A Hitchhiker's Guide to Heegaard-Floer Homology

Nicholas James

January 16, 2026

Abstract

In this expository paper based on [2] [6] and [7], the basic ideas behind Heegaard-Floer homology are explored with minimal amount of context in order to understand the essential aspects of the theory.

Contents

1 Handlebodies, Heegaard Decompositions, and Heegaard Diagrams

During our Junior seminar, I was inspired by the ideas of arcs yielding geometric information on surfaces and manifolds in general. In the case of the mapping class group, this yields extremely powerful information of how the mapping class group is generated by finitely many Dehn twists, and many other computational results. More importantly, there are several results relating the mapping class group to both symplectic geometry and Teichmüller theory. For this final paper of the geometric topology Junior seminar, I wanted to write about something adjacent but different at the same time. In Heegaard-Floer homology theory, these ideas are used to create diagrams and structures that not only relate such curves to manifolds, but also allow for the creation of interesting homology groups that encode this loop data (and a vast array of other analytic information). To kick things off in this expository set of notes, we shall assume that Y is a smooth closed oriented 3-manifold. A fundamental notion of this set of ideas is the handlebody, which we have even seen throughout the Junior Seminar.

Definition 1.0.1. A *handlebody of genus* g, which we denote as U, is the result of taking a closed 3-ball and attaching g-many 1-handles (solid cylinders) to it. Equivalently, we can say that U is diffeomorphic to a regular neighborhood of a bouquet of g-many circles in $\mathbb{R}^3$.

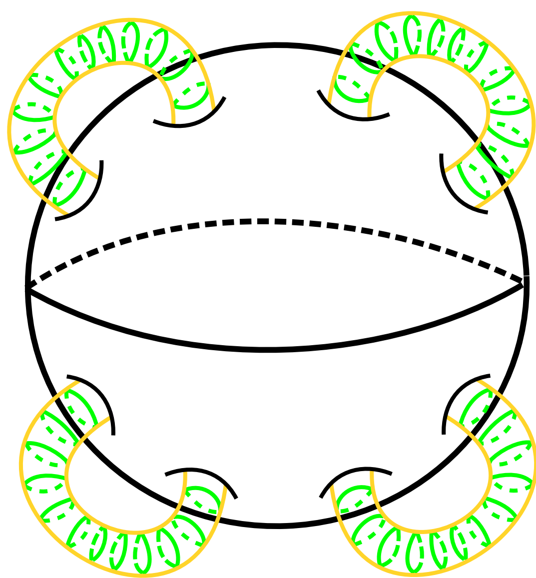

Figure 1.1. A genus 4-handlebody.

One of the central objects of Heegaard-Floer homology theory is based on a gluing procedure of such handlebodies. We can formalize this notion as follows.

Definition 1.0.2. Let U_0 and U_1 be homeomorphic g-handlebodies in $\mathbb{R}^3$. A *Heegaard decomposition* of Y is the splitting $Y = U_0 \cup_{\Sigma_g} U_1$, which is understood as the gluing of U_0 and U_1 along their shared boundary Σ_g by the map $\phi : \partial U_0 \to \partial U_1$. We say that Σ_g is a *Heegaard surface*, and it's clear that Σ_g is a genus-g surface.

Example. Consider $Y = \mathbb{S}^3$. By slicing it into two 3-balls, you get such a Heegaard decomposition.

Given that our course has been related to surfaces, there is an interesting reason why we are approaching things from a 3-manifold setting. It's essentially due to the following theorem below.

Theorem 1.0.3 (Singer)**.** *Let Y be an oriented closed three-dimensional manifold. Then Y admits a Heegaard decomposition.*

Proof. Triangulate Y, and create a graph out of its vertices and edges. Let U_0 be a sufficiently small neighborhood of this graph. By replacing each vertex with a ball, and each edge with a solid cylinder, this turns U_0 into a handlebody. It follows easily that $U_1 = Y \backslash U_0$ is also a handlebody, which is a regular neighborhood of a graph on the centers of the triangles and tetrahedra in the triangulation. This implies the desired. ■

Remark 1.0.4. It's important to note that U_0 and U_1 are clearly homeomorphic to each other. They have the same common boundary, and thus the same number of handles.

Note that we never said that this Heegaard decomposition was unique! That's because there are many different decompositions for a single manifold Y. We can absolutely formalize this a bit in the following definition.

Definition 1.0.5. Let $Y = U_0 \cup_{\Sigma_g} U_1$ be a genus-g Heegaard decomposition. Then we can define a genus-$(g+1)$ Heegaard decomposition by choosing two points in the Heegaard surface Σ_g and connecting them by a small (unknotted) arc γ in U_1. Let $\mathcal{N}$ denote a tubular neighborhood of γ such that $U_0' = U_0 \cup \mathcal{N}$ and $U_1' = U_1 \backslash \mathcal{N}$. Then the *stabilization of* $Y = U_0 \cup U_1$ is the genus-$(g+1)$ Heegaard decomposition

$$Y' = U_0' \cup_{\Sigma_{g+1}'} U_1' = (U_0 \cup \mathcal{N}) \cup_{\Sigma_{g+1}'} (U_1 \backslash \mathcal{N}).$$

In the opposite direction, we say that Y is the *destabilization* of Y'.

Remark 1.0.6. A (nontrivial) result from Singer states that any two Heegaard decompositions can be joined together by a stabilization or destabilization. In this way, stabilizations can be seen as a kind of three-manifold invariant.

Our next step is to take the idea of Heegaard decompositions and turn them into diagrams of some sort. We formalize this idea by defining a new kind of object.

Definition 1.0.7. Let U be a genus-g handlebody. A set of *attaching circles* $(\gamma_1, \ldots, \gamma_g)$ for U is a collection of closed embedded curves in $\Sigma_g = \partial U$ with the following properties:

1. The curves γ_i are disjoint from each other.
2. The remaining Heegaard surface $\Sigma_g \backslash (\gamma_1, \ldots, \gamma_g)$ is connected. In other words, $([\gamma_1], \ldots, [\gamma_g])$ are linearly independent in $H_1(\Sigma_g, \mathbb{Z})$.
3. The curves γ_i bound disjoint embedded disks in U.

Note that throughout this Junior seminar, we have spoken a similar language with the curves being manipulated by the mapping class group. A visual interpretation of an attaching circle is depicted below.

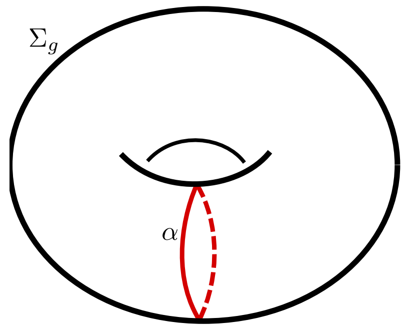

Figure 1.2. An attaching circle on a genus-1 handlebody, which is a neighborhood of a wedge of two circles in $\mathbb{R}^3$. Note that α bounds a disk, it's disjoint, and its complement is connected, so α is an attaching circle.

Heegaard decompositions can be described using special kinds of diagrams. The way we do it is by adding attaching circles to each handlebody in the decomposition.

Definition 1.0.8. Let (Σ_g, U_0, U_1) be a genus-g Heegaard decomposition for Y. A *Heegaard diagram* is a triple $\mathcal{H} = (\Sigma_g, \boldsymbol{\alpha}, \boldsymbol{\beta})$ given by Σ_g together with a collection of curves $\boldsymbol{\alpha} = \{\alpha_1, \ldots, \alpha_g\}$ and $\boldsymbol{\beta} = \{\beta_1, \ldots, \beta_g\}$ with the property that $\boldsymbol{\alpha}$ is a set of attaching circles for U_0 and $\boldsymbol{\beta}$ is a set of attaching circles for U_1.

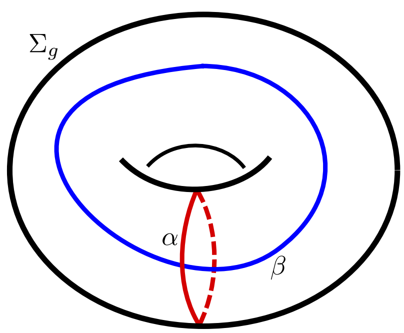

Figure 1.3. The Heegaard diagram of $\mathbb{S}^3$ on a genus 1 surface.

It is natural to inquire about the existence of such a diagram on Y. In order to show this, we make a few remarks on Morse theory although we assume familiarity. We say that a Morse function $f : Y \to \mathbb{R}$ is *self-indexing* if for all critical points p we have $f(p) = \lambda(p)$, where $\lambda(p)$ is the index of p (the maximal dimension of a subspace of T_pY on which the Hessian is negative definite, determined by the Morse Lemma).

Remark 1.0.9. The motivation for invoking Morse functions to begin with is due to the fact that in a smooth setting, they yield handle decompositions of the manifold through the sublevel and level sets that they induce, which are cobordantly related because of the critical points of the Morse function. This is rigorously treated using gradient vector fields and flows.

Self-indexing Morse functions are nice, because they allow us to arrange the handles such that they are added in increasing order of index. But the following theorem, thankfully, shows that we can always do this.

Theorem 1.0.10. *Let M be a smooth, closed, connected n-dimensional manifold. Then Y admits a self-indexing Morse function with a unique local minimum and unique local maximum.*

Now that it's clear why such a Morse theoretic approach should work in the first place, we can prove a satisfying result that we will assume throughout.

Theorem 1.0.11. *Let Y be any smooth, closed, connected, oriented three-manifold. There exists a Heegaard diagram that represents Y.*

Proof. Let f be a self-indexing Morse function with a unique maximum and a unique minimum on Y. Then f induces a handle decomposition on Y, so let Y_1 be the union of the 0-handle and all of the 1-handles in this decomposition. The Heegaard surface Σ_g is the oriented boundary of Y_1. Let the belt circles of the 1-handles be the $\boldsymbol{\alpha}$-circles (i.e. circles in $\partial Y_1 = \Sigma_g$), and the attaching circles of the 2-handles be the $\boldsymbol{\beta}$-circles (which are also circles in ∂Y_1). If f has g index-1 critical points, then Y_1 is a genus g handlebody, with Σ_g obviously being a surface of genus g (hence the notation), and we get g such $\boldsymbol{\alpha}$-curves in this way. Now consider $3 - f$ instead of f itself. This is once again a self-indexing Morse function on Y, with the same critical points. Indeed, an index k critical point of f is shifted to be an index $(3 - k)$ critical point of $3 - f$, and the attaching circles of the index 2 critical points of f are therefore the belt circles of the index 1 critical points of $3 - f$. The union of the 0-handle and the 1-handles of the handle decomposition induced by $3 - f$ give the complement of Y_1 in Y, and since the genus of the common boundary determines the number of 1-handles in the handlebody, we get that the diagram (defined previously) has g-many $\boldsymbol{\alpha}$ and g-many $\boldsymbol{\beta}$-circles. This new Heegaard diagram presents Y and we win. ■

While a single three-manifold can have many possible Heegaard diagrams, there are certain surgery-theoretic moves on a Heegaard diagram that don't change the underlying manifold itself. That is, one can manipulate the diagram such that it still depicts the same Heegaard decomposition.

Definition 1.0.12 (Isotopy)**.** Let U be a genus-g handlebody. An *isotopy* of U "homotopes" the set of curves $(\gamma_1, \ldots, \gamma_g)$ on G such that they are still disjoint.

