A Morse-theoretic approach to family Floer homology

Keeley Hoek

April 23, 2025

0. Review

- Roughly speaking, SYZ mirror symmetry begins with

$$\pi : X \to Q$$

  a (suitable) fibration by Lagrangian tori.

0. Review

- Roughly speaking, SYZ mirror symmetry begins with

$$\pi : X \to Q$$

a (suitable) fibration by Lagrangian tori.

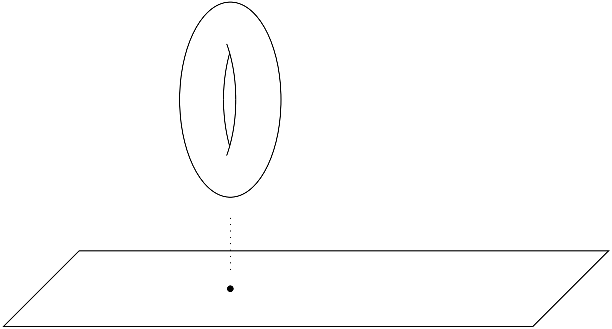

0. Review

- Roughly speaking, SYZ mirror symmetry begins with

$$\pi : X \to Q$$

a (suitable) fibration by Lagrangian tori.

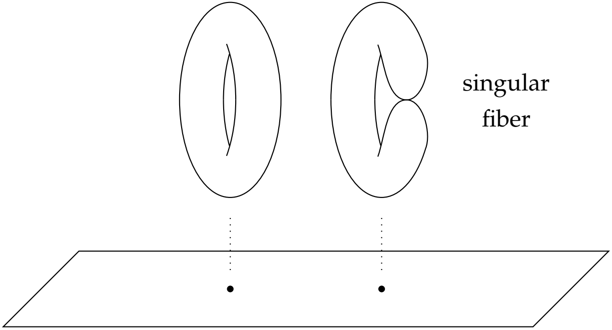

0. Review

- One then produces a dual torus fibration

$$\pi^\vee : X^\vee \to Q$$

  via a geometric recipe.
- The difficulty is that π may have singular fibers, and the construction of $X^\vee$ must be deformed accordingly.

0. Review

- One then produces a dual torus fibration

$$\pi^\vee : X^\vee \to Q$$

  via a geometric recipe.
- The difficulty is that π may have singular fibers, and the construction of $X^\vee$ must be deformed accordingly.

- On the other hand, HMS asserts

$$\mathrm{Fuk}(X) \underset{A_\infty}{\simeq} \text{“}\mathrm{D^b\,Coh}(X^\vee)\text{”}.$$

0. Review

- Family Floer theory builds a *rigid analytic mirror* $X_0^\vee$ over a local piece $Q_0 \subset Q$ as

$$X_0^\vee = \text{“moduli space of its points”} \overset{\text{(set)}}{=} \bigsqcup_{q \in Q_0} \mathrm{H}^1(F_q; U_\Lambda).$$

0. Review

- Family Floer theory builds a *rigid analytic mirror* $X_0^\vee$ over a local piece $Q_0 \subset Q$ as

$$X_0^\vee = \text{“moduli space of its points”}$$
$$\overset{\text{(set)}}{=} \bigsqcup_{q \in Q_0} \mathrm{H}^1(F_q; U_\Lambda).$$

Here $U_\Lambda = \mathrm{val}^{-1}(0) \subset \Lambda^*$ is the unitary subgroup of *Novikov field*

$$\Lambda = \left\{ \sum_{i=1}^{\infty} a_i T^{x_i} : a_i \in \Bbbk, x_i \in \mathbb{R}, \lim_{i \to \infty} x_i = \infty \right\}.$$

0. Review

- Family Floer theory builds a *rigid analytic mirror* $X_0^\vee$ over a local piece $Q_0 \subset Q$ as

$$X_0^\vee = \text{“moduli space of its points”} \overset{\text{(set)}}{=} \bigsqcup_{q \in Q_0} \mathrm{H}^1(F_q; U_\Lambda).$$

Here $U_\Lambda = \mathrm{val}^{-1}(0) \subset \Lambda^*$ is the unitary subgroup of *Novikov field*

$$\Lambda = \left\{ \sum_{i=1}^{\infty} a_i T^{x_i} : a_i \in \mathbb{k}, x_i \in \mathbb{R}, \lim_{i \to \infty} x_i = \infty \right\}.$$

- The space $X_0^\vee$ comes equipped with a comparison functor which can be used to (try to) prove HMS.

