

A BEGINNER'S INTRODUCTION TO FUKAYA CATEGORIES

DENIS AUROUX

This text is based on a series of lectures given at a Summer School on Contact and Symplectic Topology at Université de Nantes in June 2011.

The goal of these notes is to give a short introduction to Fukaya categories and some of their applications. The first half of the text is devoted to a brief review of Lagrangian Floer (co)homology and product structures. Then we introduce the Fukaya category (informally and without a lot of the necessary technical detail), and briefly discuss algebraic concepts such as exact triangles and generators. Finally, we mention wrapped Fukaya categories and outline a few applications to symplectic topology, mirror symmetry and low-dimensional topology.

These notes are in no way a comprehensive text on the subject; however we hope that they will provide a useful introduction to Paul Seidel's book [40] and other texts on Floer homology, Fukaya categories, and their applications. We assume that the reader is generally familiar with the basics of symplectic geometry, and some prior exposure to pseudo-holomorphic curves is also helpful; the reader is referred to [28, 29] for background material.

Acknowledgements. The author wishes to thank the organizers of the Nantes Trimester on Contact and Symplectic Topology for the pleasant atmosphere at the Summer School, and Ailsa Keating for providing a copy of the excellent notes she took during the lectures. Much of the material presented here I initially learned from Paul Seidel and Mohammed Abouzaid, whom I thank for their superbly written papers and their patient explanations. Finally, the author was partially supported by an NSF grant (DMS-1007177).

1. LAGRANGIAN FLOER (CO)HOMOLOGY

1.1. Motivation. Lagrangian Floer homology was introduced by Floer in the late 1980s in order to study the intersection properties of compact Lagrangian submanifolds in symplectic manifolds and prove an important case of Arnold's conjecture concerning intersections between Hamiltonian isotopic Lagrangian submanifolds [12].

Specifically, let (M, ω) be a symplectic manifold (compact, or satisfying a "bounded geometry" assumption), and let L be a compact Lagrangian submanifold of M . Let $\psi \in \text{Ham}(M, \omega)$ be a Hamiltonian diffeomorphism. (Recall that a time-dependent

The author was partially supported by NSF grant DMS-1007177.

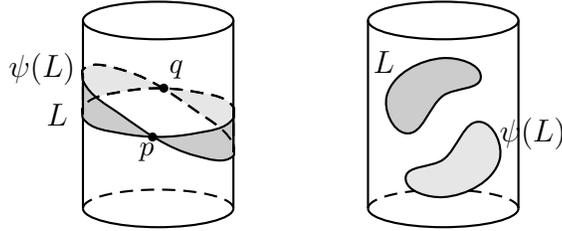


FIGURE 1. Arnold's conjecture on the cylinder $\mathbb{R} \times S^1$: an example (left) and a non-example (right)

Hamiltonian $H \in C^\infty(M \times [0, 1], \mathbb{R})$ determines a family of Hamiltonian vector fields X_t via the equation $\omega(\cdot, X_t) = dH_t$, where $H_t = H(\cdot, t)$; integrating these vector fields over $t \in [0, 1]$ yields the Hamiltonian diffeomorphism ψ generated by H .)

Theorem 1.1 (Floer [17]). *Assume that the symplectic area of any topological disc in M with boundary in L vanishes. Assume moreover that $\psi(L)$ and L intersect transversely. Then the number of intersection points of L and $\psi(L)$ satisfies the lower bound $|\psi(L) \cap L| \geq \sum_i \dim H^i(L; \mathbb{Z}_2)$.*

Note that, by Stokes' theorem, since $\omega|_L = 0$, the symplectic area of a disc with boundary on L only depends on its class in the relative homotopy group $\pi_2(M, L)$.

The bound given by Theorem 1.1 is stronger than what one could expect from purely topological considerations. The assumptions that the diffeomorphism ψ is Hamiltonian, and that L does not bound discs of positive symplectic area, are both essential (though the latter can be slightly relaxed in various ways).

Example 1.2. Consider the cylinder $M = \mathbb{R} \times S^1$, with the standard area form, and a simple closed curve L that goes around the cylinder once: then $\psi(L)$ is also a simple closed curve going around the cylinder once, and the assumption that $\psi \in \text{Ham}(M)$ means that the total signed area of the 2-chain bounded by L and $\psi(L)$ is zero. It is then an elementary fact that $|\psi(L) \cap L| \geq 2$, as claimed by Theorem 1.1; see Figure 1 left. On the other hand, the result becomes false if we only assume that ψ is a symplectomorphism (a large vertical translation of the cylinder is area-preserving and eventually displaces L away from itself); or if we take L to be a homotopically trivial simple closed curve, which bounds a disc of positive area (see Figure 1 right).

Floer's approach is to associate to the pair of Lagrangians $(L_0, L_1) = (L, \psi(L))$ a chain complex $CF(L_0, L_1)$, freely generated by the intersection points of L_0 and L_1 , equipped with a differential $\partial : CF(L_0, L_1) \rightarrow CF(L_0, L_1)$, with the following properties:

- (1) $\partial^2 = 0$, so the Floer cohomology $HF(L_0, L_1) = \text{Ker } \partial / \text{Im } \partial$ is well-defined;
- (2) if L_1 and L'_1 are Hamiltonian isotopic then $HF(L_0, L_1) \simeq HF(L_0, L'_1)$;

- (3) if L_1 is Hamiltonian isotopic to L_0 , then $HF(L_0, L_1) \simeq H^*(L_0)$ (with suitable coefficients).

Theorem 1.1 then follows immediately, since the rank of $HF(L, \psi(L)) \simeq H^*(L)$ is bounded by that of the Floer complex $CF(L, \psi(L))$, which equals $|\psi(L) \cap L|$.

Formally, Lagrangian Floer (co)homology can be viewed as an infinite-dimensional analogue of Morse (co)homology for the *action functional* on (the universal cover of) the path space $\mathcal{P}(L_0, L_1) = \{\gamma : [0, 1] \rightarrow M \mid \gamma(0) \in L_0, \gamma(1) \in L_1\}$,

$$\mathcal{A}(\gamma, [\Gamma]) = - \int_{\Gamma} \omega,$$

where $(\gamma, [\Gamma]) \in \tilde{\mathcal{P}}(L_0, L_1)$ consists of a path $\gamma \in \mathcal{P}(L_0, L_1)$ and an equivalence class $[\Gamma]$ of a homotopy $\Gamma : [0, 1] \times [0, 1] \rightarrow M$ between γ and a fixed base point in the connected component of $\mathcal{P}(L_0, L_1)$ containing γ . The critical points of \mathcal{A} are (lifts of) constant paths at intersection points, and its gradient flow lines (with respect to the natural L^2 -metric induced by ω and a compatible almost-complex structure) are pseudo-holomorphic strips bounded by L_0 and L_1 .

However, the analytic difficulties posed by Morse theory in the infinite-dimensional setting are such that the actual definition of Floer (co)homology does not rely on this interpretation: instead, the Floer differential is defined in terms of moduli spaces of pseudo-holomorphic strips.

1.2. The Floer differential. Let L_0, L_1 be compact Lagrangian submanifolds of a symplectic manifold (M, ω) , and assume for now that L_0 and L_1 intersect transversely, hence at a finite set of points.

Before we introduce the Floer complex and the Floer differential, a brief discussion of coefficients is in order. In general, Floer cohomology is defined with *Novikov coefficients* (over some base field \mathbb{K} , for example $\mathbb{K} = \mathbb{Q}$, or $\mathbb{K} = \mathbb{Z}_2$).

Definition 1.3. *The Novikov ring over a base field \mathbb{K} is*

$$\Lambda_0 = \left\{ \sum_{i=0}^{\infty} a_i T^{\lambda_i} \mid a_i \in \mathbb{K}, \lambda_i \in \mathbb{R}_{\geq 0}, \lim_{i \rightarrow \infty} \lambda_i = +\infty \right\}.$$

The Novikov field Λ is the field of fractions of Λ_0 , i.e.

$$\Lambda = \left\{ \sum_{i=0}^{\infty} a_i T^{\lambda_i} \mid a_i \in \mathbb{K}, \lambda_i \in \mathbb{R}, \lim_{i \rightarrow \infty} \lambda_i = +\infty \right\}.$$

The Floer complex is then the free Λ -module generated by intersection points: we denote by $\mathcal{X}(L_0, L_1) = L_0 \cap L_1$ the set of generators, and set

$$CF(L_0, L_1) = \bigoplus_{p \in \mathcal{X}(L_0, L_1)} \Lambda \cdot p.$$

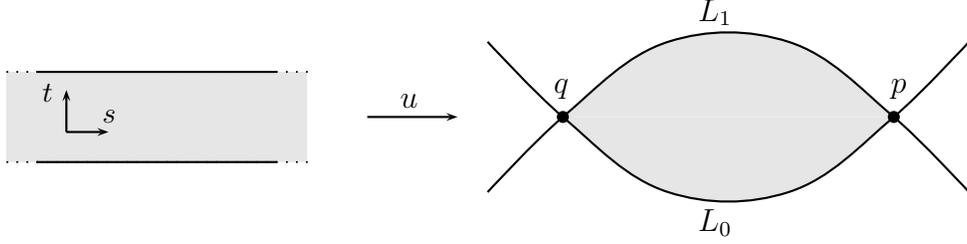


FIGURE 2. A pseudo-holomorphic strip contributing to the Floer differential on $CF(L_0, L_1)$

Equip M with an ω -compatible almost-complex structure J . (By a classical result, the space of ω -compatible almost-complex structures $\mathcal{J}(M, \omega) = \{J \in \text{End}(TM) \mid J^2 = -1 \text{ and } g_J = \omega(\cdot, J\cdot) \text{ is a Riemannian metric}\}$ is non-empty and contractible [28].)

The Floer differential $\partial : CF(L_0, L_1) \rightarrow CF(L_0, L_1)$ is defined by counting pseudo-holomorphic strips in M with boundary in L_0 and L_1 : namely, given intersection points $p, q \in \mathcal{X}(L_0, L_1)$, the coefficient of q in ∂p is obtained by considering the space of maps $u : \mathbb{R} \times [0, 1] \rightarrow M$ which solve the Cauchy-Riemann equation $\bar{\partial}_J u = 0$, i.e.

$$(1.1) \quad \frac{\partial u}{\partial s} + J(u) \frac{\partial u}{\partial t} = 0,$$

subject to the boundary conditions

$$(1.2) \quad \begin{cases} u(s, 0) \in L_0 \text{ and } u(s, 1) \in L_1 & \forall s \in \mathbb{R}, \\ \lim_{s \rightarrow +\infty} u(s, t) = p, & \lim_{s \rightarrow -\infty} u(s, t) = q, \end{cases}$$

and the *finite energy* condition

$$(1.3) \quad E(u) = \int u^* \omega = \iint \left| \frac{\partial u}{\partial s} \right|^2 ds dt < \infty.$$

(Note that, by the Riemann mapping theorem, the strip $\mathbb{R} \times [0, 1]$ is biholomorphic to $D^2 \setminus \{\pm 1\}$, the closed unit disc minus two points on its boundary; the map u then extends to the closed disc, with the boundary marked points ± 1 mapping to p and q .)

Given a homotopy class $[u] \in \pi_2(M, L_0 \cup L_1)$, we denote by $\widehat{\mathcal{M}}(p, q; [u], J)$ the space of solutions of (1.1)–(1.3) representing the class $[u]$, and by $\mathcal{M}(p, q; [u], J)$ its quotient by the action of \mathbb{R} by reparametrization (i.e., $a \in \mathbb{R}$ acts by $u \mapsto u_a(s, t) := u(s-a, t)$).

The boundary value problem (1.1)–(1.3) is a Fredholm problem, i.e. the linearization $D_{\bar{\partial}_J, u}$ of $\bar{\partial}_J$ at a given solution u is a Fredholm operator. Specifically, $D_{\bar{\partial}_J, u}$ is a $\bar{\partial}$ -type first-order differential operator, whose domain is a suitable space of sections of the pullback bundle u^*TM (with Lagrangian boundary conditions), for example $W^{1,p}(\mathbb{R} \times [0, 1], \mathbb{R} \times \{0, 1\}; u^*TM, u^*_{|_{t=0}}TL_0, u^*_{|_{t=1}}TL_1)$. The Fredholm index $\text{ind}([u]) := \text{ind}_{\mathbb{R}}(D_{\bar{\partial}_J, u}) = \dim \text{Ker } D_{\bar{\partial}_J, u} - \dim \text{Coker } D_{\bar{\partial}_J, u}$ can be computed in terms of an invariant of the class $[u]$ called the *Maslov index*, which we discuss below.

The space of solutions $\widehat{\mathcal{M}}(p, q; [u], J)$ is then a smooth manifold of dimension $\text{ind}([u])$, provided that all solutions to (1.1)–(1.3) are *regular*, i.e. the linearized operator $D_{\bar{\partial}_{J,u}}$ is surjective at each point of $\widehat{\mathcal{M}}(p, q; [u], J)$. This transversality property is one of three fundamental technical issues that need to be addressed for Floer (co)homology to be defined, the other two being the compactness of the moduli space $\mathcal{M}(p, q; [u], J)$, and its orientability (unless one is content to work over $\mathbb{K} = \mathbb{Z}_2$).

Transversality and compactness will be briefly discussed in §1.4 below. On the issue of orientations, we will only consider the case where L_0 and L_1 are oriented and spin. It is then known that the choice of spin structures on L_0 and L_1 determines a canonical orientation of the moduli spaces of J -holomorphic strips; the construction of this orientation is fairly technical, so we refer the reader to [19, 40] for details.

Assuming that all these issues have been taken care of, we observe that when $\text{ind}([u]) = 1$ the moduli space $\mathcal{M}(p, q; [u], J)$ is a compact oriented 0-manifold, i.e. a finite set of points which can be counted with signs. We can then provisionally define:

Definition 1.4. *The Floer differential $\partial : CF(L_0, L_1) \rightarrow CF(L_0, L_1)$ is the Λ -linear map defined by*

$$(1.4) \quad \partial(p) = \sum_{\substack{q \in \mathcal{X}(L_0, L_1) \\ [u]: \text{ind}([u])=1}} (\#\mathcal{M}(p, q; [u], J)) T^{\omega([u])} q,$$

where $\#\mathcal{M}(p, q; [u], J) \in \mathbb{Z}$ (or \mathbb{Z}_2) is the signed (or unsigned) count of points in the moduli space of pseudo-holomorphic strips connecting p to q in the class $[u]$, and $\omega([u]) = \int u^* \omega$ is the symplectic area of those strips.

In general, the definition needs to be modified by introducing a perturbation term into the Cauchy-Riemann equation in order to achieve transversality (see §1.4 below). Thus, the Floer differential actually counts *perturbed* pseudo-holomorphic strips connecting *perturbed* intersection points of L_0 and L_1 .

The following result is due to Floer for $\mathbb{K} = \mathbb{Z}_2$:

Theorem 1.5. *Assume that $[\omega] \cdot \pi_2(M, L_0) = 0$ and $[\omega] \cdot \pi_2(M, L_1) = 0$. Moreover, when $\text{char}(\mathbb{K}) \neq 2$ assume that L_0, L_1 are oriented and equipped with spin structures. Then the Floer differential ∂ is well-defined, satisfies $\partial^2 = 0$, and the Floer cohomology $HF(L_0, L_1) = H^*(CF(L_0, L_1), \partial)$ is, up to isomorphism, independent of the chosen almost-complex structure J and invariant under Hamiltonian isotopies of L_0 or L_1 .*

Remark 1.6. In this text we discuss the chain complex and differential for Floer cohomology, which is dual to Floer's original construction. Namely, in Floer homology, the strip shown on Figure 2 would be considered a trajectory from q to p rather than from p to q , and the grading conventions are reversed.

Remark 1.7. In general, the sum in the right-hand side of (1.4) can be infinite. However, Gromov's compactness theorem ensures that, given any energy bound E_0 ,

there are only finitely many homotopy classes $[u]$ with $\omega([u]) \leq E_0$ for which the moduli spaces $\mathcal{M}(p, q; [u], J)$ are non-empty. Thus, using Novikov coefficients and weighing counts of strips by area ensures that the sum in the right-hand side of (1.4) is well-defined.

However, it is sometimes possible to work over smaller coefficient fields. One such setting is that of *exact* Lagrangian submanifolds in an exact symplectic manifold. Namely, assume that $\omega = d\theta$ for some 1-form θ on M , and there exist functions $f_i \in C^\infty(L_i, \mathbb{R})$ such that $\theta|_{L_i} = df_i$ (for $i = 0, 1$). Then, by Stokes' theorem, any strip connecting intersection points p and q satisfies $\int u^*\omega = (f_1(q) - f_0(q)) - (f_1(p) - f_0(p))$. Thus, rescaling each generator by $p \mapsto T^{f_1(p) - f_0(p)}p$, we can eliminate the weights $T^{\omega([u])}$ from (1.4), and work directly over the coefficient field \mathbb{K} instead of Λ .

Floer's construction [17] was subsequently extended to more general settings, beginning with Oh's result on monotone Lagrangians [32], and culminating with the sophisticated methods introduced by Fukaya, Oh, Ohta and Ono for the general case [19]; however as we will see below, Theorem 1.5 does not hold in full generality, as pseudo-holomorphic discs with boundary in L_0 or L_1 "obstruct" Floer cohomology.