This doesn't even need a visualization, since this is literally the same thing as an isotopy in other forms of geometry. Nothing new here.

Definition 1.0.13 (Handlesliding)**.** Let U be a genus-g handlebody. A *handleslide* of U of γ_1 and γ_2 produces a new set of attaching circles $(\gamma_1', \gamma_2, \ldots, \gamma_g)$ where γ_1' is a simple closed curve such that

1. γ_1' is disjoint from $(\gamma_1, \ldots, \gamma_g)$
2. The triple $(\gamma_1', \gamma_1, \gamma_2)$ cobound an embedded pair of pants in $\Sigma \backslash (\gamma_3, \ldots, \gamma_g)$.

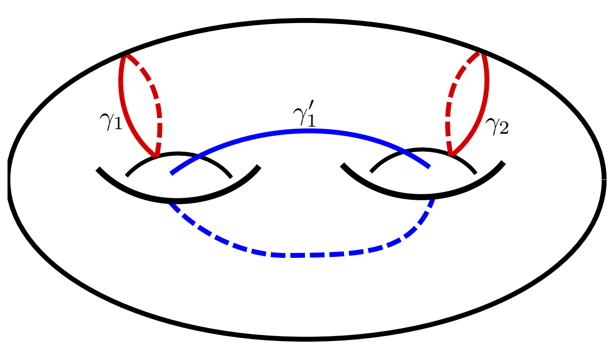

Figure 1.4. A handleslide.

Similarly to Heegaard decompositions, we can define stabilizations of Heegaard diagrams in a similar way. The main idea is to take a connected sum of the original Heegaard surface Σ_g, and raise by a genus by connecting a single torus $\mathbb{T}^2$ to it. This allows us to define two new attaching circles to the handlebodies and raise the Heegaard diagram by 1 genus in total.

Definition 1.0.14 (Stabilization)**.** Let $\mathcal{H} = (\Sigma_g, \boldsymbol{\alpha}, \boldsymbol{\beta})$ be a Heegaard diagram for Y. A *stabilization* of $\mathcal{H}$ is an operation that results in a new Heegaard diagram $\mathcal{H}' = (\Sigma_{g+1}, \boldsymbol{\alpha}', \boldsymbol{\beta}')$ such that the following are true:

1. $\Sigma_{g+1} = \Sigma_g \# \mathbb{T}^2 = \Sigma_g \# (\mathbb{S}^1 \times \mathbb{S}^1)$ is the new Heegaard surface of $\mathcal{H}'$.
2. The set of attaching circles $\boldsymbol{\alpha}' = \boldsymbol{\alpha} \cup \{\alpha_{g+1})$ and $\boldsymbol{\beta}' = \boldsymbol{\beta} \cup \{\beta_{g+1})$ such that α_{g+1} and β_{g+1} are simple closed curves in $\mathbb{T}^2$ that intersect transversally at a single point.

Conversely, the diagram $\mathcal{H}$ is referred to as the *destabilization* of $\mathcal{H}'$ depending on the context.

Handleslides, and isotopies, and stabilizations can be seen as another form of Reidemeister moves, but instead of working with knots we're working with attaching circles. Similarly to these knot theoretic analogues, we have the following theorem that allows us to relate one Heegaard diagram to another.

Theorem 1.0.15. *Let $\mathcal{H} = (\Sigma_g, \boldsymbol{\alpha}, \boldsymbol{\beta})$ and $\mathcal{H}' = (\Sigma_{g'}, \boldsymbol{\alpha}', \boldsymbol{\beta}')$ be Heegaard diagrams for Y. Then $\mathcal{H}$ and $\mathcal{H}'$ can be transformed into each other diffeomorphically by a sequence of isotopies, handleslides, and stabilizations.*

Remark 1.0.16*.* An analogous result holds for handlebodies themselves. Thinking in terms of a handlebody U, we can see why each of these three operations are necessary and sufficient to make to handlebodies U_1 and U_2 diffeomorphic. You can either (1) contort the existing set of attaching circles (2) get a new set of attaching circles and try that or (3) lower or raise the genus.

Sometimes we want to work with Heegaard diagrams that are built off of a basepoint. We define them now just for the sake of completeness.

Definition 1.0.17. A Heegaard diagram $(\Sigma_g, \boldsymbol{\alpha}, \boldsymbol{\beta})$ is said to be *pointed* if it can be written as the 4-tuple $(\Sigma_g, \boldsymbol{\alpha}, \boldsymbol{\beta}, z)$ where $z \in \Sigma_g \backslash (\boldsymbol{\alpha} \cup \boldsymbol{\beta})$ is a basepoint.

Pointed diagrams have equivalent results to the above, but you adjust for this new choice of basepoint. They are more important than just an extension, though, and this is emphasized in the next definition as follows.

Definition 1.0.18. A *doubly pointed* Heegaard diagram for a knot K in $\mathbb{S}^3$ is a 5-tuple $(\Sigma_g, \boldsymbol{\alpha}, \boldsymbol{\beta}, w, z)$ where w and z are basepoints in the complement $\Sigma_g \backslash (\boldsymbol{\alpha} \cup \boldsymbol{\beta})$ such that

- $(\Sigma_g, \boldsymbol{\alpha}, \boldsymbol{\beta})$ is a Heegaard diagram for $\mathbb{S}^3$
- K is the union of arcs $\mathcal{A}_1$ and $\mathcal{A}_2$ where $\mathcal{A}_1 \in \Sigma_g \backslash \boldsymbol{\alpha}$ connects w to z, and where $\mathcal{A}_2 \in \Sigma_g \backslash \boldsymbol{\beta}$ connects z to w.

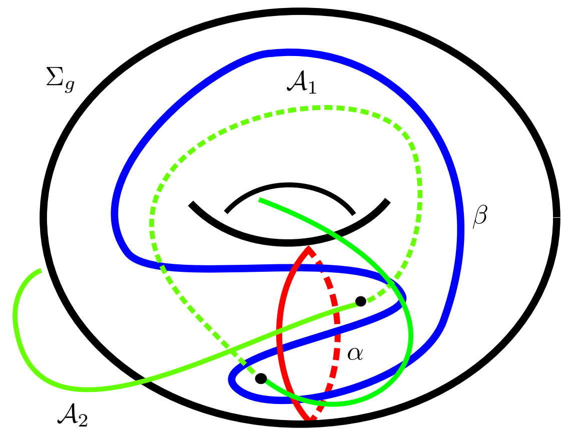

Figure 1.5. A pointed doubly pointed Heegaard Diagram for the trefoil knot.

There is another way to construct a doubly pointed Heegaard diagram for a knot K. Let D be a knot diagram for K, and suppose that D has c-many crossings. Forgetting the crossing data of the diagram D yields an immersed curve γ in the plane. The complement of γ is $c+2$ regions in the plane, one of which is unbounded. Let Σ be the boundary of a regular neighborhood of γ in $\mathbb{R}^3$; note that Σ is a surface of genus $c+1$. For each of the bounded regions in the complement of γ, we put a longitude on Σ. For each crossing of D, we put an meridian on Σ. Lastly, we add one final meridian and put two basepoints w and z on either side of it. This yields the desired doubly pointed Heegaard diagram. Like pointed Heegaard diagrams, doubly pointed Heegaard diagrams have analogous results to the above.

2 Clifford algebras, Spin structures, and Spin$^{\mathbb{C}}$ structures

An important tool used in Heegaard-Floer homology are Spin$^{\mathbb{C}}$ structures. In order to define them, we have to begin with structures called Clifford algebras, which are essentially generalizations of complex numbers to arbitrarily high dimensions.

2.1 Clifford Algebras and Spin Groups

Definition 2.1.1. Let V be a finite-dimensional vector space over a field $\mathbb{F}$ of characteristic zero endowed with a quadratic form q. We call the pair (V, q) a *quadratic space*. Let

$$\mathcal{T}(V) := \bigoplus_{k=0}^{\infty} \left(\bigotimes^{k} V \right)$$

be a tensor algebra over V, and let $\mathcal{I}_q(V)$ be a two-sided ideal in $\mathcal{T}(V)$ generated by elements of the form $(v \otimes v - q(v)1)$ where $v \in V$. Then the *Clifford algebra* of $\mathrm{Cl}(V, q)$ of the quadratic space (V, q) is defined as the quotient

$$\mathrm{Cl}(V, q) := \mathcal{T}(V)/\mathcal{I}_q(V)$$

The Clifford algebra $\mathrm{Cl}(V, q)$ is generated by the vector space V (with unity) subject to the relation that $v \cdot v = -q(v)1$ for $v \in V$. If the characteristic of $\mathbb{F}$ is not equivalent to 2, then for $v, w \in V$ we have

$$-2q(v, w) = v \cdot w + w \cdot v.$$

Using this identity we can prove the following.

Theorem 2.1.2. *Let $f : V \to \mathcal{A}$ be a linear map into an associative $\mathbb{F}$-algebra with unity such that the relation*

$$f(v) \cdot f(v) = -q(v)1$$

holds for all $v \in V$. Then f extends uniquely to an $\mathbb{F}$-algebra homomorphism

$$\tilde{f} : Cl(V, q) \to \mathcal{A},$$

and $Cl(V, q)$ is the unique associative $\mathbb{F}$-algebra with this property.

Proof. Any linear map $f : V \to \mathcal{A}$ extends to a unique algebra homomorphism $\tilde{f} : \mathcal{T}(V) \to \mathcal{A}$. Assuming that $f(v) \cdot f(v) = -q(v)1$, we have that $\tilde{f} \equiv 0$ on $\mathcal{I}_q(V)$ and so $\tilde{f}$ is induced to $\mathrm{Cl}(V, q)$ by definition. Now suppose that $\mathcal{B}$ is another associative $\mathbb{F}$-algebra with unity such that $i : V \to \mathcal{B}$ is an embedding with the property that any linear map $f : V \to \mathcal{A}$ satisfying

$$f(v) \cdot f(v) = -q(v)1$$

extends to an algebra homomorphism $\tilde{f} : \mathcal{B} \to \mathcal{A}$. Then the isomorphism $V \to i(V)$ induces an isomorphism $\mathrm{Cl}(V, q) \to \mathcal{B}$, done. ■

The above theorem allows us to take homomorphisms from Clifford bundles in the following way. Given a morphism $f : (V, q) \to (V', q')$ (i.e. $f : V \to V$ such that $f * q' = q$), we see that this induces an isomorphism $\tilde{f} : \mathrm{Cl}(V, q) \to \mathrm{Cl}(V', q')$ by Theorem 2.1. Our next definition will be of fundamental importance.

Definition 2.1.3. Let $\alpha : \mathrm{Cl}(V, q) \to \mathrm{Cl}(V, q)$ be an automorphism which extends the map $a(v) = -v$ on $v \in V$. Since $\alpha^2 = \mathrm{id}_V$, the Clifford algebra splits into the decomposition

$$\mathrm{Cl}(V, q) = \mathrm{Cl}^0(V, q) \oplus \mathrm{Cl}^1(V, q)$$

where $\mathrm{Cl}^i(V, q) = \{\varphi \in \mathrm{Cl}(V, q) \mid \alpha(\varphi) = (-1)^i \varphi\}$ are the eigenspaces of α. Note that $\alpha(\varphi_1 \varphi_2) = \alpha(\varphi_1)\alpha(\varphi_2)$, so $\mathrm{Cl}(V, q)^i \cdot \mathrm{Cl}(V, q)^j \subseteq \mathrm{Cl}(V, q)^{i+j}$, which makes $\mathrm{Cl}(V, q)$ into a $\mathbb{Z}/2\mathbb{Z}$*-graded algebra*. From this terminology, we call $\mathrm{Cl}^0(V, q)$ and $\mathrm{Cl}^1(V, q)$ the *even* and *odd* parts of $\mathrm{Cl}(V, q)$ because of this parity modulo 2.

