Fukaya's work on Lagrangian embeddings

Janko Latschev

Fachbereich Mathematik, Universität Hamburg, Bundestrasse 55, 20146 Hamburg, Germany

Abstract. In this chapter I discuss some applications of string topology to the study of Lagrangian embeddings into symplectic manifolds, as discovered by Fukaya [10].
June 2, 2014

2000 Mathematics Subject Classification:

Key words:

1 Introduction

A submanifold $L \subset (X, \omega)$ in some symplectic manifold (X, ω) is called Lagrangian if $\dim L = \frac{1}{2} \dim X$ and $\omega|_L = 0$. A simple example is given by the zero section $L \subset T^*L$ in the cotangent bundle of a smooth manifold L, and this is universal in the sense that a neighborhood of any Lagrangian embedding of a closed L into some symplectic manifold is symplectomorphic to a neighborhood of $L \subset T^*L$. Lagrangian submanifolds play a fundamental role in symplectic geometry and topology, as many constructions and objects can be recast in this form. In fact, already in a 1980 lecture (cf. [28]), A. Weinstein formulated the "symplectic creed":

EVERYTHING IS A LAGRANGIAN SUBMANIFOLD.

Today, Lagrangian submanifolds (sometimes decorated with additional structures) are for example studied as objects of the *Fukaya category*, which plays a fundamental role in Kontsevich's formulation of homological mirror symmetry. Rather than delving into such general theories, I want to concentrate here on a quite simple, and in fact basic, question:

Which closed, oriented n-manifolds admit a Lagrangian embedding into the standard symplectic space $(\mathbb{C}^n, \omega_0)$, with $\omega_0 = \sum_j dx_j \wedge dy_j$?

An excellent introduction to this question, containing a discussion of some of the relevant classical algebraic topology, as well as early results obtained by holomorphic curve methods, is [4], which I will quote freely.

For $n = 1$ there is not much to say, since S^1 is the only connected closed 1-manifold, and the Lagrangian condition $\omega_0|_L = 0$ is trivial in this case. In general, a necessary condition for an oriented closed manifold L^n to admit a Lagrangian embedding into $\mathbb{C}^n$ is that its Euler characteristic $\chi(L)$ should vanish. This is because the self-intersection number of any submanifold of $\mathbb{C}^n$ is clearly zero, but it is also equal to the Euler characteristic of the normal bundle, which for Lagrangian submanifolds is isomorphic to the cotangent bundle.

So for $n = 2$, the only orientable closed manifold that could have a Lagrangian embedding into $\mathbb{C}^2$ is $T^2 = S^1 \times S^1$, and it embeds e.g. as the product of one circle in each $\mathbb{C}$-factor. For non-orientable closed surfaces Σ, classical algebraic topology implies that a necessary condition for the existence of a Lagrangian embedding is that $\chi(\Sigma)$ is divisible by 4, and a beautiful construction by Givental [13] shows that for strictly negative Euler characteristic this is also sufficient. The embedding question was only recently completely answered, when Shevchishin showed that the Klein bottle does not have a Lagrangian embedding into $\mathbb{C}^2$ ([24], see also [23] for an alternative argument by Nemirovski).

Already for $n = 3$, elementary algebraic topology does not tell us much. It was one of the many important results in Gromov's landmark paper [14] to show that there are no exact Lagrangian embeddings into $(\mathbb{C}^n, \omega_0)$, in the sense that any global primitive λ of the symplectic form ω_0 has to restrict to a non-exact closed 1-form on the Lagrangian submanifold $L \subset \mathbb{C}^n$. This in particular rules out S^3, but of course there are plenty of closed orientable 3-manifolds with $H^1(M, \mathbb{R}) \neq 0$.

All of this and more is discussed in [4]. The goal of this chapter is to show how knowledge about string topology can be applied to give a far-reaching refinement of Gromov's result. In particular, I aim to present the overall strategy for proving the following result:

Theorem 1.1. *(Fukaya) Let L be a compact, orientable, aspherical spin manifold of dimension n which admits an embedding as a Lagrangian submanifold of $\mathbb{C}^n$. Then a finite covering space $\tilde{L}$ of L is homotopy equivalent to a product $S^1 \times L'$ for some closed $n-1$-manifold L'.*

Moreover, $\pi_1(\tilde{L}) \cong \pi_1(S^1 \times L') \subset \pi_1(L)$ is the centralizer of some element $\gamma \in \pi_1(L)$ which has Maslov class equal to 2 and positive symplectic area.

The assertion about the Maslov class is known as *Audin's conjecture*, and was originally asked for tori in $\mathbb{C}^n$, see [2]. The spin condition is a technical assumption (it is needed to make the relevant moduli spaces of holomorphic disks orientable), and I expect that it can be removed by reformulating the argument somewhat. The asphericity assumption (meaning that all higher homotopy groups of L vanish) enters the proof in a fairly transparent way, and one can imagine various replacements.

As a corollary, we obtain the following more precise statement in dimension 3.

Corollary 1.2. *(Fukaya) If the closed, orientable, prime 3-manifold L admits a Lagrangian embedding into $\mathbb{C}^3$, then L is diffeomorphic to a product $S^1 \times \Sigma$ of the circle with a closed, orientable surface.*

The fact that the product $S^1 \times \Sigma$ does embed as a Lagrangian submanifold into $\mathbb{C}^3$ follows from an elementary construction, see e.g. [4]. Basically, one starts from an isotropic embedding of Σ into $\mathbb{C}^3$, e.g. by embedding it into the Lagrangian subspace $\mathbb{R}^3 \subset \mathbb{C}^3$. Then one uses the fact that a small neighborhood necessarily is symplectomorphic to a neighborhood of the zero section in $T^*\Sigma \oplus \underline{\mathbb{C}}$, the direct sum of the cotangent bundle with a trivial symplectic vector bundle of rank 2, to embed the product $S^1 \times \Sigma$ by taking the product of the zero section in $T^*\Sigma$ with a standard small $S^1 \subset \mathbb{C}$.

The above statements are special cases of a more general result discovered by Kenji Fukaya, and first described in [10], see also [11]. As with most results involving J-holomorphic curves, the underlying idea can be traced back to Misha Gromov's foundational paper [14]. His proof of the fact that there are no exact compact Lagrangian submanifolds of $\mathbb{C}^n$ contains an important seed for Fukaya's arguments. Therefore, after discussing some aspects of moduli spaces of holomorphic disks in the next section, I begin section 3 by sketching Gromov's argument, followed by the discussion of an instructive example and the statement of Fukaya's refinement (Theorem 3.3). In a nutshell, Fukaya's important observation was that the compactification of the moduli spaces of holomorphic disks with boundary on a Lagrangian submanifold $L \subset \mathbb{C}^n$ can be expressed in terms of string topology operations, in particular the loop bracket (and possibly also its higher analogues at the chain level).

After an interlude section on basic properties of L_∞ algebras, I describe how Theorem 1.1 is a fairly straightforward consequence of Theorem 3.3. Corollary 1.2 will follow by using specific facts from 3-dimensional topology. Before finishing with a guide to the literature, I discuss a few further small observations.

At the moment of this writing, full proofs of the above theorems are not available yet. One would like to construct a chain model $C_*(\Lambda L)$ for the free loop space of the manifold L with the loop bracket operation λ_2 (and higher operations as necessary), such that $(C_*(\Lambda L), \lambda_1 = \partial, \lambda_2, \dots)$ is an L_∞-algebra. Moreover when this model is considered with coefficients in a suitable Novikov ring, the compactified moduli spaces of holomorphic disks should give rise to an element in this chain complex satisfying the Maurer-Cartan equation in this L_∞-algebra. So the principal problem is that the apparent freedom one has in building the chain model is severely constrained by the need to make it fit with the analysis of holomorphic disks, which in particular involves delicate transversality and gluing issues.

In this text, I will largely ignore the technical difficulties, by stating the key results as black boxes.

2 Moduli spaces of holomorphic disks

Basic notions

The main tool in our study of Lagrangian embeddings $L \subset \mathbb{C}^n$ will be moduli spaces of holomorphic disks with boundary on L. Before introducing the relevant spaces, I will briefly review some standard notions.

The first of these is the *Maslov index* of a loop of Lagrangian subspaces in $\mathbb{C}^n$.

The Lagrangian Grassmannian GLag_n is defined as the space of all Lagrangian subspaces of $(\mathbb{C}^n, \omega_0)$. Standard symplectic linear algebra shows that a real n-dimensional subspace $V \subset \mathbb{C}^n$ is Lagrangian if and only if it is orthogonal with respect to the standard Euclidean inner product to iV, its rotation by $i \in \mathbb{C}$. Moreover, any orthonormal basis of a Lagrangian subspace is a unitary basis of $\mathbb{C}^n$, and conversely the real linear span of a unitary frame is a Lagrangian subspace. Since $O(n)$ acts transitively on the unitary frames generating the same Lagrangian subspace, one concludes that GLag_n can be identified with $U(n)/O(n)$.

The map $\det^2 : U(n) \to S^1$ which associates to a unitary matrix the square of its determinant induces a well-defined map $\rho : \mathrm{GLag}_n \to S^1$, and for any loop $\gamma : S^1 \to \mathrm{GLag}_n$ we define *its Maslov index* as

$$\mu(\gamma) := \deg(\rho \circ \gamma).$$

Clearly $\mu(\gamma)$ so defined is an invariant of the free homotopy class of γ. Moreover, one can check that μ induces an isomorphism $\pi_1(\mathrm{GLag}_n) \cong \mathbb{Z}$.

Now given a Lagrangian immersion $L \hookrightarrow \mathbb{C}^n$, any loop $\gamma : S^1 \to L$ determines a loop in GLag_n, simply by taking the loop of tangent spaces of L at the image points of γ. In this way, we get a Maslov index

$$\mu : \pi_1(L) \to \mathbb{Z}.$$

Discussions of the Maslov index can be found in various texts, see e.g. [1, p. 116ff], [3, p. 33ff] or [21, p. 50ff]. In particular, it is easy to see that for a Lagrangian immersion of an oriented manifold L, the Maslov index takes values in $2\mathbb{Z}$.

Another useful notion is the *area* or *energy* of a disk $u : (D, \partial D) \to (\mathbb{C}^n, L)$ with boundary on a Lagrangian submanifold $L \subset (\mathbb{C}^n, \omega_0)$. It is defined as

$$E(u) := \int_D u^*\omega,$$

and one easily checks that it only depends on the free relative homotopy class of u. Indeed, just observe that given a homotopy $h : ([0,1] \times D, [0,1] \times \partial D) \to (\mathbb{C}^n, L)$ with $h_0 = u$ and $h_1 = u'$, we have

$$\int_D u'^*\omega - \int_D u^*\omega = \int_{[0,1]\times\partial D} h^*\omega = 0$$

by the assumption that L is a Lagrangian submanifold. In particular, E descends to a homomorphism $E : \pi_2(\mathbb{C}^n, L) \to \mathbb{R}$.

