From Morse Theory to Heegaard Floer Homology

Owen Brass
Advisor: Dr. Irving Dai

Abstract

We provide an exposition of the construction and uses of Heegaard Floer homology, including applications to the structure of the homology cobordism group. We begin with the development of the Morse-Smale-Witten complex, proceed to discuss Lagrangian Intersection Floer homology, and then use these constructions to motivate Heegaard Floer homology and its involutive variant. We outline a proof of Furuta's Theorem on the existence of a $\mathbb{Z}^\infty$ subgroup of the homology cobordism group using involutive Heegaard Floer homology.

Contents

1 Introduction

Morse's theory of critical points has served as a pillar of the field of topology for decades. From the h-cobordism theorem of Smale, which gave as an impressive corollary the Poincaré Conjecture in dimensions ≥ 5 and arguably gave birth to the entire field of surgery theory, to the modern development of Floer theory, which has found innumerable applications throughout low-dimensional topology, knot theory, and symplectic topology, Morse theory finds its way into nearly any major undertaking in geometric topology. One such undertaking is the Heegaard Floer theory of Oszváth and Szabó. Their construction rests heavily upon Morse theory, both in a technical sense, by utilizing the Morse theoretic structure of 3-manifolds, and in an intuitive sense, by modeling its chain complex on the Morse-Smale-Witten complex which gives rise to Morse homology. In this paper, we will track just one of the trajectories of Morse theory, from its beginnings as a method of relating the topology of a manifold to critical points of a well-chosen function, to its role as the foundation of the modern theory of Heegaard Floer homology. We will also investigate some applications of Heegaard Floer homology to the homology cobordism group, a well-studied yet still mysterious object with an impressive range of applications, from triangulation of manifolds to the construction of exotic smooth structures on 4-manifolds.

1.1 Overview of Sections

Section 2 describes the fundamentals of Morse theory, up to and including the construction of the Morse-Smale-Witten complex, with the intent of illustrating how the essential structure of a manifold can be elucidated through critical point data.

Section 3 describes how to translate much of this reasoning over to the "infinite-dimensional" case of the space of paths between two Lagrangian submanifolds to determine essential information about their intersection (packaged in Lagrangian Intersection Floer homology).

Section 4 specializes Lagrangian Intersection Floer homology to the case of a certain pair of Lagrangian submanifolds extracted from the Heegaard diagram data of a 3-manifold (a construction due to Ozváth-Szabó called Heegaard Floer homology [**OS**]), and goes on to develop the "involutive" variant due to Hendricks-Manolescu [1], which is particularly suited to treating questions about the homology cobordism group.

Section 5 discusses the homology cobordism group, and proceeds to use Involutive Heegaard Floer homology to reprove Furuta's Theorem on the existence of a $\mathbb{Z}^\infty$ subgroup (following the proof provided in [2]).

2 Morse Theory

The great insight of Morse Theory (referred to by Bott [BOTT] as "the doctrine of critical point theory") is that the critical point data of a well-chosen smooth, real-valued function f on a manifold M can give one a lot of information about the structure of M. Such "good functions" are called Morse functions. In this section, we shall see that a Morse function provides a CW decomposition of its domain manifold M, and allows for the computation of a "Morse homology" which turns out to be precisely the singular homology of M.

2.1 Morse Functions and CW Structures on Manifolds

All manifolds are smooth (that is, having a maximal atlas of charts $\{(U_\alpha, \varphi_\alpha)\}_\alpha$ such that the transition functions $\varphi_\alpha \circ \varphi_\beta^{-1}$ are smooth of class C^∞) and closed (compact, boundaryless) unless otherwise specified.

Let M be a manifold, and $f : M \to \mathbb{R}$ a smooth real-valued function. A point $p \in M$ is called a *critical point* of f if $f \circ \varphi^{-1} : \mathbb{R}^n \to \mathbb{R}$ (where (U, φ) is a chart near p) has vanishing gradient at p (this is clearly independent of the choice of chart). p is called a *nondegenerate* critical point if the Hessian, i.e. the matrix of second partial derivatives, of $f \circ \varphi^{-1}$ is nonsingular at p (it is less trivial but still not difficult to verify that this is also independent of the choice of chart). A function f with only nondegenerate critical points is called a *Morse function.*

The *index* of a critical point is defined to be the dimension of the maximal subspace on which the Hessian matrix is negative definite. Intuitively, one may imagine that the index is the number of independent directions in which the Morse function $f : M \to \mathbb{R}$ decreases starting from p. In other words, if we imagine the Morse function to be a "height function" for M locally near p, then the index is the number of directions in which the manifold "slopes downwards," with the maximal number of directions being the dimension of M. The following lemma makes this intuition precise:

Lemma 1. *Let p be an index λ critical point of a Morse function $f : M^n \to \mathbb{R}$. Then there exists a local coordinate system $(x_1, \ldots, x_n)$ near p such that f can be written in the form:*

$$f(x_1, \ldots x_n) = f(p) - x_1^2 - \cdots - x_\lambda^2 + x_{\lambda+1}^2 + \cdots + x_n^2.$$

As a corollary, it follows that the critical points of a Morse function are isolated.

Given a Morse function $f : M \to \mathbb{R}$, we define a *sublevel set* $M_c := f^{-1}((-\infty, c])$. If $c \in \mathbb{R}$ is a regular value of f (that is, a value for which $f^{-1}(c)$ contains no critical points), it follows from the implicit function theorem that M_c is a manifold with boundary $\partial M_c = f^{-1}(c)$. Using the intuition of the Morse function f as a height function for M, M_c is the part of M which lies below height c.

Classical Morse theory is built upon the following two results, which together basically say the following: Given a Morse function $f : M \to \mathbb{R}$, the manifold M does not change in between critical points, and at a critical point of index λ, its homotopy type changes by the attachment of a λ-cell (note that the actual attaching map is a more complicated issue which is discussed further in the next subsection).

Theorem 1. *Let $f : M^n \to \mathbb{R}$ be a Morse function, and suppose that there are no critical points contained in $f^{-1}([a,b])$. Then $M_a = f^{-1}((-\infty, a])$ is diffeomorphic to $M_b = f^{-1}((-\infty, b])$.*

Theorem 2. *Let p be an index λ critical point of a Morse function $f : M^n \to \mathbb{R}$. Then for all sufficiently small $\epsilon > 0$, $M_{f(p)+\epsilon}$ is homotopy equivalent to $M_{f(p)-\epsilon}$ with a λ-cell attached.*

In particular (after drawing upon some homotopy theory of CW complexes) these two results tell us that every (smooth, closed) manifold M is homotopy equivalent to a finite CW complex (of dimension equal to the dimension of M), *assuming* there exists Morse function $f : M \to \mathbb{R}$. Luckily, Morse functions are abundant in the space of smooth functions $C^\infty(M)$, so that every manifold has many. More precisely:

Theorem 3. *The space of Morse functions is an open, dense subset of the space of C^∞ functions on any a smooth manifold M.*

We can already conclude some powerful results constraining the number of critical points of a Morse function on any manifold M – these are packaged in the form of the Morse inequalities:

Theorem 4. *Let $f : M \to \mathbb{R}$ be a Morse function. Let c_i denote the number of critical points of index i, and let b_i denote the ith Betti number of M. Then, for each $\lambda \geq 0$, we have the inequality:*

$$c_\lambda - c_{\lambda-1} + \cdots \pm c_0 \geq b_\lambda - b_{\lambda-1} + \cdots \pm b_0.$$

To see the power of this inequality, one may add the Morse inequality for the case $\lambda - 1$ to the case for λ to obtain that $c_\lambda \geq b_\lambda$, for all λ.

2.2 Gradient Flow Lines and the Morse-Smale-Witten Complex

In the last subsection, we saw that the critical points of a Morse function $f : M \to \mathbb{R}$ give rise to cells in a CW complex which is homotopy equivalent to M. In order to gain some geometric intuition about this fact, we need to investigate the gradient trajectories of the Morse functions – that is, solutions to the differential equation $D_t\gamma_x(t) = -\nabla f(\gamma_x(t))$, with initial condition

$\gamma_x(0) = x \in M$. This gives rise to a smooth flow $\varphi_t(x) = \gamma_x(t)$ on M. We define the *stable manifold* of a index λ critical point p to be

$$W^s(p) := \{x \in M \mid \lim_{t\to\infty} \varphi_t(x) = p\}$$

and the *unstable manifold* of p to be

$$W^u(p) := \{x \in M \mid \lim_{t\to-\infty} \varphi_t(x) = p\}.$$

It turns out that these are both indeed smoothly embedded manifolds, and are in fact smooth cells of dimension $n-\lambda$ and λ, respectively. In analogy with treating the Morse function $f : M \to \mathbb{R}$ as a height function, one can think of $W^s(p)$ as the set of points which would "fall down to" the flat point p, and $W^u(p)$ as the set of points which would "fall directly away from" the point p.