There are many other nice properties that Clifford algebras satisfy. We list them now without going into detail.

- If V is a direct sum of two subspaces, then the Clifford algebra breaks apart into a nice $\mathbb{Z}/2\mathbb{Z}$-algebra tensor product.
- Using filtrations (see Lawson and Michelson) one can see that $\mathrm{Cl}(V, q)$ forms a graded algebra that is isomorphic to the exterior algebra of the vector space V.

Our journey continues with the study of the *multiplicative group of units* in the Clifford algebra, which is defined as the set

$$\mathrm{Cl}^\times(V, q) = \{\varphi \in \mathrm{Cl}(V, q) \mid \exists \varphi^{-1} \text{ with } \varphi\varphi^{-1} = \varphi^{-1}\varphi = 1\}.$$

Why do we care about this? The group of units always acts naturally as automorphisms of the algebra, i.e. there is an automorphism

$$\mathrm{Ad} : \mathrm{Cl}^\times(V, q) \to \mathrm{Aut}(\mathrm{Cl}(V, q))\ ,\ \mathrm{Ad}_\varphi(x) = \varphi x \varphi^{-1}$$

called the *adjoint representation* of $\mathrm{Cl}(V, q)$. Our big question that needs to be answered is the following: for which $\varphi \in \mathrm{Cl}^\times(V, q)$ does $\mathrm{Ad}_\varphi(V) = V$? The following result is a push in the right direction.

Theorem 2.1.4. *Let $v \in V \subset Cl(V, q)$ be an element with $q(v) \neq 0$. Then $\mathrm{Ad}_v(V) = V$. In fact, for all $w \in V$, the following equation holds:*

$$-\mathrm{Ad}_v(w) = w - 2\frac{q(v, w)}{q(v)} v.$$

Proof. Note that $v^{-1} = -v/q(v)$, and using an identity from before we have

$$\begin{aligned} -q(v)\mathrm{Ad}_v(w) &= -q(v) v w v^{-1} \\ &= -q(v)\left(-2q(v, w) - wv\right) v^{-1} \\ &= (q(v)2q(v, w) + q(v)wv)v^{-1} \\ &= \frac{(q(v)2q(v, w) + q(v)wv)(-v)}{q(v)} \\ &= (2q(v, w) + wv)(-v) \\ &= -2q(v, w) - wv^2 \\ &= -2q(v, w) + wq(v) \end{aligned}$$

and the result follows by dividing both sides by $q(v) \neq 0$. ■

Remark 2.1.5. The takeaway from this theorem is that the adjoint representation preserves the quadratic form q.

The above result motivates the following definition, which is intended to answer the question from before. It will serve as the structural basis to derive everything else.

Definition 2.1.6. The subgroup $P(V, q) \subseteq \mathrm{Cl}^\times(V, q)$ generated by the vectors $v \in V$ such that $q(v) \neq 0$. While not standard, we refer to $P(V, q)$ as the *preservation subgroup* of $\mathrm{Cl}^\times(V, q)$.

The preservation subgroup additionally has really important subgroups, which we present in their own definitions below.

Definition 2.1.7. The *pinor* group of (V, q), denoted by $\mathrm{Pin}(V, q)$, is the subgroup of $P(V, q)$ generated by all $v \in V$ such that $q(v) = \pm 1$. That is,

$$\mathrm{Pin}(V, q) = \{v_1, \dots, v_r \in P(V, q) \mid q(v_{1 \leq i \leq r}) = \pm 1\}.$$

Moreover, the *Spin* (or *Spinor*) group of (V, q), denoted as $\mathrm{Spin}(V, q)$, is the subgroup of $P(V, q)$ defined as the intersection of the pinor group and the even part of $\mathrm{Cl}(V, q)$. In other words,

$$\begin{aligned} \mathrm{Spin}(V, q) &= \mathrm{Pin}(V, q) \cap \mathrm{Cl}^0(V, q) \\ &= \{v_1, \dots, v_r \in \mathrm{Pin}(V, q) \mid r \equiv 0 \pmod 2\}. \end{aligned}$$

Thus, the spinor group is a subgroup of the preservation subgroup generated by an even number of vectors that have unitary length with respect to the quadratic form.

Remark 2.1.8. In Quantum Field Theory, the spin group is used to describe the symmetries of fermions, which are particles with spin 1/2. When we later complexify spin groups, we can describe the symmetries of electrically charged fermions. The most notable example would be electrons.

Let V be an n-dimensional $\mathbb{R}$-vector space, and suppose that q is a nondegenerate quadratic form on V. Then we can choose a basis on $V \cong \mathbb{R}^n$ such that

$$q(x) = \sum_{i=1}^{r} x_i^2 + \sum_{i=r+1}^{s} -x_i^2$$

where $s + r = n$. We adopt the convention of writing $\mathrm{Spin}_{r,s} = \mathrm{Spin}(V, q)$ and $\mathrm{Pin}_{r,s} = \mathrm{Pin}(V, q)$. Additionally, if the quadratic form is written as

$$q(x) = \sum_{i=1}^{n} x_i^2,$$

i.e. when $s = 0$ and $r = n$, we write $\mathrm{Spin}_n = \mathrm{Spin}_{n,0}$ and $\mathrm{Pin}_n = \mathrm{Pin}_{n,0}$. Thus, whenever we write subscripts we are assuming that V is an n-dimensional $\mathbb{R}$-vector space with a specific quadratic form.

Remark 2.1.9. In the same scenario, it is also common to write $\mathrm{Cl}(V, q) \equiv \mathrm{Cl}_{n,0}$. This notational convention will be used in later sections.

2.2 Spin structures

We first begin with orthonormal frames of vector bundles. While this seems random, this all somehow magically falls into place.

Definition 2.2.1. Let $P_{O_n}(E)$ denote the principal O_n-bundle whose fiber F at each point $x \in X$ is the set of orthonormal bases for $E_x = \pi^{-1}(x)$. We call $P_{O_n}(E)$ the *bundle of orthonormal frames* in E. Then we define the *bundle of orientations* of the total space E as $\mathrm{Or}(E) = P_{O_n}(E)/\mathrm{SO}_n$, where two bases of E_x are identified if the orthogonal matrix transforming one to the other has determinant 1.

When working with bundles like these, it's helpful to have an equivalent version to work with to bring these ideas into the realm of topology. We present an important definition (which is actually a theorem) now that will help with this.

Definition 2.2.2. Let $\pi : E \to X$ be a covering map and X be a connected topological space. Then for any $x_1, x_2 \in X$, the preimages $\pi^{-1}(x_1)$ and $\pi^{-1}(x_2)$ have the same cardinality, which we call the *degree* of the covering map p which we denote as $\deg(p)$. If $\deg(p) = 2$, we call this a *double cover* (or a *2-sheeted cover*).

Example. The bundle of orientations $\mathrm{Or}(E)$ is a 2-sheeted covering space of X, and E is orientable if $\mathrm{Or}(E)$ is trivial.

Using the above definition, we start off with a very important theorem.

Theorem 2.2.3. *Let $\mathrm{Cov}_2(X)$ denote the equivalence class of 2-sheeted covering spaces of a connected space X. Then there is a natural isomorphism $\mathrm{Cov}_2(X) \cong H^1(X, \mathbb{Z}/2\mathbb{Z})$.*

Proof. Since X is connected, we can decompose the above isomorphism to get

$$H^1(X; \mathbb{Z}/2\mathbb{Z}) \xrightarrow{\cong_1} \mathrm{Hom}(H_1(X), \mathbb{Z}/2\mathbb{Z}) \xrightarrow{\cong_2} \mathrm{Hom}(\pi_1(X); \mathbb{Z}/2\mathbb{Z}) \xrightarrow{\cong_3} \mathrm{Cov}_2(X).$$

The second isomorphism $\cong_2$ is because $H_1(X) \cong \mathrm{Ab}(\pi_1(X))$, and the third isomorphism comes from the fact that the 2-sheeted coverings of X correspond to the index-2 subgroups of $\pi_1(X)$. ∎

Because of this theorem, we see that the 2-sheeted covering space $\mathrm{Or}(E)$ of X determines a cohomology class $w_1(E) \in H^1(X; \mathbb{Z}/2\mathbb{Z})$ called the *first Stiefel-Whitney class* of E. More generally, we define these cohomology classes as follows.

Definition 2.2.4. Let ξ be a vector bundle. The *Stiefel-Whitney classes* of ξ are a sequence of cohomology classes

$$w_i(\xi) \in \mathrm{H}^i(B(\xi); \mathbb{Z}/2\mathbb{Z})$$

where the following axioms are satisfied:

1. The class $w_0(\xi)$ is equal to the unit element $1 \in \mathrm{H}^0(B(\xi); \mathbb{Z}/2\mathbb{Z})$.
2. If ξ is an $\mathbb{R}^n$-bundle and $i > n$, then $w_i(\xi) = 0$.
3. If $f : B(\xi) \to B(\eta)$ is covered by a bundle map from ξ to η, then $w_i(\xi) = f^* w_i(\eta)$.
4. If ξ and η are vector bundles over B, then

$$w_k(\xi \oplus \eta) = \sum_{i=0}^{k} w_i(\xi) \smile w_{k-i}(\eta) = \sum_{i+j=k} w_i(\xi) w_j(\eta).$$

5. If γ_1^1 is the line bundle over the circle $\mathbb{P}^1$, the first Stiefel-Whitney class $w_1(\gamma_1^1) \neq 0$.

Remark 2.2.5. Here we are using the convention that $\xi = E$, which is a slight abuse of notation.

Using the above discussion about the first Stiefel-Whitney class, this immediately implies a result worth mentioning.

Theorem 2.2.6. *A vector bundle $\xi = (E, \pi, X)$ over X is orientable if and only if $w_1(\xi) = 0$. Furthermore if $w_1(\xi) = 0$, then the distinct orientations on the total space E are in one-to-one correspondence with elements of $H^0(X, \mathbb{Z}/2\mathbb{Z})$.*

Proof. If $\xi = E$ is orientable, then $\mathrm{Or}(E)$ is trivial (or equivalently that $P_{O_n}(E)$ is disconnected) and thus $w_1(\xi) = 0$. ■

Remark 2.2.7. The latter half of the theorem's statement says that there are two possible orientations of E over each connected component of X.

Let $\xi = (E, \pi, X)$ be an oriented n-dimensional vector bundle over a manifold X equipped with a *Riemannian structure*, which is a positive definite inner product defined on each fiber $\pi^{-1}(x)$ for $x \in X$, and let $P_{SO_n}(E)$ be its bundle of oriented orthonormal frames. Note that there is a universal covering homomorphism $\psi_0 : \mathrm{Spin}_n \to \mathrm{SO}_n$ for $n \geq 3$ (proof omitted), which motivates the following definition.

