EXERCISES FOR MINI-COURSE ON THE FUKAYA CATEGORY

MAKSIM MAYDANSKIY AND NICK SHERIDAN

1. COMBINATORIAL FLOER HOMOLOGY

Let $L_0, L_1 \subset \Sigma$ be transversely intersecting, simple, closed, homotopically non-trivial curves in a symplectic 2-manifold (Σ, ω).

(1) Give a purely combinatorial definition of the Floer cochain complex $(CF^\bullet(L_0, L_1), \partial)$ over $\Lambda_{\mathbb{Z}_2}$. In particular, prove that $\partial^2 = 0$.
(2) Now do the same, but over $\Lambda_{\mathbb{K}}$ where $\mathbb{K}$ is an arbitrary field (not necessarily of characteristic 2).
(3) How far can the conditions on the L_i be relaxed? Can they be immersed curves, or nullhomotopic curves?
(4) More generally, let $L_0, \ldots, L_k$ be transversely-intersecting curves as above. Give a purely combinatorial definition of the A_∞ structure map
$$m^k : CF(L_{k-1}, L_k) \otimes \ldots \otimes CF(L_0, L_1) \to CF(L_0, L_k),$$
and prove that the maps defined in this way satisfy the A_∞ relations (start by working over a Novikov field of characteristic 2 as before, then generalize to arbitrary characteristic).

References for this section include [Che02], [dSRS14], [Sei08, §13b].

We will consider gradings in this case in §4.

2. THE ACTION FUNCTIONAL (SEE [Auroux, §1.1])

Let L_0, L_1 be Lagrangian submanifolds of (M, ω). Define the path space
$$\mathcal{P} := \{\gamma \in C^\infty([0, 1], M) : \gamma(0) \in L_0, \gamma(1) \in L_1\}.$$
We can think of this as an infinite-dimensional manifold with tangent space
$$T_\gamma \mathcal{P} = \{\eta \in \Gamma(\gamma^* TM) : \eta(0) \in T_{\gamma(0)} L_0, \eta(1) \in T_{\gamma(1)} L_1\}.$$

If J is an almost-complex structure on M compatible with ω, i.e., such that $g(\cdot, \cdot) := \omega(\cdot, J\cdot)$ is a Riemannian metric, then we can define a Riemannian metric $g_{\mathcal{P}}$ on $\mathcal{P}$ by setting

$$g_{\mathcal{P}}(\eta, \xi) := \int_0^1 g(\eta_t, \xi_t)dt.$$

Now, we define a function $\mathcal{A}$ on $\tilde{\mathcal{P}}$, the universal cover of $\mathcal{P}$, by

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

Here $\gamma \in \mathcal{P}$ and $\Gamma : (s,t) \in [0,1] \times [0,1] \mapsto \Gamma(s,t) \in M$ is a homotopy from a fixed base point β in the connected component of $\mathcal{P}$ to γ (that is $\Gamma(0,\cdot) = \beta, \Gamma(1,\cdot) = \gamma$; Γ is oriented with $(\partial s, \partial t)$ being positive), and $[\Gamma]$ is its homotopy class (relative endpaths β and γ). This is called the *action functional.*

Remark: This is the opposite of the sign convention common in physics (where the action is the integral of pdq).

Remark: To make what follows more than an analogy one would need to pick some regularity class of paths to set up relevant Banach manifolds etc. However, attempts to make rigorous parts of this exercise about gradient flowlines run into severe analytical difficulties. Reformulating the gradient flow equation as a Cauchy-Riemann equation to get a better-behaved theory is one of the insights underlying Floer cohomolgy (see [Fl87], [Fl88]).

(1) Prove that the critical points of the action functional are (lifts of) constant paths, i.e., intersection points between L_0 and L_1.

Let's compute the first variation of the action under a deformation of the path γ by η - in other words $d\mathcal{A}_\gamma(\eta)$. For sufficiently small ϵ, we can make a homotopy Γ^η from β to $\gamma + \exp(\epsilon\eta)$ by concatenating (in s direction) Γ (from β to γ) with $\Gamma_\epsilon : [0,\epsilon] \times [0,1] \to M, \Gamma_\epsilon(s,t) = \gamma(t) + \exp(s\eta)$ (from γ to $\gamma(t) + \exp(\epsilon\eta)$) and get:

$$d\mathcal{A}(\eta) = \lim_{\epsilon \to 0} -\frac{1}{\epsilon} \int_{[0,\epsilon] \times [0,1]} \omega\left((\Gamma_\epsilon)_*\left(\frac{\partial}{\partial s}\right), (\Gamma_\epsilon)_*\left(\frac{\partial}{\partial t}\right)\right) dsdt$$

$$= -\int_{[0,1]} \omega\left(\eta_t, \gamma_*\left(\frac{\partial}{\partial t}\right)\right) dt = -\int_{[0,1]} \omega(\eta_t, \dot{\gamma}(t))dt$$

The critical point conditions says that this should be 0 for all η_t. Of course if γ is constant it holds. A standard trick of calculus of variations to show the converse is to say that if $v = \dot{\gamma}(T) \neq 0$ for some $t_0 \in (0,1)$ then there is a $u \in TM_{\gamma(T)}$ such that $\omega(u,v) < 0$ and for small enough ϵ and a bump function ρ supported on $[T-\epsilon, T+\epsilon]$ the vector field $\eta(t) = \rho u$ provides the variation with $d\mathcal{A}(\eta) > 0$.

(2) Let $u : \mathbb{R} \to \tilde{\mathcal{P}}$ be a Morse cohomological flowline of the action functional with respect to the Riemannian metric $g_{\mathcal{P}}$. I.e., $u'(s) = -\nabla \mathcal{A}$ where $g_{\mathcal{P}}(\nabla \mathcal{A}, \cdot) = d\mathcal{A}$, and $\lim_{s \to -\infty} u(s) = q$, $\lim_{s \to \infty} u(s) = p$. Prove that such Morse flowlines u are in one-to-one correspondence with (lifts of) J-holomorphic strips contributing to the coefficient of q in Floer cohomology differential ∂p. Thus, Lagrangian Floer cohomology can be thought of as a version of Morse cohomology for the action functional.

Remark: There is some variation in conventions, but the one we will use is one in which Morse *homology* of f has the differential lowering the value of f; this can be done either by using for the coefficient of q in ∂p the moduli space of flowlines going down (i.e. $\gamma' = -\nabla f$) from p to q (i.e. $\lim_{s \to -\infty} \gamma(s) = p$, $\lim_{s \to \infty} \gamma(s) = q$) or the moduli space of flowlines going up ($\gamma' = \nabla f$) from q to p (with $\lim_{s \to -\infty} \gamma(s) = q$, $\lim m_{s \to \infty} \gamma(s) = p$); of course these are the same via $\gamma(s) \leftrightarrow \gamma(-s)$. Correspondingly, Morse *cohomology* has the differential increasing the value of f, and can also be defined in two ways, with coefficient of q in ∂p using either flowlines going up from p to q, or down from q to p.
   In our Floer cohomology convention, a holomorphic strip contributing to the coefficient of q in ∂p viewed as a path of paths goes from constant path at q to a constant path at p, and since the area of holomorphic strip is positive, it increases the area of homotopy from a basepoint to the path, thus decreases the action, so the action at the end point p is lower, and the differential increases action, corresponding to Morse cohomology of $\mathcal{A}$, as stated (in its "down from q to p" version; the other version would corresponds to counts of anti-holomorphic strips "up from p to q").

   For any $u : \mathbb{R} \to \tilde{\mathcal{P}}$ we can forget about the homotopy piece of the data, and think of each $u(s)$ as a path in $\mathcal{P}$. We define $U(s, t) = u(s)(t)$.
   Thus the goal is to show that U is a holomorphic strip contributing to coefficient of q in ∂p precisely when u is a negative gradient flowline starting at q and ending at p.
   Of course we always have for a fixed s, $\frac{\partial U}{\partial t} = \dot{u}(s)$, $\frac{\partial U}{\partial s} = u'(s)$ (here $\dot{u}(s)$ is the tangent vectorfield to the path $u(s)$, and $u'(s)$ is the vector field along $u(s)$ obtained by varying s).
   Continuing on the formal level as before, we have at any $\gamma \in \mathcal{P}$

$$d\mathcal{A}(\eta) = -\int_0^1 \omega(\eta_t, \dot{\gamma}(t))dt = -\int_0^1 g(\eta_t, -J\dot{\gamma}(t))dt = g_{\mathcal{P}}(\eta, J\dot{\gamma})$$

   Hence $\nabla \mathcal{A}_\gamma = J\dot{\gamma}$.
   So, again for fixed s, $-\nabla \mathcal{A}_{u(s)} = -J\frac{\partial U}{\partial t}$. As we always have $u'(s) = \frac{\partial U}{\partial s}$, the condition for u being flowline, $u' = -\nabla \mathcal{A}$, is thus equivalent to $\frac{\partial U}{\partial s} = -J\frac{\partial U}{\partial t}$, which is holomorphicity condition for U. This is what we wanted.

We have identified critical points of $\mathcal{A}$ with constant paths in part 1 of this exercise, so the $\lim s \to \pm\infty$ conditions correspond as well.

(3) Suppose that (M, ω) is exact (i.e., $\omega = d\alpha$), and L_0 and L_1 are exact Lagrangians (i.e., $\alpha|_{L_i} = dh_i$ for $i = 0, 1$).
   (a) Show that $\mathcal{A}(\tilde{p}) = \mathcal{A}(p)$ does not depend on the lift $\tilde{p}$ of p.

   We compute by Stokes theorem using exactness of ω on M:

   $$\mathcal{A}(\tilde{p}_1) - \mathcal{A}(\tilde{p}_2) = -\int_{\Gamma_1 \cup -\Gamma_2} \omega = -\int_C \alpha$$

   Here C is the union of the 4 curves traced out by the top and bottom endpoints: $C_{1,b}(s) = \Gamma_1(s, 0)$, $C_{1,t}(s) = \Gamma_1(s, 1)$ and $C_{2,b}(s) = \Gamma_2(s, 0)$, $C_{2,t}(s) = \Gamma_2(s, 1)$ (the other 4 pieces of the boundary of $\Gamma_1 \cup -\Gamma_2$ cancel in 2 pairs). Note that these glue to two piecewise smooth loops - C_b in L_0, for bottom endpoints, and and C_t in L_1 for top ones. Now by Stokes again, using exactness of α on L_0 and L_1, $\int_C \alpha = \int_{C_t} \alpha + \int_{C_b} \alpha = \int_\emptyset h_0 + \int_\emptyset h_1 = 0$.

