Math 91r final paper

Jessica J. Zhang

May 7, 2023

My goal in this paper is to explain (in somewhat loose, intuitive terms) the relationship between knot Floer homology and grid homology. In particular, we begin by presenting the original definition of knot Floer homology. This original definition involves fixing a Heegaard diagram of a knot, then counting so-called pseudo-holomorphic disks. After defining the simplest version of knot Floer homology, we define grid homology and, in doing so, show how it can be seen as the knot Floer homology of a particular Heegaard diagram where the surface is a torus. While many of the statements here are mentioned without proof, my hope is that this gives some sense of how the two homology theories are related.

1 Knot Floer homology

1.1 Heegaard diagrams

Consider a Morse function $f : X \to \mathbb{R}$, i.e., one whose critical points are all nondegenerate on some 3-manifold X. We may ask that f is self-indexing so that $f(p) = \operatorname{index}(p)$ for each critical point p. We may also ask, if X is connected and has no boundary, that f has a unique index-0 and index-3 critical point. Let $\Sigma = f^{-1}(3/2)$ be some level surface in between the index-1 and index-2 critical points. It separates X into two pieces $U_0 = f^{-1}((-\infty, 3/2])$ and $U_1 = f^{-1}([3/2, \infty))$. If Σ has genus g, then there will be exactly g index-1 and g index-2 critical points.

Suppose X has some Riemannian metric so that, in particular, the gradient vector field ∇f makes sense. Then consider the trajectories of $-\nabla f$. As seen in Figure 1, there are g red circles on Σ, each of which bound a disk in U_1. These red circles are the points of Σ that flow backwards

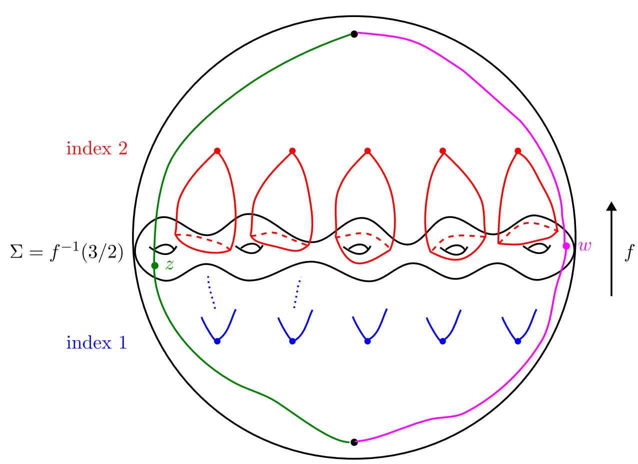

Figure 1: Getting a Heegaard diagram from a Morse function

to an index-2 critical point. Similarly, there are blue circles on Σ, bounding disks in U_0, which flow down to an index-1 critical point. All other points flow up to the index-3 critical point at the top, and down to the index-0 critical point at the bottom. Because of this, we know that the complement of the red disks in U_1 is just a ball, and similarly for the complement of the blue disks in U_0.

Now consider some points $z, w \in \Sigma$ which are not contained in either the red or blue curves. By following the gradient flow up and down to the index-3 and index-0 critical points, respectively, we get a knot K in X. This knot intersects Σ exactly at z and w.

When $X = S^3$, then we call this a *doubly-pointed Heegaard diagram* of the knot $K \subset S^3$. In particular, a doubly-pointed Heegaard diagram satisfies the following properties:

- There is a surface $\Sigma \subset X = S^3$ with genus $g \geq 0$ which separates S^3 into two handlebodies U_0 and U_1. (We ask that Σ is oriented as the boundary of U_0.)
- There are collections $\boldsymbol{\alpha} = \{\alpha_1, \ldots, \alpha_g\}$ and $\boldsymbol{\beta} = \{\beta_1, \ldots, \beta_g\}$ of pairwise disjoint, simple closed curves on Σ. The α_i's bound disks in U_0, while the β_i's bound disks in U_1. The complement of the α_i-disks in U_0 is a ball B^α, and similarly the complement of the β_i's in U_1 is a ball B^β.
- There are distinct points w and z disjoint from the α_i and β_i's.

This is a Heegaard diagram for a knot $K \subset S^3$ if $K \cap \Sigma = \{w, z\}$, where the intersection is positively oriented at w and negatively oriented at z. Furthermore, we ask that $K \cap B^\alpha$ and $K \cap B^\beta$ are intervals.