Definition 2.2.8. Let $n \geq 3$. Then a *spin structure* on ξ is a principal Spin_n-bundle $P_{\mathrm{Spin}_n}(E)$ together with a 2-sheeted covering

$$\psi : P_{\mathrm{Spin}_n}(E) \to P_{\mathrm{SO}_n}(E)$$

such that $\psi(pg) = \psi(p)\psi_0(g)$ for all $p \in P_{\mathrm{Spin}_n}(E)$ and $g \in \mathrm{Spin}_n$.

Definition 2.2.9. For $n = 2$, the definition is the same except that we write SO_2 instead of Spin_n such that $\psi_0 : \mathrm{SO}_2 \to \mathrm{SO}_2$. When $n = 1$, we have $P_{\mathrm{SO}}(E) \cong X$ so a spin structure is just a double covering on X.

Using the above definition, we can reinterpret Theorem 2.2 in the following language.

Theorem 2.2.10. *Suppose X is connected. Then the spin structures on ξ are in natural one-to-one correspondence with elements of $H^1(P_{\mathrm{SO}_n}(E); \mathbb{Z}/2\mathbb{Z})$ whose restriction to the fibre of $P_{\mathrm{SO}_n}(E)$ is nonzero.*

Using similar ideas for our interpretation of the first Stiefel-Whitney class, we have our next main result.

Theorem 2.2.11. *Let ξ be an oriented vector bundle over a manifold X . Then there exists a spin structure on ξ if and only if the second Stiefel-Whitney class of ξ is zero. Furthermore, if $w_2(\xi) = 0$, then the distinct spin structures on ξ are in one-to-one correspondence with the elements of $H^1(X; \mathbb{Z}/2\mathbb{Z})$.*

This completely determines when spin structures exist on a vector bundle with certain conditions. What is even more interesting is how this generalizes to manifolds and invariants on them in general.

Definition 2.2.12. A *spin manifold* X is an oriented Riemannian manifold with a spin structure on its tangent bundle. The Stiefel-Whitney classes of X are defined to be the Stiefel-Whitney classes of its tangent bundle TX.

Remark 2.2.13. While we won't look at which manifolds are spin, the only important thing to note is that spin manifolds are determined using characteristic classes (Stiefel-Whitney and Chern classes in particular).

2.3 $\text{Spin}^{\mathbb{C}}$ structures

We have now arrived at the end of our detour while studying Heegaard-Floer homology. Before we actually define what $\text{Spin}^{\mathbb{C}}$ structures are and how they are helpful, we need some preliminary definitions to make sense of everything.

Definition 2.3.1. Let V be an n-dimensional $\mathbb{F}$-vector space and let q be a quadratic form on V. If $K \supseteq \mathbb{F}$ is a field containing $\mathbb{F}$, then a *K-representation* of the Clifford algebra $\text{Cl}(V, q)$ is a K-algebra homomorphism

$$\rho : \text{Cl}(V, q) \to \text{Hom}_K(W, W)$$

into the algebra of linear transformations of a finite dimensional vector space W over K. In this case, we call W the *$Cl(V, q)$-module* over K, and the operation $\rho(\varphi)(w) \equiv \varphi \cdot w$ is called *Clifford multiplication*. If W can be written as a nontrivial direct sum over K

$$W = W_1 \oplus_K W_2$$

such that $\varphi \cdot W_j \subseteq W_j$ for $j = 1, 2$, then ρ is called *reducible*. If it is not reducible, we say that ρ is *irreducible*.

A real or complex representation is a K-representation with $K = \mathbb{R}$ or $K = \mathbb{C}$ respectively.

Definition 2.3.2. Consider the spin group $\text{Spin}_n \subset \text{Cl}^0(V, q) \subset \text{Cl}(V, q)$. A *real spinor representation* of Spin_n is the real representation

$$\Delta_n : \text{Spin}_n \to \text{GL}(S)$$

which can be viewed as a real representation $\rho|_{\text{Spin}_n} : \text{Cl}_n \to \text{Hom}_{\mathbb{R}}(S, S)$. Likewise, a *complex spinor representation* of Spin_n is a complex representation

$$\Delta_n^{\mathbb{C}} : \text{Spin}_n \to \text{GL}(S)$$

which can be viewed as a real representation $\rho|_{\text{Spin}_n} : \text{Cl}_n \to \text{Hom}_{\mathbb{C}}(S, S)$.

Remark 2.3.3. Note that S is the Cl_n-module over $\mathbb{R}$, so essentially an analogue of W but for the real case. It carries no additional meaning.

In simple terms, $\text{Spin}^{\mathbb{C}}$ structures are complex analogues of spin structures. Their motivation, interestingly enough, comes from complex spinor representations.

Definition 2.3.4. Suppose that the map

$$\Delta_n^{\mathbb{C}} : \text{Spin}_n \to U(n)$$

is a complex spinor representation, and let $\zeta : U(1) \to U(n)$ denote the scalar multiples of the identity. This yields a homomorphism

$$\Delta_n^{\mathbb{C}} \times \zeta : \text{Spin}_n \times U(1) \to U(n)$$

where $(-1, 1) \in \ker(\Delta_n^{\mathbb{C}} \times \zeta)$. Dividing by this element gives the quotient group

$$\text{Spin}_n^{\mathbb{C}} \equiv \text{Spin}_n \times_{\mathbb{Z}/2\mathbb{Z}} U(1) \cong (\text{Spin}_n \times U(1))/(\mathbb{Z}/2\mathbb{Z}) \cong (\text{Spin}_n \times \text{Spin}_2))/(\mathbb{Z}/2\mathbb{Z})$$

which we call the $\text{Spin}_n^{\mathbb{C}}$ *group*.

A really satisfying result is that this group satisfies the exact sequence

$$1 \to \mathbb{Z}/2\mathbb{Z} \to \mathrm{Spin}_n^{\mathbb{C}} \to \mathrm{SO}_n \times U(1) \to 1,$$

which is another way to define the $\mathrm{Spin}_n^{\mathbb{C}}$ group. Another way of understanding the $\mathrm{Spin}_n^{\mathbb{C}}$ group is to view it as tensoring a spin group with the circle group. That is, we have

$$\mathrm{Spin}_n^{\mathbb{C}} \cong \mathrm{Spin}_n \otimes \mathbb{T}$$

where $\mathbb{T} = \{z \in \mathbb{C} \mid |z| = 1\}$. The reason why we have this isomorphism is because $\mathbb{T} \cong U(1)$. Using these ideas (mainly that of the short exact sequence), we can define $\mathrm{Spin}_n^{\mathbb{C}}$ structures.

Definition 2.3.5. Let P_{SO_n} denote the principal SO_n-bundle over X. A $\mathrm{Spin}_n^{\mathbb{C}}$ *structure* on P_{SO_n} consists of a $U(1)$-bundle $P_{U(1)}$ and a principal $\mathrm{Spin}_n^{\mathbb{C}}$-bundle $P_{\mathrm{Spin}_n^{\mathbb{C}}}$ such that the map

$$\psi^{\mathbb{C}} : P_{\mathrm{Spin}_n^{\mathbb{C}}} \to P_{\mathrm{SO}_n} \times P_{U(1)}$$

is Spin_n-equivariant; that is, there exists a 2-sheeted covering

$$\psi_0^{\mathbb{C}} : \mathrm{Spin}_n^{\mathbb{C}} \to \mathrm{SO}_n \times U(1)$$

such that $\psi^{\mathbb{C}}(pg) = \psi^{\mathbb{C}}(p)\psi_0^{\mathbb{C}}(g)$ for $p \in P_{\mathrm{Spin}_n^{\mathbb{C}}}$ and $g \in \mathrm{Spin}_n^{\mathbb{C}}$. An oriented Riemannian manifold with a $\mathrm{Spin}_n^{\mathbb{C}}$-structure on its tangent (frame) bundle is called a $\mathrm{Spin}_n^{\mathbb{C}}$-*manifold*.

Remark 2.3.6. A key result to determine when a sufficiently conditioned manifold has a $\mathrm{Spin}_n^{\mathbb{C}}$ structure is when the second integral Stiefel-Whitney class vanishes. We will not elaborate on the details of this as they are unimportant in the overall context of these notes.

There is an important result that also allows us to apply $\mathrm{Spin}^{\mathbb{C}}$ structures to Heegaard-Floer homology. It is yet another reason why three manifolds are awesome.

Theorem 2.3.7. *Any compact closed and orientable three (or less) manifold is $Spin_n$, and by extension also $Spin^{\mathbb{C}}$ as well.*

3 Remarks on symplectic geometry

Before we actually start defining the key components used in Heegaard-Floer homology, we need some complex and symplectic geometry. Although this seems like a detour, this is actually the main motivation for studying the theory.

3.1 Pseudoholomorphic curves and almost complex structures

We first begin with some important definitions.

Definition 3.1.1. Let M be a smooth manifold. An *almost complex structure* J is an automorphism $J : TM \to TM$ such that $J^2 = -\mathbb{1}$. An *almost complex manifold* is a pair (M, J) consisting of a smooth manifold M and an almost complex structure J on M. The space of almost complex structures will be denoted as

$$\mathcal{J}(M) = \{J \in C^\infty(M, \mathrm{End}(TM)) \mid J^2 = -\mathbb{1}\}.$$

Remark 3.1.2. If M has a complex structure, it must have even dimension. Suppose that $\dim M = n$ for some integer n. If $J^2 = -\mathbb{1}$, then $(\det J)^2 = (-\mathbb{1})^n$. Since M is a real manifold, it follows that $\det J$ is a real number which forces n to be even.

It is important to remember that M is just a smooth manifold, and not a symplectic manifold just yet. However, we will want to explore how our almost complex structure interacts with a nondegenerate (but not closed!) 2-form on M. This is formalized with the following definition.

Definition 3.1.3. Let M be a smooth manifold, and let $\omega \in \Omega^2(M)$ be a nondegenerate 2-form such that $w^{n\times\wedge} > 0$. We say that J is *ω-tame* if $\omega(v, J(v)) > 0$ for $v \in TM$. If $\omega(J(v), J(w)) = \omega(v, w)$ for $v, w \in T_qM$ then we say that J is *ω-compatible*. We denote the spaces of these two kinds of almost complex structures as

$$\mathcal{J}(M, \omega) = \{J \in \mathcal{J}(M) \mid J \text{ is } \omega\text{-tame}\}\ ,\ \mathcal{J}_\tau(M, \omega) = \{J \in \mathcal{J}(M) \mid J \text{ is compatible with } \omega\}.$$

The triple (M, J, ω) where J is compatible with ω is called a *Kähler manifold*.

Remark 3.1.4. There is an important connection between Riemannian geometry and complex geometry that can't be overlooked here. Note that J is compatible with ω if and only if the bilinear form $\langle v, w\rangle_g := \omega(v, J(w))$ is a Riemannian metric on M. We say that such a Riemannian metric $g(v, w) = \langle v, w\rangle$ is *compatible* with J if $\langle J(v), J(w)\rangle = \langle v, w\rangle$ for $v, w \in T_qM$. Thus, there is a one to one correspondence between the nondegenerate 2-forms of M compatible with J and the Riemannian metrics on M compatible with J.

All of this technology can be summarized in the following extremely important theory from complex geometry.

Theorem 3.1.5. *Let M be a smooth manifold. Then M has an almost complex structure J if and only if it carries a nondegenerate 2-form $\omega \in \Omega^2(M)$. Moreover, the homotopy classes of almost complex structures on M are in bijective correspondence with the homotopy classes of nondegenerate 2-forms on M.*

Because of the intimate relationship between complex structures, nondegenerate 2-forms, and Riemannian metrics of a smooth manifold M, we can apply the Fundamental Theorem of Riemannian Geometry to accomplish the following.

