Floer Homology and low-dimensional topology

Jacob Rasmussen and Liam Watson

August 26, 2014

Introduction

Temporary disclaimer

These notes are still in draft form and should be considered, for the time being, a rough guide to accompany the lectures. The lecturers appologise for any and all omissions/errors/typos. We are happy to receive feedback!

Lecture 1

Floer homologies for three-manifolds

Floer homology assigns an abelian group to a closed oriented 3-manifold. There are a variety of ways to do this (Heegaard Floer homology, monopole Floer homology, ECH, Manolescu-Woodward, instanton Floer). By now it's known that the first three of these theories are isomorphic, although their relation to the original instanton Floer homology is still not well understood. These different methods all have their own advantages and disadvantages, but in these lectures we'll only use Heegaard Floer homology.

It's perhaps useful to consider the analogy with ordinary homology, where there are many different ways (singular homology, Cech cohomology, de Rham cohomology) of defining what is essentially the same invariant. For ordinary homology, we have the Eilenberg-Steenrod axioms, which spell out the key properties of ordinary homology, and which completely determine it on reasonable spaces (CW complexes). For Floer homology, we have (as yet) no such thing. But in this lecture we'll adopt an Eilenberg-Steenrod approach, and describe at least some of the basic properties that a Floer homology for three-manifolds should have. In the next lecture, we'll see how Heegaard Floer homology realizes these properties.

1.1 TQFT structure

From now on, we'll assume that all three-manifolds Y are closed, connected, and oriented. In its simplest incarnation, Floer homology assigns to such a Y an abelian group $\widehat{HF}(Y)$.

Property 1. $\widehat{HF}(Y_1 \# Y_2) = \widehat{HF}(Y_1) \otimes \widehat{HF}(Y_2)$. $\widehat{HF}(-Y) = \widehat{HF}(Y)^{*_c}$, where by $*_c$ we denote the chain level dual (i.e. if $\widehat{HF}(Y) = H_*(C)$, $\widehat{HF}(-Y) = H_*(C^*)$.)

Example 1.1. $\widehat{HF}(S^3) = \mathbb{Z}$.

Definition 1.2. A *cobordism* $W : Y_1 \to Y_2$ is an oriented smooth 4-manifold W with $\partial W = -Y_1 \coprod Y_2$.

Example 1.3. $B^4 : \emptyset \to S^3$ is a very simple cobordism. If $K \subset S^3$ is a knot, let $W_p(K) : S^3 \to K_p$ be the cobordism defined by attaching a p-framed 2-handle to K.

A cobordism $W : Y_1 \to Y_2$ induces a map $\widehat{F}_W : \widehat{HF}(Y_1) \to \widehat{HF}(Y_2)$, which is well defined up to an overall sign. (If we want to fix the sign, we'd have to choose an orientation on $H^1(W) \oplus H^2_+(W) \oplus H^1(Y_2)$.)

Example 1.4. If $W = B^4 : \emptyset \to S^3$, $\widehat{F}_W : \mathbb{Z} \to \mathbb{Z}$ is the identity map.

Example 1.5. We could equally well view W as a cobordism $r(W) : -Y_2 \to -Y_1$. The induced map $\widehat{F}_{r(W)} : \widehat{HF}(-Y_2) \to \widehat{HF}(Y_1)$ is the chain level dual of $\widehat{F}$. **NB:** There is another cobordism $-W : -Y_1 \to -Y_2$. The maps $\widehat{F}_{-W}$ and $\widehat{F}_W$ are unrelated.

If $W_1 : Y_1 \to Y_2, W_2 : Y_2 \to Y_3$ are cobordisms, their composition $W_2 \circ W_1 := W_1 \cup_{Y_2} W_2$ is a cobordism from $Y_1 \to Y_3$.

Property 2. $\widehat{F}_{W_2 \circ W_1} := \pm \widehat{F}_{W_2} \circ \widehat{F}_{W_1}$.

In somewhat fancier language, this says that $\widehat{HF}$ defines a projective functor from the $3+1$ dimensional cobordism category to the category of abelian groups and linear maps. In other words, $\widehat{HF}$ is a $3+1$ dimensional TQFT.

Remark 1.6. By viewing M^4 as a cobordism $M : \emptyset \to \emptyset$, we get a map $\widehat{F}_M : \mathbb{Z} \to \mathbb{Z}$. This is *not* the Seiberg-Witten invariant of M. Instead, it is determined completely by the ring structure on $H^*(M)$.

1.2 $Spin^c$ structures

An important property of $\widehat{HF}(Y)$ is that it decomposes as a direct sum over $Spin^c$ structures on Y. We'll leave the definition of a $Spin^c$ structure to the exercises, and instead describe a few relevant properties:

- Let M be a 3– or 4–manifold. The set $Spin^c(M)$ of $Spin^c$ structures on a manifold M is an affine copy of $H^2(M)$ (by which we mean that $H^2(M)$ acts freely and transitively on $Spin^c(M)$, but there is no natural 0, so we can't add $Spin^c$ structures.)
- There's a map $c_1 : Spin^c(M) \to H^2(M)$ which satisfies $c_1(y + \mathfrak{s}) = 2y + c_1(\mathfrak{s})$.
- If $M = Y$ is a 3-manifold, $c_1(\mathfrak{s}) \equiv 0 \mod 2$ for all $\mathfrak{s} \in Spin^c(Y)$.
- If $M = W$ is a 4-manifold, $\langle c_1(\mathfrak{s}), x\rangle \equiv x \cdot x \mod 2$ for all $x \in H_2(W)$.
- If $Y = \partial W$, there is a restriction map $|_Y : Spin^c(W) \to Spin^c(Y)$. This is compatible with restriction on cohomology, in the sense that $c_1(\mathfrak{s}|_Y) = i^*(c_1(\mathfrak{s}))$, and $(y + \mathfrak{s})|_Y = i^*(y) + \mathfrak{s}|_Y$.
- There's an involution $\mathfrak{s} \mapsto \bar{\mathfrak{s}}$ on $Spin^c(M)$. $c_1(\bar{\mathfrak{s}}) = -c_1(\mathfrak{s})$.

Example 1.7. If $H^2(Y)$ has no 2-torsion, we can canonically identify $Spin^c(Y)$ with $H^2(Y)$.

Example 1.8. Let $W = W_p(K)$. $H^2(W) = \mathbb{Z}\langle y\rangle$. If Σ is a Seifert surface for K $H_2(W) = \mathbb{Z}$ is generated by $[\widehat{\Sigma}]$, where $\widehat{\Sigma}$ is the closed surface obtained by capping Σ off with the core of the two-handle. $\Sigma \cdot \Sigma = p$, so

$$Spin^c(W) \leftrightarrow \{c_1(\mathfrak{s}) \mid \mathfrak{s} \in Spin^c(W)\} = \{ky \mid k \equiv p \mod 2\}.$$

We let $\mathfrak{s}_k$ be the $Spin^c$ structure with $c_1(\mathfrak{s}_k) = ky$.

Property 3. There is a decomposition

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

Moreover, $\widehat{HF}(Y, \mathfrak{s}) = 0$ for all but finitely many $\mathfrak{s}$.

Property 4. If $W : Y_1 \to Y_2$ and $\mathfrak{s} \in Spin^c(W)$, there is an induced map

$$\widehat{F}_{W,\mathfrak{s}} : \widehat{HF}(W, \mathfrak{s}|_{Y_1}) \to \widehat{HF}(W, \mathfrak{s}|_{Y_2}).$$

$\widehat{F}_{W,\mathfrak{s}}$ is 0 for all but finitely many $\mathfrak{s}$, and $\widehat{F}_W = \sum_{\mathfrak{s}} \widehat{F}_{W,\mathfrak{s}}$.

Property 5. (Conjugation symmetry). There are isomorphisms

$$\iota : \widehat{HF}(Y, \mathfrak{s}) \to \widehat{HF}(Y, \bar{\mathfrak{s}}),$$

and $\iota \circ \widehat{F}_{W,\mathfrak{s}} = \pm \widehat{F}_{W,\bar{\mathfrak{s}}} \circ \iota$.

Property 6. Suppose $W_1 : Y_1 \to Y_2$, $W_2 : Y_2 \to Y_3$ are cobordisms. If Y_2 is a rational homology sphere, then $F_{W_2 \circ W_1, \mathfrak{s}} = F_{W_2, \mathfrak{s}|_{W_2}} \circ F_{W_1, \mathfrak{s}|_{W_1}}$.

If Y is not a rational homology sphere, the analog of property 6 is a bit harder to state. To understand the issue, consider the Mayer-Vietoris sequence for $H^*(W)$, where $W = W_1 \cup_{Y_2} W_2$:

$$\to H^1(W_1) \oplus H^1(W_2) \xrightarrow{i_1^* \oplus i_2^*} H^1(Y_2) \to H^2(W) \to H^2(W_1) \oplus H^2(W_2) \to .$$

If the map i^* is not surjective, a $Spin^c$ structure on W will not be determined by its restriction to W_1 and W_2.

To address this problem, we must introduce a slightly more complicated theory: homology with twisted coefficients. The group $\widehat{HF}(Y, \mathfrak{s}; \mathbb{Z}[H^1(Y)])$ is the homology of a chain complex defined over the group ring $\mathbb{Z}[H^1(Y)]$. If G is a quotient of $H^1(Y)$, we can specialize to get a map

$$\widehat{HF}(Y, \mathfrak{s}; \mathbb{Z}[H^1(Y)]) \to \widehat{HF}(Y, \mathfrak{s}; \mathbb{Z}[G]).$$

If $W : Y_1 \to Y_2$ is a cobordism and $j : Y_2 \to W$ is the inclusion, there is an induced map

$$\widehat{F}_{W,\mathfrak{s}} : \widehat{HF}(Y_1, \mathfrak{s}) \to \widehat{HF}(Y_2; \mathbb{Z}[G])$$

where G is any quotient of $\operatorname{coker} j_*$.

Given W_1 and W_2 as above, we take $G = \operatorname{coker} i_1 \oplus i_2$ and form the maps

$$\widehat{F}_{W_1,\mathfrak{s}_1} : \widehat{HF}(Y_1, \mathfrak{s}_1) \to \widehat{HF}(Y_2, \mathfrak{s}_1; \mathbb{Z}[G])$$
$$\widehat{F}_{W_2,\mathfrak{s}_2} : \widehat{HF}(Y_2, \mathfrak{s}_2; \mathbb{Z}[G]) \to \widehat{HF}(Y_3, \mathfrak{s}_2)$$

where $\mathfrak{s}_1|_{Y_2} = \mathfrak{s}_1|_{Y_2}$. Then the compositions of the form

$$\widehat{F}_{W_2,\mathfrak{s}_2} \circ g\widehat{F}W_1, \mathfrak{s}_2$$

where g runs over elements of G give the maps $\widehat{F}_{W_2 \circ W_1,\mathfrak{s}}$ where $\mathfrak{s}$ runs over $Spin^c$ structures that restrict to $\mathfrak{s}_i$ on W_i.

The final topic we will mention in this section is a very important geometric property of Floer homology. In its original form it is due to Kronheimer.

Property 7. (Adjunction) Suppose $\Sigma \subset W$ is a smoothly embedded surface with $\Sigma \cdot \Sigma \geq 0$. If

$$\Sigma \cdot \Sigma + \langle c_1(\mathfrak{s}), \Sigma \rangle \geq 2g(\Sigma) - 2,$$

then $F_{W,\mathfrak{s}} = 0$.

Using the conjugation symmetry, it is easy to see this can be improved to $F_{W,\mathfrak{s}} = 0$ if $\Sigma \cdot \Sigma + |\langle c_1(\mathfrak{s}), \Sigma \rangle| \geq 2g(\Sigma) - 2$

1.3 +, − and ∞

We digress to recall the definition of S^1 equivariant homology. Suppose X is a space on which S^1 acts. Consider the universal S^1 bundle ES^1 over the universal classifying space BS^1. In plain English, this is the Hopf bundle S^∞ over $\mathbb{CP}^\infty$.

Definition 1.9. The S^1 *equivariant cohomology* $H^*_{S^1}(X) := H^*(X')$, where $X' = (X \times ES^1)/S^1$, and the quotient is by the diagonal S^1 action. The equivariant homology $H_*^{S^1}(X) = H_*(X')$.

Example 1.10. If you've never seen this construction before, it's helpful to think about the cases where the S^1 action on X is trivial, in which case $X' = X \times BS^1$, so $H^*_{S^1}(X) = H^*(X) \otimes \mathbb{Z}[U]$, and where it is free, in which case $X' = (X/S^1) \times ES^1$ and $H^*_{S^1}(X) = H^*(X/S^1)$.

There is an obvious projection $X' \to BS^1$. By pullback, this induces a map $\mathbb{Z}[U] = H^*(BS^1) \to H^*_{S^1}(X)$. Using cup and cap products, this makes $H^*_{S^1}(X)$ and $H_*^{S^1}(X)$ into $\mathbb{Z}[U]$ modules. If we use rational coefficients, the ring $\mathbb{Q}[U]$ is a PID. It follows that $H^*_{S^1}(X;\mathbb{Q})$ decomposes as a direct sum of free modules $\mathbb{Q}[U]$ and torsion modules $\mathbb{Q}[U]/(U^k)$. The free part of the this decomposition is particularly easy to describe:

Theorem 1.11. *(Localization) If* $R = \mathbb{Q}[[U^{-1}, U]$*, then*

$$H^*_{S^1}(X, R) = H^*(X^{S^1}) \otimes R,$$

where X^{S^1} *is the fixed-point locus of the* S^1 *action.*

An important fact about Floer homology is that it formally behaves like S^1 equivariant cohomology. In monopole Floer homology, this is part of the construction; in Heegaard Floer homology it is somewhat less obvious. More precisely, for each $\mathfrak{s} \in Spin^c(Y)$, there are $\mathbb{Z}[U]$–modules $HF^+(Y, \mathfrak{s}), HF^-(Y, \mathfrak{s})$, and $HF^\infty(Y, \mathfrak{s})$, which are "related" to equivariant cohomology by the following dictionary:

$$\widehat{HF}(Y, \mathfrak{s}) \longleftrightarrow H_*(X)$$
$$HF^+(Y, \mathfrak{s}) \longleftrightarrow H_*^{S^1}(X)$$
$$HF^\infty(Y, \mathfrak{s}) \longleftrightarrow H_*^{S^1}(X) \otimes R$$

When Y is a rational homology sphere, Manolescu defined a spectrum $SWF(Y, \mathfrak{s})$ which should play the role of X in the right-hand side of this dictionary. (Manolescu's recent resolution of the triangulation conjecture hinged on the fact that in the case where $\mathfrak{s}$ is fixed by conjugation, the conjugation symmetry makes $SWF(Y, \mathfrak{s})$ into a space with a $Pin(2)$ action, and one can consider its $Pin(2)$ equivariant cohomology.

Example 1.12. $HF^-(S^3) = M_- := \mathbb{Z}[[U]]$, $HF^\infty(S^3) = R$, and $HF^+(S^3) = M_+ := R/M_-$.

Property 8. There are long exact sequences

$$\xrightarrow{\partial} HF^-(Y, \mathfrak{s}) \xrightarrow{i_*} HF^\infty(Y, \mathfrak{s}) \xrightarrow{\pi_*} HF^+(Y, \mathfrak{s}) \xrightarrow{\partial}$$

$$\to \widehat{HF}(Y, \mathfrak{s}) \to HF^+(Y, \mathfrak{s}) \xrightarrow{U} HF^+(Y, \mathfrak{s}) \to \widehat{HF}(Y, \mathfrak{s}) \to .$$

If Y is a rational homology sphere, there's a spectral sequence starting at $\widehat{HF}(Y, \mathfrak{s}) \otimes M_+$ and converging to $HF^+(Y, \mathfrak{s})$

Similarly, if $W : Y_1 \to Y_2$ is a cobordism, there are induced maps

$$F^\circ_{W,\mathfrak{s}} : HF^\circ(Y_1, \mathfrak{s}|_{Y_1}) \to HF^\circ(Y_2, \mathfrak{s}|_{Y_2})$$

where $\circ \in \{+, -, \infty\}$. These maps are functorial under composition of cobordisms and commute in the obvious way with the maps in the exact sequences above.

The group $HF^\infty(Y, \mathfrak{s})$ (corresponding to the homology of the fixed-point set of the S^1 action) is much easier to understand than the other Floer groups. In fact, it is determined by purely homological information.

Property 9. If Y is a rational homology sphere, $HF^\infty(Y, \mathfrak{s}) \simeq R$.
If $W : Y_1 \to Y_2$ is a cobordism between rational homology spheres, then $F^\infty_{W,\mathfrak{s}} : HF^\infty(Y_1, \mathfrak{s}|_{Y_1}) \to HF^\infty(Y_2, \mathfrak{s}|_{Y_2})$ is an isomorphism if $b_2^+(W) = 0$ and is 0 otherwise.

In general, $HF^\infty(Y, \mathfrak{s})$ is determined by the triple cup product on $H^1(Y, \mathfrak{s})$.

Definition 1.13. $HF^{red}(Y, \mathfrak{s}) := \ker i_* \simeq \operatorname{coker} \pi_*$ where the maps as in the first exact sequence of Property 8.

We can now define an interesting invariant of four-manifolds.