We will henceforth assume that our Morse functions satisfy the Morse-Smale condition, meaning that the stable and unstable manifolds intersect transversally.

We have the following more precise result on the CW structure of a manifold M given by a Morse-Smale function $f : M \to \mathbb{R}$:

Theorem 5. *There is a CW complex X and a homotopy equivalence $g : M \to X$ such that, given a critical point of index λ, $g(W^u(p))$ is contained in the base of a unique λ-cell. In this way, g establishes a bijective correspondence between the critical points of index λ in M and the λ cells in X.*

The Morse-Smale-Witten complex of the pair (M, f) is constructed to be chain isomorphic to the cellular chain complex of X, with boundary maps that have a geometric interpretation in terms of the stable and unstable manifolds of a Morse-Smale function f.

Given two critical points q, p of index λ and $\lambda - 1$, respectively, we can consider the manifold $W(q, p) := W^u(q) \cap W^s(p)$ (which is indeed a manifold by the Morse-Smale condition). The elements $x \in W(q, p)$ of this 1-dimensional manifold are precisely those which satisfy that $\lim_{t\to-\infty} \varphi_t(x) = q$ and $\lim_{t\to\infty} \varphi_t(x) = p$. It can be shown that $W(q, p) \cup \{q, p\}$ is a compact 1-dimensional manifold, with an $\mathbb{R}$-action given by flowing for time $t \in \mathbb{R}$. It follows that $\widehat{\mathcal{M}}(q, p) := W(q, p)/\mathbb{R}$ is a compact 0-dimensional manifold. Its elements are in one-to-one correspondence with flows $\gamma : \mathbb{R} \to M$ from q to p, which can be assigned orientations in a natural way (using the orientations of the stable and unstable manifolds). We define $n(q,p)$ to be a signed count of the elements in $\widehat{\mathcal{M}}(q,p)$.

We define the *Morse-Smale-Witten complex* as follows: Let $c_i(f)$ denote the set of critical points of index i. The chain groups of the Morse-Smale-Witten complex are given by

$$C_i(M,f) := \mathbb{Z}\langle c_i(f)\rangle$$

and the differential $\partial_i : C_i(M,f) \to C_{i-1}(M,f)$ is defined by

$$\partial_i(q) = \sum_{p \in c_{i-1}(f)} n(q,p)p.$$

It turns out that ∂_i^2 vanishes (as the gradient flow lines cancel in pairs), so $(C_*(M,f), \partial_*)$ is indeed a chain complex. Its homology is called the *Morse homology* of the pair (M,f).

The Morse Homology Theorem states that the Morse homology of (M,f) is actually isomorphic to the singular homology of M – in fact, the chain complex is indeed chain isomorphic to the cellular chain complex of $X \simeq M$.

A more complete discussion of the Morse-Smale-Witten complex can be found in [3].

3 Floer Theory of Lagrangian Intersections

Morse theory has the virtue of providing constraints on the analytic properties of a manifold M (namely, the number of critical points of Morse functions $f : M \to \mathbb{R}$) using computable algebraic objects (namely, the singular homology of M). In this brief section, we will model the construction of the Morse-Smale-Witten complex to produce an algebraic object which captures geometric information about the intersection between two (Lagrangian) submanifolds (of a symplectic manifold). This will involve the development of an "infinite-dimensional analogue" of Morse theory, which will be applied to the space of paths between the two submanifolds. In general, any such infinite-dimensional Morse theory is referred to as a "Floer theory."

Let (W, ω) be a symplectic manifold with symplectic form ω (a nondegenerate closed 2-form).

Let L, L' be *Lagrangian submanifolds* of W (half-dimensional submanifolds on which the form restricts to 0). We consider the pathspace

$$\mathcal{P}(L,L') := \{\gamma : [0,1] \to W \mid \gamma(0) \in L, \gamma(1) \in L'\}$$

. The objective is to apply Morse theory to understand the structure of $\mathcal{P}(L, L')$, in such a way that the critical points are replaced by the intersection points of L and L'.

Fixing some $\gamma_0 \in \mathcal{P}(L, L')$, define for any $\gamma \in \mathcal{P}(L, L')$ a "path of paths" $\Gamma : [0,1] \times [0,1] \to W$ which restricts to γ_0 when its second component is 0 and restricts to γ when its second component is 1.

We define an *action functional* $\mathcal{A}_{\gamma_0} : \mathcal{P}(L, L') \to \mathbb{R}$ by:

$$\mathcal{A}_{\gamma_0}(\gamma) = \int_{[0,1]^2} \Gamma^* \omega.$$

It can be shown that this is independent of the choice of Γ, and its differential is independent of γ_0.

We may treat $\mathcal{A}_{\gamma_0}$ like our Morse function from the previous section. It has critical points precisely at $L \cap L'$. The gradient flow lines between intersection points are called *pseudoholomorphic strips* $u : \mathbb{R} \times [0,1] \to W$, and they satisfy:

$$\partial_s u + J(u)\partial_t u = 0, \quad u(s,0) \in L, u(s,1) \in L' \text{ for all } s \in \mathbb{R}.$$

This allows us to produce a chain complex modeling the Morse-Smale-Witten complex, generated by intersection points $L \cap L'$, with boundary maps given by counting pseudoholomorphic strips. There are many technical complications in ensuring that this is well-defined – namely, the moduli space of pseudoholomorphic strips may not be compact, and the differential may not vanish (if it is defined at all). In the case that this is well-defined, the homology of this complex is called the *Lagrangian Intersection Floer Homology*, which we'll denote $HF(L, L')$. We will make use of this framework in the following section.

4 3-Manifolds and Heegaard Floer Theory

In this section, we will employ a form of Lagrangian Intersection Floer homology to construct a 3-manifold invariant called Heegaard Floer homology. This construction is due to Ozsváth and Szabó [4]. We will furthermore develop a variant called Involutive Heegaard Floer homology (due to Manolescu and Hendricks [1]), which in recent years has been successfully applied to questions about the homology cobordism group [2] [5]. We restrict our attention to rational homology spheres throughout this section for simplicity. The case $b_1(Y) > 0$ can be treated similarly, but some technical complications arise in counting holomorphic disks.

4.1 Structure of 3-Manifolds

It is a classical fact that all topological manifolds of dimension ≤ 3 admit one and only one smooth structure (up to diffeomorphism). Thus, we may treat all 3-manifolds as smooth 3-manifolds in a unique way; in particular, we may leverage the power of Morse theory for the study of all topological 3-manifolds.

We define a *handlebody* H to be a 3-dimensional manifold with boundary obtained from the 3-ball D^3 by attaching a finite number of copies of $D^2 \times I = D^2 \times [0,1]$ (called *handles*) so that $D^2 \times 0$ and $D^2 \times 1$ are glued to disjoint disks $D^2 \subset S^2 \cong \partial D^3$. The number g of handles is the *genus* of the handlebody. The boundary of a genus g handlebody is the Riemann surface of genus g, denoted here Σ_g.

Morse theory reveals something very important about the structure of 3-manifolds: they admit *Heegaard splittings*.

