_E$
- passing through a point p on F_2 that is close to the intersection point $F_2 \cap E$
- and whose boundary output maps to the maximum point of the Morse function F_L on L.

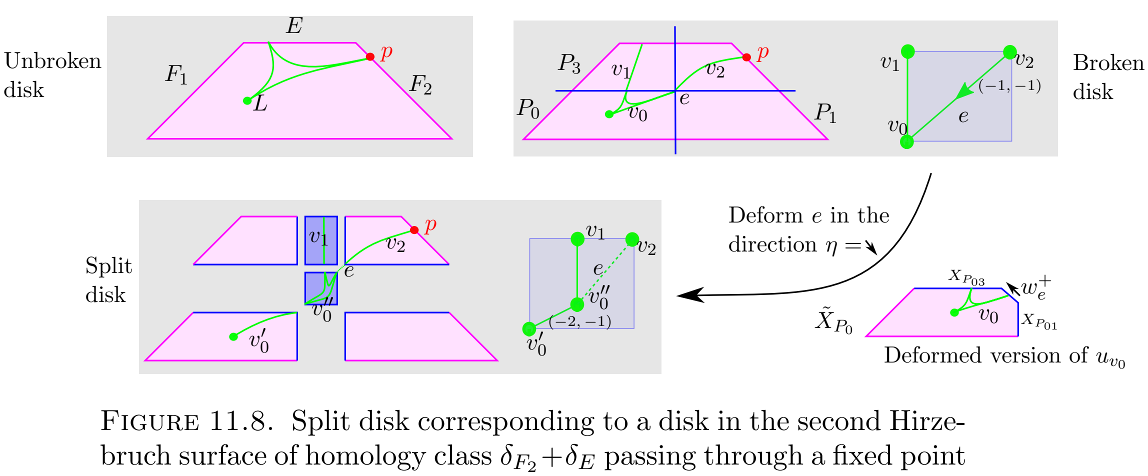

FIGURE 11.8. Split disk corresponding to a disk in the second Hirzebruch surface of homology class $\delta_{F_2} + \delta_E$ passing through a fixed point p.

The moduli space consists of a single disk with Maslov index 4. Next we apply a multiple cut on $X = H_2$ consisting of two orthogonal cuts. Each of the top-dimensional cut spaces is a first Hirzebruch surface, or in other words, a point blow-up of $\mathbb{P}^2$. Proposition 12.9 is applicable on this multiple cut, and it implies that neck-stretching preserves the potential (with an interior point constraint). Therefore, there is a single broken disk u_{brok} shown in Figure 11.8 that satisfies the constraints listed above, and whose gluing has homology class $\delta_{F_2} + \delta_E$. Next, we deform the matching condition at the edge e in the direction η. For any $\tau > 0$, in the $\tau\eta$-deformed map, the components u_{v_1}, u_{v_2} are the same as the corresponding components in the broken map u_{brok}. It is easier to understand the deformation of the component u_{v_0} if we view u_{v_0} as mapping to a different compactification of $X^{\square}_{P_0}$, namely one where the intersection of the relative divisors is blown up, and denoted by $\tilde{X}_{P_0}$. In this new compactification, the $\tau\eta$-deformation $u^{\tau}_{v_0}$ of u_{v_0} maps to a point x_τ on the exceptional divisor, and the point x_τ moves toward $X_{P_{03}}$ as $\tau \to \infty$. In the split disk, the limit of $u^{\tau}_{v_0}$, denoted by $u_{v_0'}$, intersects the divisors $X_{P_{03}}, X_{P_{01}} \subset X_{P_0}$ with multiplicity 2, 1 respectively.

11.7. Homotopy equivalence: deformed to split

In this section, we will show that deformed Fukaya algebras are A_∞ homotopy equivalent to tropical Fukaya algebras. We recall that for the tropical Fukaya algebra

$$CF_{\text{trop}}(\mathfrak{X}, L, \eta_0)$$

composition maps are counts of split disks with a generic cone direction $\eta_0 \in \mathfrak{t}^\vee$. On the other hand, defining a deformed Fukaya algebra

$$CF(\mathfrak{X}, L, \underline{\eta})$$

requires us to choose a coherent deformation datum $\underline{\eta} = (\eta_{\tilde{\Gamma}})_{\tilde{\Gamma}}$, which consists of a deformation parameter for each split edge

$$(\eta_{\tilde{\Gamma}}(e,[C]))_{e \in \text{Edge}_s(\tilde{\Gamma})} \in T_{\Gamma,\mathbb{C}} := \prod_{e \in \text{Edge}_s(\Gamma)} T_{\mathbb{C}}/T_{\mathcal{T}(e),\mathbb{C}} \tag{11.26}$$

varying smoothly with the domain curve C. To prove the homotopy equivalence we choose a sequence of deformation data $(\underline{\eta}_\nu)_\nu$ compatible with the cone direction η_0, which roughly means that for any fixed domain curve $[C]$ and any split edge e the deformation parameter $\eta_{\nu,\tilde{\Gamma}}([C]) \in T_{\mathbb{C}}/T_{\mathcal{T}(e),\mathbb{C}}$ goes to infinity in the direction η_0. Split maps with cone direction η_0 occur as limits of sequences of $\underline{\eta}_\nu$-deformed maps of bounded area. Conversely a rigid split map can be glued to yield a $\underline{\eta}_\nu$-deformed map for large enough ν. This bijective relation leads to the proof of the A_∞ homotopy equivalence.

We point out that for a converging sequence of deformed maps of index zero, the index of the limit split map is equal to $\dim(T_{\Gamma,\mathbb{C}})$ (defined in (11.26)), which is the dimension of its tropical symmetry group. The higher index of the limit split map is accounted for by the absence of matching conditions on split edges. Thus in this section, we produce a cobordism from the moduli space of framed split maps, which is a slice of the action of the tropical symmetry group on the moduli space of split maps, to the space of $\underline{\eta}$-deformed maps for any deformation datum $\underline{\eta}$.

We start by defining a sequence of deformation data compatible with the cone direction.

Definition 11.57. (Compatible deformation data for a cone direction) Given a generic cone direction $\eta_0 \in \mathfrak{t}^\vee$ and a type $(\tilde{\Gamma}, \Gamma)$ of based curves, a sequence of deformation data

$$\eta_{\tilde{\Gamma},\nu} : \mathcal{M}_{\tilde{\Gamma}} \to \mathfrak{t}_\Gamma \simeq \bigoplus_{e \in \mathrm{Edge}_s(\Gamma)} \mathfrak{t}/\mathfrak{t}_{\mathcal{T}(e)}$$

for a curve type $\tilde{\Gamma}$ with base Γ is compatible with the cone direction η_0 if

$$\eta_{\tilde{\Gamma},\nu} = c_{\tilde{\Gamma},\nu} \pi^\perp_{\mathcal{T}(e)}(\eta_0),$$

and

$$c_{\tilde{\Gamma},\nu} : \mathcal{M}_{\tilde{\Gamma}} \times \mathrm{Edge}_s(\Gamma) \to \mathbb{R}_+$$

is a continuous function such that for any converging sequence $m_\nu \to m$ in $\overline{\mathcal{M}}_{\tilde{\Gamma}}$, $(c_{\tilde{\Gamma},\nu}(m_\nu, e))_{e \in \mathrm{Edge}_s(\Gamma),\nu}$ is an increasing sequence of tuples (as in Definition 11.32).

For the gluing proof, we need a more restrictive class of deformation data.

Definition 11.58. (Uniformly continuous deformation data) A sequence of deformation data $\underline{\eta}$ compatible with a cone direction η_0 is *uniformly continuous* if for any split tropical graph $\tilde{\Gamma}$ there is a constant k such that for any two curves C, C' of type $\tilde{\Gamma}$

$$|\eta_{\tilde{\Gamma},\nu}([C]) - \eta_{\tilde{\Gamma},\nu}([C'])| \leq k \sum_{e \in \mathrm{Edge}_{\circ,-}} |\ell(T_e, [C]) - \ell(T_e, [C'])|$$

for all ν. We recall that $\ell(T_e, [C])$ is the length of the treed segment T_e in the curve C.

Next, we construct a sequence of deformation data that is coherent, compatible with the cone direction η_0 and uniformly continuous. We recall that moduli spaces of deformed maps are defined using domain-dependent deformation data that are coherent under morphisms of based curve types, and satisfy a (Marking independence) property, see Definition 11.6.

Lemma 11.59. *Given a generic cone direction $\eta_0 \in \mathfrak{t}^\vee$, there exists a compatible sequence $\{\underline{\eta}_\nu\}_\nu$ of uniformly continuous coherent deformation data (as in Definition 11.6).*

Proof. We first describe connected components of the moduli space of based curves. A connected component is determined by a tuple $\gamma := (\Gamma, m, n, \gamma_0)$ consisting of the base tropical graph Γ, the number m of boundary markings, the number n of interior markings, and a function

$$\gamma_0 : \{1, \dots, n\} \to \mathrm{Vert}(\Gamma)$$

that maps a marked point z_i on a curve component $C_v \subset C$ to $\kappa(v) \in \mathrm{Vert}(\Gamma)$, and we denote the component by $\mathcal{M}_\gamma$. The sequence of functions

$$c_{\gamma,\nu} : \mathcal{M}_\gamma \times \mathrm{Edge}_s(\Gamma) \to \mathbb{R} \tag{11.27}$$

is constructed by induction on m, n, $|\mathrm{Vert}(\Gamma)|$. We consider a connected component (Γ, n, m, γ), and assume that a sequence of coherent data has been constructed for smaller types. The data on the smaller types determines $c_{\gamma,\nu}$ on the true boundary of $\mathcal{M}_\gamma$, which consists of curves with at least a single broken boundary edge.

The deformation sequence on $\mathcal{M}_\gamma$ is constructed via a partition

$$\mathcal{M}_\gamma = \cup_{i\geq 0}\mathcal{M}^i_\gamma,$$

where $\mathcal{M}^i_\gamma$ consists of treed curves with exactly i boundary edges. On $\mathcal{M}^0_\gamma$, for any split edge e we fix $c_{\gamma,\nu}(m, e)$ to be a sequence of constant functions

$$c_{\gamma,\nu}(m, e) := c_\nu(e) \in \mathbb{R}_+$$

such that $(c_\nu(e))_{e\in\mathrm{Edge}_s(\Gamma)}$ is an increasing sequence of tuples (as in (11.10)) with respect to the ordering $\prec_{\tilde{\Gamma}}$ of split edges. Let

$$C_\nu := \sup_i \{c_\nu(m, e) : m \in \mathcal{M}^0_\gamma \cup \partial\mathcal{M}_\gamma\}.$$

We extend the deformation sequence to all of $\mathcal{M}_\gamma$ by interpolating between the codimension one strata $\partial\mathcal{M}^0_\gamma$ and $\partial\mathcal{M}_\gamma$. The boundary $\partial\mathcal{M}^0_\gamma$ partitions into sets

$$\partial\mathcal{M}^0_\gamma = \cup_{i\geq 1}\mathcal{M}^{0,i}_\gamma, \quad \mathcal{M}^{0,i}_\gamma := \overline{\mathcal{M}^0_\gamma} \cap \mathcal{M}^i_\gamma,$$

and $\mathcal{M}^i_\gamma$ is a product

$$\mathcal{M}^i_\gamma = \mathcal{M}^{0,i}_\gamma \times [0, \infty]^i, \tag{11.28}$$

such that projection to the second factor is equal to the edge length function. We first consider $\mathcal{M}^1_\gamma$. Under the splitting (11.28) we define

(11.29)
$$c_{\nu,e}(m, t) := (1 - \tau_{\nu,e}(t))c_{\nu,e}(m, 0) + \tau_{\nu,e}(t)c_{\nu,e}(m, \infty), \quad (m, t) \in \mathcal{M}^{0,1}_\gamma \times [0, \infty]$$

where $\tau_{\nu,e} : [0, \infty] \to [0, 1]$ is a diffeomorphism. In (11.29), $c_{\nu,e}$ satisfies the (Marking independence) property (11.2) because on each connected component it factors through the projection $\mathcal{M}^1_\gamma \to [0, \infty]$. Indeed $m \mapsto c_{\nu,e}(m, \infty)$ is locally constant on $\partial\mathcal{M}_\gamma \cap \mathcal{M}^1_\gamma$ by the (Marking independence) property (11.2) property for smaller strata, and $m \mapsto c_{\nu,e}(m, 0)$ is locally constant on $\partial\mathcal{M}_\gamma \cap \mathcal{M}^1_\gamma$ by definition (11.27). We obtain uniform continuity on $c_{\nu,e}$ by requiring that the derivative of $\tau_{\nu,e}$ is bounded by C^{-1}_ν for all ν. This can be arranged by allowing $d\tau_{\nu,e}$ to be supported in an interval of length C_ν. In a similar way, the functions c_{ν,e_j} are extended to $\mathcal{M}^i_\gamma$, assuming that its value on the boundary $\partial\mathcal{M}^i_\gamma$ factors through the projection to $[0, \infty]^i$. □

PROPOSITION 11.60. (Convergence) *Let $\eta_0 \in \mathfrak{t}^\vee$ be a generic cone direction, and let $\{\underline{\eta}_\nu\}_\nu$ be a coherent sequence of deformation data compatible with η_0 (as in Definition 11.57). Let $u_\nu : C_\nu \to \mathfrak{X}_\mathcal{P}$ be a sequence of $\underline{\eta}_\nu$-deformed maps with uniformly bounded area. Then, a subsequence converges to a framed split map $u_\infty : C \to \mathfrak{X}_\mathcal{P}$ of type $(\tilde{\Gamma}_\infty, \Gamma)$. The limit map is unique up to the action of the identity* *component of the framed tropical symmetry group* $T_{\text{trop,fr}}(\tilde{\Gamma}_\infty, \Gamma)$ *(see (11.20)). If the maps* u_ν *have index* 0*, and the perturbation datum for the limit is regular, then, the split tropical graph* $\tilde{\Gamma}_\infty$ *is rigid.*

If the limit split tropical graph is rigid, the limit is unique, because the framed symmetry group $T_{\text{trop,fr}}(\tilde{\Gamma}_\infty, \Gamma)$ is finite.