Finally, in order to discuss holomorphic curves, we introduce the relevant spaces of almost complex structures. Generally, an *almost complex structure* J on some manifold M is an automorphism of the tangent bundle $J : TM \to TM$ with $J^2 = -\mathrm{id}$. If M carries a symplectic form ω, then an almost complex structure J on M is said to be *tamed by* ω, if $\omega(v, Jv) > 0$ for all nonzero tangent vectors v. In other words, ω is a positive area form on each 1-dimensional J-complex subspace of each tangent space. It is a standard result that these ω-tamed almost complex structures form a contractible (in particular non-empty) space. Given such a tamed J, one can define a Riemannian metric on M by setting

$$g_J(v,w) := \frac{1}{2}\left(\omega(v,Jw) + \omega(w,Jv)\right).$$

One can also define the notion of a J-holomorphic map $u : (\Sigma, j) \to (M, J)$ from some Riemann surface (Σ, j) to M, simply by asking that it satisfy the usual Cauchy-Riemann equations with respect to the given J:

$$\bar{\partial}_J u := \frac{1}{2}(du + J \circ du \circ j) = 0. \tag{2.1}$$

In local conformal coordinates $z = s + it$ on the Riemann surface this can be written equivalently as

$$\partial_s u + J\partial_t u = 0.$$

In Riemannian geometry, the L^2-energy of a map $u : \Sigma \to (M, g_J)$ is defined as

$$E(u) = \frac{1}{2}\int_\Sigma \|du\|^2 d\mu,$$

where μ is any volume form on Σ and $\|du\|$ denotes the operator norm of $du : T_x\Sigma \to T_{u(x)}M$ with respect to the metric $\mu(\,.\,, j\,.\,)$ on Σ and the metric g_J on M. The integrand turns out to be independent of the choice of μ, as any scaling factor also appears in the operator norm with opposite exponent.

Now the importance of the taming condition stems from the following crucial fact: Suppose $u : \Sigma \to M$ is J-holomorphic, and we have chosen local conformal coordinates (s,t) on Σ. Then at x we have

$$\|du\|^2 = \omega(\partial_s u, J\partial_s u) + \omega(\partial_t u, J\partial_t u) = 2\omega(\partial_s u, \partial_t u),$$

so that

$$\frac{1}{2}\|du\|^2 ds \wedge dt = u^*\omega.$$

In particular, for J-holomorphic disks $u : (D, \partial D) \to (\mathbb{C}^n, L)$, the energy as defined above, which was a purely topological quantity, is the same as the usual L^2-energy of the map u.

Moduli spaces

Choose an almost complex structure J tamed by the standard symplectic structure ω_0 on $\mathbb{C}^n$. Given a relative homotopy class $a \in \pi_2(\mathbb{C}^n, L)$, we consider the set

$$\widetilde{\mathcal{M}}(a,J) := \{u : (D,\partial D) \to (\mathbb{C}^n, L) : \bar{\partial}_J u = 0, [u] = a \in \pi_2(\mathbb{C}^n, L)\}.$$

The real 2-dimensional group $\mathrm{Aut}(D,1) \subset PSL(2,\mathbb{R})$ of biholomorphisms of the disk fixing $1 \in S^1$ acts on $\widetilde{\mathcal{M}}(a,J)$ by precomposition, and the quotient

$$\mathcal{M}(a,J) := \widetilde{\mathcal{M}}(a,J)/\,\mathrm{Aut}(D,1)$$

is called the moduli space of holomorphic disks in the class a.

Remark 2.1. The reader should take note that in the literature moduli spaces of holomorphic curves are almost universally denoted by $\mathcal{M}$, usually with some decoration, and the precise meaning of the symbol should very carefully be checked in each case.

The equation defining $\widetilde{\mathcal{M}}(a,J)$ is elliptic, and so the linearization is a Fredholm operator. The index theorem for holomorphic curves, as discussed for example in [22, Appendix C], implies that the index of this operator, and hence the expected dimension of the moduli space $\widetilde{\mathcal{M}}(a,J)$, is $n+\mu(a)$. The assignment $u \mapsto \bar{\partial}_J u$ can be viewed as a section of a suitable Banach space bundle. Under favourable circumstances this section can be arranged to be transverse to the zero section, in which case its zero set $\widetilde{\mathcal{M}}(A,J)$ is a manifold of the expected dimension. In general, this is a serious technical difficulty beyond the scope of the present discussion, and resolving it requires substantial work.

Note that, for $a \neq 0$, the action of $\mathrm{Aut}(D,1)$ on $\widetilde{\mathcal{M}}(a,J)$ is free. So, assuming that we can arrange transversality, we conclude that $\mathcal{M}(a,J)$ is a smooth manifold of dimension

$$\dim \mathcal{M}(a,J) = n-2+\mu(a). \tag{2.2}$$

As shown by examples in [12], one generally needs the spin condition on L to be able to orient the moduli spaces $\mathcal{M}(a,J)$.

For several reasons the individual spaces $\mathcal{M}(a,J)$ tend to not be very useful for proving anything interesting about L (with the notable exception of Theorem 3.1). The first is that these spaces strongly depend on J, as simple examples show.

Example 2.2. Consider $L = S^1 \times S^1 \subset \mathbb{C}^2$ and the relative homotopy class $a = (1,0) \in \pi_2(\mathbb{C}^2, L)$ which has degree 1 in the first factor and degree 0 in the second.

Given any real number $\alpha \geq 0$, consider the almost complex structure J_α on $\mathbb{C}^2$ given by the matrix

$$J_\alpha = \begin{pmatrix} J_0 & 0 \\ A & J_0 \end{pmatrix}, \quad \text{with } A = \begin{pmatrix} \alpha & 0 \\ 0 & -\alpha \end{pmatrix} \quad \text{and } J_0 = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}.$$

If α ranges in some interval $0 \leq \alpha \leq \alpha_0$, then all the J_α are tamed by the symplectic form $\omega = \alpha_0^2\omega_0 \oplus \omega_0$ on $\mathbb{C}^2 = \mathbb{C} \oplus \mathbb{C}$.

Now for the standard split complex structure J_0 on $\mathbb{C}^2$, the moduli space $\mathcal{M}(a,J_0)$ is clearly non-empty, since for any $z_1, z_2 \in S^1$ the map $u(z) = (z_1 z, z_2)$ defines an element in it.

In general, given any J_α-holomorphic map u in the class a, note that the projection u_1 onto the first $\mathbb{C}$-factor is holomorphic in the usual sense, and since it has degree 1 it will be a biholomorphism. So, precomposing with a suitable element of $\mathrm{Aut}(D,1)$, we may assume that this projection u_1 is just a rotation by $u_1(1)$. A short computation shows that in this case for u to be J_α-holomorphic, it is necessary and sufficient that

$$\partial_s u_2 + J_0 \partial_t u_2 = \alpha \partial_t u_1.$$

Now if $u_2 : (D, \partial D) \to (\mathbb{C}, S^1)$ is a solution of this equation, then

$$\begin{aligned}\alpha = |\alpha \partial_t u_1| &= \frac{1}{\pi}\left|\int_D (\partial_s u_2 + J_0 \partial_t u_2) ds dt\right| \\ &= \frac{1}{\pi}\left|\int_D d(u_2 dt - J_0 u_2 ds)\right| \\ &= \frac{1}{\pi}\left|\int_0^{2\pi} (\cos(\theta) + \sin(\theta) J_0) u_2(e^{i\theta}) d\theta\right| \\ &\leq \frac{1}{\pi}\int_0^{2\pi} |u_2(e^{i\theta})| d\theta = 2.\end{aligned}$$

So for $\alpha > 2$ the moduli space $\mathcal{M}(a, J_\alpha)$ is empty.

The phenomenon discussed in this example is a crucial ingredient in the proof below of Gromov's Theorem 3.1.

A second, related complication is the failure of compactness, a simple instance of which is described in Example 2.5 below. Both of these apparent problems can be overcome by considering the collection of all $\{\mathcal{M}(a, J)\}_{a \in \pi_2(\mathbb{C}, L)}$ at once. In the next subsection, I briefly discuss the compactness issue.

The compactness theorem for disks

Gromov's compactness theorem, in its modern formulation, asserts that every sequence of holomorphic curves of fixed topology and uniformly bounded energy, whose images lie within a compact subset of the target, has a subsequence converging in a suitable sense to a limiting "stable curve". A proof in text book form for holomorphic spheres was given by McDuff and Salamon [22], and an exposition of Gromov's proof for the higher genus case was written up by Hummel [17]. The case of curves with boundary is for example discussed by Liu as part of her Ph.D. thesis [19]. For holomorphic disks with Lagrangian boundary conditions the precise statement and proof were worked out in U. Frauenfelder's diploma thesis, published as [9].

Here, I will describe the statement under the simplifying assumption that the target symplectic manifold (X, ω) is exact, i.e. $\omega = d\lambda$, as in the case of my main example $X = \mathbb{C}^n$. This assumption in particular implies that there are no nonconstant holomorphic spheres in X, and so the possible limiting configurations are much more restricted than in the general case. Basically, all the stable curves arising as Gromov limits will be stable trees of disks, as explained presently.

Recall that a *tree* is a (finite) set T together with an edge relation $E \subset T \times T$ satisfying the following conditions:

(symmetric) If $\alpha E \beta$ then $\beta E \alpha$.

(antireflexive) If $\alpha E \beta$ then $\alpha \neq \beta$.

(connected) For all $\alpha, \beta \in T$ with $\alpha \neq \beta$ there exist $\gamma_0, \ldots, \gamma_m \in T$ with $\gamma_0 = \alpha$, $\gamma_m = \beta$ and such that $\gamma_i E \gamma_{i+1}$ for all $0 \leq i \leq m-1$.

(no cycles) If $\gamma_0, \ldots, \gamma_m \in T$ satisfy $\gamma_i E \gamma_{i+1}$ and $\gamma_i \neq \gamma_{i+2}$ for all i then $\gamma_0 \neq \gamma_m$.

One usually draws trees by drawing the corresponding 1-dimensional CW-complexes, which have one vertex for each $\alpha \in T$ and one (unoriented) 1-cell connecting α and β whenever $\alpha E \beta$ (and $\beta E \alpha$).

It follows from these axioms that the CW-complex corresponding to a tree in this way is connected and contractible. Moreover, deleting an edge connecting two vertices $\alpha, \beta \in T$ splits the CW-complex into two connected components, and we denote the subset of vertices in the component of β by $T_{\alpha\beta}$.

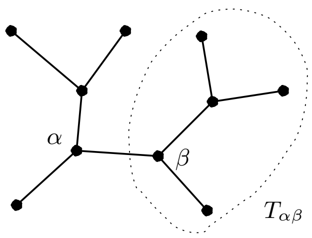

Figure 1. Two vertices α and β, and the corresponding subtree $T_{\alpha\beta}$.

A map $f : (T, E) \to (T', E')$ is called a *tree homomorphism* if for each $\alpha' \in T'$ the preimage $f^{-1}(\alpha')$ is a tree, and moreover $\alpha E \beta$ implies that either $f(\alpha) E' f(\beta)$ or $f(\alpha) = f(\beta)$. A bijective tree homomorphism is called *tree isomorphism.*

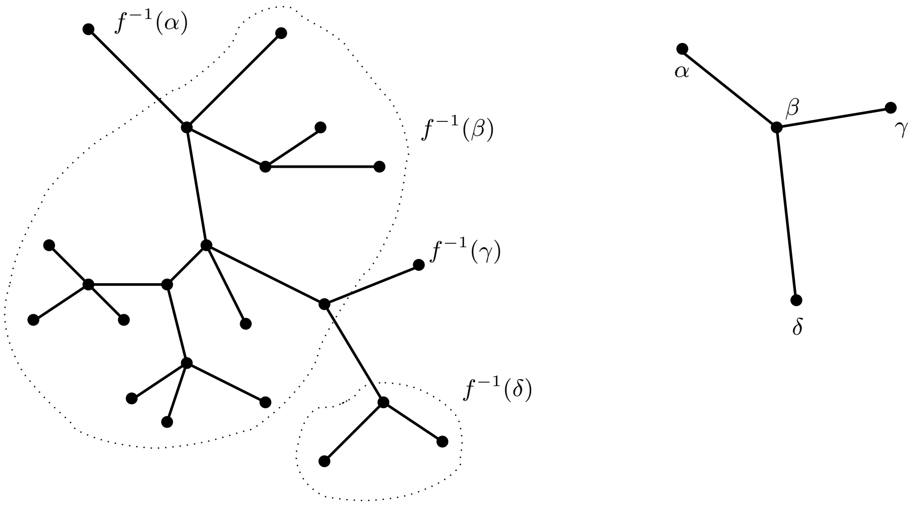

Figure 2. An example of a tree morphism.