Theorem 6. *Every closed three manifold Y is homeomorphic to the union of two genus g handlebodies $H \cup_{\Sigma_g} H'$ glued along their boundary Σ_g. Such a decomposition is called a Heegaard splitting of Y of genus g.*

Proof. It is shown in [6] that one can choose a Morse function $h : Y \to \mathbb{R}$ to be *self-indexing*, meaning that given a critical point p of index λ, $h(p) = \lambda$. One can also choose h to have exactly one index 0 and exactly one index 3 critical point (namely, the points where f achieves its minimum and maximum, respectively). Then $H := h^{-1}([0, \frac{3}{2}])$ is a handlebody of genus g, where g is the number of index 1 critical points of h, and $h^{-1}(\frac{3}{2}) = \partial H$ is a Riemann surface of genus g. Choosing $f = -h+3$, we obtain that $f^{-1}(\frac{3}{2}) = h^{-1}(\frac{3}{2}) \cong \Sigma_g$, so that by the same reasoning $H' := f^{-1}([0, 3/2]) = h^{-1}([3/2, 3])$ is also a handlebody of genus g (note that g is also the number of index 1 critical points of f, hence the number of index 2 critical points of h). It follows that $Y = H \cup_{\Sigma_g} H'$, where Σ_g is identified with $h^{-1}(\frac{3}{2})$. $\square$

The information provided by a Heegaard splitting (of genus g) can be conveniently stored in the form of a *Heegaard diagram*, as follows: From each of the g handles (copies of $D^2 \times I$) contained in each of the two handlebodies

H, H', remove a small neighborhood of $D^2 \times \{\frac{1}{2}\}$ (say, $D^2 \times (\frac{1}{2} - \delta, \frac{1}{2} + \delta)$ for some small $\delta > 0$). What remains of the handlebodies H, H' after this removal are simply 3-balls B, B'. The boundary of each disc is a copy of S^1. Enumerate the copies of S^1 arising from each handle in each handlebody; typically one denotes the curves arising from one handlebody, say H, as $\alpha_1, \ldots, \alpha_g$, and the curves arising from the other as $\beta_1, \ldots, \beta_g$. Now, onto a drawing of the genus g surface Σ_g, draw the curves $\alpha_1, \ldots, \alpha_g, \beta_1, \ldots, \beta_g$ in the way that they are attached by the Heegaard splitting. Note that this completely determines the diffeomorphism type of the manifold resulting from this Heegaard splitting, as the gluing of the 3-balls B, B' (the only remaining step in the gluing of the handlebodies) is uniquely determined up to isotopy. An example of a Heegaard diagram is provided below.

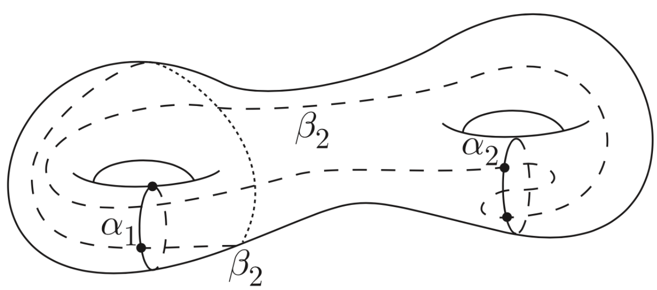

A genus 2 Heegaard splitting of S^3. Image credit: [7].

From a Morse theoretic standpoint, if the Heegaard splitting for a Heegaard diagram arises from a self-indexing Morse function $h : Y \to \mathbb{R}$ (with $H = h^{-1}([0, \frac{3}{2}])$), then an α_i curve is, up to isotopy, precisely the intersection of the stable manifold $W^s(p_i)$ for some index 1 critical point p_i with the level set $\Sigma_g = h^{-1}(\frac{3}{2})$. Similarly, a β_j curve is the intersection of the unstable manifold $W^u(q_j)$ for some index 2 critical point q_j with Σ_g. It follows that a point in $\alpha_i \cap \beta_j$ corresponds to where a gradient flow line from q_j to p_i passes through Σ_g.

Though a Heegaard diagram completely determines the diffeomorphism type of its corresponding 3-manifold, the converse does not hold (that is, the Heegaard diagram of a 3-manifold Y is not unique). Indeed, any 3-manifold admitting a genus g Heegaard splitting admits a genus k Heegaard splitting for all $k \geq g$, via a standard process of handle addition called *stabilization.* Isotopy of the curves α_i, β_j also does not change the diffeomorphism type of the resulting 3-manifold. Finally, there is a method of altering the configuration of the α_i, β_j curves called a *handleslide* which also does not alter the diffeomorphism type. As it turns out, any two Heegaard diagrams for a given 3-manifold can be related by a series of stabilizations, isotopies, and handleslides.

4.2 Heegaard Floer Homology

Our task in this subsection is to develop a topological invariant of a 3-manifold Y based on the data provided by its Heegaard diagram. The basic idea is that we want to capture information about the intersection of the α and β curves in a manner which does not depend on the chosen Heegaard splitting of Y. Unfortunately, as we remarked at the end of the last subsection, any given 3-manifold admits many Heegaard splittings, which can often look quite different.

In section 3, however, we used the model of Morse homology to develop machinery for extracting an algebraic object from intersection data. We will make use of this work here. To that end, the first step in our present construction is to produce from the Heegaard data a symplectic manifold with two intersecting Lagrangian submanifolds.

Assume Y is a compact, oriented 3-manifold, with a genus g Heegaard splitting $Y = H \cup_{\Sigma_g} H'$ induced by a self-indexing Morse function $h : Y \to \mathbb{R}$. We define the manifold

$$\mathrm{Sym}^g(\Sigma_g) := \left(\prod_g \Sigma_g \right) / S_g$$

where we are dividing by the natural action of the symmetric group S_g on the g-fold product of Σ_g given by permuting the components. In other words, $\mathrm{Sym}^g(\Sigma_g)$ consists of unordered g-tuples of points in Σ_g. It is not hard to see that $\mathrm{Sym}^g(\Sigma_g)$ is a smooth manifold, as a chart $\mathbb{C} \subset \Sigma_g$ provides a chart $\mathrm{Sym}^g(\mathbb{C}) \cong \mathbb{C}^{g-1} \subset \mathrm{Sym}^g(\Sigma_g)$, where the homeomorphism $\mathbb{C}^g \cong \mathrm{Sym}^g(\mathbb{C})$ is given by the map

$$(c_1, \dots, c_g) \mapsto \text{roots of } z^g + c_1 z^{g-1} + \cdots + c_{g-1} z + c_g.$$

It is not hard to show that a choice of complex structure J on Σ_g induces an almost complex structure $\tilde{J}$ on $\mathrm{Sym}^g(\Sigma_g)$. One can show that $\mathrm{Sym}^g(\Sigma_g)$ admits a symplectic structure (though note that this structure is not natural).

Now, as the curves $\alpha_1, \ldots, \alpha_g$ are disjoint from one another, the action of S_g has no effect on the torus $\mathbb{T}_\alpha := \alpha_1 \times \cdots \times \alpha_g \subset \prod_g \Sigma_g$ (i.e. the action identifies no distinct points), so we may consider $\mathbb{T}_\alpha$ to be a submanifold of $\mathrm{Sym}^g(\Sigma_g)$. The same reasoning applies to $\mathbb{T}_\beta := \beta_1 \times \cdots \times \beta_g$. There is no canonical choice of symplectic structure on $\mathrm{Sym}^g(\Sigma_g)$ which makes $\mathbb{T}_\alpha$ and $\mathbb{T}_\beta$ Lagrangian, but as we shall see, we do not need such a structure to carry out the construction (in fact, we will work around the "action functional" altogether).

Assuming that α_i and β_j intersect each other transversally for all i, j (which can be achieved by isotopy), it follows that $\mathbb{T}_\alpha$ and $\mathbb{T}_\beta$ intersect transversally in $\mathrm{Sym}^g(\Sigma_g)$, and hence their intersection is a 0-dimensional submanifold of $\mathrm{Sym}^g(\Sigma_g)$ (a discrete collection of points). A given intersection point $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ is a g-tuple of points $(x_1, \ldots, x_g)$ with $x_i \in \alpha_i \cap \beta_{\pi(i)}$ π, where π is a permutation of $\{1, \ldots, g\}$. As remarked in the last subsection, if our Heegaard splitting is induced by a Morse function $h : Y \to \mathbb{R}$, then a point in $\alpha_i \cap \beta_{\pi(i)}$ is precisely where a unique gradient flow line going from an index 2 critical point $q_{\pi(i)}$ to an index 1 critical point p_i intersects Σ_g. Thus, we can think of a point $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ as a selection of g gradient flow lines between distinct pairs of critical points (consisting of one index 1 and one index 2 critical point) of the Morse function h.

Fixing a complex structure on Σ_g, which induces a complex structure J on $\mathrm{Sym}^g(\Sigma_g)$, we can avoid the process of defining an "action functional" on the space of paths $\mathcal{P}(\mathbb{T}_\alpha, \mathbb{T}_\beta)$ as we did in the case of Lagrangian Intersection Floer homology, and jump straight to defining the "moduli spaces of gradient flow lines" $\mathcal{M}(\mathbf{x}, \mathbf{y})$ between intersection points $\mathbf{x}$ and $\mathbf{y}$ as follows: a map $u : \mathbb{R} \times [0, 1] \to \mathrm{Sym}^g(\Sigma_g)$ is in $\mathcal{M}(\mathbf{x}, \mathbf{y})$ if it is J-holomorphic, i.e. it satisfies

$$\partial_s u + J(u)\partial_t u = 0, \quad u(s, 0) \in \mathbb{T}_\alpha, u(s, 1) \in \mathbb{T}_\beta \text{ for all } s \in \mathbb{R},$$

and moreover, we have that, identifying $\mathbb{R} \times [0, 1]$ conformally with $\mathbb{D} \setminus \{\pm i\} \subset \mathbb{C}$, u extends to a map $\phi : \mathbb{D} \to \mathrm{Sym}^g(\Sigma_g)$ with $u(-i) = \mathbf{x}$ and $u(i) = \mathbf{y}$. In short, $\mathcal{M}(\mathbf{x}, \mathbf{y})$ consists of J-holomorphic strips from $\mathbf{x}$ to $\mathbf{y}$, just as in Lagrangian Intersection Floer homology. We will more often refer to ϕ as an element of $\mathcal{M}(\mathbf{x}, \mathbf{y})$ as opposed to u itself.