1.3. Maslov index and grading. The Maslov index plays a similar role in the index formula for pseudo-holomorphic discs to that played by the first Chern class in that for closed pseudo-holomorphic curves; in fact it can be viewed as a relative version of the Chern class.

Denote by $LGr(n)$ the Grassmannian of Lagrangian n -planes in the symplectic vector space $(\mathbb{R}^{2n}, \omega_0)$. It is a classical fact that the unitary group $U(n)$ acts transitively on $LGr(n)$, so that $LGr(n) \simeq U(n)/O(n)$, from which it follows by an easy calculation that $\pi_1(LGr(n)) \simeq \mathbb{Z}$ (see e.g. [28]). This can be understood concretely by using the square of the determinant map, $\det^2 : U(n)/O(n) \rightarrow S^1$, which induces an isomorphism on fundamental groups; the Maslov index of a loop in $LGr(n)$ is simply the winding number of its image under this map.

In a similar vein, consider two paths $\ell_0, \ell_1 : [0, 1] \rightarrow LGr(n)$ of Lagrangian subspaces in \mathbb{R}^{2n} , such that $\ell_0(0)$ is transverse to $\ell_1(0)$ and $\ell_0(1)$ is transverse to $\ell_1(1)$. The Maslov index of the path ℓ_1 relative to ℓ_0 is then the number of times (counting with signs and multiplicities) at which $\ell_0(t)$ and $\ell_1(t)$ are not transverse to each other. (More precisely, it is the intersection number of the path $(\ell_0(t), \ell_1(t))$ with the hypersurface in $LGr(n) \times LGr(n)$ consisting of non-transverse pairs of subspaces.)

We now return to our main discussion, and consider a map $u : \mathbb{R} \times [0, 1] \rightarrow M$ satisfying the boundary conditions (1.2). Since $\mathbb{R} \times [0, 1]$ is contractible, the pullback u^*TM is a trivial symplectic vector bundle; fixing a trivialization, we can view $\ell_0 = u^*_{|\mathbb{R} \times \{0\}} TL_0$ and $\ell_1 = u^*_{|\mathbb{R} \times \{1\}} TL_1$ as paths (oriented with s going from $+\infty$ to $-\infty$) in $LGr(n)$, one connecting $T_p L_0$ to $T_q L_0$ and the other connecting $T_p L_1$ to $T_q L_1$. The index of u can then be defined as the Maslov index of the path ℓ_1 relative to ℓ_0 .

An equivalent definition, which generalizes more readily to the discs that appear in the definition of product operations, is as follows. Given a pair of transverse subspaces $\lambda_0, \lambda_1 \in LGr(n)$, and identifying \mathbb{R}^{2n} with \mathbb{C}^n , there exists an element $A \in Sp(2n, \mathbb{R})$ which maps λ_0 to $\mathbb{R}^n \subset \mathbb{C}^n$ and λ_1 to $(i\mathbb{R})^n \subset \mathbb{C}^n$. The subspaces $\lambda_t = A^{-1}((e^{-i\pi t/2}\mathbb{R})^n)$, $t \in [0, 1]$ then provide a distinguished homotopy class of path connecting λ_0 to λ_1 in $LGr(n)$, which we call the *canonical short path*.

Definition 1.8. Given $p, q \in L_0 \cap L_1$, denote by λ_p the canonical short path from $T_p L_0$ to $T_p L_1$ and by λ_q that from $T_q L_0$ to $T_q L_1$. Given a strip $u : \mathbb{R} \times [0, 1] \rightarrow M$ connecting p to q , for $i \in \{0, 1\}$, denote by ℓ_i the path $u|_{\mathbb{R} \times \{i\}}^* T L_i$ oriented with s going from $+\infty$ to $-\infty$, from $T_p L_i$ to $T_q L_i$. View all these as paths in $LGr(n)$ by fixing a trivialization of $u^* TM$. The index of the strip u is then the Maslov index of the closed loop in $LGr(n)$ (based at $T_q L_0$) obtained by concatenating the paths $-\ell_0$ (i.e. ℓ_0 backwards), λ_p , ℓ_1 , and finally $-\lambda_q$.

Example 1.9. Let $M = \mathbb{R}^2$, and consider the strip u depicted in Figure 2: then it is an easy exercise to check, using either definition, that $\text{ind}(u) = 1$.

We now discuss the related issue of equipping Floer complexes with a grading. In order to obtain a \mathbb{Z} -grading on $CF(L_0, L_1)$, one needs to make sure that the index of a strip depends only on the difference between the degrees of the two generators it connects, rather than on its homotopy class. This is ensured by the following two requirements:

- (1) The first Chern class of M must be 2-torsion: $2c_1(TM) = 0$. This allows one to lift the Grassmannian $LGr(TM)$ of Lagrangian planes in TM (an $LGr(n)$ -bundle over M) to a fiberwise universal cover $\widetilde{LGr}(TM)$, the Grassmannian of *graded Lagrangian planes* in TM (an $\widetilde{LGr}(n)$ -bundle over M).

Concretely, given a nowhere vanishing section Θ of $(\Lambda_{\mathbb{C}}^n T^* M)^{\otimes 2}$, the argument of Θ associates to any Lagrangian plane ℓ a *phase* $\varphi(\ell) = \arg(\Theta|_{\ell}) \in S^1 = \mathbb{R}/2\pi\mathbb{Z}$; a graded lift of ℓ is the choice of a real lift of $\tilde{\varphi}(\ell) \in \mathbb{R}$ of $\varphi(\ell)$.

- (2) The *Maslov class* of L , $\mu_L \in \text{Hom}(\pi_1(L), \mathbb{Z}) = H^1(L, \mathbb{Z})$, vanishes. The Maslov class is by definition the obstruction to consistently choosing graded lifts of the tangent planes to L , i.e. lifting the section of $\widetilde{LGr}(TM)$ over L given by $p \mapsto T_p L$ to a section of the infinite cyclic cover $\widetilde{LGr}(TM)$. The Lagrangian submanifold L together with the choice of such a lift is called a *graded Lagrangian submanifold* of M .

Equivalently, given a nowhere vanishing section of $(\Lambda_{\mathbb{C}}^n T^* M)^{\otimes 2}$, we can associate to L its *phase function* $\varphi_L : L \rightarrow S^1$, which maps $p \in L$ to $\varphi(T_p L) \in S^1$; the Maslov class is then the homotopy class $[\varphi_L] \in [L, S^1] = H^1(L, \mathbb{Z})$, and a graded lift of L is the choice of a lift $\tilde{\varphi}_L : L \rightarrow \mathbb{R}$.

When these two assumptions are satisfied, fixing graded lifts \tilde{L}_0, \tilde{L}_1 of the Lagrangian submanifolds $L_0, L_1 \subset M$ determines a natural \mathbb{Z} -grading on the Floer

complex $CF(L_0, L_1)$ as follows. For all $p \in L_0 \cap L_1$, we obtain a preferred homotopy class of path connecting $T_p L_0$ to $T_p L_1$ in $LGr(T_p M)$ by connecting the chosen graded lifts of the tangent spaces at p via a path in $\widetilde{LGr}(T_p M)$. Combining this path with $-\lambda_p$ (the canonical short path from $T_p L_0$ to $T_p L_1$, backwards), we obtain a closed loop in $LGr(T_p M)$; the degree of p is by definition the Maslov index of this loop. It is then easy to check that any strip connecting p to q satisfies

$$(1.5) \quad \text{ind}(u) = \text{deg}(q) - \text{deg}(p).$$

In particular the Floer differential (1.4) has degree 1.

In general, if we do not restrict ourselves to symplectic manifolds with torsion $c_1(TM)$ and Lagrangian submanifolds with vanishing Maslov class, the natural grading on Floer cohomology is only by a finite cyclic group. As an important special case, if we simply assume that the Lagrangian submanifolds L_0, L_1 are oriented, then we have a $\mathbb{Z}/2$ -grading, where the degree of a generator p of $CF(L_0, L_1)$ is determined by the sign of the intersection between L_0 and L_1 at p : namely $\text{deg}(p) = 0$ if the canonical short path from $T_p L_0$ to $T_p L_1$ maps the given orientation of $T_p L_0$ to that of $T_p L_1$, and $\text{deg}(p) = 1$ otherwise.

(Another approach, which we won't discuss further, is to enlarge the coefficient field by a formal variable of non-zero degree to keep track of the Maslov indices of different homotopy classes. In the monotone case, where index is proportional to symplectic area, it suffices to give a non-zero degree to the Novikov parameter T .)

1.4. Transversality and compactness. We now discuss very briefly the fundamental technical issues of transversality and compactness.

Transversality of the moduli spaces of pseudo-holomorphic strips, i.e. the surjectivity of the linearized $\bar{\partial}$ operator at all solutions, is needed in order to ensure that the spaces $\widehat{\mathcal{M}}(p, q; [u], J)$ (and other moduli spaces we will introduce below) are smooth manifolds of the expected dimension. Still assuming that L_0 and L_1 intersect transversely, transversality for strips can be achieved by replacing the fixed almost-complex structure J in the Cauchy-Riemann equation (1.1) by a generic *family* of ω -compatible almost-complex structures which depend on the coordinate t in the strip $\mathbb{R} \times [0, 1]$.

A more basic issue is that of defining Floer cohomology for Lagrangian submanifolds which do not intersect transversely (in particular, one would like to be able to define the Floer cohomology of a Lagrangian with itself, i.e. the case $L_0 = L_1$). In view of the requirement of Hamiltonian isotopy invariance of the construction, the simplest approach is to introduce an inhomogeneous Hamiltonian perturbation term into the holomorphic curve equation: we fix a generic Hamiltonian $H \in C^\infty([0, 1] \times M, \mathbb{R})$, and consider the modified equation $(du - X_H \otimes dt)^{0,1} = 0$, i.e.

$$(1.6) \quad \frac{\partial u}{\partial s} + J(t, u) \left(\frac{\partial u}{\partial t} - X_H(t, u) \right) = 0,$$

still subject to the boundary conditions $u(s, 0) \in L_0$ and $u(s, 1) \in L_1$ and a finite energy condition. For $s \rightarrow \pm\infty$, the strip u converges no longer to intersection points but rather to trajectories of the flow of X_H which start on L_0 and end on L_1 : thus the generators of the Floer complex $CF(L_0, L_1)$ are in fact defined to be flow lines $\gamma : [0, 1] \rightarrow M$, $\dot{\gamma}(t) = X_H(t, \gamma(t))$, such that $\gamma(0) \in L_0$ and $\gamma(1) \in L_1$. Equivalently, by considering $\gamma(0)$, we set $\mathcal{X}(L_0, L_1) = L_0 \cap (\phi_H^1)^{-1}(L_1)$, where $\phi_H^1 \in \text{Ham}(M, \omega)$ is the time 1 flow generated by H . In this sense, the generators are *perturbed* intersection points of L_0 with L_1 , where the perturbation is given by the Hamiltonian diffeomorphism ϕ_H^1 .

Remark 1.10. The perturbed equation (1.6) can be recast as a plain Cauchy-Riemann equation by the following trick: consider $\tilde{u}(s, t) = (\phi_H^t)^{-1}(u(s, t))$, where ϕ_H^t is the flow of X_H over the interval $[0, t]$. Then

$$\frac{\partial \tilde{u}}{\partial t} = (\phi_H^t)_*^{-1} \left(\frac{\partial u}{\partial t} - X_H \right),$$

so Floer's equation (1.6) becomes

$$\frac{\partial \tilde{u}}{\partial s} + \tilde{J}(t, \tilde{u}) \frac{\partial \tilde{u}}{\partial t} = 0,$$

where $\tilde{J}(t) = (\phi_H^t)_*^{-1}(J(t))$. Hence solutions to Floer's equation correspond to honest \tilde{J} -holomorphic strips with boundaries on L_0 and $(\phi_H^1)^{-1}(L_1)$.

Compactness of the moduli spaces is governed by Gromov's compactness theorem, according to which any sequence of J -holomorphic curves with uniformly bounded energy admits a subsequence which converges, up to reparametrization, to a nodal *tree* of J -holomorphic curves. The components of the limit curve are obtained as limits of different reparametrizations of the given sequence of curves, focusing on the different regions of the domain in which a non-zero amount of energy concentrates ("bubbling"). In the case of a sequence of J -holomorphic strips $u_n : \mathbb{R} \times [0, 1] \rightarrow M$ with boundary on Lagrangian submanifolds L_0 and L_1 , there are three types of phenomena to consider:

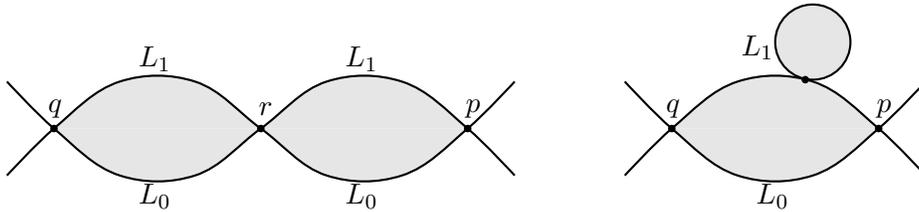


FIGURE 3. Possible limits of pseudo-holomorphic strips: a broken strip (left) and a disc bubble (right).

- (1) *strip breaking*: energy concentrates at either end $s \rightarrow \pm\infty$, i.e. there is a sequence $a_n \rightarrow \pm\infty$ such that the translated strips $u_n(s - a_n, t)$ converge to a non-constant limit strip (Figure 3 left);
- (2) *disc bubbling*: energy concentrates at a point on the boundary of the strip ($t \in \{0, 1\}$), where suitable rescalings of u_n converge to a J -holomorphic disc in M with boundary entirely contained in either L_0 or L_1 (Figure 3 right);
- (3) *sphere bubbling*: energy concentrates at an interior point of the strip, where suitable rescalings of u_n converge to a J -holomorphic sphere in M .

As we will see below, strip breaking is the key geometric ingredient in the proof that the Floer differential squares to zero, *provided that disc bubbling can be excluded*. This is not simply a technical issue – in general the Floer differential does not square to zero, as illustrated by Example 1.11 below. Another issue posed by disc and sphere bubbling is that of transversality: the perturbation techniques we have outlined above are in general not sufficient to achieve transversality for limit curves that include disc or sphere bubble components. More sophisticated techniques, such as those proposed by Fukaya et al [19]¹, or the *polyfolds* developed by Hofer-Wysocki-Zehnder [23], are needed to extend Lagrangian Floer theory to the greatest possible level of generality.

In our case, the absence of disc and sphere bubbles is ensured by the assumption that $[\omega] \cdot \pi_2(M, L_i) = 0$ in the statement of Theorem 1.5. A more general context in which the theory still works is when bubbling can be excluded for dimension reasons, for instance when all bubbles are guaranteed to have Maslov index greater than 2. (The important limit case where the minimal Maslov index is equal to 2 can also be handled by elementary methods; however, in that case disc bubbling can occur and the Floer differential does not automatically square to zero.) A common setting where an *a priori* lower bound on the Maslov index can be guaranteed is that of *monotone* Lagrangian submanifolds in monotone symplectic manifolds, i.e. when the symplectic area of discs and their Maslov index are proportional to each other [32].

1.5. Sketch of proof of Theorem 1.5. The proof that the Floer differential squares to zero (under the assumption that disc and sphere bubbling cannot occur) is conceptually similar to that for Morse (co)homology.

Fix Lagrangian submanifolds L_0 and L_1 as in Theorem 1.5, a generic almost-complex structure J and a Hamiltonian perturbation H so as to ensure transversality. Given two generators p, q of the Floer complex, and a homotopy class $[u]$ with $\text{ind}([u]) = 2$, the moduli space $\mathcal{M}(p, q; [u], J)$ is a 1-dimensional manifold. Since our assumptions exclude the possibilities of disc or sphere bubbling, Gromov compactness implies that this moduli space can be compactified to a space $\overline{\mathcal{M}}(p, q; [u], J)$ whose elements are broken strips connecting p to q and representing the total class $[u]$.

¹The cautious reader should be aware that, as of this writing, the analytic foundations of this approach are still the subject of some controversy.

Two-component broken strips of the sort depicted in Figure 3 (left) correspond to products of moduli spaces $\mathcal{M}(p, r; [u'], J) \times \mathcal{M}(r, q; [u''], J)$, where r is any generator of the Floer complex and $[u'] + [u''] = [u]$. Observe that the index is additive under such decompositions; moreover, transversality implies that any non-constant strip must have index at least 1. Thus, the only possibility is $\text{ind}([u']) = \text{ind}([u'']) = 1$, and broken configurations with more than two components cannot occur.