Definition 1.14. Suppose M is a closed 4-manifold with $b_2^+(M) > 1$. Choose a rational homology sphere Y which separates M into M_1, M_2 each with $b_2^+ > 0$, and view $M_1 : \emptyset \to Y$ and $M_2 : Y \to \emptyset$ as cobordisms. Choose also a right-inverse $\phi : HF^{red}(Y) \to HF^+(Y)$ for ∂. That is, we want $\partial \circ \phi = \mathrm{id}_{HF^{red}(Y)}$. If $\mathfrak{s} \in Spin^c(M)$, the *mixed invariant* $OS(M, \mathfrak{s})$ is

$$OS(M, \mathfrak{s}) := F^+_{M_2,\mathfrak{s}} \circ \phi \circ F^-_{M_1,\mathfrak{s}}(1)$$

1.4 Gradings

We end this lecture with a brief discussion of gradings in Floer homology.

Property 10. (Relative gradings) If Y is a rational homology sphere $HF^\circ(Y, \mathfrak{s})$ is relatively $\mathbb{Z}$ graded, and U lowers grading by 2. In general, $HF^\circ(Y, \mathfrak{s})$ is relatively $\mathbb{Z}/A$ graded, where $A = \{\langle c_1(\mathfrak{s}), x\rangle \mid x \in H_2(Y)\}$.

If $c_1(\mathfrak{s})$ is a torsion class, so $HF^\circ(Y, \mathfrak{s})$ is relatively $\mathbb{Z}$ graded, there is an absolute $\mathbb{Q}$ grading on $HF^\circ(Y\mathfrak{s})$. This means that

$$HF^\circ(Y, \mathfrak{s}) = \bigoplus_{t \in \mathbb{Q}} HF^{\circ,t}(Y, \mathfrak{s}).$$

We denote this absolute grading by gr. It is compatible with the relative $\mathbb{Z}$ grading, in the sense that if a and b are homogenous elements, then $\operatorname{gr} a - \operatorname{gr} b$ is the relative grading difference between a and b. (Note that all elements of $HF^\circ(Y, \mathfrak{s})$ have the same grading $\mod \mathbb{Z}$.)

Property 11. The absolute grading behaves predictably with respect to maps induced by cobordisms, in following sense. If $W : Y_1 \to Y_2$ is a cobordism, $\mathfrak{s} \in Spin^c(W)$, and a is a homogenous element of $HF^\circ(Y_1, \mathfrak{s})$, then $b := F^\circ_{W,\mathfrak{s}}(a)$ is homogenous and

$$\operatorname{gr} b - \operatorname{gr} a = \frac{c_1(\mathfrak{s})^2 - 2\chi(W) - 3\sigma(W)}{4}$$

Example 1.15. $HF^+(S^3) = M_+$, where $\operatorname{gr} U^{-k} = 2k$. $HF^+(S^1 \times S^2) = \langle 1, \theta \rangle \otimes M_-$, where $\operatorname{gr} 1 = -1/2$, $\operatorname{gr} \theta = 1/2$.

If Y is a rational homology sphere, it follows from Property 9 that $HF^+(Y, \mathfrak{s}) \cong M + - \oplus T$, where T is a torsion module. The $\mathbb{Z}[U]$ submodule M_+ is canonically defined by the relation

$$M_+ \bigcap_{k \in \mathbb{Z}} U^k HF^+(Y, \mathfrak{s}).$$

Definition 1.16. The d-invariant $d(Y, \mathfrak{s})$ is the minimal grading of an element of M_+.

Invariants of this type were first defined by Froyshov using Seiberg-Witten theory.

1.5 Exercises

In doing the exercises, you should take as given the properties of Floer homology stated in the lecture.

1. Let K be the unknot in S^3, and let $W = W_{-1}(K)$. Show that the map $F^-_{W,\mathfrak{s}_{\pm 1}} : HF^-(K) \to HF^-(K)$ is an isomorphism. What is the map $F^-_{W,\mathfrak{s}_{\pm k}}$ for other values of k?

2. (The blowup formula) Given a cobordism $W : Y_1 \to Y_2$, let $W' = W \oplus (-\mathbb{CP}^2)$, and let $E \in H_2(W')$ be the exceptional divisor. (That is, a sphere representing the class of a generator of $H_2(-\mathbb{CP}^2)$.) If $\mathfrak{s} \in Spin^c(W)$, show that for odd k there is a unique $\mathfrak{s}_k \in Spin^c(W')$ which agrees with $\mathfrak{s}$ on $W' - \nu(E)$, and which has $\langle c_1(\mathfrak{s}_k), [E] \rangle = k$. Use the first problem to show that $F^\circ_{W',\mathfrak{s}_{\pm 1}} = \pm F^\circ_{W,\mathfrak{s}}$. What happens if instead we take $W' = W \# \mathbb{CP}^2$?

3. If $S \subset Y$ is an embedded surface, let

$$HF^\circ(Y, S, k) = \bigoplus_{\{\mathfrak{s} \,|\, \langle c_1(\mathfrak{s}), [S] \rangle = k\}} HF^\circ(Y, \mathfrak{s}).$$

Deduce the adjunction property in the special case where $\Sigma \cdot \Sigma = 0$, $g(\Sigma) > 0$ from the fact that $HF^\circ(S^1 \times \Sigma, \Sigma, k) = 0$ for $k > 2g(\Sigma) - 2$. (We'll see how to prove this in the last lecture.)

4. Use the blow-up formula and exercise 3 to prove the adjunction property for all Σ with genus > 0. (Hint: blow up repeatedly to produce a surface Σ' with self-intersection 0.)

5. Use adjunction to show that if $S \subset Y$ is an embedded surface of genus > 0, then $HF^\circ(Y, S, k) = 0$ whenever $k > 2g(S) - 2$. Deduce that there are only finitely many $\mathfrak{s}$ for which $HF^\circ(Y, s) \neq 0$.

6. Let $W_1 : S^3 \to S^1 \times S^2$ and $W_2 : S^1 \times S^2 \to S^3$ be the cobordisms given by addition of a 1 handle and a cancelling 2-handle respectively. Use what you know about grading shifts, together with the fact that $W_2 \circ W_1$ is the identity cobordism, to determine the maps $\widehat{F}_{W_1}$ and $\widehat{F}_{W_2}$. Similarly for the cobordisms $W_2' : S^3 \to S^1 \times S^2$ and $W_3 : S^1 \times S^2 \to S^3$ given by addition of a 2-handle and a cancelling 3-handle.

7. Given that $c_1(\mathfrak{s}|_{\partial W})$ is torsion, explain how to make sense of the quantity $c_1(\mathfrak{s})^2$appearing in the formula for the degree shift. (Note that a priori $c_1^2(\mathfrak{s})$ is a class in $H^4(W) = 0$.) Compute the degree shift associated to the map $F^\circ_(W, \mathfrak{s}_k)$, where $W = W_{-p}(K)$. Verify that if $\mathfrak{s}_{k_1}$ and $\mathfrak{s}_{k_2}$ restict to the same $Spin^c$ structure on ∂W then the difference in the corresponding degree shifts is an even integer.

8. Identify the S^1-equivariant analog of the second exact sequence of Property 8 . What about the spectral sequence?

Lecture 2

Left-orderable groups

Definition 2.1. A non-trivial group G is *left-orderable* if there is a strict, total order $<$ on the elements of G that is invariant under multiplication on the left in the sense that

$$a < b \quad \Longrightarrow \quad ca < cb$$

for all $a, b, c \in G$.

Notice that the trivial group, having only a single element, quite plainly satisfies the condition above. However, our convention will be that the trivial group is *not* left-orderable. For the moment, perhaps this is best justified by the assertion that the notion of an order, in any reasonable sense, depends on having at least two objects to compare. In reality, this convention — one which is ultimately arbitrary — makes many of the statements of interest cleaner.

Among the first observations one makes about left-orderable groups is that they are necessarily infinite. Indeed:

Proposition 2.2. *Left-orderable groups are torsion free.*

Proof. Let g be any non-trivial element of finite order in a left-orderable group G and suppose that $1 < g$. Then $1 < g < \cdots < g^n$ by repeated self-multiplication on the left, hence $1 < g^n = 1$ if n is the order of g, which is impossible. A similar contradiction arises if $g < 1$. □

Of course, the class of non-left-orderable groups is much larger than finite groups or even torsion free groups. We will see lots of examples, but for the moment consider the following.

Exercise 2.3. [Calegari-Dunfield [CD03]] Show that the group

$$G = \langle a, b | bababa^{-1}b^2a^{-1}, ababab^{-1}a^2b^{-1} \rangle$$

is not-left-orderable: For a contradiction, suppose that G is left-orderable and assume, without loss of generality, that $1 < a$. Consider cases $1 < b$ (Hint: Compare a and b; the two sub-cases will prove that one of the two words in the relator is positive) and $b < 1$ (Hint: Use the relations to find an expression for 1 as a word in a and b^{-1}).

(Why is G torsion free? One way to see this is to notice that G is the fundamental group of the Weeks manifold [CD03]; this is now known to be the closed manifold with the distinction of having the smallest possible hyperbolic volume — roughly 0.942707362776 [GMM09]. The Weeks manifold being hyperbolic, its fundamental group is a torsion-free, discrete subgroup of $\mathrm{PSL}_2(\mathbb{C})$, the isometry group of hyperbolic three-space.)

It is quite reasonable to ask if there is any reason to prioritise left-multiplication over right multiplication. The answer is that there is not as illustrated below.

Exercise 2.4. Show that, given a left-order $<$ on a group G, that

$$a \mathrel{\tilde{<}} b \iff a^{-1} < b^{-1}$$

for all $a, b \in G$ defines a right-invariant, strict total order on the elements of G. Give the definition of a right-orderable group and prove that every left-order on G is equivalent to a right-order (and vice versa). In particular, left- and right-orderable groups are equivalent.

Examples

The following are some natural examples of left-orderable groups.

Example 2.5. The group $\mathbb{Z}$ is is left-orderable in the usual sense: We let $a < b$ whenever $b - a$ is a positive integer. Notice that there are precisely two such orders on this group (though the other one may not seem so natural!).

Since $\mathbb{Z}$ is abelian, the left-invariance of the order is actually a two-sided invariance. This order structure, called bi-orderability, is more general; in this case we are restricting a bi-order to a left-order. Note that not every left-order is a bi-order — a somewhat famous example, the braid group, is given below!

More generally, the group $\mathbb{Z}^n$ is left orderable: Selecting a hyperplane in $\mathbb{R}^n$ containing the origin (and no other points in $\mathbb{Z}^n$), together with a choice of distinguished half-space, $a < b$ if and only if $b - a$ is contained in the distinguished half-space. Notice that, for $n > 1$, there are now infinitely many left-orders obtained in this way.

As suggested by this example, there is an alternative means of specifying a left-order on a group by providing an exhaustive list of the positive elements in the left-order. This list is called a positive cone.

Definition 2.6. Given a group G, a *positive cone* $\mathcal{P}$ is a subset of elements of G satisfying:

(i) $\mathcal{P} \cdot \mathcal{P} \subseteq \mathcal{P}$;

(ii) $G = \mathcal{P} \sqcup \{1\} \sqcup \mathcal{P}^{-1}$; and

(iii) $\mathcal{P}$ is non-empty.

Condition (i) asserts that if $a, b \in \mathcal{P}$ then $ab \in \mathcal{P}$; in particular, $\mathcal{P}$ is a closed sub-semigroup of G. Condition (ii) asserts that given $a \in G$ exactly one of $a \in \mathcal{P}$ or $a^{-1} \in \mathcal{P}$ or $a = 1$ holds. Condition (iii) rules out the trivial group as a left-orderable group, consistent with our convention.

Exercise 2.7. Describe the positive cones implicit in Example 2.5 for any integer n (the case $n = 2$ is probably most instructive). What if the hyperplane contains some of the non-trivial group elements of $\mathbb{Z}^n$?

Exercise 2.8. Prove that the existence of a positive cone for G is equivalent to a left-order on G. Any positive cone $\mathcal{P}$ gives rise to an opposite order: Take $\mathcal{P}^{-1}$ as positive cone instead. Compare the opposite order with (the positive cone for) the right-order $\widetilde{<}$ of Exercise 2.4

Consistent with the literature, we will use these two definitions interchangeably. For instance, a statement like "the word w is a product of positive elements and, since $w = 1$, this is a contradiction" (compare Proposition 2.2 and/or Exercise 2.3) is common when proving that a given group is not-left-orderable; implicit here is the notion that a list of positive elements is available, whether by way of a positive cone $\mathcal{P}$ or by some other explicit description of $<$.

Example 2.9. Braid groups are left-orderable but not bi-orderable. This is not at all obvious, and is originally due to Dehornoy [Deh94] (that the left-orders cannot be promoted to bi-orders is proved by Rhemtulla and Rolfsen [RR02]). There are many other proofs and these are discussed at length in a book by Dehornoy, Dynikov, Rolfsen and Wiest [DDRW02]. While historically significant, this example is not central to these lectures; we leave the subject of braid orderability as recommended further reading for those interest.

Example 2.10. Let $\mathrm{Homeo}^+(\mathbb{R})$ denote the group of orientation (or, order) preserving homeomorphisms of $\mathbb{R}$ with the usual topology. The group operation is function composition and the identity is the identity homeomorphism $\mathrm{id}(x) = x$.

A family of left-orders on $\mathrm{Homeo}^+(\mathbb{R})$ depends on

- a countable dense subset $X \subset \mathbb{R}$ (for example, $\mathbb{Q} \subset \mathbb{R}$); and
- an enumeration, or counting, $X = \{x_0, x_1, x_2, \ldots\}$.

Then the positive cone

$$\mathcal{P}_X = \{f \in \mathrm{Homeo}^+(\mathbb{R}) | f(x_n) > x_n \text{ and } f(x_i) = x_i \text{ for all } i < n\}$$

defines a left-order on $\mathrm{Homeo}^+(\mathbb{R})$.

The only thing that needs a moment of thought is that this $\mathcal{P}_X$ is a closed sub-semigroup of $\mathrm{Homeo}^+(\mathbb{R})$ (the other two conditions are immediate). If $f_1, f_2 \in \mathcal{P}_X$ let n_1 and n_2 be the non-negative integers, required by $\mathcal{P}_X$, associated with f_1 and f_2, respectively. If $n_1 < n_2$ then $f_1(f_2(x_1)) > x_1$ and $f_1 \circ f_2 \in \mathcal{P}_X$, and so forth.

Note that the resulting order is extremely sensitive to both of the choices required in the construction. Indeed, specifying $X = \{\ldots, x_{i+n}, \ldots, x_i, \ldots\}$ instead of $X = \{\ldots, x_i, \ldots, x_{i+n}, \ldots\}$ should be expected to alter the left-order rather wildly.

Exercise 2.11. Write down a definition for $<$ on $\mathrm{Homeo}^+(\mathbb{R})$ as specified by $\mathcal{P}_X$ of Example 2.9.

Properties

We have already seen that left-orderable groups are torsion free, but that the converse does not hold. This is perhaps the zeroth property of left-orderable groups. A first property could be:

Proposition 2.12. *If H is a non-trivial subgroup of a left-orderable group G, then H is left-orderable as well.*

Proof. This is immediate: Given a positive cone $\mathcal{P} \subset G$ we obtain a positive cone $\mathcal{P}|_H \subset H$ by restriction. □

Proposition 2.13. *Let $G \to H$ be a surjective homomorphism with kernel K. If H and K are left-orderable groups then G is left-orderable as well.*

Proof. Let $\mathcal{P}_K$ and $\mathcal{P}_H$ be positive cones on K and H, respectively. Denoting the surjection by $f\colon G \to H$, consider the subset $\mathcal{P}_K \cup f^{-1}(\mathcal{P}_H) \subset G$. This is a positive cone (check this!), and hence G is left orderable. □

Theorem 2.14 (Vinogradov [Vin49])**.** *Let G_1 and G_2 be non-trivial groups. The free product $G_1 * G_2$ is left-orderable if and only if G_1 and G_2 are both left-orderable groups.*

To summarise, left-orderable groups as torsion free and closed (non-trivial) under subgroups, extensions and free-products.

Somewhat more subtle are quotients of left-orderable groups. Let G be a left-orderable group and suppose that $K \subset G$ is a normal subgroup. Of course, K is left-orderable; considering the short exact sequence

$$1 \longrightarrow K \longrightarrow G \longrightarrow H \longrightarrow 1$$

we'd like to know when H is a left-orderable group. This turns out to require an additional condition on K.

Definition 2.15. Let G be a group with left-order $<$. A non-trivial proper subgroup $C \subset G$ is called *convex* if, given $a, b \in C$ with $a < b$ then any $c \in G$ for which $a < c < b$ guarantees that $c \in C$.

Proposition 2.16. *Consider a short exact sequence of groups*

$$1 \longrightarrow K \longrightarrow G \longrightarrow H \longrightarrow 1$$

where G is a left-orderable group. Then H is left-orderable if and only if K is a convex subgroup (relative to some left-order on G).

Proof. First suppose that K is a convex subgroup of G relative to some left-order $<$ (with associated positive cone $\mathcal{P}$) which we will fix. We'll prove that the set $(G \smallsetminus K) \cap \mathcal{P}$ gives rise to a positive cone $\mathcal{Q}$ for $H = G/K$: Let $aK \in \mathcal{Q}$ if and only if $a \in \mathcal{P}$.

To see that this is well-defined, consider $a \in (G \smallsetminus K) \cap \mathcal{P}$ and cosets $aK = bK$ so that $b = ac$ for some $c \in K$. Notice that if $c < a^{-1} < 1$, convexity of K would show that $a^{-1} \in K$ contradicting the fact that $a \notin K$. It must be that $a^{-1} < c$ and hence $1 = aa^{-1} < ac = b$ so $bC \in \mathcal{Q}$.