Next we fix an exact symplectic manifold (X, ω) with a compact Lagrangian submanifold $L \subset X$ and an almost complex structure J tamed by ω. Given this data, we define a *stable tree of holomorphic disks with one boundary marked point* to be a tuple

$$(\mathbf{u},\mathbf{z}) = (\{u_\alpha\}_{\alpha\in T}, \{z_{\alpha\beta}\}_{\alpha E\beta}, \{\alpha_0, z_0\})$$

modelled over a tree (T,E), consisting of a collection of holomorphic maps $u_\alpha : (D,\partial D) \to (X,L)$ indexed by $\alpha \in T$, a collection of nodal points $z_{\alpha\beta} \in S^1$ indexed by directed edges $\alpha E\beta$, and a marked point $z_0 \in \partial D$ labelled by $\alpha_0 \in T$, subject to the following conditions:

(1) If $\alpha E\beta$ then $u_\alpha(z_{\alpha\beta}) = u_\beta(z_{\beta\alpha})$.

(2) For each $\alpha \in T$, the *special boundary points associated with the vertex* α, namely the points $z_{\alpha\beta}$ for different $\beta \in T$ with $\alpha E\beta$, together with z_0 for $\alpha = \alpha_0$, are all pairwise distinct.

(3) If u_α is constant, then the cardinality of the set $Y_\alpha = \{z_{\alpha\beta} : \alpha E\beta\} \cup \{z_0 : \alpha_0 = \alpha\}$ of all special points is at least 3.

As usual, two such stable trees of disks $(\mathbf{u},\mathbf{z})$ and $(\mathbf{u}',\mathbf{z}')$ modelled on trees T and T', respectively, are called *equivalent* if there exist a tree isomorphism $f : T \to T'$ and a collection of Möbius transformations $\{\phi_\alpha\}_{\alpha\in T}$ such that

$$z_{f(\alpha)f(\beta)} = \phi_\alpha(z_{\alpha\beta}), \quad \alpha_0' = f(\alpha_0), \quad z_0' = \phi_{\alpha_0}(z_0), \text{ and } \quad u'_{f(\alpha)}\phi_\alpha = u_\alpha$$

for all $\alpha, \beta \in T$ with $\alpha E\beta$.

Define the *energy* of a stable tree of holomorphic disks as the sum

$$E(\mathbf{u}) := \sum_{\alpha\in T} E(u_\alpha).$$

Now we can give the relevant notion of convergence.

Definition 2.3. Let (X,ω) be an exact symplectic manifold, let $L \subset X$ be a compact Lagrangian submanifold, and let J be an ω-tame almost complex structure on X. A sequence $u^\nu : (D,\partial D) \to (X,L)$ of holomorphic disks with one boundary marked point z_0^ν is said to *Gromov converge* to a stable tree of disks $(\mathbf{u},\mathbf{z})$ with one boundary marked point if there exist sequences of elements $\phi_\alpha^\nu \in PSL(2,\mathbb{R})$ indexed by $\alpha \in T$ such that the following statements hold:

(1) The sequence $\phi^\nu_{\alpha_0}(z_0^\nu)$ converges to z_0.

(2) For every $\alpha \in T$, the sequence of maps $u^\nu \circ \phi^\nu_\alpha$ converges to u_α uniformly on compact subsets of $D \setminus \{z_{\alpha\beta} : \alpha E\beta\}$.

(3) If $\alpha E\beta$ then $(\phi^\nu_\alpha)^{-1} \circ \phi^\nu_\beta$ converges to $z_{\alpha\beta}$ uniformly on compact subsets of $D \setminus \{z_{\beta\alpha}\}$.

(4) If $\alpha E\beta$ then
$$\sum_{\gamma\in T_{\alpha\beta}} E(u_\gamma) = \lim_{\varepsilon\to 0}\lim_{\nu\to\infty} E(u^\nu|_{\phi^\nu_\alpha(B_\varepsilon(z_{\alpha\beta}))}). \tag{2.3}$$

Intuitively, one can imagine the Möbius transformations ϕ^ν_α as microscopes, focusing on subregions in the domain to detect some piece of the limiting map. Item (3) in the definition then says that different microscopes really capture different phenomena. Item (4) is a way of phrasing that the limit captures all the essential pieces. In particular, together with (2) it implies that

$$E(\mathbf{u}) = \lim_{\nu\to\infty} E(u^\nu).$$

Now Gromov's compactness theorem for disks can be stated as follows.

Theorem 2.4 (Gromov compactness for holomorphic disks). *Let (X,ω) be an exact symplectic manifold, let $L \subset X$ be a compact Lagrangian submanifold, and let J be an ω-tame almost complex structure on X. Suppose $u^\nu : (D,\partial D) \to (X,L)$ is a sequence of holomorphic disks with bounded energy such that all images are contained in some compact subset of X. Then u^ν has a Gromov convergent subsequence.*

Moreover, for a convergent sequence u^ν the limit is unique up to equivalence.

A careful exposition of the proof of this theorem, without the simplifying assumption that X is exact, can be found in [9]. Rather than going into details here, I will illustrate the phenomenon by considering a specific example.

Example 2.5. The situation is interesting already for holomorphic disks with boundary on $S^1 \subset \mathbb{C}$ with respect to the standard complex structure J_0. By the maximum principle, these are necessarily maps $D \to D$. They exist for all non-negative degrees, and the degree d of the map equals the degree of its restriction to the boundary circle. It turns out that the Maslov index of the class $d \in \pi_1(S^1)$ is $2d$. Applying the general index formula (2.2), one finds that

$$\dim \mathcal{M}(d, J_0) = n - 2 + 2d = 2d - 1.$$

For $d = 1$, this index equals 1. Indeed, the group of holomorphic degree 1 maps, i.e. holomorphic automorphisms of D is 3-dimensional (it can be identified with $\mathrm{PSL}(2,\mathbb{R})$). In fact, any element of this group can be written uniquely as a composition $\varphi = \varphi_2 \circ \varphi_1$, where φ_2 is a rotation given by $\varphi_2(z) := \varphi(1)z$ and $\varphi_1 = \varphi_2^{-1} \circ \varphi \in \mathrm{Aut}(D,1)$ is a map fixing $1 \in S^1 = \partial D$. Since we only consider the moduli space of maps up to the equivalence relation of precomposing with an element of $\mathrm{Aut}(D,1)$, the moduli space can be identified with S^1 by recording the image of $1 \in S^1$ under any map in the equivalence class. In particular, the moduli space $\mathcal{M}(1, J_0)$ is compact.

Next consider a map of higher degree $d > 1$. Again any such map can be written as a composition $\varphi = \varphi_2 \circ \varphi_1$, where φ_1 is a product (not composition!) of d maps of degree 1, each fixing $1 \in S^1$, and φ_2 is a rotation. The map φ_1 is characterized completely in terms of its zeros, counted with multiplicities. In fact, if these zeros are $z_1, \ldots, z_d$, then we have

$$\varphi_1(z) = \prod e^{-i\theta_j} \frac{z - z_j}{-\bar{z}_j z + 1}, \quad \text{with } \theta_j = \arg\left(\frac{1 - z_j}{1 - \bar{z}_j}\right).$$

By precomposing with an appropriate $\psi \in \mathrm{Aut}(D,1)$, we may always arrange that one of the zeros, say z_d, equals 0. The coordinates of the other zeros give $2(d-1)$ free local parameters for φ_1, and together with the rotation parameter for φ_2 we get $2d - 1$ as predicted by the dimension formula above.

The moduli space $\mathcal{M}(d, J_0)$ of degree $d > 1$ self-maps of the disk D is non-compact. In fact, consider representative maps $\varphi^{(n)}$ of a sequence of points in the moduli space and orderings of the zeros $z_j^{(n)}$ of $\varphi^{(n)}$ such that $z_d^{(n)} = 0$ (as above, this can be arranged by precomposing a given representative with some element of $\mathrm{Aut}(D, 1)$, which does not change the equivalence class). After passing to a subsequence, we get convergence of the $z_j^{(n)}$ to some limiting $z_j^\infty \in D$ for each $j = 1, \ldots, d$. The formula above still makes sense in the limit, but the "zeros" z_j^∞ which lie on the circle S^1 contribute a trivial factor of 1 to the product. So if there are $0 < d' < d$ of these "phantom" zeros, the naive limiting map φ^∞ will have degree $d - d'$ and so it is not an element of $\mathcal{M}(d)$.

Notice that in the above discussion, we made two arbitrary choices: a choice of ordering $z_j^{(n)}$ of the zeros of $\varphi^{(n)}$, and the choice to always reparametrize so that $z_d^{(n)} = 0$. Suppose for definiteness that with these choices we have $z_1^\infty \in S^1$. Then there are unique maps $\psi^{(n)} \in \mathrm{Aut}(D, 1)$ such that $\psi^{(n)}(z_1^{(n)}) = 0$, and so we get a different sequence of representatives $\varphi^{(n)} \circ \psi^{(n)}$ of the same divergent sequence of points in the moduli space $\mathcal{M}(d)$. Just as above we get a, generally different, limiting map $\varphi^{\infty,*}$ for a suitable subsequence. Note that in this reparametrization, we will have $z_d^\infty \in S^1$, and so $\varphi^{\infty,*}$ has degree $d'' < d$.

If we would analyse the situation fully, we would recover, for a suitable subsequence, the existence of finitely many sequences of Möbius transformations $\psi_\alpha^{(n)}$, such that the reparametrized maps $\varphi^{(n)} \circ \psi_\alpha^{(n)}$ converge to some limiting map φ_α of degree $d_\alpha > 0$ in such a way that $\sum d_\alpha = d$. In addition, these maps fit together and form a disk tree as described in the compactness theorem above.

The compactified moduli space

Very roughly, the compactness theorem asserts that one can compactify a given space $\mathcal{M}(a, J)$ by adding pieces built out of moduli spaces $\mathcal{M}(b, J)$ with $0 < E(b) < E(a)$. This compactification is often denoted by $\overline{\mathcal{M}}(a, J)$. It admits an obvious stratification, where the stratum of codimension k corresponds to stable trees of disks modelled on trees with exactly k (unoriented) edges. Indeed, the heuristic dimension count proceeds as follows. Denote by r_α the number of special points on the disk associated to $\alpha \in T$ and by a_α its relative homotopy class. Note that $\sum r_\alpha = 2k + 1$, where k is the number of edges of the tree, since we had one marked point to start with and each edge gives rise to two nodal points. The formal dimension of the moduli space of disks associated to the vertex $\alpha \in T$ is

$$n - 3 + r_\alpha + \mu(\alpha).$$

Requiring that the nodal points corresponding to an edge in T are mapped to the same point in L gives $n = \dim L$ constraints. Putting these together, we find that the formal total dimension equals

$$(k + 1)(n - 3) + 2k + 1 + \mu(a) - kn = n - 2 + \mu(a) - k.$$

If, for a given J, all the moduli spaces appearing in the compactification $\overline{\mathcal{M}}(a, J)$ were transversely cut out, one could hope to prove a *gluing theorem*, asserting that in fact the compactified moduli space is a manifold with boundary and corners.