Theorem 3.1.6. *Let M be a smooth manifold, let ω be a nondegenerate 2-form on M, let $J \in \mathcal{J}(M, \omega)$ be a ω-compatible almost complex structure, denote by $\langle \cdot, \cdot \rangle_g := \omega(\cdot, J, \cdot)$ the Riemannian metric determined by ω and J, and let ∇ be the Levi-Civita connection of this metric. Then the following hold.*

1. $(\nabla_v J)J + J(\nabla_v J) = 0$ *for* $q \in M$ *and* $u, v, w \in T_q M$
2. $\langle (\nabla_u J)v, w \rangle + \langle v, (\nabla_u J)w \rangle$ *for* $q \in M$ *and* $u, v, w \in T_q M$
3. *For some 3-form* $\omega \in \Omega^3(M)$ *and* $q \in M$ *such that* $u, v, w \in T_q M$*, we have*

$$d\omega(u, v, w) = \langle (\nabla_u J)v, w \rangle + \langle (\nabla_v J)w, u \rangle + \langle (\nabla_w J)u, v \rangle.$$

4. *If* $\omega \in \Omega^3(M)$ *is closed and* $q \in M$ *such that* $v \in T_q M$*, we have*

$$(\nabla_{Jv} J) = -J(\nabla_v J).$$

Remark 3.1.7. A non trivial result is that if L is a totally real submanifold of (M, J) then there exists a Riemannian metric g on M such that g is compatible with J.

We now present our next important object of this section, which is definitionally obvious but has really important implications later on.

Definition 3.1.8. A *complex manifold* M of complex dimension n (or that $n = 2k$ for some k) is a smooth manifold equipped with an atlas whose coordinate charts take values in $\mathbb{C}^n$ and have holomorphic transition maps.

Example. If $n = 1$, we say that M is a *Riemann surface*. If $n = 2$, then M is a *complex surface*.

What happens when we strip the complex properties of M away but keep the same dimensions? We have the following definition immediately.

Definition 3.1.9. Let M be a $2n$-manifold. An almost complex structure J on M is *integrable* if M can be covered by coordinate charts $\phi : U \to \varphi(U) \subset \mathbb{R}^n$ such that the almost complex structure is represented by the matrix

$$J_0 := \begin{pmatrix} 0 & -\mathbb{1} \\ \mathbb{1} & 0 \end{pmatrix}$$

in local coordinates, i.e. it satisfies the *Cauchy-Riemann equation*

$$d\phi(q) \circ J_q - J_0 \circ d\phi(q) = 0$$

or all $q \in U \subseteq M$, which can be viewed as a map from $T_q M$ to $\mathbb{R}^{2n}$.

Remark 3.1.10. An almost complex structure J is integrable if and only if its Nijenhuis tensor $N_J = 0$. We won't prove this or elaborate since it's not that important in this context, but still useful to know.

We finish this section with an important definition that is a precursor of what is to come next.

Definition 3.1.11. A parametrized *pseudoholomorphic curve* in an almost complex manifold (M, J) is a holomorphic map of a Riemann surface $(\Sigma, \mathfrak{j})$ into (M, J).

There are many things omitted here of course, some of which are relatively important and others being nice anecdotes. In general it is important to note that the above is merely a summary.

3.2 The Maslov index

The Maslov index is a map which takes a loop in the fundamental group Lagrangian subspaces $\mathcal{L}(n) = L(\mathbb{R}^{2n}, \omega_0)$ and returns an integer. It turns out that this an isomorphism, so the Maslov index carries a great deal of geometric information.

Definition 3.2.1 (Symplectic Matrix)**.** The *Maslov index* is a unique function $\mu : \pi_1(\mathrm{Sp}(2n)) \to \mathbb{Z}$ which assigns an integer $\mu(\Psi)$ to every loop

$$\Psi : \mathbb{R}/\mathbb{Z} \to \mathrm{Sp}(2n)$$

in $\pi_1(\mathrm{Sp}(2n))$ of symplectic matrices and satisfies the following axioms.

1. Two loops in $\Psi_1, \Psi_2 \in \pi_1(\mathrm{Sp}(2n))$ are homotopic if and only if $\mu(\Psi_1) = \mu(\Psi_2)$.
2. For any two loops in $\Psi_1, \Psi_2 \in \pi_1(\mathrm{Sp}(2n))$, we have

$$\mu(\Psi_1\Psi_2) = \mu(\Psi_1) + \mu(\Psi_2)$$

   where we assume that the constant loop $\Psi(t) \equiv \mathbb{1}$ has $\mu(\Psi) = 0$.
3. If $k = m+n$, consider the subgroup $\mathrm{Sp}(2m) \oplus \mathrm{Sp}(2n)$ of $\mathrm{Sp}(2k)$ in the obvious manner. Then for two loops $\Psi_1 \in \pi_1(\mathrm{Sp}(2m))$ and $\Psi_2 \in \pi_1(\mathrm{Sp}(2n))$, the Maslov index is additive. That is,

$$\mu(\Psi_1 \oplus \Psi_2) = \mu(\Psi_1) + \mu(\Psi_2).$$

4. The loop $\Psi : \mathbb{R}/\mathbb{Z} \to \mathrm{Sp}(2n)$ defined by

$$\Psi(t) := \begin{pmatrix} \cos(2\pi t) & -\sin(2\pi t) \\ \sin(2\pi t) & \cos(2\pi t) \end{pmatrix}$$

   has Maslov index $\mu(\Psi) = 1$. We call this axiom the *normalization axiom*.

Proof. It first suffices to show that the inclusion $U(n) \to \mathrm{Sp}(2n)$ is a homotopy equivalence. Define the map $f : [0,1] \times \mathrm{Sp}(2n) \to \mathrm{Sp}(2n)$ by

$$f(t, \Psi) := f_t(\Psi) := \Psi(\Psi^T\Psi)^{-t/2}.$$

It's trivial to see that $\Psi^T\Psi$ is a symmetric, positive definite symplectic matrix. We want to show that $(\Psi^T\Psi)^\alpha$ is symplectic for $\alpha \geq 0$. Every positive definite symmetric matrix has positive eigenvalues and an orthonormal basis of eigenvectors. Hence there is an orthogonal decomposition

$$\mathbb{R}^{2n} = \bigoplus_{i=1}^{k} E_i \ , \ E_i = \ker(\Psi^T\Psi - \lambda_i\mathbb{1})$$

where the λ_i are the eigenvalues of $\Psi^T\Psi$. Now, choose two vectors in $\mathbb{R}^n$ written as

$$z = \sum_{i=1}^{k} z_i \ , \ w = \sum_{i=1}^{k} w_i$$

where $z_i, w_i \in E_i$. Since $\Psi^T\Psi$ is assumed to be symplectic, we have

$$\omega_0(z_i, w_j) = \omega_0(\Psi^T\Psi(z_i), \Psi^T\Psi(w_j)) = \lambda_i\lambda_j\omega_0(z_i, w_j).$$

This implies that either $\lambda_i\lambda_j = 1$ or $\omega_0(z_i, w_j) = 0$. Hence,

$$\omega_0\left\{(\Psi^T\Psi)^\alpha z, (\Psi^T\Psi)^\alpha w\right\} = \sum_{i,j=1}^{k} (\lambda_i\lambda_j)^\alpha \omega_0(z_i, w_j) = \sum_{i,j=1}^{k} \omega_0(z_i, w_j) = \omega_0(z, w)$$

for every α_0, which implies that $(\Psi^T\Psi)^\alpha$ is symplectic for $\alpha \geq 0$. We now apply this to the map f. In particular, this implies that $(\Psi^T\Psi)^{-t/2}$ is symplectic, which means that $f_t(\Psi) \in \mathrm{Sp}(2n)$ for every $t \geq 0$ and every $\Psi \in \mathrm{Sp}(2n)$. Furthermore, f is continuous, $f_0 = \mathrm{id}$, $f_1|_{U(n)} = \mathrm{id}$ for every $t \in [0, 1]$, and $f_1(\mathrm{Sp}(2n)) = U(n)$ since $f_1(\Psi)$ is symplectic and orthogonal. Thus, $f_1 : \mathrm{Sp}(2n) \to U(n)$ is the desired homotopy inverse of the inclusion $U(n) \to \mathrm{Sp}(2n)$, which implies what we want.

Now we show that $\pi_1(U(n)) \cong \mathbb{Z}$. To see this, consider the complex determinant map

$$\det_{\mathbb{C}} : U(n) \to \mathbb{S}^1,$$

which is a fibration with fiber $\mathrm{SU}(n)$. We see that there is an exact sequence of homotopy groups

$$\pi_1(\mathrm{SU}(n)) \to \pi_1(U(n)) \to \pi_1(\mathbb{S}^1) \to \pi_0(\mathrm{SU}(n)),$$

and since $\mathrm{SU}(n)$ is simply connected by induction, we have $\pi_1(U(n)) \cong \pi_1(\mathbb{S}^1) \cong \mathbb{Z}$. All that's left to do is to construct μ. Define the map $\rho : \mathrm{Sp}(2n) \to \mathbb{S}^1$ by

$$\rho(\Psi) := \det_{\mathbb{C}}(X + iY)\ ,\ \begin{pmatrix} X & -Y \\ Y & X \end{pmatrix} := \Psi(\Psi^T\Psi)^{-1/2} \in \mathrm{Sp}(2n) \cap O(2n)$$

for $\Psi \in \mathrm{Sp}(2n)$ by using the fact that $U(n) \subseteq \mathrm{Sp}(2n)$. We construct the Maslov index μ of the loop $\Psi(t) = \Psi(t + 1) \in \mathrm{Sp}(2n)$ is defined as the degree of the composition $\rho \circ \Psi : \mathbb{R}/\mathbb{Z} \to \mathbb{S}^1$ such that $\mu(\Psi) := \deg(\rho \circ \Psi)$. All of the axioms follow from the above work (except for existence, but this is easy to show). ∎

An equivalent definition can be made for loops of Lagrangian subspaces, but we won't flesh out the details here since the two definitions are literally equivalent to each other.

4 Building up to Heegaard-Floer homology

We now have all of the tools we need to get the gist of what is going on here. In order to define Heegaard-Floer homology, however, there are several objects, functions, spaces, and theorems needed so that it makes sense. This will all come together in one large "aha moment," but this moment will take quite a bit of work to reach.

4.1 Symmetric products and totally real tori

In this section, we define the "ambient space" of the theory and the central objects that are used in the theory.

Definition 4.1.1. Let Σ_g be a genus-g Riemann surface corresponding to a pointed Heegaard diagram $(\Sigma_g, \boldsymbol{\alpha}, \boldsymbol{\beta}, z)$. The *symmetric product*, denoted as $\mathrm{Sym}^g(\Sigma_g)$, is the product space

$$\mathrm{Sym}^g(\Sigma_g) = \underbrace{\Sigma_g \times \cdots \times \Sigma_g}_{g\text{-many times}} / S_g$$

where S_g is the symmetric group on g-letters. This can be visualized as the unordered set of g-tuples on Σ_g where the same point can appear more than once. We define the *diagonal* of $\mathrm{Sym}^g(\Sigma_g)$, which we denote as D, as the space of points where at least two entries coincide.

Although this is rather frustrating, there are some ambiguities in this definition and terminology that must be addressed in the two remarks below. The first is in regards to the notation we just used, since it unfortunately is exactly the same notation we used for the genus of our surface.