To see that $\mathcal{Q} \subset G/K$ is a closed sub-semigroup, consider $aK, bK \in \mathcal{Q}$. Then $a, b \in \mathcal{P}$ hence $ab \in \mathcal{P}$ also, thus $(ab)K \in \mathcal{Q}$ provided $ab \notin K$. However, were it the case that $K = (ab)K$ then $bK = (aK)^{-1} = a^{-1}K$ implying that b is negative, a contradiction.

To see that $G/K = \mathcal{Q} \sqcup \{1\} \sqcup \mathcal{Q}^{-1}$, note that K represents the identity in G/K and any coset C (other than K) either contains a positive element or else C^{-1} does.

Finally, to see that $\mathcal{Q}$ is non-empty, consult the definition of a convex subgroup, in particular, K is a proper subgroup of G so there must be a coset C different from K.

Conversely, suppose that H is a left-orderable quotient of a left-orderable group G. Let ϕ denote the projection homomorphism, and let $\mathcal{P}$, $\mathcal{P}_H$, and $\mathcal{P}|_K$ be positive cones for G, H, and K respectively.

Consider the following positive cone $\mathcal{P}_G \subset G$: An element $a \in \mathcal{P}_G$ if and only if $\phi(a) \neq 1$ and $\phi(a) \in \mathcal{P}_H$ or $\phi(a) = 1$ and $a \in \mathcal{P}|_K$. We claim that K is now a convex subgroup of G, relative to this (possibly) new left-order on G.

Indeed, consider $a < c < b$ where $a, b \in K$ and $c \in G$. Then $c^{-1}a < 1 < c^{-1}b$ (using the positive cone $\mathcal{P}_G$), and if $c \notin K$ then neither $c^{-1}a$ nor $c^{-1}b$ is an element of K either. Hence $c^{-1}b \in \mathcal{P}$ so $\phi(c^{-1}b) = \phi(c^{-1}) = \phi(c)^{-1}$ is positive while $c^{-1}a \in \mathcal{P}^{-1}$ hence $\phi(c^{-1}a) = \phi(c^{-1}) = \phi(c)^{-1}$ is (simultaneously) negative. This is a contradiction, hence $c \in K$ and K is convex in G. □

Characterisations

Theorem 2.17 (Conrad)**.** *A group G is left-orderable if and only if, given a finite subset $\{a_1, \ldots, a_n\} \subset G \setminus \{1\}$, there exist signs $\epsilon_i = \pm 1$ such that the semigroup generated by $\{a_1^{\epsilon_1}, \ldots, a_n^{\epsilon_n}\}$ does not contain the identity.*

This might be thought of as measuring the failure to obtain a contradiction when seeking to establish that a group is not left-orderable (as in Exercise 2.3, for example). The utility of this characterization is made clear in the proof of the following.

Theorem 2.18 (The Burns-Hale Criterion [BH72])**.** *A group G is left-orderable if and only if every non-trivial finitely generated subgroup of G surjects onto a left-orderable group.*

Proof. The proof is by contradiction: Suppose every non-trivial finitely generated subgroup of G surjects onto a left-orderable group, but that G is *not* left-orderable. Then by Conrad's theorem, we can find elements $a_1, \ldots, a_n \in G \setminus \{1\}$ such that, for all possible sigh choices $\epsilon_i = \pm 1$, the subsemigroup generated by $\{a_1^{\epsilon_1}, \ldots, a_n^{\epsilon_n}\}$ contains the identity.

Choose the smallest possible integer n for which this holds, and consider the subgroup $H \subset G$ with finite generating set $\{a_1, \ldots, a_n\}$. While H is not left-orderable (applying Conrad's theorem), H surjects onto a left-orderable group H/K by assumption; the kernel of the surjection K cannot contain all of the a_i since H/K is non-trivial. On the other hand, K contains *some* a_i since H/K is left-orderable but H is not.

Let $a_1, \ldots, a_r \notin K$ and $a_{r+1}, \ldots, a_n \in K$ for $0 < r < n$. Choose signs $\epsilon_i = \pm 1$ so that the subsemigroup generated by $\{a_1^{\epsilon_1}, \ldots, a_r^{\epsilon_n}\}$, denoted $(a_1^{\epsilon_1}, \ldots, a_r^{\epsilon_n})$, does not contain the identity; since $r < n$ this is possible by the minimality of n. As a result $(a_1^{\epsilon_1}, \ldots, a_r^{\epsilon_n}) \cap K = \emptyset$. Again by minimality of n, $1 \notin (a_{r+1}^{\epsilon_{r+1}}, \dots, a_n^{\epsilon_n})$ for some choices of signs. However $1 \in (a_1^{\epsilon_1}, \dots, a_n^{\epsilon_n})$, meaning we can express the identity as a word in the set $\{a_1^{\epsilon_1}, \dots, a_n^{\epsilon_n}\}$. This word must contain at least one occurrence of an $a_i^{\epsilon_i}$ for $i \le r$. This contradicts the fact that $(a_1^{\epsilon_1}, \dots, a_r^{\epsilon_n}) \cap K = \emptyset$. □

This is a fine characterisation, but it is clearly not particularly practical for establishing left-orderability for a given group. Burns and Hale referred to their (necessary and sufficient) condition for left-orderability *order indicability*, which is sensible once you recall that a *locally indicable* group G satisfies the property that every non-trivial, finitely generated subgroup surjects onto $\mathbb{Z}$. As a result, the Burns-Hale theorem proves that locally indicable groups are left-orderable; the Burns-Hale criterion weakens the requirement from $\mathbb{Z}$ (our first example of a left-orderable group) to *any* left-orderable group.

Example 2.19. This gives rise to a large family of left-orderable groups: Let K be a knot in S^3 and let $M = S^3 \smallsetminus \nu(K)$ denote the knot exterior. Then the knot group $\pi_1(K) = \pi_1(M)$ is a locally indicable group by a result of Howie and Short [HS85]; it follows that all knot groups are left-orderable.

Theorem 2.20 (see Linnell [Lin99] or Boyer-Rolfsen-Wiest [BRW05])**.** *If G is a countable group, then G is left-orderable if and only if G is a subgroup of* $\mathrm{Homeo}^+(\mathbb{R})$.

Proof. First consider the countable subgroup $\mathbb{Q} \subset \mathbb{R}$, and set $X = \mathbb{Q} \times G$; X may be ordered lexicographically by setting $(p, a) < (q, b)$ if and only if $a < b$ (in G) or $a = b$ and $p < q$ (in $\mathbb{Q}$). Notice that by viewing $p \in \mathbb{Q}$ as an equivalence class of pairs (n_p, m_p), $\mathbb{Q}$ may be ordered in precisely the same way. In fact, if $p = \frac{n_p}{m_p}$ then this recovers the standard order on $\mathbb{Q}$ inherited (by restriction) from $\mathbb{R}$. We leave as an exercise the following claim: X and $\mathbb{Q}$ are isomorphic as (countable) ordered sets. Denote this order preserving bijection by $\phi\colon X \to \mathbb{Q}$.

Now G acts on X by setting $a(p, b) = (p, ab)$; we obtain an action of G on $\mathbb{Q}$ by $ap = \phi(a(\phi^{-1}(p)))$. So to every $a \in G$ we may assign the function $f_a \in \mathrm{Homeo}^+(\mathbb{Q})$ associated with multiplication on the left by a.

Now, by continuity, the inclusion of $\mathbb{Q}$ into $\mathbb{R}$ gives rise to an extension of $f_a \in \mathrm{Homeo}^+(\mathbb{Q})$ to $f_a \in \mathrm{Homeo}^+(\mathbb{R})$. Hence G is realised as a subgroup of $\mathrm{Homeo}^+(\mathbb{R})$, as claimed. □

This has the pleasing consequence that countable left-orderable groups are familiar objects: What could be simpler than a continuous monotone function $f\colon \mathbb{R} \to \mathbb{R}$? However, as we have seen in Example 2.9, the set of left-orders on $\mathrm{Homeo}^+(\mathbb{R})$ seems to be rather complicated.

Finally, we advertise a great result of Farrell which, to the best of our knowledge, has yet to be applied to a topological problem.

Theorem 2.21 (Farrell [Far76])**.** *Suppose that X is a locally-compact, paracompact topological space (for example, let X be a manifold), and let $\widetilde{X}$ be the universal covering of X. Then $\pi_1(X)$ is left-orderable if and only if there is a embedding of topological spaces $\widetilde{X} \to X \times \mathbb{R}$ and a commutative diagram*

$$\begin{array}{ccc} \widetilde{X} & \longrightarrow & X \times \mathbb{R} \\ & \searrow \quad \swarrow & \\ & X & \end{array}$$

where $X \times \mathbb{R} \to X$ is projection to the first factor.

Three-manifold groups

Farrell's result hints that there may be more to left-orderability when restricting to a topological setting. We will be interested in three-manifold groups, that is, groups arising as the fundamental group of a three-manifold. Note that these are always countable groups, so by Theorem 2.20 left-orderable three-manifold groups are subgroups of $\mathrm{Homeo}^+(\mathbb{R})$. Continuing with previous lectures, all three-manifolds will be assumed to be connected orientable and, unless noted otherwise explicitly, compact. (Hatcher's notes are a good reference [Hat].)

This hint is born out by a result of Boyer, Rolfsen and Wiest:

Theorem 2.22 (Boyer-Rolfsen-Wiest [BRW05])**.** *Let M be an compact, connected, orientable, irreducible three-manifold. Then $\pi_1(M)$ is left-orderable if and only if $\pi_1(M)$ surjects onto a left-orderable group.*

Proof. One direction is immediate: If $\pi_1(M)$ is left-orderable then $\pi_i(M)$ surjects onto a left-orderable group, namely, itself.

Suppose then that $\pi_1(M)$ surjects onto a left-orderable group G. By the criterion of Burns and Hale it is enough to show that every non-trivial finitely generated subgroup of $\pi_1(M)$ surjects onto a left-orderable group. Let H be such a subgroup; H is either finite or infinite index in $\pi_1(M)$. Then if H has finite index in $\pi_1(M)$ it's image in G has finite index also. This image being non-trivial, it suffices to restrict $\pi_1(M) \to G$ to $H \subset \pi_1(M)$ to obtain the desired surjection from H to a left-orderable group.

So it remains to consider the case where H is an infinite index, finitely generated, subgroup of $\pi_1(M)$. Let $p\colon \widetilde{M} \to M$ be the corresponding covering space so that $p_*(\pi_1(\widetilde{M})) = H$. Note that $\widetilde{M}$ is a non-compact manifold. At this stage we appeal to a major piece of machinery from three-manifold topology, namely, the Compact Core Theorem due to Scott [Sco73]: There exists a compact submanifold $i\colon M^c \hookrightarrow \widetilde{M}$ with the property that the induced homomorphism i_* sends $\pi_1(M^c)$ isomorphically to $\pi_1(\widetilde{M})$. Since M^c is compact, it must have non-empty boundary (otherwise M^c and $\widetilde{M}$ would coincide).

Consider the case where some component of the boundary ∂M^c is a two-sphere $S \subset \partial M^c$. Since M is irreducible, so is $\widetilde{M}$ (this is not an obvious statement, see [Hat, Theorem 3.15]); hence S bounds a three-ball $B \subset \widetilde{M}$. Suppose that $M^c \subset B \subset \widetilde{M}$. It follows that i factors through B and $\pi_1(M^c)$ is not isomorphic to $\pi_1(\widetilde{M})$. This is a contradiction, so it must be that $B \cap M^c = S \subset \partial M^c$ and we may attach B to M^c along S without compromising the isomorphism between $\pi_1(M^c)$ and $\pi_1(\widetilde{M})$. Without loss of generality then, we may assume that ∂M^c has no two-sphere components.

We now claim that M^c, a compact orientable manifold with non-empty boundary containing no two-sphere components, has positive first betti number. Since each boundary component has negative Euler characteristic we have

$$0 \geq \tfrac{1}{2}\chi(\partial M^c) = \chi(M^c) = 1 - b_1(M^c) + b_2(M^c),$$

where the last equality uses the fact that $H_3(M^c) = 0$ while the first equality appeals to the double of $D(M^c)$ of M^c which, as a closed manifold, must have $0 = \chi(D(M^c)) = 2\chi(M^c) - \chi(\partial M^c)$. As such we conclude that $b_1(M^c) \geq 1 + b_2(M^c) \geq 1$.

Thus $H_1(M^c;\mathbb{Z})$ admits a surjection to $\mathbb{Z}$ by restriction to a free summad, so the composition $\pi_1(M^c) \to H_1(M^c;\mathbb{Z}) \to \mathbb{Z}$ yields a surjection $H \to \mathbb{Z}$, as required. □

This result immediately gives rise to two very large families of left-orderable three-manifold groups:

Corollary 2.23. *If M is a compact, irreducible, orientable manifold with non-empty (non-spherical) boundary, then $\pi_1(M)$ is left-orderable. In particular, if M is the exterior of a knot in the three-sphere, then $\pi_1(M)$ is left-orderable, recovering Howie and Short's result.*

Corollary 2.24. *If Y is an closed, irreducible, orientable manifold with $H_1(Y;\mathbb{Q}) \neq 0$ (i.e. Y has positive first betti number) then $\pi_1(Y)$ is left-orderable.*

From this last observation, it follows immediately that the class of rational homology spheres is of particular interest.

Further reading

The material surveyed above is decidedly idiosyncratic and far from being a complete account of left-orderable groups in low-dimensional topology. As such it is farther still from being a complete account of left-orderable group theory, a subject of study in its own right. Left-orderable groups show up in interesting and surprising places: For example, they make an appearance in Mineyev's recent resolution of the Hanna Neuman Conjecture [Min12].

The notion of endowing a group with a left-order, viewed as an auxiliary structure on a group, seems natural enough as a weakening of the notion of local indicability (as discussed in relation to the Burns-Hale criterion). A nice reference for orderability in this context is the book by Mura and Rhemtulla [BMR77]. More recently, an observation of Sikora [Sik04] has added another topological flavour to left-orderable groups by considering the set of all left-orders on a given group as a topological space.

Exercise* 2.25. Fix a left orderable group G and denote by $LO(G)$ the set of all left orders on G. This becomes a topological space with the subbasis $U_a^b = \{<: a < b\}$. To study this topology further, consider a positive cone $\mathcal{P} \in LO(G)$ (compare Exercise 2.8) and prove that there is an inclusion of sets $LO(G)$ into $2^G = \{S : S \subseteq G\}$, the *power set* of G (considered as a set). A subbasis for a topology on 2^G is given by sets

$$U_a = \{S \subset G : a \in S\} \quad \text{and} \quad U'_a = \{S \subset G : a \notin S\}.$$

You can prove that the induced relative topology on $LO(G)$ is equivalent to the topology on $LO(G)$ described in terms of the subasis $\{U_a^b\}$. Using the fact that 2^G is a compact topological space (a consequence of Tychonoff's theorem), prove that $LO(G)$ is a compact topological space. Finally, prove that $LO(G)$ is a totally disconnected topological, that is, that the only connected components are singletons.

As such, the distinction between orders (on a given group) that are isolated versus non-isolated is an interesting one that has received attention. For example, the Dehornoy ordering of the braid groups is known to be non-isolated [Nav10]. And, in the case where $LO(G)$ contains no isolated points, it is necessarily homeomorphic to a Cantor set. However, the braid groups admit left-orders that correspond to isolated points in $LO(B_n)$ [Cla10].

The fact that the braid groups are left-orderable was originally proved using techniques from logic [Deh94]. As a result, this property was not immediately accessible to topologists. However, interested in the problem of zero divisors in a group ring — it turns out that if G is left-orderable then $\mathbb{Z}G$ has no zero divisors; see Passman [Pas77] — Rolfsen was put onto the work of Dehornoy by Birman; he later encouraged the low-dimensional topology community to understand the result. New proofs followed, and this was something of a watershed moment for low-dimensional topology and left-orderable groups. For example, it was subsequently shown that many more mapping class groups — braids being mapping classes of an n-punctured disk — are left-orderable as well. This culminates with, and is collected in, the work *Why are braids orderable?* by Dehornoy, Dynikov, Rolfsen and Wiest [DDRW02].

The first comprehensive discussion of left-orderable groups in the context of three-manifold topology is the work of Boyer, Rolfsen and Wiest [BRW05] however important work concurrent with this is that of Calegari and Dunfield [CD03] and Roberts, Shareshian and Stein [RSS03]. These latter to works are principally concerned with foliations, and we will have more to say on this in a future lecture. While it seems unreasonable to reproduce the entire work of Boyer, Rolfsen and Wiest here, it should be noted that this paper is an essential reference in the subject and should be read by anyone seriously considering delving deeper into left-orderability in the context of three-manifold groups. Note that Boyer, Rolfsen and Wiest do not restrict to orientable three-manifolds as we have done here.

Lecture 3

The basic structure of Heegaard Floer homology

In this lecture we'll sketch the construction of an invariant satisfying the properties stated in the first lecture — Heegaard Floer homology. Heegaard Floer is a bit like simplicial homology, in the sense that in order to define it, we first choose a specific combinatorial model of our space (in this case, a handle decomposition), and then show that the homology does not depend on which model we choose.

3.1 Heegaard diagrams

Definition 3.1. Let Σ be a closed oriented surface. A *set of attaching circles* on Σ is a set of disjoint simple closed curves on Σ.