0. Review

- We take a Morse-theoretic approach; pick a suitable Morse function f on X.

Theorem

There is a curved A_∞-functor

$$\mathcal{C} : \mathcal{F}_{sec}(\pi, f) \to \text{mod-}\mathcal{A}(\pi, f).$$

0. Review

- We take a Morse-theoretic approach; pick a suitable Morse function f on X.

Theorem

There is a curved A_∞-functor

$$\mathcal{C} : \mathcal{F}_{sec}(\pi, f) \to \text{mod-}\mathcal{A}(\pi, f).$$

- In other words, a functor

$$\left\{ \begin{array}{c} \text{Fukaya category of} \\ \text{Lagrangian sections of } \pi \end{array} \right\} \to \left\{ \begin{array}{c} A_\infty\text{-modules for the} \\ \text{Morse–Fukaya algebra of } \pi \end{array} \right\}.$$

1. The Morse–Fukaya algebra $\mathcal{A}$

- This is an A_∞-algebra; for a single Lagrangian is due to Charest–Woodward, being in turn based on the ideas of Cornea–Lalonde and Fukaya–Oh–Ohta–Ono.

- Associated to a Lagrangian $L \subset X$ and choice of Morse function $f : L \longrightarrow \mathbb{R}$ is

$$\mathcal{A}(L, f) = \Lambda\langle \operatorname{crit} f \rangle,$$

graded by Morse index mod 2.

1. The Morse–Fukaya algebra $\mathcal{A}$

- This is an A_∞-algebra; for a single Lagrangian is due to Charest–Woodward, being in turn based on the ideas of Cornea–Lalonde and Fukaya–Oh–Ohta–Ono.

- Associated to a Lagrangian $L \subset X$ and choice of Morse function $f : L \longrightarrow \mathbb{R}$ is

$$\mathcal{A}(L, f) = \Lambda\langle \mathrm{crit}\, f \rangle,$$

  graded by Morse index mod 2.
- This algebra is equipped with a family of structure maps

$$\mu^d : \mathcal{A}^{\otimes d} \longrightarrow \mathcal{A}[2-d],$$

  which we now define.

1. The Morse–Fukaya algebra $\mathcal{A}$

- The basic objects we consider are *pseudoholomorphic treed disks*. These are continuous maps

$$u : \Delta \to X$$

from decorated domains Δ inductively built from the disk

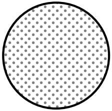

1. The Morse–Fukaya algebra $\mathcal{A}$

- The basic objects we consider are *pseudoholomorphic treed disks*. These are continuous maps

$$u : \Delta \to X$$

from decorated domains Δ inductively built from the disk

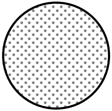

by attaching

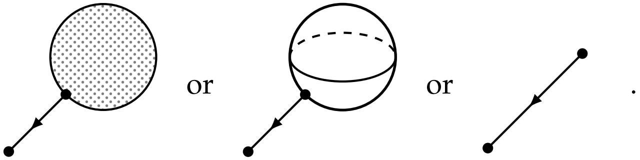

1. The Morse–Fukaya algebra $\mathcal{A}$

- An example treed disk domain:

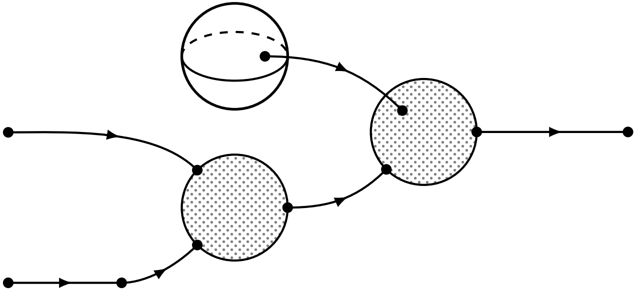

- Each edge e is attached interior-to-interior or boundary-to-boundary, and has a length $l(e) \in [0, \infty]$.

1. The Morse–Fukaya algebra $\mathscr{A}$

- We may write $\Delta = S_\Delta \cup T_\Delta$ as a union of the *surface* and *tree* parts, respectively.