The holomorphicity condition leads to $\mathcal{M}(\mathbf{x}, \mathbf{y})$ admitting an $\mathbb{R}$-action, corresponding to the holomorphic automorphisms of $\mathbb{D}$ which fix i and $-i$ (the $\mathbb{R}$-action can be seen as horizontal translation under the identification of $\mathbb{D} \setminus \{\pm i\}$ with $\mathbb{R} \times [0, 1]$). We can then define the unparametrized moduli space

$$\widehat{\mathcal{M}}(\mathbf{x}, \mathbf{y}) := \mathcal{M}(\mathbf{x}, \mathbf{y})/\mathbb{R}.$$

A continuous map $\phi : \mathbb{D} \to \mathrm{Sym}^g(\Sigma_g)$ which satisfies all of the requirements for an element of $\mathcal{M}(\mathbf{x}, \mathbf{y})$ except for the holomorphicity requirement (i.e. it maps $-i$ to $\mathbf{x}$, i to $\mathbf{y}$, takes the left edge into $\mathbb{T}_\alpha$ and the right edge into $\mathbb{T}_\beta$) is called a *Whitney disk*. The set of homotopy classes of Whitney disks between $\mathbf{x}$ and $\mathbf{y}$ is denoted $\pi_2(\mathbf{x}, \mathbf{y})$. For a Whitney disk ϕ, we define $\mathcal{M}(\mathbf{x}, \mathbf{y}; \phi)$ to be the components of the moduli space $\mathcal{M}(\mathbf{x}, \mathbf{y})$ which contain elements homotopic to ϕ. To each Whitney disk ϕ, we may assign an index $\mu(\phi)$, called the *Maslov index*, which corresponds to the "expected dimension" of $\mathcal{M}(\mathbf{x}, \mathbf{y}; \phi)$. In analogy to Morse theory, one can think of $\mu(\phi)$ as measuring the relative index between two "critical points" $\mathbf{x}$ and $\mathbf{y}$ (along the disk ϕ). In constructing our chain complex for Heegaard Floer homology, we will only count components of moduli spaces $\mathcal{M}(\mathbf{x}, \mathbf{y}; \phi)$ for which $\mu(\phi) = 1$, which is justified by the following result, due to Ozsváth and Szabó [4].

Theorem 7. *If $\mu(\phi) = 1$ (and the genus of the Heegaard splitting is greater than 2), then, under appropriate perturbations, $\widehat{\mathcal{M}}(\mathbf{x}, \mathbf{y}; \phi) := \mathcal{M}(\mathbf{x}, \mathbf{y}; \phi)/\mathbb{R}$ is a compact, oriented, zero-dimensional manifold.*

In particular, the elements of $\widehat{\mathcal{M}}(\mathbf{x}, \mathbf{y}; \phi)$ can be counted meaningfully. At this point, we have enough information to define a complex with generators given by the intersection points in $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$ and with boundary map $\partial \mathbf{x}$ counting (with sign) the components of $\widehat{\mathcal{M}}(\mathbf{x}, \mathbf{y}; \phi)$ over all $\mathbf{y}$ and all ϕ with $\mu(\phi) = 1$ (c.f. the Morse-Smale-Witten complex). In essence, this is precisely what we shall do, but more structure is needed to make the homology of the complex a useful invariant.

4.2.1 Basepoints and Spin^c structures

The first step in enriching the structure of the complex generated by elements of the intersection $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$ is to observe that the complex admits a natural differential preserving grading by Spin^c structures. More precisely, we mean that the complex is the direct sum of summands which are indexed by Spin^c structures, such that the boundary map sends each summand into itself. We provide a brief review of Spin^c structures here, before explaining how they grade the Heegaard Floer complex.

Recall that the tangent bundle of an oriented, closed 3-manifold Y has structure group $SO(3)$, which can be identified with $U(2)/U(1)$, where $U(1)$

is embedded in $U(2)$ as the diagonal subgroup. The projection $p : U(2) \to SO(3)$ is a principal $U(1) = S^1$ bundle. Now, putting a Riemannian metric on Y, we may consider the frame bundle $f : Fr(Y) \to Y$, a principal $SO(3)$-bundle, over Y. A Spin^c *structure* on a smooth, closed, oriented 3-manifold Y is a lift of the frame bundle f to a principal $U(2)$-bundle. As such, a Spin^c structure consists of a principal $U(2)$-bundle $F \to Y$ such that $F/U(1)$ (obtained by the applying the projection $p : U(2) \to SO(3)$ fiberwise) is isomorphic to $Fr(Y)$.

There is an alternative definition of a Spin^c structure on an oriented, closed 3-manifold Y. We say that two nonvanishing vector fields u and v on Y are *homologous* if they are homotopic through nonvanishing vector fields on the complement of a ball $D^3 \subset Y$. Turaev [8] shows that there is a canonical bijection between the homology classes of nonvanishing vector fields on Y and the set of Spin^c structures on Y. We henceforth define a Spin^c structure on Y to be a choice of homology class of nonvanishing vector fields on Y. We will denote the set of Spin^c structures on Y by $\mathcal{S}(Y)$.

Recall that we want the differential for the Heegaard Floer complex to essentially count holomorphic disks between elements of the intersection $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$. If we partition $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$ according to the equivalence relation $\mathbf{x} \sim \mathbf{y}$ if and only if there exists a Whitney disk between $\mathbf{x}$ and $\mathbf{y}$ (i.e. if and only if $\pi_2(\mathbf{x}, \mathbf{y}) \neq \emptyset$), then it is clear that this partition induces a differential preserving grading on the complex. However, this is a challenging condition to verify explicitly – we need a (slightly weaker) computable condition.

A point $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ uniquely corresponds to a choice of g gradient trajectories (of the Morse function inducing the Heegaard splitting) between pairs of index 2 and index 1 critical points. This provides a 1-chain, denoted $\gamma_\mathbf{x}$, in $H_1(Y)$. It is not hard to show that, if there is a Whitney disk between $\mathbf{x}$ and $\mathbf{y}$, the 1-chain $\gamma_\mathbf{x} - \gamma_\mathbf{y}$ must be trivial in $H_1(Y)$ (i.e. null-homologous). This is explained in [7].

Now, we add a basepoint $z \in \Sigma_g \setminus \{\alpha_1, \ldots, \alpha_g, \beta_1, \ldots, \beta_g\}$ to our Heegaard diagram. We define a map

$$s_z : \mathbb{T}_\alpha \cap \mathbb{T}_\beta \to \text{Spin}^c(Y)$$

from the set of intersection points to the set of Spin^c structures on Y as follows: a point $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ specifies g gradient trajectories from index 2 to index 1 critical points, so that it "pairs up" index 2 critical points with index 1 critical points in such a way that no point (of either index) is paired twice. We can remove small neighborhoods (diffeomorphic to balls) around these trajectories, so that the only remaining points at which the gradient vector field ∇h (where h is the Morse function generating the Heegaard splitting) vanishes is at the minimum and maximum of h. There is exactly one gradient flow line from the maximum (index 3 critical point) to the minimum (index 0 critical point) passing through the chosen point z. Remove a small neighborhood (diffeomorphic to a ball) around this trajectory, so that what ∇h is nonvanishing on what remains of Y. Since, in each removed ball, the sum of the indices of the critical points is 3, the vector field ∇h can be modified in the balls to a nonvanishing vector field, say $v_{\mathbf{x}}$. We set $s_z(\mathbf{x})$ to be the homology class of $v_{\mathbf{x}}$, which is well-defined since it is equal to ∇h outside of a union of disjoint balls (note that a union of disjoint balls can always be isotoped to fit into a larger ball). This construction has the virtue that $v_{\mathbf{x}}$ and $v_{\mathbf{y}}$ are homologous if and only if $\gamma_{\mathbf{x}}$ and $\gamma_{\mathbf{y}}$ are homologous, and thus $s_z(\mathbf{x}) = s_z(\mathbf{y})$ if and only if $\gamma_{\mathbf{x}} - \gamma_{\mathbf{y}} = 0$.