   (b) Show that for any holomorphic strip u from q to p (counted in Floer cohomology in the coefficient of q in ∂p), we have

   $$\omega(u) = \mathcal{A}(q) - \mathcal{A}(p).$$

   First of all, changing the base point in $\mathcal{P}$ (from the definition of $\mathcal{A}$) changes $\mathcal{A}$ by an overall constant and does not affect the difference $\mathcal{A}(q) - \mathcal{A}(p)$.
   Thus we are free to pick the base point in the component of contractible paths to be the constant path at q; then we get $\mathcal{A}(q) = 0$. Now to compute $\mathcal{A}(p)$ we can view the strip u as a choice of homotopy Γ from constant path at q to constant path at p (after rescaling the time s to run over $[0, 1]$ instead of $\mathbb{R}$, which does not affect the relevant integrals). Then we get tautologically $\mathcal{A}(p) = -\omega(u)$. This does it.
   Remark: In the exact case one could also just define the action of a path $\gamma \in \mathcal{P}$ to be $h_1(\gamma(1)) - h_0(\gamma(0))$. Stokes theorem implies this differs from the original definition by a locally constant function, and so has same gradient.

   (c) Show that in this case, one can define $CF(L_0, L_1; \mathbb{Z}_2)$ by counting holomorphic strips with coefficients in $\mathbb{Z}_2$ rather than Λ (i.e., show that one doesn't get any infinite sums); and that the map

   $$(CF(L_0, L_1; \mathbb{Z}_2), \partial_{\mathbb{Z}_2}) \otimes_{\mathbb{Z}_2} \Lambda \to (CF(L_0, L_1; \Lambda), \partial)$$
   $$p \mapsto T^{\mathcal{A}(p)} \cdot p$$

   is an isomorphism of cochain complexes.

The previous item says that the energy of holomorphic strip from p to q is independent of the curve; so it is bounded, and Gromov compactness applies to show that the set of curves counted in the differential is a *compact* 0-dimensional manifold, and thus is finite. This lets us define Floer cohomology over $\mathbb{Z}_2$.

More precisely, similarly to the lecture, we define, at least for transverse L_0 and L_1, $CF(L_0, L_1; \mathbb{Z}_2) = \mathbb{Z}_2\langle L_0 \cap L_1 \rangle$, and the differential

$$\partial_{\mathbb{Z}_2} p = \sum_{q \in L_0 \cap L_1, \beta} \#\mathcal{M}(p, q, \beta, J)q$$

Here we keep track of the class β of the disc only to make the sum run over the 0-dimensional moduli spaces, but not in the coefficients or the differential.

Now, since we are summing over the same moduli spaces, the rescaling r of generators by their action identifies the complexes $(CF(L_0, L_1; \mathbb{Z}_2), \partial_{\mathbb{Z}_2}) \otimes_{\mathbb{Z}_2} \Lambda$ and $(CF(L_0, L_1; \Lambda), \partial)$. In formulas:

$$\partial(r(p)) = \partial(T^{\mathcal{A}(p)} \cdot p) = T^{\mathcal{A}(p)} \cdot \partial p =$$

$$T^{\mathcal{A}(p)} \cdot \left(\sum_{q,\beta} \#\mathcal{M}(p, q, \beta, J) T^{\mathcal{A}(q) - \mathcal{A}(p)} q \right) =$$

$$\sum_{q,\beta} \left(\#\mathcal{M}(p, q, \beta, J) T^{\mathcal{A}(q)} q \right) = \sum_{q,\beta} \left(\#\mathcal{M}(p, q, \beta, J) r(q) \right) = r(\partial_{\mathbb{Z}_2 \otimes \Lambda}(p))$$

Remark: The last exercise shows that in the exact case, the use of the Novikov field is unnecessary. For that reason one often simply defines $CF(L_0, L_1)$ with $\mathbb{Z}_2$ coefficients.

Remark: A compact symplectic manifold can never be exact (because $\omega^{\wedge top}$ must represent a non-zero class in the top degree of cohomology). All steps in the construction of Floer cohomology of closed Lagrangian submanifolds of a non-compact symplectic manifold work as in the compact case, except for compactness: one needs to impose some extra conditions on the manifold to ensure that the holomorphic curves remain confined inside a compact region, typically by some version of the maximum principle.

3. Gromov Compactness

The following definitions are in [MS04], 5.1.1, 5.2.1 (with minor notational modifications).

Definition (Stable map of genus zero with no marked points). Given (M, ω) symplectic and J compatible, a *stable holomorphic map into M modeled on a tree $T = \{V, E\}$* is 1) a holomorphic sphere $u_\alpha : S^2 \to M$ for each vertex $\alpha \in V$ 2) a special point $z_{\alpha\beta} \in S^2$

for each edge $\alpha\beta \in E$ (and since the tree T is undirected, we automatically have $\beta\alpha \in E$, but of course in general $z_{\beta\alpha} \neq z_{\alpha\beta}$ - morally, the first on lives in the domain of u_β and the second in the domain of u_α), subject to two conditions. Firstly, the maps u_α should "glue to a continuous map" - formally $u_\alpha(z_{\alpha\beta}) = u_\beta(z_{\beta\alpha})$. Secondly, the map should be *stable*, meaning it should have a finite automorphism group - this is achieved by requiring that if u_α is constant then $\deg(\alpha) \geq 3$.

We denote such a stable map by $(\mathbf{u}, \mathbf{z})$.

Convergence of holomorphic maps is largely controlled by energy bounds; with this in mind, we introduce the following definition. Given a stable map, we can split the tree into two pieces at a nodal point given by the edge $\alpha\beta$, and we then define $m_{\alpha\beta}(\mathbf{u})$ to be the total energy of the subtree containing the vertex β (and thus $m_{\beta\alpha}(\mathbf{u})$ is the total energy of the other subtree, the one containing the vertex α).

Definition (Convergence of a sequence of holomorphic spheres to a stable map of genus zero). Given a sequence of holomorphic maps $u_n : S^2 \to M$ we say that it *converges to stable map* $(\mathbf{u}, \mathbf{z})$ if there exists, for each $\alpha \in V$, a sequence of reparametrizations $\phi_n^\alpha \in Aut(S^2, j)$ such that 1) reparametrized maps $u_n^\alpha := u_n \cdot \phi_n^\alpha$ converge to u_α uniformly (with all derivatives) on compact subsets of $S^2 \setminus \bigcup_{\{\beta | \alpha\beta \in E\}} z_{\alpha\beta}$ 2) for every edge $\alpha\beta \in E$, the sequence $\phi_n^{\alpha\beta} := (\phi_n^\alpha)^{-1} \cdot \phi_n^\beta$ converges to $z_{\alpha\beta}$ uniformly (with all derivatives) on $S^2 \setminus z_{\beta\alpha}$ and 3) $m_{\alpha\beta}(\mathbf{u}) = \lim_{\epsilon \to 0} \lim_{n\to\infty} E(u_n^\alpha, B_\epsilon(z_{\alpha\beta}))$.

These mean, roughly, that one can "focus" on parts of u_n converging to each bubble in the bubble tree of $\mathbf{u}$; that when one focuses on bubble α, "refocusing" to the bubble β via $\phi_{\alpha\beta}$ looks like zooming in on a neighbourhood of the bubbling point $z_{\alpha\beta}$; and that these "changes of perspective" are done in a way that captures all of the energy of u_n, without missing any bubble components (see [MS04], p.115 for discussion).

A precise definition of Gromov convergence for maps of discs with a single lagrangian boundary condition, similar to the above, is in [Fr08], Definition 3.1 (for strips see also [FOOO14], sections 2.1.2 and 3.7). We shall not delve into this and thus will keep corresponding parts of the exercise informal.

(1) Let J be the standard complex structure on $\mathbb{CP}^1$.
  (a) Find a sequence of J-holomorphic maps $u_n : \mathbb{CP}^1 \to \mathbb{CP}^1$ which converge to a nodal sphere.

   Consider the set of rational maps $\mathbb{CP}^1 \to \mathbb{CP}^1$ of degree 2. Geometrically these are double branched covers; if the branching points are distinct then Euler characteristic implies that there are 2 of them. These are the two critical points of

our rational map. The Gromov compactification adds in the nodal curves as the limit where the two branched points come together (that is, the critical *values* converge to each other). One of the simplest such things is $p_n(z) = z^2 + a_n$ where the second branch point/critical value a_n moves towards the one at infinity. With benefit of hindsight to make algebra simpler we take $a_n = -n^2$. Then we have $u_n([z : w]) = ([z^2 - n^2 w^2 : w^2])$. The goal now is to find some reparametrizations $\phi_n, \psi_n : \mathbb{CP}^1 \to \mathbb{CP}^1$ such that the two sequences of maps $u_n \cdot \phi_n$ and $u_n \cdot \psi_n$ converge (uniformly with all derivatives on compact subsets of compliment of the single point where the other bubble is attached).

To simplify our task we first write $u = z - nw, v = z + nw$, so that $z = \frac{u+v}{2}, w = \frac{v-u}{2n}$ and the maps become $u_n([u : v]) = ([uv; (\frac{u-v}{2n})^2])$.

This is better because some things are explicitly vanishing (in fact this converges with all derivatives on the complement of $\{u = 0\} \cap \{v = 0\}$ to the constant map to $[1 : 0]$; this is *not* allowed in Gromov convergence since any constant component of a stable curve must have at least 3 special points). We just need to find ϕ_n, ψ_n so that the reparametrized sequences converge to non-constant, degree 1 maps. Inspired by the uv term and looking to give $(\frac{u-v}{2n})^2$ a non-vanishing piece, we consider $\phi_n([u : v]) = ([2nu : \frac{v}{2n}])$. Then $u_n \cdot \phi_n([u : v]) = [uv : u^2 + (\frac{uv}{4n^2} + \frac{v^2}{16n^4})]$, which converges to $b([u, v]) = [v : u]$ on the complement of $u = 0$ (uniformly with all derivatives on compact subsets of that complement). Similarly, the rescalings $\psi_n([u, v]) = [\frac{u}{2n} : 2nv]$ make $u_n \cdot \psi_n([u, v])$ converge to $b([u, v]) = [u : v]$ outside $\{v = 0\}$.