Remark 4.1.2. Here, the use of g in $\mathrm{Sym}^g(\Sigma_g)$ does not necessarily equal the genus of Σ_g - it is just a placeholder, but we will soon examine the case of when g is precisely the genus of Σ_g.

The other remark that must be made is why exactly $\mathrm{Sym}^g(\Sigma_g)$ is a "space" at all. It turns out it is a manifold in its own right, and we explain why.

Remark 4.1.3. The symmetric product $\mathrm{Sym}^g(\Sigma_g)$ is a manifold by the Fundamental Theorem of Algebra. On a local level, one can see $\mathrm{Sym}^g(\mathbb{C}) \cong \mathbb{C}^n$. But by approaching this globally by considering $\mathrm{Sym}^g(\Sigma_g)$ as an algebraic variety (and using the complex structure on Σ_g), one can apply the Fundamental Theorem of Algebra in the same way to conclude that $\mathrm{Sym}^g(\Sigma_g)$ is a manifold.

The symmetric product will act as the ambient space of the homology theory we're going to construct. What's convenient is that the complex structure on Σ_g induces a complex structure on $\mathrm{Sym}^g(\Sigma_g)$, which allows us to apply symplectic topology to the symmetric product. Our next definition will be about special subspaces of $\mathrm{Sym}^g(\Sigma_g)$.

Definition 4.1.4. Let $(\Sigma, \mathfrak{j})$ be a complex manifold for some smooth manifold Σ. A submanifold $L \subset \Sigma$ is called *totally real* of $\dim L = \frac{1}{2}\dim \Sigma$ and the tangent space of L does not contain a complex line, i.e.

$$TL \cap \mathfrak{j}(TL) = \{0\}.$$

With this definition, we define half dimensional "tori" of $\mathrm{Sym}^g(\Sigma_g)$ as the subspaces which are products of g-many attaching circles from our pointed Heegaard diagram. That is,

$$\mathbb{T}_{\boldsymbol{\alpha}} = \alpha_1 \times \cdots \times \alpha_g \ , \ \mathbb{T}_{\boldsymbol{\beta}} = \beta_1 \times \cdots \beta_g .$$

In this definition, each attaching circle $\alpha_i \pitchfork \beta_i$ which means $\mathbb{T}_{\boldsymbol{\alpha}} \pitchfork \mathbb{T}_{\boldsymbol{\beta}}$.

Remark 4.1.5. It follows immediately by definition that every Lagrangian submanifold of Σ is totally real, so our theory aligns perfectly with the machinery of Floer homology.

While it is tempting to state the tori as "totally real," we can't make such a drastic claim without proof. Unlike most sources, we actually prove this instead of magically stating it.

Theorem 4.1.6. *The half dimensional tori $\mathbb{T}_{\boldsymbol{\alpha}}$ and $\mathbb{T}_{\boldsymbol{\beta}}$ contained in $\mathrm{Sym}^g(\Sigma_g)$ are totally real submanifolds, and thus can be called totally real tori.*

Proof. Consider the projection $\pi : \Sigma^{\times g} \to \mathrm{Sym}^g(\Sigma_g)$, which is a holomorphic local diffeomorphism that misses that diagonals of $\mathrm{Sym}^g(\Sigma_g)$. Because $\mathbb{T}_{\boldsymbol{\alpha}}$ and $\mathbb{T}_{\boldsymbol{\beta}}$ miss the diagonals of $\mathrm{Sym}^g(\Sigma_g)$, the result follows from the fact that

$$\alpha_1 \times \cdots \times \alpha_g \subset \Sigma^{\times g} \ , \ \beta_1 \times \cdots \times \beta_g \subset \Sigma^{\times g}$$

are totally real submanifolds of $\Sigma^{\times g}$, which follows trivially from definition. ■

A really important detail that is often overlooked is how $\mathbb{T}_{\boldsymbol{\alpha}} \pitchfork \mathbb{T}_{\boldsymbol{\beta}}$. The reason why is because our entire homology theory depends on choosing arbitrary intersection points in $\mathbb{T}_{\boldsymbol{\alpha}} \cap \mathbb{T}_{\boldsymbol{\beta}}$, so we need these two totally real tori to be inseparable from each other.

4.2 Loops and homotopy theory of symmetric products

Let $x, y \in \mathbb{T}_{\boldsymbol{\alpha}} \cap \mathbb{T}_{\boldsymbol{\beta}}$ be intersection points, which exist because of the transversality condition imposed on the totally real tori. Choose a pair of paths $a : [0,1] \to \mathbb{T}_{\boldsymbol{\alpha}}$ and $b : [0,1] \to \mathbb{T}_{\boldsymbol{\beta}}$ going from x to y, i.e. $a(0) = b(0) = x$ and $a(1) = b(1) = y$.

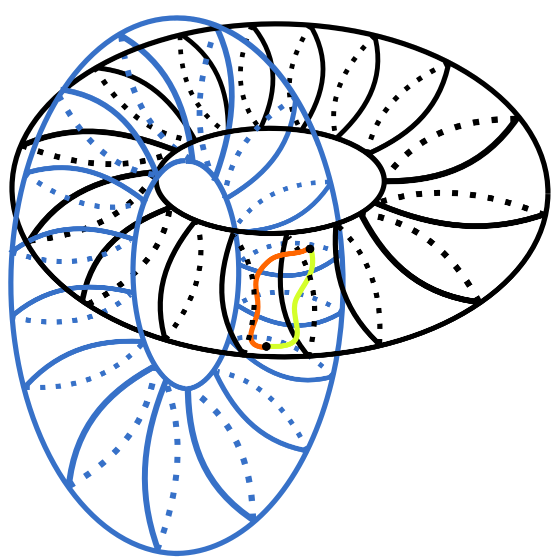

Figure 4.1. A visual depiction of what a and b might look like in action.

If we take the difference of the paths $a - b$, this forms a loop in $\text{Sym}^g(\Sigma_g)$. With this idea of "loops in the intersection of the two tori," we arrive at our next main object of our homology theory.

Definition 4.2.1. The *image of the loop* $a - b$ under the map

$$\frac{H_1(\text{Sym}^g(\Sigma_g))}{H_1(\mathbb{T}_{\boldsymbol{\alpha}}) \oplus H_1(\mathbb{T}_{\boldsymbol{\beta}})}$$

is denoted symbolically as $\epsilon(x, y)$.

Is there a better way to represent $\epsilon(x, y)$? Yes, and to show this we need to compute the fundamental group of $\text{Sym}^g(\Sigma_g)$.

Theorem 4.2.2. *The fundamental group of the symmetric product*

$$\pi_1(\text{Sym}^g(\Sigma_g)) \cong H_1(\text{Sym}^g(\Sigma_g)) \cong H_1(\Sigma_g).$$

Using this theorem from Szabó et al., we can conclude that

$$\frac{H_1(\text{Sym}^g(\Sigma_g))}{H_1(\mathbb{T}_{\boldsymbol{\alpha}}) \oplus H_1(\mathbb{T}_{\boldsymbol{\beta}})} \cong \frac{H_1(\Sigma_g)}{[\alpha_1], \ldots, [\alpha_g], [\beta_1], \ldots, [\beta_g]} \cong H_1(Y; \mathbb{Z})$$

where the third isomorphism is the result of applying Mayer-Vietoris repeatedly to each 2-handle addition on Σ_g. Note that

$$\epsilon(x, z) = \epsilon(x, y) + \epsilon(y, z),$$

so ϵ allows us to partition the intersection points in $\mathbb{T}_{\boldsymbol{\alpha}} \cap \mathbb{T}_{\boldsymbol{\beta}}$ into equivalence classes where $x \sim y$ if $\epsilon(x, y) = 0$. In other words, $x = y$ if there is no loop connecting x and y.

Definition 4.2.3. Let X be a connected space endowed with a basepoint $x \in X$. We denote $\pi_2'(X)$ as the quotient of $\pi_2(X, x)$ by the action of $\pi_1(X, x)$.

We finish this section with an important result that we will use frequently in the upcoming sections. While not apparent now, the use of this theorem will have vast applications later on.

Theorem 4.2.4. *Let* Σ_g *be a Riemann surface of genus* $g > 1$. *Then* $\pi_2'(\text{Sym}^g(\Sigma_g)) \cong \mathbb{Z}$. *If* $g > 2$, *then* $\pi_1(\text{Sym}^g(\Sigma_g))$ *acts trivially on* $\pi_2(\text{Sym}^g(\Sigma_g))$ *and thus* $\pi_2(\text{Sym}^g(\Sigma_g)) \cong \mathbb{Z}$.

4.3 Spin$^{\mathbb{C}}$ structures on three-manifolds

The reason why we covered Spin$^{\mathbb{C}}$ structures so much in chapter 2 is because we want to construct a map

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

where $z \in \Sigma_g \backslash (\boldsymbol{\alpha} \cup \boldsymbol{\beta})$. The secret to doing this for 3 manifolds is by using Turaev's formulation (see [10]) of Spin$^{\mathbb{C}}$ structures on three-manifolds, which is to basically view them as homology classes of vector fields on Y.

Theorem 4.3.1 ([10]). *Let* Y *be a closed connected oriented* 3*-manifold equipped with a Riemannian metric. Then there is a canonical* $H_1(Y)$*-equivariant bijection from the homology classes of vector fields on* Y*, denoted as* $\text{vect}(Y)$*, to* $\text{Spin}^{\mathbb{C}}(Y)$.

Remark 4.3.2. In short, all this theorem is saying is that homology classes of vector fields, in the three dimensional case, are the same thing as Spin$^{\mathbb{C}}$ structures.

But how does this formulation agree with our previous work above? The correspondence with the more traditional definition of $\text{Spin}^{\mathbb{C}}$ structures is given by associating to the vector v the "canonical" $\text{Spin}^{\mathbb{C}}$ structure associated to the reduction of the structure group of TY to SO(2). The motivation for turning $\text{Spin}^{\mathbb{C}}$ structures into homology classes of vector fields is because of two reasons.

- Similar to Seiberg-Witten theory, we want to work with $\text{Spin}^{\mathbb{C}}$ structures and understand their behavior on Y.
- However, in Heegaard-Floer homology, there is a twist. We will soon be looking at flow trajectories on Y, and we want to relate our $\text{Spin}^{\mathbb{C}}$ structures to those flow trajectories using vector fields.

In order to define s_z, we need a definition that relates Morse theory with Heegaard diagrams. Using Morse theory we can describe flow trajectories and then apply $\text{Spin}^{\mathbb{C}}$ structures to define what we want.

Definition 4.3.3. Let $f : Y \to [0, 3]$ be a self-indexing Morse function on Y with one minimum and one maximum. Then f induces a Heegaard decomposition with Heegaard surface $\Sigma_g = f^{-1}(3/2)$, and with handlebodies $U_0 = f^{-1}(0, 3/2)$ and $U_1 = f^{-1}(3/2, 3)$. The attaching circles $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ are the intersections of Σ_g with the ascending and descending manifolds for the index-one and index-two critical points respectively (with respect to some choice of Riemannian metric over Y). We call such a Morse function f *compatible* with the Heegaard diagram $(\Sigma_g, \boldsymbol{\alpha}, \boldsymbol{\beta})$.