Conversely, a *gluing theorem* states that every broken strip is locally the limit of a unique family of index 2 strips, and $\overline{\mathcal{M}}(p, q; [u], J)$ is a 1-dimensional manifold with boundary, with

$$(1.7) \quad \partial \overline{\mathcal{M}}(p, q; [u], J) = \coprod_{\substack{r \in \mathcal{X}(L_0, L_1) \\ [u'] + [u''] = [u] \\ \text{ind}([u']) = \text{ind}([u'']) = 1}} \left(\mathcal{M}(p, r; [u'], J) \times \mathcal{M}(r, q; [u''], J) \right)$$

Moreover, the choice of orientations and spin structures on L_0 and L_1 equips all these moduli spaces with natural orientations, and (1.7) is compatible with these orientations (up to an overall sign). Since the total (signed) number of boundary points of a compact 1-manifold with boundary is always zero, we conclude that

$$(1.8) \quad \sum_{\substack{r \in \mathcal{X}(L_0, L_1) \\ [u'] + [u''] = [u] \\ \text{ind}([u']) = \text{ind}([u'']) = 1}} (\#\mathcal{M}(p, r; [u'], J)) (\#\mathcal{M}(r, q; [u''], J)) T^{\omega([u']) + \omega([u''])} = 0.$$

Summing over all possible $[u]$, the left-hand side is precisely the coefficient of q in $\partial^2(p)$; therefore $\partial^2 = 0$.

When L_0 and/or L_1 bound J -holomorphic discs, the sum (1.8) no longer cancels, because the boundary of the 1-dimensional moduli space $\mathcal{M}(p, q; [u], J)$ also contains configurations with disc bubbles. The following example shows that this is an issue even in the monotone case.

Example 1.11. Consider again the cylinder $M = \mathbb{R} \times S^1$, and let L_0 be a simple closed curve that goes around the cylinder once, and L_1 a homotopically trivial curve intersecting L_0 in two points p and q , as shown in Figure 4 left. Then L_0 and L_1

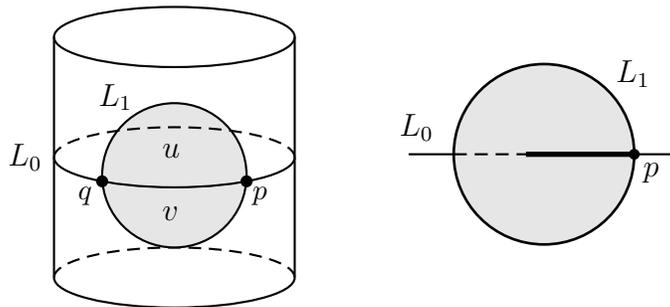


FIGURE 4. A counterexample to $\partial^2 = 0$

bound precisely two holomorphic strips of index 1, denoted by u and v in Figure 4. (There are other holomorphic discs with boundary on L_0 and L_1 but those have higher index.) Comparing with the convention depicted in Figure 2, u is a trajectory from p to q , while v is a trajectory from q to p : thus we have

$$\partial p = \pm T^{\omega(u)} q \quad \text{and} \quad \partial q = \pm T^{\omega(v)} p,$$

and $\partial^2 \neq 0$. To understand why $\partial^2(p) \neq 0$, consider the moduli space of index 2 holomorphic strips connecting p to itself. The images of these strips exactly cover the disc bounded by L_1 , with a slit along L_0 , as shown in Figure 4 right.

We can give an explicit description using local coordinates in which L_0 corresponds to the real axis and L_1 to the unit circle: using the upper half-disc minus the points ± 1 as domain of our maps instead of the usual $\mathbb{R} \times [0, 1]$ (to which it is biholomorphic), one easily checks that any index 2 strip connecting p to itself can be parametrized as

$$u_\alpha(z) = \frac{z^2 + \alpha}{1 + \alpha z^2}$$

for some $\alpha \in (-1, 1)$ (corresponding to the end point of the slit).

The two ends of this moduli space are different: when $\alpha \rightarrow -1$, energy concentrates at $z = \pm 1$, and the index 2 strips u_α converge to a broken strip whose nonconstant components are the index 1 strips u and v ; whereas for $\alpha \rightarrow 1$ the maps u_α exhibit disc bubbling at $z = i$, the limit being a constant strip at p together with a disc bubble whose image is the disc bounded by L_1 . Thus, broken strips do not cancel in pairs in the manner needed for $\partial^2 = 0$ to hold.

Once the Floer differential is shown to square to zero, it remains to prove that Floer cohomology does not depend on the choice of almost-complex structure and Hamiltonian perturbation. Recall that the spaces of such choices are contractible. Thus, given two choices (H, J) and (H', J') (for which we assume transversality holds), let $(H(\tau), J(\tau))$, $\tau \in [0, 1]$ be a (generically chosen) smooth family which agrees with (H, J) for $\tau = 0$ and (H', J') for $\tau = 1$. One can then construct a *continuation map* $F : CF(L_0, L_1; H, J) \rightarrow CF(L_0, L_1; H', J')$ by counting solutions to the equation

$$(1.9) \quad \frac{\partial u}{\partial s} + J(\tau(s), t, u) \left(\frac{\partial u}{\partial t} - X_H(\tau(s), t, u) \right) = 0,$$

where $\tau(s)$ is a smooth function of s which equals 1 for $s \ll 0$ and 0 for $s \gg 0$. Unlike (1.6), the equation (1.9) is not invariant under translations in the s direction. Given generators $p \in \mathcal{X}(L_0, L_1; H)$ and $p' \in \mathcal{X}(L_0, L_1; H')$ of the respective Floer complexes, the coefficient of p' in $F(p)$ is defined as a count of index 0 solutions to (1.9) which converge to p at $s \rightarrow +\infty$ and to p' at $s \rightarrow -\infty$ (weighted by energy as usual).

The proof that F is a chain map, i.e. satisfies $\partial' \circ F = F \circ \partial$ (again assuming the absence of bubbling), comes from studying spaces of index 1 solutions to (1.9). These spaces are 1-dimensional manifolds, whose end points correspond to broken

trajectories where the main component is an index 0 solution to (1.9), either preceded by an index 1 J -holomorphic strip with perturbation data H (if energy concentrates at $s \rightarrow +\infty$), or followed by an index 1 J' -holomorphic strip with perturbation data H' (if energy concentrates at $s \rightarrow -\infty$). The composition $F \circ \partial$ counts the first type of limit configuration, while $\partial' \circ F$ counts the second type of limit configuration, and the equality between these two maps follows again from the statement that the total (signed) number of end points of a compact 1-manifold with boundary is zero.

Using the reverse homotopy, i.e., considering (1.9) with $\tau(s) = 0$ for $s \ll 0$ and 1 for $s \gg 0$, one similarly defines a chain map $F' : CF(L_0, L_1; H', J') \rightarrow CF(L_0, L_1; H, J)$. The chain maps F and F' are quasi-inverses, i.e. $F' \circ F$ is homotopic to identity (and similarly for $F \circ F'$). An explicit homotopy can be obtained by counting index -1 solutions to a one-parameter family of equations similar to (1.9) but where $\tau(s)$ is 0 near $\pm\infty$ and is nonzero over an interval of values of s of varying width.

1.6. The Floer cohomology $HF(L, L)$. The Floer cohomology of a Lagrangian submanifold with itself is of particular interest in the context of Arnold's conjecture. By Weinstein's Lagrangian neighborhood theorem, a neighborhood of a Lagrangian submanifold L in (M, ω) is symplectomorphic to a neighborhood of the zero section of the cotangent bundle T^*L with its standard symplectic form. In light of this, we first consider the model case of the cotangent bundle.

Example 1.12. Let N be a compact real n -dimensional manifold, and consider the cotangent bundle T^*N , with its standard exact symplectic form (given locally by $\omega = \sum dq_i \wedge dp_i$, where q_i are local coordinates on N and p_i are the dual coordinates on the fibers of the cotangent bundle). Let L_0 be the zero section, and given a Morse function $f : N \rightarrow \mathbb{R}$ and a small $\epsilon > 0$, denote by L_1 the graph of the exact 1-form ϵdf . Then L_0, L_1 are exact Lagrangian submanifolds of T^*N , Hamiltonian isotopic to each other (the Hamiltonian isotopy is generated by $H = \epsilon f \circ \pi$ where $\pi : T^*N \rightarrow N$ is the bundle map); L_0 and L_1 intersect transversely at the critical points of f .

Choosing a graded lift of L_0 , and transporting it through the Hamiltonian isotopy to define a graded lift of L_1 , we obtain a grading on the Floer complex $CF(L_0, L_1)$; by an explicit calculation, a critical point p of f of Morse index $i(p)$ defines a generator of the Floer complex of degree $\deg(p) = n - i(p)$. Thus, the grading on the Floer complex agrees with that on the complex $CM^*(f)$ which defines the Morse cohomology of f .

The Morse differential counts index 1 trajectories of the gradient flow between critical points of f , and depends on the choice of a Riemannian metric g on N , which we assume to satisfy the Morse-Smale transversality condition. A result of Floer [18] is that, for a suitable choice of (time-dependent) almost-complex structure J on T^*N , solutions of Floer's equation

$$\frac{\partial u}{\partial s} + J(t, u) \frac{\partial u}{\partial t} = 0$$

with boundary on L_0 and L_1 are regular and in one-to-one correspondence with gradient flow trajectories

$$\dot{\gamma}(s) = \epsilon \nabla f(\gamma(s))$$

on N , the correspondence being given by $\gamma(s) = u(s, 0)$. (Note: an ascending gradient flow line with $\gamma(s)$ converging to p as $s \rightarrow +\infty$ and q as $s \rightarrow -\infty$ counts as a trajectory from p to q in the Morse differential.)

To understand this correspondence between moduli spaces, observe that, at any point x of the zero section, the natural almost-complex structure on T^*N induced by the metric g maps the horizontal vector $\epsilon \nabla f(x) \in T_x N \subset T_x(T^*N)$ to the vertical vector $X_H(x) = \epsilon df(x) \in T_x^* N \subset T_x(T^*N)$. This allows us to construct particularly simple solutions of (1.6) for this almost-complex structure and the Hamiltonian perturbation $-H$, with both boundaries of the strip mapping to L_0 : for any gradient flow line γ of f , we obtain a solution of (1.6) by setting $u(s, t) = \gamma(s)$. Floer's construction of strips with boundary on L_0 and L_1 is equivalent to this via Remark 1.10.

Thus, for specific choices of perturbation data, after a rescaling of the generators by $p \mapsto T^{\epsilon f(p)} p$, the Floer complex of (L_0, L_1) is isomorphic to the Morse complex of f , and the Floer cohomology $HF^*(L_0, L_1)$ is isomorphic to the Morse cohomology of f (with coefficients in Λ). Using the independence of Floer cohomology under Hamiltonian isotopies and the isomorphism between Morse and ordinary cohomology, we conclude that $HF^*(L_0, L_0) \simeq HF^*(L_0, L_1) \simeq H^*(L_0; \Lambda)$.

(Since we are in the exact case, by Remark 1.7 one could actually work directly over \mathbb{K} rather than over Novikov coefficients.)

Now we consider the general case of a compact Lagrangian submanifold L in a symplectic manifold (M, ω) , under the assumption that $[\omega] \cdot \pi_2(M, L) = 0$. Energy estimates then imply that, for a sufficiently small Hamiltonian perturbation, the pseudo-holomorphic strips that determine the Floer cohomology $HF^*(L, L)$ must all be contained in a small tubular neighborhood of L , so that the calculation of Floer cohomology reduces to Example 1.12, and we get the following result (due to Floer in the exact case and for $\mathbb{K} = \mathbb{Z}_2$):

Proposition 1.13. *If $[\omega] \cdot \pi_2(M, L) = 0$, then $HF^*(L, L) \simeq H^*(L; \Lambda)$.*

Together with Theorem 1.5, this implies Arnold's conjecture (Theorem 1.1).

Example 1.14. Let L be the zero section in $T^*S^1 = \mathbb{R} \times S^1$ (see Figure 1 left), and consider the Hamiltonian perturbation depicted in the figure, which comes from a Morse function on $L = S^1$ with a maximum at p and a minimum at q . Then L and $\psi(L)$ bound two index 1 holomorphic strips (shaded on the figure), both connecting p to q , and with equal areas. However, the contributions of these two strips to the Floer differential cancel out (this is obvious over $\mathbb{K} = \mathbb{Z}_2$; when $\text{char}(\mathbb{K}) \neq 2$ a verification of signs is needed). Thus, $\partial p = 0$, and $HF^*(L, \psi(L)) \simeq H^*(S^1)$, as expected from Proposition 1.13.

Things are different when L bounds pseudo-holomorphic discs, and the Floer cohomology $HF^*(L, L)$ (when it is defined) is in general smaller than $H^*(L; \Lambda)$. For example, let L be a monotone Lagrangian submanifold in a monotone symplectic manifold, with minimal Maslov index at least 2; this is a setting where $HF^*(L, L)$ is well defined (though no longer \mathbb{Z} -graded), as disc bubbles either do not occur at all or occur in cancelling pairs. Using again a small multiple ϵf of a Morse function f on L as Hamiltonian perturbation, the Floer complex differs from the Morse complex $CM^*(f)$ by the presence of additional terms in the differential; namely there are index 1 Floer trajectories representing a class in $\pi_2(M, L)$ of Maslov index k and connecting a critical point p of Morse index $i(p)$ to a critical point q of index $i(q) = i(p) + k - 1$. This situation was studied by Oh [32, 33], who showed that the Floer complex is filtered by index (or equivalently energy), and there is a *spectral sequence* starting with the Morse cohomology $HM^*(f)$ (or equivalently the ordinary cohomology of L), whose successive differentials account for classes of increasing Maslov index in $\pi_2(M, L)$, and converging to the Floer cohomology $HF^*(L, L)$.

It is often easier to study honest pseudo-holomorphic discs with boundary on L , rather than solutions of Floer's equation with a Hamiltonian perturbation, or strips with boundary on L and its image under a small isotopy. This has led to the development of alternative constructions of $HF^*(L, L)$. For instance, another model for the Floer cohomology of a monotone Lagrangian submanifold is the *pearl complex* first introduced in [34] (see also [15]). In this model, the generators of the Floer complex are again the critical points of a Morse function f on L , but the differential counts “pearly trajectories”, which arise as limits of Floer trajectories of the sort considered above as $\epsilon \rightarrow 0$. Namely, a pearly trajectory between critical points p and q of f consists of $r \geq 0$ pseudo-holomorphic discs in M with boundary in L , connected to each other and to p and q by $r + 1$ gradient flow lines of f in L . (When there are no discs, a pearly trajectory is simply a gradient flow line between p and q .) Yet another model, proposed by Fukaya-Oh-Ohta-Ono [19], uses a chain complex where $CF(L, L) = C_*(L)$ consists of chains in L , and the differential is the sum of the classical boundary map and a map defined in terms of moduli spaces of pseudo-holomorphic discs with boundary on L . This model is computationally convenient, but requires great care in its construction to address questions such as exactly what sort of chains are considered and, in the general (non-monotone) case, how to achieve transversality of the evaluation maps.

2. PRODUCT OPERATIONS

2.1. The product. Let L_0, L_1, L_2 be three Lagrangian submanifolds of (M, ω) , which we assume intersect each other transversely and do not bound any pseudo-holomorphic discs. We now define a product operation on their Floer complexes, i.e. a map

$$CF(L_1, L_2) \otimes CF(L_0, L_1) \longrightarrow CF(L_0, L_2).$$

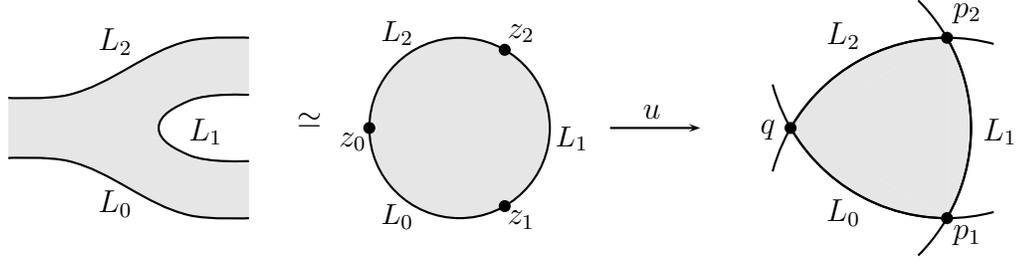


FIGURE 5. A pseudo-holomorphic disc contributing to the product map.

Given intersection points $p_1 \in \mathcal{X}(L_0, L_1)$, $p_2 \in \mathcal{X}(L_1, L_2)$, and $q \in \mathcal{X}(L_0, L_2)$, the coefficient of q in $p_2 \cdot p_1$ is a weighted count of pseudo-holomorphic discs in M with boundary on $L_0 \cup L_1 \cup L_2$ and with corners at p_1, p_2, q . More precisely, let D be the closed unit disc minus three boundary points, say for instance $z_0 = -1, z_1 = e^{-i\pi/3}, z_2 = e^{i\pi/3}$, and observe that a neighborhood of each puncture in D is conformally equivalent to a strip (i.e., the product of an infinite interval with $[0, 1]$).

Given an almost-complex structure J on M and a homotopy class $[u]$, we denote by $\mathcal{M}(p_1, p_2, q; [u], J)$ the space of finite energy J -holomorphic maps $u : D \rightarrow M$ which extend continuously to the closed disc, mapping the boundary arcs from z_0 to z_1 , z_1 to z_2 , z_2 to z_0 to L_0, L_1, L_2 respectively, and the boundary punctures z_1, z_2, z_0 to p_1, p_2, q respectively, in the given homotopy class $[u]$ (see Figure 5).