Given two sets of attaching circles $\alpha = \{\alpha_1, \ldots, \alpha_j\}$ and $\beta = \{\beta_1, \ldots, \beta_k\}$ on Σ, we can form an oriented 3-manifold $Y_{\alpha\beta}$ by starting with $\Sigma \times I$ (with its standard orientation) and attaching two handles along the circles $\alpha_i \times 0$ and $\beta_i \times 1$. If any components of the boundary of the resulting manifold are spheres, we fill them in with D^3. The resulting manifold is $Y_{\alpha\beta}$.

Definition 3.2. A *generalized Heegaard diagram* for $Y_{\alpha\beta}$ is a quadruple $(\Sigma, \alpha, \beta, Z)$ where $Z = \{z_1, \ldots z_l\}$ is a set of points in $\Sigma - \alpha - \beta$. The diagram is *balanced* if the number of circles in α and β are the same and each component of $\Sigma - \alpha$ and $\Sigma - \beta$ contains exactly one z_i.

If $|\alpha| = |\beta| = g(\Sigma)$, the generalized Heegaard diagram is a Heegaard diagram in the usual sense. In this case, being balanced means there is exactly one basepoint.

Elementary Morse theory shows that any compact orientable 3-manifold can be represented by a generalized Heegaard diagram, and that any closed 3-manifold can be represented by a Heegaard diagram.

Example 3.3. Some Heegaard diagrams of S^3. Standard diagram of $L(p, q)$. Complement of the trefoil.

3.2 Lagrangian Floer homology

Given a 3-manifold Y, choose a Heegaard diagram $(\Sigma, \alpha, \beta, z)$ representing Y. We'll define a chain complex $\widehat{CF}(\Sigma, \alpha, \beta, z)$ whose homology is $\widehat{HF}(Y)$. In brief, this complex is the Lagrangian Floer chain complex $CF(\mathbb{T}_\alpha, \mathbb{T}_\beta)$ associated to a pair of Lagrangians in $\mathrm{Sym}^g(\Sigma - z)$. (In the case when $b_1(Y) > 0$, we'll have to impose another condition, called *admissability* on our Heegaard diagram.)

In order to make sense of this description, we give a brief sketch of Lagrangian Floer homology. (Please be aware that in doing so I've collapsed several illustrious careers worth of work in analysis into a few sentences.) Let (M, ω) be a symplectic manifold, and choose a *compatible almost complex structure* $J : TM \to TM$. By definition, this means that the two-form $g(v, w) := \omega(v, Jw)$ defines a Riemannian metric on M.

Suppose that L_+, L_- are two transversely intersecting Lagrangians M (i.e. $\omega|_{L_\pm} = 0$), and fix two points $\mathbf{x}, \mathbf{y} \in L_+ \cap L_-$.

Definition 3.4. A *compatible map* $\varphi : D^2 \to M$ is a continuous map for which $\varphi(-i) = \mathbf{x}, \varphi(i) = \mathbf{y}$, and $\varphi(\partial D^2_\pm) \subset L_\pm$, where

$$\partial D^2_+ = \{(x, y) \in S^1 \,|\, x \geq 0\}$$

and similarly for ∂D^2_-. We let $\pi_2(\mathbf{x}, \mathbf{y})$ be the set of homotopy classes of compatible maps.

If $\phi \in \pi_2(\mathbf{x},\mathbf{y})$ and $\psi \in \pi_2(\mathbf{y},\mathbf{z})$, there is a natural class $\phi\#\psi \in \pi_2(\mathbf{x},\mathbf{z})$, as well as a class $-\phi \in \pi_2(\mathbf{y},\mathbf{x})$. Thus there is an equivalence relation $\sim$ on $L_1 \cap L_2$ defined by $\mathbf{x} \sim \mathbf{y}$ if $\pi_2(\mathbf{x},\mathbf{y}) \neq \emptyset$.

A compatible map φ is *J-holomorphic* if it intertwines the standard complex structure on the disk with the almost-complex structure J. In other words, $d\varphi(iv) = J(d\varphi(v))$ for all $v \in T_z D^2$, $z \in \text{int } D^2$.

Definition 3.5. Given $\phi \in \pi_2(\mathbf{x},\mathbf{y})$, let $\mathcal{M}(\phi)$ be the set of all compatible J-holomorphic maps $\varphi : D^2 \to M$ which are in the homotopy class φ.

Theorem 3.6. *For a generic choice of J, the set $\mathcal{M}(\phi)$ is a manifold. Its dimension $\mu(\phi)$ is called the* Maslov index *of ϕ.*

The mod 2 Maslov index is determined by topological data: if $\mathbf{x}$ and $\mathbf{y}$ have the same intersection sign, $\mu(\phi)$ is even, and otherwise it is odd. The index is additive in the following sense. If $\phi \in \pi_2(\mathbf{x},\mathbf{y})$ and $\psi \in \pi_2(\mathbf{y},\mathbf{z})$, there is a natural class $\phi\#\psi \in \pi_2(\mathbf{x},\mathbf{z})$ and $\mu(\phi\#\psi) = \mu(\phi) + \mu(\psi)$.

There is a 1-parameter group of holomorphic maps $g_t : D^2 \to D^2$ which fix $\pm i$. We can use this family to define an action of $\mathbb{R}$ on $\mathcal{M}(\phi)$ via $t \cdot \varphi = \varphi \circ g_t$. This action is free unless $\mathbf{x} = \mathbf{y}$ and φ is the constant map. We denote the quotient $\mathcal{M}(\phi)/\mathbb{R}$ by $\overline{\mathcal{M}}(\phi)$.

After fixing some additional choices (which we'll leave unspecified), we get an orientation on $\mathcal{M}(\phi)$. When $\mu(\phi) = 1$, the quotient $\overline{\mathcal{M}}(\phi)$ is an oriented, 0-dimensional manifold, so it consists of a set set of signed points. We let $\#\overline{\mathcal{M}}(\phi)$ be the signed number of points in $\overline{\mathcal{M}}(\phi)$.

Definition 3.7. The Lagrangian Floer chain complex $CF(L_+, L_-)$ is the chain complex generated over $\mathbb{Z}$ by $L_+ \cap L_-$, and with differential given by

$$d\mathbf{x} = \sum_{\mathbf{y} \in L_+ \cap L_-} \sum_{\substack{\phi \in \pi_2(\mathbf{x},\mathbf{y}) \\ \mu(\phi)=1}} \#\overline{\mathcal{M}}(\phi)\mathbf{y}$$

It is immediate from the definition that $CF(L_+, L_-)$ decomposes as a direct sum

$$CF(L_+, L_-) = \bigoplus_s CF(L_+, L_-, s) \tag{3.1}$$

where s runs over the set of equivalence classes of generators, and $CF(L_+, L_-, s)$ is generated by those $\mathbf{x} \in s$.

To ensure that the sum is finite and we have $d^2 = 0$, we must impose additional conditions on M and $L_\pm$. (The relevant conditions for Heegaard Floer homology will be discussed below.) The proof that $d^2 = 0$ proceeds by studying the ends of moduli spaces with $\mu(\phi) = 2$. Specifically, the coefficient of $\mathbf{z}$ in $d^2\mathbf{x}$ is

$$\sum_{\substack{\phi\in\pi_2(\mathbf{x},\mathbf{y})\\ \mu(\phi)=1}} \sum_{\substack{\psi\in\pi_2(\mathbf{y},\mathbf{z})\\ \mu(\psi)=1}} \#\overline{\mathcal{M}}(\phi)\#\overline{\mathcal{M}}(\psi) = \sum_{\substack{\chi\in\pi_2(\mathbf{x},\mathbf{z})\\ \mu(\chi)=2}} \sum_{\phi\#\psi=\chi} \#\overline{\mathcal{M}}(\phi)\#\overline{\mathcal{M}}(\psi)$$

The inner sum on the RHS is the same as the signed number of ends in the noncompact 1-dimensional moduli space $\overline{\mathcal{M}}(\chi)$.

Gradings: Suppose that for all $\mathbf{x}, \mathbf{y} \in L_+ \cap L_-$ $\pi_2(\mathbf{x}, \mathbf{y})$ contains a unique element $\phi_{\mathbf{xy}}$. Then the Maslov index defines a relative $\mathbb{Z}$ grading on $CF(L_+, L_-)$ by $\operatorname{gr} x - \operatorname{gr} y = \mu(\phi_{\mathbf{xy}})$. This is well defined, since

$$\begin{aligned}(\operatorname{gr}\mathbf{x} - \operatorname{gr}\mathbf{y}) + (\operatorname{gr}\mathbf{y} - \operatorname{gr}\mathbf{z}) &= \mu(\phi_{\mathbf{xy}}) + \mu(\phi_{\mathbf{yz}})\\ &= \mu(\phi_{\mathbf{xy}}\#\phi_{\mathbf{yz}}) = \mu(\phi_{\mathbf{xz}}) = \operatorname{gr}\mathbf{x} - \operatorname{gr}\mathbf{z}.\end{aligned}$$

In general, the best we can expect to get is a $\mathbb{Z}/a$ grading, where

$$a = \gcd\{\mu(\phi) \,|\, \pi \in \pi_2(\mathbf{x}, \mathbf{x}).$$

3.3 Definition of $\widehat{HF}$

We now return to the definition of $\widehat{HF}$. Let $(\Sigma, \alpha, \beta, z)$ be a Heegaard diagram of Y, and let g be the genus of Σ. The symplectic manifold we will use to define $\widehat{CF}(\Sigma, \alpha, \beta, z)$ is $\mathrm{Sym}^g(\Sigma - z)$.

The symmetric product $\mathrm{Sym}^g\Sigma := \Sigma^g/S_g$, where the symmetric group S_g acts by permuting the factors. A point of $\mathrm{Sym}^g(\Sigma - z)$ is an unordered g-tuple of points of Σ. $\mathrm{Sym}^g(\Sigma - z)$ is an open subset of $\mathrm{Sym}^g\Sigma$; its complement is the divisor in $\mathrm{Sym}^g\Sigma$ consisting of unordered g-tuples of points, at least one of which is z.

To define the Lagrangian $\mathbb{T}_\alpha$, note that $\alpha_1 \times \ldots \times \alpha_g$ is an embedded torus in Σ^g. Since all the α_i's are disjoint, the projection maps this torus injectively to $\mathrm{Sym}^g(\Sigma)$. Its image is $\mathbb{T}_\alpha$. (It's not completely obvious how to choose the symplectic structure on $\mathrm{Sym}^g(\Sigma - z)$ so that $\mathbb{T}_\alpha$ is Lagrangian; see Perutz/Lekili for an explanation.)

The dimension of $\mathrm{Sym}^g(\Sigma)$ can be quite high, but it is possible to describe the generators of $\widehat{CF}(\Sigma, \alpha, \beta, z)$ and compatible disks between them quite explicitly in terms of Σ. We explain how to do this now.

Generators: The complex $\widehat{CF}(\Sigma, \alpha, \beta, z)$ is generated by the points of $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$. An intersection point of this form is an unordered g-tuple of points $\{x_1, \ldots, x_k\}$, where each $x_i \in \alpha_j \cap \beta_k$ for some j, k, and each α_j and β_k contain precisely one x_i.

Compatible Maps: Given $\mathbf{x}, \mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ and $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$, choose a representative $\varphi : D^2 \to \mathrm{Sym}^g(\Sigma)$ of ϕ, and let

$$S_\varphi := \{(z, w) \in D^2 \times \Sigma \,|\, w \in \varphi(z)\}.$$

There are obvious projections $\pi : S_\varphi \to D^2$ and $p : S_\varphi \to \Sigma$. For most φ (those which are transverse to the big diagonal in $\mathrm{Sym}^g\Sigma)$), the map π is a degree g branched covering map. The map π takes ∂S to $\alpha \cup \beta$.

Conversely, given a surface S and maps $\pi : S \to D^2$ and $p : S \to \Sigma$ satisfying the above conditions, there is a map $\varphi_S : D^2 \to \mathrm{Sym}^g\Sigma$ given by $\varphi(z) = p(\pi^{-1}(z))$, where $\pi^{-1}(z)$ is to be interpreted as a multiset.

Definition 3.8. Let $[S_\varphi]$ be the generator of $H_2(S_\varphi, \partial S_\varphi)$. The class $p_*([S_\varphi]) \in H_2(\Sigma, \alpha \cup \beta)$ depends only on the homotopy class ϕ of φ. It is called the *domain* of ϕ and denoted by $\mathcal{D}(\phi)$. Usually we represent $\mathcal{D}(\phi)$ by writing numbers (multiplicities) in the components of $\Sigma - (\alpha \cup \beta)$.

Let $x_{\alpha,j}$ be the point of $\mathbf{x}$ which lies on α_j. Suppose we are given a domain $\mathcal{D}$ whose boundary $\partial\mathcal{D} \in H_1(\alpha \cup \beta)$ can be represented by a chain of the form $\sum[\gamma_{\alpha_j}] - \sum[\gamma_{\beta_j}]$, where γ_{α_j} is a path on α_j from $x_{\alpha,j}$ to y_{α_j}, and γ_{β_j} is a similar path on β. (Let's call this condition $*$.) Then we can construct an appropriate S, π and p as above. We have thus established a correspondence between elements of $\pi_2(\mathbf{x}, \mathbf{y})$ and domains which satisfy condition $*$.

To compute $\widehat{HF}$, we should use only compatible holomorphic disks whose images are contained in $\mathrm{Sym}^g(\Sigma - z)$. This condition is also easy to describe. If we use a J sufficiently near to an honest complex structure on $\mathrm{Sym}^g\Sigma$, then any intersections of the holomorphic disk $\varphi(D^2)$ with the divisor which is the complement of $\mathrm{Sym}^g(\Sigma - z)$ will be positive. There will be no intersections if and only if the algebraic intersection number of the disk with the divisor is 0. It is easy to see that this intersection number is just the multiplicity of the domain $\mathcal{D}(\phi)$ in the region containing z, which we denote by $n_z(\phi)$.

3.4 *Spin^c* structures

We now investigate how the set of generators $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$ partitions into equivalence classes under the relation $\sim$. Given generators $\mathbf{x}, \mathbf{y}$, we try to find a domain $\mathcal{D}$ relating them. To do so, we first pick any paths $\gamma_{\alpha,j}$ from $x_{\alpha,j}$ to y_{α_j} and $\gamma_{\beta,j}$ from $x_{\beta,j}$ to y_{β_j}. We can find a domain $\mathcal{D}$ with

$$\partial\mathcal{D} = \sum[\gamma_{\alpha_j}] - \sum[\gamma_{\beta_j}]$$

if and only if the homology class on the right-hand side is 0 in $H_2(\Sigma)$. If this is not the case, we may modify our original choice of $\gamma_{\alpha,j}$ by adding copies of α_j, and similarly for $\gamma_{\beta,j}$. The net result is that the obstruction to $\mathbf{x}$ and $\mathbf{y}$ being related by a domain is

$$\epsilon(\mathbf{x}, \mathbf{y}) = \sum[\gamma_{\alpha_j}] - \sum[\gamma_{\beta_j}] \in \frac{H_1(\Sigma)}{\langle\alpha, \beta\rangle} = H_1(Y).$$

It is not hard to see (exercise) that $\epsilon(\mathbf{x}, \mathbf{y}) + \epsilon(\mathbf{y}, \mathbf{z}) = \epsilon(\mathbf{x}, \mathbf{z})$, and thus that the set of equivalence classes may be identified with a subset of an affine copy of $H_1(Y) \simeq H^2(Y)$. It's tempting (but incorrect) to identify equivalence classes with *Spin^c* structures on Y. The correct assertion is that a generator $\mathbf{x}$ and a basepoint z together determine a *Spin^c* structure $\mathfrak{s}_z(\mathbf{x})$. If $\mathbf{x} \sim \mathbf{y}$, then $\mathfrak{s}_z(\mathbf{x}) = \mathfrak{s}_z(\mathbf{y})$, and more generally $\mathfrak{s}_z(\mathbf{x}) - \mathfrak{s}_z(\mathbf{y}) = \epsilon(\mathbf{x}, \mathbf{y})$. The decomposition (3.1) thus gives the desired decomposition

$$\widehat{CF}(Y) = \bigoplus \widehat{CF}(Y, \mathfrak{s})$$

.

3.5 Manifolds with $b_1 > 0$

Next we describe the structure of $\pi_2(\mathbf{x}, \mathbf{y})$. The # operation makes $\pi_2(\mathbf{x}, \mathbf{x})$ into an abelian group, which acts freely and transitively (again by #) on $\pi_2(\mathbf{x}, \mathbf{y})$. Thus it suffices to determine the structure of the group $\pi_2(\mathbf{x}, \mathbf{x})$.

Lemma 3.9. $\pi_2(\mathbf{x}, \mathbf{x}) \simeq \mathbb{Z} \oplus H_2(Y)$.

Proof. There is a surjective homomorphism $\rho : \pi_1(\mathbf{x}, \mathbf{x}) \to \mathbb{Z}$ given by $\phi \mapsto n_z(\phi)$. An class $\phi \in \ker \rho$ is determined by $\partial\phi$, which is a linear combination of α's and β's. Such a linear combination is the boundary of a domain if and only if it is null-homologous in Σ, so $\ker \phi = \ker \iota \simeq H_2(Y)$, where $\iota : H_1(\alpha) \oplus H_1(\beta) \to H_1(\Sigma)$. □

If $\phi \in \ker \rho$, its domain $\mathcal{D}(\phi)$ is called a *periodic domain.* In order to define $\widehat{HF}(Y)$ for manifolds with $b_1 > 0$, we must use an *admissible* Heegaard diagram; that is one in which every periodic domain has both positive and negative coefficients.

Example 3.10. The standard Heegaard diagram for $S^1 \times S^2$ has no generators and is not admissible. An admissible diagram has two generators a, b and $da = b - b = 0$, $db = 0$.