Thus, with a fixed choice of basepoint z, the set of Spin^c structures on Y provide a natural differential preserving grading on $\mathbb{Z}\langle \mathbb{T}_\alpha \cap \mathbb{T}_\beta \rangle$ by partitioning $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$ into classes of the form $s_z^{-1}(\mathfrak{s})$, iterating over all $\mathfrak{s} \in \text{Spin}^c(Y)$.

Before we complete the definition of our chain complex(es), there is one final piece of data which we want to keep track of: intersection of the holomorphic disks with the hypersurface $\{z\} \times \text{Sym}^{g-1}(\Sigma_g) \subset \text{Sym}^g(\Sigma_g)$, whose Poincaré dual generates $H^2(\text{Sym}^g(\Sigma_g))$ (for $g > 1$). We denote by $n_z(\phi) \in \mathbb{Z}$ the oriented intersection number of a holomorphic disk $\phi : \mathbb{D} \to \text{Sym}^g(\Sigma_g)$ with the hypersurface $\{z\} \times \text{Sym}^{g-1}(\Sigma_g)$. Note that $n_z(\phi)$ is always nonnegative for holomorphic ϕ, since $\{z\} \times \text{Sym}^{g-1}(\Sigma_g)$ is a *complex* hypersurface (holomorphic trajectories always meet complex hypersurfaces positively, according to the natural orientation induced by the complex structure).

4.2.2 The Heegaard Floer Chain Complexes

We are now ready to construct the "infinity" Heegaard Floer chain complex $CF^\infty(Y, \mathfrak{s})$. We assume a fixed choice of basepoint $z \in \Sigma_g \backslash \{\alpha_1, \ldots, \alpha_g, \beta_1, \ldots, \beta_g\}$. We define $CF^\infty(Y, \mathfrak{s})$ to be the free abelian group with generators of the form $[\mathbf{x}, i]$ where $\mathbf{x} \in s_z^{-1}(\mathfrak{s}) \subset \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ and $i \in \mathbb{Z}$. We define a relative grading on this group as follows:

$$\text{gr}([\mathbf{x}, i], [\mathbf{y}, j]) = \mu(\phi) - 2n_z(\phi) + 2(i - j)$$

where $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ is any holomorphic disk between $\mathbf{x}$ and $\mathbf{y}$. It turns out that the number $\mu(\phi) - 2n_z(\phi)$ is independent of the choice of holomorphic representative, so this grading is well-defined. The differential ∂^∞ is given by

$$\partial^\infty([\mathbf{x}, i]) = \sum_{\mathbf{y} \in s_z^{-1}(\mathfrak{s})} \sum_{\phi \in \pi_2(\mathbf{x},\mathbf{y}) | \mu(\phi)=1} n(\mathbf{x}, \mathbf{y}; \phi)[\mathbf{y}, i - n_z(\phi)]$$

where $n_z(\phi)$ is a signed count of the elements in the unparametrized moduli space of holomorphic representatives homotopic to ϕ, $\widehat{\mathcal{M}}(\mathbf{x}, \mathbf{y}; \phi) := \mathcal{M}(\mathbf{x}, \mathbf{y}; \phi)/\mathbb{R}$. By work of Ozsváth and Szabó, it turns out that $(\partial^\infty)^2$ vanishes, so this is indeed a chain complex.

It is natural to consider the action U on this complex, given by

$$U \cdot [\mathbf{x}, i] = [\mathbf{x}, i - 1].$$

With this in mind, we may consider $CF^\infty(Y, \mathfrak{s})$ to be a $\mathbb{Z}[U]$-module. The homology of the complex $(CF^\infty(Y, \mathfrak{s}), \partial^\infty)$ is denoted $HF^\infty(Y, \mathfrak{s})$. We will refer to it as the "infinity" variant of Heegaard Floer Homology.

We can form other variants by taking advantage of the action U. Firstly, since $n_z(\phi)$ is always nonnegative, we have that the complex generated by elements of the form $[\mathbf{x}, i]$ with $i < 0$ is closed under the differential ∂^∞, and thus is a subcomplex of $(CF^\infty(Y, \mathfrak{s}), \partial^\infty)$. It is also closed under the action of U. We obtain the "minus" Heegaard Floer chain complex, $(CF^-(Y, \mathfrak{s}), \partial^-)$, where $CF^-(Y, \mathfrak{s})$ is considered as a $\mathbb{Z}[U]$-module, and ∂^- is simply the restriction of ∂^∞. Its homology, denoted $HF^-(Y, \mathfrak{s})$, is the "minus" variant of Heegaard Floer homology.

The quotient $CF^\infty(Y, \mathfrak{s})/CF^-(Y, \mathfrak{s})$ is called the "plus" Heegaard Floer chain complex, denoted $CF^+(Y, \mathfrak{s})$, with a corresponding "plus" variant of Heegaard Floer homology $HF^+(Y, \mathfrak{s})$. It is also a $\mathbb{Z}[U]$-module – in this case, the U-action has a nontrivial kernel, generated by elements of the form $[\mathbf{x}, 0]$. We call this kernel the "hat" variant of Heegaard Floer homology, $\widehat{HF}(Y, \mathfrak{s})$. It can equivalently be constructed by simply considering the chain complex generated by elements $\mathbf{x} \in s_z^{-1}(\mathfrak{s}) \subset \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, with differential given by counting only the disks which do not intersect $\{z\} \times \mathrm{Sym}^{g-1}(\Sigma_g)$:

$$\widehat{\partial}(\mathbf{x}) = \sum_{\mathbf{y} \in s_z^{-1}(\mathfrak{s})} \sum_{\phi \in \pi_2(\mathbf{x},\mathbf{y}) | \mu(\phi)=1, n_z(\phi)=0} n(\mathbf{x}, \mathbf{y}; \phi)\mathbf{y}.$$

With HF° denoting any of the above variants, we write:

$$HF^\circ(Y) := \bigoplus_{\mathfrak{s}\in \mathrm{Spin}^c(Y)} HF^\circ(Y, \mathfrak{s}).$$

Heegaard Floer homology is an umbrella term which can apply to any of HF^∞, HF^-, HF^+, or $\widehat{HF}$. Henceforth, if the particular variant is unspecified, then it is safe to assume that the discussion applies to all four variants.

4.2.3 Naturality of Heegaard Floer Homology

It is not at all obvious that the various Heegaard Floer homologies constructed above are topological invariants of the oriented, closed 3-manifold Y. In principle, two very important choices were made: (1) a Heegaard splitting for Y, consisting of the data $(\Sigma_g, \alpha, \beta, z)$, where $\alpha = (\alpha_1, \ldots, \alpha_g)$, $\beta = (\beta_1, \ldots, \beta_g)$, Σ_g has the orientation induced by identifying it as the level set $h^{-1}(\frac{3}{2})$, and $z \in \Sigma_g$ is the chosen basepoint disjoint from the α and β curves; and (2) a complex structure on Σ_g.

As it turns out, as shown by Ozsváth and Szabó [4], Heegaard Floer homology is well-defined up to isomorphism.

Theorem 8. *All of the Heegaard Floer homology variants defined above are well-defined up to isomorphism for a given oriented, closed 3-manifold Y and spin structure $\mathfrak{s}$ on Y.*

Thus, the notation $HF^\circ(Y, \mathfrak{s})$ (where "$\circ$" stands in for any of ∞, $+$, $-$, or a "hat") is justified and unambiguous, in the sense that it refers to a single isomorphism class of $\mathbb{Z}[U]$-modules.

However, there remains a question of naturality (or "higher-order" well-definedness): is there a canonical choice of isomorphism induced on Heegaard Floer homology for a given change in the Heegaard data (consisting of both the Heegaard splitting information and the chosen complex structure)? Juhász, Thurston, and Zemke [9] obtained an answer in the affirmative. In fact, they obtained a stronger result on the chain level, which we present below (after a clarification on notation).

If we choose a Heegaard splitting $G = (\Sigma_g, \alpha, \beta, z)$ and complex structure J on $\mathrm{Sym}^g(\Sigma_g)$, we can notationally contain all of this data in the form of a *Heegaard pair* $\mathcal{H} = (G, J)$. Then, we denote by $CF^\circ(\mathcal{H}, \mathfrak{s})$ the corresponding Heegaard Floer chain complex, and by $HF^\circ(\mathcal{H}, \mathfrak{s})$ the corresponding Heegaard Floer homology.

Proposition 1. *For a fixed $(Y, z, \mathfrak{s})$, and for any two Heegaard pairs $\mathcal{H}$, $\mathcal{H}'$, there is a distinguished chain homotopy equivalence (unique up to chain homotopy)*

$$\Phi(\mathcal{H}, \mathcal{H}') : CF^\circ(\mathcal{H}, \mathfrak{s}) \to CF^\circ(\mathcal{H}', \mathfrak{s})$$

and the maps given by Φ satisfy that, for any $\mathcal{H}, \mathcal{H}', \mathcal{H}''$,

(1) $\Phi(\mathcal{H}, \mathcal{H}) \sim id_{CF^\circ(\mathcal{H},\mathfrak{s})}$,

(2) $\Phi(\mathcal{H}', \mathcal{H}'') \circ \Phi(\mathcal{H}, \mathcal{H}') \sim \Phi(\mathcal{H}, \mathcal{H}'')$

where $\sim$ indicates chain homotopic maps.

It follows that the same result holds on the level of homology $HF^\circ(\mathcal{H}, \mathfrak{s})$, but we can replace the chain homotopy equivalences with isomorphisms and chain homotopy with equality. We say that the groups $HF^\circ(\mathcal{H}, \mathfrak{s})$ *form a transitive system.*

Though this discussion of naturality may seem pedantic at first, it is actually critical to the construction of involutive Heegaard Floer homology in the subsection which follows. In fact, there are "higher order" questions of naturality – e.g. is there a canonical choice of chain homotopy between any two choices of $\Phi(\mathcal{H}, \mathcal{H}')$ within its chain homotopy class? – which would assist in even more powerful constructions if answered in the affirmative. See the discussion in [1] for more detail.

4.3 The Involutive Variant

In this subsection, we will build a variant of Heegaard Floer homology by leveraging an additional piece of information: the involution induced on the Heegaard Floer complex by switching from the Morse function $h : Y \to \mathbb{R}$ which induces the Heegaard splitting to its negative $-h : Y \to \mathbb{R}$ (to make $-h$ self-indexing, we need only shift it by $+3$; we will assume this shift in the notation $-h$). We follow the original construction due to Hendricks and Manolescu [1]. Let the Heegaard splitting induced by h be given by $G = (\Sigma_g, \alpha, \beta, z)$. The Morse function $-h$ induces the *conjugate Heegaard splitting* $\overline{G} := (-\Sigma_g, \beta, \alpha, z)$, where $-\Sigma_g$ denotes Σ_g with orientation reversed. This flipping of Morse functions also induces conjugation on the almost complex structure J on $\mathrm{Sym}^g(\Sigma_g)$, so that the Heegaard pair $\mathcal{H} = (G, J)$ gets sent to $\overline{\mathcal{H}} := (\overline{G}, \overline{J})$.

It was shown in [10] that there is a canonical isomorphism between chain complexes

$$\eta : CF^\circ(\mathcal{H}, \mathfrak{s}) \to CF^\circ\left(\overline{\mathcal{H}}, \bar{\mathfrak{s}}\right).$$

From the previous subsection, we know that there is a distinguished chain homotopy equivalence

$$\Phi(\overline{\mathcal{H}}, \mathcal{H}) : CF^\circ\left(\overline{\mathcal{H}}, \bar{\mathfrak{s}}\right) \to CF^\circ\left(\mathcal{H}, \bar{\mathfrak{s}}\right)$$

We define their composition $\iota := \Phi(\overline{\mathcal{H}}, \mathcal{H}) \circ \eta$. It is shown in [1] that this is a chain homotopy involution, meaning that $\iota^2 \sim \mathrm{id}_{CF^\circ(\mathcal{H},\mathfrak{s})}$, where $\sim$ denotes chain homotopy. As such, $\mathcal{J} := \iota_* : HF^\circ(\mathcal{H}, \mathfrak{s}) \to HF^\circ(\mathcal{H}, \bar{\mathfrak{s}})$ is an involution on the level of Heegaard Floer homology.

We define the *involutive Heegaard Floer complex* (of type $\circ$, where $\circ$ stands for any of $+$, $-$, ∞, or "hat"), denoted $CFI^\circ(\mathcal{H}, \mathfrak{s})$, to be the mapping cone complex

$$CF^\circ(\mathcal{H}, \mathfrak{s}) \xrightarrow{Q(1+\iota)} Q \cdot CF^\circ(\mathcal{H}, \mathfrak{s})[-1]$$

where Q is a formal variable with degree -1, satisfying $Q^2 = 0$, and the bracketed $[-1]$ denotes a -1 grading shift. More explicitly, the complex consists of a module over $\mathbb{Z}_2[Q, U]/(Q^2)$ (note that we are using $\mathbb{Z}_2$-coefficients for simplicity) given by

$$CF^\circ(\mathcal{H}, \mathfrak{s})[-1] \otimes \mathbb{Z}_2[Q]/(Q^2)$$

with differential

$$\partial^\iota = \partial^\circ + Q(1 + \iota)$$

where ∂° is the standard Heegaard Floer differential of type $\circ$.

Hendricks and Manolescu show that the above chain complex is well-defined up to quasi-isomorphism for a fixed choice of 3-manifold and Spin^c structure, and thus the homology of the complex, called the *involutive Heegaard Floer homology* $HFI^\circ(\mathcal{H}, \mathfrak{s})$, is well-defined up to isomorphism.

The involutive Heegaard Floer complex turns out to be particularly well-behaved under connected sum:

Proposition 2. *There is a chain homotopy equivalence*

$$CF^-(Y_1 \# Y_2, \mathfrak{s}_1 \# \mathfrak{s}_2) \simeq CF^-(Y_1, \mathfrak{s}_1) \otimes_{\mathbb{Z}_2[U]} CF^-(Y_2, \mathfrak{s}_2)[-2]$$

such that the "conjugation" homotopy involution ι on $CF^-(Y_1 \# Y_2, \mathfrak{s}_1 \# \mathfrak{s}_2)$ is chain homotopy equivalent to $\iota_1 \otimes \iota_2$, where ι_1 and ι_2 are the conjugation homotopy involutions on $CF^-(Y_1, \mathfrak{s}_1)$ and $CF^-(Y_2, \mathfrak{s}_2)$, respectively.

4.3.1 ι-Complexes and Local Equivalence

The proposition at the end of the last section suggests a natural way to make the set of complexes C with a chain homotopy involution ι (satisfying some basic properties) into a group under tensor product.

We define an ι-complex to be a pair $\mathcal{C} = (C, \iota)$, where C is a finitely generated, free chain complex over $\mathbb{Z}_2[U]$ with $\mathbb{Z}$-grading such that there is a relatively graded isomorphism between $U^{-1}H_*(C)$ (the localization of $H_*(C)$ at U) and $\mathbb{Z}_2[U, U^{-1}]$, and $\iota : C \to C$ is grading preserving and satisfies $\iota^2 \simeq \mathrm{id}_C$, where $\simeq$ denotes chain homotopy. Note that $(CF^-(Y, \mathfrak{s}), \iota)$, where Y is a rational homology sphere and ι is the conjugation chain homotopy involution, is an ι-complex.

We define an equivalence relation, called *local equivalence*, on the set of ι-complexes as follows: we say two ι-complexes (C, ι), (C', ι') are equivalent if there exist chain maps $f : C \to C'$, $g : C' \to C$ which homotopy commute with the involutions – that is, $f \circ \iota \simeq \iota' \circ f$ and $\iota \circ g \simeq g \circ \iota'$ – and which induce isomorphisms on $U^{-1}H_*$.

The motivation for this notion of equivalence comes from the following fact:

Proposition 3. *Let $(Y_1, \mathfrak{s}_1)$ and $(Y_2, \mathfrak{s}_2)$ be spin rational homology cobordant. Then the ι-complexes $(CF^-(Y_1), \mathfrak{s}_1)$ and $(CF^-(Y_2), \mathfrak{s}_2)$ are locally equivalent.*

The set of ι-complexes modulo local equivalence forms an abelian group, denoted $\mathfrak{I}$, under tensor product:

$$\mathcal{C} \otimes \mathcal{C}' := \left(C \otimes_{\mathbb{Z}_2[U]} C'[-2], \iota \otimes \iota'\right).$$

The identity element is the class of $(\mathbb{Z}_2[U], \mathrm{id})$, with $\mathrm{gr}(1) = -2$, which is precisely the ι-complex associated to the standard 3-sphere S^3 (the grading is by convention). Inversion in $\mathfrak{I}$ is given by $[(C, \iota)] \mapsto [(C^*, \iota^*)]$, where $C^* := \mathrm{Hom}_{\mathbb{Z}_2[U]}(C, \mathbb{Z}_2[U])$ and ι^* is dual to ι. The grading convention for C^* is that the dual generator x^* to a generator x of C has grading $\mathrm{gr}(x^*) = -\mathrm{gr}(x) - 4$.