As in the case of strips, the expected dimension of $\mathcal{M}(p_1, p_2, q; [u], J)$ is given by the index of the linearized Cauchy-Riemann operator $D_{\bar{\partial}_{J,u}}$. This index can be expressed in terms of the Maslov index, exactly as in Definition 1.8: we now concatenate the paths given by the tangent spaces to L_0, L_1, L_2 going counterclockwise along the boundary of u , together with the appropriate canonical short paths at p_1, p_2, q , to obtain a closed loop in $LGr(n)$ whose Maslov index is equal to $\text{ind}(u)$. If $c_1(TM)$ is 2-torsion and the Maslov classes of L_0, L_1, L_2 vanish, then after choosing graded lifts of the Lagrangians we have \mathbb{Z} -gradings on the Floer complexes, and one checks that

$$(2.1) \quad \text{ind}(u) = \text{deg}(q) - \text{deg}(p_1) - \text{deg}(p_2).$$

Remark 2.1. The apparent lack of symmetry in the index formula (2.1) is due to the difference between the gradings on $CF(L_0, L_2)$ and $CF(L_2, L_0)$. Namely, the given intersection point $q \in L_0 \cap L_2$ defines generators of both complexes, whose degrees sum to n (the dimension of L_i). In fact, the Floer complexes $CF(L_0, L_2)$ and $CF(L_2, L_0)$ and the differentials on them are dual to each other, provided that the almost-complex structures and perturbations are chosen suitably. For instance, the strip depicted in Figure 2 is a trajectory from p to q in the Floer complex $CF(L_0, L_1)$, and from q to p in $CF(L_1, L_0)$.

Assume that transversality holds, so that the moduli spaces $\mathcal{M}(p_1, p_2, q; [u], J)$ are smooth manifolds; if $\text{char}(\mathbb{K}) \neq 2$, assume moreover that orientations and spin

structures on L_0, L_1, L_2 have been chosen, so as to determine orientations of the moduli spaces. Then we define:

Definition 2.2. *The Floer product is the Λ -linear map $CF(L_1, L_2) \otimes CF(L_0, L_1) \rightarrow CF(L_0, L_2)$ defined by*

$$(2.2) \quad p_2 \cdot p_1 = \sum_{\substack{q \in \mathcal{X}(L_0, L_2) \\ [u]: \text{ind}([u])=0}} (\#\mathcal{M}(p_1, p_2, q; [u], J)) T^{\omega([u])} q.$$

As in the previous section, in general this construction needs to be modified by introducing domain-dependent almost-complex structures and Hamiltonian perturbations to achieve transversality. We discuss this below, but for now we assume transversality holds without further perturbations and examine the properties of the Floer product.

Proposition 2.3. *If $[\omega] \cdot \pi_2(M, L_i) = 0$ for all i , then the Floer product satisfies the Leibniz rule (with suitable signs) with respect to the Floer differentials,*

$$(2.3) \quad \partial(p_2 \cdot p_1) = \pm(\partial p_2) \cdot p_1 \pm p_2 \cdot (\partial p_1),$$

and hence induces a well-defined product $HF(L_1, L_2) \otimes HF(L_0, L_1) \rightarrow HF(L_0, L_2)$. Moreover, this induced product on Floer cohomology groups is independent of the chosen almost-complex structure (and Hamiltonian perturbations) and associative.

(However, the chain-level product on Floer complexes is *not* associative, as we will see below.)

We now sketch the geometric argument behind the Leibniz rule, which relies on an examination of index 1 moduli spaces of J -holomorphic discs and their compactification. Namely, consider a triple of generators p_1, p_2, q as above, and let $[u]$ be a homotopy class with $\text{ind}([u]) = 1$. Then (still assuming transversality) $\mathcal{M}(p_1, p_2, q; [u], J)$ is a smooth 1-dimensional manifold, and by Gromov compactness admits a compactification $\overline{\mathcal{M}}(p_1, p_2, q; [u], J)$ obtained by adding nodal trees of J -holomorphic curves.

Since our assumptions exclude bubbling of discs or spheres, the only phenomenon that can occur is strip-breaking (when energy concentrates at one of the three ends of the punctured disc D). Since transversality excludes the presence of discs of index less than 0 and nonconstant strips of index less than 1, and since the sum of the indices of the limit components must be 1, there are only three types of limit configurations to be considered, all consisting of an index 0 disc with boundary on L_0, L_1, L_2 and an index 1 strip with boundary on two of these three submanifolds; see Figure 6.

The three types of configurations contribute to the coefficient of $T^{\omega([u])}q$ in $\partial(p_2 \cdot p_1)$ (Figure 6 left), $(\partial p_2) \cdot p_1$ (middle), and $p_2 \cdot (\partial p_1)$ (right) respectively. On the other hand, a gluing theorem states that every such configuration arises as an end of $\mathcal{M}(p_1, p_2, q; [u], J)$, and that the compactified moduli space is a 1-dimensional compact manifold with boundary. Moreover, the orientations agree up to overall sign

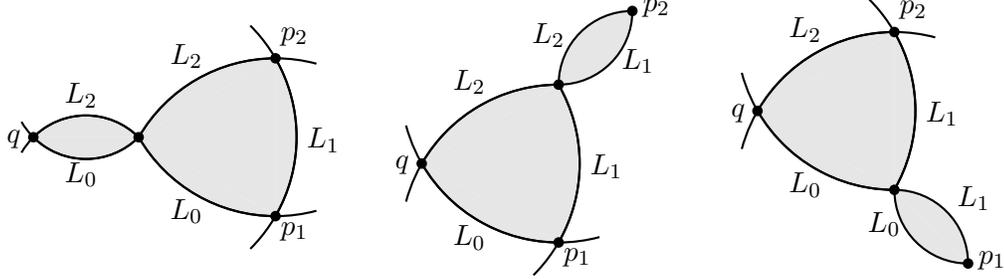


FIGURE 6. The ends of a 1-dimensional moduli space $\mathcal{M}(p_1, p_2, q; [u], J)$.

factors depending only on the degrees of p_1 and p_2 . Since the (signed) total number of boundary points of $\overline{\mathcal{M}}(p_1, p_2, q; [u], J)$ is zero, the Leibniz rule (2.3) follows.

Before moving on to higher products, we briefly discuss the issue of transversality and compatibility in the choice of perturbations. As in the case of strips, even without assuming that L_0, L_1, L_2 intersect transversely, we can ensure transversality by introducing domain-dependent almost-complex structures and Hamiltonian perturbations; however, for the Leibniz rule to hold, these need to be chosen suitably near the punctures z_0, z_1, z_2 . Fix once and for all “strip-like ends” near the punctures, i.e. biholomorphisms from $\mathbb{R}_+ \times [0, 1]$ (resp. $\mathbb{R}_- \times [0, 1]$) to neighborhoods of the punctures z_1 and z_2 (resp. z_0) in D ; we denote by $s + it$ the natural complex coordinate in each strip-like end. Also fix a 1-form $\beta \in \Omega^1(D)$, such that $\beta|_{\partial D} = 0$ and $\beta = dt$ in each strip-like end. Now, given L_0, L_1, L_2 , we choose a family of ω -compatible almost-complex structures depending smoothly on $z \in D$, i.e. $J \in C^\infty(D, \mathcal{J}(M, \omega))$, and a family of Hamiltonians $H \in C^\infty(D \times M, \mathbb{R})$, with the property that in each strip-like end $J(z)$ and $H(z)$ depend only on the coordinate $t \in [0, 1]$. We then perturb the Cauchy-Riemann equation to

$$(2.4) \quad \left(du - X_H \otimes \beta \right)_J^{0,1} = 0,$$

which in each strip-like end reduces to (1.6).

For $0 \leq i < j \leq 2$, denote by $H_{ij} \in C^\infty([0, 1] \times M, \mathbb{R})$ and $J_{ij} \in C^\infty([0, 1], \mathcal{J}(M, \omega))$ the time-dependent Hamiltonians and almost-complex structures on the strip-like end whose boundaries map to L_i and L_j . The solutions of (2.4) converge no longer to intersection points of $L_i \cap L_j$, but to trajectories of the time 1 flow generated by H_{ij} which begin on L_i and end on L_j , i.e. generators of the perturbed Floer complex of (L_i, L_j) with respect to the Hamiltonian perturbation H_{ij} . Moreover, when strip breaking occurs, the main component remains a solution of (2.4), while the strip component that breaks off is a solution of (1.6) with respect to H_{ij} and J_{ij} .

Thus, by considering the moduli spaces of solutions to the perturbed equation (2.4) and proceeding as in Definition 2.2, we obtain a product map

$$CF(L_1, L_2; H_{12}, J_{12}) \otimes CF(L_0, L_1; H_{01}, J_{01}) \longrightarrow CF(L_0, L_2; H_{02}, J_{02})$$

on the perturbed Floer complexes, and Proposition 2.3 still holds (with respect to the perturbed Floer differentials).

2.2. Higher operations. Given $k + 1$ Lagrangian submanifolds L_0, \dots, L_k , a construction similar to those above allows us to define an operation

$$\mu^k : CF(L_{k-1}, L_k) \otimes \cdots \otimes CF(L_1, L_2) \otimes CF(L_0, L_1) \longrightarrow CF(L_0, L_k)$$

(of degree $2 - k$ in the situation where the Floer complexes are graded), where μ^1 is the Floer differential and μ^2 is the product.

Given generators $p_i \in \mathcal{X}(L_{i-1}, L_i)$ ($i = 1, \dots, k$) and $q \in \mathcal{X}(L_0, L_k)$, the coefficient of q in $\mu^k(p_k, \dots, p_1)$ is a count (weighted by area) of (perturbed) pseudo-holomorphic discs in M with boundary on $L_0 \cup \cdots \cup L_k$ and corners at p_1, \dots, p_k, q .

Specifically, one considers maps $u : D \rightarrow M$ whose domain D is the closed unit disc minus $k + 1$ boundary points $z_0, z_1, \dots, z_k \in S^1$, lying in that order along the unit circle. The positions of these marked points are not fixed, and the moduli space $\mathcal{M}_{0,k+1}$ of conformal structures on the domain D , i.e., the quotient of the space of ordered $(k + 1)$ -tuples of points on S^1 by the action of $\text{Aut}(D^2)$, is a contractible $(k - 2)$ -dimensional manifold.

Given an almost-complex structure J on M and a homotopy class $[u]$, we denote by $\mathcal{M}(p_1, \dots, p_k, q; [u], J)$ the space of J -holomorphic maps $u : D \rightarrow M$ (where the positions of z_0, \dots, z_k are not fixed a priori) which extend continuously to the closed disc, mapping the boundary arcs from z_i to z_{i+1} (or z_0 for $i = k$) to L_i , and the boundary punctures z_1, \dots, z_k, z_0 to p_1, \dots, p_k, q respectively, in the given homotopy class $[u]$, up to the action of $\text{Aut}(D^2)$ by reparametrization. (Or, equivalently, one can avoid quotienting and instead take a slice for the reparametrization action by fixing the positions of three of the z_i .)

For a fixed conformal structure on D , the index of the linearized Cauchy-Riemann operator is again given by the Maslov index, as previously. Thus, accounting for deformations of the conformal structure on D , assuming transversality, the expected dimension of the moduli space is

$$(2.5) \quad \dim \mathcal{M}(p_1, \dots, p_k, q; [u], J) = k - 2 + \text{ind}([u]) = k - 2 + \deg(q) - \sum_{i=1}^k \deg(p_i).$$

Thus, assuming transversality, and choosing orientations and spin structures on L_0, \dots, L_k if $\text{char}(\mathbb{K}) \neq 2$, we define:

Definition 2.4. *The operation $\mu^k : CF(L_{k-1}, L_k) \otimes \cdots \otimes CF(L_0, L_1) \rightarrow CF(L_0, L_k)$ is the Λ -linear map defined by*

$$(2.6) \quad \mu^k(p_k, \dots, p_1) = \sum_{\substack{q \in \mathcal{X}(L_0, L_k) \\ [u]: \text{ind}([u]) = 2 - k}} (\#\mathcal{M}(p_1, \dots, p_k, q; [u], J)) T^{\omega([u])} q.$$

Remark 2.5. As before, in general this construction needs to be modified by introducing domain-dependent almost-complex structures and Hamiltonian perturbations to achieve transversality. Thus, we actually count solutions of a perturbed Cauchy-Riemann equation similar to (2.4), involving a domain-dependent almost-complex structure $J \in C^\infty(D, \mathcal{J}(M, \omega))$ and Hamiltonian $H \in C^\infty(D \times M, \mathbb{R})$. As before, compatibility with strip-breaking requires that, in each of the $k + 1$ strip-like ends near the punctures of D , the chosen J and H depend only on the coordinate $t \in [0, 1]$ and agree with the almost-complex structures and Hamiltonians used to construct the Floer complexes $CF(L_i, L_{i+1})$ and $CF(L_0, L_k)$. An additional compatibility condition comes from the possible degenerations of the domain D to unions of discs with fewer punctures, as discussed below: we need to require that, when D degenerates in such a way, H and J are translation-invariant in the strip-like regions connecting the components and agree with the choices made in the construction of the Floer complexes $CF(L_i, L_j)$, while in each component H and J agree with the choices made for that moduli space of discs with fewer punctures. This forces the choices of H and J to further depend on the conformal structure of D . We refer the reader to [40] for a detailed construction (and proof of existence) of compatible and consistent choices of perturbation data (H, J) .

The algebraic properties of μ^k follow from the study of the limit configurations that arise in compactifications of 1-dimensional moduli spaces of (perturbed) pseudo-holomorphic discs; besides strip breaking, there are now other possibilities, corresponding to cases where the domain D degenerates. The moduli space of conformal structures $\mathcal{M}_{0,k+1}$ admits a natural compactification to a $(k - 2)$ -dimensional polytope $\overline{\mathcal{M}}_{0,k+1}$, the Stasheff *associahedron*, whose top-dimensional facets correspond to nodal degenerations of D to a pair of discs $D_1 \cup D_2$, with each component carrying at least two of the marked points z_0, \dots, z_k ; and the higher codimension faces correspond to nodal degenerations with more components.

Example 2.6. $\overline{\mathcal{M}}_{0,4}$ is homeomorphic to a closed interval, whose end points correspond to configurations where two adjacent marked points come together (Figure 7). For example, fixing the positions of z_0, z_1, z_2 on the unit circle and letting z_3 vary

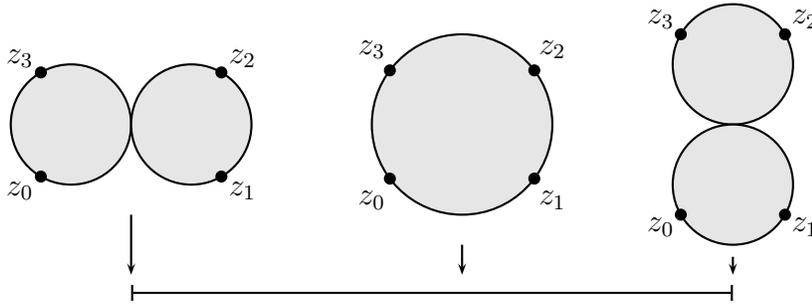


FIGURE 7. The 1-dimensional associahedron $\overline{\mathcal{M}}_{0,4}$.

along the arc from z_2 to z_0 , the right end point corresponds to the case where z_3 approaches z_2 ; the “main” component of the limit configuration carries the marked points z_0 and z_1 , while the component carrying z_2 and z_3 arises from rescaling by suitable automorphisms of the disc. Equivalently up to automorphisms of the disc, one could instead fix the positions of z_1, z_2, z_3 , and let z_0 vary along the arc from z_3 to z_1 ; the right end point then corresponds to the case where z_0 approaches z_1 .

Proposition 2.7. *If $[\omega] \cdot \pi_2(M, L_i) = 0$ for all i , then the operations μ^k satisfy the A_∞ -relations*

$$(2.7) \quad \sum_{\ell=1}^k \sum_{j=0}^{k-\ell} (-1)^* \mu^{k+1-\ell}(p_k, \dots, p_{j+\ell+1}, \mu^\ell(p_{j+\ell}, \dots, p_{j+1}), p_j, \dots, p_1) = 0,$$

where $*$ = $j + \deg(p_1) + \dots + \deg(p_j)$.

The case $k = 1$ of (2.7) is the identity $\partial^2 = 0$, while $k = 2$ corresponds to the Leibniz rule (2.3). For $k = 3$, it expresses the fact that the Floer product μ^2 is associative up to an explicit homotopy given by μ^3 :

$$(2.8) \quad \pm (p_3 \cdot p_2) \cdot p_1 \pm p_3 \cdot (p_2 \cdot p_1) = \pm \partial \mu^3(p_3, p_2, p_1) \pm \mu^3(\partial p_3, p_2, p_1) \pm \mu^3(p_3, \partial p_2, p_1) \pm \mu^3(p_3, p_2, \partial p_1).$$

More generally, each operation μ^k gives an explicit homotopy for a certain compatibility property among the preceding ones.

The proof of Proposition 2.7 again relies on an analysis of 1-dimensional moduli spaces of (perturbed) J -holomorphic discs and their compactification. Fix generators p_1, \dots, p_k, q and a homotopy class $[u]$ with $\text{ind}([u]) = 3 - k$, and assume that J and H are chosen generically (so as to achieve transversality) and compatibly (see Remark 2.5). Then the moduli space $\mathcal{M}(p_1, \dots, p_k, q; [u], J)$ compactifies to a 1-dimensional manifold with boundary, whose boundary points correspond either to an index 1 (perturbed) J -holomorphic strip breaking off at one of the $k + 1$ marked points, or to a degeneration of the domain to the boundary of $\overline{\mathcal{M}}_{0, k+1}$, i.e. to a pair of discs with each component carrying at least two of the marked points. The first case corresponds to the terms involving μ^1 in (2.7), while the second case corresponds to the other terms.