Both limits map $\mathbb{CP}^1 \setminus p$ to $\mathbb{CP}^1 \setminus \{[1 : 0]\}$, and both have continuous extensions that map their "missing points" to $[1 : 0]$ (exactly where our geometric argument leads us to expect the double point of our nodal curve to map), thus they glue to a stable map.

We also have $\phi_n^{-1} \cdot \psi_n([u : v]) = [\frac{u}{4n^2} : 4n^2 v]$, converging to $[0 : 1]$ outside $[1 : 0]$ and $\psi_n^{-1} \cdot \phi_n([u : v]) = [4n^2 u : \frac{v}{4n^2}]$ converges to $[1 : 0]$ outside $[0 : 1]$, as required by Gromov convergence.

Finally, the energy condition holds as well, since $u_n \cdot \phi_n([u, v])$ converges outside $[0 : 1]$ to degree 1 map, the energy on complement of $B_\epsilon([0 : 1])$ converges to twice the area of the line, and since the energy of each u_n is twice the area of a line, the energy on the ball itself converges to the area of a line as well, as wanted (and similarly for $u_n \cdot \psi_n([u, v])$).
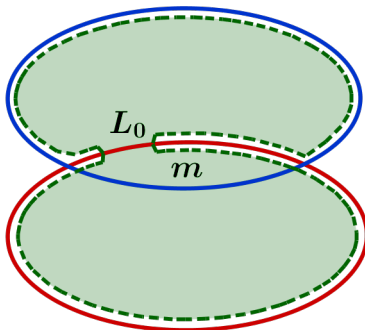
(b) Now let $L_0, L_1 \subset \mathbb{CP}^1$ be simple closed curves (i.e., Lagrangians) which intersect, transversely. Find a sequence of J-holomorphic strips with boundary on L_0 and L_1 which converges to a non-constant J-holomorphic strip with a disc bubble attached at a boundary point.

See the next part.

(c) Now do the same, but with a sphere bubble attached at an interior point.

For any such curves which intersect there is an embedded topological strip m : $\mathbb{R} \times [0,1] \to \mathbb{CP}^1$ with $m(s,i) \in L_i$ between a pair of intersection points p and q. Now, for part c), postcompose the maps u_n in part a) with (a fixed) rotation taking $[1:0]$ to a point r in the interior of (image of) m; call the result v_n. Now v_ns bubble in the interior of m, so all we need to do is cut their domain S^2 down to a strip, to get a sequence of holomorphic strips bubbling at r. This is not hard to do. For large enough n both branched points of v_n are inside m, so on the complement of m the map v_n is double covering. Excizing one of the preimages of that complement from the domain we get a piece of S^2 biholomorphic to a strip. Composition of that biholomorphism ρ_n with restriction of v_n to ρ_n's image gives a holomrphic strip w_n. Now w_n converges to a non-constant strip with a bubble at r.

For part b), we have to choose rotations that move the two branch points to points on the L_0 boundary of m, and be a bit more careful about cutting the domain. The picture is in figure below.

(2) Let J be the standard integrable complex structure on $\mathbb{CP}^2$.
   (a) Consider the J-holomorphic curves (without boundary)

$$u_\epsilon : \mathbb{CP}^1 \to \mathbb{CP}^2$$
$$u_\epsilon([z:w]) = [z^2 : \epsilon z w : w^2].$$

What is the limit of this sequence of curves as $\epsilon \to 0$? As $\epsilon \to \infty$?

As $\epsilon \to 0$ the maps converge to $[z^2 : 0 : w^2]$ (since there is no bubbling, there is nothing else to check besides convergence).

As $\epsilon \to \infty$ the image lies in the curve $z_0 z_2 = \frac{1}{\epsilon^2} z_1^2$ which degenerates to the nodal subvariety $z_0 z_2 = 0$. Thus we expect the maps to converge to a stable map with two components, one mapping to $z_0 = 0$ and the other to $z_2 = 0$.

Consider the reparametrizations $\phi_\epsilon([z : w]) = [z : \epsilon w]$. The compositions $u_\epsilon \cdot \phi_\epsilon([z : w]) = [z^2 : \epsilon^2 zw : \epsilon^2 w^2] = [\frac{z^2}{\epsilon^2} : zw : w^2]$ which converges to $[0 : z : w]$ outside $[1 : 0]$; the limit sends $[1 : 0]$ to $[0 : 1 : 0]$.

Similarly, we take $\psi_\epsilon([z : w]) = [\epsilon z : w]$. The compositions $u_\epsilon \cdot \psi_\epsilon([z : w]) = [\epsilon^2 z^2 : \epsilon^2 zw : w^2] = [z^2 : zw : \frac{w^2}{\epsilon^2}]$ which converges to $[z : w : 0]$ outside $[0 : 1]$; the limit sends $[0 : 1]$ to $[0 : 1 : 0]$ - thus the two limits glue to a stable map as expected.

We check that $\psi_\epsilon^{-1} \cdot \phi_\epsilon([u : v]) = [z : \epsilon^2 w]$ converges to $[0 : 1]$ outside $[1 : 0]$ and similarly for $\phi_\epsilon^{-1} \cdot \psi_\epsilon$; an argument similar to one in part 1a of this exercise shows that the the energy condition also holds and completes the proof of Gromov convergence.

(b) Use the above example to construct a sequence of J-holomorphic discs $u :$ $(\mathbb{D}, \partial\mathbb{D}) \to (\mathbb{CP}^2, \mathbb{RP}^2)$ whose limit is two discs joined at a boundary node.

Taking ϵ real (and increasing to infinity) and restricting the above u_ϵ to, say, upper hemisphere $\Im(z\bar{w}) \geq 0$ we get a sequence of maps of the disc with boundary $\mathbb{R}P^1$ mapping to $\mathbb{R}P^2$ and with reparametrizations preserving $\mathbb{R}P^1$, and bubbling occuring at real points $[0 : 1]$ and $[1 : 0]$; so these maps indeed Gromov-converge to a pair of discs joined at the boundary.

4. GRADING

(1) Prove that the Maslov index for paths is uniquely characterized by the properties stated in lecture.

We want to show that any μ satisfying the properties of homotopy invariance, product, and normalization given in the lecture computes the same thing on any path p of Lagrangians with transverse endpoints.

First we reduce to the case $p(0) = \mathbb{R}^n$. For this, we connect $p(0)$ to $\mathbb{R}^n$ by a path in $\mathcal{L}$, lift this connecting path to a path $\gamma(s)$ in $U(n)$, then rotate our p by $\gamma(s)$ to get a homotopy with transverse endpoints from p to p_n with $p_n(0) = \mathbb{R}^n$. By the homotopy invariance, $\mu(p) = \mu(p_n)$.

Next, we homotope $p_n(1)$ to a product Lagrangian through Lagrangians transverse to $\mathbb{R}^n$. In fact, all subspaces of dimension n transverse to $\mathbb{R}^n$ are given as graphs of maps from $\mathbb{R}^n$ to $i\mathbb{R}^n$ (aka $n \times n$ matrices) and the Lagrangian condition is the symmetry of the matrix. The set of such matrices is connected and we simply connect the "endpoint" matrix corresponding to $p_n(1)$ to the identity matrix, thus connecting $p_n(1)$ to $i\mathbb{R}^n$.

Now we homotope the resulting path q of Lagrangians to a path of product Lagrangians relative endpoints. This is always possible, because π_1 of the subset of
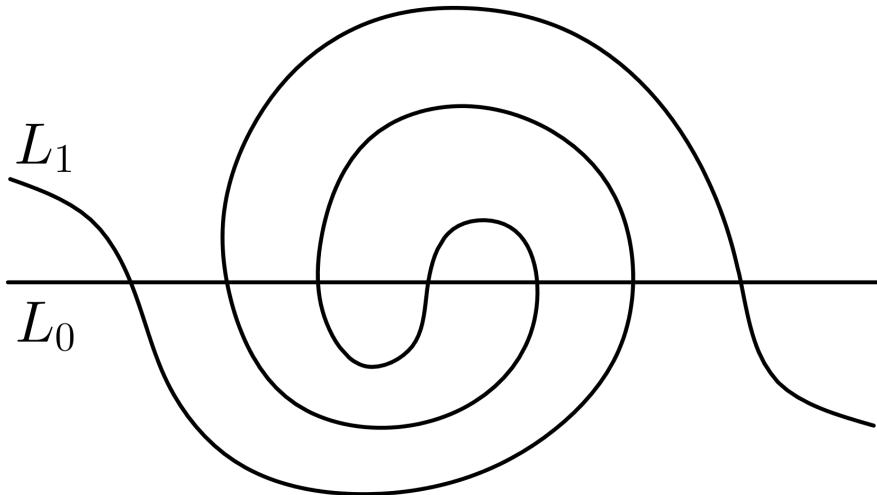
product Lagrangians $\mathcal{L}_{pr}$ surjects onto the π_1 of $\mathcal{L}$ (and thus the long exact sequence of homotopy groups for pair $(\mathcal{L}, \mathcal{L}_{pr})$ implies $\pi_1(\mathcal{L}, \mathcal{L}_{pr}) = 0$, which then implies what we want). Finally each of the component paths of Lagrangians is homotopic to one of the standard paths that appear the normalization axiom (since $\mathcal{L}(1)$ is a circle).

The end result is that any path with transverse ends is homotopic in the class of such paths to a product of standard paths whose index is given by axioms of normalization and product. Thus the index is uniquely determined by these axioms.