This is generally too much to ask. In [12], Fukaya, Oh, Ohta and Ono describe a procedure to put a so-called *Kuranishi structure* on the compactified moduli spaces. Without going into details, this roughly means that these spaces admit fundamental chains *that make them function as if they were manifolds with corners.* Theorem 3.3 below should be understood in this sense.

Presumably, the ongoing polyfold project of Hofer, Wysocki and Zehnder (cf. [16]) will eventually lead to an alternative approach to the problem of putting enough structure on the compactified moduli space to prove a statement like Theorem 3.3.

3 Gromov's Theorem and Fukaya's refinement

No exact Lagrangian submanifolds in $\mathbb{C}^n$

As already mentioned, it is instructive to review the proof for the following well-known theorem of Gromov.

Theorem 3.1 (Gromov, 1985). *If a compact manifold L admits a Lagrangian embedding into $\mathbb{C}^n$, then $H^1(L;\mathbb{R}) \neq 0$.*

Proof. (Sketch) I sketch the proof of this theorem given in [22, Section 9.2], slightly rephrasing the end of the argument in order to make the relation to the following discussion even more apparent.

Fix a Lagrangian embedding $L \subset \mathbb{C}^n$, and choose a vector $a \in \mathbb{C}^n$ with $\|a\| \geq 2\sup_{z\in L} \|z\|$. Consider the set $\mathcal{H} \subset C^\infty([0,1] \times D \times \mathbb{C}^n)$ of Hamiltonian functions such that

$$H^0_{s,t}(z) = 0, \quad H^1_{s,t}(z) = \langle a, z\rangle.$$

The idea is to consider, for a fixed $H \in \mathcal{H}$, the moduli space $\mathcal{N}$ of maps $u : (D^2, \partial D^2) \to (\mathbb{C}^n, L)$ satisfying the following conditions:

- $\partial_s u + J_0 \partial_t u = \nabla H^\lambda_{s,t}(u)$ for some $\lambda \in [0,1]$, and
- the relative homotopy class $[u] \in \pi_2(\mathbb{C}^n, L)$ vanishes.

So for $\lambda = 0$, we are considering holomorphic disks with boundary on the Lagrangian L, and since the relative homotopy class vanishes, these are precisely the constant maps. Note that we do not divide out any automorphisms here, since for positive λ these have no reason to preserve the solution space to the equation. One can prove that for fixed small $\lambda > 0$, there is still a compact n-dimensional family of solutions. In fact, for generic choice of the Hamiltonian $H \in \mathcal{H}$, standard transversality techniques show that $\mathcal{N}$ is a smooth $n+1$-dimensional manifold, whose boundary consists of those elements with $\lambda \in \{0,1\}$.

On the other hand, a straightforward computation as in Example 2.2 shows that, for our choice of a, there are no solutions to the equation with $\lambda = 1$. So if $\mathcal{N}$ was compact, it would give a smooth cobordism from L to the empty set. Now consider the evaluation map

$$ev : \mathcal{N} \to \Lambda L, \quad u \mapsto u|_{\partial D^2}.$$

From what we said above, the boundary of this $(n+1)$-chain in the free loop space of L is the cycle of constant loops $[L] \in C_n(\Lambda L)$. Since this cycle is nontrivial in homology, $\mathcal{N}$ cannot be compact.

It follows from elliptic regularity theory (cf. [22, Theorem 4.1.1]) that if compactness fails, there is a sequence $u_n \in \mathcal{N}$ such that $|du_n|_\infty \to \infty$ as $n \to \infty$. Appropriately rescaling such a sequence and applying removal of singularities as in [22, section 4.2], one finds either a nonconstant holomorphic sphere or a nonconstant holomorphic disk with boundary on L. Since $\mathbb{C}^n$ does not contain nonconstant holomorphic spheres (such spheres would be contractible and have positive energy ,contradicting Stokes' theorem), the only possibility is the existence of some nonconstant holomorphic disk $v : (D^2, \partial D^2) \to (\mathbb{C}^n, L)$.

Now the standard symplectic form $\omega = \sum_{j=1}^n dx_j \wedge dy_j$ is positive on all complex lines in $\mathbb{C}^n$, and hence on all the tangent planes to the image of v, so we have

$$\int_{D^2} v^*(\omega) > 0.$$

Moreover, $\omega = d\eta$, where e.g. $\eta = \sum_{j=1}^n x_j dy_j$, and so Stokes' theorem implies that

$$\int_{S^1} v^*(\eta) > 0.$$

On the other hand, the Lagrangian condition states that $\omega = d\eta$ vanishes pointwise when restricted to L. Combining these observations, it follows that $\eta|_L$ is a closed 1-form representing a nonzero class in $H^1(L; \mathbb{R})$, and this proves the theorem. $\square$

The technical outcome of analysing holomorphic disks

What was the essence of the proof of Gromov's theorem? Basically, the point is that the space $\mathcal{N}$ has a single boundary component, corresponding to the space of constant disks, and hence for topological reasons it cannot be compact. Analysing the breakdown of compactness, we found holomorphic disks.

The elements of $\mathcal{N}$ do not appear to be holomorphic curves, due to the nonzero right hand side $\nabla H^\lambda_{s,t}$ of the equation. However, they can in fact be viewed as holomorphic maps into $\mathbb{C}^n \times D$ with respect to a family of almost complex structures which have an off-diagonal term built out of this right hand side, projecting holomorphically and with degree 1 to the disk. This basic phenomenon was already present in Example 2.2.

The stable map compactification in this particular case is given by bubble trees of disks with exactly one main component, satisfying the original equation, and all other bubbles being strictly holomorphic (in the graph picture just mentioned, each of these has constant projection to the disk). Fukaya's insight was to see that the new boundary can be described in terms of string topology operations.

Namely, for each $a \in \pi_2(\mathbb{C}^n, L)$, consider the compactified moduli space $\overline{\mathcal{M}}(a) := \overline{\mathcal{M}}(a, J_0)$ of holomorphic disks in the relative homotopy class a. Similarly, denote by $\mathcal{N}(a)$ the space of solutions $u : (D, \partial D) \to (\mathbb{C}^n, L)$ to the equation

$$\partial_s u + J_0 \partial_t u = \nabla H^\lambda_{s,t} \quad \text{for some } \lambda \in [0, 1]$$

in the relative homotopy class $a \in \pi_2(\mathbb{C}^n, L)$. Pretending as always that transversality holds, this space is a manifold of dimension

$$\dim \mathcal{N}(a) = n + 1 + \mu(a),$$

and we denote its compactification by $\overline{\mathcal{N}}(a)$.

Assuming the analysis can be made to work, both $\overline{\mathcal{M}}(a)$ and $\overline{\mathcal{N}}(a)$ can be thought of as chains on ΛL, simply by associating to each map of the disk its restriction to the boundary circle. In fact, for $\overline{\mathcal{M}}$ there is a slight ambiguity, since its elements are only well-defined up to precomposition by $\varphi \in \mathrm{Aut}(D, 1)$. But one can easily get around this point, for example by replacing the actual map by a parametrization proportional to arc length.

To arrive at a clean statement, I will introduce some further notation. Suppose we are given a suitable model $C_*(\Lambda L)$ for the chains on the free loop space of L with coefficients in $\mathbb{Q}$, such that for each $a \in \pi_2(\mathbb{C}^n, L)$ the compactified spaces $\overline{\mathcal{M}}(a)$ and $\overline{\mathcal{N}}(a)$, with their respective evaluation maps to ΛL, define elements in it. Note that ΛL is a disjoint union over its connected components $\Lambda_\gamma L$, which can be identified with conjugacy classes γ of elements of $\pi_1(L)$. It is convenient to introduce a new complex $\mathcal{C}$ whose underlying vector space is $C_*(\Lambda L)$, but with grading shifted according to the Maslov index, i.e. an element in $C_k(\Lambda_\gamma L)$ will have degree $k - \mu(\gamma)$ in $\mathcal{C}$.

The complex $\mathcal{C}$ comes with a filtration by the symplectic area as follows. It is shown in [22, Prop. 4.1.4] that the infimum

$$\hbar := \inf\{E(u) \mid u : (D, \partial D) \to (\mathbb{C}^n, L) \text{ nonconstant and holomorphic}\} \tag{3.1}$$

of the symplectic energy is strictly positive. Since each loop on L bounds a disk in $\mathbb{C}^n$, we can view the energy as a map $E : \pi_0(\Lambda L) \to \mathbb{R}$. The energy of $c = \sum c_i \in \mathcal{C}$ is now defined as $E(\sum c_i) := \sum E(c_i)$, where $E(c_i)$ is the energy of the free homotopy class of the loops parametrized by the chain c_i (assumed to have connected domain of definition). Then the filtration $\{\mathcal{F}^k\}_{k\in\mathbb{Z}}$ on $\mathcal{C}$ is given by

$$\mathcal{F}^k := \{c \in \mathcal{C} : E(c) \geq k\hbar\}, \quad k \in \mathbb{Z}.$$

Now consider the completion $\widehat{\mathcal{C}}$ of $\mathcal{C}$ with respect to this filtration. This means that an element in $\widehat{\mathcal{C}}$ will be a possibly infinite sum $c = \sum c_i$ of chains $c_i \in \mathcal{C}$, provided that for each $k \in \mathbb{Z}$ there are only a finite number of summands satisfying $c_i \notin \mathcal{F}^k$.

With this definition, Gromov compactness and our grading convention imply that

$$\overline{\mathcal{M}} := \sum_{a \in \pi_2(\mathbb{C}^n, L) \setminus \{0\}} \overline{\mathcal{M}}(a)$$

is a well-defined element of $\widehat{\mathcal{C}}$ of degree $n-2$. In fact, it is contained in the submodule $\mathcal{F}^1 \subset \widehat{\mathcal{C}}$ of chains with strictly positive area. Similarly,

$$\overline{\mathcal{N}} := \sum_{a \in \pi_2(\mathbb{C}^n, L)} \overline{\mathcal{N}}(a)$$

is an element of $\widehat{\mathcal{C}}$ of degree $n+1$, as follows by applying the analogue of (3.1) to the graph of elements of $\mathcal{N}(a)$ in $D \times \mathbb{C}^n$.

As explained at the end of section 2, the compactification of $\overline{\mathcal{M}}(a)$ is obtained by adding lower dimensional strata built as fiber products of other such moduli spaces along evaluation maps. In particular, the codimension 1 pieces $\lambda_1(\overline{\mathcal{M}}(a))$ are build from configurations of two holomorphic disks for which suitable boundary points are mapped to the same point in L. A more careful analysis reveals that, on the level of boundary values of the holomorphic maps, these configurations correspond to loop brackets $\lambda_2(\overline{\mathcal{M}}(a_1), \overline{\mathcal{M}}(a_2))$ with $a_1 + a_2 = a$.