Example 2.8. For $k = 3$, limit configurations consisting of an index 1 strip together with an index -1 disc with 4 marked points account for the right-hand side in (2.8), while those consisting of a pair of index 0 discs with 3 marked points (when the domain degenerates to one of the two end points of $\overline{\mathcal{M}}_{0,4}$, see Figure 7) account for the two terms in the left-hand side.

2.3. The Fukaya category. There are several variants of the Fukaya category of a symplectic manifold, depending on the desired level of generality and a number of

implementation details. The common features are the following. The objects of the Fukaya category are suitable Lagrangian submanifolds, equipped with extra data, and morphism spaces are given by Floer complexes, endowed with the Floer differential. Composition of morphisms is given by the Floer product, which is only associative up to homotopy, and the Fukaya category is an A_∞ -category, i.e. the differential and composition are the first two in a sequence of operations

$$\mu^k : \text{hom}(L_{k-1}, L_k) \otimes \cdots \otimes \text{hom}(L_0, L_1) \rightarrow \text{hom}(L_0, L_k)$$

(of degree $2 - k$ when a \mathbb{Z} -grading is available), satisfying the A_∞ -relations (2.7).

Given the setting in which we have developed Floer theory in the preceding sections, the most natural definition is the following:

Definition 2.9. *Let (M, ω) be a symplectic manifold with $2c_1(TM) = 0$. The objects of the (compact) Fukaya category $\mathcal{F}(M, \omega)$ are compact closed, oriented, spin Lagrangian submanifolds $L \subset M$ such that $[\omega] \cdot \pi_2(M, L) = 0$ and with vanishing Maslov class $\mu_L = 0 \in H^1(L, \mathbb{Z})$, together with extra data, namely the choice of a spin structure and a graded lift of L . (We will usually omit those from the notation and simply denote the object by L .)*

For every pair of objects (L, L') (not necessarily distinct), we choose perturbation data $H_{L, L'} \in C^\infty([0, 1] \times M, \mathbb{R})$ and $J_{L, L'} \in C^\infty([0, 1], \mathcal{J}(M, \omega))$; and for all tuples of objects (L_0, \dots, L_k) and all moduli spaces of discs, we choose consistent perturbation data (H, J) compatible with the choices made for the pairs of objects (L_i, L_j) , so as to achieve transversality for all moduli spaces of perturbed J -holomorphic discs. (See [40, §9] for the existence of such perturbation data.)

Given this, we set $\text{hom}(L, L') = CF(L, L'; H_{L, L'}, J_{L, L'})$; and the differential μ^1 , composition μ^2 , and higher operations μ^k are given by counts of perturbed pseudo-holomorphic discs as in Definition 2.4. By Proposition 2.7, this makes $\mathcal{F}(M, \omega)$ a Λ -linear, \mathbb{Z} -graded, non-unital (but cohomologically unital [40]) A_∞ -category.

One can also consider other settings: for example, we can drop the requirement that $2c_1(TM) = 0$ and the assumption of vanishing of the Maslov class if we are content with a $\mathbb{Z}/2$ -grading; spin structures can be ignored if we work over a field of characteristic 2; and Novikov coefficients are unnecessary if we restrict ourselves to exact Lagrangian submanifolds in an exact symplectic manifold.

As is obvious from the definition, the actual chain-level details of the Fukaya category depend very much on the choice of perturbation data; however, the A_∞ -categories obtained from various choices of perturbation data are *quasi-equivalent* (i.e., they are related by A_∞ -functors which induce equivalences, in fact in this case isomorphisms, at the level of cohomology) [40].

We finish this section with a few remarks.

Remark 2.10. One can recover an honest category from an A_∞ -category by taking the cohomology of morphism spaces with respect to the differential μ^1 ; the A_∞ -relations imply that μ^2 descends to an associative composition operation on cohomology. The cohomology category of $\mathcal{F}(M, \omega)$, where $\text{hom}(L, L') = HF(L, L')$ and composition is given by the cohomology-level Floer product, is sometimes called the Donaldson-Fukaya category. However, the higher operations contain important information that gets lost when passing to the cohomology category, and it is usually much better to work with the chain-level A_∞ -category (see for instance the next section).

Remark 2.11. In the context of homological mirror symmetry, one is naturally led to consider a slightly richer version of the Fukaya category, whose objects are Lagrangian submanifolds equipped with *local systems*, i.e. flat vector bundles $\mathcal{E} \rightarrow L$ with unitary holonomy (over the Novikov field over $\mathbb{K} = \mathbb{C}$). In this situation, we define

$$CF((L_0, \mathcal{E}_0), (L_1, \mathcal{E}_1)) = \bigoplus_{p \in \mathcal{X}(L_0, L_1)} \text{hom}(\mathcal{E}_{0|p}, \mathcal{E}_{1|p}),$$

and modify the definition of μ^k as follows. Fix objects $(L_0, \mathcal{E}_0), \dots, (L_k, \mathcal{E}_k)$, intersections p_1, \dots, p_k, q , and a homotopy class $[u]$. Set $p_0 = p_{k+1} = q$ for simplicity. Parallel transport along the portion of the boundary of $[u]$ that lies on L_i yields an isomorphism $\gamma_i \in \text{hom}(\mathcal{E}_{i|p_i}, \mathcal{E}_{i|p_{i+1}})$ for each $i = 0, \dots, k$. Now, given elements $\rho_i \in \text{hom}(\mathcal{E}_{i-1|p_i}, \mathcal{E}_{i|p_i})$ ($i = 1, \dots, k$), the composition of all these linear maps defines an element $\eta_{[u], \rho_k, \dots, \rho_1} = \gamma_k \cdot \rho_k \cdot \dots \cdot \gamma_1 \cdot \rho_1 \cdot \gamma_0 \in \text{hom}(\mathcal{E}_{0|q}, \mathcal{E}_{k|q})$. Then we set

$$\mu^k(\rho_k, \dots, \rho_1) = \sum_{\substack{q \in \mathcal{X}(L_0, L_k) \\ [u]: \text{ind}([u]) = 2-k}} (\#\mathcal{M}(p_1, \dots, p_k, q; [u], J)) T^{\omega([u])} \eta_{[u], \rho_k, \dots, \rho_1}.$$

Remark 2.12. It is in principle possible to lift the assumption $[\omega] \cdot \pi_2(M, L) = 0$ we have made throughout, at the expense of considerable analytic and algebraic difficulties in situations where disc bubbling occurs. Analytically, disc bubbles pose transversality problems that cannot be solved with the techniques we have described above. Algebraically, they lead to a *curved* A_∞ -category, i.e. for each object L we have an element $\mu_L^0 \in \text{hom}(L, L)$ which encodes a weighted count of J -holomorphic discs bounded by L . The A_∞ -relations (2.7) are then modified by allowing the case $\ell = 0$ in the sum. For example, the relation for $k = 1$ becomes

$$\mu^1(\mu^1(p)) + (-1)^{\deg p} \mu^2(\mu_{L_1}^0, p) + \mu^2(p, \mu_{L_0}^0) = 0,$$

where the last two terms correspond to disc bubbling along either edge of an index 2 strip. To regain some sanity, one usually considers not arbitrary objects, but *weakly unobstructed* objects, i.e. those for which μ_L^0 is a scalar multiple of the (cohomological) unit of $\text{hom}(L, L)$ (this multiple is sometimes called “central charge” or “superpotential” in the context of mirror symmetry); this happens for instance when the minimal Maslov index of a holomorphic disc with boundary on L is equal to two and Maslov

index 2 discs are regular. Weakly unobstructed objects of fixed central charge then form an honest A_∞ -category. The curious reader is referred to [19].

3. EXACT TRIANGLES AND GENERATORS

While it is usually impossible to classify all Lagrangian submanifolds of a given symplectic manifold, or even to directly compute Floer cohomology for all those we can find, it is often possible to understand the whole Fukaya category in terms of a small subset of *generating objects* – *provided that we understand not only differentials and products but also higher operations among those generators*. To understand how this comes about, a healthy dose of homological algebra is necessary; in this section we give a very brief and informal overview of exact triangles, twisted complexes and generators, in general and as they pertain to Fukaya categories in particular. The first part of [40] fills in the many details that we omit here, and more.

3.1. Exact triangles and mapping cones. An exact triangle

$$\begin{array}{ccc} A & \xrightarrow{f} & B \\ & \searrow^{[1]} & \downarrow g \\ & h & C \end{array}$$

in an A_∞ -category \mathcal{A} consists of a triple of objects A, B, C and closed morphisms $f \in \text{hom}^0(A, B)$, $g \in \text{hom}^0(B, C)$, $h \in \text{hom}^1(C, A)$ such that C is (up to quasi-isomorphism) a *mapping cone* of $f : A \rightarrow B$, with g and h the natural maps to and from it. We will clarify the meaning of this definition in the next section; for now, we simply mention some key features and motivate the concept.

Exactness means that the compositions $\mu^2(g, f)$, $\mu^2(h, g)$ and $\mu^2(f, h)$ are exact, i.e. in the cohomology category $H(\mathcal{A})$ the maps compose to zero. (However, their triple Massey product is typically nontrivial.) An exact triangle induces long exact sequences on morphism spaces in the cohomology category: for every test object T , we have a long exact sequence

$$(3.1) \quad \cdots \rightarrow H^i \text{hom}(T, A) \xrightarrow{f} H^i \text{hom}(T, B) \xrightarrow{g} H^i \text{hom}(T, C) \xrightarrow{h} H^{i+1} \text{hom}(T, A) \xrightarrow{f} \cdots$$

where $H^i \text{hom}(T, A)$ is the cohomology of $\text{hom}(T, A)$ with respect to the differential μ^1 , and the maps are given by composition (in the cohomology category) with f, g , and h ; and similarly (in the contravariant direction) for morphisms from A, B, C to T . Moreover, as T varies these long exact sequences fit together naturally with respect to the multiplicative action of the groups $H^* \text{hom}(T', T)$, i.e. (3.1) fits into an exact sequence of modules over $H(\mathcal{A})$.

Exact triangles can also be characterized as images under A_∞ -functors of a “universal” abstract exact triangle living in an A_∞ -category with three objects [40, §3g].

The A_∞ -category \mathcal{A} is said to be *triangulated* if every closed morphism $f : A \rightarrow B$ can be completed to an exact triangle (and the shift functor $[1]$ acting on \mathcal{A} by change of gradings is a quasi-equivalence); or, in other terms, if all morphisms in \mathcal{A} have mapping cones. Here it is important to point out a key difference with the case of ordinary triangulated categories, where the triangles are an additional piece of structure on the category: the A_∞ -structure is rich enough to “know” about triangles, and triangles automatically satisfy an analogue of the usual axioms. In the same vein, A_∞ -functors are always exact, i.e. map exact triangles to exact triangles.

Before saying more about mapping cones in A_∞ -categories, let us discuss some classical motivating examples.

Example 3.1. The mapping cone of a continuous map $f : X \rightarrow Y$ between topological spaces is, by definition, the space obtained from $X \times [0, 1]$ by attaching Y to $X \times \{1\}$ via the map f and collapsing $X \times \{0\}$ to a point:

$$\text{Cone}(f) = ((X \times [0, 1]) \sqcup Y) / (x, 0) \sim (x', 0), (x, 1) \sim f(x) \quad \forall x, x' \in X.$$

We then have a sequence of maps

$$X \xrightarrow{f} Y \xrightarrow{i} \text{Cone}(f) \xrightarrow{p} \Sigma X \xrightarrow{\Sigma f} \Sigma Y \rightarrow \dots,$$

where i is the inclusion of Y into the mapping cone, and p is the projection to the suspension of X obtained by collapsing Y . The composition of any two of these maps is nullhomotopic, and the induced maps on (co)homology form a long exact sequence.

Example 3.2. The notion of mapping cone in the category of chain complexes is directly modelled on the previous example: let $A = (\bigoplus A^i, d_A)$ and $B = (\bigoplus B^i, d_B)$ be two chain complexes, and let $f : A \rightarrow B$ be a chain map (i.e., a collection of maps $f^i : A^i \rightarrow B^i$ satisfying $d_B f^i + f^{i+1} d_A = 0$). Then the mapping cone of f is, by definition, the chain complex $C = A[1] \oplus B$ (i.e., $C^i = A^{i+1} \oplus B^i$), equipped with the differential

$$d_C = \begin{pmatrix} d_A & 0 \\ f & d_B \end{pmatrix}.$$

The map f , the inclusion of B into C as a subcomplex, and the projection of C onto the quotient complex $A[1]$ then fit into an exact sequence.

Example 3.3. Let A be an algebra (resp. differential graded algebra or A_∞ -algebra), and consider the category of differential graded modules (resp. A_∞ -modules) over A . Recall that such a module M is a chain complex equipped with a degree 1 differential d_M and a multiplication map $A \otimes M \rightarrow M$, $(a, m) \mapsto a \cdot m$, satisfying the Leibniz rule and associative (up to homotopies given by higher structure maps $\mu_M^{k|1} : A^{\otimes k} \otimes M \rightarrow M[1 - k]$, in the case of A_∞ -modules). The mapping cone of a module homomorphism $f : M \rightarrow N$ can then be defined essentially as in the previous example. In the differential graded case, f is a chain map compatible with the multiplication, and the mapping cone of f as a chain complex inherits a natural module

structure. For A_∞ -modules, recalling that an A_∞ -homomorphism is a collection of maps $f^{k|1} : A^{\otimes k} \otimes M \rightarrow N[-k]$ (where the linear term $f^{0|1}$ is a chain map compatible with the product $\mu^{1|1}$ up to a homotopy given by $f^{1|1}$, and so on), the structure maps $\mu_K^{k|1} : A^{\otimes k} \otimes K \rightarrow K[1-k]$ ($k \geq 0$) of the mapping cone $K = M[1] \oplus N$ are given by $\mu_K^{k|1}(a_1, \dots, a_k, (m, n)) = (\mu_M^{k|1}(a_1, \dots, a_k, m), f^{k|1}(a_1, \dots, a_k, m) + \mu_N^{k|1}(a_1, \dots, a_k, n))$.

3.2. Twisted complexes. When an A_∞ -category \mathcal{A} is not known to be triangulated, it is often advantageous to embed it into a larger category in which mapping cones are guaranteed to exist. For example, one can always do so by using the *Yoneda embedding* construction into the category of A_∞ -modules over \mathcal{A} (in which mapping cones always exist, cf. Example 3.3); see e.g. [40, §1]. A milder construction, which retains more features of the original category \mathcal{A} , involves *twisted complexes*. We give a brief outline, and refer the reader to [40, §3] for details.

Definition 3.4. A twisted complex (E, δ^E) consists of:

- a formal direct sum $E = \bigoplus_{i=1}^N E_i[k_i]$ of shifted objects of \mathcal{A} (i.e., a finite collection of pairs (E_i, k_i) where $E_i \in \text{ob } \mathcal{A}$ and $k_i \in \mathbb{Z}$);
- a strictly lower triangular differential $\delta^E \in \text{End}^1(E)$, i.e. a collection of maps $\delta_{ij}^E \in \text{Hom}^{k_j - k_i + 1}(E_i, E_j)$, $1 \leq i < j \leq N$, satisfying the equation

$$(3.2) \quad \sum_{k \geq 1} \mu^k(\delta^E, \dots, \delta^E) = 0,$$

$$\text{i.e., } \sum_{k \geq 1} \sum_{i=i_0 < i_1 < \dots < i_k=j} \mu^k(\delta_{i_{k-1}i_k}^E, \dots, \delta_{i_0i_1}^E) = 0 \text{ for all } 1 \leq i < j \leq N.$$

A degree d morphism of twisted complexes is simply a degree d map between the underlying formal direct sums, i.e. if $E = \bigoplus E_i[k_i]$ and $E' = \bigoplus E'_j[k'_j]$ then an element of $\text{Hom}^d(E, E')$ is by definition a collection of morphisms $a_{ij} \in \text{Hom}^{d+k'_j - k_i}(E_i, E'_j)$.

Finally, given twisted complexes $(E_0, \delta^0), \dots, (E_k, \delta^k)$, $k \geq 1$, and morphisms $a_i \in \text{Hom}(E_{i-1}, E_i)$, we set

$$\mu_{\text{Tw}}^k(a_k, \dots, a_1) = \sum_{j_0, \dots, j_k \geq 0} \mu^{k+j_0+\dots+j_k}(\underbrace{\delta^k, \dots, \delta^k}_{j_k}, a_k, \dots, \underbrace{\delta^1, \dots, \delta^1}_{j_1}, a_1, \underbrace{\delta^0, \dots, \delta^0}_{j_0}).$$

(The sum is finite since each δ^i is strictly lower triangular).

Proposition 3.5. The above construction defines a triangulated A_∞ -category which we denote by $\text{Tw } \mathcal{A}$, and into which \mathcal{A} embeds fully faithfully.