To define homology with twisted coefficients, we fix one generator $\mathbf{x}$, and for each other generator $\mathbf{y} \sim \mathbf{x}$, we choose a class $\phi_{\mathbf{y}} \in \pi_2(\mathbf{x}, \mathbf{y})$ with $n_z(\phi_y)$. Then we can identify $\pi_2(\mathbf{y}, \mathbf{z})$ with $\pi_2(\mathbf{x}, \mathbf{x})$ via $\psi \mapsto \phi_{\mathbf{y}} + \psi - \phi_{\mathbf{z}}$. Given $\psi \in \pi_2(\mathbf{y}, \mathbf{z})$ with $n_z(\phi) = 0$, we get a class $t(\psi) \in H_2(Y)$ which is the image of $\phi_{\mathbf{y}} + \psi - \phi_{\mathbf{z}}$ in $H_2(Y)$.

Definition 3.11. The chain complex $\widehat{CF}(Y; \mathbb{Z}[H_1(Y)])$ is generated by $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$. The differential is given by

$$d\mathbf{x} = \sum_{\mathbf{y}} \sum_{\substack{\phi \in \pi_2(\mathbf{x}, \mathbf{y}) \\ \mu(\phi)=1 \\ n_z(\phi)=0}} \#\overline{\mathcal{M}}(\phi)[t(\phi)]\mathbf{y}.$$

Example 3.12. The standard Heegaard diagram for $Y = S^1 \times S^2$ gives a complex $\widehat{CF}(Y, \mathbb{Z}[H_1(Y)] = \widehat{CF}(Y, \mathbb{Z}[T^{\pm 1}])$ with $da = (T-1)b$ which is the standard complex for the twisted homology of S^1.

3.6 $+, -$, and ∞

To define the equivariant Floer groups, we use all domains in $\pi_2(\mathbf{x}, \mathbf{y})$ (not just those with $n_z(\phi) = 0$ and take twisted cohomology with respect to the map $\rho : \pi_2(\mathbf{x}, \mathbf{x}) \to \mathbb{Z}$.

Definition 3.13. $CF^\infty(\Sigma,\alpha,\beta,z)$ is the complex over $\mathbb{Z}[U^{\pm 1}]$ generated by $\mathbb{T}_\alpha\cap\mathbb{T}_\beta$ and with differential

$$dx = \sum_{\mathbf{y}} \sum_{\substack{\phi\in\pi_2(\mathbf{x},\mathbf{y})\\ \mu(\phi)=1}} \#\overline{\mathcal{M}}(\phi)U^{n_z(\phi)}\mathbf{y}.$$

Since the Maslov index of the domain $\mathcal{D}_\Sigma$ is 2 (exercise), U has homological grading 2.

The fact that $n_z(\phi)\geq 0$ whenever $\overline{\mathcal{M}}(\phi)\neq\emptyset$ implies that the $\mathbb{Z}[U]$ submodule of $CF^\infty(Y)$ generated by $\mathbf{x}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta$ is a subcomplex.

Definition 3.14. $CF^-(\Sigma,\alpha,\beta,z)$ is the complex over $\mathbb{Z}[U]$ generated by $\mathbb{T}_\alpha\cap\mathbb{T}_\beta$ and with differential

$$d\mathbf{x} = \sum_{\mathbf{y}} \sum_{\substack{\phi\in\pi_2(\mathbf{x},\mathbf{y})\\ \mu(\phi)=1}} \#\overline{\mathcal{M}}(\phi)U^{n_z(\phi)}\mathbf{y}.$$

It is a subcomplex (over $\mathbb{Z}$) of $CF^\infty(\Sigma,\alpha,\beta,z)$, and we define

$$CF^+(\Sigma,\alpha,\beta,z) := \frac{CF^\infty(\Sigma,\alpha,\beta,z)}{CF^-(\Sigma,\alpha,\beta,z)}$$

Remark 3.15. If Y is a rational homology sphere, fix $\mathbf{x}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta$ and let $\phi_{\mathbf{y}}$ be a class in $\pi_2(\mathbf{x},\mathbf{y})$. The map $\mathbf{y}\mapsto U^{n_w(\phi_{\mathbf{y}})-n_z(\phi_{\mathbf{y}})}$ defines an isomorphism between $CF^\infty(Y,\alpha,\beta,z)$ and $CF^\infty(Y,\alpha,\beta,w)$. It follows that $HF^\infty(Y,\mathfrak{s}_1)\simeq HF^\infty(Y,\mathfrak{s}_2)$ for all $\mathfrak{s}_1,\mathfrak{s}_2\in Spin^c(Y)$. This should make the first part of Property 9 from the first lecture at least plausible.

Generalized Heegaard Diagrams: We can also compute $HF^-(Y)$ starting from a balanced generalized Heegaard diagram (Σ,α,β,Z). To do so, we work over the ring $R_Z := \mathbb{Z}[U_1,\dots,U_k]$, where $Z=\{z_1,\dots,z_k\}$. We let

$$U_Z^n(\phi) := \prod_{i=1}^k U_i^{n_{z_i}}(\phi)$$

and define $CF^-(\Sigma,\alpha,\beta,Z)$ to be the complex over R_Z generated by $\mathbb{T}_\alpha\cap\mathbb{T}_\beta$ with differential

$$d\mathbf{x} = \sum_{\mathbf{y}} \sum_{\substack{\phi\in\pi_2(\mathbf{x},\mathbf{y})\\ \mu(\phi)=1}} \#\overline{\mathcal{M}}(\phi)U^{nZ(\phi)}\mathbf{y}.$$

Ozsváth and Szabó showed that if the generalized Heegaard diagram is admissible, the homology of this complex is $HF^-(Y)$, and that each U_i acts as U.

We can also consider the hat version of this construction, in which we only count domains with $n_{z_i}(\phi) = 0$ for all i. The homology of the resulting complex is $\widehat{HF}(Y) \otimes H^*(T^{k-1})$.

3.7 More about μ

We can compute the Maslov index $\mu(\phi)$ directly from the domain of ϕ.

Definition 3.16. If S is a component of $\Sigma - \alpha - \beta$, we define the Euler measure $m_e(S) := \chi(S) + n/4$, where n is the number of corners of S. If $\mathcal{D}$ is a domain in Σ, we define

$$m_e(\mathcal{D}) := \sum_S n_S(\mathcal{D}) m_e(S)$$

where the sum runs over all components of $\Sigma - \alpha - \beta$.

For $x \in \alpha \cap \beta$, let

$$n_x(\phi) := (1/4) \sum_{i=1}^{4} n_{S_i}(\phi),$$

where the sum runs over the 4 components of $\Sigma - \alpha - \beta$ adjacent to x. If $\mathbf{x} = \{x_1 \dots x_k\} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, we define $n_{\mathbf{x}}(\phi) = \sum_{i=1}^k n_{x_i}(\phi)$.

Theorem 3.17. *(Lipshitz's formula)* $\mu(\phi) = m_e(\mathcal{D}(\phi)) + n_{\mathbf{x}}(\phi) + n_{\mathbf{y}}(\phi)$.

A theorem of Lipshitz and Lee says that if Y is a rational homology sphere, we can use this formula to compute the difference in absolute gradings for generators in *different Spin^c* structures. Suppose that $\epsilon(\mathbf{x}, \mathbf{y})$ has order p in $H_1(Y)$. Then we can find a domain $\mathcal{D}$ in $\Sigma - z$ whose boundary has the form

$$\partial \mathcal{D} = \sum_{i=1}^{p} \sum_{j=1}^{g} \gamma_{\alpha_j, i} - \sum_{i=1}^{p} \sum_{j=1}^{g} \gamma_{\beta_j, i}$$

where each $\gamma_{\alpha_j,i}$ is a path from $\mathbf{x}_{\alpha,j}$ to $\mathbf{y}_{\alpha,j}$ along α_j, and similarly for the $\gamma_{\beta_j,i}$'s.

Theorem 3.18. *(Lipshitz -Lee)* $\operatorname{gr} \mathbf{x} - \operatorname{gr} \mathbf{y} = \mu(\mathcal{D})/p$.

3.8 More about $\overline{\mathcal{M}}$

The first question most people ask when they see the definition of $\widehat{HF}$ is "How can I compute $\overline{\mathcal{M}}(\phi)$?" In general, there is no easy answer. However there are a few things that it is good to keep in mind. (Here we assume we are using a an almost complex structure J induced from a complex structure on Σ.)

- If the multiplicity of any region in $\mathcal{D}(\phi)$ is negative, $\overline{\mathcal{M}}(\phi)$ is empty.
- If $\mathcal{D}(\phi)$ can be split into two disjoint components $\mathcal{D}_1$ and $\mathcal{D}_2$ and $\mu(\phi) = 1$, then $\mathcal{M}(\phi) = \mathcal{M}(\mathcal{D}_1) \times \mathcal{M}(\mathcal{D}_2)$. Since $\mu(\phi) = \mu(\mathcal{D}_1) + \mu(\mathcal{D}_2)$, wlog $\mu(\mathcal{D}_1) \leq 0$. Then either the map associated to $\mathcal{D}_1$ is the constant map, in which case $\mathcal{D}_1$ is a trivial domain consisting of isolated points and $\mathcal{M}(\phi) = \mathcal{M}(\mathcal{D}_2)$, or $\mathcal{M}(\mathcal{D}_1) = \emptyset$, so $\overline{\mathcal{M}}(\phi) = \emptyset$.
- There are some domains for which it can be shown that $\#\overline{\mathcal{M}}(\phi) = \pm 1$ for any choice of complex structure on Σ. In contrast, there are other domains for which $\overline{\mathcal{M}}(\phi)$ varies with the complex structure.
- A region of a Heegaard diagram is said to be *nice* if every component of $\Sigma - \alpha - \beta$ has nonnegative Euler measure (i.e. is a bigon, a rectangle, or an annulus.) A beautiful theorem of Sarkar shows that if $\mu(\phi) = 1$, $\mathcal{D}(\phi)$ is nice and $\mathcal{M}(\phi)$ is nonempty, then $\mathcal{D}(\phi)$ is either a bigon or a rectangle. This observation is the starting point for most combinatorial methods of computing Heegaard Floer homology.

3.9 Exercises

1. Show that $\mathrm{Sym}^n \mathbb{C} = \mathbb{C}^n$ (Hint: consider the correspondence which assigns to a monic polynomial of degree n its roots.) Deduce that $\mathrm{Sym}^n \Sigma$ is a complex manifold. Show that $\mathrm{Sym}^n \mathbb{CP}^1 = \mathbb{CP}^n$.
2. Show that $\mathcal{D}(\phi \# \psi) = \mathcal{D}(\phi) + \mathcal{D}(\psi)$.
3. Show that $\epsilon(\mathbf{x}, \mathbf{y}) + \epsilon(\mathbf{y}, \mathbf{z}) = \epsilon(\mathbf{x}, \mathbf{z})$.
4. Let $\phi_\Sigma \in \pi_2(\mathbf{x}, \mathbf{x})$ correspond to the domain $\mathcal{D}_\Sigma$ which has multiplicity 1 everywhere in Σ. Show that $\mu(\phi_\Sigma) = 2$.

5. Let J be a complex structure on $\mathrm{Sym}^g\Sigma$ induced by a complex structure on Σ. Use the Riemann mapping theorem to show that if $\mathcal{D}(\phi)$ is a bigon, then $\#\overline{\mathcal{M}}(\phi) = \pm 1$.

6. Suppose $\mathcal{D}(\phi)$ is a convex rectangle. If φ is a holomorphic representative of ϕ, describe S_ϕ, p and π. Show that $\#\overline{\mathcal{M}}(\phi) = \pm 1$.

7. Using the Heegaard diagram of $S^3 - \nu(T)$ drawn in lecture, draw a Heegaard diagram for T_0 (0-surgery on the trefoil). List the generators and partition them into equivalence classes. What happens if we do a different surgery on T?

8. Show that $\widehat{HF}$ satisfies the Property 1 from the first lecture. (Hint: for the connected sum, put the basepoint in the connected sum region.)

9. Construct the exact sequences and spectral sequences of Property 8 from the first lecture.

10. Use Lipshitz and Lee's theorem to compute the differences in absolute grading for the generators of $\widehat{HF}(L(5,1))$. Do the same thing for $L(5,2)$.

3.10 Exercises for Lectures 3 and 4

1. Suppose (M,γ) is a sutured manifold and $D \hookrightarrow M$ is a properly embedded disk with $|\partial D \cap \gamma| = 2$. Show that the sutured manifold (M',γ') obtained by decomposing (M,γ) along D does not depend on which orientation we give D. (The manifold (M',γ') is said to be obtained from (M,γ) by a *disk decomposition.*) Use the decomposition formula to show that $SFH(M',\gamma') \simeq SFH(M,\gamma)$.

2. Prove the result of the first exercise directly from the definition of SFH.

3. A sutured manifold (M,γ) is *disk-decomposable* if it can be reduced to the trivial sutured manifold (B^3,γ_0) be repeated disk decompositions. Show (without appealing to Floer homology) that if (M,γ) is disk decomposable, it is a product manifold. Find a disk decomposition for the complement of the trefoil knot, and deduce that it is fibred. Do the same for each knot in S^3 with ≤ 7 crossings and monic Alexander polynomial.

4. Let K_n be the n-twist knot in S^3, and let (M_n, γ) be sutured manifold given by the complement of its standard Seifert surface with a longitudinal suture. Use a disk-decomposition to compute $SFH(M_n, \gamma_n)$.

5. As in the last problem, but using the pretzel knot $P(3,3,3)$. Draw an arc diagram for (M, γ) and use it to compute $SFH(M, \gamma)$.

6. Let K be the $(2,5)$ torus knot. Use the mapping cone to compute the homology of p surgery on K for $p = -3, 0, 3$.

7. Let $K \subset S^3$ be a knot, and let $C^+(n, K)$ be the mapping cone for n surgery on K. Show that if $n > 0$, the map $G^+_{n,\mathfrak{s}_k} : CF^+(S^3) \to C^+(n, K)$ vanishes on all elements of sufficiently high degree. Conversely show that if $n < 0$, show that it is an isomorphism in high degrees.

8. With notation as in the previous problem, compare the relative gradings of the images of $G_{-1,\mathfrak{s}_i}$ for differing values of i. Check that it is compatible with the expected shift in the absolute grading corresponding to the map $F^+_{W_{-1}(K),\mathfrak{s}_i}$.

9. Suppose $K \subset Y$ is a knot in a homology sphere, and let $g(K_0)$ be the minimal genus of a representative of $H_2(K_0)$. Show by example that we can have $g(K_0) < g(K)$. (Hint: start with a 2 component link whose algebraic linking number is 1 but whose geometric linking number is > 1.)

10. Suppose $K \subset S^3$ is a knot, and that K_n is an L-space for some $n > 0$. Show that each of the groups $\widehat{A}_i$ appearing in the mapping cone must be $\mathbb{Z}$. If each $\widehat{A}_i \simeq \mathbb{Z}$, show that K_n is an L -space if and only $n \geq 2g(K) - 1$.

Lecture 4

What is an L-space?

We have seen that, to a rational homology sphere Y and a $Spin^c$-structure $\mathfrak{s} \in Spin^c(Y)$ Heegaard Floer homology assigns a group $\widehat{HF}(Y, \mathfrak{s})$. For the purposes of this lecture, we will restrict attention to coefficients in the two element field $\mathbb{F}$; as a result, the invariant is just a vector space. In it's simplest form, the homological grading is a $\mathbb{Z}/2\mathbb{Z}$-grading and, relative to this grading, these homology groups enjoy the property that $\chi(\widehat{HF}(Y, \mathfrak{s})) = 1$. Said another way,

$$\dim(\widehat{HF}(Y, \mathfrak{s})) \geq 1.$$

Definition 4.1. An L-space is a rational homology sphere with simplest-possible Heegaard Floer homology, in the sense that $\dim(\widehat{HF}(Y, \mathfrak{s})) = 1$ for every $\mathfrak{s} \in Spin^c(Y)$. Equivalently, L-spaces satisfy the equality

$$\dim \widehat{HF}(Y) = |H_1(Y; \mathbb{Z})|.$$

Notice that S^3 is a first example of an L-space since $\widehat{HF}(S^3, \mathfrak{s}_0) = \widehat{HF}(S^3) \cong \mathbb{F}$ where $\mathfrak{s}_0$ is the unique $Spin^c$-structure on S^3. Consequently, the *simplicity* of the Heegaard Floer homology associated with L-spaces may be viewed as an inability to distinguish S^3 from Y in any given $Spin^c$-structure, at least as a graded group. (Of course, the number of $Spin^c$-structures alone —that is, the order of the group $H_1(Y; \mathbb{Q}) = H^2(Y; \mathbb{Q})$ — may well be enough to distinguish Y from S^3.)

A large class of examples of L-spaces is given by lens spaces where, in keeping with our restriction to rational homology spheres, $S^2 \times S^1$ is not included as a lens space.

Exercise 4.2. Working from the definition, compute the Heegaard Floer homology of a lens space $L(p, q)$.

In fact, L-space is short for *Heegaard Floer homology lens space.* From the longer moniker, it should be clear that there is a direct analogy with homology spheres with respect to $H_1(-;\mathbb{Z})$, or rational homology spheres with respect to $H_1(-;\mathbb{Q})$. Namely, these are the objects that the invariant of interest fails to separate from some standard, simple object.

Recall that $\widehat{HF}(Y)$ may be derived form a somewhat richer invariant $HF^+(Y)$; this latter object is a module over $\mathbb{F}[U]$ and $\widehat{HF}(Y)$ is the cone on the map $U\colon\ HF^+(Y) \to HF^+(Y)$.

Definition 4.3. The reduced Heegaard Floer homology of Y is identified with the torsion of $HF^+(Y)$ as a $\mathbb{F}[U]$-module.

This gives rise to an alternate definition for L-spaces.

Exercise 4.4. Show that a rational homology sphere Y is an L-space if and only if $HF_{\text{red}}(Y)$ vanishes.

As a result, given an L-space Y we have that $HF^+(Y, \mathfrak{s}) \cong M_+$ for every $\mathfrak{s} \in Spin^c(Y)$.

There is some additional structure that will not feature heavily in this lecture, but which is extremely important to the theory more generally. While an L-space is indistinguishable from a lens space, it may be distinguished as a graded group by the d-invariants, or, correction terms. Recall that his is a rational number $d(Y, \mathfrak{s})$ assigned to the element in lowest U-degree in $HF^+(Y, \mathfrak{s}) \cong M_+$.

Example 4.5. The Poicaré conjecture, now known to be true by Perelman's work, was originally (falsely) stated by Poincaré for homology groups and not homotopy groups. Poincaré quickly corrected his own mistake, and produced an example of a three-manifold Y for which $H_1(Y; bZ) = 0$ — this is now known as the Poincaré sphere [Poi10]. There are many constructions known for this manifold; see Rolfsen for an illustrative tour [Rol76].

In order to see the fundamental group immediately,the Poincaré may be obtained by considering the quotient of $\mathrm{SO}(3)/I$, where I is the icosahedral group, that is, the rotational symmetry group of the regular icosahedron. The universal cover of SO(3) is S^3 and the perfect double-cover of I in S^3 is the binary icosahedra group, denoted $\tilde{I}$. The Poincaré homology sphere is $Y \cong S^3/\tilde{I}$, and $\pi_1(Y)$ is therefore isomorphic to the (perfect) group $\tilde{I}$.

This manifold turns out to be an L-space, so that $\widehat{HF}(Y) \cong \widehat{HF}(Y)$. However, $d(S^3) = 1$ and $d(Y) = \pm 2$ (depending on the orientation).

Among integer homology spheres, L-spaces appear to be quite rare. In fact, we have now recorded a list of known irreducible examples.

Question 4.6. *Are the only irreducible L-space integer homology spheres the Poincare homology sphere, its mirror image, and the three-sphere?*

The answer to this question is conjectured to be yes; this conjecture appears to be due, by general consensus, to Ozsváth and Szabó.

While, at first blush, the study of L-spaces may seem to be simply a question of better understanding the behaviour of a given — in this case relatively new — invariant, there is much more going on. Indeed, L-spaces come up frequently in applications of Heegaard Floer homology and, consequently, better understanding this class of manifolds lies at the heart of a great many problems whose solution might apply Heegaard Floer homology. We will illustrate this through an exercise. First, we construct an important class of three-manifolds: Two-fold branched covers.

Recall that the knot group $\pi_1(K)$ is identified with the group $\pi_1(S^3 \smallsetminus \nu(K))$, where $\nu(K)$ is an open tubular neighbourhood of K. Since $H_1(S^3 \smallsetminus \nu(K);\mathbb{Z}) \cong \mathbb{Z}$, there is a unique non-trivial projection $\pi_1(K) \to \mathbb{Z}/2\mathbb{Z}$. Consider the two-fold cover $M \to S^3 \smallsetminus \nu(K)$ determined by this projection, that is, $\pi_1(M) \cong \ker(\pi_1(K) \to \mathbb{Z}/2\mathbb{Z})$.

Exercise 4.7. Show that there is a unique extension of M to a closed manifold $\mathbf{\Sigma}_K$ by attaching a solid torus $D^2 \times S^1$ in such a way that the resulting cover $\mathbf{\Sigma}_K \to S^3$ is one-to-one along the knot K.

The closed manifold $\mathbf{\Sigma}_K$ is called the two-fold branched cover of K. It is the two-fold branched cover of S^3, branched along the knot K. Note that this is sometimes called the *double branched cover* or the *branched double-cover*. Now for the exercise.

Exercise★ 4.8. Use Heegaard Floer homology to prove that reduced Khovanov homology detects the trivial knot, assuming Conjecture 4.6. To do this, you will make use of the fact that (1) there is a spectral sequence with $E_2 = \widetilde{Kh}(K)$ and converging to $E_\infty = \widehat{HF}(-\mathbf{\Sigma}_K)$ [OS05b]; and (2) that the Poincaré homology sphere (indeed, any two-fold branched cover with finite fundamental group) is realised as a two-fold branched cover of S^3 in a unique way [Wat12]. In this case, the branch set is the $(3,5)$-torus knot; you might try and convince yourself of this.

A recent result of Kronheimer and Mrowka establishes that $\widetilde{Kh}(K) \cong \mathbb{F}$ if and only if K is the trivial knot [KM11]. This uses a variant of the spectral sequence used in the previous exercise, which converges instead to instanton Floer homology.

Examples

The following are some natural classes of L-spaces.

Example 4.9. Suppose Y has elliptic geometry, that is, $Y = S^3/\Gamma$ where Γ is a subgroup of the $O(4)$, the isometry group of the sphere. For example, the Poincaré sphere admits elliptic geometry. Note that manifolds admitting eliptic geometry necessarily have finite fundamental group; the converse to this assertion holds as a result of Perelman's proof of geometrization. Work of Ozsváth and Szabó establishes that every such Y is an L-space [OS05a].

Example 4.10. If K is an alternating knot (or, more generally, a non-split alternating link), then $\mathbf{\Sigma}_K$ is an L-sapce [OS05b]. This immediately gives rise to hyperbolic examples of L-spaces.

Note that the proof of this fact singles out an apparently more general class of links. First recall that given a crossing in ╳ in some fixed diagram of a link there are two *resolutions* of the crossing that yield diagrams with one fewer crossing:

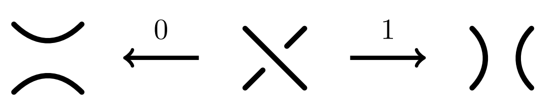

The set of quasi-alternating links $\mathcal{Q}$ is the smallest set of links containing the trivial knot that is closed under the following relation: if L admits a projection with distinguished crossing $L($╳$)$ so that $\det(L($╳$)) =$

$\det(L(\asymp)) + \det(L()\,()))$ where $L(\asymp), L()\,() \in \mathcal{Q}$, then $L = L(\times) \in \mathcal{Q}$ as well.

Those familiar with Khovanov homology will recognize these resolutions as a first step in constructing a chain complex. It is a result of Manolescu and Ozsváth that $\widetilde{Kh}(L)$ is a supported in a single diagonal. It follows from this fact that $\dim \widetilde{Kh}(L) = \det(L) = |H_1(\mathbf{\Sigma}_L;\mathbb{Z})|$. Now the spectral sequence used (as a black box!) in Exercise 4.8 shows that $\dim \widehat{HF}(\mathbf{\Sigma}_K) = |H_1(\mathbf{\Sigma}_L;\mathbb{Z})|$ whenever L is quasi-alternating.

Exercise 4.11. Prove that alternating knots are quasi-alternating.

Exercise* 4.12. Applying the surgery formula for Heegaard Floer homology to the two-fold branched covers of $\{L(\times), L(\asymp), L()\,()\}$, prove directly that if $L \in \mathcal{Q}$ then $\mathbf{\Sigma}_L$ is an L-space.

A conjectured connection with the fundamental group

It is natural to ask if there is a characterization of L-spaces that is topological and, in particular, makes no reference to Heegaard Floer homology. Closely related to this is the question of whether a connection between Heegaard Floer homology and the fundamental group exists.

Examples correlating Heegaard Floer homology and left-orderability began to appear (in print) in 2009 [Pet, Wat09]; examples were likely noticed prior. The following is formalized in the work of Boyer, Gordon and Watson.

Conjecture 4.13. *Let Y be an irreducible, closed, connected, orientable three-manifold. Then Y is an L-space if and only if $\pi_1(Y)$ is not left-orderable.*

Notice that, since irreducibility has been required, both the non-left-orderability of $\pi_1(Y)$ as well as the definition of L-space restrict Y to the class of rational homology spheres.

Exercise 4.14. Irreducibility is required: If $HF_{\mathrm{red}}(Y) \neq 0$ then $HF_{\mathrm{red}}(Y\#Y') \neq 0$ but, on the other hand, if $\pi_1(Y')$ is not-left-orderable, then neither is $\pi_1(Y\#Y') \cong \pi_1(Y) * \pi_1(Y')$. Can you give an explicit example?

In this context, the class of non-left-orderable groups presents a natural expansion of the class of finite groups. In the same way, L-spaces include the class of manifolds with elliptic geometry but are apparently a much larger class.

Exercise★ 4.15. Prove that the Weeks manifold (compare Exercise 2.3) is an L-space. Hint: The Weeks manifold is surgery on the Whitehead link.

Seifert fibred spaces

A Seifert structure on a three-manifold is a foliation by circles; a simple example is given by $D^2 \times S^1$. This example is a little misleading, since not every Seifert structure is a circle bundle over a surface. Indeed, for a given pair of relatively prime integers (p, q), with $p \geq 1$, we may define

$$V_{p,q} = (D^2 \times I)/\{(x, 1) = (e^{2\pi i \frac{p}{q}}, 0)\}.$$

This quotient is homeomorphic to a solid torus, though a $\frac{p}{q}$-twist has been added. As a result the induced foliation by circles is non-standard: The core circle of $D^2 \times S^1$ is now a *singular fibre* of order p whenever $p > 1$ (the $V_{1,p}$ are, in fact, circle bundles).

We can take this to be the definition of a singular fibre in a Seifert structure in general, sine a result of Epstein establishes that every Seifert structure on the solid torus is fibre-preserving diffeomorphic to one of these $V_{p,q}$ [Eps72]. That is, a Seifert structure on an arbitrary three-manifold is a foliation by circles such that the tubular neighbourhood of any fibre is fibre-preserving diffeomorphic to some $V_{p,q}$. A manifold with a fixed Seifert structure is called a Seifert fibred space.

Note that the index $p \geq 1$ of the fibred solid torus is the index of the fibre in the definition above; the fibre is *singular* or *exceptional* if $p > 1$. In general, the leaf space (or, orbit space) of a Seifert fibred space is a surface together with a finite collection of points, corresponding to the image of the singular fibres, labeled by the index of these fibres. This leaf space is called the *base orbifold* of the Seifert fibred space. Note that the underlying surface need not be orientable, even if the three-manifold is; fibres may be coherently oriented locally, but need not admit a global orientation that is coherent with an orientation on the manifold.

The conjecture is known to hold for Seifert fibred spaces.

Theorem 4.16. *Suppose Y is a closed, connect, orientable Seifert fibred space. Then $\pi(Y)$ is left-orderable if and only if Y is an L-space.*

Proof. Let Y be a Seifert fibred space. There is a short exact sequence

$$1 \longrightarrow \langle\langle\varphi\rangle\rangle \longrightarrow \pi_1(Y) \longrightarrow \pi_1^{\mathrm{orb}}(\mathcal{B}) \longrightarrow 1$$

where the subgroup $\langle\langle\varphi\rangle\rangle$ the normal closure of a regular fibre φ (see [Sco83, Lemma 3.2] for details). The quotient $\pi_1^{\mathrm{orb}}(\mathcal{B})$ is the orbifold fundamental group of the leaf space $\mathcal{B}$ with underlying closed surface B (that it, B is obtained by ignoring the marked points recording the singular fibres in $\mathcal{B}$). As a result, both $\pi_1(B)$ and $H_1(B;\mathbb{Z})$ are quotients of $\pi_1^{\mathrm{orb}}(\mathcal{B})$. This places strong restrictions on B whenever $H_1(Y;\mathbb{Q}) = 0$. Indeed, surjectivity is preserved under abelianization and hence the surjection $\pi_1(Y) \to \pi_1(B)$ gives a surjection $H_1(Y;\mathbb{Z}) \to H_1(B;\mathbb{Z})$ so it must be that $H_1(B;\mathbb{Z})$ is finite hence B is one of S^2 or RP^2.

As a result, there are two cases to consider when the given Seifert fibred space is a rational homology sphere, as in the present setting. We appeal then to the work of Boyer, Rolfsen and Wiest [BRW05, Theorem 1.3]:

(1) If Y has base orbifold with underlying surface S^2, then $\pi_1(Y)$ is left-orderable if and only if Y admits a co-orientable taut foliation; and

(2) if Y has base orbifold with underlying surface RP^2, then $\pi_1(Y)$ is not left-orderable.

In the first case, work of Lisca and Stipsicz completely characterises non-L-spaces in terms of existence of co-orientable taut foliations [LS07]. In the second case, a calculation using the surgery exact sequence is required to establish that every Seifert fibred rational homology sphere with leaf space RP^2 is an L-space [BGW13]. In particular, this class of manifolds may be constructed inductively from circle bundles over RP^2 and Dehn surgery on a regular fibres. □

This proof hints at another structure on three-manifolds that we will return to later: a co-orientable taut foliation is a foliation of a three-manifold by surfaces that are cooriented and for which there is a closed curve meeting every leaf of the foliation transversely.

Note that Seifert fibred spaces account for 6 of the 8 geometries arising on geometric three-manifolds [Sco83]. Recall that a general three-manifold is first decomposed along spheres and tori before each of the pieces is endowed with a geometric structure.

In the general setting, an exciting aspect of the conjecture — whether or not you think it is a reasonable one — is the predictions it makes based on the fact that two apparently unrelated concepts are brought into equivalence with one another. Simplicity in Heegaard Floer theory, on the one hand, and the non-existence of an auxiliary structure on the fundamental group, on the other. We'll expand on two instances of this in the following sections.

Something L-spaces predict about left-orderable groups

An immediate question whenever a statement is proved for Seifert fibred spaces is "what about the hyperbolic case?", since hyperbolic three-manifolds constitute (in some sense) the generic three-manifold (at least in the geometric setting). Fortunately, we have seen a very large and easy to describe class of L-spaces that contain many hyperbolic examples:

Theorem 4.17 (Ozsváth-Szabó [OS05b]). *If K is an alternating knot then $\boldsymbol{\Sigma}_K$ is an L-space.*

In this case, the conjecture predicts that $\pi_1(\boldsymbol{\Sigma}_K)$ is not-left-orderable. In the absense of the conjecture, there is no reason to expect that this should be true, especially in light of the fact that many discrete, torsion free subgroups of $\mathrm{PSL}_2(\mathbb{C})$ will arise from this construction. The fact that $\boldsymbol{\Sigma}_K$ is a hyperbolic three-manifold, in some generic sense, is due to the tricotomy for knots in S^3: If K is a torus knot or a satellite knot then $\boldsymbol{\Sigma}_K$ is a Seifert fibred space or contains an essential torus, respectively. Neither of these is hyperbolic. Thus, among hyperbolic knots, we have that either K is Montesinos (this includes two-bridge knots) or it is not — both classes are infinite. The later class will have hyperbolic two-fold branched covers. We note that, by restricting to knots, Sol geometry does not arise (see the next section, particularly Exercise 4.26).

Nevertheless, the conjecture holds up in this setting.

Theorem 4.18 (Boyer-Gordon-Watson [BGW13]). *If K is an alternating knot then $\pi_1\boldsymbol{\Sigma}_K$ is not left-orderable.*

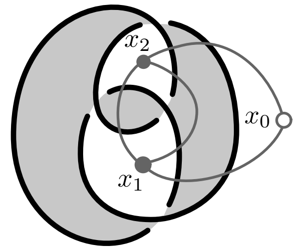

The obvious approach to proving this theorem is to find a presentation for $\pi_1(\mathbf{\Sigma}_K)$ that depends on the diagram and assuming that the group is left-orderable, show that a contradiction is reached when the diagram is alternating. The original proof makes use of a presentation called Wada's group that depends on the diagram [BGW13]; Greene gives a similar proof using a group presentation appealing to graph theoretic data that is essentially dual to Wada's [Gre]. In both cases, the group presentation is indexed by the black graph associated with the knot diagram. This is a signed graph in general, but the signs all agree for alternating knots (up to taking mirrors, we may assume that the setup is as in the diagram for the figure eight knot shown ion the right). In particular, assign a vertex to each white region and an edge to each crossing. Label the vertices $\{x_0, x_1, \dots, x_n\}$. Then Greene's presentation has generators x_i and relations

$$\prod_{\text{edges } (x_j, x_i) \text{ incident to } x_i} (x_j^{-1} x_i)$$

and the additional relation x_0 (that is, there is a root vertex that is set to the identity in the group). Greene proves that the group obtained in this way is isomorphic to $\pi_1(\mathbf{\Sigma}_K)$ [Gre13].

For the example shown, we obtain

$$\begin{aligned}\pi_1(\mathbf{\Sigma}_{4_1}) &\cong \langle x_0, x_1, x_2 | x_0, x_2^{-1}x_1x_2^{-1}x_1x_0^{-1}x_1, x_1^{-1}x_2x_1^{-1}x_2x_0^{-1}x_2\rangle \\ &\cong \langle x_1, x_2 | (x_2^{-1}x_1)^2x_1, (x_1^{-1}x_2)^2x_2\rangle\end{aligned}$$

which can't be left orderable: The presentation is symmetric so if $x_1 < x_2$ and x_2 is positive then $(x_1^{-1}x_2)^2x_2$ is positive, a contradiction. This should set you at ease, as the two-fold branched cover of a two-bridge knot is always lens space, and finite cyclic groups ($\mathbb{Z}/5\mathbb{Z}$, in this example) cannot be left-ordered. There's a little more work to do, but the real point here is that we can choose a maximal, positive element among the generators — it's a finite set, and we can switch to the opposite order if needed.

Exercise 4.19. Using Greene's presentation, complete the proof of Theorem 4.18.

We will outline a different proof that illustrates just how special the class of two-fold branched covers of alternating knots is, even among L-spaces. This is due to Levine and Lewallen [LL].

First, consider an abelian group on 4 symbols $X = \{0, +, -, *\}$ where multiplication is defined according to (1) $0\epsilon = \epsilon 0 = 0$ for $\epsilon \in X$; (2) $++ = -- = +$; (3) $+- = -+ = -$; and (4) $\epsilon * = *\epsilon = *$ for $\epsilon \in \{+, -, *\}$. Given a group presentation $\langle x_1, \ldots, x_m | r_1, \ldots, r_n \rangle$ form a matrix (ϵ_{ij}) with entries in X according the the rules:

$$e_{ij} = \begin{cases} 0 & \text{if neither } x_i \text{ nor } x_i^{-1} \text{ occur in the relator } r_j; \\ + & \text{if } x_i \text{ appears in the relator } r_j, \text{ but } x_i^{-1} \text{ does not;} \\ - & \text{if } x_i^{-1} \text{ appears in the relator } r_j, \text{ but } x_i \text{ does not; and} \\ * & \text{if both } x_i \text{ and } x_i^{-1} \text{ occur in the relator } r_j. \end{cases}$$

With this matrix assigned to the group presentation for G in hand, Levine and Lewallen prove the following criterion:

Exercise 4.20. Suppose that for any $d_1, \ldots, d_m \in \{0, +, -\}$, not all zero, the result of multiplying the i^{th} row of the matrix (ϵ_{ij}) by d_i has a non-zero column with entries only in $\{+, -\}$. Prove that G is not left-orderable. Trick: Left-order G and choose d_i according to the signs of the generators.

This criterion is cleverly applied to a particular class of groups:

Proposition 4.21. *Let $G \cong \langle x_1, \ldots, x_n | r_1, \ldots, r_n \rangle$ with associated matrix (ϵ_{ij}) satisfying*

(1) There exists a permutation $\sigma_0 \in S_n$ such that the entries $\{\epsilon_{i\sigma_0(i)}\}$ are non-zero for all $1 \le i \le n$;

(2) For any permutation $\sigma \in S_n$ with $\{\epsilon_{i\sigma(i)}\}$ all non-zero, we have $\epsilon_{i\sigma(i)} \in \{+, -\}$ for all $1 \le i \le n$; and

(3) For any two permutations $\sigma, \sigma' \in S_n$ satisfying (2), we have

$$\text{sign}(\sigma) \prod_{i=1}^{n} \epsilon_{i\sigma(i)} = \text{sign}(\sigma') \prod_{i=1}^{n} \epsilon_{i\sigma'(i)}.$$

Then G is not left-orderable.

Proof. Up to reordering the generators, we may assume that σ_0 is the identity so that, in application of (2), ϵ_{ii} is either $+$ or $-$ for every $1 \le i \le n$. To see that G is not left-orderable we will show that (ϵ_{ij}) satisfies the hypothesis laid out in Exercise 4.20.

For a contradiction, select the d_i and construct a matrix (m_{ij}) by multiplying the i^{th} row of (e_{ij}) by d_i where every column contains an off-diagonal entry that is neither $+$ nor $-$. We use this information to construct a permutation σ satisfying (1) and (2) but violating (3) when compared to the identity.

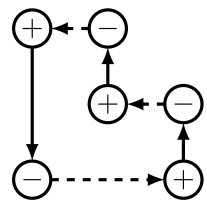

This is a proof whose key ideas can be illustrated on a simple example: Suppose that $m_11 = +$. Notice that there is an index i_1 for which $m_{i_1}1 = -$ (this is by our assumption on (m_{ij}) together with property (2)). But now $m_{i_1 i_1} = +$ (or perhaps $-$, but this is not so important), and we can repeat the process until we get a closed loop. Notice that some row/columns might be skipped in the process of carrying this out. However, the result is that if the loop closes up in k-steps then we may form a k-cycle σ that takes $1 \mapsto i_1$, $2 \mapsto i_2$, and so forth. The sign of this permutation is $(-1)^{k-1}$ and there are k occurrences of $-$. So we have our contradiction: You can check that σ satisfies (1) and (2) but we have, buy construction, that

$$(-1)^{2k-1}\,\text{sign}(\sigma)\prod_{i=1}^{n} \epsilon_{i\sigma(i)} = \text{sign}(\text{id})\prod_{i=1}^{n} \epsilon_{ii}.$$

Of course, it might not be possible to start with m_11, but the indices can be reordered after the fact, once you have constructed the loop. □

Definition 4.22. A strong L-space is a rational homology sphere admitting a Heegaard diagram for which $\dim \widehat{CF}(Y) = |H_1(Y;\mathbb{Z})|$. That is, strong L-spaces can be realised as L-spaces on the chain level.

Exercise★ 4.23. Given a strong Heegaard diagram for a strong L-space, show that the resulting group presentation satisfies the hypothesis of Proposition 4.21.

It is an observation of Greene that alternating diagrams of knots give rise to Heegaard diagrams for the two-fold branched cover for which there is a single generator in each *Spin*^{c}-structure [Gre13]. Interestingly, there are no known examples of strong L-spaces that are *not* the two-fold branched cover of an alternating knot (or, non-split alternating link).

Something left-orderable groups predict about L-spaces

In the other direction, a great deal is known about the left-orderability of certain classes of three-manifold groups, due to the work of Boyer, Rolfsen and Wiest. For example:

Theorem 4.24 (Boyer-Rolfsen-Wiest)**.** *If Y is a rational homology sphere admitting Sol geometry the $\pi_1(Y)$ is not left-orderable.*

In this case, we have a new suggested class of L-spaces. As expected, at least if one trusts the conjecture, we have:

Theorem 4.25 (Boyer-Gordon-Watson)**.** *If Y is a rational homology sphere admitting Sol geometry then Y is an L-space.*

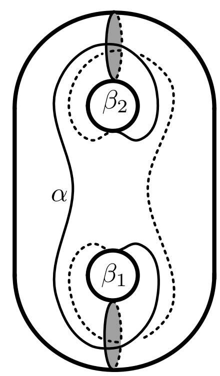

We will outline the proof as it serves to advertise some additional structure in Heegaard Floer theory that is not emphasized in these notes. The first fact we make use of is a characterization of Sol rational homology spheres: They are all of the form $Y = M \cup_h M$ where M is the twisted I-bundle over the Klein bottle, and h is a homeomorphism identifying the torus boundaries of M. The twisted I-bundle over the Klein bottle is an orientable three-manifold admitting a pair of Seifert structures: One with base orbifold a disk with two cone points each of order 2; the other with base orbifold a Möbius band. A Heegaard diagram describing this manifold is shown on the right. The β-curves bound the collapsing disks in the handlebody shown, and notice that there is one fewer α-curve corresponding to the fact that M has a torus boundary.

Exercise 4.26. Compute the fundamental group of M based on the given Heegaard diagram, and show that it is isomorphic to the Klein bottle group. Use the description given to exhibit both Seifert structures on M. Hint: Decompose the handlebody into solid tori first. Finally, can you make use of the obvious symmetry to realise M as a two-fold branched cover?

Recent work of Lipshitz, Ozsváth and Thurston has introduced a gluing theorem to Heegaard Floer homology [LOTa]. This comes from a theory for three-manifolds-with-boundary called bordered Heegaard Floer homology, which ultimately allows us to decompose along the torus and obtain $\widehat{HF}(Y)$ as the homology of a chain complex of the form

$$\widehat{CFA}(N) \boxtimes \widehat{CFD}(N).$$

Rather than providing all the details of this theory we will just use the example at hand to introduce all of the players. For the interested reader there is a good overview of the theory [LOT11] and an excellent set of notes from a recent summer school [LOTb]

The first hint that the above tensor-like statement is an oversimplification is that the dependence on h has vanished. In bordered Floer homology, the boundary data is kept track of by an algebra $\mathcal{A}$, in this setting called the *torus algebra*. How this algebra acts on two flavours of homological invariants in a way that depends on, and is highly sensitive to, a parametrization of the boundary. In the present setting, we have a very natural choice of basis for $H_1(\partial M\ bZ)$ given by the two fiber slopes (that is, regular fibers that we choose in the boundary). Let φ be a regular fiber for the Seifert structure over the disk, and let λ be the regular fiber for the Seifert structure over the Möbius band.

Exercise 4.27. Find a pair of curves in the handlebody describing M that avoid the α arc and intersect once transversally. Argue that this describes a basis for $H_1(\partial M\ bZ)$ and show that you can make your choice coincide with (φ, λ).

Recall that the longitude of a knot K in S^3 is specified by the intersection of a Seifert surface for K with $\partial(S^3 \smallsetminus \nu(K))$. The fiber slope λ plays the role of the longitude, in the following sense:

Exercise 4.28. Prove that, as an element of $H_1(M;\mathbb{Z})$, $[\lambda]$ has order 2.

For the purposes of this discussion (in particular, this only makes sense in the case of torus boundary!), a bordered structure on M is given by the ordered triple (M, λ, φ). Now $\widehat{CFD}(M, \lambda, \varphi)$ is a left differential module over $\mathcal{A}$ and (a type D structure) and $\widehat{CFA}(M, \lambda, \varphi)$ is a right A_∞-module over $\mathcal{A}$ and (a type A structure). These homological objects represent different packaging of equivalent data; they are dual to each other in an appropriate sense. In this notation

$$\widehat{CFA}(M, \lambda, \varphi) \boxtimes \widehat{CFD}(M, \lambda, \varphi)$$

corresponds to the homeomorphism determined by $\{h(\lambda) = \varphi, h(\varphi) = \lambda\}$ (on homology, $h_* = \left(\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right)$) and the product $\boxtimes$ is a version of the tensor product taken over the idempotent subring of $\mathcal{A}$.

The main point here is that, in order to alter the identification h, we are forced to re-parametrise the boundary data. For example, to calculate the result of identifying $\{h(\lambda) = \varphi + \lambda, h(\varphi) = \lambda\}$ (that is, $h_* = \left(\begin{smallmatrix}1 & 1\\ 1 & 0\end{smallmatrix}\right)$), we would want

$$\widehat{CFA}(M, \lambda, \varphi) \boxtimes \widehat{CFD}(M, \lambda, \varphi + \lambda).$$

At this stage a striking property of the bordered invariants for the twisted I-bundle over the Klein bottle is observed by a direct calculation: There is a homotopy equivalence between the differential modules

$$\widehat{CFD}(M, \lambda, \varphi + \lambda) \cong \widehat{CFD}(M, \lambda, \varphi)$$

meaning that the chain complexes

$$\widehat{CFA}(M, \lambda, \varphi) \boxtimes \widehat{CFD}(M, \lambda, \varphi + \lambda) \cong \widehat{CFA}(M, \lambda, \varphi) \boxtimes \widehat{CFD}(M, \lambda, \varphi)$$

are homotopy equivalent and hence have the same homology.

Now in the case where one fibre slope is glued to another, it is possible to extend the Seifert structures on both manifolds to a Seifert structure on the closed manifold. In the case at hand, this means that $N \cup_h N$ is actually Seifert fibred for certain identifications, and in these cases we already know that the conjecture holds (see Theorem 4.16). In fact, for any matrix representative of h_* for which there is a zero in a corner, the resulting $N \cup_h N$ is Seifert fibered. On the other hand, at the level of Heegaard Floer homology, as a result of this property of the bordered invariants of the twisted I-bundle over the Klein bottle, there is no distinction between identifying with a homeomorphism of the form $\left(\begin{smallmatrix}* & *\\ 1 & *\end{smallmatrix}\right)$ versus one of the form$\left(\begin{smallmatrix}0 & *\\ 1 & *\end{smallmatrix}\right)$.

This reduces to the case of Seifert fibred L-spaces, and ultimately provides a base case for induction to deal with homeomorphisms of the form $\left(\begin{smallmatrix}* & *\\ n & *\end{smallmatrix}\right)$.

Remark 4.29. The property that $\widehat{CFA}(M, \lambda, \varphi) \boxtimes \widehat{CFD}(M, \lambda, \varphi + \lambda)$ as differential modules over $\mathcal{A}$ establishes that M is a *Heegaard Floer homology solid torus*, by analogy with Heegaard Floer homology lens spaces. It turns out that, just as in the case of L-spaces versus lens spaces, there are lots (infinitely many, in fact) of Heegaard Floer homology solid tori that are not from $D^2 \times S^1$ [Wat].

Further reading

The first examples illustrating the correspondence in the conjecture were very likely those provided by Seifert fibred spaces with orientable base orbifold, as this follows immediately on juxtaposition of the results of Boyer, Rolfsen and Wiest [BRW05] and Lisca and Stipsicz [LS07]. This is recorded in print by Peters [Pet] and the full statement for Seifert fibered spaces appears in [Wat09]. Peters gives a range of other interesting examples that we have not include here corresponding to families of p-fold cyclic branched covers. Another class of examples is studied by Li and Watson [LW14]: If Y is an L-space admitting a genus one open book decomposition then $\pi_1(Y)$ is not left-orderable. This implies, for instance, that the two-fold branched covers of quasi-alternating links with braid index at most three have non-left-orderable fundamental group.

There is an important body of computational evidence for the conjecture thanks to work of Dunfield that should not go unmentioned. This is mentioned on the blog *Low Dimensional Topology* [Dun11] however there have been developments since then that will limey appear in print soon. Dunfield's work represents a rather intensive study of low-volume hyperbolic three-manifolds with regards to the conjecture, and the take home point is that the conjecture holds up on this class (there are roughly 11,031 closed three-manifolds in this sample) with surprising odds: The naive probability that the data found supports the conjecture (if one takes the view that it is false and looks for a counterexample) is 10^{-29}. With some adjustments for the samples, Dunfield has estimated that the odds are probably more like 10^{-18}, though this is admittedly splitting hairs.

An interesting fact is that every example we have discussed, as well as those considered by Dunfield, admits a strong inversion (that is, an involution with one-dimensional fixed point set). It tursn out that this is not a necessary condition for L-spaces: Dunfield, Hoffman and Licata have very recently given examples of L-spaces with trivial symmetry group [DHL].

Collected Exercises

4.1 Lecture 1

In doing the exercises, you should take as given the properties of Floer homology stated in the lecture.

1. Let K be the unknot in S^3, and let $W = W_{-1}(K)$. Show that the map $F^-_{W,\mathfrak{s}_{\pm 1}} : HF^-(K) \to HF^-(K)$ is an isomorphism. What is the map $F^-_{W,\mathfrak{s}_{\pm k}}$ for other values of k?

2. (The blowup formula) Given a cobordism $W : Y_1 \to Y_2$, let $W' = W \oplus (-\mathbb{CP}^2)$, and let $E \in H_2(W')$ be the exceptional divisor. (That is, a sphere representing the class of a generator of $H_2(-\mathbb{CP}^2)$.) If $\mathfrak{s} \in Spin^c(W)$, show that for odd k there is a unique $\mathfrak{s}_k \in Spin^c(W')$ which agrees with $\mathfrak{s}$ on $W' - \nu(E)$, and which has $\langle c_1(\mathfrak{s}_k), [E] \rangle = k$. Use the first problem to show that $F^\circ_{W',\mathfrak{s}_{\pm 1}} = \pm F^\circ_{W,\mathfrak{s}}$. What happens if instead we take $W' = W \# \mathbb{CP}^2$?

3. If $S \subset Y$ is an embedded surface, let

$$HF^\circ(Y, S, k) = \bigoplus_{\{\mathfrak{s} \,|\, \langle c_1(\mathfrak{s}), [S] \rangle = k\}} HF^\circ(Y, \mathfrak{s}).$$

   Deduce the adjunction property in the special case where $\Sigma \cdot \Sigma = 0$, $g(\Sigma) > 0$ from the fact that $HF^\circ(S^1 \times \Sigma, \Sigma, k) = 0$ for $k > 2g(\Sigma) - 2$. (We'll see how to prove this in the last lecture.)

4. Use the blow-up formula and exercise 3 to prove the adjunction property for all Σ with genus > 0. (Hint: blow up repeatedly to produce a surface Σ' with self-intersection 0.)

5. Use adjunction to show that if $S \subset Y$ is an embedded surface of genus > 0, then $HF^\circ(Y, S, k) = 0$ whenever $k > 2g(S) - 2$. Deduce that there are only finitely many $\mathfrak{s}$ for which $HF^\circ(Y, s) \neq 0$.

6. Let $W_1 : S^3 \to S^1 \times S^2$ and $W_2 : S^1 \times S^2 \to S^3$ be the cobordisms given by addition of a 1 handle and a cancelling 2-handle respectively. Use what you know about grading shifts, together with the fact that $W_2 \circ W_1$ is the identity cobordism, to determine the maps $\widehat{F}_{W_1}$ and $\widehat{F}_{W_2}$. Similarly for the cobordisms $W_2' : S^3 \to S^1 \times S^2$ and $W_3 : S^1 \times S^2 \to S^3$ given by addition of a 2-handle and a cancelling 3-handle.

7. Given that $c_1(\mathfrak{s}|_{\partial W})$ is torsion, explain how to make sense of the quantity $c_1(\mathfrak{s})^2$appearing in the formula for the degree shift. (Note that a priori $c_1^2(\mathfrak{s})$ is a class in $H^4(W) = 0$.) Compute the degree shift associated to the map $F^\circ_(W, \mathfrak{s}_k)$, where $W = W_{-p}(K)$. Verify that if $\mathfrak{s}_{k_1}$ and $\mathfrak{s}_{k_2}$ restict to the same $Spin^c$ structure on ∂W then the difference in the corresponding degree shifts is an even integer.

8. Identify the S^1-equivariant analog of the second exact sequence of Property 8 . What about the spectral sequence?

4.2 Lecture 2

1. Show that, given a left-order $<$ on a group G, that

$$a \,\widetilde{<}\, b \iff a^{-1} < b^{-1}$$

for all $a, b \in G$ defines a right-invariant, strict total order on the elements of G. Give the definition of a right-orderable group and prove that every left-order on G is equivalent to a right-order (and vice versa). In particular, left- and right-orderable groups are equivalent.

2. Describe the positive cones implicit in Example 2.5 for any integer n (the case $n = 2$ is probably most instructive). What if the hyperplane contains some of the non-trivial group elements of $\mathbb{Z}^n$?

3. Prove that the existence of a positive cone for G is equivalent to a left-order on G. Any positive cone $\mathcal{P}$ gives rise to an opposite order: Take $\mathcal{P}^{-1}$ as positive cone instead. Compare the opposite order with (the positive cone for) the right-order $\widetilde{<}$ of Exercise 2.4

4. Write down a definition for $<$ on $\mathrm{Homeo}^+(\mathbb{R})$ as specified by $\mathcal{P}_X$ of Example 2.9.

5. ★ Fix a left orderable group G and denote by $LO(G)$ the set of all left orders on G. This becomes a topological space with the subbasis $U_a^b = \{<: a < b\}$. To study this topology further, consider a positive cone $\mathcal{P} \in LO(G)$ (compare Exercise 2.8) and prove that there is an inclusion of sets $LO(G)$ into $2^G = \{S : S \subseteq G\}$, the *power set* of G (considered as a set). A subbasis for a topology on 2^G is given by sets

$$U_a = \{S \subset G : a \in S\} \quad \text{and} \quad U_a' = \{S \subset G : a \notin S\}.$$

You can prove that the induced relative topology on $LO(G)$ is equivalent to the topology on $LO(G)$ described in terms of the subasis $\{U_a^b\}$. Using the fact that 2^G is a compact topological space (a consequence of Tychonoff's theorem), prove that $LO(G)$ is a compact topological space. Finally, prove that $LO(G)$ is a totally disconnected topological, that is, that the only connected components are singletons.

4.3 Lecture 3

1. Show that $\mathrm{Sym}^n\mathbb{C} = \mathbb{C}^n$ (Hint: consider the correspondence which assigns to a monic polynomial of degree n its roots.) Deduce that $\mathrm{Sym}^n\Sigma$ is a complex manifold. Show that $\mathrm{Sym}^n\mathbb{CP}^1 = \mathbb{CP}^n$.

2. Show that $\mathcal{D}(\phi\#\psi) = \mathcal{D}(\phi) + \mathcal{D}(\psi)$.

3. Show that $\epsilon(\mathbf{x}, \mathbf{y}) + \epsilon(\mathbf{y}, \mathbf{z}) = \epsilon(\mathbf{x}, \mathbf{z})$.

4. Let $\phi_\Sigma \in \pi_2(\mathbf{x}, \mathbf{x})$ correspond to the domain $\mathcal{D}_\Sigma$ which has multiplicity 1 everywhere in Σ. Show that $\mu(\phi_\Sigma) = 2$.

5. Let J be a complex structure on $\mathrm{Sym}^g\Sigma$ induced by a complex structure on Σ. Use the Riemann mapping theorem to show that if $\mathcal{D}(\phi)$ is a bigon, then $\#\overline{\mathcal{M}}(\phi) = \pm 1$.

6. Suppose $\mathcal{D}(\phi)$ is a convex rectangle. If φ is a holomorphic representative of ϕ, describe S_ϕ, p and π. Show that $\#\overline{\mathcal{M}}(\phi) = \pm 1$.

7. Using the Heegaard diagram of $S^3 - \nu(T)$ drawn in lecture, draw a Heegaard diagram for T_0 (0-surgery on the trefoil). List the generators and partition them into equivalence classes. What happens if we do a different surgery on T?

8. Show that $\widehat{HF}$ satisfies the Property 1 from the first lecture. (Hint: for the connected sum, put the basepoint in the connected sum region.)

9. Construct the exact sequences and spectral sequences of Property 8 from the first lecture.

10. Use Lipshitz and Lee's theorem to compute the differences in absolute grading for the generators of $\widehat{HF}(L(5,1))$. Do the same thing for $L(5,2)$.

4.4 Lecture 4

1. Show that a rational homology sphere Y is an L-space if and only if $HF_{\mathrm{red}}(Y)$ vanishes.

2. Show that there is a unique extension of M to a closed manifold $\mathbf{\Sigma}_K$ by attaching a solid torus $D^2 \times S^1$ in such a way that the resulting cover $\mathbf{\Sigma}_K \to S^3$ is one-to-one along the knot K.

3. ⋆ Use Heegaard Floer homology to prove that reduced Khovanov homology detects the trivial knot, assuming Conjecture 4.6. To do this, you will make use of the fact that (1) there is a spectral sequence with $E_2 = \widetilde{Kh}(K)$ and converging to $E_\infty = \widehat{HF}(-\mathbf{\Sigma}_K)$ [OS05b]; and (2) that the Poincaré homology sphere (indeed, any two-fold branched cover with finite fundamental group) is realised as a two-fold branched cover of S^3 in a unique way [Wat12]. In this case, the branch set is the $(3,5)$-torus knot; you might try and convince yourself of this.

4. Prove that alternating knots are quasi-alternating.

5. Irreducibility is required: If $HF_{\mathrm{red}}(Y) \neq 0$ then $HF_{\mathrm{red}}(Y \# Y') \neq 0$ but, on the other hand, if $\pi_1(Y')$ is not-left-orderable, then neither is $\pi_1(Y \# Y') \cong \pi_1(Y) * \pi_1(Y')$. Can you give an explicit example?

6. ⋆ Prove that the Weeks manifold (compare Exercise 2.3) is an L-space. Hint: The Weeks manifold is surgery on the Whitehead link.

7. Using Greene's presentation, complete the proof of Theorem 4.18.

8. Suppose that for any $d_1, \ldots, d_m \in \{0, +, -\}$, not all zero, the result of multiplying the i^{th} row of the matrix (ϵ_{ij}) by d_i has a non-zero column

with entries only in $\{+,-\}$. Prove that G is not left-orderable. Trick: Left-order G and choose d_i according to the signs of the generators.

9. ★ Given a strong Heegaard diagram for a strong L-space, show that the resulting group presentation satisfies the hypothesis of Proposition 4.21.

10. Find a pair of curves in the handlebody describing M [the twisted I-bindle over the Klein bottle] that avoid the α arc and intersect once transversally. Argue that this describes a basis for $H_1(\partial M\ bZ)$ and show that you can make your choice coincide with (φ, λ).

11. Prove that, as an element of $H_1(M;\mathbb{Z})$, $[\lambda]$ has order 2.

4.5 Lectures 5 and 6

1. Suppose (M,γ) is a sutured manifold and $D \hookrightarrow M$ is a properly embedded disk with $|\partial D \cap \gamma| = 2$. Show that the sutured manifold (M',γ') obtained by decomposing (M,γ) along D does not depend on which orientation we give D. (The manifold (M',γ') is said to be obtained from (M,γ) by a *disk decomposition.*) Use the decomposition formula to show that $SFH(M',\gamma') \simeq SFH(M,\gamma)$.

2. Prove the result of the first exercise directly from the definition of SFH.

3. A sutured manifold (M,γ) is *disk-decomposable* if it can be reduced to the trivial sutured manifold (B^3,γ_0) be repeated disk decompositions. Show (without appealing to Floer homology) that if (M,γ) is disk decomposable, it is a product manifold. Find a disk decomposition for the complement of the trefoil knot, and deduce that it is fibred. Do the same for each knot in S^3 with ≤ 7 crossings and monic Alexander polynomial.

4. Let K_n be the n-twist knot in S^3, and let (M_n,γ) be sutured manifold given by the complement of its standard Seifert surface with a longitudinal suture. Use a disk-decomposition to compute $SFH(M_n,\gamma_n)$.

5. As in the last problem, but using the pretzel knot $P(3,3,3)$. Draw an arc diagram for (M,γ) and use it to compute $SFH(M,\gamma)$.

6. Let K be the $(2,5)$ torus knot. Use the mapping cone to compute the homology of p surgery on K for $p = -3, 0, 3$.

7. Let $K \subset S^3$ be a knot, and let $C^+(n, K)$ be the mapping cone for n surgery on K. Show that if $n > 0$, the map $G^+_{n,\mathfrak{s}_k} : CF^+(S^3) \to C^+(n, K)$ vanishes on all elements of sufficiently high degree. Conversely show that if $n < 0$, show that it is an isomorphism in high degrees.

8. With notation as in the previous problem, compare the relative gradings of the images of $G_{-1,\mathfrak{s}_i}$ for differing values of i. Check that it is compatible with the expected shift in the absolute grading corresponding to the map $F^+_{W_{-1}(K),\mathfrak{s}_i}$.

9. Suppose $K \subset Y$ is a knot in a homology sphere, and let $g(K_0)$ be the minimal genus of a representative of $H_2(K_0)$. Show by example that we can have $g(K_0) < g(K)$. (Hint: start with a 2 component link whose algebraic linking number is 1 but whose geometric linking number is > 1.)

10. Suppose $K \subset S^3$ is a knot, and that K_n is an L-space for some $n > 0$. Show that each of the groups $\widehat{A}_i$ appearing in the mapping cone must be $\mathbb{Z}$. If each $\widehat{A}_i \simeq \mathbb{Z}$, show that K_n is an L -space if and only $n \geq 2g(K) - 1$.

Bibliography

[BGW13] Steven Boyer, Cameron McA. Gordon, and Liam Watson, *On L-spaces and left-orderable fundamental groups*, Math. Ann. **356** (2013), no. 4, 1213–1245.

[BH72] R. G. Burns and V. W. D. Hale, *A note on group rings of certain torsion-free groups*, Canad. Math. Bull. **15** (1972), 441–445.

[BMR77] Roberta Botto Mura and Akbar Rhemtulla, *Orderable groups*, Marcel Dekker, Inc., New York-Basel, 1977, Lecture Notes in Pure and Applied Mathematics, Vol. 27.

[BRW05] Steven Boyer, Dale Rolfsen, and Bert Wiest, *Orderable 3-manifold groups*, Ann. Inst. Fourier (Grenoble) **55** (2005), no. 1, 243–288. MR MR2141698 (2006a:57001)

[CD03] Danny Calegari and Nathan M. Dunfield, *Laminations and groups of homeomorphisms of the circle*, Invent. Math. **152** (2003), no. 1, 149–204.

[Cla10] Adam Clay, *Isolated points in the space of left orderings of a group.*, Groups Geom. Dyn. **4** (2010), no. 3, 517–532.

[DDRW02] Patrick Dehornoy, Ivan Dynnikov, Dale Rolfsen, and Bert Wiest, *Why are braids orderable?*, Panoramas et Synthèses [Panoramas and Syntheses], vol. 14, Société Mathématique de France, Paris, 2002.

[Deh94] Patrick Dehornoy, *Braid groups and left distributive operations.*, Trans. Am. Math. Soc. **345** (1994), no. 1, 115–150.

[DHL] Nathan M. Dunfield, Neil R. Hoffman, and Joan E. Licata, *Asymmetric hyperbolic L-spaces, Heegaard genus, and Dehn filling*, Preprint, arXiv:1407.7827.

[Dun11] Nathan Dunfield, *L-spaces and left-orderability of 3-manifold groups*, `http://ldtopology.wordpress.com/2011/07/26/`, 2011.

[Eps72] D. B. A. Epstein, *Periodic flows on three-manifolds*, Ann. of Math. (2) **95** (1972), 66–82.

[Far76] F.Thomas Farrell, *Right-orderable deck transformation groups.*, Rocky Mt. J. Math. **6** (1976), 441–447.

[GMM09] David Gabai, Robert Meyerhoff, and Peter Milley, *Minimum volume cusped hyperbolic three-manifolds.*, J. Am. Math. Soc. **22** (2009), no. 4, 1157–1215.

[Gre] Joshua Evan Greene, *Alternating links and left-orderability*, Preprint, arXiv:1107.5232.

[Gre13] Joshua Evan Greene, *A spanning tree model for the Heegaard Floer homology of a branched double-cover*, J. Topol. **6** (2013), no. 2, 525–567.

[Hal50] M. Hall, *Review of Vinogradov's paper* On the free product of ordered groups, MathSciNet Mathematical Reviews, 1950.

[Hat] Allen Hatcher, *Notes on basic 3-manifold topology*, available electronically at `http://www.math.cornell.edu/~hatcher`.

[HS85] James Howie and Hamish Short, *The band-sum problem*, J. London Math. Soc. (2) **31** (1985), no. 3, 571–576.

[KM11] P. B. Kronheimer and T. S. Mrowka, *Khovanov homology is an unknot-detector*, Publ. Math. Inst. Hautes Études Sci. (2011), no. 113, 97–208. MR 2805599

[Lin99] Peter A. Linnell, *Left ordered amenable and locally indicable groups.*, J. Lond. Math. Soc., II. Ser. **60** (1999), no. 1, 1.

[LL] Adam Simon Levine and Sam Lewallen, *Strong L-spaces and left orderability*, Preprint, arXiv:1110.0563.

[LOTa] Robert Lipshitz, Peter Ozsváth, and Dylan Thurston, *Bordered Heegaard Floer homology*, Preprint, arXiv:0810.0687.

[LOTb] ———, *Notes on bordered Floer homology*, Preprint, arXiv:1211.6791.

[LOT11] Robert Lipshitz, Peter S. Ozsváth, and Dylan P. Thurston, *Tour of bordered Floer theory*, Proc. Natl. Acad. Sci. USA **108** (2011), no. 20, 8085–8092. MR 2806643 (2012e:57027)

[LS07] Paolo Lisca and András I. Stipsicz, *Ozsváth-Szabó invariants and tight contact 3-manifolds. III*, J. Symplectic Geom. **5** (2007), no. 4, 357–384.

[LW14] Yu Li and Liam Watson, *Genus one open books with non-left-orderable fundamental group*, Proc. Amer. Math. Soc. **142** (2014), no. 4, 1425–1435.

[Min12] Igor Mineyev, *Groups, graphs, and the Hanna Neumann conjecture.*, J. Topol. Anal. **4** (2012), no. 1, 1–12.

[Nav10] Andrés Navas, *On the dynamics of (left) orderable groups.*, Ann. Inst. Fourier **60** (2010), no. 5, 1685–1740.

[OS05a] Peter Ozsváth and Zoltán Szabó, *On knot Floer homology and lens space surgeries*, Topology **44** (2005), no. 6, 1281–1300. MR MR2168576 (2006f:57034)

[OS05b] ———, *On the Heegaard Floer homology of branched double-covers*, Adv. Math. **194** (2005), no. 1, 1–33.

[Pas77] Donald S. Passman, *The algebraic structure of group rings*, Pure and Applied Mathematics, Wiley-Interscience [John Wiley & Sons], New York-London-Sydney, 1977.

[Pet] Thomas Peters, *On L-spaces and non left-orderable 3-manifold groups*, Preprint, arXiv:0903.4495.

[Poi10] Henri Poincaré, *Papers on topology*, History of Mathematics, vol. 37, American Mathematical Society, Providence, RI; London Mathematical Society, London, 2010, ıt Analysis situs and its five supplements, Translated and with an introduction by John Stillwell.

[Rol76] Dale Rolfsen, *Knots and links*, Publish or Perish, Inc., Berkeley, Calif., 1976, Mathematics Lecture Series, No. 7.

[RR02] Akbar Rhemtulla and Dale Rolfsen, *Local indicability in ordered groups: braids and elementary amenable groups*, Proc. Amer. Math. Soc. **130** (2002), no. 9, 2569–2577 (electronic).

[RSS03] R. Roberts, J. Shareshian, and M. Stein, *Infinitely many hyperbolic 3-manifolds which contain no Reebless foliation*, J. Amer. Math. Soc. **16** (2003), no. 3, 639–679 (electronic).

[Sco73] G. P. Scott, *Compact submanifolds of* 3*-manifolds*, J. London Math. Soc. (2) **7** (1973), 246–250.

[Sco83] Peter Scott, *The geometries of* 3*-manifolds*, Bull. London Math. Soc. **15** (1983), no. 5, 401–487.

[Sik04] Adam S. Sikora, *Topology on the spaces of orderings of groups.*, Bull. Lond. Math. Soc. **36** (2004), no. 4, 519–526.

[Vin49] A. A. Vinogradov, *On the free product of ordered groups*, Mat. Sbornik N.S. **25(67)** (1949), 163–168.

[Wat] Liam Watson, *Heegaard Floer homology solid tori*, In preparation; notes available from `http://www.maths.gla.ac.uk/~lwatson/`.

[Wat09] ———, *Involutions on 3-manifolds and Khovanov homology*, PhD, Université du Québec à Montréal, 2009.

[Wat12] ———, *Surgery obstructions from Khovanov homology*, Selecta Math. (N.S.) **18** (2012), no. 2, 417–472.