Proof of Proposition 11.60. As in the convergence of broken maps, the first step is to find the component-wise limit map. By Proposition 8.36, there are finitely many types of deformed maps that satisfy an area bound. Therefore, after passing to a subsequence, we can assume that the maps u_ν have a ν-independent type $\tilde{\Gamma} \to \Gamma$, with deformation datum

$$\eta_\nu(e) := \eta_{\tilde{\Gamma},\nu}([C_\nu], e) \in \mathfrak{t}/\mathfrak{t}_{\mathcal{T}(e)}, \quad \forall e \in \operatorname{Edge}_s(\Gamma).$$

We apply Gromov convergence for broken maps (Theorem 8.3) on each connected component of $\tilde{\Gamma} \setminus \operatorname{Edge}_s(\Gamma)$. The collection of the limit maps, denoted by u_∞, is modelled on a quasi-split tropical graph $\tilde{\Gamma}_\infty \to \Gamma$, equipped with a tropical edge collapse morphism $\kappa : \tilde{\Gamma}_\infty \setminus \operatorname{Edge}_s(\Gamma) \to \tilde{\Gamma} \setminus \operatorname{Edge}_s(\Gamma)$. The translation sequence

$$t_\nu(v) \in \operatorname{Cone}(\kappa, v), \quad v \in \operatorname{Vert}(\tilde{\Gamma}_\infty)$$

satisfies the (Slope) condition for all non-split edges $e \in \operatorname{Edge}(\tilde{\Gamma}_\infty) \setminus \operatorname{Edge}_s(\Gamma)$.

The limit map u_∞ is a quasi-split map, and it remains to show that the cone condition (11.9) is satisfied. The sequence of relative translations required by the cone condition will be a refinement of the translation sequences t_ν. For any split edge e, the sequences

$$((t_\nu(v_+) - t_\nu(v_-)) \mod \mathcal{T}(e)) - \eta_\nu(e), \quad e = (v_+, v_-) \in \operatorname{Edge}_s(\Gamma) \tag{11.30}$$

are uniformly bounded. This fact can be proved in the same way as the proof of boundedness of the sequences in (8.38) as part of the proof of convergence of broken maps. The next step is to adjust the translation sequences by a uniformly bounded amount so that the discrepancy at the split edges is equal to the deformation parameters, i.e. the sequences in (11.30), vanish. The adjustment to the translation sequences is by the iterative process used in Step 2 of the proof of Lemma 8.25. We run the iteration on the vertex sequence $\{t_\nu(v)\}_{v,\nu}$ and the edge sequence $\{\eta_\nu(e)\}_{e,\nu}$. In each step a fastest growing sequence is subtracted, while the quantities in (11.30) are left unchanged. At any step, if both a vertex sequence and an edge sequence grow at the same rate, and are both fastest growing, then we break the tie by subtracting the edge sequence. The resulting sequence in the i-th step is denoted by $\{t^i_\nu(v)\}_v$, $\{\eta^i_\nu(e)\}_e$. The iteration ends when the sequences corresponding to vertices and split edges are uniformly bounded, say, after k steps. Our tie breaking method, together with the fact that the sequence corresponding to any pair of split edges grows at a different rate (since $(c_{e,\nu})_{e,\nu}$ is a sequence of increasing tuples), ensures that $\eta^k_\nu(e) = 0$ for all split edges e and all ν. We subtract the uniformly bounded final sequence $\{t^k_\nu(v)\}_{v,\nu}$ from the initial sequence $\{t_\nu(v)\}_{v,\nu}$ and denote the result by $\{\bar{t}_\nu(v)\}_{v,\nu}$, which satisfies

$$((t_\nu(v_+) - t_\nu(v_-)) \mod \mathcal{T}(e)) - \eta_\nu(e) = 0, \quad e = (v_+, v_-) \in \operatorname{Edge}_s(\Gamma). \tag{11.31}$$

Since for any vertex v the translation sequences $t_\nu(v)$, $\bar{t}_\nu(v)$ differ by a uniformly bounded amount, a subsequence of $e^{-\bar{t}_\nu}u_\nu$ converges to a limit $\overline{u}_{\infty,v}$. The limit map $\overline{u}_\infty := (\overline{u}_{\infty,v})_v$ satisfies the matching condition on all edges, including split edges, because by (11.31) the translations $\bar{t}_\nu$ compensate for the deformation parameters. Finally $\tilde{\Gamma} \to \Gamma$ is a split graph because

$$\mathrm{Diff}(\bar{t}_\nu) = \eta_\nu \implies \eta_\nu \in \mathrm{Disc}(\tilde{\Gamma}, \Gamma) \quad \forall \nu.$$

We have shown that $\overline{u}_\infty$ is a framed split map. The proof of uniqueness of limits is the same as in the convergence of broken maps. The last statement about the rigidity of the limit split map follows from a dimensional argument, see Remark 11.44. □

Gluing a rigid split map with a cone direction η_0 yields a sequence of deformed maps for any sequence of uniformly continuous deformation data compatible with η_0. We assume that the perturbation datum for the deformed maps is equal to the perturbation datum of the split map.

Remark 11.61. In Lemma 11.59, we constructed a sequence of uniformly continuous deformation data $\{\underline{\eta}_\nu\}_{\nu \in \mathbb{Z}_+}$ that is compatible with a cone direction η_0. By the same proof, we can also a construct a family of uniformly continuous deformation data $\{\underline{\eta}_\nu\}_{\nu \in \mathbb{R}_+}$ such that for any sequence $\nu_i \to \infty$, $\{\underline{\eta}_{\nu_i}\}_i$ is compatible with the cone direction η_0.

Proposition 11.62. (Gluing a split map) *Suppose u is a regular rigid framed split map of type $(\tilde{\Gamma}, \Gamma)$ with a generic cone direction $\eta_0 \in \mathfrak{t}^\vee$. Let $(\underline{\eta}_\nu)_{\nu \in \mathbb{R}_+} = \{\eta_{\tilde{\Gamma},\nu}\}_{\tilde{\Gamma},\nu}$ be a family of uniformly continuous coherent deformation data compatible with η_0 (as in Definition 11.57). Then, there exists ν_0 such that for $\nu \geq \nu_0$*

- *(a)* (Existence of glued family) *there is a regular rigid $\underline{\eta}_\nu$-deformed map u_ν such that the family $\{u_\nu\}_\nu$ of deformed maps converges to u as $\nu \to \infty$.*
- *(b)* (Surjectivity of gluing) *For any sequence of $\underline{\eta}_\nu$-deformed maps u'_ν that converges to the framed split map u, for large enough ν the map u'_ν is contained in the glued family constructed in (a).*

Proof. We first describe the type of the deformed maps in the glued family. Given that the type of the split map u is $\kappa : \tilde{\Gamma} \to \Gamma$, the type of the glued deformed maps is $(\tilde{\Gamma}_d, \Gamma)$. That is, the base tropical graph Γ stays the same, and $\tilde{\Gamma}_d$ is obtained by tropically collapsing the edges $e \in \mathrm{Edge}_{\mathrm{trop}}(\tilde{\Gamma}) \setminus \mathrm{Edge}(\Gamma)$. Thus, the tropical edges in $\tilde{\Gamma}_d$ are exactly the edges of the base tropical graph Γ, and the edge collapse map κ factors as

$$\tilde{\Gamma} \xrightarrow{\kappa_0} \tilde{\Gamma}_d \xrightarrow{\kappa_1} \Gamma.$$

The tree part of $\tilde{\Gamma}_d$ is the same as that of $\tilde{\Gamma}$. We will construct a family of deformed maps of type $\tilde{\Gamma}_d \to \Gamma$.

The gluing construction is very similar to the gluing of broken maps in Theorem 9.1, so we only point out the differences at each step of the proof.

Step 1: Construction of an approximate solution: An approximate solution is constructed using relative vertex positions corresponding to a deformation parameter. The domain C_ν of the approximate solution is constructed by gluing some of the interior nodes in C, and therefore the edge lengths for treed segments T_e, $e \in \mathrm{Edge}_{\circ,-}(\tilde{\Gamma})$ in C_ν are the same as C. By the (Marking independence property) of deformation data, $\eta_{\nu,\tilde{\Gamma}_d}(C) = \eta_{\nu,\tilde{\Gamma}_d}(C_\nu)$, and we denote

$$\eta_\nu := \eta_{\nu,\tilde{\Gamma}_d}(C) \in \mathfrak{t}_\Gamma,$$

where we recall from (11.1) that $\mathfrak{t}_\Gamma = \bigoplus_{e \in \mathrm{Edge}_s(\Gamma)} \mathfrak{t}/\mathfrak{t}_{\mathcal{T}(e)}$. The sequence

$$(\eta_\nu(e))_{e \in \mathrm{Edge}_s(\tilde{\Gamma}),\nu}$$

is an increasing sequence of tuples (as in Definition 11.32). Therefore the (Cone condition) for the split tropical graph $\tilde{\Gamma}$ implies that for large enough ν, η_ν lies in the discrepancy cone of $\tilde{\Gamma}$, and so, there is a relative translation $\mathcal{T}_\nu \in w(\tilde{\Gamma}, \Gamma)$ satisfying

$$\pi^\perp_{\mathcal{T}(e)}(\mathcal{T}_\nu(v_+) - \mathcal{T}_\nu(v_-)) = \eta_\nu(e), \quad e = (v_+, v_-) \in \mathrm{Edge}_s(\Gamma). \tag{11.32}$$

Since the split tropical graph $\tilde{\Gamma}$ is rigid, the relative translation $\mathcal{T}_\nu \in w(\tilde{\Gamma}, \Gamma)$ is uniquely determined by (11.32). Indeed, if there is another solution $\mathcal{T}'$ of (11.32), the difference $\mathcal{T} - \mathcal{T}'$ satisfies the slope condition on all edges, including split edges, and therefore generates a non-trivial subgroup $\exp((\mathcal{T} - \mathcal{T}')(\cdot))$ in $T_{\mathrm{trop,fr}}(\tilde{\Gamma})$ contradicting the rigidity of $\tilde{\Gamma}$. For future use, we point out that the inverse

$$\mathrm{Diff}^{-1}_{\tilde{\Gamma}} : \mathfrak{t}_\Gamma \to w^\pm(\tilde{\Gamma}, \Gamma) := \mathbb{R}\langle w(\tilde{\Gamma}, \Gamma)\rangle, \tag{11.33}$$

which maps deformation parameters to elements in the $\mathbb{R}$-span of relative translations, is a well-defined linear map. For any edge $e = (v_+, v_-)$ that is collapsed by $\tilde{\Gamma} \to \tilde{\Gamma}_d$, we can assign a length $l_\nu(e) > 0$ satisfying

$$\mathcal{T}_\nu(v_+) - \mathcal{T}_\nu(v_-) = l_\nu(e)\mathcal{T}(e).$$

We now describe the domain and target spaces for the approximate solution. The domain is a nodal curve C_ν of type $\tilde{\Gamma}_d$ which is obtained from C by replacing each node corresponding to an edge $e \in \mathrm{Edge}_{\mathrm{trop}}(\tilde{\Gamma}) \setminus \mathrm{Edge}_{\mathrm{trop}}(\tilde{\Gamma}_d)$ by a neck of length $l_\nu(e)$. A component of the glued curve corresponds to a vertex v of $\tilde{\Gamma}_d$. The map $u|C_{\kappa_0^{-1}(v)}$ is a broken map with relative marked points, whose pieces map to $X_{\overline{Q}}$, $Q \subseteq P(v)$. Our goal is to glue at the nodes of $u|C_{\kappa_0^{-1}(v)}$ to produce a curve lying in $X_{\overline{P(v)}}$.