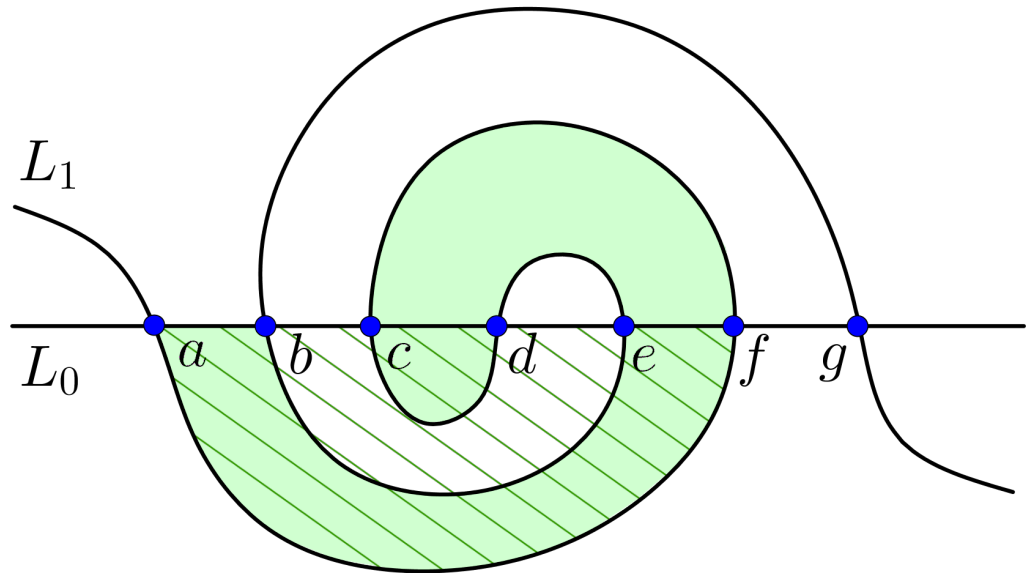(2) Prove the claim made in the lecture: the Floer differential has degree $+1$.

The lemma giving the index of a holomorphic strip in ungraded case gives the formula $i(\beta) = \mu(\tilde{p}) - \mu(p_p) - \mu(p_q)$; once we have gradings we are free to take the p_p and p_q that lift to paths connecting $\widetilde{T_pL_0}$ to $\widetilde{T_pL_1}$ and $\widetilde{T_qL_0}$ to $\widetilde{T_qL_1}$ respectively. Then the loop $\tilde{p}$ obtained by concatenating p_q to the path given by $u(s, 1)$ in L_1 to the reverse of p_p and to the path given by $u(s, 0)$ in L_0 (ran backwards) as in the lecture actually lifts to a path from $\widetilde{T_qL_0}$ to itself. This means that $\tilde{p}$ is contractible in $\mathcal{L}$, and hence has $\mu(\tilde{p}) = 0$. Now the formula $i(\beta) = \mu(\tilde{p}) - \mu(p_p) + \mu(p_q)$ reduces to $i(\beta) = deg(q) - deg(p)$.

(3) Work out the graded Floer complex $CF^\bullet(L_0, L_1)$ for the following two Lagrangians in the plane:

Let's label the intersection points from left to right a, b, c, d, e, f, g as in the picture below. The tangent bundle to the plane is canonically trivialized so it's particularly easy to use absolute gradings. L_0 has all tangent spaces horizontal which we can lift to 0 in $\mathbb{R} = \tilde{\mathcal{L}}$ covering $S^1 = \mathcal{L}$. We have a choice of liftings for L_1 and we can choose one for which $\deg(a) = 0$ (this corresponds to $\widetilde{T_aL_1} \approx -0.4$). Then it follows from the normalization axiom and our definitions that any other intersection point q will get grading equal to the number of times the tangent of L_1 crosses the horizontal

on the way from a to q. Thus $\deg(f) = 1$, $\deg(c) = 2$, $\deg(d) = 3$, $\deg(e) = 2$, $\deg(b) = 1$ and $\deg(g) = 0$. By Riemann mapping theorem the holomorphic disc count in the differential coincide with the count of topological discs between points of index difference 1. We find these discs directly in the picture. There is a lower-half disc looking one from a to f, marked in the picture (note that in our conventions - we are doing Floer *co*homology - the strip goes from right intersection point to left intersection point when L_0 is on the bottom and from left to right when L_0 is on top, as in the present case.) There is the "tongue" from a to b, also marked. There is "small tongue" from b to c and a "small half-disc" from b to e; there is am "extra small half disc" from c to d. Symmetrically there are unique discs from g to b and to f, from f to e and to c, and from e to d.

Putting it all together we see that the differential is exact in all degrees except 0, where the cohomology is generated by $a + g$. This is the same answer that we would get by making a Hamiltonian isotopy that removes all intersection points to the right of a; the resulting graded complex (with gradings induced from the isotopy, which we can take to be identity near a) would have a single generator in degree 0, thus computing the same cohomology.
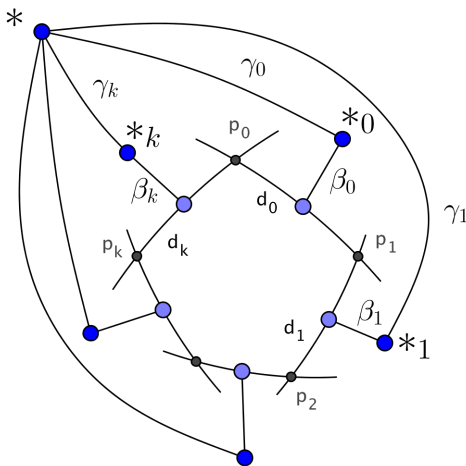
(4) Let $L \subset M$ be a Lagrangian such that the image of $\pi_1(L, *_L)$ in $\pi_1(M, *)$ is trivial. An *anchoring* for L is a homotopy class of paths γ in M from $*$ to $*_L$ (see [FOOO10]). Now suppose that L_0 and L_1 are anchored; for each $p \in L_0 \cap L_1$, assign a grading $deg(p) \in \pi_1(M, *)$ by following the chosen paths from $*$ to $*_{L_0}$, to p (inside L_0), to $*_{L_1}$ (inside L_1), to $*$. Prove that the A_∞ products m^k respect this grading, in the

sense that

$$deg(m^k(p_1, \ldots, p_k)) = \sum_{i=1}^{k} deg(p_i).$$

First one makes an obvious check that the resulting $\deg p$ is independent of choices (using that the image of $\pi_1(L_j, *_{L_j})$ in $\pi_1(M, *)$ is trivial for all j). It is also clear that $[\deg p_i \in CF(L, K)] = -[\deg p_i \in CF(K, L)]$.

Now, if $p_0 \in L_0 \cap L_k$ is a summand of $m^k(p_1, \ldots, p_k)$ then there is a disc with pieces of the boundary d_0 connecting p_0 to p_1 in L_0, d_1 connecting p_1 to p_2 in L_1 etc. and finally d_k connecting p_k to p_0 in L_k. Choose paths β_j connecting $*_{L_j} = *_j$ to some point on d_j (see figure below). Using these β_js together with with pieces of d_j as the paths in L_j connecting $*_{L_j}$, to p_j in the above recipe for $\deg p_j$, we get that $\sum_{i=0}^{k} \deg(p_i) = \sum_{i=0}^{k} d_j = 0$ in $\pi_1(M, *)$, since $\sum d_i$ is the boundary of the disc. This shows what we want.

(5) Show that there is a natural isomorphism of cochain complexes

$$(CF^\bullet(L_0, L_1), \partial) \cong (CF^{n-\bullet}(L_1, L_0)^\vee, \partial^\vee).$$

This isomorphism (or more precisely, the induced isomorphism on the cohomology level) is called *Poincaré duality*.

First, we note that if p is a path in $\mathcal{L}$ and $-p$ is the same path running packwards then $\mu(p) + \mu(-p) = n$. To see this, observe that this is true in dimension 1, since for non-integer N (corresponding to transverse endpoints) $(\lfloor N \rfloor + 1) + (\lfloor -N \rfloor) + 1) = 1$, thus it is true for paths of product Lagrangians, and since every p is homotopic to such (see part 1 of this exercise), and the homotopy of p to a product path pr is also homotpy of $-p$ to $-pr$, we conclude this is true for all p.

In the graded case, if the path p lifts to a path connecting $\widetilde{T_{pt}L_0}$ to $\widetilde{T_{pt}L_1}$, then $-p$ lifts to a path connecting $\widetilde{T_aL_1}$ to $\widetilde{T_aL_0}$. Thus the gradings of a in $CF(L_0, L_1)$ and in $CF(L_1, L_0)$ add to n. This identifies $CF^\bullet(L_0, L_1) = CF^{n-\bullet}(L_1, L_0)^\vee$ (at this stage, i.e. without the differential, $\vee$ is optional).

In both complexes the differential counts the same strips, only the strip going from p to q in $CF(L_0, L_1)$ is a strip going from q to p in $CF(L_1, L_0)$ (strictly speaking this is only true for the case of time-independent J; otherwise we should use $\tilde{J}_t = J_{1-t}$ in $CF(L_1, L_0)$ if we are using J_t in $CF(L_0, L_1)$). This gives precisely the identification $(CF^\bullet(L_0, L_1), \partial) \cong (CF^{n-\bullet}(L_1, L_0)^\vee, \partial^\vee)$. (Think of the differential ∂ between $CF^i(L_0, L_1)$ and $CF^{i+1}(L_0, L_1)$ as a matrix D (with Novikov field entries); then ∂ between $CF^{n-(i+1)}(L_1, L_0)$ and $CF^{n-i}(L_1, L_0)$ is matrix D^T, so that $\partial^\vee$ between $CF^{n-i}(L_1, L_0)^\vee$ and $CF^{n-(i+1)}(L_1, L_0)^\vee$ is again the matrix D, but now it is a homological differential of degree -1.)

Remark: One has to be a bit more careful in the case of wrapped Floer cohomolgy, where one is passing to limits of a directed system of complexes or is dealing with infinite-dimensional complexes.

5. Holomorphic discs and displacement energy.

Let ϕ be a compactly supported Hamiltonian symplectomorphism of (M, ω). Define the Hofer norm of ϕ to be the infimum, over all Hamiltonians generating ϕ, of their time-averaged variations; that is define

$$\rho(\phi) = \inf_{H_t} \int_0^1 \left(\max H_t - \min H_t \right) dt.$$

where the infimum is taken over all compactly supported $H_t : [0, 1] \times M \to \mathbb{R}$ with $\phi^1 = \phi$ (this was proved to be a metric on $Ham^c(M, \omega)$ by Lalonde and McDuff [LM]).