- We require that $u : \Delta \to X$ obeys:
  1. *Pseudoholomorphic on the surface part*–we have
  $$J \circ \mathrm{D}u = \mathrm{D}u \circ j \quad \text{on } S_\Delta.$$
  2. *A Morse gradient flow on the tree part*–we have
  $$\frac{\mathrm{d}u}{\mathrm{d}t} = \nabla f \quad \text{on } T_\Delta.$$

1. The Morse–Fukaya algebra $\mathcal{A}$

- We may write $\Delta = S_\Delta \cup T_\Delta$ as a union of the *surface* and *tree* parts, respectively.
- We require that $u : \Delta \to X$ obeys:
  1. *Pseudoholomorphic on the surface part*–we have
  $$J \circ \mathrm{D}u = \mathrm{D}u \circ j \quad \text{on } S_\Delta.$$
  2. *A Morse gradient flow on the tree part*–we have
  $$\frac{\mathrm{d}u}{\mathrm{d}t} = \nabla f \quad \text{on } T_\Delta.$$
- Of course, in practice we will actually introduce domain-dependent perturbations of (J, f) into the equations to avoid transversality issues which arise.

1. The Morse–Fukaya algebra $\mathcal{A}$

- The basic idea, originally due to Cieliebak–Mohnke, is to solve this problem via *stabilizing divisors*.

1. The Morse–Fukaya algebra $\mathscr{A}$

- The basic idea, originally due to Cieliebak–Mohnke, is to solve this problem via *stabilizing divisors*.

Theorem (Charest–Woodward, Auroux–Muñoz–Presas)

Under suitable rationality assumptions on X and L, there exists a codimension 2 symplectic $D \subset X - L$, such that any J-holomorphic disk $u : (\mathbb{D}, \partial\mathbb{D}) \to (X, L)$ with $\omega([u]) > 0$ intersects D.

1. The Morse–Fukaya algebra $\mathscr{A}$

- The basic idea, originally due to Cieliebak–Mohnke, is to solve this problem via *stabilizing divisors*.

Theorem (Charest–Woodward, Auroux–Muñoz–Presas)

Under suitable rationality assumptions on X and L, there exists a codimension 2 symplectic $D \subset X - L$, such that any J-holomorphic disk $u : (\mathbb{D}, \partial\mathbb{D}) \to (X, L)$ with $\omega([u]) > 0$ intersects D.

Proof sketch.

Take an approximately holomorphic section of an ample line bundle on X concentrated on L, then perturb—the zero section gives D. □

1. The Morse–Fukaya algebra $\mathcal{A}$

- In particular, pseudoholomorphic treed disks $u : \Delta \to X$ will be:

1. The Morse–Fukaya algebra $\mathcal{A}$

- In particular, pseudoholomorphic treed disks $u : \Delta \to X$ will be:
  1. *stable*—disk and sphere components have "enough" special points, e.g. if $Du($$) = 0$ then  has at least 3 special points. In order to facilitate this, we introduce interior marked points ×, e.g.

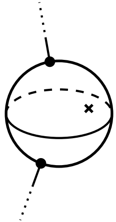

1. The Morse–Fukaya algebra $\mathscr{A}$

- In particular, pseudoholomorphic treed disks $u : \Delta \to X$ will be:
  1. *stable*—disk and sphere components have "enough" special points, e.g. if $Du($$) = 0$ then  has at least 3 special points. In order to facilitate this, we introduce interior marked points ×, e.g.

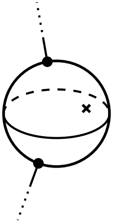

  2. *adapted to D*—each marked point × maps to D, and connected component of $u^{-1}(D)$ contains a marked point.

1. The Morse–Fukaya algebra $\mathcal{A}$

Definition

Fixing $\mathbf{x} = (x_0, \ldots, x_d) \in \operatorname{crit} f$ and $\beta \in \mathrm{H}_2(X, L)$ we may form

$$\mathcal{M} = \mathcal{M}(L, D, \mathbf{x}, \beta),$$

the *moduli space of all adapted stable pseudoholomorphic treed disks* $u : \Delta \to X$ which

- *have correct boundaries*—

$$u(\partial\Delta) \subset L \quad \text{for} \quad \partial\Delta = T_\Delta \cup \bigcup_{\substack{\mathbb{D} \subset \Delta \\ \text{a disk}}} \partial\mathbb{D},$$

- *have correct I/O*—$u(v_i) = x_i$ for v_i the ith bdry point, and
- *represent* β—

$$\sum_{C \subset \Delta} [u|_C] = \beta.$$

1. The Morse–Fukaya algebra $\mathcal{A}$

- We know that the expected dimension of the moduli space of pseudoholomorphic disks with n marked points and which represent $\beta \in \mathrm{H}_2(X, L)$ is

$$(n-3) + \mu(\beta) + (d+1),$$

essentially by the definition of the *Maslov class* $\mu(\beta)$. So, treed disks of this type contribute to a counting operation of degree $2 - d - \mu(\beta)$.