It turns out that any knot can be represented by a doubly-pointed Heegaard diagram. However, we will want a more general version, known as *multi-pointed (or $2k$-pointed) Heegaard diagrams*. Such a diagram is denoted $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$. In this case, we still have a surface Σ of genus g. However, our collections $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ each contain $g + k - 1$ curves, where $k \geq 1$. The complements are now the union of balls $B_1^\alpha, \ldots, B_k^\alpha$ and $B_1^\beta, \ldots, B_k^\beta$, respectively. Furthermore, we now have sets $\mathbf{w} = \{w_1, \ldots, w_k\}$ and $\mathbf{z} = \{z_1, \ldots, z_k\}$ of points on Σ which correspond to the positively and negatively oriented intersection points of $K \cap \Sigma$, respectively. We ask that each $K \cap B_i^\alpha$ and $K \cap B_i^\beta$ is an interval.

1.2 The knot Floer complex

At this point, we can briefly give some of the definitions of knot Floer homology. Let $d = g + k - 1$ be the number of α-curves on Σ. Then define

$$\mathrm{Sym}^d(\Sigma) = \Sigma^d / S_d,$$

where S^d acts on the Cartesian product Σ^d by permuting factors. Within the $2d$-dimensional manifold $\mathrm{Sym}^d(\Sigma)$, we may define the d-dimensional tori

$$\mathbb{T}_\alpha = \alpha_1 \times \cdots \times \alpha_d, \quad \mathbb{T}_\beta = \beta_1 \times \cdots \times \beta_d.$$

(Technically $\mathbb{T}_\alpha$ and $\mathbb{T}_\beta$ are the *projections* of these products onto $\mathrm{Sym}^d(\Sigma)$.)

It turns out that we can equip $\mathrm{Sym}^d(\Sigma)$ with a complex structure such that $\mathbb{T}_\alpha, \mathbb{T}_\beta$ are totally real tori. In this setting, we can define something called the *Lagrangian Floer homology* of $(\mathbb{T}_\alpha, \mathbb{T}_\beta)$. Roughly speaking, the complex in question is generated by the points $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, while the differential is obtained by counting so-called pseudo-holomorphic disks in $\mathrm{Sym}^d(\Sigma)$ whose boundaries are on $\mathbb{T}_\alpha$ and $\mathbb{T}_\beta$.

First, we briefly describe the complex structure on $\mathrm{Sym}^d(\Sigma)$. We consider Σ to be a Riemann surface. Then $\mathrm{Sym}^d(\Sigma) = \Sigma^d / S_d$ inherits a complex manifold structure from Σ. Loosely speaking, the local image is as follows: A small neighborhood in $\mathrm{Sym}^d(\Sigma)$ looks like a small neighborhood in $\mathrm{Sym}^d(\mathbb{C})$ since Σ locally looks like $\mathbb{C}$. But $\mathrm{Sym}^d(\mathbb{C}) \cong \mathbb{C}^d$ thanks to the fundamental theorem of algebra; in particular, a degree-d polynomial in $\mathbb{C}$ is uniquely determined by an unordered d-tuple of roots, which belongs to $\mathrm{Sym}^d(\mathbb{C})$, or by an ordered d-tuple of coefficients, which belongs to $\mathbb{C}^d$.

This complex manifold structure induces an almost-complex structure J on the tangent bundle. The tori $\mathbb{T}_\alpha, \mathbb{T}_\beta$ are then *totally real* in the sense that $J(T_{\mathbf{x}}\mathbb{T}_\alpha) \cap T_{\mathbf{x}}\mathbb{T}_\alpha = 0$ for each $\mathbf{x} \in \mathbb{T}_\alpha$, and similarly for $\mathbb{T}_\beta$.

As for the generators of the complex, note that any point on $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$ corresponds to a d-tuple of points in the intersections $\alpha_i \cap \beta_i$ for $i = 1, \ldots, d$. We may assume that the α- and β-curves intersect transversely, so that there are finitely many points of intersection $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$.