We extend the definition of the deformation parameter to all tropical edges of $\tilde{\Gamma}_d$ by defining

$$\eta_\nu(e) = 0 \quad \forall e \in \mathrm{Edge}_{\mathrm{trop}}(\tilde{\Gamma}_d) \setminus \mathrm{Edge}_s(\Gamma).$$

Thus by definition a η_ν-deformed map satisfies the edge matching condition for non-split edges of Γ.

Next, we construct the approximate solution which will shown to be η_ν-deformed. Since u is a framed split map, it satisfies a matching condition on the split edges:

For any edge $e = (v_+, v_-)$ in $\tilde{\Gamma}$, and domain coordinates in the neighborhood of $w_e^{\pm}$ that respect the framing, the tropical evaluation maps (defined in (4.19)) satisfy

$$\mathrm{ev}_{w_e^+}^{\mathcal{T}(e)}(u_{v_+}) = \mathrm{ev}_{w_e^-}^{\mathcal{T}(e)}(u_{v_-}). \tag{11.34}$$

For a vertex v in $\tilde{\Gamma}_d$, and a vertex v' in the collapsed graph $\kappa_0^{-1}(v)$, the translation $\mathcal{T}_\nu(v')$ gives an identification (see (3.41))

$$e^{\mathcal{T}_\nu(v')} : X^{\square}_{\overline{P}(v')} \to X^{\square}_{\overline{P}(v)}.$$

The translated map

$$u_{v',\eta_\nu} := e^{\mathcal{T}_\nu(v')} u_{v'} : C^\circ_{v'} \to X^{\square}_{\overline{P}(v)} \tag{11.35}$$

is well-defined on the complement of nodal points on $C_{v'}$. For any uncollapsed tropical edge $e = (v_+, v_-)$ in $\tilde{\Gamma} \to \tilde{\Gamma}_d$, the matching condition (11.34) and the translation in (11.35) together imply a deformed matching condition

$$\mathrm{ev}_{w_e^+}^{\mathcal{T}(e)}(u_{v_+,\eta_\nu}) = e^{\eta_\nu(e)}\, \mathrm{ev}_{w_e^-}^{\mathcal{T}(e)}(u_{v_-,\eta_\nu}), \tag{11.36}$$

because $\eta_\nu(e) = \mathcal{T}_\nu(v_+) - \mathcal{T}_\nu(v_-)$. As in the proof of Theorem 9.1, for each vertex v, translated maps $(u_{v',\eta_\nu})_{v' \in \kappa_0^{-1}(v)}$ can be glued at the nodal points to yield an approximate solution for the holomorphic curve equation which is denoted by

$$u_\nu^{\mathrm{pre}} = (u_{v,\eta_\nu}^{\mathrm{pre}})_{v \in \mathrm{Vert}(\tilde{\Gamma}_d)}, \quad u_{v,\eta_\nu}^{\mathrm{pre}} : C_{\eta_\nu, v} \to X_{\overline{P}(v)}.$$

Since the patching does not alter the maps u_{v,η_ν} away from the collapsed nodes of $\tilde{\Gamma} \to \tilde{\Gamma}_d$, we conclude by (11.36) that the pre-glued map u_ν^{pre} is η_ν-deformed.

Step 2: Fredholm theory: The Sobolev norms carry over entirely from the proof of Theorem 9.1. On the curves C_ν, we use Sobolev weights on both neck regions created by gluing, and on nodal points corresponding to interior edges in $\tilde{\Gamma}_d$. The map $\mathcal{F}_\nu$ in Theorem 9.1 incorporated the holomorphicity condition, marked points mapping to divisors and matching at disk nodes. Now we additionally require a deformed matching condition on interior nodes, given by

$$\mathrm{ev}_e(u) \in \Delta_e \subset (X_{\overline{P}(e)})^2, \quad e \in \mathrm{Edge}_{\mathrm{trop}}(\tilde{\Gamma}_d),$$

where Δ_e is the diagonal and ev_e is the tropical evaluation map on the nodal lifts twisted by the deformation parameter, that is,

$$\begin{aligned} \mathrm{ev}_e : \mathcal{M}_{\mathrm{def},\tilde{\Gamma}_d} \times \mathrm{Map}(C_\nu, \mathfrak{X})_{1,p,\lambda} &\to (X^{\square}_{\overline{P}(e)})^2, \\ (m, u) &\mapsto (\mathrm{ev}_{w_+(e)}^{\mathcal{T}(e)}(u_{v_+})), \exp(\eta_{\nu,\tilde{\Gamma}_d}(m,e))(\mathrm{ev}_{w_-(e)}^{\mathcal{T}(e)}(u_{v_-})). \end{aligned}$$

Incorporating these conditions we obtain a map

$$\mathcal{F}_\nu : \mathcal{M}^i_{\mathrm{def},\tilde{\Gamma}_d} \times \Omega^0(C_\nu, (u_\nu^{\mathrm{pre}})^* T\mathfrak{X})_{1,p} \to \Omega^{0,1}(C_\nu, (u_\nu^{\mathrm{pre}})^* T\mathfrak{X})_{0,p} \oplus \mathrm{ev}_{\tilde{\Gamma}}^* T\mathfrak{X}(\tilde{\Gamma}_d)/\Delta(\tilde{\Gamma}_d)$$

whose zeros correspond to η_ν-deformed pseudoholomorphic maps near the approximate solution u_ν^{pre}, and where the notations $\mathfrak{X}(\tilde{\Gamma}_d)$ and $\Delta(\tilde{\Gamma}_d)$ are defined in (6.29).

Step 3: Error estimate: The error estimate for the approximate solution is produced in the same way as Theorem 9.1. Indeed the only contribution to the error estimate is from the failure of holomorphicity, as the approximate solution satisfies the matching conditions at boundary and interior nodes.

The next few steps of the proof, namely, the construction of a uniformly bounded right inverse, the proof of quadratic estimates, and the application of Picard's Lemma are the same as in Theorem 9.1.

Step 7: Surjectivity of gluing: Compared to the corresponding step in Theorem 9.1, in the gluing of a split map, one additionally has to deal with domain-dependent deformation parameters. Consider a sequence $u'_\nu : C'_\nu \to \mathfrak{X}$ of $\underline{\eta}_\nu$-deformed maps that converges to the split map u. To prove that the maps u'_ν lie in the image of the gluing map of u, it is enough to show that u'_ν is close enough to the pre-glued map $u_\nu^{pre} : C_\nu \to \mathfrak{X}$. Indeed, similar to the proof of Theorem 9.1, one of the conclusions of Picard's Lemma is that the glued map u_ν is the unique $\underline{\eta}_\nu$-deformed map in an ϵ-neighborhood of the pre-glued map u_ν^{pre}, and ϵ is independent of ν. The sequences of domains $(C_\nu)_\nu$ and $(C'_\nu)_\nu$ converge to C in the compactified moduli space $\overline{\mathcal{M}}_{\mathrm{def},\tilde{\Gamma}_d}$, and so, the edge lengths $\ell(C_\nu, T_e)$, $\ell(C'_\nu, T_e)$ of the treed segments converge to $\ell(C, T_e)$ for all boundary edges $e \in \mathrm{Edge}_{\circ,-}(\tilde{\Gamma}_d)$. The edge length $\ell(C, T_e)$ in the split map u is finite by the rigidity of u. Therefore the differences converge

$$|\ell(C'_\nu, T_e) - \ell(C_\nu, T_e)| \to 0.$$

The deformation parameters for u'_ν are $\eta'_\nu := \eta_{\tilde{\Gamma},\nu}(C'_\nu)$. On a fixed stratum $\tilde{\Gamma}_d$ of curves, the deformation parameter depends only on the lengths of the treed edges by the (Marking independence) property (11.2), see (11.2). Together with the uniform continuity of the deformation datum, we conclude $|\eta_\nu - \eta'_\nu| \to 0$. By the proof of convergence of deformed maps, the translation sequence $\{t'_\nu(v)\}_{\nu,v}$ for the convergence of $\{u'_\nu\}_\nu$ can be chosen to satisfy

$$t_\nu(v_+) - t_\nu(v_-) = \eta'_\nu(e), \quad \forall e = (v_+, v_-),$$

see (11.31). The inverse $\mathrm{Diff}_{\tilde{\Gamma}}^{-1}$ is a well-defined map (see (11.33)) that maps

$$\eta_\nu \mapsto \mathcal{T}_\nu, \quad \eta'_\nu \mapsto t_\nu.$$

Hence $\mathcal{T}_\nu - t_\nu \to 0$. As in Theorem 9.1 the closeness of the relative translations $\mathcal{T}_\nu$, t_ν implies that the difference between the gluing parameters of C_ν, C'_ν (that is, the gluing parameters used to construct C_ν, C'_ν from C) converge to 0. The rest of the proof is the same as the proof of Theorem 9.1. □

PROPOSITION 11.63. (Bijection of moduli spaces) *Let $\eta_0 \in \mathfrak{t}^\vee$ be a generic cone direction and let $\{\underline{\eta}_\nu\}_\nu$ be a family of uniformly continuous coherent deformation data compatible with η_0. For any set of disk input/outputs $\underline{x}$ and area level $E > 0$, there exists ν_0 such that for $\nu \geq \nu_0$ there is a bijection between the moduli space of rigid framed split disks $\mathcal{M}_{\mathrm{fr}}^{\mathrm{split}}(\underline{\mathfrak{p}}, \eta_0)_{<E}$ and the moduli space $\mathcal{M}_{\mathrm{def}}(\underline{\mathfrak{p}}, \underline{\eta}_\nu)_{<E}$ of rigid $\underline{\eta}_\nu$-deformed disks of area less than E.*

PROOF. The proposition follows from the convergence and gluing results for split maps, namely Proposition 11.60 and Proposition 11.62. □

Proposition 11.64. *Let $\eta_0 \in \mathfrak{t}^\vee$ be a generic cone direction and let $\{\underline{\eta}_\nu : \nu \in \mathbb{R}_+\}$ be a family of uniformly continuous coherent deformation data compatible with η_0. Let $\mathfrak{p}^\infty$ be a regular perturbation datum for split maps on $\mathfrak{X}$. For any $E_0 > 0$, there exists $\nu_0(E_0)$ such that the following holds. For any $\nu \in \mathbb{Z}_{>0}$, there exists a regular perturbation datum $\mathfrak{p}^\nu$ for $\underline{\eta}_\nu$-deformed maps, and a perturbation morphism $\mathfrak{p}^{\nu,\nu+1}$ extending $\mathfrak{p}^\nu$ and $\mathfrak{p}^{\nu+1}$ such that for all $E_0 > 0$ and $\nu \geq \nu_0(E_0)$, the A_∞ morphism induced by $\mathfrak{p}^{\nu,\nu+1}$*

$$\phi_\nu : CF^{\mathrm{def}}(L, \underline{\mathfrak{p}}^\nu, \underline{\eta}_\nu) \to CF^{\mathrm{def}}(L, \underline{\mathfrak{p}}^{\nu+1}, \underline{\eta}_{\nu+1})$$

is equal to the identity modulo q^{E_0}. Similarly there is an A_∞ morphism

$$\psi_\nu : CF^{\mathrm{def}}(L, \underline{\mathfrak{p}}^{\nu+1}, \underline{\eta}_{\nu+1}) \to CF^{\mathrm{def}}(L, \underline{\mathfrak{p}}^\nu, \underline{\eta}_\nu)$$

that is identity modulo q^{E_0}.