It is instructive to see how twisted complexes relate to ordinary chain complexes:

Example 3.6. Given objects A, B, C of \mathcal{A} and $f \in \text{hom}^0(A, B)$, $g \in \text{hom}^0(A, C)$, we can consider $(A[2] \oplus B[1] \oplus C, \delta = f + g)$, conventionally denoted by

$$\{A \xrightarrow{f} B \xrightarrow{g} C\}.$$

This forms a twisted complex if and only if $\mu^1(f) = \mu^1(g) = 0$ and $\mu^2(g, f) = 0$, i.e. f and g are closed morphisms and their composition is zero. However, we can also introduce an extra term $h \in \text{hom}^{-1}(A, C)$ into the differential δ , in which case the last condition becomes $\mu^2(g, f) + \mu^1(h) = 0$: thus it is sufficient for the composition of f and g to be exact, with a homotopy given by h .

Definition 3.7. Given twisted complexes $(E, \delta), (E', \delta') \in \text{Tw } \mathcal{A}$ and a closed morphism $f \in \text{hom}^0(E, E')$ (i.e., such that $\mu_{\text{Tw}}^1(f) = 0$), the abstract mapping cone of f is the twisted complex

$$\text{Cone}(f) = \left(E[1] \oplus E', \begin{pmatrix} \delta & 0 \\ f & \delta' \end{pmatrix} \right).$$

Given objects A, B, C of \mathcal{A} and a closed morphism $f \in \text{hom}^0(A, B)$, we say that C is a mapping cone of f if, in the category of twisted complexes $\text{Tw } \mathcal{A}$, the object C is quasi-isomorphic to the abstract mapping cone of f , $\{A \xrightarrow{f} B\} = (A[1] \oplus B, f)$.

When C is a mapping cone of $f : A \rightarrow B$, by composing the inclusion of B into the abstract mapping cone (resp. the projection to $A[1]$) with the given quasi-isomorphism from the abstract mapping cone to C (resp. its quasi-inverse) we obtain morphisms $i : B \rightarrow C$ and $p : C \rightarrow A[1]$, which sit with f in an exact triangle.

3.3. Exact triangles in the Fukaya category. The reader may legitimately wonder about the relevance of the above discussion to Fukaya categories. It turns out that at least some mapping cones in the Fukaya category of a symplectic manifold can be understood geometrically. There are two well-known sources of these: Dehn twists, and Lagrangian connected sums.

3.3.1. Dehn twists. The symplectic geometry of Dehn twists was first considered by Arnold, and later studied extensively by Seidel [42, 40]. The local model is as follows. In the cotangent bundle T^*S^n equipped with its canonical symplectic form, a Hamiltonian of the form $H(p, q) = h(\|p\|)$ (where p is the fiber coordinate and $\|\cdot\|$ is the standard metric) generates a rescaled version of geodesic flow. Choosing $h : [0, \infty) \rightarrow \mathbb{R}$ so that $h'(0) = \pi$, $h'' \leq 0$, and h is constant outside of a neighborhood of zero, we obtain a Hamiltonian diffeomorphism of the complement of the zero section $T^*S^n \setminus S^n$, which can be extended across the zero section by the antipodal map on S^n to obtain a symplectomorphism of T^*S^n (see Figure 8).

Now, given a Lagrangian sphere S in a symplectic manifold (M, ω) , by Weinstein's theorem a neighborhood of S in M is symplectomorphic to a neighborhood of the zero section in T^*S^n ; thus, performing the above construction inside the standard

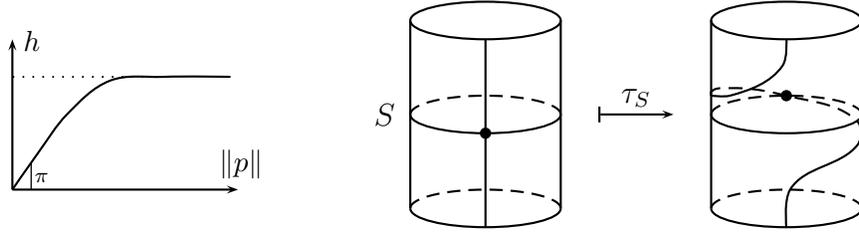


FIGURE 8. The generating Hamiltonian on the complement of the zero section in T^*S^n , and the action of the Dehn twist on a cotangent fiber.

neighborhood of S , we obtain a symplectomorphism τ_S , the *Dehn twist* about S , which is supported in a neighborhood of S and maps S to itself antipodally. (Note: τ_S depends on the choices made in the construction, but its isotopy class doesn't.)

Theorem 3.8 (Seidel [42, 40]). *Given a Lagrangian sphere S and any object L of $\mathcal{F}(M, \omega)$, there is an exact triangle in $\text{Tw } \mathcal{F}(M, \omega)$,*

$$(3.3) \quad \begin{array}{ccc} HF^*(S, L) \otimes S & \xrightarrow{ev} & L \\ & \searrow [1] & \downarrow \\ & & \tau_S(L) \end{array}$$

In other terms, the object $\tau_S(L)$ of $\mathcal{F}(M, \omega)$ is quasi-isomorphic in $\text{Tw } \mathcal{F}(M, \omega)$ to the abstract mapping cone of ev .

In (3.3), $HF^*(S, L) \otimes S$ is a direct sum of shifted copies of S , with one summand for each generator of $HF^*(S, L)$, and ev is a tautological evaluation map, mapping each summand to L by a closed morphism representing the given generator of $HF^*(S, L) = H^*\text{Hom}(S, L)$.

Given a test object T , the corresponding long exact sequence (3.1) is Seidel's long exact sequence in Floer cohomology [42] associated to the Dehn twist τ_S for all T, L :

$$\dots \rightarrow HF^*(S, L) \otimes HF^*(T, S) \xrightarrow{\mu^2} HF^*(T, L) \longrightarrow HF^*(T, \tau_S(L)) \xrightarrow{[1]} \dots$$

3.3.2. Lagrangian connected sums. Given two Lagrangian submanifolds L_1, L_2 which intersect transversely in a single point p , we can form the *Lagrangian connected sum* (or *surgery* in the terminology of [38] and [20]) $L_1 \# L_2$. One possible construction is as follows. For $\epsilon > 0$, the graph of the 1-form $\epsilon d \log \|x\|$ on \mathbb{R}^n , given by the equations $y_i = \epsilon x_i / \|x\|^2$, is a Lagrangian submanifold of $T^*\mathbb{R}^n \simeq \mathbb{C}^n$ which is asymptotic to the zero section (i.e., $\mathbb{R}^n \subset \mathbb{C}^n$) as $\|x\| \rightarrow \infty$ and to the cotangent fiber over zero (i.e., $(i\mathbb{R})^n \subset \mathbb{C}^n$) as $\|y\| \rightarrow \infty$; using suitable cut-off functions, we can modify this Lagrangian so that it agrees with $\mathbb{R}^n \cup (i\mathbb{R})^n$ outside of a small neighborhood of the origin. Pasting this local model into a suitable Darboux chart centered at the

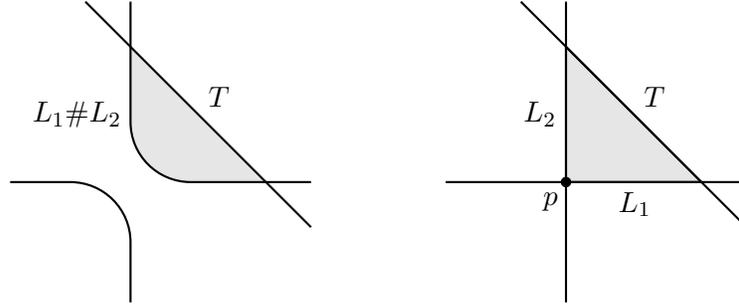


FIGURE 9. The Lagrangian connected sum $L_1\#L_2$ vs. $L_1 \cup L_2$

intersection point p and chosen so that $T_p L_1 = \mathbb{R}^n$ and $T_p L_2 = (i\mathbb{R})^n$ yields $L_1\#L_2$. (Note that, for a single connected sum operation, the end result is independent of the size parameter ϵ and other choices up to Hamiltonian isotopy; not so when summing at multiple points. Also note that $L_2\#L_1$ is *not* isotopic to $L_1\#L_2$.)

Remark 3.9. When L_2 is a sphere, $L_1\#L_2$ is Hamiltonian isotopic to $\tau_{L_2}(L_1)$; this provides the basis for an alternative description of the connected sum operation.

Given some other Lagrangian submanifold T (in generic position relatively to L_1 and L_2), choosing ϵ small enough in the above construction ensures that the intersections of T with $L_1\#L_2$ are the same as with $L_1 \cup L_2$. Fukaya-Oh-Ohta-Ono [20] have studied the moduli spaces of J -holomorphic discs bounded by $L_1\#L_2$ and T . Their main result is that, for suitable J and small enough ϵ , J -holomorphic strips with boundary on T and $L_1\#L_2$ connecting an intersection in $T \cap L_2$ to one in $T \cap L_1$ are in bijection with J -holomorphic triangles bounded by T , L_2 and L_1 with a corner at p , whereas the counts of rigid strips in the other direction vanish. This is elementary in dimension 1, as illustrated by Figure 9, but much harder in higher dimensions.

The outcome is that, as a chain complex, $CF(T, L_1\#L_2)$ is the mapping cone of the map $\mu^2(p, \cdot) : CF(T, L_2) \rightarrow CF(T, L_1)$ given by multiplication by the generator p of $CF(L_2, L_1)$. Hence, the short exact sequence

$$0 \rightarrow CF(T, L_1) \rightarrow CF(T, L_1\#L_2) \rightarrow CF(T, L_2) \rightarrow 0$$

induces a long exact sequence

$$\dots \rightarrow HF(T, L_1) \longrightarrow HF(T, L_1\#L_2) \longrightarrow HF(T, L_2) \xrightarrow{\mu^2(p, \cdot)} HF(T, L_1) \rightarrow \dots$$

By an analogous argument for higher structure maps, one expects that this long exact sequence can be upgraded to an exact triangle in the Fukaya category,

$$(3.4) \quad \begin{array}{ccc} L_2 & \xrightarrow{p} & L_1 \\ & \searrow [1] & \downarrow \\ & & L_1\#L_2 \end{array}$$

i.e., $L_1 \# L_2$ is quasi-isomorphic to the twisted complex $\text{Cone}(p) = \{L_2 \xrightarrow{p} L_1\}$.
 (If L_2 is a sphere, this is Seidel's exact triangle for the Dehn twist of L_1 about L_2 .)

Remark 3.10. Recall that, by definition, the differential μ_{Tw}^1 on $\text{hom}(T, \text{Cone}(p))$ involves not only the original Floer differential μ^1 , but also multiplication by the differential of the twisted complex, i.e. $\mu^2(p, \cdot)$. This is exactly consistent with the above description of J -holomorphic strips with boundary on T and $L_1 \# L_2$. Thus, replacing Lagrangian submanifolds by quasi-isomorphic twisted complexes built out of simpler Lagrangians, while computationally powerful, comes at the expense of having to consider higher operations on their Floer complexes (in this case, the expression for μ_{Tw}^1 involves μ^2 , and similarly that for μ_{Tw}^2 involves μ^3).

3.4. Generation and Yoneda embedding.

3.4.1. Generators and split-generators.

Definition 3.11. The objects G_1, \dots, G_r are said to generate the A_∞ -category \mathcal{A} if, in $\text{Tw } \mathcal{A}$, every object of \mathcal{A} is quasi-isomorphic to a twisted complex built from copies of G_1, \dots, G_r . (In other terms, every object of \mathcal{A} can be obtained from G_1, \dots, G_r by taking iterated mapping cones.)

The objects G_1, \dots, G_r are said to split-generate \mathcal{A} if every object of \mathcal{A} is quasi-isomorphic to a direct summand in a twisted complex built from copies of G_1, \dots, G_r .

Example 3.12. Consider the Fukaya category of the torus T^2 with its standard area form. Starting from the standard curves α and β along the two factors of the torus, by taking iterated mapping cones we can obtain simple closed curves representing all nontrivial primitive elements in $\pi_1(T^2) = \mathbb{Z}^2$. For instance, the loop $\tau_\alpha(\beta) \simeq \beta \# \alpha$ (Figure 10 left) is quasi-isomorphic to the mapping cone of $p \in \text{Hom}(\alpha, \beta)$; further applications of the Dehn twists τ_α and τ_β (which generate the mapping class group of T^2) eventually yield simple closed curves in all primitive homotopy classes. However, the objects obtained in this manner all satisfy a certain ‘‘balancing’’ condition: given a 1-form $\theta \in \Omega^1(T^2 \setminus \{pt\})$ with $d\theta = \omega$ and such that $\int_\alpha \theta = \int_\beta \theta = 0$, θ also integrates to zero on all iterated mapping cones built from α and β . For instance, all the

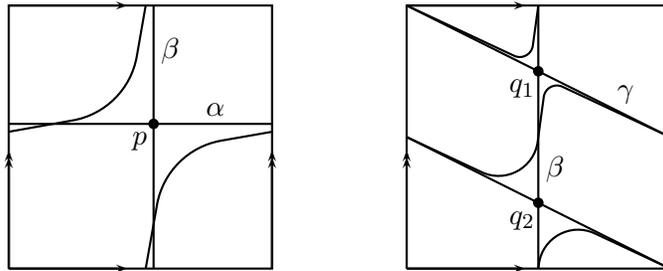


FIGURE 10. Split-generating the Fukaya category of T^2

simple closed curves that can be obtained in a given homotopy class are Hamiltonian isotopic to each other. Thus, α and β generate the subcategory of $\mathcal{F}(T^2)$ consisting of Lagrangians which are balanced with respect to θ , but not all of $\mathcal{F}(T^2)$.

On the other hand, given the two loops β and γ shown on Figure 10 right, the mapping cone of $T^{a_1}q_1 + T^{a_2}q_2 \in \text{Hom}(\gamma, \beta)$ can be interpreted geometrically as the connected sum of β and γ at their two intersection points q_1 and q_2 , with different gluing parameters. This mapping cone is therefore quasi-isomorphic to the direct sum of two simple closed curves in the homotopy class of α , but whose Hamiltonian isotopy classes depend on a_1 and a_2 . Thus, by considering direct summands in mapping cones we can obtain all nontrivial simple closed curves up to Hamiltonian isotopy, rather than only those that are balanced: α and β split-generate $\mathcal{F}(T^2)$.

3.4.2. *Yoneda embedding.* Let G_1, \dots, G_r be split-generators of the A_∞ -category \mathcal{A} . Then the endomorphism algebra of $G_1 \oplus \dots \oplus G_r$,

$$\mathcal{G} = \bigoplus_{i,j=1}^r \text{hom}(G_i, G_j)$$

is an A_∞ -algebra (with structure maps given by the operations μ^k of \mathcal{A}). Next, given any object L of \mathcal{A} ,

$$\mathcal{Y}(L) = \bigoplus_{i=1}^r \text{hom}(G_i, L)$$

is a (right) A_∞ -module over \mathcal{G} , with differential given by μ^1 , multiplication μ^{11} given by the operations

$$\text{hom}(G_j, L) \otimes \text{hom}(G_i, G_j) \xrightarrow{\mu^2} \text{hom}(G_i, L),$$

and so on (the structure map μ^{1k} of $\mathcal{Y}(L)$ is given by μ^{k+1}).

Moreover, to a morphism $a \in \text{hom}(L, L')$ we can associate an A_∞ -homomorphism $\mathcal{Y}(a) \in \text{hom}_{\text{mod-}\mathcal{G}}(\mathcal{Y}(L), \mathcal{Y}(L'))$, whose linear term is given by composition with a .

The assignment $L \mapsto \mathcal{Y}(L)$, $a \mapsto \mathcal{Y}(a)$ is in turn the linear term of an A_∞ -functor \mathcal{Y} , which is the restriction to the given set of objects G_1, \dots, G_r of the A_∞ *Yoneda embedding* $\mathcal{A} \rightarrow \text{mod-}\mathcal{A}$ (see e.g. [40, §1]):

Proposition 3.13. *The above construction extends to an A_∞ -functor \mathcal{Y} from \mathcal{A} to $\text{mod-}\mathcal{G}$. Moreover, if G_1, \dots, G_r split-generate \mathcal{A} then this A_∞ -functor is a fully faithful quasi-embedding.*

4. THE WRAPPED FUKAYA CATEGORY, EXAMPLES AND APPLICATIONS

In this section we assume that (M, ω) is a *Liouville manifold*, i.e. an exact symplectic manifold such that the Liouville vector field Z associated to the chosen primitive $\theta \in \Omega^1(M)$ of the symplectic form (i.e., the conformally symplectic vector field defined

by $\iota_Z\omega = \theta$) is complete and outward pointing at infinity. More precisely, we require that M contains a compact domain M^{in} with boundary a smooth hypersurface ∂M on which $\alpha = \theta|_{\partial M}$ is a contact form, and Z is positively transverse to ∂M and has no zeroes outside of M^{in} . The flow of Z can then be used to identify $M \setminus M^{in}$ with the positive symplectization $(1, \infty) \times \partial M$ equipped with the exact symplectic form $\omega = d(r\alpha)$ and the Liouville field $Z = r\frac{\partial}{\partial r}$.