We also define the displacement energy of $K, L \subset M$ of two compact subsets of M to be the infimum of Hofer norms of all Hamiltonian symplectomorphisms disjoining K and L from each other:
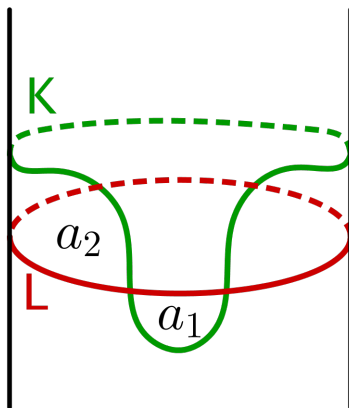
$$e(K, L) = \inf_{\{\phi | \phi(K) \cap L = \emptyset\}} \rho(\phi).$$

When $K = L$ then $e(K) := e(K, K)$ is known as the displacement energy of K.

(1) Prove $e(K, L) = e(L, K)$.

Of course $\phi(K) \cap L = \emptyset \iff K \cap \phi^{-1}(L) = \emptyset$. So we just need to see that $\rho(\phi) = \rho(\phi^{-1})$. This is the kind of thing which is clear once one remembers a proof of the fact that the group of Hamiltonian symplectomorphisms is a group, where one had to provide a Hamiltonian function for ϕ^{-1}. One way is to show that if F_t is the (time-dependent) Hamiltonian generating ϕ_t, then $-F_t \cdot \phi_t$ generates ϕ_t^{-1} (another is to use $-F(1-t)$). Of course F_t and $-F_t \cdot \phi_t$ (or $-F(1-t)$) have the same time-averaged variations.

(2) Consider K and L two (Lagrangian) curves in the cylinder:

Suppose the areas of discs between K and L are a_1 and a_2 (with $a_1 < a_2$). Compute $e(K, L)$. (Warning: This is not entirely elementary.)

First, we show how one can displace K from L with a Hamiltonian of Hofer norm arbitrarily close to a_1.

We can view the cylinder as $C = \mathbb{R} \times S^1$ with coordinates r, θ with $L = 0 \times S^1$ and $K = \{f(\theta), \theta\}$ for some function f and symplectic form $\omega = dr \wedge d\theta$ (if we view the cylinder as $T^\star S^1$ this is $\omega = d\lambda = dp \wedge dq$). Let the intersection points be $(0, \theta_1) = (0, 0)$ and $(0, \theta_2)$.

The point is that a Hamiltonian $H(\theta)$ depending only on the angle coordinate will generate a flow preserving that coordinate. In fact, $\omega(\cdot, X) = dH$ is $X_\theta dr = 0$, $-X_r d\theta = \frac{\partial H}{\partial \theta} d\theta$, so the flow will in time 1 move K by $-\frac{\partial H}{\partial \theta}$.

Thus we can consider a Hamiltonian H independent of time and the r coordinate, with $H(\theta) = -\int_{\theta_1}^{\theta} (f(t) + \epsilon) dt$ for $\theta \in [\theta_1, \theta_2]$ and for all other θ choose $H(\theta)$ smoothly extending the above in such a way that $f(\theta) - H'(\theta) > 0$ and $H'(\theta) \geq 0$ (for $\theta \notin [\theta_1, \theta_2]$; this is possible for small enough ϵ since $a_2 > a_1$). Then we see that $\phi(K)$ is completely above L, and the Hofer energy of H is $a_1 + \epsilon(\theta_2 - \theta_1)$, which is arbitrarily close to a_1.

The converse is more difficult. We will use the most basic version of *energy-capacity inequality* in $\mathbb{R}^2$, which states that in displacement energy of a ball in the plane is at least its area.

Remark: In general, an energy-capacity inequality bounds displacement energy of a subset in terms of its symplectic capacity (of some type). There are many proofs of such inequalities, but they all use some non-trivial ingredient, be it some variational analysis, non-squeezing theorem or Floer theory (see [MS98], p. 384). No elementary proof seems to be known, even in the lowest dimension.

Existence of covering map $(\mathbb{R}^2, \omega_{std}) \to (C, \omega)$ means that the energy-capacity inequality for balls in $\mathbb{R}^2$ implies the same for balls in C: if H_t displaces a ball in C, pullback of H displaces a ball in $\mathbb{R}^2$. (Existence of embedding $(\mathbb{R}^2, \omega_{std}) \to (C, \omega)$ shows the implication the other way, so the energy-capacity inequality in C and $\mathbb{R}^2$ are actually equivalent).

Now any Hamiltonian disjoining K from L will displace the disc between them of area a_1 off of itself. Thus it must have energy at least a_1, which is what we wanted.

(3) In general, for two Lagrangians K and L for which the Floer cohomology is well defined, give a bound on $e(K, L)$ in terms of this Floer cohomology (over the Novikov *ring*), and give an idea of a proof.

If Floer cohomology over the Novikov ring is non-torsion, then the Floer homology over Novikov field is non-zero, and the two Lagrangians can not be displaced from each other by *any* Hamiltonian flow; thus the diplacement energy is infinite.

If Floer cohomology over the Novikov ring is torsion, then for any e such that it is "non-zero modulo T^e" (formally, a torsion cohomology group is isomorphic to a sum $\oplus \Lambda_0 / T_i^\mu \Lambda_0$; the μ_i's are called *torsion exponents* and we ask that the biggest torsion exponent μ is larger than e), then the displacement energy is at least e. To see this, one writes down the continuation map from $CF(L, K; \Lambda_0)$ to $CF(\phi(L), K; \Lambda_0)$ and checks that it does not change the largest torsion exponent by more than $|\phi|$. Thus if ϕ disjoints K and L the largest torsion exponent must go from mu to 0, so the energy of ϕ must be at least μ. A more refined version giving bounds on number of intersection points of L_0 and $\phi(L_1)$ in terms of torsion exponents is [FOOO14], Theorem J. The largest torsion exponent is related to the notion of *boundary depth* introduced by Usher; this and other quntitative invariants of Floer complexes are captured by their *barcodes* (with a caveat that one has to be extra careful studying persistence modules over Novikov field; see [UZ16]).

6. MORSE TRAJECTORIES AND HOLOMORPHIC STRIPS (SEE [Auroux, §1.6]).

Consider a Morse function f on smooth manifold N. Given a metric g such that the pair (f, g) satisfy the Morse-Smale transversality condition, one obtains a Morse-Witten complex generated by critical points of f and with differential counting gradient trajectories of f.

Let L_0 be the zero section in $T^\star N$, and L_1 the graph in $T^\star N$ of df. We denote the projection from T^*N to N by π.

(1) Show that the Hamiltonian $H = f \cdot \pi$ isotopes L_0 to L_1.
    Remark: This exercise follows conventions of [Auroux, §1.6], and in particular we take $\omega = \sum_j dq_j \wedge dp_j$ (while keeping $\omega(X, \cdot) = -dH$). This makes Hamilton's equations of motion come out with non-standard sign (see [Wend16]), but keeps things consistent with our reference materials.

    We compute the Hamiltonian vector filed $X = X_H$ for $H = f \cdot \pi$:

$$\omega(X, \cdot) = -d(f(q_1, \ldots, q_n))$$
$$-X_{p_j} dq_j + X_{q_j} dp_j = -\frac{\partial f}{\partial q_j} dq_j - 0 dp_j$$
$$X_{q_j} = 0, X_{p_j} = \frac{\partial f}{\partial q_j}$$

    Thus the flow of H is translation in the fiber with constant speed; at time one the point $(q_1, \ldots, q_n, 0, \ldots, 0)$ reaches $(q_1, \ldots, q_n, \frac{\partial f}{\partial q_1}, \ldots, \frac{\partial f}{\partial q_n})$ as expected.

(2) Show that L_0 and L_1 intersect transversely at the critical points of f.
    This identifies the generators of Morse-Witten and Floer complexes in this setting.

    The identification of intersection points and critical points is obvious form the previous part. Transversality of df to the zero section just means (in coordinates as above) that the derivative of map $(q_1, \ldots, q_n) \to (\frac{\partial f}{\partial q_1}, \ldots, \frac{\partial f}{\partial q_n})$ is surjective, which is to say the Hessian of f has full rank - precisely what f being Morse guarantees.

(3) Use g to get a compatible almost complex structure for $T^\star N$, and then identify $T_{(p,q)}(T^\star N) = \mathbb{C} \otimes T_q N$, with vertical subbundle purely imaginary and horizontal purely real. Show that this gives a trivialization of $(\Lambda_{\mathbb{C}}^n(T^\star T^\star N))^2$ as in the lecture.

    To describe the almost-complex associated to g, we start by noting that we have a subbundle $T_w^v(T^*N) = \ker D\pi_w \simeq T^*N_{\pi(w)} \subset T(T^*N)$. A choice of Ehresmann connection for T^*N is by definition a choice of complimentary subspace $T_w^h(T^*N)$, which is then isomorphic to $T_{\pi(w)}N$ via $D\pi$. The canonical such choice is given by

dual of Levi-Civita connection of g. Putting together these identifications we get an isomorphism $i : T_w(T^*N) \xrightarrow{\sim} T_{\pi(w)}N \oplus T^*_{\pi(w)}N$ (note that we put horizintal part first). Both sides have canonical (up to sign) symplectic forms and one checks (for example by choosing normal coordinates on N around $\pi(w)$ so that the splitting at w becomes trivial) that i is a symplectomorphism (thus picking the sign on the right hand side). Finally we use g to identify $T^*_{\pi(w)}N$ with $T_{\pi(w)}N$ (and by abuse of notation drop this identification from notation), and then define J by $J(a,b) = (-b,a)$ (recall this is for the symplectic form on T^*N equal to $\sum dq_j \wedge dp_j$). It is then easy to check that the metric compatible with J, ω is the one that restricts to g on both pieces of the decomposition and makes them orthogonal to each other. This provides the identification $T_{(p,q)}(T^\star N) = \mathbb{C} \otimes T_q N$ that we want.

To describe the trivialization $(\Lambda^n_{\mathbb{C}}(T^\star T^\star N))^2$ that this gives, pick any (real) orthonormal basis in the horizontal subspace. Any complex basis is obtained from it by a complex matrix, and we send it to the square of the determinant of that matrix (linearly dependent sets of n tangent vectors are sent to 0). This gives a section of $(\Lambda^n_{\mathbb{C}}(T^\star T^\star N))^2$; since basis change matrices between different real orthonormal bases are in $O(n)$, the resulting section is well-defined and independent of the choice of the real basis.