1. The Morse–Fukaya algebra $\mathscr{A}$

- We know that the expected dimension of the moduli space of pseudoholomorphic disks with n marked points and which represent $\beta \in \mathrm{H}_2(X, L)$ is

$$(n-3) + \mu(\beta) + (d+1),$$

essentially by the definition of the *Maslov class* $\mu(\beta)$. So, treed disks of this type contribute to a counting operation of degree $2 - d - \mu(\beta)$.

- The expected dimension of $\mathcal{M}$ is then

$$\dim \mathcal{M} = d - 2 + I(x_0) - \sum_{i=1}^{d} I(x_i) + \sum_{C \subset \Delta} I(u|_C).$$

1. The Morse–Fukaya algebra $\mathcal{A}$

- We could now proceed in the customary way to define the operations μ^k, if say L was equipped with a local system—if you have seen the definition of a Fukaya category before, you'll know that we are tantalizingly close.

- We are going to go in a slightly different direction.

2. A family version of $\mathcal{A}$

- The natural way to construct a family version of $\mathcal{A}$ is to consider $u : \Delta \to X$ with each disk boundary constrained to a (possibly different) fiber of π:

2. A family version of $\mathcal{A}$

- The natural way to construct a family version of $\mathcal{A}$ is to consider $u : \Delta \to X$ with each disk boundary constrained to a (possibly different) fiber of π:

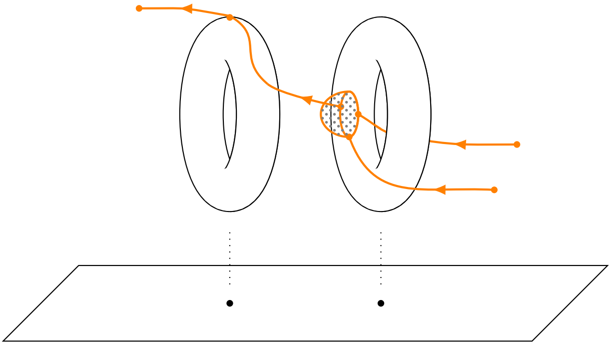

2. A family version of $\mathcal{A}$

- Suppose instead that we had chosen a Morse function f on all of X, and arranged that f lifted a Morse function on B.
- Also for simplicity, let's work over a simply connected compact piece $Q_0 \subset Q$, away from the singular fibers of π.

2. A family version of $\mathcal{A}$

- Suppose instead that we had chosen a Morse function f on all of X, and arranged that f lifted a Morse function on B.
- Also for simplicity, let's work over a simply connected compact piece $Q_0 \subset Q$, away from the singular fibers of π.
- We arrange a cellular decomposition $P^{[k]}$ of Q_0 such that:
  1. each k-cell $\sigma \in P^{[k]}$ contains a unique $q_\sigma \in \mathrm{crit}_k f$, and
  2. the union of the descending manifolds of all critical points contained in σ is σ itself.

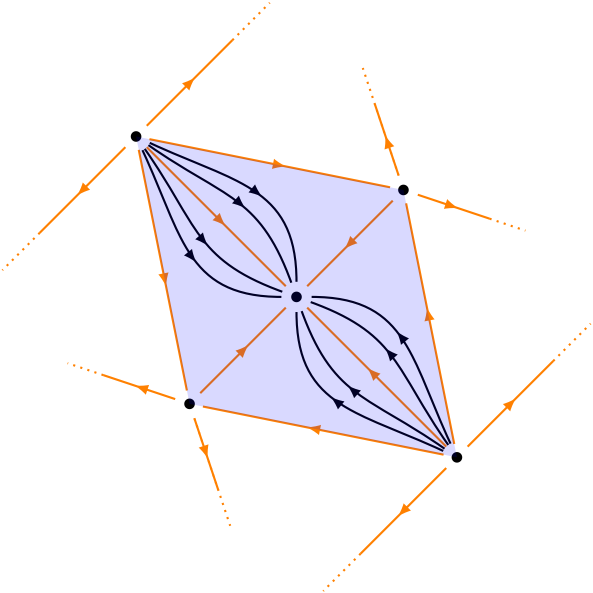

2. A family version of $\mathcal{A}$

- We need one final piece: the *Floer-theoretic weights*

$$z^\beta = T^{\omega(\beta)} \cdot \mathrm{hol}(\partial\beta)$$

  are analytic functions on $X_0^\vee$ for each $\beta \in \pi_2(X, F_q)$ by parallel transport $q \to p$.
- Recall that according to us, points of $X_0^\vee$ are elements of $\mathrm{H}^1(F_q; U_\Lambda)$, so hol is just fancy notation for evaluation.