In this setting it is natural to consider not only compact exact Lagrangian submanifolds as we have done above, but also some noncompact ones with suitable behavior at infinity. There are two different types of such noncompact Fukaya categories, depending on the manner in which perturbations at infinity are used to define Floer complexes. One possibility is to perform “small” perturbations at infinity, restricting oneself to a smaller set of “admissible” objects which go to infinity along well-controlled directions. Two constructions that follow this philosophy are the “infinitesimal” Fukaya category first defined by Nadler and Zaslow for cotangent bundles [31] and later extended to Liouville manifolds equipped with a choice of Lagrangian skeleton; and Fukaya categories of Lefschetz fibrations as constructed by Seidel [40, 43], and their putative generalization to Landau-Ginzburg models, in which the behavior at infinity is controlled by a projection to the complex plane. Here we focus on the other approach, which is to consider *large* perturbations at infinity, leading to the *wrapped Fukaya category* of Abouzaid and Seidel [11, 3]. For completeness we mention the nascent subject of *partially wrapped* Fukaya categories, which attempt to interpolate between these two approaches (cf. e.g. [13]).

4.1. The wrapped Fukaya category. The objects of the wrapped Fukaya category $\mathcal{W}(M)$ of a Liouville manifold $(M, \omega = d\theta)$ are exact Lagrangian submanifolds $L \subset M$ which are conical at infinity, i.e. invariant under the flow of the Liouville vector field outside of a compact subset, and such that the exact 1-form $\theta|_L$ vanishes outside of a compact set. In other terms, if L is noncompact then at infinity it must coincide with the cone $(1, \infty) \times \partial L$ over some Legendrian submanifold ∂L of ∂M .

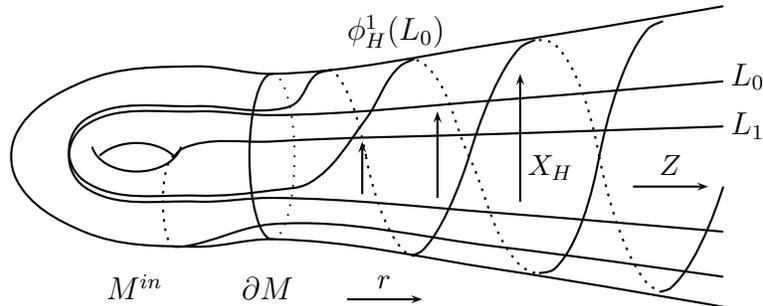


FIGURE 11. Wrapping by a quadratic Hamiltonian

The Hamiltonian perturbations used to define Floer complexes in the wrapped setting are very specific: namely, we only consider Hamiltonians $H : M \rightarrow \mathbb{R}$ which, outside of a compact subset of M , satisfy $H = r^2$ where $r \in (1, \infty)$ is the radial coordinate of the symplectization $(1, \infty) \times \partial M$. Thus, outside of a compact set the Hamiltonian vector field X_H is equal to $2r$ times the Reeb vector field R_α of the contact form α on ∂M .

Given two objects L_0, L_1 , the generating set $\mathcal{X}(L_0, L_1)$ of the *wrapped Floer complex* $CW(L_0, L_1) = CW(L_0, L_1; H)$ consists of time 1 trajectories of the flow of X_H which start on L_0 and end on L_1 , i.e. points of $\phi_H^1(L_0) \cap L_1$. More concretely, these consist of (perturbed) intersections between L_0 and L_1 in the interior M^{in} on one hand, and Reeb chords (of arbitrary length) from ∂L_0 to ∂L_1 on the other hand (see Figure 11). Thus, wrapped Floer cohomology is closely related to Legendrian contact homology. (Of course, we need to assume that $\phi_H^1(L_0)$ intersects L_1 transversely, and in particular that the Reeb chords from ∂L_0 to ∂L_1 are non-degenerate; otherwise a small modification of H is required.)

The differential on the wrapped Floer complex counts solutions to Floer's equation (1.6), i.e. perturbed J -holomorphic strips with boundary on L_0 and L_1 , as in §1. (Note: due to exactness we can work directly over the field \mathbb{K} , without resorting to Novikov coefficients.) As in Remark 1.10, these can equivalently be viewed as $(\phi_H^{1-t})_* J$ -holomorphic strips with boundary on $\phi_H^1(L_0)$ and L_1 . The assumptions made on the objects of $\mathcal{W}(M)$ and on the Hamiltonian H ensure that, for suitably chosen J , perturbed J -holomorphic strips are well-behaved: an *a priori* energy estimate ensures that all solutions of (1.6) which converge to a given generator $p \in \mathcal{X}(L_0, L_1)$ as $s \rightarrow +\infty$ remain within a bounded subset of M (see e.g. [3]). Thus, ∂p is a finite linear combination of generators of the wrapped Floer complex.

A subtlety comes up when we attempt to define the product operation on wrapped Floer complexes,

$$(4.1) \quad CW(L_1, L_2; H) \otimes CW(L_0, L_1; H) \rightarrow CW(L_0, L_2; H).$$

For the perturbed Cauchy-Riemann equation (2.4) to be well-behaved and satisfy *a priori* energy estimates in spite of the non-compactness of M , one needs the 1-form β that appears in the perturbation term $X_H \otimes \beta$ to satisfy $d\beta \leq 0$ (cf. [11, 3]). In other terms, the naturally defined product map would take values in $CW(L_0, L_2; 2H)$, and the usual continuation map from this complex to $CW(L_0, L_2; H)$ fails to be well-defined. This can be remedied using the following rescaling trick alluded to in [21] and systematically developed in [3].

Recall that the flow of the Liouville vector field is conformally symplectic and, in the symplectization $(1, \infty) \times \partial M$ where $Z = r \frac{\partial}{\partial r}$, simply amounts to rescaling in the r direction. For $\rho > 1$, denote by ψ^ρ the time $\log \rho$ flow of Z , which rescales r by a factor of ρ . Then there is a natural isomorphism

$$(4.2) \quad CW(L_0, L_1; H, J) \cong CW(\psi^\rho(L_0), \psi^\rho(L_1); \rho^{-1}H \circ \psi^\rho, \psi_*^\rho J).$$

Moreover, our assumptions imply that $\psi^\rho(L_i)$ is exact Lagrangian isotopic to L_i by a compactly supported isotopy, and $\rho^{-1}H \circ \psi^\rho$ coincides with ρH at infinity. Abouzaid shows that these properties ensure the existence of a well-defined product map

$$(4.3) \quad CW(L_1, L_2; H, J) \otimes CW(L_0, L_1; H, J) \rightarrow CW(\psi^2(L_0), \psi^2(L_2); \frac{1}{2}H \circ \psi^2, \psi_*^2 J),$$

determined by counts of index 0 finite energy maps $u : D \rightarrow M$ from a disc with three strip-like ends to M , mapping the three components of ∂D to the images of the respective Lagrangians under suitable Liouville rescalings, and solving the perturbed Cauchy-Riemann equation

$$(du - X_{\tilde{H}} \otimes \beta)_{\tilde{J}}^{0,1} = 0,$$

where β is a closed 1-form on D with $\beta|_{\partial D} = 0$ which is standard in the strip-like ends (modelled on dt for the input ends, $2dt$ for the output end), and \tilde{H} and \tilde{J} are obtained from H and J by suitable rescalings ($\tilde{H} = H$ and $\tilde{J} = J$ near the input punctures; $\tilde{H} = \frac{1}{4}H \circ \psi^2$ and $\tilde{J} = \psi_*^2 J$ near the output puncture; see [3]). The map (4.3), composed with the isomorphism (4.2), yields the desired product map (4.1). The higher products

$$\mu^k : CW(L_{k-1}, L_k; H) \otimes \cdots \otimes CW(L_0, L_1; H) \rightarrow CW(L_0, L_k; H)$$

are constructed in the same manner [3]. These structure maps make $\mathcal{W}(M)$ an A_∞ -category, the wrapped Fukaya category of the Liouville manifold M .

Remark 4.1. The rescaling trick can be informally understood as follows. As mentioned above, the naturally defined product map on wrapped Floer complexes takes values in $CW(L_0, L_2; 2H)$; while the usual construction of a continuation map cannot be used to map this complex to $CW(L_0, L_2; H)$, the fact that $\frac{1}{2}H \circ \psi^2 = 2H$ at infinity and the assumptions made on L_0 and L_2 imply that there is a well-defined continuation map to $CW(\psi^2(L_0), \psi^2(L_2); \frac{1}{2}H \circ \psi^2)$, which by (4.2) is isomorphic to $CW(L_0, L_2; H)$. (Note: while this is a slightly simpler way to describe the cohomology-level product, it lacks the compatibility and consistency features needed to construct the chain-level A_∞ -structure, hence the slightly more complicated construction in [3]).

Remark 4.2. Since compact exact Lagrangian submanifolds of M^{in} are not affected by the wrapping at infinity, $\mathcal{W}(M)$ contains the ordinary Fukaya category (of compact exact Lagrangian submanifolds) as a full A_∞ -subcategory.

4.2. An example. Let $M = T^*S^1 = \mathbb{R} \times S^1$, equipped with the standard Liouville form $r d\theta$ and the wrapping Hamiltonian $H = r^2$, and consider the exact Lagrangian submanifold $L = \mathbb{R} \times \{pt\}$. We can label the intersection points of $\phi_H^1(L)$ with L by integers, $\mathcal{X}(L, L) = \{x_i, i \in \mathbb{Z}\}$, in increasing order along the real axis, where x_0 is the intersection occurring at the minimum of H ; in other terms, x_0 is an interior intersection of L with a small pushoff of it, while the other generators correspond to

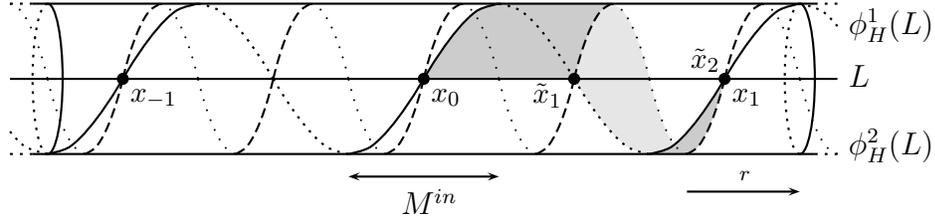


FIGURE 12. The wrapped Floer cohomology of $L = \mathbb{R} \times \{pt\}$ in $\mathbb{R} \times S^1$

Reeb chords from $\partial L = \{pt\} \sqcup \{pt\}$ to itself in the contact manifold $\partial M = S^1 \sqcup S^1$ (see Figure 12).

Recall that the differential on $CW(L, L)$ counts rigid pseudo-holomorphic strips (for a t -dependent almost-complex structure) with boundary on L and $\phi_H^1(L)$. Since there are no such strips (see Figure 12), the Floer differential on $CW(L, L)$ vanishes identically, and $HW(L, L) \simeq CW(L, L) = \text{span} \{x_i, i \in \mathbb{Z}\}$. (This can also be seen by observing that all generators of $CW(L, L)$ have degree 0 for the natural \mathbb{Z} -grading.)

The product structure on $CW(L, L)$ counts perturbed pseudo-holomorphic discs with three strip-like ends, as explained above; in the present case, L is invariant under the Liouville flow $\psi^\rho : (r, \theta) \mapsto (\rho r, \theta)$, while $H \circ \psi^\rho = \rho^2 H$. Thus, the rescaling trick only affects the almost-complex structure (i.e., ψ^2 intertwines $CW(L, L; H, J)$ and $CW(L, L; 2H, \psi_*^2 J)$), and otherwise simply amounts to identifying $\mathcal{X}(L, L; 2H) = \phi_H^2(L) \cap L$ with $\mathcal{X}(L, L; H) = \phi_H^1(L) \cap L$ via the radial rescaling $r \mapsto 2r$.

Proceeding as in Remark 1.10, the perturbed pseudo-holomorphic discs with boundary on L which determine the product on $CW(L, L)$ can then be reinterpreted as genuine pseudo-holomorphic discs (with respect to a modified family of almost-complex structures) with boundaries on $\phi_H^2(L)$, $\phi_H^1(L)$ and L . Specifically, the coefficient of a generator $q \in \mathcal{X}(L, L)$ in the product $p_2 \cdot p_1$ of two generators $p_1, p_2 \in \mathcal{X}(L, L)$ is given by a count of index 0 pseudo-holomorphic discs with boundaries on $\phi_H^2(L)$, $\phi_H^1(L)$ and L , and with strip-like ends converging to the intersection points $\phi_H^1(p_1) \in \phi_H^2(L) \cap \phi_H^1(L)$, $p_2 \in \phi_H^1(L) \cap L$, and $\tilde{q} \in \phi_H^2(L) \cap L$, where \tilde{q} corresponds to $q \in \phi_H^1(L) \cap L$ under the Liouville rescaling.

With this understood, the product structure can be determined directly by looking at Figure 12. Observe that any two input intersections $\phi_H^1(x_i) \in \phi_H^2(L) \cap \phi_H^1(L)$ and $x_j \in \phi_H^1(L) \cap L$ are the vertices of a unique immersed triangle, whose third vertex is $x_{i+j} \in \phi_H^2(L) \cap L$. (This is easiest to see by lifting the diagram of Figure 12 to the universal cover of M .) These triangles are all regular, and we conclude that

$$x_j \cdot x_i = x_{i+j}.$$

(Recall that thanks to exactness we are working over \mathbb{K} and not keeping track of symplectic areas.) For example, the triangle shaded in Figure 12 illustrates the identity $x_0 \cdot x_1 = x_1$. In other terms, renaming the generator x_i to x^i , we have a ring

isomorphism

$$(4.4) \quad CW(L, L) \simeq \mathbb{K}[x, x^{-1}].$$

Furthermore, the higher products on $CW(L, L)$ are all identically zero, as can be checked either by drawing the successive images of L under the wrapping flow and looking for rigid holomorphic polygons (there are none), or more directly by recalling that $\deg(x^i) = 0$ for all $i \in \mathbb{Z}$ whereas $\deg(\mu^k) = 2 - k$. Thus (4.4) is in fact an isomorphism of A_∞ -algebras.

4.3. Cotangent bundles. The previous example is the simplest case of a general result about cotangent bundles. Let N be a compact spin manifold, and let $M = T^*N$ equipped with its standard Liouville form $p dq$ and the wrapping Hamiltonian $H = \|p\|^2$ (for some choice of Riemannian metric on N). Then we have:

Theorem 4.3 (Abouzaid [4]). *Let $L = T_q^*N$, the cotangent fiber at some point $q \in N$. Then there is a quasi-isomorphism of A_∞ -algebras*

$$(4.5) \quad CW^*(L, L) \simeq C_{-*}(\Omega_q N)$$

*between the wrapped Floer complex of $L = T_q^*N$ and chains on the based loop space $\Omega_q N$ equipped with (an A_∞ -refinement of) the usual Pontryagin product.*

(The corresponding statement for cohomology is an earlier result of Abbondandolo and Schwarz [2].)

For instance, in the case of $N = S^1$, the based loop space $\Omega_q S^1$ has countably many components, each of which is contractible, thus $\Omega_q S^1 \sim \mathbb{Z}$, and (4.5) reduces to (4.4). In fact, the assumption that N is spin can be removed; in that case, $CW^*(L, L)$ is related to chains on $\Omega_q N$ twisted by the \mathbb{Z} -local system determined by $w_2(N)$ [4].

Furthermore, Abouzaid has shown that the fiber $L = T_q^*N$ generates the wrapped Fukaya category $\mathcal{W}(T^*N)$ [5]. Using Yoneda embedding (cf. §3.4.2), we conclude:

Corollary 4.4 (Abouzaid). *The wrapped Fukaya category $\mathcal{W}(T^*N)$ quasi-embeds fully faithfully into the category of A_∞ -modules over $C_{-*}(\Omega_q N)$.*

(Here again, when N is not spin a twist by a suitable local system is required.)

This and other related results can be viewed as the culmination of over a decade of investigations of the deep connections between the symplectic topology of T^*N and the algebraic topology of the loop space of N , as previously studied by Viterbo [48], Salamon-Weber [39], Abbondandolo-Schwarz [1, 2], Cieliebak-Latschev [16], etc.

At the same time, studying Fukaya categories of cotangent bundles has led to much progress on Arnold's conjecture on exact Lagrangian submanifolds:

Conjecture 4.5 (Arnold). *Let N be a compact closed manifold: then any compact closed exact Lagrangian submanifold of T^*N (with its standard Liouville form) is Hamiltonian isotopic to the zero section.*

Theorem 4.6 (Fukaya-Seidel-Smith [21], Nadler-Zaslow [31], Abouzaid [7], Kragh [26]). *Let L be a compact connected exact Lagrangian submanifold of T^*N . Then as an object of $\mathcal{W}(T^*N)$, L is quasi-isomorphic to the zero section, and the restriction of the bundle projection $\pi|_L : L \rightarrow N$ is a homotopy equivalence.*

Abouzaid has further shown that Floer theory detects more than purely topological information about exact Lagrangians in cotangent bundles: certain exotic spheres (in dimensions ≥ 9) do not admit Lagrangian embeddings into T^*S^{4k+1} [8].

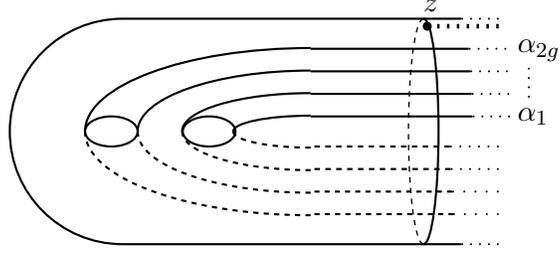
However, in spite of all the recent progress, Conjecture 4.5 appears to remain out of reach of current technology.

4.4. Homological mirror symmetry. Kontsevich's homological mirror symmetry conjecture [24] asserts that the main manifestation of the phenomenon of mirror symmetry is as a derived equivalence between the Fukaya category of a symplectic manifold and the category of coherent sheaves of its mirror. While this conjecture was initially stated for compact Calabi-Yau manifolds (and recently proved for the quintic 3-fold by Sheridan [46]), it also holds (and is often easier to prove) for non-compact manifolds (in which case one should consider the *wrapped* Fukaya category), and outside of the Calabi-Yau case (in which case the mirror is a *Landau-Ginzburg model*, for which one should consider Orlov's derived category of singularities [35, 36] rather than the ordinary derived category of coherent sheaves).

The calculation we have performed in §4.2, together with Abouzaid's generation statement, essentially proves homological mirror symmetry for the cylinder $\mathbb{C}^* = T^*S^1$, and its mirror $\mathbb{C}^* = \text{Spec } \mathbb{C}[x^{\pm 1}]$. Namely, coherent sheaves over \mathbb{C}^* are the same thing as finite rank $\mathbb{C}[x^{\pm 1}]$ -modules. However, since the object L considered in §4.2 generates the wrapped Fukaya category, $\mathcal{W}(T^*S^1)$ quasi-embeds into the category of modules over $CW(L, L) \simeq \mathbb{C}[x^{\pm 1}]$, and the image can be characterized explicitly enough to prove the desired equivalence between $\mathcal{W}(T^*S^1)$ and $D^b\text{Coh}(\mathbb{C}^*)$.

This general approach extends to other examples, with the caveat that in general there are infinitely many non-trivial higher A_∞ -operations; one then needs to rely on an algebraic classification result in order to determine which structure coefficients need to be computed in order to fully determine the A_∞ -structure up to homotopy. Symplectic manifolds whose Fukaya categories have been determined in this manner include (but are not limited to) pairs of pants [9], genus 2 curves [44], and Calabi-Yau hypersurfaces in projective space [46].

4.5. An application to Heegaard-Floer homology. Heegaard-Floer homology associates to a closed 3-manifold Y a graded abelian group $\widehat{HF}(Y)$. This invariant is constructed by considering a Heegaard splitting $Y = Y_1 \cup_{\bar{\Sigma}} Y_2$ of Y into two genus g handlebodies Y_i , each of which determines a product torus T_i in the g -fold symmetric product of the Heegaard surface $\bar{\Sigma} = \partial Y_1 = -\partial Y_2$. Deleting a marked point z from

FIGURE 13. Generating $\mathcal{W}(\mathrm{Sym}^g(\Sigma))$

$\bar{\Sigma}$ to obtain an open surface Σ , $\widehat{HF}(Y)$ is then defined as the Floer cohomology of the Lagrangian tori T_1, T_2 in the symplectic manifold $\mathrm{Sym}^g(\Sigma)$, see [37].

In this context it is natural to study the Fukaya category (ordinary or wrapped) of $\mathrm{Sym}^g(\Sigma)$ (equipped with a Kähler form which agrees with the product one away from the diagonal). It turns out that the wrapped category has a particularly nice set of generators. Namely, consider a collection of $2g$ disjoint properly embedded arcs $\alpha_1, \dots, \alpha_{2g}$ in Σ such that $\Sigma \setminus (\alpha_1 \cup \dots \cup \alpha_{2g})$ is homeomorphic to a disc, see e.g. Figure 13. Given a g -element subset $s \subseteq \{1, \dots, 2g\}$, the product $D_s = \prod_{i \in s} \alpha_i$ is an exact Lagrangian submanifold of $\mathrm{Sym}^g(\Sigma)$, and we have:

Theorem 4.7. [13, 14] *The Lagrangian submanifolds $D_s = \prod_{i \in s} \alpha_i$, $s \subseteq \{1, \dots, 2g\}$, $|s| = g$ generate $\mathcal{W}(\mathrm{Sym}^g(\Sigma))$.*

Thus, by Yoneda embedding, Lagrangian submanifolds of $\mathrm{Sym}^g(\Sigma)$ can be viewed as modules over the A_∞ -algebra $\bigoplus_{s, s'} \mathrm{hom}(D_s, D_{s'})$.

Determining this A_∞ -algebra is not completely hopeless, as the wrapping Hamiltonian H on $\mathrm{Sym}^g(\Sigma)$ can be chosen in a manner compatible with the product structure so that $\phi_H^1(D_s) = \prod_{i \in s} \phi_h^1(\alpha_i)$, where h is a Hamiltonian on Σ that grows quadratically in the cylindrical end, and pseudo-holomorphic discs in the symmetric product can be viewed by projecting them to Σ as is customary in Heegaard-Floer theory; nonetheless, things are complicated by the presence of many nontrivial A_∞ -products.

It is easier to study a *partially wrapped* version of the Fukaya category, in which the wrapping “stops” along a ray $\{z\} \times (1, \infty)$ in the cylindrical end of Σ ; i.e., the Hamiltonian is again chosen to be compatible with the product structure away from the diagonal, but the effect on each component is to push the ends of the arc α_i in the positive direction towards the ray $\{z\} \times (1, \infty)$, without ever crossing it: see [13]. Theorem 4.7 continues to hold in this setting: the product Lagrangians D_s also generate the partially wrapped Fukaya category. Furthermore, in the partially wrapped case the A_∞ -algebra $\mathcal{A} = \bigoplus_{s, s'} \mathrm{hom}(D_s, D_{s'})$ turns out to be a finite-dimensional dg-algebra (i.e., $\mu^k = 0$ for $k \geq 3$) which admits a simple explicit combinatorial description [13]; in fact, \mathcal{A} is precisely the *strands algebra* first introduced by Lipshitz, Ozsváth and Thurston [27].

By Yoneda embedding, Lagrangian submanifolds of $\mathrm{Sym}^g(\Sigma)$, such as the product tori associated to genus g handlebodies in Heegaard-Floer theory, can be viewed as A_∞ -modules over the strands algebra. Moreover, the same holds true for *generalized Lagrangian submanifolds* of $\mathrm{Sym}^g(\Sigma)$ (i.e., formal images of Lagrangian submanifolds under sequences of Lagrangian correspondences, cf. [49]), such as those associated to arbitrary 3-manifolds with boundary $\bar{\Sigma}$ (not just handlebodies) according to ongoing work of Lekili and Perutz. This provides a symplectic geometry interpretation of Lipshitz-Ozsváth-Thurston's *bordered Heegaard-Floer homology* [27], which associates to a 3-manifold Y with boundary $\partial Y = \bar{\Sigma}$ an A_∞ -module $\widehat{CFA}(Y)$ over the strands algebra. Namely, Lekili and Perutz's construction associates to such a 3-manifold a generalized Lagrangian submanifold of $\mathrm{Sym}^g(\Sigma)$, whose image under Yoneda embedding (as in §3.4.2, but using *quilted Floer cohomology* of Lagrangian correspondences) is the A_∞ -module $\widehat{CFA}(Y)$; see [13, 14].

4.6. A closing remark. The methods available to calculate Floer cohomology and Fukaya categories are still evolving rapidly. Besides the use of algebraic generation statements such as those in [3] and [40] to reduce to a simpler set of Lagrangian submanifolds, there are at least two key ideas that have made calculations possible.

On one hand, it is often possible to find holomorphic projection maps (to the complex plane or to other Riemann surfaces) under which the given Lagrangians project to arcs or curves, in which case holomorphic discs can be studied by looking at their projections to the base and by reducing to the symplectic geometry of the fiber; this is e.g. the guiding principle of Seidel's work on Lefschetz fibrations [40, 43] and the various calculations done using that framework.

At the same time, since such holomorphic projections are easier to come by on open manifolds, another idea that nicely complements this one is to carry out calculations for an exact open subdomain M^0 of the given symplectic manifold M obtained by deleting some complex hypersurface, and then use abstract deformation theory to view the Fukaya category of M as an A_∞ -deformation of that of M^0 (cf. [41]). The Hochschild cohomology class that determines the deformation is then often determined by symmetry considerations and/or by studying specific A_∞ -structure maps (i.e., certain counts of holomorphic discs in M). See e.g. [44, 46] for an illustration of this approach. (One guiding principle which might explain why this approach is so successful is that algebraic deformations of Fukaya categories are often *geometric*: natural “closed-open” maps from the quantum or symplectic cohomology of M to the Hochschild cohomology of its ordinary or wrapped Fukaya category often turn out to be isomorphisms [10, 22].)

Going forward, there is hope that sheaf-theoretic methods will lead to completely new methods of computation of Fukaya categories (at least for Liouville manifolds) in terms of the topology of a Lagrangian “skeleton”. This is an idea that to our

knowledge originated with Kontsevich [25], and was subsequently developed by various other authors (see e.g. [45, 6, 47, 30]); the ultimate goal being to bypass the analysis of pseudo-holomorphic curves in favor of algebraic and topological methods. It is too early to tell how successful these approaches will be, but it is entirely possible that they will ultimately supplant the techniques we have described in this text.

REFERENCES

- [1] A. Abbondandolo, M. Schwarz, *On the Floer homology of cotangent bundles*, Comm. Pure Appl. Math. **59** (2006), 254–316.
- [2] A. Abbondandolo, M. Schwarz, *Floer homology of cotangent bundles and the loop product*, Geom. Topol. **14** (2010), 1569–1722.
- [3] M. Abouzaid, *A geometric criterion for generating the Fukaya category*, Publ. Math. IHÉS **112** (2010), 191–240.
- [4] M. Abouzaid, *On the wrapped Fukaya category and based loops*, J. Symplectic Geom. **10** (2012), 27–79.
- [5] M. Abouzaid, *A cotangent fibre generates the Fukaya category*, Adv. Math. **228** (2011), 894–939.
- [6] M. Abouzaid, *A topological model for the Fukaya categories of plumbings*, J. Differential Geom. **87** (2011), 1–80.
- [7] M. Abouzaid, *Nearby Lagrangians with vanishing Maslov class are homotopy equivalent*, Invent. Math. **189** (2012), 251–313.
- [8] M. Abouzaid, *Framed bordism and Lagrangian embeddings of exotic spheres*, Ann. Math. **175** (2012), 71–185.
- [9] M. Abouzaid, D. Auroux, A. I. Efimov, L. Katzarkov, D. Orlov, *Homological mirror symmetry for punctured spheres*, arXiv:1103.4322.
- [10] M. Abouzaid, K. Fukaya, Y.-G. Oh, H. Ohta, K. Ono, in preparation.
- [11] M. Abouzaid, P. Seidel, *An open string analogue of Viterbo functoriality*, Geom. Topol. **14** (2010), 627–718.
- [12] V. Arnol’d, *Sur une propriété topologique des applications globalement canoniques de la mécanique classique*, C. R. Acad. Sci. Paris **261** (1965), 3719–3722.
- [13] D. Auroux, *Fukaya categories of symmetric products and bordered Heegaard-Floer homology*, J. Gökova Geom. Topol. **4** (2010), 1–54.
- [14] D. Auroux, *Fukaya categories and bordered Heegaard-Floer homology*, Proc. International Congress of Mathematicians (Hyderabad, 2010), Vol. II, Hindustan Book Agency, 2010, 917–941.
- [15] P. Biran, O. Cornea, *A Lagrangian quantum homology*, New Perspectives and Challenges in Symplectic Field Theory, CRM Proc. Lecture Notes **49**, Amer. Math. Soc., Providence, 2009, pp. 1–44.
- [16] K. Cieliebak, J. Latschev, *The role of string topology in symplectic field theory*, New Perspectives and Challenges in Symplectic Field Theory, CRM Proc. Lecture Notes **49**, Amer. Math. Soc., Providence, 2009, pp. 113–146.
- [17] A. Floer, *Morse theory for Lagrangian intersections*, J. Differential Geom. **28** (1988), 513–547.
- [18] A. Floer, *Witten’s complex and infinite-dimensional Morse theory*, J. Differential Geom. **30** (1989), 207–221.
- [19] K. Fukaya, Y.-G. Oh, H. Ohta, K. Ono, *Lagrangian intersection Floer theory: anomaly and obstruction I and II*, AMS/IP Studies in Advanced Math., vol. 46, Amer. Math. Soc. and International Press, 2009.

- [20] K. Fukaya, Y.-G. Oh, H. Ohta, K. Ono, *Lagrangian surgery and holomorphic discs*, Chapter 10 of preprint version of [19].
- [21] K. Fukaya, P. Seidel, I. Smith, *The symplectic geometry of cotangent bundles from a categorical viewpoint*, Homological Mirror Symmetry, Lecture Notes in Phys. **757**, Springer, Berlin, 2009, pp. 1–26.
- [22] S. Ganatra, *Symplectic cohomology and duality for the wrapped Fukaya category*, Ph.D. thesis, MIT, 2012.
- [23] H. Hofer, *Polyfolds and a general Fredholm theory*, arXiv:0809.3753.
- [24] M. Kontsevich, *Homological algebra of mirror symmetry*, Proc. International Congress of Mathematicians (Zürich, 1994), Birkhäuser, Basel, 1995, pp. 120–139.
- [25] M. Kontsevich, *Symplectic geometry of homological algebra*, preprint, 2009.
- [26] T. Kragh, *Parametrized ring-spectra and the nearby Lagrangian conjecture*, arXiv:1107.4674.
- [27] R. Lipshitz, P. Ozsváth, D. Thurston, *Bordered Heegaard Floer homology: invariance and pairing*, arXiv:0810.0687.
- [28] D. McDuff, D. Salamon, *Introduction to symplectic topology*, Oxford University Press, New York, 1998.
- [29] D. McDuff, D. Salamon, *J-holomorphic curves and symplectic topology*, Amer. Math. Soc. Colloq. Publ. **52**, Amer Math Soc., Providence, 2004.
- [30] D. Nadler, *Fukaya categories as categorical Morse homology*, arXiv:1109.4848.
- [31] D. Nadler, E. Zaslow, *Constructible sheaves and the Fukaya category*, J. Amer. Math. Soc. **22** (2009), 233–286.
- [32] Y.-G. Oh, *Floer cohomology of Lagrangian intersections and pseudo-holomorphic disks I and II*, Comm. Pure Appl. Math. **46** (1993), 949–993 and 995–1012.
- [33] Y.-G. Oh, *Floer cohomology, spectral sequences, and the Maslov class of Lagrangian embeddings*, Internat. Math. Res. Notices (1996), 305–346.
- [34] Y.-G. Oh, *Relative Floer and quantum cohomology and the symplectic topology of Lagrangian submanifolds*, Contact and Symplectic Geometry (Cambridge, 1994), Publ. Newton Inst. **8**, Cambridge Univ. Press, Cambridge, 1996, pp. 201–267.
- [35] D. Orlov, *Triangulated categories of singularities and D-branes in Landau-Ginzburg models*, Proc. Steklov Inst. Math. **246** (2004), 227–249.
- [36] D. Orlov, *Matrix factorizations for nonaffine LG-models*, Math. Ann. **353** (2012), 95–108.
- [37] P. Ozsváth, Z. Szabó, *Holomorphic disks and topological invariants for closed three-manifolds*, Ann. of Math. (2) **159** (2004), 1027–1158.
- [38] L. Polterovich, *The surgery of Lagrange submanifolds*, Geom. Funct. Anal. **1** (1991), 198–210.
- [39] D. Salamon, J. Weber, *Floer homology and the heat flow*, Geom. Funct. Anal. **16** (2006), 1050–1138.
- [40] P. Seidel, *Fukaya categories and Picard-Lefschetz theory*, Zurich Lect. in Adv. Math., European Math. Soc., Zürich, 2008.
- [41] P. Seidel, *Fukaya categories and deformations*, Proc. International Congress of Mathematicians, Vol. II (Beijing, 2002), Higher Ed. Press, Beijing, 2002, pp. 351–360.
- [42] P. Seidel, *A long exact sequence for symplectic Floer cohomology*, Topology **42** (2003), 1003–1063.
- [43] P. Seidel, *Fukaya A_∞ structures associated to Lefschetz fibrations I*, arXiv:0912.3932.
- [44] P. Seidel, *Homological mirror symmetry for the genus two curve*, J. Algebraic Geometry **20** (2011), 727–769.
- [45] P. Seidel, *Some speculations on pairs of pants decompositions and Fukaya categories*, arXiv:1004.0906.

- [46] N. Sheridan, *Homological mirror symmetry for Calabi-Yau hypersurfaces in projective space*, arXiv:1111.0632.
- [47] N. Sibilla, D. Treumann, E. Zaslow, *Ribbon graphs and mirror symmetry I*, arXiv:1103.2462.
- [48] C. Viterbo, *Functors and computations in Floer homology with applications II*, preprint, 1996.
- [49] K. Wehrheim, C. Woodward, *Functoriality for Lagrangian correspondences in Floer theory*, *Quantum Topol.* **1** (2010), 129–170.

DEPARTMENT OF MATHEMATICS, UC BERKELEY, BERKELEY CA 94720-3840, USA

E-mail address: `auroux@math.berkeley.edu`