(4) Show that L_0 is graded (by a constant lift!), and we can use the Hamiltonian isotopy above to grade L_1.

We note that since the splitting is given by horizontal subspaces of Levi-Civita connection, on the zero section it is simply equal to the splitting of $T(T^*N)$ to the tangents of fiber and zero section, and thus tangent to the zero section is horizontal, i.e. real. Hence L_0 is always sent to 0 by the above section, and thus has a constant (0) lift. As for the last bit, any Hamiltonian isotopy lifts to isotopy in $\mathfrak{L}M$ and hence an isotopy of a graded Lagrangian lifts to $\tilde{\mathfrak{L}}M$.

(5) Show that Maslov and cohomological Morse gradings of all $p \in L_0 \cap L_1 = crit(f)$ coincide.
This identifies Morse-Witten and Floer complexes as graded vector spaces.

We want to see that the two gradings of an intersection point $p \in L_0 \cap L_1 = crit f$ coincide. Picking coordinates on N in which Hessian of f is diagonal with entries ± 1, we see that the isotopy lift from the previous part is the path of Lagrangian submanifolds at the same point $(0, q)$ which is split. Each $+1$ eigenspace contributes as a split component path of Lagrangians specified by symmtric 1 by 1 matrix t, which has index 1; each negative eigenspace contributes as a path of Lagrangians given by path of matrices $-t$. Thus by our axioms of Maslov index for paths, index of a critical point is equal to the dimension of the positive eigenspace, which is exactly the cohomological Morse index, as wanted.

(6) Recall that Floer equation perturbs the Cauchy-Riemann holomorphicity condition by adding a Hamiltonian term X_H:

$$\frac{\partial u}{\partial s} + J_t(u)\left(\frac{\partial u}{\partial t} - X_{H_t}(u)\right) = 0$$

Show that $u(s,t)$ solves the above perturbed equation if and only if $\tilde{u}(s,t) = (\phi_H^t)^{-1}(u(s,t))$ solves the original Cauchy-Riemann equation with $\tilde{J}_t = (\phi_H^t)^{-1}(J_t)$. The two versions, one perturbing the boundary L_1 (by ϕ^{-1}) and one perturbing the equation, are therefore equivalent.

This is just chain rule. In fact $\frac{\partial \tilde{u}}{\partial s}|_{s,t} = (\phi_H^t)_\star^{-1}\left(\frac{\partial u}{\partial s}|_{s,t}\right)$ and $\frac{\partial \tilde{u}}{\partial t}|_{s,t} = (\phi_H^t)_\star^{-1}\left(\frac{\partial u}{\partial t}|_{s,t}\right) + (\phi_H^t)_\star^{-1}\left(-X_{H_t}\right)$ and the result follows.

(7) Using the compatible J from part (3) of this exercise and Hamiltonian perturbation $H_t = -H$ (thus corresponding to moving L_1 by H, as in part (1)), see that $u(s,t) = \gamma(s)$ is a solution for perturbed equation whenever γ is a Morse flowline of f.

Of course $\frac{\partial \gamma}{\partial t} = 0$, and so Floer equation reduces to $\frac{\partial \gamma}{\partial s} = -JX_H$. Since H is pulled back from N, dH is zero on all vertical vectors, so X_H is vertical. Then $-JX_H$ is horizontal, and since we are at the zero section, $-JX_H \in TN \subset T(T^*N)$. Moreover for any $v \in TN$ (at the same point of N), and since ω is compatible with the metric which restricts to g on TN, $g(v, -JX_H) = \omega(v, X_H) = dH(v) = df(v)$, so $-JX_H = \nabla f$. Floer equation is now reduced to $\frac{\partial \gamma}{\partial s} = \nabla f$, which γ solves being a gradient flowline.

One can show that for f sufficiently C^2 small the converse is also true - any solution of Floer equation with this perturbation data comes form a gradient trajectory ([Fl89], Theorem 2).

(8) Rescale the generators $p \to T^{f(p)}p$ to identify (for C^2 small f and modulo regularity issues which we ignore) Morse and Floer complexes over the Novikov ring.

Note that our Lagrangians are exact and $h_0 = 0$, $h_1 = f$ so by the remark in 2.3b we can take $\mathcal{A}(p) = f(p)$. Then the rescalings appearing here and in that problem coincide. Since we identified generators and the differential, the Floer chaincomplex of (L_0, L_1) over $\mathbb{Z}_2$ is identified with cohomological Morse complex of f over $\mathbb{Z}_2$. The rescaling identifies the Λ versions as in 2.3c.

Remark: In 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 imply that, for a sufficiently small Hamiltonian perturbation, the pseudo-holomorphic strips that determine the Floer cohomology $HF^\star(L, L)$ must all be contained in a small tubular neighborhood of L, so that the calculation of Floer cohomology reduces to the above, and we get that $HF^\star(L, L) = H(L, \Lambda)$ (this is originally due to Floer in the exact case for $K = \mathbb{Z}_2$, [Fl89]).

7. The Stasheff associahedra and holomorphic discs

The Stasheff associahedron K_k [Sta63] is a polyhedron whose faces F_T are indexed by planar trees T with $k+1$ labeled, cyclically ordered semi-infinite leaves (with no finite leaves, every vertex of degree at least 3, and up to planar isomorphism, of course); and $F_{T_1} \subset F_{T_2}$ if and only if T_2 is obtained from T_1 by contracting some edges.

(1) Draw K_3, K_4 and K_5. Describe the natural action of $\mathbb{Z}/k+1$ on K_k in these examples (the natural action arises by rotating the planar trees, aka cyclically shifting the labels).

  Let's say that T_1 is obtained from T_2 by expanding an edge $e \in E(T_1)$ if T_2 is obtained from T_1 by contracting the (finite) edge e. If T_1 is so obtained then it can be split into two trees by cutting the expanded edge e. We will call these e-subtrees of T_1. Observe that every tree we consider for K_k is obtained by repeated expansions from the star tree S_{k+1} (all infinite leaves incident on a single internal vertex of degree $k + 1$), which corresponds to the interior face.
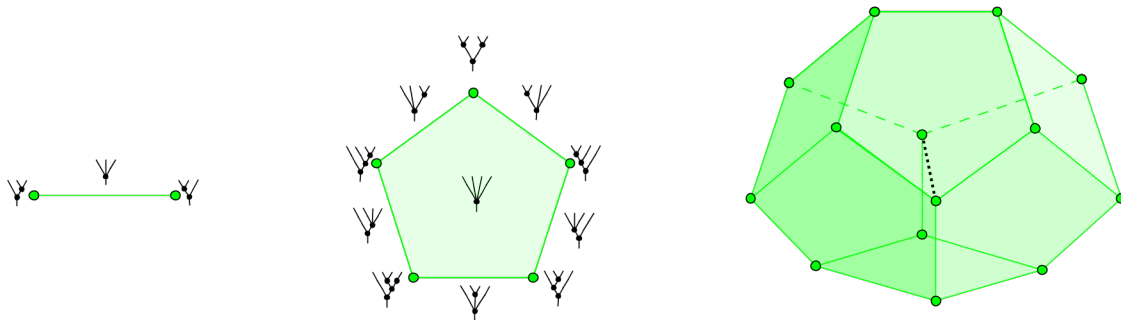 We label the infinite leaves starting from 0.
  We note that S_3 is not expandable, corresponding to the fact that K_2 is a point.
  For K_3 there are two ways to expand the S_4, corresponding to splitting the infinite leaves into groups $0 - 1, 2 - 3$ or $1 - 2, 3 - 0$, at which point the expansion stops. We get a segment. The action of $\mathbb{Z}_4$ is by rotation by π, so twice the generator acts trivially (renaming leaves $0 \to 2$, $1 \to 3$ etc does not in fact change the isomorphism type of the tree in this case).
  For K_4 there are 5 ways to make initial expansion of S_5, corresponding to splitting the leaves into two groups of 2 and 3 cyclically consecutive leaves. For any such expansion, one of the resulting e-subtrees is an S_4 and one is S_3. An S_4 can then be expanded 2 ways as in the case of K_3, and S_3 is non-expandable (as in the case of K_2). Thus we get 5 segments (which we think of as 5 copies of $K_3 \times K_2$), and the whole thing is a pentagon (see figure below), with $\mathbb{Z}_5$ acting by rotation by $4\pi/5$.
  Finally, for K_5, the first expansion of S_6 is either into a S_4 and an S_4 (3 ways), or an S_5 and an S_3 (6 ways). In the first case 2 more expansions give a copy of $K_3 \times K_3$ - a quadrilateral; in the second case a copy of $K_4 \times K_2$ - a pentagon. So K_5 has 6 pentagonal facets and 3 quadrilateral ones (see picture below). These are glued according to which edges (and vertices) correspond to isomorphic trees.
  The resulting combinatorial polytope has a 3-fold symmetry which in our picture is rotation by third full turn around the axis shown, and a 2-fold reflection symmetry with respect to the perpendicular bisector plane to that axis. These commute, and give the action of $\mathbb{Z}/6 = \mathbb{Z}/2 \times \mathbb{Z}/3$.

Remark: One other description which is convenient is to note that we can always rotate the tree (keeping the labels) so that the 0-th leaf is "down" (if our semi-infinite leaves ended in vertices, so that the trees were trees in the sense of graph theory, this would be choose the vertex of the 0-th leaf as the root of the tree.) Now we can do the following. We read the tree "from the rooot up" and upon encountering a vertex we have some number of child subtrees and we insert a pair of parenthesis around the group of leaf labels of each subtree. Proceeding this way until we have processed every vertex, we get an (possibly incomplete) parenthesizing of the k labels $1, \cdot, k$. The open face corresponds to $(1, 2, \cdot, k)$ and the vertices correspond to complete parenthesizings. It is not hard to see that conversely, each (possibly incomplete) parenthesizing one gets a "parsing tree". This gives alternative combinatorial labels to the faces of the associahedron.