Example 3.2. The mechanism just described can be seen in Example 2.5, but maybe it is slightly easier to visualize for the standard Lagrangian torus $T^2 = S^1 \times S^1 \subset \mathbb{C}^2$. Any class $a \in \pi_2(\mathbb{C}^2, T^2)$ is characterized by two integers (d_1, d_2) giving the degrees of the projections to the two coordinate disks. For a moduli space $\overline{\mathcal{M}}((d_1, d_2), J_0)$ with respect to the standard complex structure J_0 on $\mathbb{C}^2$ to be nontrivial we need $d_j \geq 0$. Leaving aside the constant maps, the simplest moduli spaces $\mathcal{M}((1,0), J_0)$ and $\mathcal{M}((0,1), J_0)$ are compact, and in fact one can identify both with T^2. Indeed, the equivalence classes of the maps $u_{z_1,z_2} : D^2 \to \mathbb{C}^2$ given by $u_{z_1,z_2}(z) = (z_1 z, z_2)$ for $z_1, z_2 \in S^1$ represent all elements in $\mathcal{M}((1,0), J_0)$, and similarly the maps $v_{z_1,z_2}(z) = (z_1, z_2 z)$ represent all elements in $\mathcal{M}((0,1), J_0)$. Note that because of the symmetries in the problem, in this particularly simple example the evaluation maps at 1 are submersions, so that the geometric definition of the loop bracket $\lambda_2(\mathcal{M}((0,1)), \mathcal{M}((1,0)))$ can be used.

To illustrate the discussion above, we want to argue that

$$\begin{aligned} \partial\overline{\mathcal{M}}((1,1), J_0) &= \lambda_2(\mathcal{M}((0,1)), \mathcal{M}((1,0))) \\ &= \frac{1}{2}\Big[\lambda_2(\mathcal{M}((0,1)), \mathcal{M}((1,0))) + \lambda_2(\mathcal{M}((1,0)), \mathcal{M}((0,1)))\Big]. \end{aligned}$$

Every element of the space $\widetilde{\mathcal{M}}((1,1), J_0)$ is a map $u : D^2 \to \mathbb{C}^2$ of the form $u(z) = (z_1 \phi_1(z), z_2 \phi_2(z))$ with $z_j \in S^1$ and $\phi_j \in \mathrm{Aut}(D, 1)$. The space $\mathcal{M}((1,1), J_0)$ is obtained as the quotient by the diagonal action of $\mathrm{Aut}(D,1)$, so the equivalence class of u as above is alternatively represented by both $u'(z) = (z_1 z, z_2 \psi(z))$ or $u''(z) = (z_1 \psi^{-1}(z), z_2 z)$, where $\psi = \phi_2 \circ \phi_1^{-1}$ is uniquely associated with the equivalence class of u.

Now consider a sequence $u_n \in \widetilde{\mathcal{M}}((1,1), J_0)$ with z_1 and z_2 fixed but $\phi_{1,n}$ and $\phi_{2,n}$ varying. Assume that the projection of the sequence to $\mathcal{M}((1,1), J_0)$ leaves every compact subset, meaning that the corresponding sequence $\psi_n \in \mathrm{Aut}(D,1)$ in the above notation does the same. Elementary considerations now show that for a suitable subsequence ψ_{n_k} there will be a point $w \in S^1$ such that

(i) $\psi_{n_k} \longrightarrow w$ uniformly on compact subsets of $D \setminus \{1\}$ and $\psi_{n_k}^{-1} \longrightarrow 1$ uniformly on compact subsets of $D \setminus \{w\}$, or

(ii) $\psi_{n_k} \longrightarrow 1$ uniformly on compact subsets of $D \setminus \{w\}$ and $\psi_{n_k}^{-1} \longrightarrow w$ uniformly on compact subsets of $D \setminus \{1\}$.

In both cases, the corresponding subsequences u'_{n_k} and u''_{n_k} converge to elements $u'_\infty \in \mathcal{M}((1,0), J_0)$ and $u''_\infty \in \mathcal{M}((0,1), J_0)$, respectively, and the pair (u'_∞, u''_∞) represents a boundary point of $\overline{\mathcal{M}}((1,1), J_0)$. One checks that as z_1 and z_2 and the sequence ψ_{n_k} vary, one obtains all boundary points from this construction. In case (i), $u''_\infty(1) = (z_1, z_2)$ and $u'_\infty(1) = (z_1, z_2 w)$, which is the unique intersection point of $u'_\infty(S^1)$ and $u''_\infty(S^1)$, and in case (ii) the roles are reversed. In particular, the boundary loops of the two limit disks concatenate to represent points in the loop bracket $\lambda_2(\mathcal{M}((0,1), J_0), \mathcal{M}((1,0), J_0))$.

We now return to the general discussion. Similarly to the case of $\overline{\mathcal{M}}$, the codimension 1 stratum for $\overline{\mathcal{N}}(a)$ corresponds to stable maps consisting of one component satisfying the perturbed equation and one holomorphic disk, and so it is described by the loop brackets of the form $\lambda_2(\overline{\mathcal{N}}(a_1), \overline{\mathcal{M}}(a_2))$. Depending on the precise technical implementation, the gluing along lower dimensional boundary strata might actually introduce more terms, corresponding to higher operations.

The main technical assertions which should come out of such an implementation can be formulated as the following theorem. To get a cleaner statement, I have chosen to state it in slightly stronger form than is strictly necessary. The concept of a filtered L_∞ algebra which appears in the statement is discussed in detail in the following section, where I also give some algebraic perspective on the equations (3.2) and (3.3).

Theorem 3.3. *Let $L \subset \mathbb{C}^n$ be a closed, oriented, spin Lagrangian submanifold.*

Then on the filtered, degree-shifted chain complex $\widehat{\mathcal{C}}$ associated to a suitable chain model $C_*(\Lambda L)$ for the free loop space ΛL there exists a filtered L_∞-algebra structure $\{\lambda_k\}_{k \geq 1}$ of degree $1-n$, whose bracket on homology coincides with the loop bracket of string topology, and such that

(1) *the union of moduli spaces $\overline{\mathcal{M}}$ gives rise to an element $\alpha \in \widehat{\mathcal{C}}_+$ of degree $n-2$ satisfying*
$$\sum_{k=1}^{\infty} (-1)^{\frac{(k-1)k}{2}} \frac{1}{k!} \lambda_k(\alpha, \cdots, \alpha) = 0. \tag{3.2}$$

(2) *the union of moduli spaces $\overline{\mathcal{N}}$ gives rise to an element $\beta \in \widehat{\mathcal{C}}$ of degree $n+1$ satisfying*
$$\sum_{k=1}^{\infty} (-1)^{\frac{(k-2)(k-1)}{2}} \frac{1}{(k-1)!} \lambda_k(\beta, \alpha, \cdots, \alpha) = [L], \tag{3.3}$$
where $[L] \in C_n(\Lambda L, \mathbb{Q})$ denotes the chain of constant loops.

I will treat this theorem as a black box, and deduce the main results in the introduction from it by using abstract algebraic arguments and some 3-manifold topology.

4 Some algebraic properties of L_∞ algebras

For a graded vector space $C = \oplus_{d\in\mathbb{Z}} C_d$ we denote by $C[n]$ the vector space with grading shifted by n, i.e. $C[n]_d = C_{d+n}$. On the k-fold tensor product $C \otimes \cdots \otimes C$, we consider two actions of the permutation group S_k. In the first one, a permutation $\rho \in S_k$ acts on some tensor product of elements $c_i \in C$ of pure degrees $|c_i|$ via

$$\rho\cdot(c_1\otimes\cdots\otimes c_k) = \varepsilon(\rho;c_1,\dots,c_k)\cdot c_{\rho(1)}\otimes\cdots\otimes c_{\rho(k)}$$

with $\varepsilon(\rho;c_1,\dots,c_k) = (-1)^{\sum_{i<j,\rho(i)>\rho(j)}|c_i|\cdot|c_j|}$. The quotient is the kth symmetric power S^kC of C, whose decomposable elements we write as $c_1\cdots c_k$. The second action is the first one twisted by the sign representation,

$$\rho\cdot(c_1\otimes\cdots\otimes c_k) = \mathrm{sgn}(\rho)\varepsilon(\rho;c_1,\dots,c_k)\cdot c_{\rho(1)}\otimes\cdots\otimes c_{\rho(k)}.$$

The quotient is the kth exterior power Λ^kC of C, whose elements are usually denoted by $c_1\wedge\cdots\wedge c_k$. With these definitions, for an element $c\in C$ of odd degree we have $c\cdot c = 0$, but $c\wedge c\neq 0$.

Definition 4.1. An L_∞ algebra of degree 0 consists of a graded vector space C and a sequence of multilinear operations

$$\lambda_k : \Lambda^kC \to C, \quad k\geq 1$$

of degree $|\lambda_k| = k-2$ satisfying the sequence of quadratic relations

$$\sum_{\substack{k_1+k_2=k+1,\\ \rho\in S_k}} \pm\frac{1}{k_1!(k-k_1)!}\lambda_{k_2}(\lambda_{k_1}(c_{\rho(1)},\dots,c_{\rho(k_1)}),c_{\rho(k_1+1)},\dots,c_{\rho(k)}) = 0 \qquad (4.1)$$

for each $k\geq 1$. (The signs are made explicit below.)

More generally, an L_∞ algebra structure of degree d on C is defined to be an L_∞ structure of degree 0 on $C[-d]$.

Remark 4.2. If the vector space C of an L_∞ algebra of degree 0 is concentrated in degree 0, then for degree reasons the only possibly nontrivial operation is λ_2, and the relation for $k=3$ turns out to be the Jacobi identity for λ_2, so we recover Lie algebras as a special case.

Remark 4.3. If $\lambda_k = 0$ for $k\geq 3$, then we recover the definition of a dg Lie algebra. Indeed, the first relation reads $\lambda_1\circ\lambda_1 = 0$. The second relation shows that λ_1 is a derivation of λ_2, and the third relation is again the Jacobi identity. In general, the Jacobi identity holds "up to homotopy" given by λ_3, so it always holds for the induced bracket on $H_*(C,\lambda_1)$.