5 Application to the Homology Cobordism Group

In this section, we will put involutive Heegaard Floer Homology to use in uncovering a piece of the structure of the (3-dimensional) *homology cobordism* *group* $\Theta^3_{\mathbb{Z}}$ – the group of homology 3-spheres modulo smooth (or piecewise-linear; they are equivalent in this case) homology cobordism.

5.1 Introduction to the Homology Cobordism Group

A *$\mathbb{Z}$-homology 3-sphere*, or, briefly, a *homology sphere*, is a 3-manifold Y satisfying $H_*(Y) = H_*(S^3)$. Given two homology spheres Y_1 and Y_2, a *$\mathbb{Z}$-homology cobordism*, or just a homology cobordism, between Y_1 and Y_2 is a smooth (or, equivalently, PL), oriented 4-manifold with boundary W satisfying $\partial W = Y_1 \sqcup -Y_2$ and $H_*(W, Y_1; \mathbb{Z}) = H_*(W, Y_2; \mathbb{Z}) = 0$. The set of homology spheres modulo homology cobordism forms a group under connected sum. This group is denoted $\Theta^3_{\mathbb{Z}}$.

The 3-dimensional homology cobordism group is notable for the fact that the analogous group defined in any other dimension (using PL homology spheres and homology cobordisms) is the trivial group. Likewise, the "topological" 3-dimensional homology cobordism group, defined using not-necessarily-smooth homology cobordisms, is trivial (that is, every homology 3-sphere bound a topological homology ball). However, as we shall see, the group $\Theta^3_{\mathbb{Z}}$ has some highly nontrivial structure. In an imprecise sense, the 3-dimensional homology cobordism group is a measure of the degree to which the topological and smooth/PL categories diverge in dimension 4.

The first known structural fact about the homology cobordism group was the existence of a surjective homomorphism $\mu : \Theta^3_{\mathbb{Z}} \to \mathbb{Z}_2$, called the *Rokhlin homomorphism*. The well-known Poincaré homology sphere has Rokhlin invariant 1, and thus is a nontrivial element of $\Theta^3_{\mathbb{Z}}$.

The structure of the homology cobordism group is closely connected to the problem of triangulating high-dimensional (≥ 5) manifolds. It was shown by Galewski and Stern [11] that the triangulation conjecture for dimensions ≥ 5 is equivalent to the existence of a 2-torsion element $\Sigma \in \Theta^3_{\mathbb{Z}}$ with Rokhlin invariant 1. In [12], Manolescu used monopole Floer theory to disprove the triangulation conjecture by showing that such an element does not exist.

The homology cobordism group has also found applications in studying knot concordance and smooth 4-dimensional topology.

5.2 Homology Cobordism Invariants from Involutive Heegaard Floer Homology

By the construction of the group $\mathfrak{I}$ of ι-complexes modulo local equivalence, we have a well-defined homomorphism

$$\Theta^3_{\mathbb{Z}} \to \mathfrak{I}$$

given by sending $(Y, \mathfrak{s})$ to the class of its ι-complex $[(CF^-(Y, \mathfrak{s}), \iota)]$. Note that we are considering *spin* homology spheres (that is, homology spheres with Spin^c structures), modulo spin homology cobordism. Our objective in this subsection is to uncover some structural aspects of $\mathfrak{I}$. The above homomorphism will allow us to leverage our knowledge of the structure of $\mathfrak{I}$ to prove an important fact about the structure of $\Theta^3_{\mathbb{Z}}$.

For an ι-complex $\mathcal{C} = (C, \iota)$, we may consider its set of *self-local equivalences*, which are maps $f : C \to C$ which commute with ι and induce isomorphisms on $U^{-1}H_*$. This set can be assigned a pre-order according to containment of kernel – that is, we say $f \leq g$ if $\ker f \subseteq \ker g$. We then have a notion of *maximal self-local equivalence*, according to this pre-order.

It is shown by Hendricks, Hom, and Lidman [2] that a maximal self-local equivalence always exists for any ι-complex, and that, given a maximal self-local equivalence $f : C \to C$, $f\,|_{\mathrm{im}\, f} : \mathrm{im}\ f \to \mathrm{im}\ f$ is an isomorphism of chain complexes. We can then define $\iota_f := f \circ \iota \circ (f\,|_{\mathrm{im}\, f})^{-1}$, and then $(\mathrm{im}\ f, \iota_f)$ is an ι-complex which turns out to have the same local equivalence type as $\mathcal{C}$. It can be shown that two choices of maximal self-local equivalence give homotopy equivalent ι-complexes of the form $(\mathrm{im}\ f, \iota_f)$. We call this the *connected complex*, denoted $\mathcal{C}_{\mathrm{conn}} = (C_{\mathrm{conn}}, \iota_{\mathrm{conn}})$ associated to $\mathcal{C}$.

Proposition 4. *Locally equivalent ι-complexes have chain isomorphic connected complexes.*

Thus, the connected complex is an invariant of local equivalence. Now, given an ι-complex $\mathcal{C}$, we can define the *connected homology of* $\mathcal{C}$ to be

$$H_{\mathrm{conn}}(\mathcal{C}) := H_{\mathrm{red}}(\mathcal{C}_{\mathrm{conn}})$$

where *reduced homology* H_{red} is defined by

$$H_{\mathrm{red}}(C) := \ker \left(U^N : H_*(C) \to H_*(C)\right)[-1]$$

for large N.

We can now define the *connected Heegaard Floer homology* of a spin rational homology sphere $(Y, \mathfrak{s})$ by

$$HF_{\text{conn}}(Y, \mathfrak{s}) := H_{\text{conn}}(CF^-(Y, \mathfrak{s}), \iota).$$

Proposition 5. *Connected Heegaard Floer homology is a homology cobordism invariant.*

Proof. This follows immediately from the fact that the connected complex is a local equivalence invariant and the fact that spin homology cobordant homology spheres have locally equivalent ι-complexes. $\square$

Connected homology provides powerful insight into the structure of the group $\mathfrak{I}$. In particular, it provides a filtration, as follows: Let $\mathcal{P}$ denote the set of subsets of $\mathbb{N}$. Let $P \in \mathcal{P}$, and define:

$$\mathcal{F}_P^{\mathfrak{I}} := \{[\mathcal{C}] \in \mathfrak{I} \mid H_{\text{conn}}(\mathcal{C}) \cong \bigoplus_i \mathbb{Z}_2[U]/U^{n_i}\mathbb{Z}_2[U], n_i \in P\}.$$

Note that the above isomorphism is ungraded.

Theorem 9. *The subsets $\mathcal{F}_P^{\mathfrak{I}}$ provide a filtration of $\mathfrak{I}$ by $\mathcal{P}$; that is, each $\mathcal{F}_P^{\mathfrak{I}}$ is a subgroup of $\mathfrak{I}$, and if $P_1 \subseteq P_2$, then $\mathcal{F}_{P_1}^{\mathfrak{I}} \leq \mathcal{F}_{P_2}^{\mathfrak{I}}$.*

We conclude this section by defining an invariant associated to an ι-complex which measures the nilpotence of H_{conn}.

$$\omega(\mathcal{C}) := \min\{n \mid U^n H_{\text{conn}}(\mathcal{C}) = 0\}.$$

This quantity satisfies the useful property that

$$\omega(\mathcal{C} \otimes \mathcal{C}') \leq \max\{\mathcal{C}, \mathcal{C}'\}.$$

5.3 A Proof Outline of Furuta's Theorem

In this subsection, we will outline a proof of Furuta's Theorem, continuing to follow Hendricks, Hom, and Lidman [2].