Note that this begins to explain the name associahedron. The vertices are labeld by complete parenthecisings, which are counted by Catalan numbers (and we don't mean un, dos, tres, quatre, cinc, sis...), as are more than six dosen other things ([Stan99]). One face is contained in the other if its label can be obtained by adding some more parenthesis to the other's label. This description produces the same polytope, of course. (Perhaps a deeper, albeit related, reason for the name is the fact that Stasheff associahedra form a topological non-Σ operad, which is a resolution of the associative operad; thus modules over this operad are precisely the spaces that have a "homotopy invariant version" of an associative product on them; see [MSS07]). Remark: The statement of this exercise gives a description of a poset; of course it is not a priori clear that there exists a polytope uderlying this poset. One can verify that this poset satisfies axioms of an abstract polytope, and therefore can be realized by a general construction in some high dimensional euclidean space. However this is a moot point, since there are several direct constructions of realizations of K_n as convex polytope in $\mathbb{R}^{n-2}$ (see [CZ12]).

(2) Let $\mathcal{R}_k := \{(x_1, \ldots, x_k) \in \mathbb{R}^k : x_1 < \ldots < x_k\}/Stab_\infty(SL(2, \mathbb{R}))$, where the subgroup $Stab_\infty(SL(2, \mathbb{R}))$ is the stabilizer of infinity inside $SL(2, \mathbb{R})$, viewed as Möbius transformations, acts by translations and (positive) scalings. Show that there is a

well-defined embedding

$$i : \mathcal{R}_k \hookrightarrow [0, \infty]^{\binom{k}{3}},$$

where the coordinate indexed by the subset $\{\ell < m < n\} \subset \{1, \ldots, k\}$ is given by the corresponding *cross ratio* (see [MS04, Appendix D.1]), $(x_n - x_m)/(x_m - x_\ell)$. Remark: This is the same as moduli space of $k + 1$ cyclically ordered points on $S^1 = \partial \mathbb{D}^2$ up to holomorphic reparametrizations - we simply fixed x_0 to be at infinity.

First, we check that the map is well-defined. Indeed translations and scalings do not change the cross ratio $(x_n - x_m)/(x_m - x_\ell)$, so the map is defined on the quotient. Second, we check that it's an embedding. For this it is convenient to fix a representative of the equivalence class (we will somtimes refer to this as "choosing gauge"). One way to do this is to note that we can always translate so that $x_1 = 0$ and then rescale so that $x_k = 1$ (from the perspective of holomorphic discs this is zooming in on the disk containing $x_0 = \infty$). Then the cross ratio $(x_k - x_j)/(x_j - x_1) = \frac{1}{x_j} - 1$ so the image determines all x_j's, and the map is injective.

(This is of course not the only way. We could fix $x_1 = 0$ and $x_2 = 1$, thus "focusing on the disc containing x_1"; that works as well.)

(3) Define $\bar{\mathcal{R}}_k$ to be the closure of the image of i. This is the *Deligne–Mumford compactification* of $\mathcal{R}_k$. Prove that it is isomorphic to K_k in a way that preserves the natural stratifications.

To see what the closure is, consider a sequence $\vec{x}^j \in \mathcal{R}_k$ with converging images $i(\vec{x}^k)$. We can switch to different "choices of gauge" to understand the convergence. To begin, we "fix gauge" by replacing each $\vec{x}^j$ with y^j in the same orbit of $Stab_\infty(SL(2, \mathbb{R}))$ with $y_1^j = 0$ and $y_k^j = 1$, as in the previous part. (This is "the view of the bubble tree from the point of view of the marked point at infinity".) Then, again as in the previous part, the convergence of cross ratios implies that for each fixed j the sequence y_j^k converges. Let the corresponding limits be y_j^∞ (each in $[0, 1]$). For those tipples of indexes $l < m < n$ for which y_l^∞ and y_n^∞ are different, the corresponding cross-ratios are guaranteed to converge (possibly to 0 or ∞). Thus all the indexes are grouped in consecutive bunches $(j + 1, \ldots, j + t)$ with all y_{j+r}^∞ equal for $r = 0, \ldots, t$ and different from y_{j-1}^∞ and y_{j+t+1}^∞. This allows us to start associating to our convergent sequence x^k a planar tree by assigning infinite leaves to the "singleton" bunches and interior vertices to bunches with more than one index. Then for each bunch with more than one index, we can pass to new gauge, in which $y_{j+1}^k = 0$ and $y_{j+t}^k = 1$ for all k and get convergence of coordinates of subbunches (the coordinates of all indexes not in the bunch will converge to infinity), to which we can then assign infinite leaves and internal vertices etc. Proceeding by induction, we see that any sequence $\vec{x}^k$ with convergent images $i(\vec{x}^k)$ determines a planar tree, according to which x_j^{k}'s are coming together. For indexes corresspoding to leaves incident on different internal vertices of the tree the cross ratios will converge to zero

or infinity, and for the ones incident to the same vertex they will converge to a finite value. Thus for a fixed tree structure one gets an open ball of limit points, and for each innermost bunch of size t we have choice of $t-2$ coordinates in $(0,1)$, thus the ball is of dimension $k-1$ minus number of internal vertices, which matches dimension of the corresponding stratum in K_k (alternatively, simply note that dimensions of the strata associated to the star S_{k+1} match and then drop by one for each expansion in both cases). This gives a cell complex identification between $\bar{\mathcal{R}}_k$ and K_k, as expected.

(4) Now consider the analogous construction, but where we do not impose an ordering condition on the x_i: call the resulting topological space $\bar{\mathcal{R}}_k^{unord}$. Is the resulting i_{unord} an embedding? Identify the topological space $\bar{\mathcal{R}}_4^{unord}$.

No, i_{unord} is not an embedding, since $i_{unord}(x) = i_{unord}(-x)$. In fact, i_{unord} is 2-1 (on the space of orbits of $Stab_\infty(SL(2,\mathbb{R}))$). Indeed, if $r = i_{unord}(x) = i_{unord}(-x)$ suppose $r = i_{unord}(y)$. Then y_1 and y_k are in the same order as in either x_1, x_k or $-x_1, -x_k$. Then by fixing gauge as in previous part, we see that y is either in the orbit of x or in the orbit of $-x$.

We see that $\bar{\mathcal{R}}_4^{unord}$ is made from $4! = 24$ copies of $\bar{\mathcal{R}}_4$, corresponding to the reorderings of the 4 x_ks. These are identified in pairs in the image of i_{unord}, giving 12 pentagons. These glue to a non-orientable surface of Euler characteristic -3. This is a 4-point blow up of $\mathbb{R}P^2$ or a 3 point blow up of a torus - corresponding to this being the real part of the moduli space $\bar{\mathcal{M}}_{0,5} = \mathbb{C}P^2 \# 4\overline{\mathbb{C}P}^2 = (\mathbb{C}P^1 \times \mathbb{C}P^1)\#3\overline{\mathbb{C}P}^2$, see [MS04], section D7 and [Kap93], Example 4.9.

The significance of this for Fukaya categories is as follows. We regard $\mathbb{R}$ as the boundary of the upper half-plane $\mathbb{H} := \{z \in \mathbb{C} : im(z) \geq 0\}$. Any complex disc with $k+1$ ordered boundary marked points is biholomorphic to $\mathbb{H}\cup\{\infty\}$, with the first boundary marked point corresponding to ∞ and the remaining ones corresponding to points $x_1 < \ldots < x_k \in \mathbb{R} = \partial\mathbb{H}$. Thus $\mathcal{R}_k$ is isomorphic to the moduli space of complex discs with $k+1$ boundary marked points, which are precisely the domains of the J-holomorphic maps whose count defines an A_∞ structure map m^k in the Fukaya category.

The codimension-1 boundary components of the Gromov compactification of a one dimensional moduli space of J-holomorphic discs are of two types. Firstly, the complex structure on the domain may degenerate: i.e., it may approach a point in the boundary $\bar{\mathcal{R}}_k \setminus \mathcal{R}_k \cong K_k \setminus int(K_k)$. Since we are considering a one-dimensional moduli space, generically it will avoid the faces of codimension 2 and only run into the faces of codimension 1. These correspond precisely to the terms in the A_∞ relations which do not involve $m^1 = \partial$. Secondly, we may have concentration of energy at a boundary puncture, as happened in the proof that $\partial^2 = 0$: these boundary points correspond to the remaining terms in the A_∞ relations.

8. Floer theory in toric varieties

(1) Calculate the disc potential function for Lagrangian torus fibres in the following cases:
   (a) $\mathbb{C}P^1$.

   The moment polytope is a segment. It has two faces (aka vertices), and indeed there are 2 discs of Maslov index 2, with areas adding to the total area of $\mathbb{C}P^1$ which we denote by A. If the coordinate on the segment is r, then $\lambda = T^r + T^{A-r} = z + \frac{T^A}{z}$. (In this and later parts of this exercise, we omit the proof that these discs are regular.)

   (b) $\mathbb{C}P^1 \times \mathbb{C}P^1$ (with equal areas for the two factors).

   The moment polytope is a square; if the coordinates on it are r_1 and r_2 then there are 4 discs of Maslov index 2 and their areas are $r_1, A - r_1, r_2$ and $A - r_2$. Then $\lambda = T^{r_1} + T^{A-r_1} + T^{r_2} + T^{A-r_2} = z_1 + \frac{T^A}{z_1} + z_2 + \frac{T^A}{z_2}$

(2) Hence work out which torus fibres have non-vanishing Floer cohomology in these cases.

   For $\mathbb{C}P^1$:
   $\frac{\partial \lambda}{\partial z} = 1 - \frac{T^A}{z^2}$; critical point is $z = \pm T^{A/2}$. Thus the critical point is the equatorial fiber with 2 local systems - one trivial and one with monodromy -1.
   For $\mathbb{C}P^1 \times \mathbb{C}P^1$:
   $\frac{\partial \lambda}{\partial z_1} = 1 - \frac{T^a}{z_1^2}$, $z_1 = \pm T^{a/2}$
   $\frac{\partial \lambda}{\partial z_2} = 1 - \frac{T^a}{z_2^2}$, $z_2 = \pm T^{a/2}$
   Thus the critical point is the "middle" (monotone) fiber (with 4 local systems).

(3) In the case of $\mathbb{C}\mathbb{P}^1$, verify your answer to the previous question using the combinatorial model for $HF^\bullet(L, L) \cong HF^\bullet(L, \psi(L))$ from §1.

   The equatorial circle has a Hamiltonian perturbation which intersects itself at 2 points, and there are 2 lunes/discs between these intersection points which cancel in the differential of the Floer complex. Thus it's Floer cohomology is of rank 2.
   All other circles are Hamiltonian displaceable (see below), so have vanishing Floer cohomology.

(4) Prove that every torus fibre in $\mathbb{C}\mathbb{P}^1$ with vanishing self Floer cohomology (i.e., such that $HF^\bullet(L, L) \cong 0$) can be displaced by a Hamiltonian isotopy.

   One simple way to do this is to consider as a Hamiltonian a height function - that is linear function on $\mathbb{R}^3$ restricted to $S^2 = \mathbb{C}\mathbb{P}^1$ - chosen in such a way that its zero set (in $\mathbb{R}^3$) runs through the center (in $\mathbb{R}^3$) of the circle we are trying to displace.

Such a Hamiltonian generates a rotation flow, which will eventually rotate the sphere by π and displace all non-central fibers from themselves.

(5) In fact, prove that every torus fibre in $\mathbb{CP}^1$ with vanishing self Floer cohomology can be displaced by a Hamiltonian isotopy supported in $\mathbb{CP}^1 \setminus 0$ or $\mathbb{CP}^1 \setminus \infty$ (but not both).

For any point p one has that $S^2 \setminus p$ is symplectomorphic to $\mathbb{D}^2$ of the same symplectic area. Given a toric fiber L which is not the central one, one of the choices of $p = 0$ or $p = \infty$ this symplectomorphism takes L to a circle in $\mathbb{D}^2$ enclosing less than half the area of $\mathbb{D}$, and the other choice takes it to a circle enclosing area more than half of the area of $\mathbb{D}^2$. The one with smaller area is displaceable (this is kind of obvious, but see [AS], Proposition A 1), and the one with the bigger is not (this is obvious).

(6) Use the previous exercise to show that every torus fibre in $\mathbb{CP}^1 \times \mathbb{CP}^1$ with vanishing self Floer cohomology can be displaced by a Hamiltonian isotopy.

This is quite easy to show directly - all fiber tori in $\mathbb{CP}^1 \times \mathbb{CP}^1$ are products of fiber tori in the two $\mathbb{CP}^1$ components. If L is not central, then at least one of its factors is not central in its $\mathbb{CP}^1$, and so can be displaced by a Hamiltonian H. Pulling H back to $\mathbb{CP}^1 \times \mathbb{CP}^1$ under projection gives a Hamiltonian displacing L.

Another way to do this, which generalizes well, is to say that our L_p - fiber over p in the moment polytope (which in this case is a square) - has p that is closer than half way to one of the faces of the square; let's say the top one (the other cases are similar). Then we consider the preimage K under momentum map of the half-open vertical segment passing throgh p, connecting it to the top face, and extending down to the bottom face but not including the point on the bottom face itself. This K is diffeomorphic to $(\mathbb{CP}^1 \setminus pt) \times S^1 = \mathbb{D}^2 \times S^1$ with the restriction of the symplectic form equal to the pill back of $dx \wedge dy$ from the $\mathbb{D}^2$ factor. L is then a product of the S^1 factor with a circle enclosing area less than half in $\mathbb{D}^2$. The pull back of the compactly suppported Hamiltonian on (the open) $\mathbb{D}^2$ can then be extended to a Hamiltonian on all of $\mathbb{CP}^1 \times \mathbb{CP}^1$ which displaces L (with the flow keeping it inside K, in fact).

(7) Use the same idea to show that every torus fibre in $\mathbb{CP}^2$ with vanishing self Floer cohomology can be displaced by a Hamiltonian isotopy.

Suppose p is "less than half way to the left face" - meaning the horizontal cut of the moment polytope triangle through p has less than half of it's length on the left. Then the same argument as above implies it's displaceable. Now for any non-central L_p we can change coordinates (recall the moment polytope is only defined up to integral affine equivalence) so that it is "less than half way to the left face". This means all non-central fibers are displaceable, as wanted.

Ramark: These non-central fibers can also be displaced by a Hamiltonians coming from action of $U(3)$ on $\mathbb{C}P^2$ (in fact the permutation matrices displace all non-central fibers).

This method of displacing torus fibres is called the method of 'probes' [McD11].

9. THE ELLIPTIC CURVE

This exercise is about the Fukaya category of an elliptic curve; be warned that it is more computationally intensive than the other exercises! It is based on [Zas05].

Let $(T, \omega) = (\mathbb{R}^2/\mathbb{Z}^2, \tau \cdot dx \wedge dy)$ be a symplectic torus of area τ. Consider the Lagrangians $L_k = \{y = kx\}$ for $k \in \mathbb{Z}$. Observe that the isomorphism

$$T \to T$$
$$(x, y) \mapsto (x, y + x)$$

takes L_k to L_{k+1} for all k. Therefore, we can define a map

$$HF^\bullet(L_0, L_j) \otimes HF^\bullet(L_0, L_k) \xrightarrow{\mathrm{id} \otimes T^j} HF^\bullet(L_0, L_j) \otimes HF^\bullet(L_j, L_{j+k}) \xrightarrow{m^2} HF^\bullet(L_0, L_{j+k}).$$

(1) Show that this product defines a bigraded algebra $\mathcal{A}_{\bullet,\bullet}$, where
$$\mathcal{A}_{j,k} := HF^j(L_0, L_k).$$

(2) Show that the corresponding graded algebra
$$\mathcal{C}_k := \mathcal{A}_{0,3k}$$
is isomorphic to $\Lambda[X_0, X_1, X_2]/p(X_0, X_1, X_2)$, where $|X_i| = 1$ and
$$p_\tau(X_0, X_1, X_2) = a_\tau X_0^3 + b_\tau X_1^3 + c_\tau X_2^3 + d_\tau X_0 X_1 X_2.$$
Compute $a_\tau, b_\tau, c_\tau, d_\tau$.

(3) Calculate the j-invariant of the elliptic curve defined by $\{p_\tau = 0\}$, as a function of τ.

REFERENCES

[AS]   M. Akveld and D. Salamon. Loops of lagrangian submanifolds and pseudoholomorphic discs. Geom. Funct. Anal. 11(4) (2001) 609650

[Auroux] Denis Auroux. A beginners introduction to Fukaya categories. Contact and symplectic topology. Springer International Publishing, 2014. 85-136.

[CZ12]   Cesar Ceballos and Günter Ziegler. Realizing the associahedron: Mysteries and questions. Associahedra, Tamari lattices and related structures. Springer Basel, 2012. 119-127.

[Che02]   Yuri Chekanov. Diffrential algebra of Legendrian links. Invent. Math., 150(3):441-483, 2002.

[dSRS14] Vin de Silva, Joel Robbin, and Dietmar Salamon. Combinatorial Floer homology. Mem. Amer. Math. Soc., 230(1080):v+114, 2014.

[Fl87]   Floer, Andreas. The unregularized gradient flow of the symplectic action. Communications on Pure and Applied Mathematics 41.6 (1988): 775-813.

[Fl88]   Andreas Floer. Morse theory for the symplectic action. J. Diff. Geom 28 (1988): 513-547.

[Fl89]   Andreas Floer. Witten's complex and infinite-dimensional Morse theory. J. Differential Geom., 30(1):207-221, 1989.

[Fr08]   Urs Frauenfelder. Gromov convergence of pseudoholomorphic disks. Journal of Fixed Point Theory and Applications 3.2 (2008): 215-271.

[FOOO10] Kenji Fukaya, Yong-Geun Oh, Hiroshi Ohta, and Kaoru Ono. Anchored Lagrangian submanifolds and their Floer theory. Contemp. Math 527 (2010): 15-54.

[FOOO14] Kenji Fukaya, Yong-Geun Oh, Hiroshi Ohta, and Kaoru Ono. Lagrangian Intersection Theory, Anomaly and Obstruction, Parts I and II. AMS/IP Studies in Advanced Mathematics, Amer. Math. Soc. and Internat. Press (2014).

[Kap93]   Mikhail M. Kapranov. The permutoassociahedron, Mac Lane's coherence theorem and asymptotic zones for the KZ equation. Journal of Pure and Applied Algebra 85.2 (1993): 119-142.

[LM]   Francois Lalonde and Dusa McDuff. The geometry of symplectic energy. Annals of Mathematics 141.2 (1995): 349-371.

[MSS07]   Martin Markl, Steven Shnider, and James D. Stasheff. Operads in algebra, topology and physics. No. 96. American Mathematical Soc., 2007.

[McD11]   Dusa McDuff. Displacing Lagrangian toric fibers via probes. In Low-dimensional and symplectic topology, Proc. Sympos. Pure Math., volume 82, pages 131-160. Amer. Math. Soc., Providence, RI, 2011.

[MS98]   Dusa McDuff and Dietmar Salamon. Introduction to symplectic topology. Oxford University Press, 1998.

[MS04]   Dusa McDuff and Dietmar Salamon. J-holomorphic Curves and Symplectic Topology. Amer. Math. Soc., 2004.

[PR14]   Leonid Polterovich and Daniel Rosen. Function theory on symplectic manifolds. American Mathematical Society, 2014.

[Sei08]   Paul Seidel. Fukaya categories and Picard-Lefschetz Theory. Eur. Math. Soc., 2008.

[Stan99] Richard P. Stanley, Exercises on Catalan and Related Numbers excerpted from Enumerative Combinatorics, vol. 2, Cambridge University Press 1999. Available online at .

[Sta63]   James D Stasheff. Homotopy associativity of H-spaces. I, II. Trans. Amer. Math. Soc., 108:293-312, 1963.

[UZ16]   Michael Usher and Jun Zhang. Persistent homology and Floer-Novikov theory. arXiv preprint arXiv:1502.07928 (2015).

[Wend16] Chris Wendl. Signs (or how to annoy a symplectic topologist). Blog post at https://symplecticfieldtheorist.wordpress.com/2015/08/23/

[Zas05]   Eric Zaslow. Seidel's mirror map for the torus. Advances in Theoretical and Mathematical Physics, 9(6):999-1006, 2005.