To make the signs in the quadratic relations as well as other signs below explicit, it is useful to give an alternative description. First observe the graded linear isomorphism

$$\begin{aligned}\sigma_k : (\Lambda^k C)[-k] &\to S^k(C[-1])\\ c_1 \wedge \cdots \wedge c_k &\mapsto (-1)^{\sum (k-i)|c_i|} c_1 \cdots c_k,\end{aligned}$$

where $|c_i|$ denotes the degree in C. Next introduce operations $\ell_k : S^k(C[-1]) \to C[-1]$ as $\ell_k = \sigma_1 \circ \lambda_k \circ \sigma_k^{-1}$ and note that with this degree shift these are all of degree -1. Set $S(C[-1]) := \oplus_{k\geq 1} S^k(C[-1])$ and observe that each of these operations can be extended to a map $\hat{\ell}_k : S(C[-1]) \to S(C[-1])$ defined as

$$\hat{\ell}_k(c_1 \cdots c_r) = \begin{cases} 0 & \text{if } r < k \\ \sum_{\rho\in S_r} \frac{\varepsilon(\rho, c_1, \ldots, c_r)}{k!(r-k)!} \ell_k(c_{\rho(1)} \cdots c_{\rho(k)}) c_{\rho(k+1)} \cdots c_{\rho(r)} & \text{if } r \geq k, \end{cases}$$

where $\varepsilon(\rho, c_1, \ldots, c_r)$ is the sign introduced above

Finally, one defines $\hat{\ell} := \sum_{k\geq 1} \hat{\ell}_k : S(C[-1]) \to S(C[-1])$. Then the quadratic relations (4.1) (with the correct signs) are equivalent to the single equation

$$\hat{\ell} \circ \hat{\ell} = 0. \tag{4.2}$$

The above passage from the operations λ_k on C to the operation $\hat{\ell}$ on $S(C[-1])$ is called the *bar construction.* Conceptually, one views $S(C[-1])$ as a coalgebra via the comultiplication $\Delta : S(C[-1]) \to S(C[-1]) \otimes S(C[-1])$ given by

$$\Delta(c_1 \cdots c_r) = \sum_{r_1=1}^{r-1} \sum_{\rho\in S_r} \frac{\varepsilon(\rho, c_1, \ldots, c_r)}{r_1!(r-r_1)!} c_{\rho(1)} \cdots c_{\rho(r_1)} \otimes c_{\rho(r_1+1)} \cdots c_{\rho(r)}.$$

This map has the coassociativity property

$$(\mathbb{1} \otimes \Delta) \circ \Delta = (\Delta \otimes \mathbb{1}) \circ \Delta,$$

and it also turns out to be cocommutative in the sense that $\tau \circ \Delta = \Delta$, where $\tau : S(C[-1]) \otimes S(C[-1]) \to S(C[-1]) \otimes S(C[-1])$ is the signed permutation of the two factors. Then $\hat{\ell}_k$ is the unique way to extend ℓ_k as a coderivation, i.e. as a map satisfying the co-Leibniz rule

$$\Delta \hat{\ell}_k = (\hat{\ell}_k \otimes \mathbb{1} + \mathbb{1} \otimes \hat{\ell}_k)\Delta.$$

Conversely, one can prove that any coderivation $D : S(C[-1]) \to S(C[-1])$ is completely determined by its *linear part* $\pi_1 \circ D : S(C[-1]) \to C[-1]$. So $\hat{\ell}$ is the unique coderivation of degree -1 on $S(C[-1])$ such that the restriction of its linear part to $S^k(C[-1])$ equals ℓ_k.
It is also easy to see that the commutator $[D_1, D_2] := D_1 \circ D_2 - (-1)^{|D_1||D_2|} D_2 \circ D_1$ of two homogeneous coderivations is a coderivation, and so in our example above $\hat{\ell} \circ \hat{\ell} = \frac{1}{2}[\hat{\ell}, \hat{\ell}]$ has this property. These remarks explain why the relation (4.2) is equivalent to the sequence of relations (4.1), since this sequence is obtained by restricting the linear part of $\hat{\ell} \circ \hat{\ell}$ to $S^k(C[-1])$ for each $k \geq 1$ (and precomposing with σ_k).

So in summary, an L_∞ structure on a graded vector space C is the same as a coderivation of square zero on the symmetric tensor coalgebra $S(C[-1])$.

Remark 4.4. I have adopted homological conventions here, whereas often in the literature one finds cohomological conventions, where the λ_k have degrees $2-k$, and in the bar construction one shifts degrees by 1 instead of -1.

Definition 4.5. Given two L_∞ algebras $\mathcal{C} = (C, \{\lambda_k\}_{k\geq 1})$ and $\mathcal{C}' = (C', \{\lambda'_k\}_{k\geq 1})$, a morphism from $\mathcal{C}$ to $\mathcal{C}'$ consists of a sequence of maps $\phi_k : \Lambda^k C \to C'$ of degrees $|\phi_k| = k-1$ satisfying the sequence of relations

$$\sum_{k_1 k_2 = k+1} \pm \phi_{k_1} \circ \hat{\lambda}_{k_2} = \sum_{k_1+\cdots+k_r=k} \pm \frac{1}{r!} \lambda'_r \circ (\phi_{k_1} \otimes \cdots \otimes \phi_{k_r}) \tag{4.3}$$

for $k \geq 1$.

Again, to state the signs correctly, it is useful to pass to the associated maps $f_k : S^k(C[-1]) \to C'[-1]$ of degree 0 given by $f_k = \sigma_1 \circ \phi_k \circ \sigma_k^{-1}$. Any such collection of linear maps determines a unique morphism of coalgebras $e^f : S(C[-1]) \to S(C'[-1])$, given by

$$e^f(c_1 \cdots c_k) = \sum_{k_1+\cdots+k_r=k} \sum_{\rho \in S_k} \frac{\varepsilon(\rho, c_1, \ldots, c_k)}{r! k_1! \cdots k_r!} (f_{k_1} \otimes \cdots \otimes f_{k_r})(c_{\rho(1)} \cdots c_{\rho(k)}).$$

The fact that $\{\phi_k\}$ is a morphism of L_∞ algebras can now be stated equivalently (including the correct signs) as

$$e^f \hat{\ell} = \hat{\ell}' e^f. \tag{4.4}$$

The first important result about L_∞ algebras asserts that the structure of an L_∞ algebra can be transferred from a complex C to its homology with respect to λ_1, without the loss of any essential information. More precisely, it is formulated as follows.

Theorem 4.6. *Suppose $\mathcal{C} = (C, \{\lambda_k\}_{k\geq 1})$ is an L_∞ algebra over a field of characterstic 0. Then there exists an L_∞ algebra structure $\mathcal{H} = (H_*(C, \lambda_1), \{\lambda'_k\}_{k\geq 2})$ on the homology which is homotopy equivalent to $\mathcal{C}$.*

Here a *homotopy equivalence* between L_∞ algebras is the essentially obvious generalization of the classical notion. In particular, it is an L_∞ morphism which induces an isomorphism in the homology of the underlying complexes. It is a theorem that every such map admits a homotopy inverse. For detailed definitions and a proof of these assertions, including the theorem, see e.g. [20].

The construction of the homotopy equivalence starts with a linear homotopy equivalence $\iota : H_*(C, \lambda_1) \to C$ given by choosing a cycle in each homology class, which has a homotopy inverse $\pi : C \to H_*(C, \lambda_1)$ given by projection along a complement of the image of ι. One sets $\phi_1 = \iota$ and $\lambda'_2 = \pi \circ \lambda_2 \circ \iota \otimes \iota$, and constructs the higher maps $\phi_k, k \geq 2$ and operations λ'_k, $k \geq 3$ simultaneously by induction. The fact that the homologies of the two complexes agree is used to prove that all relevant obstructions vanish.

Now let $\mathcal{C} = (C, \{\lambda_k\}_{k\geq 1})$ be an L_∞ algebra. Suppose that C is the completion of some complex C' with respect to a doubly infinite filtration $C' = \cup_{k\in\mathbb{Z}} \mathcal{F}'_k$ with

$\mathcal{F}'_k \supset \mathcal{F}'_{k+1}$, so that elements of C are (possibly infinite) sums of elements of C' of the form

$$c = c_{-r} + \cdots + c_{-1} + c_0 + c_1 + \ldots, \quad c_k \in \mathcal{F}'_k.$$

Denote the induced filtration on C by $\{\mathcal{F}_k\}$.

Definition 4.7. The L_∞ structure on C is called *filtered* if

$$\lambda_k(\mathcal{F}_{d_1}, \ldots, \mathcal{F}_{d_k}) \subset \mathcal{F}_{d_1+\cdots+d_k}.$$

In the following discussion, it is convenient to denote by $\bar{c} = \sigma(c)$ the image of an element under the identity map $\sigma : C \to C[-1]$ of degree $+1$. An element $a \in \mathcal{F}_1$ of degree -1 satisfying the equation

$$\sum_{k\geq 1} \frac{1}{k!} \ell_k(\bar{a}, \ldots, \bar{a}) = 0 \quad \text{in } C[-1]$$

is called a *Maurer-Cartan element of* $\mathcal{C}$. Since $a \in \mathcal{F}_1$, the left hand side of this equation is indeed a well-defined element of $\mathcal{C}$. Note that (3.2) is an instance of this equation. Also, this equation is equivalent to

$$\hat{\ell}(e^{\bar{a}}) = 0.$$

Moreover, an easy calculation yields

Lemma 4.8. *If $a \in C$ is a Maurer-Cartan element of $\mathcal{C}$ and $b \in C$ is arbitrary, then*

$$\hat{\ell}(\bar{b}e^{\bar{a}}) = \sum_{k\geq 1} \frac{1}{(k-1)!} \ell_k(\bar{b}, \bar{a}, \ldots, \bar{a}) e^{\bar{a}}.$$

In particular, the map $\hat{\ell}^a : C[-1] \to C[-1]$ given by

$$\hat{\ell}^a(\bar{b}) = \hat{\ell}(\bar{b}e^{\bar{a}})e^{-\bar{a}} = \sum_{k\geq 1} \frac{1}{(k-1)!} \ell_k(\bar{b}, \bar{a}, \ldots, \bar{a})$$

is a differential. □

Finally, we describe what happens to equations (3.2) and (3.3) under morphisms.

Proposition 4.9. *Suppose $\{\phi_k\}_{k\geq 1}$ is a morphism between L_∞ algebras $\mathcal{C}$ and $\mathcal{D}$ preserving filtrations as above.*

(1) *If $a \in C$ is a Maurer-Cartan element for $\mathcal{C}$, then $a' \in D$ with $\bar{a}' = (\sum_k \frac{1}{k!} f_k(\bar{a}, \ldots, \bar{a}))$ is a Maurer-Cartan element for $\mathcal{D}$.*

(2) *If $a \in C$ is a Maurer-Cartan element and $b, c \in C$ satisfy*

$$\hat{\ell}^{\mathcal{C}}(\bar{b}e^{\bar{a}}) = \bar{c}e^{\bar{a}},$$

then the elements a', b' and c' with

$$\bar{a}' = \left(\sum_k \frac{1}{k!} f_k(\bar{a}, \ldots, \bar{a})\right),$$
$$\bar{b}' = \left(\sum_k \frac{1}{(k-1)!} f_k(\bar{b}, \bar{a}, \ldots, \bar{a})\right) \quad \text{and}$$
$$\bar{c}' = \left(\sum_k \frac{1}{(k-1)!} f_k(\bar{c}, \bar{a}, \ldots, \bar{a})\right)$$

satisfy

$$\hat{\ell}^{\mathcal{D}}(\bar{b}' e^{\bar{a}'}) = \bar{c}' e^{\bar{a}'}.$$

Proof. To prove the first assertion, just observe that for a Maurer-Cartan element $a \in C$ one has

$$0 = e^f \hat{\ell}^{\mathcal{C}}(e^{\bar{a}}) = \hat{\ell}^{\mathcal{D}} e^f(e^{\bar{a}}) = \hat{\ell}^{\mathcal{D}}(e^{\bar{a}'}),$$

where the equality $e^f(e^{\bar{a}}) = e^{\bar{a}'}$ follows directly from the definitions.

To prove the second assertion, one first checks that for any elements $\bar{x}, \bar{y} \in C[-1]$

$$e^f(\bar{x} e^{\bar{y}}) = \left(\sum_{k \geq 1} \frac{1}{(k-1)!} f_k(\bar{x}, \bar{y}, \ldots, \bar{y})\right) e^{\sum \frac{1}{r!} f_r(\bar{y}, \ldots, \bar{y})}.$$

Using this, we compute

$$e^f \hat{\ell}^{\mathcal{C}}(\bar{b} e^{\bar{a}}) = \hat{\ell}^{\mathcal{D}} e^f(\bar{b} e^{\bar{a}}) = \hat{\ell}^{\mathcal{D}}(\bar{b}' e^{\bar{a}'}).$$

On the other hand,

$$e^f \hat{\ell}^{\mathcal{C}}(\bar{b} e^{\bar{a}}) = e^f(\bar{c} e^{\bar{a}}) = \bar{c}' e^{\bar{a}'}.$$

□

5 The proofs of Theorem 1.1 and Corollary 1.2

Looking back at Theorem 3.3, we see that it asserts that the holomorphic disks with boundary on the Lagrangian submanifold give rise to a Maurer-Cartan element α in the L_∞ structure on $\widehat{\mathcal{C}}$ such that with respect to the twisted differential the element $[L] \in \widehat{\mathcal{C}}$ becomes exact. Since $[L]$ is never exact with respect to the ordinary boundary operator $\partial = \lambda_1$, this tells us that both the Maurer-Cartan element α and at least one of the operations λ_k with $k \geq 2$ must be nontrivial, since otherwise the twisted differential coincides with the untwisted boundary operator ∂. This observation lies at the core of Fukaya's proof of Theorem 1.1. Before I discuss that, I will state two purely topological facts that will turn out to be useful.

Lemma 5.1. *Let $\gamma : S^1 \to L$ be a loop, and denote by $Z \subset \pi_1(L)$ the centralizer of γ, i.e. the set of all elements commuting with γ. Let $\pi : \tilde{L} \to L$ be a connected covering of L associated to the subgroup Z, and let $\tilde{\gamma}$ be a lift of γ. Then the projection π induces a homeomorphism $\Pi : \Lambda_{\tilde{\gamma}}\tilde{L} \to \Lambda_\gamma L$ between the components of $\tilde{\gamma}$ and γ in the respective free loop spaces.*

Proof. Since $\pi : \tilde{L} \to L$ is a covering, any free homotopy $h : [0,1] \times S^1 \to L$ with $h|_{\{0\}\times S^1} = \gamma$ admits a (unique) lift $\tilde{h}$ to $\tilde{L}$ with $\tilde{h}|_{\{0\}\times S^1} = \tilde{\gamma}$, and so in particular $h|_{\{1\}\times S^1}$ is the image of $\tilde{h}|_{\{1\}\times S^1}$ under the map Π induced by the projection. This proves surjectivity of Π.

To prove injectivity, assume that $\Pi(\tilde{\delta}_1) = \Pi(\tilde{\delta}_2) = \delta$. Note that our two lifts $\tilde{\delta}_1$ and $\tilde{\delta}_2$ of δ are related by a deck transformation, i.e. by the action of some homeomorphism $g : \tilde{L} \to \tilde{L}$ satisfying $\pi \circ g = \pi$. If $\tilde{h}_1$ is a free homotopy from $\tilde{\gamma}$ to $\tilde{\delta}_1$, then $g \circ \tilde{h}_1$ is a homotopy from $g \circ \tilde{\gamma}$ to $\tilde{\delta}_2$. Since by assumption $\tilde{\delta}_2$ is also freely homotopic to $\tilde{\gamma}$, we conclude that if Π is not injective, then γ has at least two preimages, namely $\tilde{\gamma}$ and $g \circ \tilde{\gamma}$.

Now suppose $\tilde{\gamma}$ and $g \circ \tilde{\gamma}$ are freely homotopic for some deck transformation g, so that they are both preimages of γ under Π. A free homotopy $\tilde{h}$ from $\tilde{\gamma}$ to $g \circ \tilde{\gamma}$ can be reinterpreted as a based homotopy from $\tilde{\gamma}$ to $\chi * (g \circ \tilde{\gamma}) * \chi^{-1}$, where $\chi = \tilde{h}|_{[0,1]\times\{1\}}$ is the path travelled by the base point under the homotopy. Note that χ projects to a closed loop in L representing $\hat{g} \in \pi_1(L)$. In particular, the projection of the homotopy $\tilde{h}$ yields that

$$\gamma \cong \hat{g}\gamma\hat{g}^{-1} \quad \text{in } \pi_1(L).$$

But $Z = \pi_*(\pi_1(\tilde{L}))$ was chosen to be the centralizer of γ in $\pi_1(L)$, so $\hat{g} \in \pi_*(\pi_1(\tilde{L}))$. In other words, this implies that χ was a closed loop and so $\tilde{\gamma} = g \circ \tilde{\gamma}$, i.e. any two preimages of γ coincide. Together with the previous observation this shows that Π is injective, completing the proof of the lemma. $\square$

Lemma 5.2. *In the situation of the previous lemma, assume moreover that L (and so $\tilde{L}$ as well) is aspherical. Then evaluation at the base point $ev : \Lambda_{\tilde{\gamma}}\tilde{L} \to \tilde{L}$ is a homotopy equivalence.*

Proof. The fiber of the map $ev : \Lambda_{\tilde{\gamma}}\tilde{L} \to \tilde{L}$ at $\tilde{\gamma}(0)$ is the space $\Omega_{\tilde{\gamma}}\tilde{L}$ of loops which are based at $\tilde{\gamma}(0)$ and *freely* homotopic to $\tilde{\gamma}$. As in the previous proof, we observe that any free homotopy between $\delta \in \Omega_{\tilde{\gamma}}\tilde{L}$ and $\tilde{\gamma}$ can be reinterpreted as a based homotopy between δ and $\chi * \tilde{\gamma} * \chi^{-1}$. But $\tilde{\gamma}$ is central in $\pi_1(\tilde{L})$, so that $\chi * \tilde{\gamma} * \chi^{-1}$ is *based* homotopic to $\tilde{\gamma}$. So we conclude that in fact $\Omega_{\tilde{\gamma}}\tilde{L}$ is the component of $\tilde{\gamma}$ in the based loop space of $\tilde{L}$, which is contractible since $\tilde{L}$ is aspherical. So ev is a fibration with contractible fibers, and hence a homotopy equivalence. $\square$

Corollary 5.3. *If L is an aspherical manifold, then every component of the free loop space ΛL has the homotopy type of a CW complex of dimension at most* $\dim L$. $\square$

After these preliminaries, I come to the proof of the main theorem.

Proof. (of Theorem 1.1) Recall that the Maurer-Cartan element $\alpha \in \widehat{\mathcal{C}}$ is built from the moduli spaces $\{\overline{\mathcal{M}}(a)\}_{a\in\pi_2(\mathbb{C}^n,L)}$, which have geometric dimensions

$$\dim \overline{\mathcal{M}}(a) = n - 2 + \mu(a),$$

and the element $\beta \in \widehat{\mathcal{C}}$ is built from the spaces $\{\overline{\mathcal{N}}(a)\}_{a\in\pi_2(\mathbb{C}^n,L)}$ with geometric dimensions

$$\dim \overline{\mathcal{N}}(a) = n + 1 + \mu(a).$$

Denote by $\widehat{\mathcal{H}}$ the homology with respect to the usual boundary operator of $\widehat{\mathcal{C}}$. Recall that we denote by $\bar{c}$ the image of $c \in \widehat{\mathcal{C}}$ under the degree shift $\widehat{\mathcal{C}} \to \widehat{\mathcal{C}}[-1]$. According to Theorem 4.6, the L_∞ structure on $\widehat{\mathcal{C}}$ pushes forward to an L_∞ structure on $\widehat{\mathcal{H}}$ under a homomorphism e^f, and by Proposition 4.9, this homomorphism maps the elements α, β and $[L]$ in $\widehat{\mathcal{C}}$ to elements α', β' and $[L]$ in $\widehat{\mathcal{H}}$ satisfying the equation

$$\sum_{k=2}^{\infty} \frac{1}{(k-1)!} \ell_k^{\mathcal{H}}(\bar{\beta}', \bar{\alpha}', \cdots, \bar{\alpha}') = \overline{[L]}.$$

Writing

$$\alpha' = \sum_{a\in\pi_2(\mathbb{C}^n,L)} \alpha'(a), \quad \beta' = \sum_{a\in\pi_2(\mathbb{C}^n,L)} \beta'(a),$$

the part of this equation corresponding to the trivial relative homotopy class can be written more explicitly as

$$\sum_{k=2}^{\infty} \frac{1}{(k-1)!} \sum_{a=a_1+\cdots+a_{k-1}} \ell_k^{\mathcal{H}}(\bar{\beta}'(-a), \bar{\alpha}'(a_1), \cdots, \bar{\alpha}'(a_{k-1})) = \overline{[L]}. \tag{5.1}$$

Since the homomorphism between the L_∞ structures preserves degrees, the geometric degrees of $\alpha'(a_i)$ and $\beta'(-a)$ are $n-2+\mu(a_i)$ and $n+1+\mu(-a) = n+1-\mu(a)$, respectively.

By the assumption that L is aspherical, Corollary 5.3 implies that the homology $\widehat{\mathcal{H}}$ is concentrated in geometric degrees $0 \le d \le n$. Combining this observation with (5.1) and the fact that the Maslov index is even for orientable Lagrangian submanifolds L, we find that for the term $\ell_k^{\mathcal{H}}(\bar{\beta}'(-a), \bar{\alpha}'(a_1), \cdots, \bar{\alpha}'(a_{k-1}))$ to be nonzero we must have

$$2 \le \mu(a) \le n+1, \quad \text{and} \quad 2-n \le \mu(a_i) \le 2.$$

The first equation immediately implies that μ is not identically zero. Moreover, if $\mu(a_i) \le 0$ for all $i = 1\dots, k-1$, it follows that $\mu(a) \le 0$, again contradicting the first equation. Thus we conclude that some $\alpha(a_i)$ with $\mu(a_i) = 2$ must be nonzero, implying that the corresponding moduli space is nonempty. So a_i is represented by a holomorphic disk, and hence must have positive symplectic energy.

Set $\gamma := \partial(a_i) \in \pi_1(L)$ and let $C \subset \pi_1(L)$ denote the centralizer of γ. Notice that we have a short exact sequence

$$0 \to \mathrm{Ker}(\mu|_Z) \to C \xrightarrow{\frac{1}{2}\mu} \mathbb{Z} \to 0,$$

in which the last map admits an inverse sending 1 to γ. It follows that the map $\rho : \mathbb{Z} \times \mathrm{Ker}(\mu|_C) \to C$, defined by $\rho(k, g) = \gamma^k \cdot g$, is an isomorphism ($\rho$ is indeed a group homomorphism because γ commutes with all elements of C). Since L is a $K(\pi, 1)$, the covering space $\tilde{L}$ of L with $\pi_1(\tilde{L}) = C$ is a $K(\mathbb{Z} \times \mathrm{Ker}(\mu|_C), 1)$, so it is homotopy equivalent to $S^1 \times L'$ for a $K(\mathrm{Ker}(\mu|_C), 1)$ space L'.

To complete the proof of the theorem, it remains to show that L' is closed or, equivalently, that $\tilde{L} \to L$ is a finite covering space.

Note that the class a_i with $\partial a_i = \gamma$ had the property that $\mu(a_i) = 2$, and moreover $\alpha'(a_i)$ is a nonzero element of geometric degree n in $\widehat{\mathcal{H}}$. So the homology in degree n of $\Lambda_\gamma(L)$ must be nonzero. But combining Lemma 5.1 and Lemma 5.2, we see that $\Lambda_\gamma L$ is homotopy equivalent to the n-manifold $\tilde{L}$. The nonvanishing of its top-dimensional homology now implies that $\tilde{L}$ is closed, which in turns means that $\tilde{L} \to L$ is a finite covering space. □

The more precise statement in dimension 3 can be proven with some specific results from 3-dimensional topology. I wish to thank K. Fukaya, K. Honda and S. Maillot for helpful correspondence, which lead to the following proof of Corollary 1.2.