Theorem 10. *The homology cobordism group $\Theta^3_{\mathbb{Z}}$ has a $\mathbb{Z}^\infty$ subgroup.*

Proof. From the filtration on $\mathfrak{I}$ stated in the previous subsection, along with the homomorphism $\Theta^3_{\mathbb{Z}} \to \mathfrak{I}$, $[Y] \mapsto [(CF^-(Y), \iota)]$, it follows that $\Theta^3_{\mathbb{Z}}$ has a filtration given by the subgroups:

$$\mathcal{F}_P = \{[Y] \mid HF_{\text{conn}}(Y) \cong \bigoplus_i \mathcal{T}_{a_i}(n_i) \mid n_i \in P\}$$

where $\mathcal{T}_{a_i}(n_i)$ is, in an ungraded sense, isomorphic to $\mathbb{Z}_2[U]/U^{n_i}\mathbb{Z}_2[U]$, but satisfies the grading requirement $\text{gr}(1) = a_i$.

Our next step is to show that there is an infinite sequence of *properly* nested subgroups of $\Theta^3_{\mathbb{Z}}$. Consider the collection of subgroups Θ^3_N generated by $\{S_{1/n}(K) \mid g(K) < N, n \in \mathbb{Z}\}$; that is, generated by homology 3-spheres obtained from surgery on a knot of genus bounded above by $N > 0$. It is shown in [13] that

$$U^{g(K)+g(K)/2+1} HF_{\text{red}}(S_{1/n}(K)) = 0$$

so that each Θ^3_N is contained in $\mathcal{F}_{\{1,\ldots,3N/2\}}$. It is shown in [2] that $HF_{\text{conn}}(S^3_{-1}(K), [0]) = \mathcal{T}_{-1}(V_0)$, where V_0 is the concordance invariant, and K is an L-space knot. But there exist L-space knots with any concordance invariant value V_0, so that $\mathcal{F}_{\{1,\ldots,m\}}/\mathcal{F}_{\{1,\ldots,m-1\}}$ is nontrivial for all m, and so in particular Θ^3_N is a proper subgroup of $\Theta^3_{\mathbb{Z}}$ for all N.

At this point, we have proven that $\Theta^3_{\mathbb{Z}}$ is infinitely generated. To finish the proof of the statement, we need to produce an infinite set of linearly independent elements. We claim that the collection $\{S^3_{-1}(T_{2,4n+1})\}_n$ is linearly independent. It is shown in [14] that the knots $T_{2,4n+1}$ have concordance invariant given by $V_0(T_{2,4n+1}) = n$. It is shown in [2] that $\omega(kS^3_{-1}(T_{2,4n+1})) = n$ for any k nonzero. Recall that the invariant ω defined in the previous subsection has the property that

$$\omega\left(k_1 S^3_{-1}(T_{2,5}))\#k_2 S^3_{-1}(T_{2,9}))\#\ldots\#k_n S^3_{-1}(T_{2,4n+1}))\right) \leq \max_{k_i \neq 0} i$$

(where we used the fact that connected sum maps to tensor product under the homomorphism $\Theta^3_{\mathbb{Z}} \to \mathfrak{I}$). If, for some n, a nontrivial sum were equal to 0, then we would obtain:

$$n = \omega\left(-k_n S^3_{-1}(T_{2,4n+1})\right) = \omega\left(k_1 S^3_{-1}(T_{2,5}))\#k_2 S^3_{-1}(T_{2,9}))\#\ldots\#k_{n-1} S^3_{-1}(T_{2,4(n-1)+1}))\right) \leq n-1$$

which is a contradiction. Therefore, this is a linearly independent set, and as such it generates a $\mathbb{Z}^\infty$ subgroup of $\Theta^3_{\mathbb{Z}}$. □

This proof, combined with the constructions of the previous subsections, demonstrates the general strategy which has been used (with great effectiveness) to uncover facts about the structure of the homology cobordism group: use Floer theory to construct a map from $\Theta^3_{\mathbb{Z}}$ to some other object, determine structural aspects of the other object, and then leverage the map to transfer some of that structure over to $\Theta^3_{\mathbb{Z}}$.

Though Furuta's Theorem was originally proven using a different Floer theory – namely, Instanton Floer homology – involutive Heegaard Floer homology provides another very useful, and in many ways more powerful, avenue of proof. Indeed, involutive Heegaard Floer homology, together with the notion of local equivalence, was used by Dai, ... [DAI] to show that $\Theta^3_{\mathbb{Z}}$ actually contains a $\mathbb{Z}^\infty$ summand, which is a meaningful improvement upon Furuta's result. As of the writing of this paper, Heegaard Floer theory is the only Floer theory which has been used to prove this result.

6 Next Steps in Heegaard Floer Theory

Heegaard Floer theory has proven to be both a powerful and computationally tractable tool for low-dimensional topologists, but its development phase is far from over. Currently, its main drawback in comparison with other Floer theories, such as Seiberg-Witten Floer homology, is that the Heegaard Floer chain complex is not known to possess the "higher-order naturality" which would be necessary for constructing, for instance, a "Pin(2)-equivariant" Heegaard Floer homology [1]. Establishing higher order naturality results, such as identifying a canonical class of chain homotopy between any two representatives of the chain homotopy equivalence $\Phi(\mathcal{H}, \mathcal{H}')$ for two different sets of Heegaard data $\mathcal{H}, \mathcal{H}'$, would be an important next step in the field.

7 Acknowledgements

I'd like to thank my advisor, Dr. Irving Dai, for all of his support on this project. His reading recommendations, intuitive explanations, and overall guidance have been invaluable to me. I'd also like to thank the director of SURIM, Dr. Lernik Asserian. Without her amazing work organizing the program and its events, this project would not have been possible.

8 References

[1] Kristen Hendricks and Ciprian Manolescu. "Involutive Heegaard Floer homology". In: *Duke Mathematical Journal* 166.7 (May 2017). DOI: 10.1215/00127094-3793141. URL: https://doi.org/10.1215%2F00127094-3793141.

[2] Kristen Hendricks, Jennifer Hom, and Tye Lidman. "Applications of Involutive Heegaard Floer Homology". In: *Journal of the Institute of Mathematics of Jussieu* 20.1 (Apr. 2019), pp. 187–224. DOI: 10.1017/s147474801900015x. URL: https://doi.org/10.1017%2Fs147474801900015x.

[3] Augustin Banyaga and David Hurtubise. *Lectures on Morse homology.* Springer, 2011.

[4] Peter Ozsvath and Zoltan Szabo. *Holomorphic disks and topological invariants for closed three-manifolds.* 2001. DOI: 10.48550/ARXIV.MATH/0101206. URL: https://arxiv.org/abs/math/0101206.

[5] Irving Dai et al. *An infinite-rank summand of the homology cobordism group.* 2018. DOI: 10.48550/ARXIV.1810.06145. URL: https://arxiv.org/abs/1810.06145.

[6] John W. Milnor, R. Wells, and M. Spivak. *Morse theory.* Princeton University Press, 1973.

[7] Dusa Mcduff. "Floer theory and low dimensional topology". In: *Bulletin of the American Mathematical Society* 43 (2005), pp. 25–42.

[8] Vladimir Turaev. "Torsion invariants of Spin^c-structures on 3-manifolds". In: *Mathematical Research Letters* 4 (1997), pp. 679–695.

[9] András Juhász, Dylan P. Thurston, and Ian Zemke. *Naturality and mapping class groups in Heegaard Floer homology.* 2012. DOI: 10.48550/ARXIV.1210.4996. URL: https://arxiv.org/abs/1210.4996.

[10] Peter Ozsvath and Zoltan Szabo. *Holomorphic disks and three-manifold invariants: properties and applications.* 2001. DOI: 10.48550/ARXIV.MATH/0105202. URL: https://arxiv.org/abs/math/0105202.

[11] David Galewski and Ronald Stern. "Classification of Simplicial Triangulations of Topological Manifolds". In: *Annals of Mathematics* 111 (1980), pp. 1–34.

[12] Ciprian Manolescu. *Pin(2)-equivariant Seiberg-Witten Floer homology and the Triangulation Conjecture*. 2013. DOI: 10.48550/ARXIV.1303.2354. URL: https://arxiv.org/abs/1303.2354.

[13] Fyodor Gainullin. "The mapping cone formula in Heegaard Floer homology and Dehn surgery on knots in S3". In: *Algebraic &amp Geometric Topology* 17.4 (Aug. 2017), pp. 1917–1951. DOI: 10.2140/agt.2017.17.1917. URL: https://doi.org/10.2140%2Fagt.2017.17.1917.

[14] Peter Ozsvá th and Zoltán Szabó. "Heegaard Floer homology and alternating knots". In: *Geometry Topology* 7.1 (Mar. 2003), pp. 225–254. DOI: 10.2140/gt.2003.7.225. URL: https://doi.org/10.2140%2Fgt.2003.7.225.