2. A family version of $\mathcal{A}$

- Actually, essentially the same construction gives analytic charts on $X_0^\vee$: for a basis $\gamma_1, \ldots, \gamma_n$ of $\mathrm{H}_1(F_q)$, for each i parallel transport $q \to p$ causes γ_i to trace out a sheet α_i, to which we in turn associate

$$\left(T^{\omega(\alpha_1)} \operatorname{hol}(\gamma_1), \ldots, T^{\omega(\alpha_n)} \operatorname{hol}(\gamma_n)\right) \in (\Lambda^*)^n.$$

2. A family version of $\mathcal{A}$

- Actually, essentially the same construction gives analytic charts on $X_0^\vee$: for a basis $\gamma_1, \ldots, \gamma_n$ of $\mathrm{H}_1(F_q)$, for each i parallel transport $q \to p$ causes γ_i to trace out a sheet α_i, to which we in turn associate

$$\left(T^{\omega(\alpha_1)}\,\mathrm{hol}(\gamma_1), \ldots, T^{\omega(\alpha_n)}\,\mathrm{hol}(\gamma_n)\right) \in (\Lambda^*)^n.$$

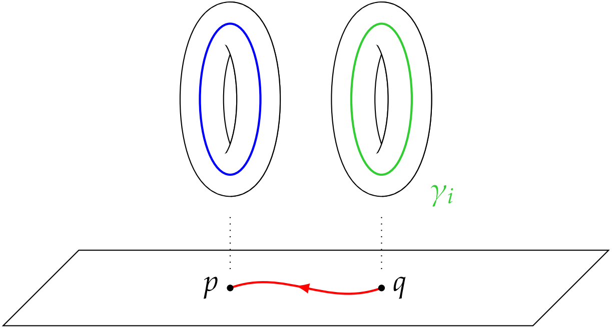

2. A family version of $\mathcal{A}$

- Actually, essentially the same construction gives analytic charts on $X_0^\vee$: for a basis $\gamma_1, \dots, \gamma_n$ of $\mathrm{H}_1(F_q)$, for each i parallel transport $q \to p$ causes γ_i to trace out a sheet α_i, to which we in turn associate

$$\left(T^{\omega(\alpha_1)}\,\mathrm{hol}(\gamma_1), \dots, T^{\omega(\alpha_n)}\,\mathrm{hol}(\gamma_n)\right) \in (\Lambda^*)^n.$$

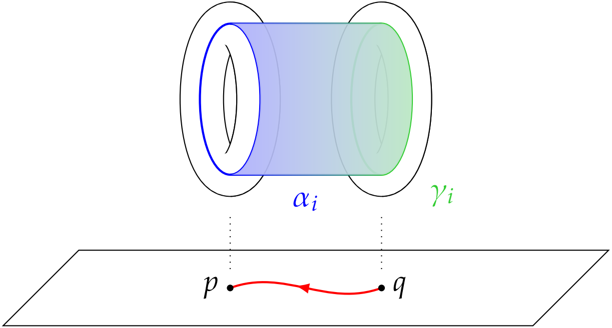

2. A family version of $\mathcal{A}$

- By suitably refining P by perturbing f, we can arrange that the collection of functions on $\pi^{-1}(\mathrm{star}(\sigma))$ assemble into a sheaf of universal weights

$$\mathcal{O}_{\mathrm{an}} = \pi^{\vee}_{*}(\mathcal{O}_{X_0^{\vee}}).$$

2. A family version of $\mathcal{A}$

- By suitably refining P by perturbing f, we can arrange that the collection of functions on $\pi^{-1}(\mathrm{star}(\sigma))$ assemble into a sheaf of universal weights

$$\mathcal{O}_{\mathrm{an}} = \pi_*^\vee(\mathcal{O}_{X_0^\vee}).$$

- Our algebra $\mathcal{A}$ is now an $\mathcal{O}_{\mathrm{an}}$-module.

2. A family version of $\mathcal{A}$

Definition

For $\mathbf{x} = (x_1, \ldots, x_n) \in (\text{crit } f)^n$ set

$$\mu^d(\mathbf{x}) := \sum_{x_0,\beta} \#\mathcal{M}_{d+1}(x_0, \mathbf{x}, \beta) \cdot z^\beta x_0,$$

where it is understood that the sum is taken over all (x_0, β) for which $\dim \mathcal{M}_{d+1}(x_0, \mathbf{x}, \beta) = 0$.

2. A family version of $\mathcal{A}$

Definition

For $\mathbf{x} = (x_1, \ldots, x_n) \in (\operatorname{crit} f)^n$ set

$$\mu^d(\mathbf{x}) := \sum_{x_0,\beta} \#\mathcal{M}_{d+1}(x_0, \mathbf{x}, \beta) \cdot z^\beta x_0,$$

where it is understood that the sum is taken over all (x_0, β) for which $\dim \mathcal{M}_{d+1}(x_0, \mathbf{x}, \beta) = 0$.

Theorem

The operations μ^d endow $\mathcal{A}$ with the structure of a (curved) A_∞-algebra, i.e. for homogeneous $a_1, \ldots, a_d$ we have

$$0 = \sum_{m+n\leq d} (-1)^\heartsuit \mu^{d+1-n}(a_1, \ldots, \mu^n(a_{m+1}, \ldots, a_{m+n}), \ldots, a_d)$$

with $\heartsuit = (-1)^{m+\sum_{i=1}^m |a_i|}$.

2. A family version of $\mathcal{A}$

Proof.

Analyze the boundary strata of the 1-dimensional moduli spaces $\mathcal{M}$; one shows that the only possible strata are of the type

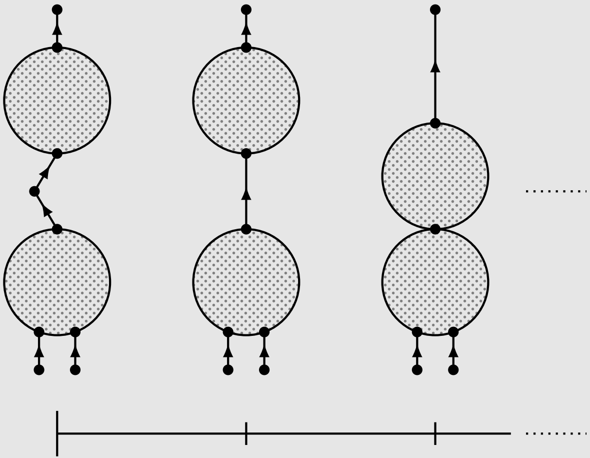

□

2. A family version of $\mathscr{A}$

- Switching to a family setting poses some significant technical challenges—just for example, no stabilizing divisor is disjoint from every fiber of π.
- So, we develop a scheme whereby divisors are turned on and off via a system of weights.

3. An HMS comparison functor

Definition

The category $\mathcal{F}_{\text{sec}}$ is the full subcategory of $\mathcal{F} = \text{Fuk}(X)$ of Lagrangian sections of π.

- Concretely and for simplicity, let $\{L_i\} \subset \mathcal{F}_{\text{sec}}$ be a finite family intersecting pairwise transversely. We set

$$\text{Hom}(L_i, L_j) = \begin{cases} \Lambda\langle L_i \cap L_j \rangle & i \neq j \\ \mathcal{A}(L_i) & i = j \end{cases}.$$

3. An HMS comparison functor

- In $\mathcal{F}_{\text{sec}}$ we compose $p_1 \in \text{Hom}(L_1, L_2)$ and $p_2 \in \text{Hom}(L_2, L_3)$ in the usual way:

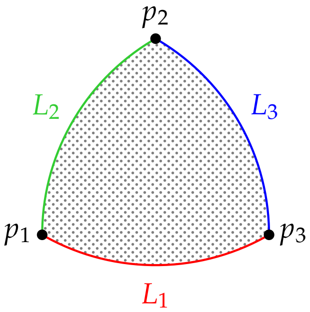

3. An HMS comparison functor

- The functor C on *objects*—set

$$L \in \mathcal{F}_{\mathrm{sec}} \longmapsto \mathcal{O}_{\mathrm{an}}\langle \mathrm{crit}\, f|_L \rangle.$$