Proof. This result is an analogue of Proposition 10.21 (b) for maps on neck-stretched manifolds. The parameter ν is analogous to the neck length parameter in Proposition 10.21. For any integer $\nu \in \mathbb{Z}_+$ the perturbation $\mathfrak{p}^\nu$ is taken to be equal to $\mathfrak{p}^\infty$ for strata of low area, that is, the strata for which the bijection between the moduli spaces of split maps and deformed maps (from Proposition 11.63) holds. The perturbation datum $\mathfrak{p}^\nu$ is extended to higher area strata by the transversality result for deformed maps (Proposition 11.10). To define the A_∞ morphism ϕ_ν, the family of deformation data interpolating between $\underline{\eta}_\nu$ and $\underline{\eta}_{\nu+1}$ is taken to be the restriction of the family $\{\underline{\eta}_\mu : \mu \in \mathbb{R}_+\}$ given in the hypothesis of the Proposition to $\mu \in [\nu, \nu+1]$. The perturbation morphism $\mathfrak{p}^{\nu,\nu+1}$ is defined to be equal to $\mathfrak{p}^\infty$ for low area strata, and extended in a regular coherent way to strata of higher area. Following the same arguments as in the proof of Proposition 10.21, we conclude that there are no non-constant $\mathfrak{p}^{\nu,\nu+1}$-adapted deformed quilted disks with area at most E_0, and therefore,

$$\phi_1^\nu = \mathrm{Id} \mod q^{E_0}, \quad \phi_k^\nu \mod q^{E_0} = 0, \;\; k \neq 1.$$

Here the identity term in ϕ_1^ν counts constant quilted maps with a single input. □

Proposition 11.65. (Homotopy equivalence) *Let η_0 be a generic cone direction, and let $\underline{\eta}_\nu$ be a compatible sequence of coherent deformation data that is uniformly continuous. For any $\nu_0 \in \mathbb{Z}_+$, let $\underline{\mathfrak{p}}_{\nu_0}$ be the perturbation defined in Proposition 11.64 for $\underline{\eta}_\nu$-deformed maps. Then, there are strictly unital convergent A_∞ morphisms*

$$\phi : CF_{\mathrm{def}}(L, \underline{\eta}_{\nu_0}, \underline{\mathfrak{p}}_{\nu_0}) \to CF_{\mathrm{trop}}(L, \eta_0, \underline{\mathfrak{p}}^\infty), \quad \psi : CF_{\mathrm{trop}}(L, \eta_0, \underline{\mathfrak{p}}^\infty) \to CF_{\mathrm{def}}(L, \underline{\eta}_{\nu_0}, \underline{\mathfrak{p}}_{\nu_0}).$$

such that $\psi \circ \phi$ and $\phi \circ \psi$ are A_∞ -homotopy equivalent to identity.

Proof. The A_∞ morphisms ψ, ϕ are defined as the limits of the morphisms ψ_ν, ϕ_ν in Proposition 11.64. The definition of the limit and the justification of its existence are exactly as in the proof of Proposition 10.22, which is the analogous result for broken maps.

As in the proof of Proposition 11.64, the homotopy equivalence of $\psi \circ \phi$ and $\phi \circ \psi$ to identity uses twice-quilted disks. The definition of twice-quilted disks is extended to twice-quilted deformed disks in a standard way. Indeed, the 'distance to the seam' function is constant on any disk and its connected sphere components, and these determine the deformation parameter to be used for any split edge. With this principle, the discussion on homotopy equivalences in [**18**, Theorem 5.10] carries over to the case of deformed maps, and one can show that there are homotopies satisfying

$$\phi_k \circ \psi_k - \mathrm{Id} = m_1(h_k), \quad \psi_k \circ \phi_k - \mathrm{Id} = m_1(g_k).$$

As in the proof of Proposition 11.64, the infinite composition

$$h = \lim_{k\to\infty} h_k, \quad g = \lim_{k\to\infty} g_k$$

exists and gives a homotopy equivalence between $\phi \circ \psi$ resp. $\psi \circ \phi$ and the identities. □

11.8. The cone displacement formula

The cone condition on split tropical graphs resembles the Fulton-Sturmfels formula [**33**] for the degeneration of the diagonal in a product $X \times X$ of a toric variety X. We describe the Fulton-Sturmfels formula here.

First, we describe a degeneration of the diagonal in the special case of a toric variety possessing an invariant Morse function. For a compact oriented manifold X with a Morse-Smale pair (f, g), the homology class of the diagonal $\Delta_X \subset X \times X$ admits a decomposition

$$[\Delta_X] = \sum_{x\in \mathrm{crit}(f)} [\overline{W}^s(x)] \times [\overline{W}^u(x)],$$

where $W^s(x), W^u(x) \subset X$ are respectively the stable and unstable submanifolds of a critical point x of the Morse function $f : X \to \mathbb{R}$. Now suppose that X is a toric variety with moment polytope P, and a component of the moment map of X, namely

$$\langle \Phi, \eta\rangle : X \to \mathbb{R}, \quad \eta \in \mathfrak{t},$$

is a Morse function. Then any critical point is a torus-fixed point $X_Q \in X_P$, and its stable and unstable manifolds are torus-invariant submanifolds X_{Q_+}, X_{Q_-} and correspond to faces Q_+, Q_- of the polytope P. Thus there is a decomposition of the diagonal

$$[\Delta_X] = \sum_{Q\in P : X_Q \in \mathrm{crit}(f)} [X_{Q_+}] \times [X_{Q_-}]. \tag{11.37}$$

Here we note that the homology classes of the torus-invariant submanifolds $X_Q \subset X$, $Q \subset P$ span $H_*(X)$. There are many ways to express the diagonal Δ_X as the sum of products of the classes $[X_Q]$; the choice of a generic $\eta \in \mathfrak{t}$ gives a decomposition of the form (11.37).

EXAMPLE 11.66. In the case of projective space, the decomposition (11.37) of the diagonal is just the usual Künneth decomposition

$$[\Delta_{\mathbb{P}^n}] = \sum_{i=1}^{n} [\mathbb{P}^i] \times [\mathbb{P}^{n-i}].$$

However, generic components of moment maps in toric varieties do not always form a Morse-Smale pair. For example, in a two-point blow-up of $\mathbb{P}^2$, Φ_ξ is not Morse-Smale (when paired with a torus-invariant metric) for any $\xi \in \mathfrak{t}$. Indeed for any ξ, there will be a gradient flow line connecting two critical points of index 2.

Given a generic vector $\eta \in \mathfrak{t}$, the Fulton-Sturmfels formula is a generalization of the decomposition of the diagonal in (11.37) to the cases when a generic component of a moment map does not form a Morse-Smale pair. For a toric variety X there is a decomposition formula corresponding to every generic vector $\eta \in \mathfrak{t}$: Let $\operatorname{Cone}(Q_\pm) \subset \mathfrak{t}$ denote the cones dual to the faces $Q_\pm \subset P := \Phi(X)$. We write

$$(\operatorname{Cone}(Q_-) + \eta) \cap \operatorname{Cone}(Q_+) = \{\mathrm{pt}\} \tag{11.38}$$

if the intersection between the displaced cones is a transversely-cut-out point, and if so, let

$$n(Q_-, Q_+) = \frac{(\operatorname{span}\operatorname{Cone}(Q_-) \cap \mathfrak{t}_{\mathbb{Z}}) + (\operatorname{span}\operatorname{Cone}(Q_+) \cap \mathfrak{t}_{\mathbb{Z}})}{\mathfrak{t}_{\mathbb{Z}}} \tag{11.39}$$

be the quotient of lattices in $\mathfrak{t}$, which is necessarily finite.

THEOREM 11.67. (Cone displacement formula, [**33**]) *There exists an equivalence in the Chow group*

$$\Delta \sim \sum_{(\operatorname{Cone}(Q_-)+\eta)\cap\operatorname{Cone}(Q_+)=\{\mathrm{pt}\}} n(Q_-, Q_+) X_{Q_-} \times X_{Q_+}.$$

REMARK 11.68. (Relation to the cone condition for split maps) We describe the relation between the Fulton-Sturmfels formula and the cone condition on split tropical graphs. The transverse intersection condition (11.38) for pairs $\operatorname{Cone}(Q_+), \operatorname{Cone}(Q_-) \subset \mathfrak{t}$ in the Fulton-Sturmfels formula implies that the cones are transverse, their difference $\operatorname{Cone}(Q_+) - \operatorname{Cone}(Q_-)$ contains the generic vector η, and consequently the difference $\operatorname{Cone}(Q_+) - \operatorname{Cone}(Q_-)$ generates the torus T. We may say that a toric variety assigns, via the Fulton-Sturmfels formula, a number $n(Q_+, Q_-)$ to every pair of transverse cones in $\mathfrak{t}$ that generate T, and whose difference contains η. In the split map situation, the analogues of $\operatorname{Cone}(Q_+)$, $\operatorname{Cone}(Q_-)$ are the discrepancy cones for the components of the split tropical graph $\tilde{\Gamma} \setminus \operatorname{Edge}_s(\Gamma)$. The quotient $n(Q_+, Q_-)$ is replaced by the multiplicity of the split tropical graph.

For a split map u with a single split edge and discrepancy cones $\mathcal{C}_\pm$, if there is a toric variety X and a critical point $Q \in \operatorname{crit}(\Phi_\eta)$ whose cones $\operatorname{Cone}(Q_\pm)$ are equal to $\mathcal{C}_\pm$ then we expect the multiplicity of the split map u to be equal to the number $n(Q_+, Q_-)$ from the Fulton-Sturmfels formula applied to X. However, even the case of a single split edge exhibits more complicated phenomena than the Fulton-Sturmfels formula for toric varieties: The toric variety is a compactification of the torus $(\mathbb{C}^\times)^n$, and the choice of compactification determines the numbers $n(Q_+, Q_-)$

for pairs of transverse cones Q_+, Q_-. On the other hand there is no fixed compactification of ends when we study convergence of deformed maps. As a result there is a larger set of cone pairs that occur as discrepancy cones of a split tropical graph with a single split edge. This set of cones grows larger as the area of the maps increases.

When there are multiple split edges, our cone direction is chosen so that the ratios between the vectors for different split edges are arbitrarily small. Splitting the tropical graph with this increasing condition is equivalent to splitting one edge at a time (as explained in Example 11.55), and the discussion in the previous paragraph is valid at each of the splitting steps individually.

CHAPTER 12

Disk potentials and unobstructedness

In this chapter we finish the proofs of applications of the theory given in Chapters 1 and 2.

12.1. Unobstructedness

We apply the tropical Fukaya algebra to the question of obstructedness. We introduce the following notation. Let $P_0 \in \mathcal{P}$ be a top-dimensional polytope and let $L \subset X_{P_0} \subset \mathfrak{X}_{\mathcal{P}}$ be an embedded Lagrangian submanifold.

DEFINITION 12.1. The Lagrangian L is called a *tropical moment fiber* if

- X_{P_0} is a compact toric manifold, and the tropical moment map $\Phi_{P_0} : X_{P_0} \to \mathfrak{t}^\vee$ is an honest moment map, and
- $L = \Phi_{P_0}^{-1}(\lambda)$ for a point λ in the interior of the polytope P_0,

We prove the following Corollary which was stated in Chapter 1. We refer the reader to Section 1.6 for the definition of weak unobstructedness of a Lagrangian brane, and the consequences of being weakly unobstructed.

COROLLARY 1.7. (Unobstructedness of a tropical torus) *Suppose that $L \subset X_{P_0}$ is a tropical torus equipped with a brane structure and that all the facets of P_0 are elements of $\mathcal{P}$. For any generic cone direction $\eta_0 \in \mathfrak{t}^\vee$, we have*

$$m^0_{CF_{\mathrm{trop}}(\mathfrak{X},L,\eta_0)} = wx^{\blacktriangledown}$$

where $x^{\blacktriangledown}$ is the unique maximum point of the Morse function on L and $w \in \Lambda_{>0}$ is an element of the positive part of the Novikov ring. In the tropical Fukaya algebra $CF_{\mathrm{trop}}(\mathfrak{X}, L, \eta_0)$,

$$b := wx^{\blacktriangledown} \in MC(L)$$

is a solution of the Maurer-Cartan equation, and the potential of the b-deformed A_∞ algebra $CF_{\mathrm{trop}}(\mathfrak{X}, L, \eta_0, b)$ is

$$W(b) := w1_L.$$

Consequently $CF_{\mathrm{trop}}(\mathfrak{X}, L, \eta_0)$, and hence $CF(X, L)$, are weakly unobstructed.

PROOF. We prove unobstructedness for the tropical Fukaya algebra $CF_{\mathrm{trop}}(L)$. The tropical Fukaya algebra is homotopy equivalent to the unbroken Fukaya algebra $CF(L)$ by Theorem 1.4. Unobstructedness is preserved under homotopy equivalence, see [**18**, Lemma 5.2].

We first describe the perturbation datum for the moduli space of split maps. Let J_0 be the standard almost complex structure on $X_{\overline{Q}}$ for all $Q \subseteq P_0$. Since torus-invariant divisors are relative divisors, Proposition 6.32 implies that J_0-holomorphic disks and spheres in X_{P_0} are regular. Therefore, there is a coherent regular perturbation datum $\underline{\mathfrak{p}}$ on $\mathfrak{X}$ for which the almost complex structure on X_{P_0} is standard. Next, we show that $m^0(1)$ is a multiple of the geometric unit $x^{\blacktriangledown}$. Consider a split disk $[u]$ of type $\tilde{\Gamma} \to \Gamma$ contributing to $m_0(1)$, and whose boundary output asymptotes to $x_0 \in \mathcal{I}(L)$. Let

$$\tilde{\Gamma}_1 \subset \tilde{\Gamma} \backslash \operatorname{Edge}_s(\Gamma)$$

be the connected component containing the disk components, and let $u_1 := u|_{\tilde{\Gamma}_1}$. Let u' be the split disk obtained by forgetting the boundary output leaf, and let $\tilde{\Gamma}' \to \Gamma$ be the type of u'. For any torus element $t \in T$, the split map

$$(u'_t)_v := \begin{cases} tu'_v, & v \in \operatorname{Vert}(\tilde{\Gamma}_1), \\ u_v, & v \in \operatorname{Vert}(\tilde{\Gamma}) \backslash \operatorname{Vert}(\tilde{\Gamma}_1). \end{cases}$$

is not contained in the $T_{\mathrm{trop}}(\tilde{\Gamma}, \Gamma)$-orbit of u'. Indeed the tropical symmetry group has a trivial action on X_{P_0}, which contains the disk components. By the same reason, for any pair $t_1 \neq t_2 \in T$, the maps u'_{t_1}, u'_{t_2} lie in distinct $T_{\mathrm{trop}}(\tilde{\Gamma}', \Gamma)$-orbits. The regularity of u implies u'_t is regular for all $t \in T$, Thus the tropical isomorphism classes $\{[u'_t] : t \in T\}$ form a $\dim(T)$-dimensional family in the reduced moduli space, and therefore.

$$\dim(\mathcal{M}^{\mathrm{split}}_{\tilde{\Gamma}',\mathrm{red}}(L)) \geq \dim(T).$$

Since for the original split map u, the dimension of the reduced moduli space of split maps $\mathcal{M}^{\mathrm{split}}_{\tilde{\Gamma},\mathrm{red}}(L)$ is zero, the output is necessarily the geometric unit $x^{\blacktriangledown}$. So, $m_0(1) = W x^{\blacktriangledown}$ for some $W \in \Lambda_{>0}$.

The existence of a solution to the projective Maurer-Cartan equation now follows: We first claim that $m_1(x^{\triangledown})$ only has zero order terms. If not, let u be a split map with non-zero area contributing to $m_1(x^{\triangledown})$. By the locality axiom, and the standardness of the perturbation datum on X_{P_0}, we conclude that forgetting the input in u produces a regular split disk u' with no inputs. Since the index of u' can not be negative, we conclude u can not exist for dimension reasons. Indeed, the index of u is two more than the index of u' : one from the choice of a boundary incoming marking, and one from the weight on the incoming leaf. Therefore,

$$m_1(x^{\triangledown}) = x^{\triangledown} - x^{\blacktriangledown}.$$

By a similar argument,

$$m_d(x^{\triangledown}, \ldots, x^{\triangledown}) = 0 \quad d \geq 2,$$

and consequently, $W x^{\triangledown}$ is a solution of the projective Maurer-Cartan equation. □

The unobstructedness result in tropical Fukaya algebras gives an alternate proof of the Fukaya-Oh-Ohta-Ono result [35] on unobstructedness of toric Lagrangians in toric manifolds.

COROLLARY 12.2. (Unobstructedness in a toric manifold) *Suppose X is a symplectic toric manifold with an action of a compact torus T, and moment map $\Phi : X \to \mathfrak{t}^\vee$. Then any Lagrangian moment fiber in X is weakly unobstructed.*

PROOF. Let X be a T-toric variety whose moment polytope is

$$\Delta := \Phi(X) = \{x \in \mathfrak{t}^\vee : \langle \mu_i, x\rangle \leq c_i, i = 1, \ldots, N\}$$

where $\mu_i \in \mathfrak{t}$ is the primitive outward pointing normal of the i-th facet of Δ, and $c_i \in \mathbb{R}$. To show that a moment fiber $\Phi^{-1}(\lambda)$ is unobstructed, we consider a broken manifold $\mathfrak{X}_{\mathcal{P}}$ obtained by the following cuts

$$\langle x, \mu_i\rangle = c_i - \epsilon_i, \quad i = 1, \ldots, N$$

where $\epsilon_i > 0$ is a small constant. In particular, we assume that $\langle \lambda, \mu_i\rangle \leq c_i - \epsilon_i$ and the piece $\{\langle x, \mu_i\rangle \geq c_i - \epsilon_i\}$ is a $\mathbb{P}^1$-fibration for all i. The fiber $\Phi^{-1}(\lambda) \subset X_{P_0}$ is indeed a tropical torus because a toric invariant divisor of X_{P_0} is of the form

$$\{\, x \mid \langle x, \mu_i\rangle = c_i - \epsilon_i \,\},$$

and is a relative divisor. Therefore, by Corollary 1.7 the Lagrangian $\Phi^{-1}(\lambda)$ is weakly unobstructed. □

12.2. Tropical disk potentials

In a toric manifold leading order terms in the Batyrev-Givental potential ([**8**], [**38**]) of a toric Lagrangian correspond to index two Maslov disks. In the notation of the proof of Corollary 12.2 the Batyrev-Givental potential for the Lagrangian $L := \Phi^{-1}(\lambda)$ is given by

$$W_{BG}(y_1, \ldots, y_n) = \sum_{i=1}^{N} \left(\prod_j y_j^{\mu_{i,j}} \right) q^{c_i - \langle \lambda, \mu_i\rangle}, \tag{12.1}$$

where $\mu_i = (\mu_{i,1}, \ldots, \mu_{i,n})$ is the primitive outward normal of the i-th facet of the moment polytope. The next result shows that for a tropical Lagrangian L in a broken manifold $\mathfrak{X}$ the potential $W_{CF_{\mathrm{trop}}(\mathfrak{X},L)}$ of the tropical Fukaya algebra $CF_{\mathrm{trop}}(\mathfrak{X}, L)$ contains all the terms of the terms of the Batyrev-Givental potential, and any other term in $W_{CF_{\mathrm{trop}}(\mathfrak{X},L)}$ is higher order in the formal variable q compared to at least one of the terms in the Batyrev-Givental potential. In the statement of the following proposition, the 'leading order term' of a polynomial $W \in \Lambda_{>0}$ refers to the non-zero terms with the least exponent of q.

PROPOSITION 12.3. (Leading order terms of the potential) *Let $\mathcal{P}$ be the polyhedral decomposition of the toric manifold X in the proof of Corollary 12.2, and let $\mathfrak{X}_{\mathcal{P}}$ be the corresponding broken manifold. The leading order terms of the potential of $CF_{\mathrm{trop}}(\mathfrak{X}_{\mathcal{P}}, L)$ coincide with the leading order terms of the Batyrev-Givental potential.*

REMARK 12.4. The tropical disk potential $W_{CF_{\mathrm{trop}}(\mathfrak{X},L)}$ may contain higher order terms that are not present in the Batyrev-Givental potential. For example, in

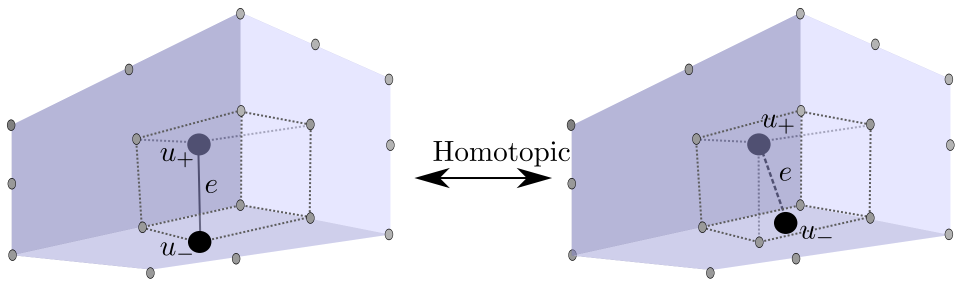

FIGURE 12.1. Proposition 12.3 : The tropical graph of the broken and split map in the dual polytope. The edge e is a split edge.

the semi-Fano toric surface X^ϵ (see Figure 2.1) that arises by deforming the symplectic form in the cubic surface, the Batyrev-Givental potential only has 9 terms, corresponding to each of the toric divisors. The tropical potential $W_{CF_{\mathrm{trop}}(\mathfrak{x},L)}$ has twenty one terms.

PROOF OF PROPOSITION 12.3. The terms in the Batyrev-Givental potential correspond to disks of Maslov index 2 that have an intersection of 1 with one of the toric divisors, and do not intersect the other toric divisors. We will show that the potential of the tropical Fukaya algebra consists of a term corresponding to each of these disks.

To describe the polytopes in $\mathcal{P}$, we recall the multiple cut from Corollary 12.2. We assume that the moment polytope of the T-toric variety X is

$$\Delta := \Phi(X) = \{x \in \mathfrak{t}^\vee : \langle \mu_i, x\rangle \leq c_i, i = 1, \dots, N\},$$

where $\mu_i \in \mathfrak{t}$ is the primitive outward pointing normal of the i-th facet of Δ, and $c_i \in \mathbb{R}$. The multiple cut $\mathcal{P}$ consists of simultaneous single cuts along the hypersurfaces

$$\langle x, \mu_i\rangle = c_i - \epsilon_i, \quad i = 1, \dots, N$$

where $\epsilon_i > 0$ is a small constant. There is a top-dimensional polytope $P_F \in \mathcal{P}$ corresponding to every face $F \subseteq \Delta$, and two of these polytopes P_{F_1}, P_{F_2} intersect exactly if either $F_1 \subseteq F_2$ or $F_2 \subseteq F_1$. Further, any polytope in $\mathcal{P}$ is of the form

$$P_{F_1 F_2} = P_{F_1} \cap P_{F_2}, \quad F_1 \subseteq F_2$$

for a pair of faces F_1, F_2 in Δ, and $\dim(P_{F_1F_2}) = \mathrm{codim}_{F_2}(F_1)$.

We use the following shorthand notation for cut spaces

$$X_{F_1F_2} := X_{P_{F_1F_2}}, \quad X_F := X_{P_F}$$

for faces F, F_1, F_2 of Δ, and a further abbreviation for some cut spaces

$$X_0 := X_\Delta, \quad X_i := X_{D_i}, \quad X_{0i} := X_{(D_i,\Delta)}.$$

where D_i is a facet of Δ, and $1 \leq i \leq N$. The dual complex $B^\vee$ consists of a top-dimensional polytope for every corner in Δ, and a zero-dimensional cell for every face of Δ. For example, if Δ is a three-dimensional cube, the dual polytope $B^\vee$ is as in Figure 12.1, where the grey dots are 0-cells, and the dotted lines bound one of the top-dimensional cells.

Corresponding to every toric divisor of X there is a broken disk whose glued type has Maslov index two, which we now describe. The perturbation datum used to define the broken Fukaya algebra on $\mathfrak{X}_{\mathcal{P}}$ is such that the almost complex structure on X_0 is standard, and the pieces $X_1, \ldots, X_N$ are fibrations

$$\pi_i : X_i \to X_{0i} \tag{12.2}$$

whose fibers are holomorphic spheres, each of which intersects the relative divisor $X_{0i} \simeq X_i \cap X_0$ at a single point. The index two disk incident on the i-th toric divisor has two components : $u_{\text{brok}} = (u_+, u_-)$. In this pair $u_+ : \mathbb{D} \to X_0$ is an index two Maslov disk intersecting the relative divisor X_{0i} at $p_i \in X_{0i}$, and u_- maps to X_i and is a fiber of π_i in (12.2).

The broken disk u_{brok} is homotopic to a family of deformed disks

$$T_{\mathbb{C}}/T_{\mathcal{T}(e),\mathbb{C}} \ni \tau \mapsto u^\tau, \quad u_0 = u.$$

Here $u^\tau := (u_+, u_-^\tau)$ is a τ-deformed map, and u_-^τ is a sphere in X_{P_i} homotopic to u_- that satisfies

$$p_i := \operatorname{ev}_{X_{0i}}(u_+) = e^\tau \operatorname{ev}_{X_{0i}}(u_-^\tau), \tag{12.3}$$

where $\operatorname{ev}_{X_{0i}}$ is the ordinary evaluation map at the lift of the nodal point $w_\pm(e)$ mapping to X_{0i}. The component u_+ stays constant under variation of τ because it is a disk of Maslov index two in the toric variety X_0, and is therefore rigid.

The split map is the limit of a sequence of deformed broken maps whose deformation parameters τ_ν approach the infinite end of the torus in a generic direction. As $\nu \to \infty$, the sequence of points $e^{-\tau_\nu} p_i \in X_{0i}$ approaches a T-fixed point $X_{p,D_i} \in X_{0i}$, where $p \in D_i$ is a vertex on the facet D_i. Here we have used the observation that all T-fixed points in X_{0i} are relative submanifolds. Consequently, the sequence u^τ converges to a split map $u^\infty = (u_+, u_-^\infty)$ where u_-^∞ maps to the neck piece $X_{\overline{P}_{(p,D_i)}}$. The tropical graph underlying u^∞ satisfies the cone condition for the following reason. The quasi-split tropical graph $\tilde{\Gamma}$ underlying u^∞ has vertices v_+, v_- corresponding to u_+, u_-^∞, which are connected by a split edge e. The vertex v_- is free to move in the dual polytope $P^\vee_{(p,D_i)}$ which is $n-1$-dimensional (see Figure 12.1). Thus the set of relative vertex positions $\mathcal{W}(\tilde{\Gamma}, \Gamma)$ is $(\dim(\mathfrak{t})-1)$-dimensional. The area and boundary holonomy of u^∞ are equal to the area and boundary holonomy of u, and therefore, u^∞ makes the expected contribution to the potential $m_0(1)$ of $CF_{\text{trop}}(\mathfrak{X}_{\hat{P}}, L)$.