Say $\mathbf{x}, \mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$. Consider the set $\pi_2(\mathbf{x}, \mathbf{y})$ of homotopy classes of disks $u : D^2 \to \mathrm{Sym}^d(\Sigma)$ with $u(-1) = \mathbf{x}$, $u(+1) = \mathbf{y}$, and u taking the lower half of the boundary ∂D^2 to $\mathbb{T}_\alpha$ and the upper half to $\mathbb{T}_\beta$. If ϕ is some class in $\pi_2(\mathbf{x}, \mathbf{y})$, then a *pseudo-holomorphic representative* for ϕ is a disk

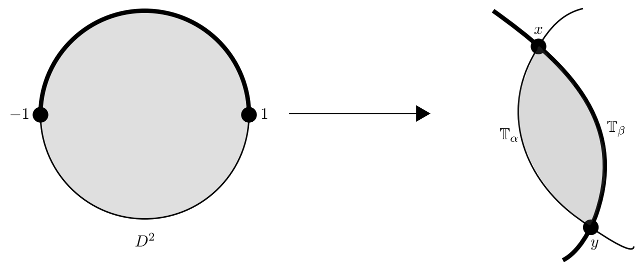

Figure 2: A pseudo-holomorphic disk

u with $\phi = [u]$ which satisfy the nonlinear Cauchy–Riemann equations for some path $J = (J_t)$ of almost-complex structures.

Now consider the moduli space $\mathcal{M}(\phi)$ of pseudo-holomorphic representatives for ϕ. There is an associated *Maslov index* $\mu(\phi) \in \mathbb{Z}$. It turns out that $\mathcal{M}(\phi)$ is actually a smooth manifold of dimension $\mu(\phi)$. Since we are interested primarily in counting distinct pseudo-holomorphic disks, we quotient out by the automorphisms of D^2 which fix ± 1. In particular, we get a space $\widehat{\mathcal{M}}(\phi)$ of dimension $\mu(\phi) - 1$. When $\mu(\phi) = 1$, this is just a discrete set of points. It turns out that it is also compact, however, and so it is finite, and we can count it. We now have a count $\#\widehat{\mathcal{M}}(\phi) \in \mathbb{Z}/2$ of pseudo-holomorphic disks. (If one accounts for sign, we can actually count this in $\mathbb{Z}$ instead.)

Each basepoint $v \in \{w_i, z_i\}$ gives us a $(2d-2)$-manifold $R_v = \{v\} \times \mathrm{Sym}^{d-1}(\Sigma) \subset \mathrm{Sym}^d(\Sigma)$. If $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ for intersection points $\mathbf{x}, \mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, then we define $n_v(\phi)$ to be the intersection number (in $\mathbb{Z}/2$) between ϕ and R_v.

At this point, one may define a bigrading on this complex, consisting of the *Maslov grading* and the *Alexander grading*. We will skip this part in this write-up, and continue on to defining the knot Floer complex in the case of crossing no basepoints.

1.3 The knot Floer complex, crossing no basepoints

There are several versions of the knot Floer complex, which correspond to different ways of keeping track of pseudo-holomorphic disks. We only introduce the simplest one, which is where we only count pseudo-holomorphic disks that don't pass through any of the basepoints w_i, z_i.

In particular, the complex $\widetilde{CFK}(\mathcal{H})$ of a Heegaard diagram $\mathcal{H}$ is generated as a $\mathbb{Z}/2\mathbb{Z}$-module by the intersection points $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$. Given $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, we define its differential to be

$$\partial \mathbf{x} = \sum_{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} \sum_{\substack{\phi \in \pi_2(\mathbf{x}, \mathbf{y}) \\ \mu(\phi)=1, n_{z_i}(\phi)=n_{w_i}(\phi)=0}} \#\widehat{\mathcal{M}}(\phi) \cdot \mathbf{y}.$$

Because $\mu(\phi) = 1$, we know that $\widehat{\mathcal{M}}(\phi)$ is 0-dimensional, and hence contains finitely many points. Of course, we have not defined the Maslov and Alexander gradings, but it turns out that the differential preserves the Alexander grading and drops the Maslov grading by 1.

It turns out that $\partial^2 = 0$, so this does indeed define a complex. We denote its homology $\widetilde{HFK}(\mathcal{H})$. It turns out that the bigrading from the Maslov and Alexander gradings descends to homology, so we write $\widetilde{HFK}_d(\mathcal{H}, s)$ for the part of $\widetilde{HFK}(\mathcal{H})$ with Maslov grading d and Alexander grading s.

If the Heegaard diagram is doubly pointed, i.e., if $k = 1$, then we denote the complex $\widehat{CFK}(\mathcal{H})$ and the homology $\widehat{HFK}(\mathcal{H})$.

Theorem 1.1. *The bigraded $\mathbb{Z}/2\mathbb{Z}$-module $\widehat{HFK}(\mathcal{H})$ is an invariant of the knot $K \subset S^3$. In particular, we may write $\widehat{HFK}(K)$ instead.*

In the case of multi-pointed Heegaard diagrams, we have the following theorem instead.

Theorem 1.2. *Let V be the bigraded $\mathbb{Z}/2\mathbb{Z}$-module freely generated by an element in bidegree $(-1, -1)$ and an element in bidegree $(0, 0)$. Then, with k the number of w-basepoints in the Heegaard diagram $\mathcal{H}$, we have*

$$\widetilde{HFK}(\mathcal{H}) = \widehat{HFK}(K) \otimes V^{\otimes(k-1)}.$$

As mentioned before, by choosing suitable orientations and counting everything with sign, we can make $\widetilde{HFK}(\mathcal{H})$ an abelian group, instead of just a $\mathbb{Z}/2\mathbb{Z}$-module.

2 Grid homology

We now turn to defining a combinatorial version of knot Floer homology known as grid homology, though we again only do it in the simplest version. A *planar grid diagram* G of size n is an $n \times n$ grid in the plane with X- and O-markings such that each row and each column contains exactly one X and exactly one O. To get from a grid to an oriented knot (or link), draw horizontal lines between the two marked squares in each row and vertical lines between the two marked squares in each column, as seen in Figure 3. At each crossing, we ask that the vertical segment goes over the

O		X		
	O		X	
		O		X
X			O	
	X			O

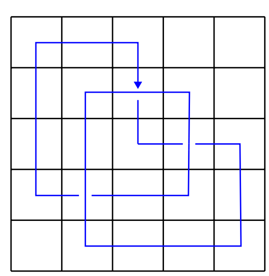

Figure 3: Getting a knot (here, a trefoil knot) from a grid

horizontal segment. We orient the knot so that we go from X to O in the columns, and from O to X in the rows. It turns out that any knot can be put into a grid diagram.

In grid homology, we prefer to think of these grids as *toroidal grid diagrams*, so that we glue the top and bottom edges, as well as the left and right edges. Any $n \times n$ toroidal grid diagram can be cut up into n^2 different planar grid diagrams; all these planar grid diagrams turn out to give isotopic knots. Connecting this back to knot Floer homology, we see that the vertical and horizontal lines correspond to α- and β-circles on the Heegaard surface $\Sigma = S^1 \times S^1$. The O- and X-markings correspond to the w- and z-basepoints.

Now the generators $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ are simply ways to put dots on the vertices of the grid so that there is one dot in each row and each column. That is, we choose points so that there is one on each α_i and one on each β_i. Any such arrangement of dots is called a *grid state*. The set of all grid states associated to a grid G is denoted $\mathbb{S}(G)$.

At this point, to continue our combinatorial construction, it is important to understand how we may count the pseudo-holomorphic disks between two grid states $\mathbf{x}$ and $\mathbf{y}$ using only the grid itself. First, suppose $\mathbf{x}$ and $\mathbf{y}$ are grid states. Then a *rectangle* between $\mathbf{x}$ and $\mathbf{y}$ is a rectangle in G, as seen in Figure 4, whose edges are on the horizontal and vertical gridlines such that the bottom left and top right corners are in $\mathbf{x}$, while the bottom right and top left corners are in $\mathbf{y}$. Furthermore, we ask that the other $n-2$ points of $\mathbf{x}$ and of $\mathbf{y}$ are the same.

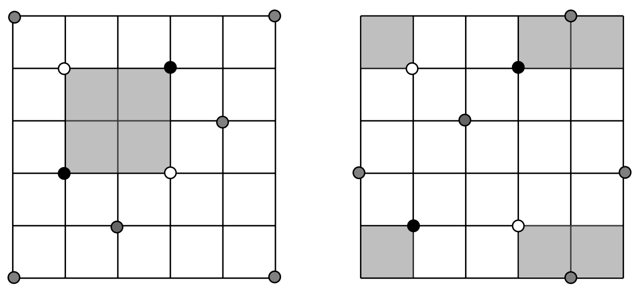

Figure 4: A rectangle on a grid diagram between $\mathbf{x}$, which consists of the black and gray dots, and $\mathbf{y}$, which consists of the white and gray dots.

Because these grid diagrams are toroidal, our rectangles can wrap around, as seen on the right side of Figure 4. If the interior of the rectangle contains no other dots in the grid state $\mathbf{x}$ (or, equivalently, in the grid state $\mathbf{y}$), as in the left side of Figure 4, then we call it an *empty rectangle*.

We denote the set of all rectangles between $\mathbf{x}$ and $\mathbf{y}$ to be $\mathrm{Rect}(\mathbf{x},\mathbf{y})$, while the set of all *empty* such rectangles is $\mathrm{Rect}^\circ(\mathbf{x},\mathbf{y})$. Each empty rectangle $r \in \mathrm{Rect}^\circ(\mathbf{x},\mathbf{y})$ has an associated $\phi \in \pi_2(\mathbf{x},\mathbf{y})$. In fact, it turns out that every $\phi \in \pi_2(\mathbf{x},\mathbf{y})$ with Maslov index 1 has an empty rectangle as its underlying domain, while the number of pseudo-holomorphic representatives of each $r \in \mathrm{Rect}^\circ(\mathbf{x},\mathbf{y})$ is odd. This implies that the differential operator in knot Floer homology is equal to the following differential operator:

$$\partial \mathbf{x} = \sum_{\mathbf{y} \in \mathbb{S}(G)} \sum_{\substack{r \in \mathrm{Rect}^\circ(\mathbf{x},\mathbf{y}) \\ O_i(r)=X_i(r)=0}} \mathbf{y},$$

where $O_i(r)$ (respectively, $X_i(r)$) is 0 if the i-th O-marking (respectively, X-marking) is contained in the interior of r and is 1 otherwise. (We can just label the O- and X-markings in any order from 1 to n; note that $O_i(r)$ and $X_i(r)$ correspond to $n_{w_i}(\phi)$ and $n_{z_i}(\phi)$.)

The *fully blocked grid chain complex* $\widetilde{GC}(G)$ is the $\mathbb{Z}/2\mathbb{Z}$-module generated by all $n!$ possible grid states, and whose differential is given by the equation above. The *fully blocked grid homology* $\widetilde{GH}(G)$ is just the homology of this complex.

It turns out that there is also a combinatorial way to define and calculate both the Maslov and the Alexander grading; we do not do so here, but using this one can show that $\widetilde{GH}(G)$ and

$\widetilde{HFK}(\mathcal{H})$ are indeed the same, where $\mathcal{H}$ is the Heegaard diagram represented by the grid G.

Then Theorem 1.2 implies the following:

Theorem 2.1. *If G is an $n \times n$ grid of a knot K, then*

$$\widetilde{GH}(G) = \widehat{HFK}(K) \otimes V^{\otimes(n-1)},$$

where V is the same bigraded, rank two $\mathbb{Z}/2\mathbb{Z}$-module as before.

There is in fact a purely combinatorial way to write $\widetilde{HFK}(K)$, known as the *simply blocked grid homology* $\widehat{GH}(K)$ of K. For this homology, the differential counts rectangles which may not contain any X-markings and may not contain the n-th O-marking, but may contain any of the other O-markings.

References

[Man16] Ciprian Manolescu. *An introduction to knot Floer homology.* 2016. DOI: 10.1090/conm/680/13701. URL: https://doi.org/10.1090/conm/680/13701.

[MOS09] Ciprian Manolescu, Peter Ozsváth, and Sucharit Sarkar. "A combinatorial description of knot Floer homology". In: *Annals of Mathematics* 169.2 (Mar. 2009), pp. 633–660. DOI: 10.4007/annals.2009.169.633. URL: https://doi.org/10.4007/annals.2009.169.633.

[OS07] Peter S. Ozsváth and Zoltán Imre Szabó. "An introduction to Heegaard Floer homology". In: 2007.

[OSS15] Peter Ozsváth, András Stipsicz, and Zoltán Szabó. *Grid Homology for Knots and Links.* American Mathematical Society, Dec. 2015. DOI: 10.1090/surv/208. URL: https://doi.org/10.1090/surv/208.