3. An HMS comparison functor

- The structure maps

$$\triangleleft^{d-1} : \mathcal{C}(L) \otimes \mathcal{A}^{d-1} \to \mathcal{C}(L)[2-d]$$

now count pictures of the form (e.g. to compute $y_1 \triangleleft^1 x_1$):

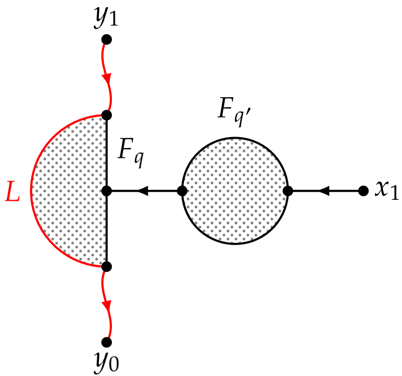

3. An HMS comparison functor

- The functor $\mathcal{C}$ on *morphisms*—given $p_i \in \text{Hom}(L_i, L_{i+1})$ we must specify

$$\mathcal{C}^n(p_1, \ldots, p_n)^{d-1} : \mathcal{C}(L_1) \otimes \mathcal{A}^{\otimes d-1} \to \mathcal{C}(L_n)[1-n-d].$$

For example, given $p \in \text{Hom}(L_1, L_2)$, compute $\mathcal{C}^1(p)^0(y_1)$ by counting:

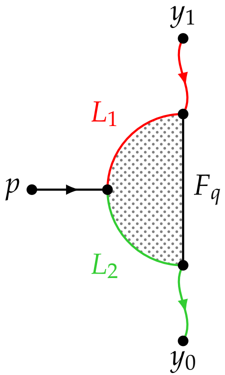

3. An HMS comparison functor

- One subtlety is that, in order to obtain the functor maps we desire, we must actually replace honest critical points of $f|_L$ with *anchors*.
- Fixing a distinguished $L_* \in \operatorname{ob} \mathcal{F}_{\text{sec}}$, an anchor (path) $\gamma : [0, 1] \to F_q$ is just a path from $\gamma(0) \in \operatorname{crit} f|_L$ to $\gamma(1) \in L_*$ contained wholly in F_q.

3. An HMS comparison functor

- Each input x_i (from $\mathcal{A}$) induces a *base flow path*:

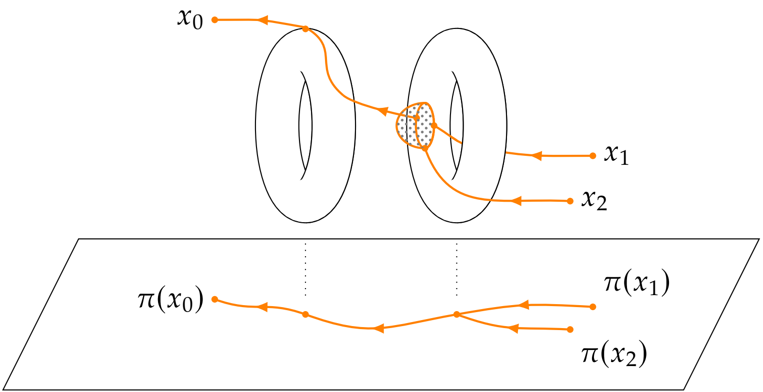

3. An HMS comparison functor

- Each input x_i (from $\mathcal{A}$) induces a *base flow path*:

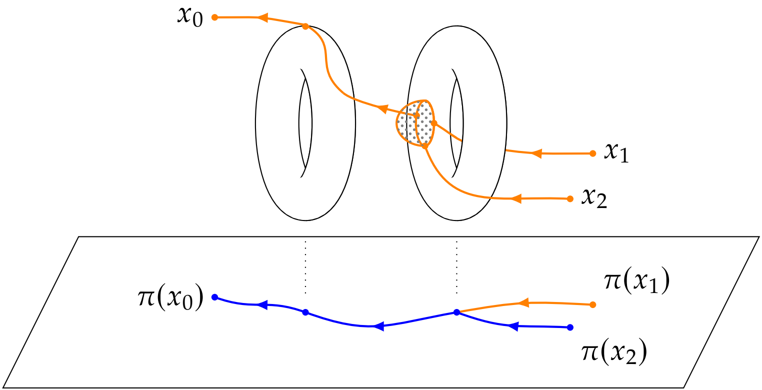

3. An HMS comparison functor

- Base paths act on anchors by parallel transport through fibers. We insert a correction by $T^{\omega(\alpha)}$, the area of the swept sheet:

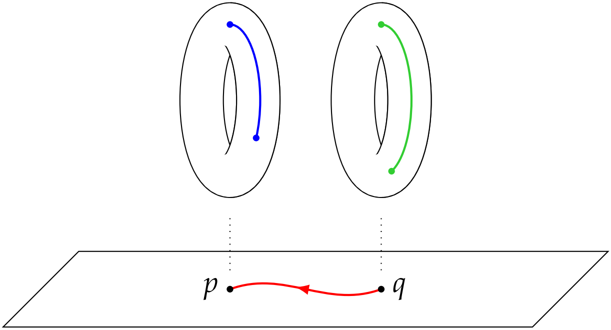

3. An HMS comparison functor

- Base paths act on anchors by parallel transport through fibers. We insert a correction by $T^{\omega(\alpha)}$, the area of the swept sheet:

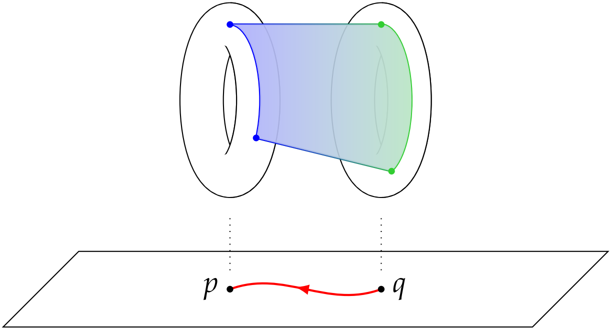

3. An HMS comparison functor

Theorem

There is a curved A_∞-functor

$$\mathcal{C} : \mathcal{F}_{sec}(\pi, f) \to \text{mod-}\mathcal{A}(\pi, f).$$

3. An HMS comparison functor

Theorem

There is a curved A_∞-functor

$$\mathcal{C} : \mathcal{F}_{sec}(\pi, f) \to \text{mod-}\mathcal{A}(\pi, f).$$

Proof.

We again verify the A_∞-relations by examining boundary strata. As an example, in the case of $\mu^0 = 0$, the module map $y_1 \triangleleft^1 (x_1, x_2)$ gives a homotopy between

$$(y_1 \triangleleft^1 x_1) \triangleleft^1 x_2 \quad \text{and} \quad y_1 \triangleleft^1 \mu^2(x_1, x_2).$$

(continued)

Proof (continued).

This corresponds to the two possible breakings:

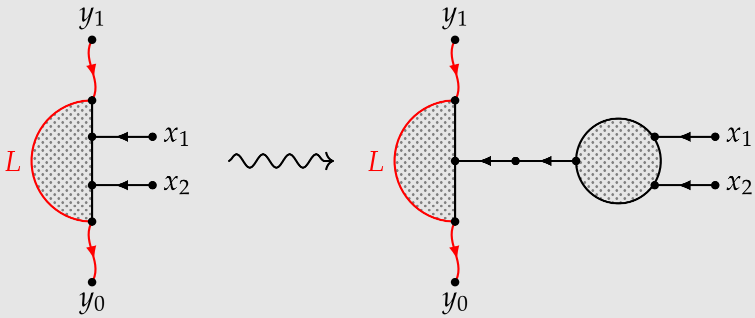

(continued)

Proof (continued).

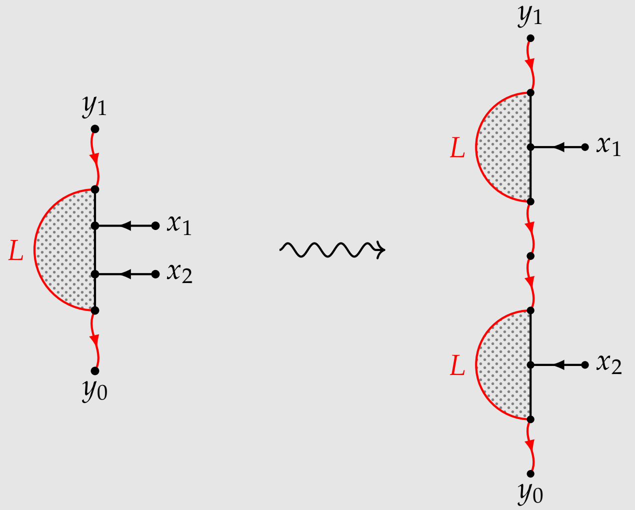

End