Finally we show that all other terms in the potential of $CF_{\text{trop}}(\mathfrak{X}_{\hat{P}}, L)$ are of higher order in q than at least one of the disks in the last paragraph. Denote the split disk intersecting the divisor $D_i \subset X$ (from the last paragraph) by $u_i := (u_{i,+}, u_{i,-}^\infty)$. Suppose u is a split disk that contributes to the potential, and is not equal to u_i for any i. There are two possibilities:

(a) The disk part of u, denoted by $u_\circ : C \to X_0$ is not a Blaschke disk of Maslov index two. Since all the torus-invariant divisors of X_0 are relative divisors, we may conclude that $u_\circ$ intersects more than one torus-invariant divisors, say X_{0i} and X_{0j}, transversely. Since the constants ϵ_i, ϵ_j are small (or in other words the cut is close to the toric divisor of X), we conclude that the area of $u_\circ$ is larger than both the split disks u_i, u_j, and thus u contributes to a higher order term.

(b) The disk part of u, denoted by $u_\circ : C \to X_0$, is a Blaschke disk of Maslov index two intersecting the relative divisor X_{0i}, but the other component(s) is not the fiber sphere in X_i. Again, this configuration has a larger area than the split disk $u_i = (u_{i,+}, u^\infty_{i,-})$, since $u_\circ = u_{i,+}$, $u^\infty_{i,-}$ is homologous to a fiber sphere of area ϵ_i in X_i, and in the manifold X_i the fiber sphere has the smallest area amongst all non-constant symplectic spheres.

This finishes the proof of Proposition 12.3. □

12.3. Potential on semi-Fano toric surfaces

Recall that in Section 2.1 we counted disks in a cubic surface by deforming the symplectic form on the cubic surface to that of a *semi-Fano* toric surface. A complex manifold X is semi-Fano if the first Chern number $c_1(TX|S)$ is non-negative on any holomorphic curve $S \subset X$. In this section, we show that the disk potential of a toric Lagrangian in a semi-Fano toric surface is well-defined if we use a perturbation datum for which the almost complex structure is close enough to a divisor-preserving almost complex structure – these are almost complex structures for which the torus-invariant divisors are holomorphic. We also show that under some homological hypotheses, the potential is preserved by a multiple cut on a semi-Fano toric surface. These results justify our strategy in Section 2.1 to obtain the potential by counting broken disks.

Definition 12.5. Let X be a toric symplectic manifold. A tamed almost complex structure J on X is called *divisor-preserving* if all toric divisors Y of X are J-holomorphic, that is, $J(TY) = TY$.

Definition 12.6. (Potential for a non-regular almost complex structure) Suppose J_0 is a compatible almost complex structure on a symplectic manifold (M, ω) and $L \subset M$ be a Lagrangian with a brane structure. Let $\underline{F} = (F_\Gamma)_\Gamma$ be a coherent regular domain-dependent Morse function on L. The manifold (X, J_0, L) with Morse datum $\underline{F}$ has potential

$$W_0(J_0)1^\blacktriangledown, \quad W_0(J_0) \in \Lambda_{>0} \tag{12.4}$$

if the following holds: Given an area level E_0, there exists $\epsilon > 0$ such that for any coherent regular perturbation $\underline{\mathfrak{p}} = (\underline{J}, \underline{F})$ such that $\underline{J}$ maps to an ϵ-neighborhood of J_0,

$$\mathfrak{m}^0_{CF(L,\underline{\mathfrak{p}})} = W_0 1^\blacktriangledown \mod q^{E_0}.$$

Proposition 12.7. (Potential on semi-Fano toric manifolds) *Let X be a semi-Fano symplectic toric manifold and $L \subset X$ is a toric Lagrangian. Let J_0 be a divisor-preserving almost complex structure, and let $\underline{F}$ be a generic coherent domain-dependent Morse function on L. The potential for the manifold (X, J_0, L) is well-defined in the sense of Definition 12.6.*

PROOF. We first show that one can define m^0 for the divisor-preserving almost complex structure J_0.

CLAIM. There is an element $m^0_{J_0} \in \Lambda\langle \mathrm{Crit}(F)\rangle$ such that for any area level E_0 there exists $\epsilon > 0$ such that for any coherent regular perturbation $\underline{\mathfrak{p}}$ that is ϵ-close to J_0,

$$\mathfrak{m}^0_{CF(L,\underline{\mathfrak{p}})} = m^0_{J_0} \mod q^{E_0}.$$

To prove the Claim, let us assume the contrapositive. Then there is a sequence of coherent regular perturbations $\underline{\mathfrak{p}}^\nu$ that converge uniformly to the constant J_{std}, and such that

$$\mathfrak{m}^0_{CF(L,\underline{\mathfrak{p}}^\nu)} \neq \mathfrak{m}^0_{CF(L,\underline{\mathfrak{p}}^{\nu+1})} \mod q^{E_0}$$

for all ν. Choose a regular perturbation morphism

$$\underline{\mathfrak{p}}^{\nu,\nu+1} = (\underline{J}^{\nu,\nu+1}, \underline{F})$$

with end points $\underline{\mathfrak{p}}^\nu$, $\underline{\mathfrak{p}}^{\nu+1}$, such that the sequence $\underline{J}^{\nu,\nu+1}$ converges uniformly to the constant J_0 as $\nu \to \infty$. Since the potentials are unequal for $\underline{\mathfrak{p}}^\nu$ and $\underline{\mathfrak{p}}^{\nu+1}$, there is a $\underline{\mathfrak{p}}^{\nu,\nu+1}$-quilted disk u_ν with no inputs, having index zero, and area at most E_0, which contains a non-constant quilted surface component S_v. Indeed, the boundary of the one-dimensional part of the moduli space of quilted disks with no inputs consists of quilted disks with a broken edge e (that is, $\ell(e) = \infty$), and the only configurations where the map is constant on the quilted disk are those counted in $\mu^0_{CF(P^\nu)}$ and $\mu^0_{CF(P^{\nu+1})}$.

The quilted disk u_ν is of one of the forms shown in Figure 12.2, and we now rule out both possibilities by a dimension argument. A subsequence of the sequence u_ν converges to a J_0-holomorphic treed disk u_∞ of index -1. The limit u^∞ also has a broken edge, and let u^i_∞, $i = 0, 1, \cdots$, be the maps obtained by cutting the broken edges in u_∞. One of these maps, say u^j_∞, has index -1 and the map on the surface part is non-constant. The surface component in u^j_∞ is J_0-holomorphic, and so, it consists of

- disks that have positive isolated intersections with the toric divisors, and so, have Maslov index at least 2;
- and spheres that have $c_1 \geq 0$ by the semi-Fano condition.

Consequently, if u^j_∞ does not have inputs, the index $i(u^j_\infty) \geq 0$, which rules out configurations of the form (b) in Figure 12.2 for large enough ν. Suppose an input leaf of u^j_∞ is the output of a treed disk u^k_∞ as in Figure 12.2 (a). By the same reasoning used for u^j_∞ the surface part of u^k_∞ has Maslov index ≥ 2, and so, the output is necessarily the maximum point $x^\blacktriangledown$. Therefore, in Figure 12.2 (a), any input to the (-1)-index disk u^j_∞ is $x^\blacktriangledown$, and consequently, the disk u^j_∞ does not exist for degree reasons since the surface component has Maslov index at least 2. Therefore the configuration (a) is also ruled out for large ν. Thus the Claim is proved.

Finally, $m^0_{J_0}$ is a multiple of $x^\blacktriangledown$ as follows: For any sequence of regular perturbations $\underline{\mathfrak{p}}^\nu = (\underline{J}^\nu, \underline{F})$ that uniformly converge to $(J_0, \underline{F})$, and a sequence of treed

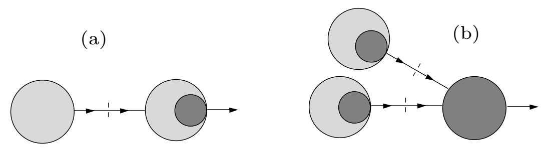

FIGURE 12.2. Codimension one strata in a moduli space of quilted disks with no inputs.

disks $\underline{\mathfrak{p}}_\nu$ disks u_ν with no inputs, of index zero, and uniformly bounded area, a subsequence converges to a $(J_0, \underline{F})$-disk u_∞ of index 0. By the semi-Fano condition the Maslov index of u_∞ is at least 2, and since there are no inputs, a dimension count dictates that $x^\blacktriangledown$ is the only possible output. □

PROPOSITION 12.8. *For a semi-Fano toric symplectic manifold, the potential $W(J_0)$ does not depend on the choice of divisor-preserving almost complex structure J_0.*

OUTLINE OF PROOF. Assuming the contrapositive, there is a sequence of quilted disks that converge to a J_t-holomorphic disk of index -1 for some $t \in [0,1]$, where J_t is a divisor-preserving almost complex structure. The existence of such a disk is ruled out as in the proof of Proposition 12.7. □

The following Proposition gives a sufficient condition under which the potential is not altered by multiple cutting of a semi-Fano toric surface.

PROPOSITION 12.9. (When breaking does not alter the potential) *Let X be a semi-Fano toric surface, $\mathcal{P}$ a toric multiple cut, and $L \subset X_{P_0}$ a toric Lagrangian in a component X_{P_0} of the broken manifold $\mathfrak{X}_{\mathcal{P}}$. Let $\mathfrak{J}_0$ be a broken divisor-preserving tamed almost complex structure on $\mathfrak{X}$. Suppose for any broken $\mathfrak{J}_0$-holomorphic disk u in $\mathfrak{X}$, $I(u_{\text{glue}}) \geq 2$. Then the potential for $(\mathfrak{X}, \mathfrak{J}_0, L)$ is well-defined and is the same as the toric potential for X equipped with a divisor preserving almost complex structure.*

PROOF. Let J_0^ν be the divisor preserving almost complex structure obtained from $\mathfrak{J}_0$ by gluing the cylindrical ends of $\mathfrak{X}$. Following the proof of Proposition 12.7, if the Proposition were not true, there would exist

- a sequence of regular perturbations $\underline{\mathfrak{p}}_\nu$ on X^ν that converge to the constant broken almost complex structure J_0,
- perturbation morphisms $\underline{\mathfrak{p}}^{\nu,\nu+1}$ that also converge to J_0,
- a sequence u_ν of $\underline{\mathfrak{p}}^{\nu,\nu+1}$-quilted disks of index zero, uniformly bounded area, each having a broken edge e (that is, $\ell(e) = \infty$), and a quilted surface component on which the map is non-constant,

and consequently, after passing to a subsequence, the sequence $\{u_\nu\}_\nu$ converges to a $\mathfrak{J}_0$-holomorphic broken treed disk u_∞ of index -1. Indeed, this is the only case to consider, since in the proof of Proposition 12.7 we have already ruled out quilted disks of the above form that converge to a J_0^ν-holomorphic disk for a finite ν. The quilted disks in the sequence ν are of one of the two forms shown in Figure 12.2. Copying the arguments from the proof of Proposition 12.7, each of the cases is ruled out using the fact that for any $\mathfrak{J}_0$-holomorphic disk u, the Maslov index $I(u_{\mathrm{glue}})$ is ≥ 2. □

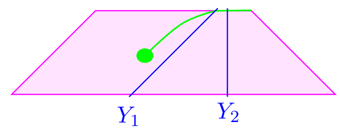

FIGURE 12.3. A broken disk with Maslov index 0 in the second Hirzebruch surface with the multiple cut $\mathcal{P}_2$ from Example 1.5.

REMARK 12.10. (A counter-example) In Example 1.5 we saw that there is a multiple cut $\mathcal{P}_2$ on the second Hirzebruch surface $X := \mathcal{H}_2$ for which the potential of the broken manifold $\mathfrak{X}_{\mathcal{P}_2}$ differs from that of the unbroken one X. The broken manifold $\mathfrak{X}_{\mathcal{P}_2}$ does not satisfy the hypothesis of Proposition 12.9 because there is a broken disk u (shown in Figure 12.3) whose glued Maslov index $I(u_{\mathrm{glue}})$ is 0. The map u contains a disk of Maslov index 4 and two spheres with self-intersection -1.

In the following result we verify that the multiple cut $\mathcal{P}$ on the cubic surface given in Figure 12.4 satisfies the hypothesis of Proposition 12.9, and use it to conclude that the multiple cut preserves the potential of the toric Lagrangian in the cubic surface.

PROPOSITION 12.11. *In the multiply cut cubic surface $\mathfrak{X}$ as in Figure 12.4 equipped with a divisor-preserving domain-dependent almost complex structure $\mathfrak{J}_0$ and a toric Lagrangian $L \subset X_{P_0}$, for any $\mathfrak{J}_0$-holomorphic broken disk u, the glued Maslov index $I(u_{\mathrm{glue}})$ is at least 2. Consequently the multiple cut of Figure 12.4 preserves the potential of the toric Lagrangian L.*

PROOF. We first consider disks whose tropical nodes do not map to orbifold singularities, and prove that their glued Maslov index is at least 2. We recall from (6.22) that for a rigid broken map $u : C \to \mathfrak{X}$ of type Γ, the Maslov index of the glued type Γ_{glue} is

$$I(\Gamma_{\mathrm{glue}}) = \sum_{v \in \mathrm{Vert}(\Gamma)} \overline{I}(u_v), \quad \overline{I}(u_v) := I(u_v) - \textstyle\sum_{e \in \mathrm{Edge}_{\mathrm{trop}}(\Gamma): e \ni v} 2\mu_{e,v}, \tag{12.5}$$

where $\mu_{e,v}$ is the sum of the intersection multiplicities with all the relative divisors at the lift of the node w_e on C_v, and $I(u_v)$ is the Maslov index of the disk/sphere u_v. Via a case-by-case analysis of the components of the map, we will show that $I(\Gamma_{\mathrm{glue}}) \geq 2$. In the sequel, we say that $v \in \mathrm{Vert}(\Gamma)$ is a *descendent* of $w \in \mathrm{Vert}(\Gamma)$ if there is an edge between v and w, and v is further from the root than w.

CASE 1: If the map u_v does not map to a toric divisor of $X_{\overline{P}(v)}$ then $\overline{I}(u_v) \geq 0$. Indeed, in this case $\overline{I}(u_v)$ is equal to the number of intersections of the map u_v with non-boundary toric divisors counted with multiplicity, which is a non-negative

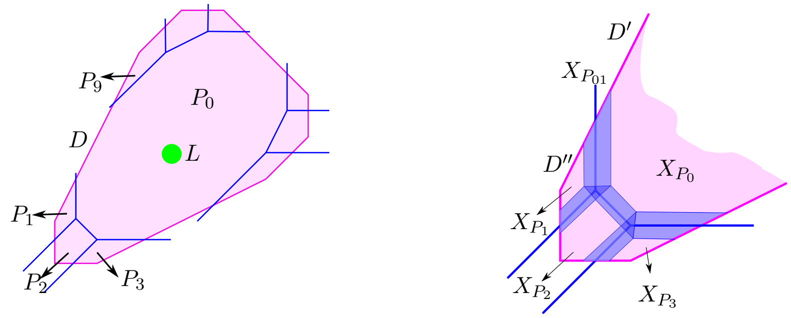

FIGURE 12.4. Left: Multiple cut $\mathcal{P}$ on the cubic surface X. Right: Components of the broken cubic surface $\mathfrak{X}$.

integer. An inspection of the multiply cut cubic surface $\mathfrak{X}$ tells us that this case covers all maps u_v where $\mathrm{codim}(P(v)) > 0$.

CASE 2: Suppose $P(v)$ is top-dimensional, $P(v) \neq P_0$, and u_v maps to a non-boundary toric divisor $Y \subset X_{P(v)}$. Then Y is a (-1)-curve and one of the ends of Y intersects a non-boundary toric divisor $Y_1 \subset X_{P(v)}$. Therefore, $\overline{I}(u_v) = 0$.

CASE 3: Suppose $P(v) = P_0$ and u_v maps to a long divisor, say D'_1. In this case $\overline{I}(u_v)$ is negative, but we show that its contribution is cancelled by positive contributions from some descendent vertices of v. For the moment, we assume u_v is a simple cover of D'_1. Then v has a descendent v_1 such that

Case A: $P(v_1) = P_1$ or P_{01} or
Case B: $P(v_1) = P_4$ or P_{04}.

Since (A) and (B) are symmetric, we only consider (A). If $P(v_1) = P_1$, u_{v_1} intersects either D''_1 or E'_1, both of which are non-relative toric divisors of X_{P_1}; and if $P(v_1) = P_{01}$, u_{v_1} has an intersection with the thickening of $D'_1 \cap D''_1$, which is a non-boundary toric divisor of $X_{\overline{P}_{01}}$. In both cases $I(u_{v_1}) = 2$, which cancels the negative contribution $I(u_v) = -2$. In case u_v is a k-cover of D'_1, then a similar cancellation argument applies, using descendent vertices of v.

So far, we have shown that the sum of $\overline{I}(u_v)$ over all vertices is non-negative. We now show that the sum is positive. Consider a disk component $v \in \mathrm{Vert}_\circ(\Gamma)$ on which the map is non-constant. The disk u_v either intersects a long divisor, in which case $\overline{I}(u_v) \geq 2$; otherwise it intersects the relative divisor, say $X_{P_{01}}$. In the latter case, there is a descendent vertex v_1 of v for which $\overline{I}(u_{v_1}) \geq 2$ using the same argument as in Case 3. This finishes the proof in the case when the tropical nodes of the broken disk u do not map to orbifold singularities.

In the case when a tropical node on a component u_v maps to an orbifold singularity in $X_{\overline{P}(v)}$, we consider a toric resolution $\tilde{X}_v$ of $X_{\overline{P}(v)}$ and a lift $\tilde{u}_v : C_v \to \tilde{X}_v$. The formula (12.5) is applicable if the Maslov index $I(u_v)$ and the sum of intersection multiplicities $\mu_{e,v}$ are replaced by the corresponding quantities for the lift $\tilde{u}_v$. The other arguments in the proof now carry over. In fact, the components u_v

mapping to orbifold points are covered by Case 1, since in the multiply cut cubic surface, toric divisors in $X^{\square}_{P(v)}$ do not contain orbifold singularities in their closure.

We have shown that in the multiply cut cubic surface $\mathfrak{X}$, the glued Maslov index of any broken disk is at least 2. Therefore, we may apply Proposition 12.9 on $\mathfrak{X}$ and conclude that the potential of the toric Lagrangian L in the cubic surface is preserved under the multiple cut $\mathcal{P}$. This finishes the proof of Proposition 12.11. □

REMARK 12.12. A curious reader may wonder whether the potential has the same terms if the Lagrangian is in some other top-dimensional piece of the multiply cut cubic surface. The potential has the same terms as the unbroken case when L is a toric Lagrangian in any top-dimensional piece X_{P_i}, $i = 0, \dots, 9$ in the broken manifold. Indeed, the proof of Proposition 12.11 can be replicated, and there are no disks of Maslov index < 2 in each of the cases. However, if $L \subset X_{P_i}$, $i \neq 0$, the disks can not be counted in a straightforward way, since some of the broken disks contain components that are multiple covers of (-1)-spheres. For example, for the disk u in Figure 12.5, the gluing u_{glue} is homologous to the Blaschke disk of Maslov index two that intersects the short divisor E_5. A calculation using obstruction bundles shows that this disk makes a negative contribution to the disk count. There is another broken disk u' in the same homology class, where the sphere u_{v_3} is replaced by two spheres u_{v_3}, $u_{v_3'}$, each of homology class $[D_1']$ (see Figure 12.6), which makes a positive contribution. We do not perform the obstruction bundle computations here.

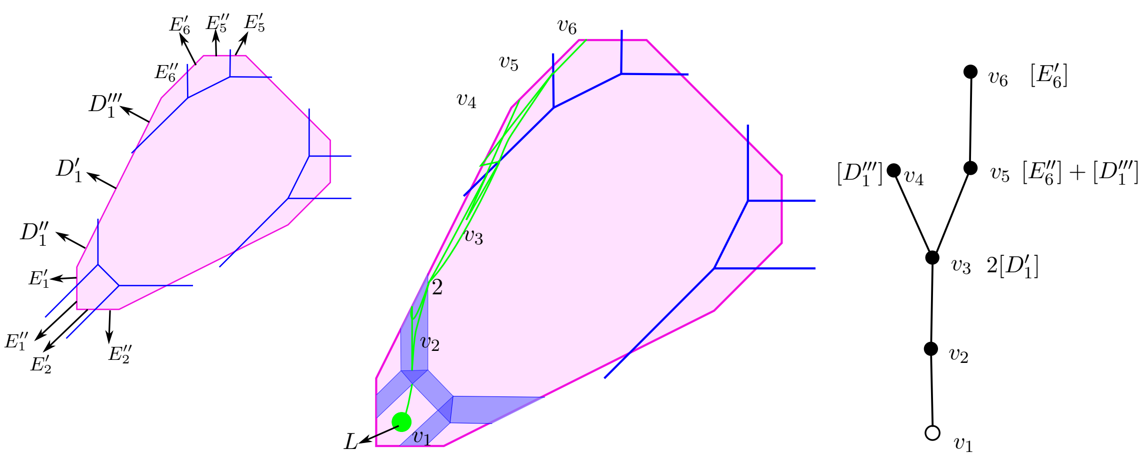

FIGURE 12.5. A broken map u whose gluing is a disk of Maslov index two and intersects the short divisor E_5, and which makes a negative contribution to the disk count. The node (v_1, v_2) (whose v_1-end maps to an orbifold singularity) is similar to the node in Example 6.26 (b).

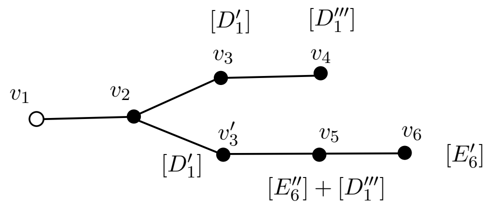

FIGURE 12.6. The graph of a broken map u' whose gluing is homologous to u, and which makes a positive contribution to the disk count.

Bibliography

[1] M. Abouzaid. Framed bordism and Lagrangian embeddings of exotic spheres. *Ann. of Math. (2)*, 175(1):71–185, 2012.

[2] D. Abramovich, Q. Chen, M. Gross, and B. Siebert. Decomposition of degenerate Gromov-Witten invariants. arXiv:1709.09864

[3] Byung Hee An, Yunhyung Cho, and Jang Soo Kim. On the f-vectors of Gelfand-Cetlin polytopes. *European J. Combin.*, 67:61–77, 2018.

[4] D. Auroux. A remark about Donaldson's construction of symplectic submanifolds. *J. Symplectic Geom.*, 1(3):647–658, 2002.

[5] D. Auroux, D. Gayet, and J.-P. Mohsen. Symplectic hypersurfaces in the complement of an isotropic submanifold. *Math. Ann.*, 321(4):739–754, 2001.

[6] S. Bardwell-Evans and M. W. Cheung, H. Hong, Y.-S. Lin. Scattering Diagrams from Holomorphic Discs in Log Calabi-Yau Surfaces. arXiv:2110.15234

[7] L. J. Barrott. Explicit equations for mirror families to log Calabi-Yau surfaces arxiv:1810.08356.

[8] V. V. Batyrev. Quantum cohomology rings of toric manifolds. *Astérisque*, (218):9–34, 1993. Journées de Géométrie Algébrique d'Orsay (Orsay, 1992).

[9] P. Biran and O. Cornea. Quantum structures for Lagrangian submanifolds. arxiv:0708.4221.

[10] J. M. Boardman and R. M. Vogt. *Homotopy invariant algebraic structures on topological spaces*. Lecture Notes in Mathematics, Vol. 347. Springer-Verlag, Berlin, 1973.

[11] F. Bourgeois, Y. Eliashberg, H. Hofer, K. Wysocki, and E. Zehnder. Compactness results in symplectic field theory. *Geom. Topol.*, 7:799–888 (electronic), 2003.

[12] F. Bourgeois. A Morse-Bott approach to contact homology. In *Symplectic and contact topology: interactions and perspectives (Toronto, ON/Montreal, QC, 2001)*, volume 35 of *Fields Inst. Commun.*, pages 55–77. Amer. Math. Soc., Providence, RI, 2003.

[13] J. Bryan and N. C. Leung. The enumerative geometry of K3 surfaces and modular forms. *J. Amer. Math. Soc.* 13:371–410, 2000.

[14] Marco Castronovo. Fukaya category of Grassmannians: rectangles. *Adv. Math.*, 372:107287, 40, 2020.

[15] K. Chan, C.-H. Cho, S.-C. Lau, H.H. Tseng. Gross fibrations, SYZ mirror symmetry, and open Gromov-Witten invariants for toric Calabi-Yau orbifolds. *J. Differential Geom.* 103:207–288, 2016.

[16] Kwokwai Chan and Siu-Cheong Lau. Open Gromov-Witten Invariants and Superpotentials for Semi-Fano Toric Surfaces. *International Mathematics Research Notices*, 2014(14):3759–3789, 03 2013.

[17] F. Charest and C. Woodward. Floer trajectories and stabilizing divisors. *Jour. Fixed Point Theory.* 19:1165–1236, 2017. arXiv:1401.0150.

[18] F. Charest and C. Woodward. Floer trajectories and flips. To appear in *Memoirs of the American Mathematical Society* arxiv:1508.01573.

[19] Cheol-Hyun Cho. Holomorphic discs, spin structures, and Floer cohomology of the Clifford torus. *Int. Math. Res. Not.*, (35):1803–1843, 2004.

[20] C.-H. Cho and Y.-G. Oh. Floer cohomology and disc instantons of Lagrangian torus fibers in Fano toric manifolds. *Asian J. Math.*, 10(4):773–814, 2006.

[21] Yunhyung Cho and Yoosik Kim. Monotone Lagrangians in Flag Varieties. arxiv:1801.07554.

[22] K. Cieliebak and K. Mohnke. Symplectic hypersurfaces and transversality in Gromov-Witten theory. *J. Symplectic Geom.*, 5(3):281–356, 2007.

[23] K. Cieliebak and K. Mohnke. Compactness for punctured holomorphic curves. *J. Symplectic Geom.*, 3(4):589–654, 2005. Conference on Symplectic Topology.

[24] O. Cornea and F. Lalonde. Cluster homology: An overview of the construction and results. *Electron. Res. Announc. Amer. Math. Soc.*, (12):1-12, 2006.

[25] T. Delzant. Hamiltoniens périodiques et images convexes de l'application moment. *Bull. Soc. Math. France*, 116:315–339, 1988.

[26] S. K. Donaldson. Symplectic submanifolds and almost-complex geometry. *J. Differential Geom.*, 44(4):666–705, 1996.

[27] J. J. Duistermaat. On global action-angle coordinates. *Comm. Pure Appl. Math.*, 33(6):687–706, 1980.

[28] Jonathan David Evans. Lectures on lagrangian torus fibrations. *ArXiv e-prints*, 2021.

[29] Mohammad F. Tehrani. Pseudoholomorphic Curves Relative to a Normal Crossings Symplectic Divisor: Compactification. arXiv:1710.00224

[30] Mohammad F. Tehrani. Towards a Degeneration Formula for the Gromov-Witten Invariants of Symplectic Manifolds. arXiv:1710.00599

[31] Mohammad F. Tehrani. Deformation Theory of Log Pseudo-holomorphic Curves and Logarithmic Ruan-Tian Perturbations. arXiv:1910.05201

[32] A. Floer. The unregularized gradient flow of the symplectic action. *Comm. Pure Appl. Math.* 41, no. 6, 775813 (1988).

[33] W. Fulton, B. Sturmfels. Intersection theory on toric varieties. *Topology* 36:335-353, 1997.

[34] K. Fukaya, Y.-G. Oh, H. Ohta, and K. Ono. *Lagrangian intersection Floer theory: anomaly and obstruction.*, volume 46 of *AMS/IP Studies in Advanced Mathematics*. American Mathematical Society, Providence, RI, 2009.

[35] K. Fukaya, Y.-G. Oh, H. Ohta, and K. Ono. Lagrangian Floer theory on compact toric manifolds. I. *Duke Math. J.*, 151(1):23–174, 2010.

[36] S. Galkin and A. Usnich. Mutations of potentials. Preprint IPMU, 2010. http://member.ipmu.jp/sergey.galkin/papers/ipmu-10-0100.pdf

[37] Fabio Gironella, Vicente Muñoz, Zhengyi Zhou. Asymptotically holomorphic theory for symplectic orbifolds. arXiv:1910.05201

[38] A. Givental. Homological geometry and mirror symmetry. In *Proceedings of the International Congress of Mathematicians, Vol. 1, 2 (Zürich, 1994)*, pages 472–480, Basel, 1995. Birkhäuser.

[39] A. Givental. Stationary phase integrals, quantum Toda lattices, flag manifolds and the mirror conjecture. In *Topics in singularity theory* American Mathematical Society Translations Ser. 2, AMS, 1997.

[40] R. Gompf. A new construction of symplectic manifolds. *Annals of Math.* 142:527–595, 1995.

[41] L. Göttsche and R. Pandharipande. The quantum cohomology of blow-ups of P2 and enumerative geometry. *J. Differential Geom.* 48:61–90, 1998.

[42] M. Gross. Toric degenerations and Batyrev–Borisov duality. Math. Ann. 333:645–688, 2005.

[43] M. Gross. Mirror symmetry and the Strominger-Yau-Zaslow conjecture. Current developments in mathematics 2012, 133191, Int. Press, Somerville, MA, 2013.

[44] Mark Gross. *Tropical geometry and mirror symmetry*, volume 114 of *CBMS Regional Conference Series in Mathematics*. Published for the Conference Board of the Mathematical Sciences, Washington, DC; by the American Mathematical Society, Providence, RI, 2011.

[45] Mark Gross and Bernd Siebert. From real affine geometry to complex geometry. *Ann. of Math. (2)*, 174(3):1301–1428, 2011.

[46] Mark Gross and Bernd Siebert. An invitation to toric degenerations. In *Surveys in differential geometry. Volume XVI. Geometry of special holonomy and related topics*, volume 16 of *Surv. Differ. Geom.*, pages 43–78. Int. Press, Somerville, MA, 2011.

[47] V Guillemin and S Sternberg. The gelfand-cetlin system and quantization of the complex flag manifolds. *Journal of Functional Analysis*, 52(1):106–128, 1983.

[48] M. Harada and K. Kaveh. Integrable systems, toric degenerations and Okounkov bodies. *Invent. Math.*, 202:927–985, 2015.

[49] Helmut Hofer. Pseudoholomorphic curves in symplectizations with applications to the weinstein conjecture in dimension three. *Invent Math*, (114):515–563, 1993.

[50] E.-N. Ionel. GW invariants relative to normal crossing divisors. Adv. Math. 281 (2015), 40–141.

[51] E.-N. Ionel and T. H. Parker. Relative Gromov-Witten invariants. *Ann. of Math. (2)*, 157(1):45–96, 2003.

[52] Maxim Kontsevich and Yan Soibelman. Homological mirror symmetry and torus fibrations. In *Symplectic geometry and mirror symmetry (Seoul, 2000)*, pages 203–263. World Sci. Publ., River Edge, NJ, 2001.

[53] M. Kontsevich and Y. Soibelman. Notes on A_∞-algebras, A_∞-categories and non-commutative geometry. In *Homological mirror symmetry*, volume 757 of *Lecture Notes in Phys.*, pages 153–219. Springer, Berlin, 2009.

[54] E. Lerman. Symplectic cuts. *Math. Res. Letters*, 2:247–258, 1995.

[55] J. Li. A Degeneration formula of GW-invariants *J. Differential Geom.* 60–199–293, 2002.

[56] R. B. Lockhart and R. C. McOwen. Elliptic differential operators on noncompact manifolds. *Ann. Scuola Norm. Sup. Pisa Cl. Sci. (4)*, 12(3):409–447, 1985.

[57] E. Meinrenken. Symplectic Surgery and the Spin^c Dirac Operator. *Advances in Mathematics* 134 240–277, 1998.

[58] S. Ma'u and C. Woodward. Geometric realizations of the multiplihedra. *Compos. Math.*, 146(4):1002–1028, 2010.

[59] Dusa McDuff and Dietmar Salamon, *Introduction to symplectic topology*, second ed., Oxford Mathematical Monographs, The Clarendon Press, Oxford University Press, New York, 1998. MR 1698616

[60] D. McDuff and D. Salamon. *J-holomorphic curves and symplectic topology*, volume 52 of *American Mathematical Society Colloquium Publications*. American Mathematical Society, Providence, RI, 2004.

[61] Grigory Mikhalkin. Enumerative tropical algebraic geometry in $\mathbb{R}^2$. *J. Amer. Math. Soc.*, 18(2):313–377, 2005.

[62] D. Mumford. Algebraic geometry. I. Complex projective varieties. Grundlehren der Mathematischen Wissenschaften, No. 221. Springer–Verlag, Berlin–New York, 1976.

[63] T. Nishinou, Y. Nohara, and K. Ueda. Toric degenerations of Gelfand-Cetlin systems and potential functions. *Adv. Math.*, 224(2):648–706, 2010.

[64] Y. Nohara and K. Ueda. Toric degenerations of integrable systems on Grassmannians and polygon spaces. *Nagoya Math. J.* 214: 125–168, 2014.

[65] Y.-G. Oh. Floer cohomology, spectral sequences, and the Maslov class of Lagrangian embeddings. *Internat. Math. Res. Notices*, (7):305–346, 1996.

[66] Y.-G. Oh. Riemann-Hilbert problem and application to the perturbation theory of analytic discs. *Kyungpook Math. J.*, 35(1):3975, 1995.

[67] J. Palmer and C. Woodward. Immersed Floer cohomology and Maslov flow. Preprint, March 2018.

[68] Brett Parker. Gluing formula for Gromov-Witten invariants in a triple product. arXiv:1511.0779.

[69] Brett Parker. Notes on exploded manifolds and a tropical gluing formula for Gromov-Witten invariants. *Gromov-Witten theory, gauge theory and dualities*, volume 48 of *Proc. Centre Math. Appl. Austral. Nat. Univ.*, page 17. Austral. Nat. Univ., Canberra, 2019.

[70] Brett Parker. Tropical enumeration of curves in blowups of the projective plane. arXiv:1411.5722.

[71] Brett Parker. De Rham theory of exploded manifolds. *Geom. Topol.*, 22(1):1–54, 2018.

[72] Brett Parker. Holomorphic curves in exploded manifolds: regularity. *Geom. Topol.*, 23(4):1621–1690, 2019.

[73] Brett Parker. Exploded manifolds. *Adv. Math.*, 229(6):3256–3319, 2012.

[74] Brett Parker. Log geometry and exploded manifolds. *Abh. Math. Sem. Hamburg*, 82:43–81, 2012. arxiv:1108.3713.

[75] Brett Parker. Holomorphic curves in exploded manifolds: Kuranishi structure. arXiv:1301.4748, 2013.

[76] Brett Parker. Holomorphic curves in exploded manifolds: compactness. *Adv. Math.*, 283:377–457, 2015. arXiv:0911.2241.

[77] J. Pascaleff and D. Tonkonog. The wall-crossing formula and Lagrangian mutations arxiv:1711.03209.

[78] Matthias Schwarz. *Morse homology*, volume 111 of *Progress in Mathematics*. Birkhäuser Verlag, Basel, 1993.

[79] P. Seidel. Graded Lagrangian submanifolds. *Bull. Soc. Math. France*, 128(1):103–149, 2000.

[80] P. Seidel. Homological mirror symmetry for the genus two curve. *J. Algebraic Geom.* 20:727–769, 2011.

[81] P. Seidel. A_∞-subalgebras and natural transformations. *Homology, Homotopy Appl.*, 10(2):83–114, 2008.

[82] P. Seidel. *Fukaya categories and Picard-Lefschetz theory*. Zurich Lectures in Advanced Mathematics. European Mathematical Society (EMS), Zürich, 2008.

[83] N. Sheridan. On the Fukaya category of a Fano hypersurface in projective space. —em Pub. math. de l'IHÉS, 124(1)165–317, 2016.

[84] B. Siebert. Gromov-Witten invariants of general symplectic manifolds. arXiv:dg-ga/9608005.

[85] J. Stasheff. *H-spaces from a homotopy point of view*. Lecture Notes in Mathematics, Vol. 161. Springer-Verlag, Berlin, 1970.

[86] Margaret Symington. Four dimensions from two in symplectic topology. In *Topology and geometry of manifolds (Athens, GA, 2001)*, volume 71 of *Proc. Sympos. Pure Math.*, pages 153–208. Amer. Math. Soc., Providence, RI, 2003.

[87] K. Wehrheim and C.T. Woodward. Orientations for pseudoholomorphic quilts. arXiv:1503.07803.

[88] Y. Wu. Splitting of Gromov-Witten Invariants with Toric Gluing Strata. arXiv:2103.14780.

[89] F. Zhang. Matrix theory: basic results and techniques. Universitext. Springer, New York, NY (2011).

[90] A. Zinger. Notes on J-holomorphic curves. arXiv:1706.00331.

Index

Acknowledgements : We thank Mohammad F. Tehrani for discussing with us and pointing out an issue with framings in an earlier version. We thank Nick Sheridan for sharing his picture of disks on the cubic surface.