Proof. (of Corollary 1.2) Let L be a compact, orientable, prime 3-manifold. It is well-known (see e.g. [15, chapter 3]) that either $L \cong S^1 \times S^2$ or L is irreducible, meaning that every embedded two-sphere in L bounds a ball in L.

If an irreducible 3-manifold L admits a Lagrangian embedding into $\mathbb{C}^3$, then by Gromov's Theorem 3.1 it has infinite first homology, and hence infinite fundamental group, and so its universal cover $\tilde{L}$ is non-compact. Moreover, by the sphere theorem (see [15, chapter 4]), an irreducible 3-manifold has trivial second homotopy group. It follows that $H_k(\tilde{L}) = 0$ for $k \geq 1$, and so by Hurewicz's theorem $\pi_k(\tilde{L}) = 0$ for $k \geq 1$, implying that L itself is aspherical. Now by Theorem 1.1, a finite cover of L is homotopy equivalent to $S^1 \times \Sigma$ for some closed oriented surface Σ, and a result of Waldhausen [27, Corollary 6.5] implies that this homotopy equivalence can be improved to a homeomorphism.

Recall from the proof above that the fundamental group C of the cover arises as the centralizer of an element $\gamma \in \pi_1(L)$ with $\mu(\gamma) = 2$. Now I will argue that in fact γ is central in $\pi_1(L)$, so that the covering projection is actually a homeomorphism. Indeed, consider the exact sequence

$$0 \to K \to \pi_1(L) \xrightarrow{\frac{1}{2}\mu} \mathbb{Z} \to 0,$$

where $K = \ker \mu$. Since the centralizer C of γ is of finite index, $K' = C \cap K$ is of finite index in K. From the above proof of the theorem, we see that K' is finitely generated (it is the fundamental group of Σ), so K is also finitely generated. Then by Stallings' fibration theorem ([25], see also [15, Theorem 11.6]), we deduce that K is the fundamental group of a compact surface S, which under our current assumptions must be closed of genus at least 1.

Now the proof concludes with the following observation:

Lemma 5.4. *Any automorphism φ of the fundamental group K of a closed oriented surface which is trivial on some finite index subgroup K' is trivial.*

Proof. If the surface is a sphere there is nothing to prove, so we consider the case that the genus of the surface is at least 1. One knows that the fundamental group K of a closed surface has no torsion. Let $g \in K$ be given and consider the infinite cyclic subgroup $Z = \langle g \rangle$ generated by g. The subgroup $\varphi(Z) \cap Z$ contains the finite index subgroup $K' \cap Z$, on which φ acts trivially. But any automorphism of an infinite cyclic group fixing some nontrivial subgroup must be the identity, so $\varphi(g) = g$. Since this applies to any $g \in K$, the lemma is proven. □

Applying the lemma to the action of γ on K by conjugation, which clearly fixes all elements of $K' = K \cap C$, we finally conclude that γ is central in $\pi_1(L)$, so that $C = \pi_1(L)$, which finishes the proof of the theorem. □

6 Reflections

Above, I have presented Fukaya's elegant arguments leading to some substantial new results about Lagrangian submanifolds in $\mathbb{C}^n$. At first glance, it seems that string topology is really essential to the approach. However, on further inspection, one discovers that there may be a way to avoid it almost entirely.

The basic idea is the following. By Viterbo's theorem (see [26] and chapter chap:Viterbo), the homology of the free loop space can be described in symplectic terms as the symplectic homology of the cotangent bundle. This theory can be defined in more general situations, for example for exact symplectic manifolds with contact-type boundary. Moreover, for *exact* codimension 0 embeddings $U \hookrightarrow W$ one has restriction maps $SH_*(W) \to SH_*(U)$. It seems reasonable to expect (and is the subject of current work) that every algebraic structure that exists on the homology of the free loop space can also be defined on symplectic homology in general, even if the underlying domain is not a cotangent bundle. In the exact case, the restriction homomorphism should respect all these structures.

But even more should be true. In the case of a non-exact embedding $U \hookrightarrow W$, there will be a Maurer-Cartan element in $SH_*(U)$ such that after twisting all the structures by this Maurer-Cartan element we get a morphism from $SH_*(W)$ to the twisted version $SH_*^{\mathrm{twisted}}(U)$. This expectation is consistent with (and gives one of several possible conceptual explanations for) the results of Fukaya for Lagrangians in $\mathbb{C}^n$.

Indeed, with U being a small neighborhood of the zero section in the cotangent bundle of L and $W = \mathbb{C}^n$, we are exactly in the situation just described. What I have argued in earlier sections is that, after twisting by a Maurer-Cartan element coming from the embedding, the unit $[L] \in H_*(\Lambda L)$ with respect to the loop product has become exact, which by a standard argument will force the twisted homology to vanish completely. This is good news, because only in this case can we even expect to have a morphism from $SH_*(\mathbb{C}^n) = 0$ to this ring. The prediction is that this morphism can indeed be defined in a suitable chain version of the theory.

Once the above argument has been made to work, it extends the applicability of Fukaya's approach in several directions. Notice that string topology only enters indirectly, via Viterbo's isomorphism. As long as the algebraic operations can be defined and the morphism associated to a codimension zero embedding respects them, one does not even need to know that the operations on symplectic homology are the same as those in string topology (although this is of course expected to be true). Moreover, one can study non-exact codimension 0 embeddings of general exact symplectic manifolds with contact boundary by this method.

7 Guide to the literature

The basic source for this chapter are of course Fukaya's papers [10, 11]. Versions of Theorem 1.1 under additional assumptions, like monotonicity of the Lagrangian submanifold, are much easier to achieve, see e.g. [5, 6, 8] and the references therein.

For an introduction to symplectic topology the book [21] is recommended. It covers a lot more than is necessary to understand the problem discussed here, and it gives some hints why Lagrangian submanifolds are so central in symplectic topology. To learn something more specific about Lagrangian embeddings and immersions, the excellent survey [4] is still the best place to start. Recently, new results have appeared which suggest that the problem in higher dimensions is more flexible than previously expected [7].

A chain complex $\mathcal{C}$ for the free loop space on which the loop bracket is fully defined, and which therefore might serve in the implementation of Theorem 3.3, has recently been proposed by Irie [18].

Finally, the reader who really wants to appreciate the discussion in this chapter needs to know quite a bit about holomorphic curves. One good source which thoroughly covers a lot of the basics, including a version of Gromov compactness and a complete proof of Gromov's Theorem 3.1, is [22]. Many aspects of the theory are also covered in the earlier book [3]. With these as a guide, the monumental [12] will hopefully look less daunting.

References

[1] V. I. Arnold and A. B. Givental, *Symplectic geometry*, Dynamical systems, IV, 1–138, Encyclopaedia Math. Sci., 4, Springer, 1990

[2] M. Audin, *Fibrés normaux d'immersions en dimension double, points doubles dimmersions lagragiennes et plongements totalement réels*, Comment. Math. Helv. 63 (1988), no. 4, 593–623.

[3] M. Audin, J. Lafontaine (eds.), *Holomorphic cuves in symplectic geometry*, Birkhäuser, Progress in Mathematics vol. 117, 1994

[4] M. Audin, F. Lalonde and L. Polterovich, *Symplectic rigidity: Lagrangian submanifolds*, in [3]

[5] L. Buhovsky, *The Maslov class of Lagrangian tori and quantum products in Floer cohomology*, J. Topol. and Anal. 2 (2010),no. 1, 57–75.

[6] M. Damian, *Floer homology on the universal cover, Audins conjecture and other constraints on Lagrangian submanifolds*, Comment. Math. Helv. 87 (2012), no. 2, 433–463

[7] T. Ekholm, Y. Eliashberg, E. Murphy and I. Smith, *Constructing exact Lagrangian immersions with few double points*, preprint, arXiv:1303.0588

[8] J. D. Evans, J. Kedra, *Remarks on monotone Lagrangians in* $\mathbb{C}^n$, preprint, arXiv:1110.0927

[9] U. Frauenfelder, *Gromov convergence of pseudoholomorphic disks*, J. Fixed Point Theory Appl. 3 (2008), no. 2, 215–271.

[10] K. Fukaya, *Application of Floer homology of Langrangian submanifolds to symplectic topology*, in: Paul Biran et al. (eds.), Morse Theoretic Methods in Nonlinear Analysis and in Symplectic Topology, Springer, 2006, 231-276.

[11] K. Fukaya, *Loop space and holomorphic disc -summary-*, preprint, 2007, downloaded from: `https://www.math.kyoto-u.ac.jp/~fukaya/fukaya.html`

[12] K. Fukaya, Y.-G. Oh, H. Ohta and K. Ono, *Lagrangian intersection Floer theory: anomaly and obstruction*, AMS/IP Studies in Advanced Mathematics, 2009

[13] A. Givental, *Lagrange imbeddings of surfaces and the open Whitney umbrella*, Funktsional. Anal. i Prilozhen. 20 (1986), no. 3, 35–41, 96

[14] M. Gromov, *Pseudoholomorphic curves in symplectic geometry*, Invent. Math. 82 (1985), no. 2, 307–347.

[15] J. Hempel, *3-manifolds*, Annals of Mathematics Studies 86, Princeton University Press, 1976

[16] H. Hofer, *A general Fredholm theory and applications*, in: Current developments in mathematics, 2004, 1–71, Int. Press, Somerville, MA, 2006

[17] C. Hummel, *Gromov's compactness theorem for pseudoholomorphic curves*, Birkhäuser Verlag, Progress in Mathematics 151, Basel, 1997

[18] K. Irie, *Transversality problems in string topology and de Rham chains*, preprint, arXiv:1404.0153

[19] C.-C. M. Liu, *Moduli of J-holomorphic curves with Lagrangian boundary conditions and open Gromov-Witten invariants for an S^1-equivariant pair*, Ph.D. Thesis, Harvard 2002, arXiv:math.SG/0210257

[20] J.-L. Loday and B. Vallette, *Algebraic operads*, Grundlehren der mathematischen Wissenschaften 346, Springer-Verlag Berlin, 2012. xxiv+634 pp.

[21] D. McDuff and D. Salamon, *Introduction to symplectic topology*, second edition, Oxford University Press, 1998

[22] D. McDuff and D. Salamon, *J-holomorphic curves and symplectic topology* , American Mathematical Society Colloquium Publications, 52. American Mathematical Society, Providence, RI, 2004. xii+669 pp. ISBN: 0-8218-3485-1

[23] S. Yu. Nemirovski, *Homology class of a Lagrangian Klein bottle*, Isvestiya Math. 73, 2009, no. 4, 689–698.

[24] V. V. Shevchishin, *Lagrangian embeddings of the Klein bottle and combinatorial properties of mapping class groups*, Isvestiya Math. 73, 2009, no. 4, 797–859.

[25] J. Stallings, *On Fibering Certain 3-Manifolds*, in: Topology of 3-Manifolds and Related Topics, M. K. Fort, Jr. (ed.), Prentice Hall, 1962

[26] C. Viterbo, *Functors and computations in Floer homology with applications I*, Geom. Funct. Anal. 9, 1999, no. 5, 985–1033.

[27] F. Waldhausen, *On irreducible 3-manifolds which are sufficiently large*, Ann. of Math. (2), 87, 1968, 56–88.

[28] A.Weinstein, *Symplectic geometry*, Bull. Amer. Math. Soc. (N.S.) 5 (1981), no. 1, 1–13.