Let f be a Morse function on Y compatible with $(\Sigma_g, \boldsymbol{\alpha}, \boldsymbol{\beta})$. Then each $x \in \mathbb{T}_{\boldsymbol{\alpha}} \cap \mathbb{T}_{\boldsymbol{\beta}}$ determines a g-tuple of trajectories for the gradient flow of f connecting the index-one critical points to the index-two critical points of f. Moreover, the basepoint z determines a single gradient flow trajectory of f connecting the index-zero critical point of f to the index-three critical point of f. Delete the tubular neighborhoods surrounding these $g + 1$-many gradient flow trajectories of f. What remains is a subset of Y where the gradient vector field ∇f does not vanish. Since each trajectory connects critical points of different parities (one and two, zero and three), the gradient vector field has index 0 on all the boundary spheres of the subset, so it can be extended as a nowhere vanishing vector field over the entire three-manifold Y.

Definition 4.3.4. The homology class of the nowhere vanishing vector field obtained by this surgery gives the $\text{Spin}^{\mathbb{C}}$ structure s_z, defined as a map

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

Remark 4.3.5. A result from Szabó et al. shows that s_z does not depend on f or the extension of ∇f to the balls on Y, but instead it depends on the choice of intersection point $x \in \mathbb{T}_{\boldsymbol{\alpha}} \cap \mathbb{T}_{\boldsymbol{\beta}}$ and the choice of basepoint z.

4.4 Holomorphic disks, additive assignments, and periodic domains

Motivated from Gromov's 1985 work, Szabó and Osváth wanted to look at properties of (pseudo) holomorphic disks. We define these in the following way.

Definition 4.4.1. Let $\mathbb{D} = \{z \in \mathbb{C} \mid |z| \leq 1\}$ denote the unit disk in $\mathbb{C}$, and let $e_1 \in \partial\mathbb{D}$ and $e_2 \in \partial\mathbb{D}$ denote the arcs of $\mathbb{D}$ such that $\text{Re}(z) \geq 0$ and $\text{Re}(z) \leq 0$ respectively. For $g > 2$, define $\pi_1(x, y)$ as the set of homotopy classes of *Whitney disks* connecting x and y

$$\pi_2(x, y) = \left\{ u : \mathbb{D} \to \text{Sym}^g(\Sigma_g) \;\middle|\; u(-i) = x, u(i) = y, u(e_1) \subset \mathbb{T}_{\boldsymbol{\alpha}}, u(e_2) \subset \mathbb{T}_{\boldsymbol{\beta}} \right\}.$$

Remark 4.4.2. If $g = 2$, then $\pi_2(x, y) = \pi_2(x, y)/\sim$ where $u \sim u'$ if they are homotopic by slicing a fixed sphere in $\pi_2(\text{Sym}^g(\Sigma_g))$. We assume that $\pi_2(x, y) = 0$ if $\epsilon(x, y) \neq 0$.

The set of Whitney disks is equipped with an algebraic structure which we describe here. First, there is a *natural splicing action* defined by

$$\pi_2'(\text{Sym}^g(\Sigma_g)) \star \pi_2(x, y) \to \pi_2(x, y).$$

This operation "$\star$" basically splices two whitney disks connecting x to y and y to z to create a new disk connecting x to z. We write it as

$$\star : \pi_2(x, y) \times \pi_2(y, z) \to \pi_2(x, z)$$

which is trivially associative. Moreover, if $x = y$ then $(\star, \pi_2(x, y))$ is actually a group. Using this star operation, we can define two key components of Heegaard-Floer homology.

Definition 4.4.3. An *additive assignment* is a collection of maps $\mathcal{A} = \{A_{x,y} : \pi_2(x, y) \to \mathbb{Z}\}_{x,y \in \mathbb{T}_{\boldsymbol{\alpha}} \cap \mathbb{T}_{\boldsymbol{\beta}}}$ satisfying the relation

$$A_{x,y}(u) + A_{y,z}(u') = A_{x,z}(u \star u')$$

for holomorphic disks $u \in \pi_2(x, y)$ and $u' \in \pi_2(y, z)$.

A geometric interpretation of an additive assignment is that it's a really well-behaved intersection number. In fact, this is precisely why we use additive assignments, and we'll see an example of this in the next definition.

Definition 4.4.4. Let $z \in \Sigma_g \backslash (\boldsymbol{\alpha} \cup \boldsymbol{\beta})$ be a fixed basepoint. Then the intersection number $n_z(u)$ of a Whitney disk $u \in \pi_2(x, y)$ is the additive assignment

$$n_z : \pi_2(x, y) \to \mathbb{Z} \ , \ n_z(u) = \# u^{-1}(\{z\} \times \text{Sym}^{g-1}(\Sigma_g)).$$

In this case, our additive assignment n_z is counting the number of times the Whitney disk intersects $\text{Sym}^g(\Sigma_g)$ about z. We finish this section with a definition and a theorem.

Definition 4.4.5. Let $\mathcal{D}_1, \dots, \mathcal{D}_m$ denote the closures of the components of $\Sigma_g \backslash (\boldsymbol{\alpha} \cup \boldsymbol{\beta})$. Given a Whitney disk $u \in \pi_2(x, y)$, the domain $\mathcal{D}(u)$ of u is the formal linear combination of the domains $\{\mathcal{D}_i\}_{i=1}^m$ with coefficients $n_{z_i}(u)$. That is,

$$\mathcal{D}(u) = \sum_{i=1}^{m} n_{z_i}(u) \mathcal{D}_i.$$

For a pointed diagram $(\Sigma_g, \boldsymbol{\alpha}, \boldsymbol{\beta}, z)$, a *periodic domain* is a linear combination

$$\mathcal{P} = \sum_{i=1}^{m} a_i \mathcal{D}_i$$

for arbitrary coefficients a_i, whose boundary is a sum of α and β curves such that $n_z(\mathcal{P}) = 0$. For an intersection point $x \in \mathbb{T}_{\boldsymbol{\alpha}} \cap \mathbb{T}_{\boldsymbol{\beta}}$, a Whitney disk $u \in \pi_2(x, x)$ with $n_z(u) = 0$ is called a *periodic class* and we denote the subgroup of periodic classes of x as $\prod_x(z)$.

All of this time we were working with Whitney disks, but we still don't know what they are related to algebraically. The next theorem answers this question for certain genuses.

Theorem 4.4.6. *If $g > 1$, then $\pi_2(x,x) \cong \mathbb{Z} \oplus H^1(Y;\mathbb{Z})$ which identifies the periodic classes $\prod_x(z) \cong H^1(Y;\mathbb{Z})$. In general,*

$$\pi_2(x,y) \cong \mathbb{Z} \oplus H^1(Y;\mathbb{Z})$$

as principal $\pi_2'(\mathrm{Sym}^g(\Sigma_g)) \times \prod_x(z)$ spaces.

Now that we have defined Whitney disks properly, we finish this section with a result that will become relevant when working with Heegaard-Floer chain operators.

Theorem 4.4.7. *Let $S \in \pi_2'(Sym^g(\Sigma_g))$ be a positive generator. Then for a Whitney disk $u \in \pi_2(x,y)$, we have*

$$\mu(u + k[S]) = \mu(u) + 2k$$

where μ is the Maslov index. Moreover, if $O_x \in \pi_2(x,x)$ is the homotopy class of the constant map, then the Maslov index is

$$\mu(O_x + kS) = 2k$$

where k is of course an arbitrary constant.

Why does the Maslov index make an appearance again here? This will become clearer later when the moduli space of some $\phi \in \pi_2(x,y)$ is defined, but in general it's considered the dimension of such a moduli space. It's defined by an elliptic operator here, and if you perturb the almost complex structure of $\mathrm{Sym}^g(\Sigma_g)$ in an n-dimensional family, the new dimension of this perturbed moduli space is $n + \mu(\phi)$ around solutions that are smoothly cut out by the defining equation.

Remark 4.4.8. This is also closely related to the idea of domains defined above.

4.5 The moduli space

In this section, we define the final ingredients and theorems that we'll need to define Heegaard-Floer homology. We begin with a definition.

Definition 4.5.1. Fix a Kähler structure $(\mathfrak{j}, \eta)$ over a Riemann surface Σ_g, a finite collection of basepoints

$$\{z_i\}_{i=1}^m \subset \Sigma_g \backslash (\boldsymbol{\alpha} \cup \boldsymbol{\beta}),$$

and an open set $V \subset \mathrm{Sym}^g(\Sigma_g)$ such that

$$\left(\{z_i\}_{i=1}^m \times \mathrm{Sym}^{g-1}(\Sigma_g) \cup D\right) \subset V \ , \ \overline{V} \cap (\mathbb{T}_{\boldsymbol{\alpha}} \cap \mathbb{T}_{\boldsymbol{\beta}}) = \oslash.$$

An almost-complex structure J on $\mathrm{Sym}^g(\Sigma_g)$ is called $(\mathfrak{j}, \eta, V)$-*nearly symmetric* if

1. If $\pi : \Sigma^{\times g} \to \mathrm{Sym}^g(\Sigma_g)$ is a quotient map, J tames the Kähler form $\pi_*(\omega_0)$ over $\mathrm{Sym}^g(\Sigma_g)\backslash D$ where $\omega_0 = \eta^{\times g}$.
2. J agrees with $\mathrm{Sym}^g(\mathfrak{j})$ over V.

The space of such structures will be denoted by $\mathcal{J}(\mathfrak{j}, \eta, V)$.

The reasoning for this definition will not become clear until we propose our next main definition, that of the moduli space for our homology theory.

Definition 4.5.2. Let $\mathbb{D} = [0,1] \times i\mathbb{R} \subset \mathbb{C}$ be a strip in the complex plane. Fix a path J_s of almost-complex structures over $\mathrm{Sym}^g(\Sigma_g)$. We define the moduli space of *pseudoholomorphic strips* connecting intersection points x and y to be the set

$$\mathcal{M}_{J_s}(x,y) = \left\{ u : \mathbb{D} \to \mathrm{Sym}^g(\Sigma_g) \;\middle|\; \begin{array}{c} u(\{1\} \times \mathbb{R}) \subset \mathbb{T}_{\boldsymbol{\alpha}} \\ u(\{0\} \times \mathbb{R}) \subset \mathbb{T}_{\boldsymbol{\beta}} \\ \lim_{t \to -\infty} u(s,t) = x \\ \lim_{t \to \infty} u(s,t) = y \\ \frac{du}{ds} + J(s)\frac{du}{dt} = 0 \end{array} \right\}.$$

Given a Whitney disk $\phi \in \pi_2(x,y)$, the space $\mathcal{M}_{J_s}(\phi)$ denotes the subset of pseudoholomorphic strips $\mathcal{M}_{J_s}(x,y)$ that represent ϕ as a homotopy class. We define the space of *pseudoholomorphic J_s-disks* representing the Whitney disk ϕ as the quotient

$$\widehat{M_{J_s}}(\phi) = \mathcal{M}_{J_s}(\phi)/\mathbb{R},$$

since the translation action on $\mathbb{D}$ endows the moduli space $\mathcal{M}_{J_s}(\phi)$ with an $\mathbb{R}$ action.

Remark 4.5.3. The word "disk" is used, in view of the holomorphic identification of the strip with the unit disk in the complex plane with two boundary points removed (and maps in the moduli space extend across these points, in view of the asymptotic conditions).

There are many properties of the moduli space which we will not need here necessarily. Just assume that the moduli space admits proper orientations, is compact, is smoothly cut out by its defining equations, and that

$$\dim(\widehat{\mathcal{M}}(\phi)) = \mu(\phi) - 1.$$

A small notational note is that when accounting for orientations, it is commonplace to include the notation $\mathfrak{o}$ which denotes a *coherent system of orientations* of the $\mathrm{Spin}^{\mathbb{C}}$ structure $\mathfrak{s}$ for the Heegaard-Floer homology.

5 Heegaard-Floer Homology

We are now ready to define Heegaard-Floer homology and start doing stuff with it. But before we begin, a small disclaimer is in order: the Betti number $b_1(Y)$ determines how complicated our homology theory is. For now, we assume that $b_1(Y) = 0$. Fixing the first Betti number at 0, let $(\Sigma_g, \boldsymbol{\alpha}, \boldsymbol{\beta}, z)$ be a pointed Heegaard diagram with $g > 0$, and choose a $\text{Spin}^{\mathbb{C}}$ structure $\mathfrak{s} \in \text{Spin}^{\mathbb{C}}(Y)$. Define the set of intersection points

$$\mathfrak{S} = \{x \in \mathbb{T}_{\boldsymbol{\alpha}} \cap \mathbb{T}_{\boldsymbol{\beta}} \mid s_z(x) = \mathfrak{s}\}.$$

We fix the following auxiliary data for our homology theories.

1. In the case that our homology has integral coefficients, we assume that it possesses a coherent orientation system $\mathfrak{o}$. This will not be necessary when we're considering $\mathbb{Z}/2\mathbb{Z}$ coefficients.

2. We assume that Σ_g possesses a "generic" complex structure $\mathfrak{j}$ such that each intersection point $x \in \mathbb{T}_{\boldsymbol{\alpha}} \cap \mathbb{T}_{\boldsymbol{\beta}}$ is disjoint from the $\text{Sym}^g(\mathfrak{j})$-holomorphic spheres in $\text{Sym}^g(\Sigma_g)$.

3. We finally assume that $\text{Sym}^g(\Sigma_g)$ has a "generic" path of nearly symmetric almost complex structures J_s over $\text{Sym}^g(\Sigma_g)$, contained in an open contractible neighborhood $\mathcal{U} \subset \text{Sym}^g(\mathfrak{j})$.

Before we define any homology theories, there is one last ingredient we need to make our definitions precise.

Definition 5.0.1. A *relatively graded Abelian group G* is one which is generated by elements partitioned into equivalence classes $\mathfrak{S}$, with a relative grading function

$$\text{gr} : \mathfrak{S} \times \mathfrak{S} \to \mathbb{Z}$$

satisfying $\text{gr}(x, y) + \text{gr}(y, w) = \text{gr}(x, w)$ for $x, y, z \in \mathfrak{S}$.

Whenever we want to construct a Heegaard-Floer homology theory, we have to introduce a relative grading. While it's an extra step, it will make sense when implemented in our first definition.

5.1 Hat Heegaard-Floer homology

Definition 5.1.1. Let $\widehat{CF}(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s})$ denote the free abelian group generated by the points $\mathfrak{S} \subset \mathbb{T}_{\boldsymbol{\alpha}} \cap \mathbb{T}_{\boldsymbol{\beta}}$. This group can be endowed with the relative grading

$$\text{gr}(x, y) = \mu(u) + 2n_z(u)$$

where $u \in \pi_2(x, y)$ is a Whitney disk. Write the boundary operator as

$$\widehat{\partial} : \widehat{CF}(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s}) \to \widehat{CF}(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s}),$$

and define the moduli space $\widehat{\mathcal{M}}_{J_s}(\phi) := \widehat{\mathcal{M}}^0_{J_s}(x,y)$ for the element $\phi \in \pi_2(x,y)$ where $\mu(\phi) = 1$ and $n_z(\phi) = 0$. Then we define the boundary operator as

$$\widehat{\partial}_{J_s} x = \sum_{\{y \in \mathfrak{S} \ | \ \mathrm{gr}(x,y)=1\}} \# \left(\widehat{\mathcal{M}^0_{J_s}}(x,y) \right) y.$$

Remark 5.1.2. By Theorem 4.4.7, there is one and only one such choice of $\phi \in \pi_2(x,y)$.

By analyzing the Gromov compactification of $\widehat{M}(\phi)$ for $n_z(\phi) = 0$ and $\mu(\phi) = 2$, we achieve the following.

Theorem 5.1.3 (Szabó et al., 2003)**.** *When $b_1(Y) = 0$, the pair $\left(\widehat{CF}(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s}), \widehat{\partial}\right)$ is a chain complex with $\widehat{\partial}^2 = 0$.*

Proofs of the proceeding theorems have been omitted since they rely on completely separate details. In spite of this, however, we can now define our first Heegaard-Floer homology theory.

Definition 5.1.4. The *Hat Heegaard-Floer Homology* (HHFH) groups $\widehat{HF}(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s})$ are the homology groups of the chain complex $\left(\widehat{CF}(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s}), \widehat{\partial}\right)$.

These are the simplest kinds of Heegaard-Floer homology groups, and note that quite a lot of the analytical detail to formalize the Gromov compactness and energy is omitted simply because this paper would be quite long if included. However, there are many other Heegaard-Floer homologies to consider.

5.2 Infinite Heegaard-Floer homology

Before we move on to other kinds of Heegaard-Floer homology theories, we have to remark on a very specific condition that we used before. When defining HHFH, we were counting pseudoholomorphic J_s-disks which avoid the subvariety

$$\{z\} \times \mathrm{Sym}^{g-1}(\Sigma_g)$$

since we were assuming that $n_z(\phi) = 0$ by definition. This isn't necessarily an issue, it's just restrictive. Our next Heegaard-Floer homology theory will remove this restriction to a much broader setting.

Definition 5.2.1. Let $CF^\infty(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s})$ be the free abelian group generated by the pairs $[x,i]$ where $s_z(x) = \mathfrak{s}$ and $i \in \mathbb{Z}$. These generators have a relative grading

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

Let the map between these free abelian groups be

$$\partial^\infty : CF^\infty(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s}) \to CF^\infty(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s})$$

such that we can write it as the double sum

$$\partial^\infty[x,i] = \sum_{y \in \mathbb{T}_{\boldsymbol{\alpha}} \cap \mathbb{T}_{\boldsymbol{\beta}}} \sum_{\{\phi \in \pi_2(x,y) \ | \ \mu(\phi)=1\}} \# \left(\widehat{\mathcal{M}}(\phi) \right) [y, i - n_z(\phi)]$$

Of course, we can construct a chain complex from these boundary operators, although the way to prove this is quite complicated.

Theorem 5.2.2. *If $b_1(Y) = 0$, the pair $(CF^\infty(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s}), \partial^\infty)$ is a chain complex with $(\partial^\infty)^2 = 0$. Its corresponding homology groups are denoted by $HF^\infty(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s})$ called infinite Heegaard-Floer homology groups (IHFH).*

What's kind of disappointing is that these don't encode very good invariants. Here's why. Consider the chain map

$$U : CF^\infty(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s}) \to CF^\infty(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s})\ ,\ U([x, i]) = [x, i-1]$$

which lowers the degree of the chain by two. It turns out that $HF^\infty(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s}) \cong \mathbb{Z}[U, U^{-1}]$ in all cases, so it's not fun to compute this kind of Heegaard-Floer homology. To make this more interesting, let $CF^-(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s}) \subseteq CF^\infty(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s})$ be a free subgroup generated by the pairs $[x, i]$ for $i < 0$, and denote the quotient group

$$CF^+(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s}) := CF^\infty(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s})/CF^-(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s}).$$

Using a lemma from Osváth and Szabó's original paper, we can prove the following.

Theorem 5.2.3. *There is an exact sequence*

$$0 \longrightarrow CF^-(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s}) \xrightarrow{incl} CF^\infty(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s}) \xrightarrow{\pi} CF^+(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s}) \longrightarrow 0.$$

Proof. All you have to show is that $CF^-(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s})$ is a subcomplex of $CF^\infty(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s})$, which follows from a theorem not mentioned. ■

We define the Heegaard-Floer homology groups $HF^+(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s})$ and $HF^-(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s})$ in the obvious way, and they are all $\mathbb{Z}[U]$ modules. But there is a slight remark that must be made here.

Remark 5.2.4. The choice of orientations and complex structures is irrelevant, although this should really be proved. The reason why it doesn't matter is because of the fact that we are assuming that $b_1(Y) = 0$. If we don't do this, we must specify our choice of orientation and complex structure.

We finish this chapter with a beautiful summary of the theory of Heegaard-Floer homology, and what takes almost 50 pages for the original creators of this genre to prove.

Theorem 5.2.5. *All of the Heegaard diagrams mentioned above are invariant under the choice of Heegaard diagram, are isomorphic to each other for any three-manifold Y with suitable hypotheses, and are topological invariants.*

Remark 5.2.6. Note that this is true for $b_1(Y) > 0$. The reason why this is an issue in the first place is that $\pi_2(x, y)$ is larger, so now we can have infinitely many homotopy classes $\phi \in \pi_2(x, y)$ with Maslov index one. To keep our sums finite, one has to redefine "admissibility" on the $\mathrm{Spin}^{\mathbb{C}}$ structure and introduce modified Heegaard diagrams while accounting for orientations and complex structures.

There are many other Heegaard-Floer homology theories. They have a variety of knot-theoretic analogues [3] with most recent additions to the theory being in bordered Heegaard-Floer homology theory, the pong algebra, and the wrapped Fukaya category [1] [4] [5]. The exposition of such theories, of course, would be enormous.

Bibliography

[1] Lipshitz, Robert, Peter S. Ozsváth, and Dylan P. Thurston. "Bordered Heegaard Floer Homology: Invariance and Pairing." *arXiv*, 2008, https://arxiv.org/abs/0810.0687

[2] Ozsváth, Peter, and Zoltán Szabó. "Holomorphic Disks and Topological Invariants for Closed Three-Manifolds." *arXiv*, 2001, https://arxiv.org/abs/math/0101206v4

[3] Ozsváth, Peter, and Zoltán Szabó. "An Overview of Knot Floer Homology." *arXiv*, 2017, https://arxiv.org/abs/1706.07729

[4] Ozsváth, Peter S., and Zoltán Szabó. "The Pong Algebra." *arXiv*, 2022, https://arxiv.org/abs/2212.11885

[5] Ozsváth, Peter, and Zoltán Szabó. "The Pong Algebra and the Wrapped Fukaya Category." *arXiv*, 2022, https://arxiv.org/abs/2212.14420

[6] Ozsváth, Peter, and Zoltán Szabó. "Holomorphic Disks and Three-Manifold Invariants: Properties and Applications." *arXiv*, 2001, https://arxiv.org/abs/math/0105202v4

[7] Ozsváth, Peter S., András I. Stipsicz, and Zoltán Szabó. *Heegaard Floer Homology: Version 4*. arXiv, 2025, https://web.math.princeton.edu/ petero/konyv4.pdf

[8] Lawson, H. Blaine, Jr., and Marie-Louise Michelsohn. *Spin Geometry*. Princeton University Press, 1989.

[9] McDuff, Dusa, and Dietmar Salamon. *Introduction to Symplectic Topology*. 3rd ed., Oxford University Press, 2017.

[10] Turaev, Vladimir. *Torsions of 3-Dimensional Manifolds*. Springer Basel AG, 2002.

Honor Code Pledge: I pledge my honor that I have not violated the honor code while writing this document.

Signature: