Closed Smooth 4-Manifolds **Spring 2025**

Part I

Abhishek Shivkumar
Professor Lisa Piccirillo Illustrations by Elise Brod

Overview

Lecture 1: January 14th

All our manifolds will be 4-dimensional, compact, oriented, and smooth unless stated otherwise. There is a thermodynamic situation in the problem of classifying manifolds: in dim ≤ 3 (the solid phase), the problem is solved, and in dim ≥ 5 (the gaseous phase), the problem in some sense is not solvable, but there are tools such as surgery theory which give partial results (for instance, simply-connected closed 5-manifolds are classified). However (for instance) the Poincaré conjecture is open in (say) dimension 126. In dimension 4 (the phase boundary) there is no conjectural classification, and the tools that work in higher dimensions generally fail.

How can we come up with a conjectural classification? The basic strategy is to look for a set of "features" these manifolds have, and check whether they completely determine the manifolds in question. The first such feature to consider is the fundamental group. If we restrict to $\pi_1 = 1$, we may ask whether this is complete; evidently, it is not, as H_2 and its intersection form Q_X are not detected by π_1. Per Freedman, these additional data are sufficient for classifying simply connected closed 4-manifolds up to homeomorphism; per Donaldson, these data are insufficient in the smooth category via gauge/Floer theoretic invariants. These gauge and Floer theoretic data also cannot complete our classification per Ren-Willis[1] (at least for manifolds with boundary).

If we allow for manifolds with boundary, conjecturing a classification is more difficult as, in addition to the features of the manifold itself, the features of the boundary and the features of $\iota : \partial X \hookrightarrow X$ are also essential to consider, which results in complication that we will seek for the moment to avoid.

[1] Ren and Willis, *Khovanov homology and exotic 4-manifolds*.

Returning for a moment to Donaldson's results, which give multiple smooth structures corresponding to *some* intersection forms Q_X; it is then also interesting to ask whether Q_X is insufficient to determine X (smoothly) for *all* Q_X. That is, are there exotic copies of all closed, smooth 4-manifolds? This will be one of our guiding questions in this course.

How To Construct Exotica

The procedure for building exotica is as follows:

Step I Build candidate manifolds X, X' (one often gets X for free)

Step II Show that $X \underset{\text{top}}{\cong} X'$

Step III Show that $X \underset{\text{sm}}{\not\cong} X'$

Step 2. usually involves calculating the classical algebraic invariants of the manifolds, step 3. usually involves calculating some non-classical gauge/Floer/lasagna-theoretic invariants. We will primarily be focused on step 1. in this class, though all three can be quite difficult.

The outline for the course is as follows:

Part I We will begin with constructive basics, in the setting of manifolds with boundary. We will introduce Floer TQFT methods and build some exotic homotopy S^2s (i.e., knot traces) and B^4s.

Part II We will apply Floer TQFT methods to closed manifolds, develop the theory of Lefschetz fibrations (these give Donaldson's original exotic construction of $\mathbb{CP}^2\#_9\overline{\mathbb{CP}^2}$). We will consider various surgery operations such as torus surgeries and knot surgeries, and cover the relevant TQFT gluing formulae relevant to these surgeries. We will sketch many many examples of exotica (down to $\mathbb{CP}^2\#_5\overline{\mathbb{CP}^2}$), and discuss Lefschetz fibration geography.

Part III We will sketch the record-setting construction of an exotic $\mathbb{CP}^2\#_3\overline{\mathbb{CP}^2}$ due to Akhmedov-Park,[2] and the construction of small $\sigma = 0$ exotica due to Baykur-Hamada.[3]

[2] Akhmedov and Park, *Exotic smooth structures on small 4-manifolds*.

[3] Baykur and Hamada, *Exotic 4-manifolds with signature zero*.

Framings

An n-dimensional k-handle is D^n thought of as $D^k \times D^{n-k}$ (with corners appropriately smoothed), which we attach to some manifold X via a smooth embedding $f : \partial D^k \times D^{n-k} \to \partial X$.

Lemma 1.2.1

$X \cup_f$ handle is well-defined up to isotopy of f.

The essential data of a handle attachment is where $\partial D^k \times \{0\}$ goes (called the *attaching sphere*) in ∂X. The attaching map turns out to constitute only a mild decoration of this data, corresponding to a choice of trivialization for the tubular neighborhood of the attaching sphere.

Definition 1.2.2: Framings

For a $k-1$-sphere $S \hookrightarrow Y^{n-1}$ with trivial normal bundle, a *framing* is a choice of diffeomorphism (or bundle isomorphism) $f : S^{k-1} \times D^{n-k} \to \nu(S)$.

From a handle decomposition, by deformation retracting every handle onto its core, we obtain a CW complex homotopy equivalent to it; this process is insensitive to framings i.e. is determined by the data of the attaching spheres alone.

A framed sphere is enough info to attach a handle. With the title of our class in mind, we don't need to be coy and work in generality; set $n = 4$ and $k = 2$, so we are interested in embeddings $K : S^1 \hookrightarrow Y^3$ and framings thereof.

Lemma 1.2.3

The set of framings in this setting (up to isotopy of the framing map) is *non-canonically* in bijection with $\mathbb{Z} = \pi_1(\mathrm{SO}(2))$.

PROOF : Choose any framing f and declare $\varphi(f) = 0$; for any other framing g, consider $f^{-1} \circ g : S^1 \times D^2 \to S^1 \times D^2$. Using the fact that the mapping class group $\mathrm{Diffeo}^+(S^1 \times D^2)/\text{isotopy}$ is isomorphic to $\mathbb{Z}$ and generated by the Dehn twist element τ, set $\varphi(g) = [f^{-1} \circ g] \in \mathbb{Z}$. φ is our desired bijection. ■

In general, the set of framings for an n-dimensional k-handle are a $\pi_{k-1}\,\mathrm{O}(n-k)$ torsor.

Exercise 1.2.4

1. Show that φ above is a bijection
2. Show that f is determined by $f(S^1 \times \{1\})$
3. Bonus: prove that handle attachment is determined by the isotopy class of the attaching map
4. Bonus: exhibit the relationship between framings in this setting and $\pi_1(\mathrm{SO}(2))$

This is not a particularly satisfactory situation (e.g., 3 is non-canonically 7 in torsor land); fortunately, we can do a little better if we assume that our attaching circles are nullhomologous.

Definition 1.2.5: Meridians and Longitudes

For a (classical) knot $K : S^1 \to Y^3$, the *meridian* of K is $\mu_K : S^1 \hookrightarrow \partial\nu(K)$ the boundary of a disc fiber of the normal bundle. A *longitude* is $\lambda_K \to S^1 \to \partial\nu(K)$ such that $[\lambda_K]$ generates $H_1(\nu(K))$ (equivalently, λ_K intersects μ_K geometrically once).

Corollary 1.2.6

Framings are determined by a choice of longitude.

Lisa says that this is the only foolproof notion of framing, and what you should return to if/when you become confused.

Definition-Proposition 1.2.7

If K is nullhomologous in Y, then there is a unique longitude λ_0 in $\partial\nu(K)$ s.t. $[\lambda_0] = 0 \in H_1(Y \setminus \nu(K))$, called the *Seifert longitude*. Thus, for K nullhomologous, we have a canonical bijection for framings.

Exercise 1.2.8

Prove the above proposition. As a bonus, show that if $\Gamma^2 \hookrightarrow Y$ is a surface with $\partial\Gamma = Y$, show that $\lambda_n \cap \Gamma$ in n points (where λ_n is the longitude with $\varphi(\lambda_n) = n$).

Lecture 2: January 16[th]

Proof : We will invoke without proof a theorem of Thom:

Theorem 1.2.9: Thom

If $n \leq 4$, then any class in $H_*(X^n)$ can be represented by a submanifold. Moreover, if $Y^m \hookrightarrow X^n$ with $m \leq n-1$, with $[Y] = 0 \in H_m(X)$ then $Y = \partial Z^{m+1} \hookrightarrow X$.

To see the result from this, since $[K] = 0$, there exists a Seifert surface $\Sigma^2 \hookrightarrow Y$ with $\partial\Sigma = K$. A pushoff of K into Σ gives λ with $[\lambda] = 0 \in H_1(Y \setminus \nu(K))$.

To prove uniqueness of λ, we claim that $[\mu_K]$ is infinite order in $H_1(Y \setminus \nu(K))$. Suppose there exists λ' with $[\lambda'] = 0$, then $[\lambda - \lambda'] = [0]$, and $[\lambda - \lambda'] = n[\mu_K] \in H_1(Y \setminus \nu(K))$ so $n = 0$ and $\lambda = \lambda'$.

To see that $[\mu_K]$ is non-torsion, note that 3-manifolds have a bilinear intersection pairing $H_1(Y) \times H_2(Y) \to \mathbb{Z}$ given by counting the number of intersections; this implies that meridians are infinite order. To see this, note that if $n[\mu_K] = 0 \in H_1(Y \setminus \nu(K))$, then $n[\mu_K]$ bounds a surface $\Gamma \hookrightarrow Y \setminus \nu(K)$, then in Y we can obtain a new surface Γ' by gluing in the obvious disc bounded by each of the n parallel copies of μ_K. In Y, $|\Gamma' \cap K| = n$ but $[K] = 0 \in H_1(Y)$ so $n = 0$. ■

One can also prove this result using Mayer-Vietoris. Lisa remarks that in other settings, there might be a different preferred longitude given by some additional data (e.g. a surface the knot lies on).

Definition-Proposition 1.2.10: Linking Pairing

For $K, K' \subseteq Y$ a pair of nullhomologous knots, we may consider $K' \subseteq Y \setminus \nu(K)$. We claim that $H_1(Y) = H_1(Y \setminus \nu(K))/\langle \mu_K \rangle$. Since $[K'] = 0 \in H_1(Y)$, $[K'] = n[\mu_K] \in H_1(Y \setminus \nu(K))$. Thus, we may define the *linking pairing* $\mathrm{lk}(K, K') := n$.

Exercise 1.2.11

Show that $H_1(Y) = H_1(Y \setminus \nu(K))/\langle \mu_K \rangle$. There are two bonus exercises: first, give a diagrammatic definition of lk for knots in S^3 or (double bonus) arbitrary 3-manifolds; second, prove that $\mathrm{lk}(K, K') = \mathrm{lk}(K', K)$.

One can also do this with Mayer-Vietoris or van Kampen.

Proof : $Y = Y \setminus \nu(K) \cup (S^1 \times D^2)$. We may decompose $S^1 \times D^2$ as a 3-ball (A)

with one 1-handle attached (B), so, instead, we can write

$$Y = (Y \setminus \nu(K) \cup_g A) \cup B$$

Gluing in A along g is essentially just gluing in a disc along its boundary, and B gets glued in along its entire S^2 boundary (which does not affect H_1). Gluing in a disc picks us up a relator in homology, so, $H_1(Y) = H_1(Y \setminus \nu(K))/\langle g(\gamma)\rangle$ where γ is the boundary of the disc, and $g(\gamma) = \mu_K$ by construction. ■

By convention X denotes a 4-manifold, Y a 3-manifold, Σ and Γ 2-manifolds. The mnemonic for this is to count the number of leaves of the graph defined by each letter; X has 4 leaves, Y, 3, etc.

Suppose X is a 2-handlebody (which, for our purposes, means that it is built from a single 0-handle and some collection of 2-handles) described as (K_i, f_i).

Question 1.2.12

What are the λ_{f_i}? Why do we care, if we have this apparently simpler description in terms of an integer framed link diagram?

Exercise 1.2.13

Suppose X is a 2-handlebody with two 2-handles (K_1, f_1) and (K_2, f_2). Suppose further that we do some slide of the K_2 handle over K_1. What is the resulting attaching circle for the new handle diagram? Bonus: prove that the effect on the boundary of 2-handle attachment is λ-framed Dehn surgery.

Note that λ_{f_1} bounds a disc in the *new* boundary created by attaching h_1, i.e., in

$$\partial(X \cup h_1) = (\partial X \setminus S^1 \times D^2) \cup D^2 \times S^1$$

The moral of this story is that it pays to know where the non-obvious discs are; this will come up repeatedly in this course.

PROOF : K_2' is a band sum of K_2 with λ_{f_1}. The point is that a handle slide is just an isotopy, so we're taking a bight of K_2 and sliding it over some disc in the boundary. Since λ_{f_1} is the image of $\partial D^2 \times \{1\}$ under the attaching map for the 2-handle, it bounds the embedded disc $D^2 \times \{1\} \hookrightarrow \partial(X \cup h_1)$; this is the disc we are sliding over. The effect of this isotopy is a band sum. ■

Note that, in general, our attaching circles may not appear nullhomologous in our handle diagram (and so our framing numbers *a priori* do not make sense); the convention for 2-handlebodies is that we are considering the framing number for $K_i \subseteq S^3$ (i.e., ignore everything but the one knot in question).

Lecture 3: January 23rd

Recall *Seifert's algorithm*, which from a diagram of a knot K in S^3 produces an orientable surface $\Sigma \hookrightarrow S^3$ with $\partial\Sigma = K$.

Definition 1.2.14: Blackboard Framing

For a knot diagram, the blackboard framing λ_{bb} is the longitude that lies in the blackboard.

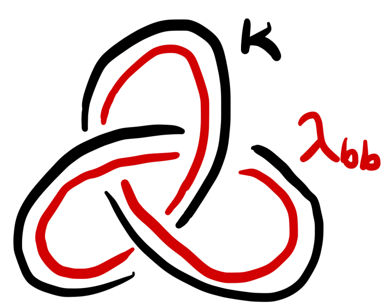

Figure 1.1: The blackboard framing for the standard drawing of the right-handed trefoil.

Exercise 1.2.15

Pick any nontrivial knot, and draw both λ_{bb} and λ_0 (by running Seifert's algorithm). Prove that $\lambda_0 = \lambda_{\text{bb}} - \omega$ where ω is the *writhe* of the diagram, i.e., the number of positive minus the number of negative crossings.

Positive Crossing

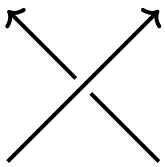

Negative Crossing

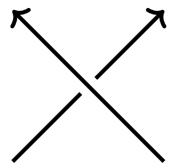

The right hand rule is useful for remembering the signs of crossings.

As a bonus, consider how integer framings change under handle slides (the relevant players are f_1, f_2 and $\text{lk}(K_1, K_2)$).

Proof : To convince yourself of the formula, it suffices to consider the various framings at a single crossing. The Seifert longitude is obtained by pushing λ_{bb}

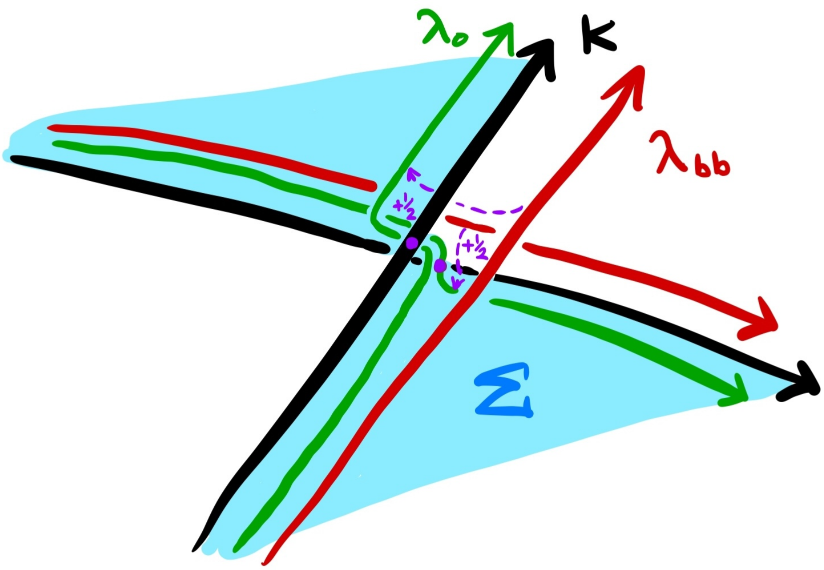

Figure 1.2: Pushing λ_{bb} onto the Seifert surface.

onto the Seifert surface Σ whose boundary is K; each $\pm$ crossing of the diagram of K induces two $\mp\frac{1}{2}$ twists, which gives us the desired formula after summing over all crossings. ∎

Dotted Circles

A 4-dimensional 1-handle is $D^1 \times D^3$ attached along $S^0 \times D^3$ which is just a pair of 3-balls. One way to draw a 1-handle attachment, then, is just to draw a pair of balls in the handle diagram. This has some downsides: for one, this is just cumbersome when there's multiple 1-handles since they need to be labeled; for another, this does not automatically give us a Dehn surgery presentation of the boundary which we had for 2-handlebodies. Thankfully, there is a remedy.

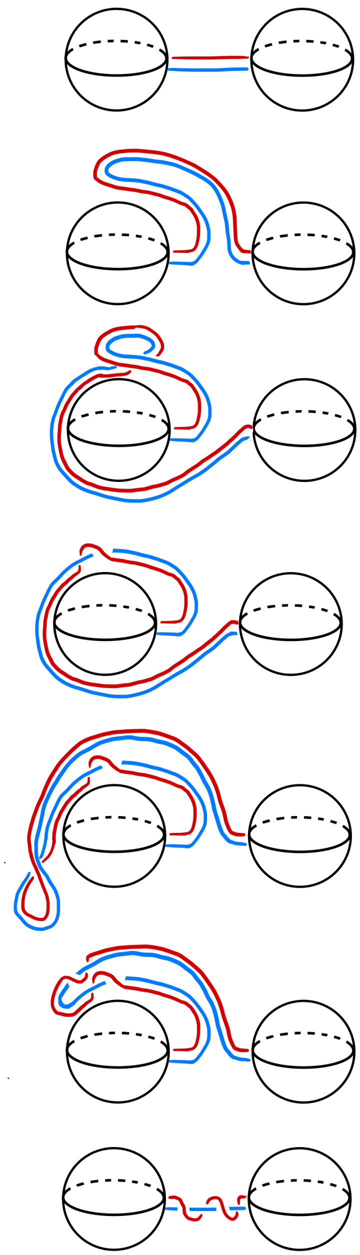

Figure 1.3: Planar isotopy that changes 2-handle framing by 2. Since this is a picture of a cancelling pair, this shouldn't be too surprising; as we will discuss below, 1-handle attachment does 0-framed Dehn surgery on the boundary, so this is essentially a drawing of the standard fact that the $(0,0)$-framed Hopf link is the same (as a handle diagram) as the $(0,2n)$-framed Hopf link.

Remark 1.3.1

One objection to the "pair of 3-balls" notation is that 2-handle framings are not well-defined as in Figure 1.3, where the 2-handle framing can change by ± 2 via planar isotopy. In fact, this is only a mental problem, and it is perfectly fine for framings to change along a planar isotopy.

Proposition 1.3.2

If a 2-handle h_2 runs over a 1-handle h_1, such that the attaching sphere of h_2 *geometrically* intersects the belt sphere of h_1 in a single point, then h_1 and h_2 are a *cancelling pair*, i.e., $X \cup h_1 \cup h_2 = X$.

One mental model for this is just a thickened disc filling in a genus, but this is too simple for our 4-dimensional needs. The point will of course be that a dotted circle meridian to any 2-handle (with any framing) forms a cancelling pair.

Lemma 1.3.3

Let $Y^3 \hookrightarrow \partial Z^4$, then $Z \cup_\tau (Y \times I) \underset{\text{sm}}{\cong} Z$ where $Y \times \{0\} \xrightarrow{\tau} \partial Z$ is the embedding we already had.

PROOF : Essentially, we are just attaching reverse collar neighborhoods only to $Y \subseteq \partial Z$. Taking an existing collar neighborhood of Y, we can build a diffeomorphism from Z that is the identity away from this collar and just stretches $Y \times I$ out by a factor of 2 on the collar. ■

Lemma 1.3.4

For any manifold M,

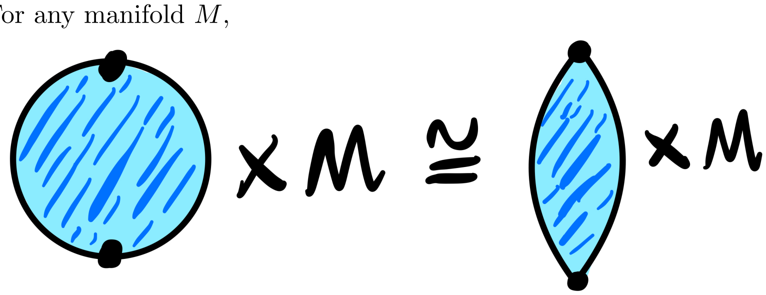

This is barely a lemma and does not deserve a proof.

Exercise 1.3.5

Prove the above proposition using the above two lemmas. As a bonus, using your proof, give a handle diagram for the diagram in the margin after cancelling the appropriate handles.

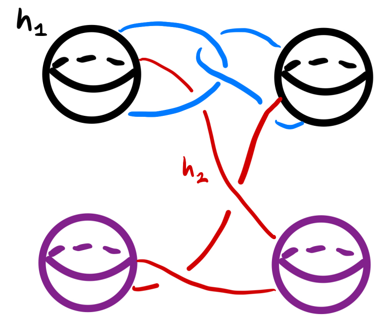

Figure 1.4: A tricky handle diagram.

PROOF : $h_2 = D^2 \times D^2$ attached along a thickened S^1; the point of the second lemma above is that we should (here) think of a disc bounded by a circle as guiding an interpolation between two paths (given by splitting the circle into two arcs) rather than as guiding a nullhomotopy.

In our case, we want to partition our attaching circle for our 2-handle into two arcs, one of which is visible in the handle diagram, and the other runs over the 1-handle (see below). Then we use the non-obvious disc between these two arcs (the core disc of the 2-handle) to guide an isotopy that places the two arcs almost parallel, and makes the 1-handle into a bulge on the boundary.

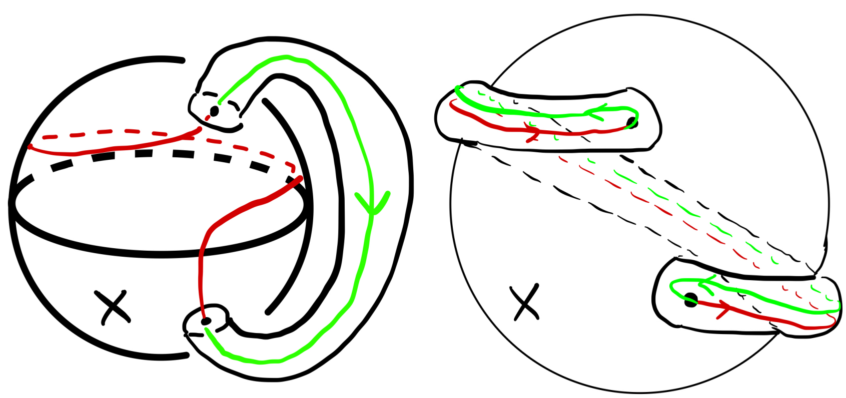

Figure 1.5: Guiding an isotopy using a hidden disc.

We may then apply the first of the two above lemmas (bulges don't matter) to finish. ■

Another example of it paying to know where your hidden discs are.

Thus $X \cong X \cup h_1 \cup h_2 \cong X \cup (D^3 \times D^1)$; if we want to recover $X \cup h_1$ from this, we need to remove a thickened disc. This is the key observation that leads to dotted circle notation for 1-handles: that 1-handle attachment is the same as deleting a disc.

Lecture 4: January 28th

To make this explicit, we consider $X \cong X \cup (D^3 \times D^1)$ with a cancelling pair attached, thought of as a 1-handle "lying down" as in the above picture. In particular, thinking of $D^3 \times D^1$ as a 1-parameter family (indexed by time) of 3-balls, glued to ∂X along $D^2 \times D^1$ (the $D^2 \subseteq D^3$ being one half of the boundary sphere), we may obtain $X \cup h_2$ by deleting a judiciously chosen disc (see Figure 1.6). In this setup, the evident disc to delete is $D_0 = D^2 \times \{0\} \subseteq D^3 \times [-1, 1]$. Note that $\partial D_0 \subseteq \partial X$ is unknotted (i.e., bounds a disc in ∂X) by construction, and this disc is boundary parallel. ∂D_0 in fact bounds D_0, so this belabors the point a little since

Recall that a submanifold is boundary parallel if there is an ambient isotopy (fixing the boundary) taking the submanifold to the boundary.

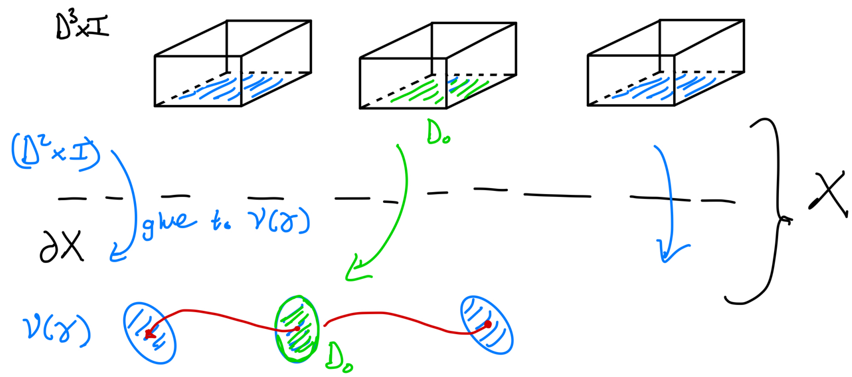

Figure 1.6: We can produce the correct dimensional "hole" (as is created by attaching a 1-handle) by deleting an appropriate neighborhood of D_0.

D_0 automatically lies in ∂X but these are precisely the criteria for a circle in the boundary to define a unique disc, and therefore prescribe a 1-handle attachment.

Exercise 1.3.6

We know that $X \cup h_1 \cong X \natural S^1 \times B^3$. Where is $S^1 \times \{\bullet\}$ in the dotted circle notation (where $\bullet \in \partial B^3$)? How can we convert from the dotted circle notation back to standard notation?

My answer to this was that $S^1 \times \{\bullet\}$ is where we would attach the 2-handle that cancels the 1-handle, so it must be a meridian of the dotted circle. This is a little circular, and you can come up with the same answer by somewhat carefully studying the above diagram. The standard fact that the boundary of $(D^2, S^1) \cap (D^2, S^1) \subseteq (D^4, S^3)$ intersecting transversely in a single point has boundary the Hopf link may be relevant.

Proposition 1.3.7

The effect on the boundary of 1-handle attachment is 0-framed Dehn surgery on the dotted circle (*not* the attaching circle, which is S^0). Bonus: prove this using $Q_{S^4} = (\,0\,)$.

PROOF : Since the effect of attaching a handle is local, we will study the effect of 1-handle attachment on B^4. $\partial(B^4 \cup h_1) = \partial(B^4 \setminus \nu(D_0)) = (S^3 \setminus \nu(U)) \cup D_0 \times S^1$ so deleting a disc evidently results in Dehn surgery, and it remains only to determine the surgery coefficient. Observe that we know that $\partial(B^4 \cup h_1) = S^1 \times S^2$, so $H_1(B^4 \cup h_1) = \mathbb{Z}$. This will imply that the surgery coefficient is 0.

Recall that $H_1(S^3) = 0 \cong H_1(S^3 \setminus \nu(U))/\langle \mu_0 \rangle$ and μ_0 is infinite order; the key point is that $H_1(S^3_{\frac{p}{q}}(U)) = \langle \mu_0 \rangle / \langle p\mu_0 + q\lambda_0 \rangle = \mathbb{Z}/p\mathbb{Z}$, so we know that $p = 0$. Then $q = 1$ since other values of q prescribe attaching S^1 to disjoint copies of circles ($q = -1$ is disallowed to preserve orientation). ■

For the bonus, note that we may split S^4 into two copies of B^4 glued along their common S^3 boundary. In one of these copies of B^4, we may remove $D_0 \times D^2$, i.e., attach a 1-handle, and glue this $D_0 \times D^2$ to the other B^4. Taking the standard surface associated to this 2-handle (core disc union Seifert surface), this surface must have self-intersection 0 since $Q_{S^4} = (\,0\,)$, so the framing of the 2-handle curve is 0. But since $\partial(B^4 \setminus D_0 \times D^2) = \partial(B^4 \cup D_0 \times D^2)$, the surgery coefficient for 1-handle attachment is 0.

This also gives an argument that no 4-manifolds (with boundary) with nonzero intersection form can embed in S^4.

Proposition 1.3.8

We can slide 2-handles "over" (some say "under") 1-handle dotted circles.

Since the effect on the boundary of a dotted circle is 0-framed Dehn surgery, sliding over a 1-handle has to be identical to sliding over a 0-framed 2-

handle. The effect on framing under such a slide is also given by pretending dotted circles are 0-framed 2-handles.

Remark 1.3.9

There are many proscriptions for dotted circle notation. You can only dot unknots (a knot, even a slice knot, will not prescribe a unique disc). You cannot slide a dotted circle over over another dotted circle in general as the result may no longer bound a boundary parallel disc (i.e. it may no longer be unknotted). Dotted circles may not be linked (this screws up the discs).

Relative Handle Calculus

Lecture 5: January 30^{th}

Exercise 1.4.1

Build simply connected closed 4-manifolds out of at most (say) six handles (including the 0 and 4-handles) that are not obviously $\#^k S^2 \times S^2$ or $m\mathbb{CP}^2 \# n\overline{\mathbb{CP}^2}$. Can you do it with $b_2 = 0$? $b_2 = 4$?

The point of this exercise is that constructions are quite hard, and this is a big problem in the field. The two dual problems are that it's hard to close things up, and that when we can, closing tends to collapse everything back to one of our standard examples.

For intersection forms of size less than 8, we also know from number theory that they are a direct sum of a ± 1 diagonal matrix plus some hyperbolic forms $\left(\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right)$, which would correspond topologically to $\#^k S^2 \times S^2 \# m\mathbb{CP}^2 \# n\overline{\mathbb{CP}^2}$, so this exercise was to produce small exotica explicitly from handles.

Relative handle decompositions allow us to build interesting closed manifolds out of handles by eliminating the problem of closing things up. The idea is to build two 4-manifolds X and X' with the same boundary and then glue them together to get a closed 4-manifold. The tricky part is that we have to turn X' "upside down" in order to have a unified handle diagram for $X \cup_\partial X'$. An alternative way to think of these data is to consider the handle diagram of X' relative to that of X.

A relative handle diagram starts from some 3-manifold Y (the common boundary of X and X' in the above setup), which we stretch to a 4-manifold $X = Y \times I$, which we then add more handles to. The left boundary is $\partial^- X = Y$, and $\partial^+ X$ depends on the handles we add. To describe X we have to describe the attaching regions of handles in Y. For example, if we have a Heegaard splitting for Y, we may draw our 2-handle attaching curves on the pair of handlebodies. Integer framings are tricky in this setting, so we need to provide an explicit longitude (see Figure 1.7 below). We may also draw 1-handles as dotted circles provided they bound a disc in Y (i.e. they cannot wrap the topology of the handlebodies).

We may also start with a Dehn surgery presentation for Y and attach 4-

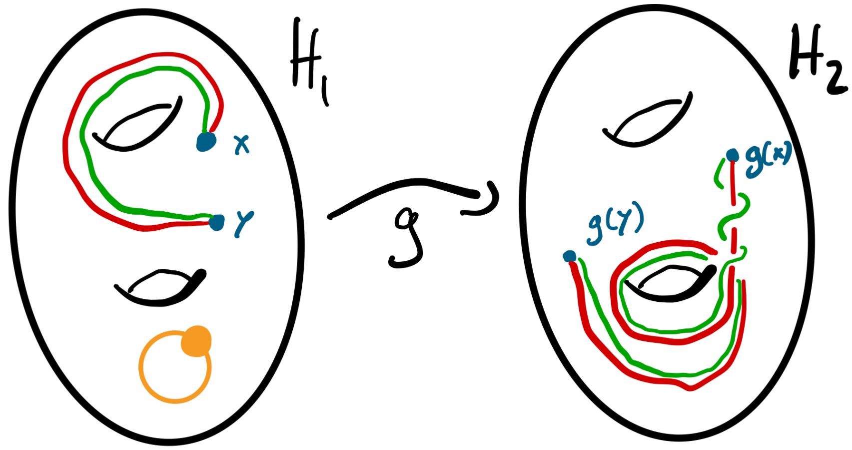

Figure 1.7: A relative handle diagram starting from a genus 2 Heegaard splitting of Y. The red curve is the attaching curve for a 2-handle, the green curve is its framing longitude. The dotted circle is appropriately contained in a 3-ball.

dimensional handles from there. The Dehn surgery curves are distinguished from the handle attaching curves by bracketed framings. We can slide the 4-manifold 2-handles over the bracketed handles but not vice versa, as the bracketed handles give us Y, and Y cannot be modified in this setting. We may also slide the bracketed handles over each other to modify our Dehn surgery presentation for Y.

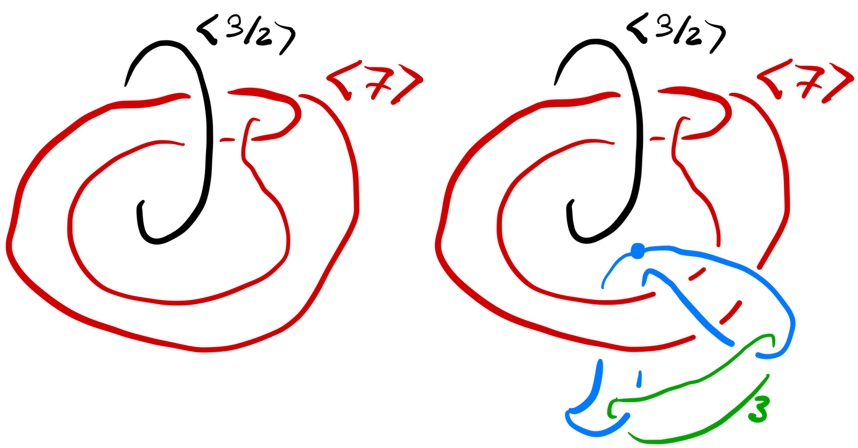

Figure 1.8: A relative handle diagram from a Dehn surgery presentation. Y is on the left, before the handle attachments, and X is on the right.

Exercise 1.4.2

Let X be the 4-manifold described by the relative handle diagram in Figure 1.9. Turn X upside down, and then draw a simpler relative handle diagram for X. Bonus: build three other manifolds X_i with $\partial X_i = \partial X$.

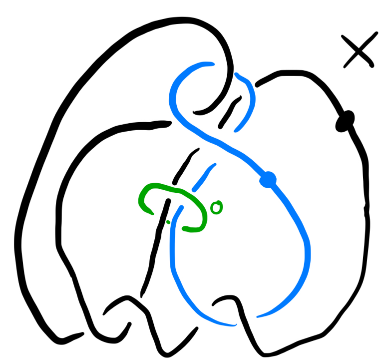

Figure 1.9: X, which will imminently be upside down

The 1-handles upside down become 3-handles so we don't need to draw them (using Laudenbach-Poénaru — 3-handle attachment is unique). A 2-handle upside down becomes a 0-framed meridian to its attaching curve, so a relative handle diagram for X upside down is the first of the following diagrams:

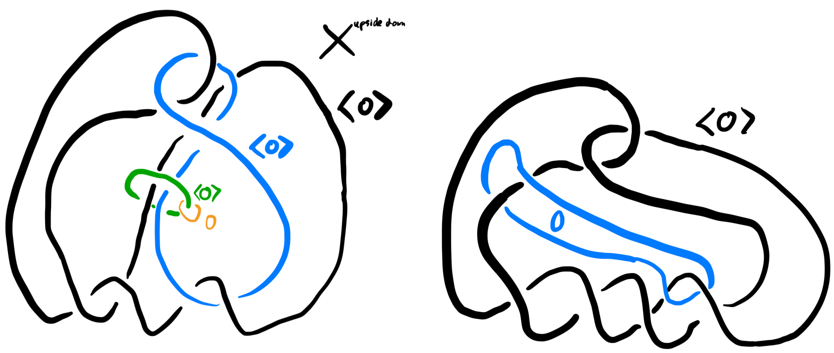

Figure 1.10: Two relative handle diagrams for X turned upside down, 3 and 4-handles not depicted.

Our original handles for X are now surgery curves for $\partial X = Y = \partial X^{\text{upside down}}$. In order to draw a simpler diagram for X, we need the following fact:

One can also derive this lemma using the slam-dunk move from Kirby calculus.

Lemma 1.4.3

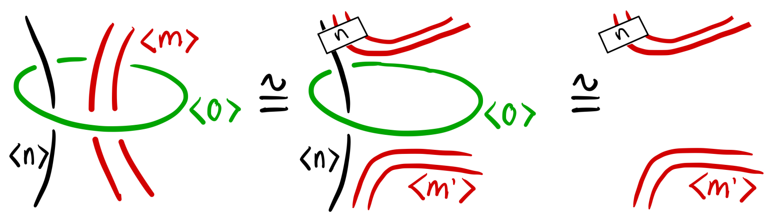

This holds for 3-manifold surgery presentations (as indicated by the bracketed framings), not for 4-manifold handle diagrams, and follows by sliding the red strands over the black strand and then cancelling black with green (this makes sense if we think of our 3-manifold as the boundary of a 4-manifold, where the black strand is a 2-handle and the green circle is a cancelling 1-handle). We apply this lemma to the above diagram by sliding black, then yellow over blue (so that neither links green), then cancelling blue and yellow. The result of these slides is the second diagram in Figure 1.10 (where the yellow 2-handle has been redrawn in blue).

For the bonus, one option is to swap $\langle 0 \rangle$ and 0 in our simplified diagram. I'm not quite sure how to come up with others.

Now, if we can find X' and $f : \partial X \xrightarrow{\sim} \partial X'$, we can form $Z = X \cup_f X'$, a closed 4-manifold using the following lemma:

Lemma 1.4.4

Given $\varphi : Y \xrightarrow{\sim} Y'$,

$$(Y \times I) \cup_A \text{handles} \cong (Y' \times I) \cup_{\varphi(A)} \text{handles}$$

where A denotes the attaching regions.

This lemma just ensures that our gluing operation is actually well-defined, and that a diffeomorphism of the boundary is enough to glue our data together (a possible counterfactual would be dependence on the specific surgery presentation of the common boundary Y).

Lecture 6: February 4th

Note that ∂X as drawn in the right-hand side of Figure 1.10 is just 0-surgery on a single knot, which turns out to be 6_1. This is because, for relative handlebodies, $\partial^- X$ is the 3-manifold prescribed (here) by the Dehn surgery curves, before any 2, 3, and 4-handles (in this case, $\partial^+ = \emptyset$); see Figure 1.11.

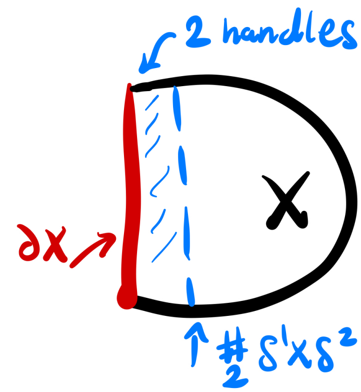

Figure 1.11: $\partial^- X$ is never affected by handle attachments.

In particular, ∂X has the same boundary as the 0-*trace* of 6_1:

Definition 1.4.5: n-Trace of a Knot

Given a knot $K \subseteq S^3$, its n-*trace* $X_n(K)$ is the 4-manifold given by attaching an n-framed 2-handle to B^4 along K.

Thus, we set $Z = X \cup_\partial X_0(6_1)$ and obtain a closed 4-manifold. Z is simply connected since there are no 1-handles — we can combine the natural handle decomposition for $X_0(6_1)$ with the above relative handle decomposition for X to obtain an absolute handle decomposition for Z with no 1-handles. The data of the map identifying ∂X with $\partial(X_0(6_1))$ tells us (in principle) where the relative handle attachment curves go. In this case, since we have manipulated the relative handle diagram for X to look exactly like the natural handle diagram for $X_0(6_1)$, the identification is superimposition, and thus, the handle diagram for Z we obtain in Figure 1.12 is just the relative handle diagram for X with $\langle 0 \rangle$ replaced by 0.

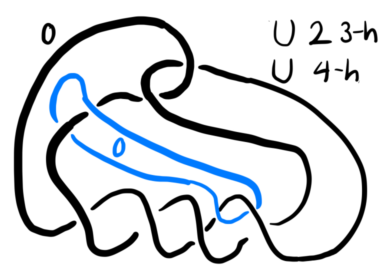

Figure 1.12: The handle diagram for Z obtained from the relative handle diagram for X (upside down).

Exercise 1.4.6

Show that $Z \underset{\text{sm}}{\cong} S^4$. Bonus: for any $f' : \partial X \to \partial X_0(6_1)$, $\pi_1(Z_{f'}) = 1$ where $Z_{f'}$ is obtained by gluing X to $X_0(6_1)$ along f'. Moreover, for any such f', $H_*(Z_{f'}) \cong H_*(S^4)$ so $Z_{f'}$ is a homotopy 4-sphere.

When saying that $\pi_1 = 1$ and an isomorphism on H_* are enough to imply homotopy equivalence, we are invoking a mildly nontrivial application of Whitehead's theorem; Freedman then implies that homotopy 4-spheres are homeomorphic to S^4.

PROOF : Sliding black over blue in the above handle diagram produces two unlinked 0-framed unknots which are then cancelled by the 3-handles:

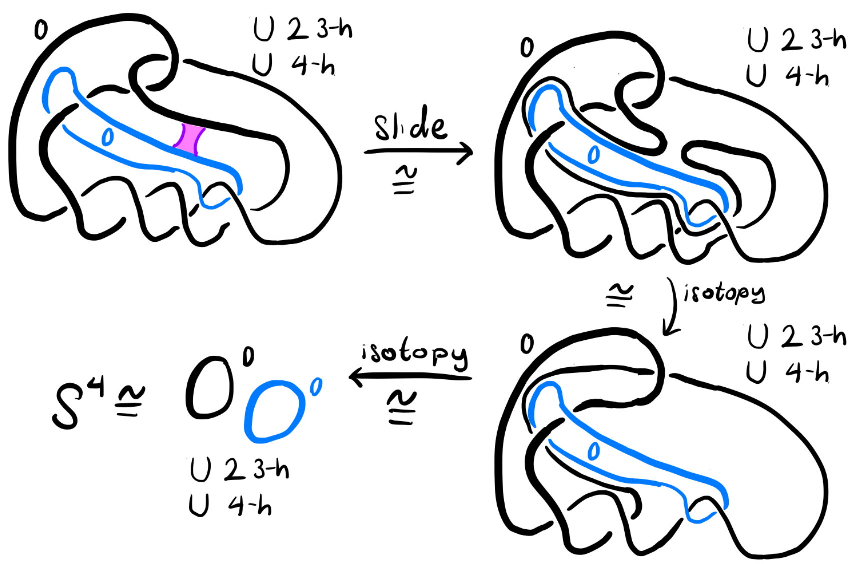

I wonder if this example has anything to do with sliceness of 6_1, this almost looks like a Kirby diagram for a ribbon disc exterior.

For the first bonus, note that the above argument still suffices; there are no 1-handles. Any gluing also produces a homology 4-sphere by counting handles; the problem is that, in principle, f' might produce a very complicated handle diagram for Z. However, since we know $\pi_1(Z_{f'}) = 1$, we know that H_2 is free and $H_3 = 0$, so it follows that the 2 and 3-handles must homologically cancel, as otherwise the 3-handles would generate nontrivial H_3 (one can also see this by calculating $\chi(X) = 2 + b_2(X)$ from the given handle presentation and noting that there can be no torsion in H_2). ■

Since homotopy 4-spheres are exciting (pending resolution of the smooth Poincaré conjecture in dimension 4), we should be interested in finding $f' \in \text{MCG}(S^3_0(6_1))$ different from our f above. To that end, we have the following definition:

Definition 1.4.7: Generalized Dehn Twists

Suppose we have $\iota : W^{n-1} \hookrightarrow Y^n$ and $\varphi : S^1 \to \text{Diff}^+(W)$ based at the identity. Then define the (ι, φ) Dehn twist of Y to be $\Phi : Y \to Y$ supported on $\nu(W) \cong W \times I$ given by

$$\Phi|_{W \times I}(w, t) = (\varphi(t)(w), t)$$

where $t \in S^1 = [0, 1]/(0 \sim 1)$. Then Φ is a self-diffeomorphism of Y.

I suppose the point(s) are that 1) the mapping class group in general is hard to understand and 2) having such a concrete element of an automorphism in terms of a codimension 1 submanifold allows one to (hopefully) apply some kind of gluing formula to compute Floer invariants.

Example 1.4.8

Consider a circle in a genus 2 surface which splits it into two genus 1 surfaces, and $S^1 \to \text{Diff}^+(S^1)$ given by θ mapping to rotation through θ. This is an ordinary Dehn twist.

Exercise 1.4.9

Find $T^2 \hookrightarrow S^3_0(6_1)$. Use that T^2 to define a Dehn twist homeomorphism f' of $S^3_0(6_1)$. Bonus: give a handle diagram for $Z_{f'}$.

Floer Homology

spin^c structures

Lecture 7: February 6$^{\text{th}}$

Our coverage of Floer homology will be very high-level, and we will mostly engage with the formal properties of the invariants which are often enough to do nontrivial computations with them. In particular, we don't really even need to know what a spin^c structure *means*, so we won't try to motivate them for the umpteenth time.

An orientation is extra structure on a manifold; one may or may not exist, and if one exists, there are several. spin^c structures can be thought of analogously, with the same properties. For 4-manifolds X, if $H_1(X)$ has no 2-torsion, then

$$\begin{aligned}\text{Spin}^c(X) = \{\text{spin}^c \text{ structures on } X\} = \\ \{\alpha \in H_2(X, \partial) : \forall \beta \in H_2(X, \partial) \quad \alpha \cdot \beta \cong \beta \cdot \beta \pmod 2\}\end{aligned}$$

i.e. spin^c structures are in bijection with the set of characteristic elements of second homology.

To see that a characteristic vector/element always exists, note that 0 is characteristic in the even case. In the odd case, pick an odd orthogonal basis for H_2 (one can always do this by starting with any orthogonal basis then adding any odd basis element (one must exist) to the even basis elements) and set α to be the sum of all basis elements. Then $\alpha \cdot \beta$ is the sum of the coefficients of β in this basis, and $\beta \cdot \beta$ is the sum of the squares of these coefficients. Since $x \equiv x^2 \pmod 2$, $\alpha \cdot \beta \equiv \beta \cdot \beta \pmod 2$.

For 3-manifolds Y, if $H_1(Y)$ has no 2-torsion, and Y is closed,

$$\text{Spin}^c(Y) = \{\alpha \in H_1(Y) := 2\beta \text{ for some } \beta \in H_1(Y)\}$$

A spin^c manifold is then a pair $(M, \mathfrak{s})$ of a manifold and a spin^c structure on it.

Note that we are not interested *per se* in spin^c 4-manifolds, it just turns out that the invariants we want to study are invariants of spin^c manifolds. Thus, if we want to show that $M \underset{\text{sm}}{\not\cong} M'$ using an invariant σ of spin^c manifolds, then we need to show that $\sigma(M, \mathfrak{s}) \neq \sigma(M', \mathfrak{s}')$ for any $\mathfrak{s}'$ which could arise as $f_*(s)$.

Note that we are interested in 4-manifolds with boundary and closed 3-manifolds which will arise as boundaries in our cases of interest; in particular, a spin^c structure on X will induce a spin^c structure on $\partial X = Y$ via the long exact sequence

$$\cdots \to H_2(X) \to H_2(X, Y) \xrightarrow{\partial} H_1(Y) \to \cdots$$

Exercise 1.5.1

Show that a characteristic element is sent to a 2-divisible element by the boundary map.

Example 1.5.2

If X is B^4 with a single 2-handle (i.e. a knot trace) then the cocore disc D of the 2-handle gives a class $[D] \in H_2(X, \partial)$, hence a spin^c structure on X ($[D]$ is a generator hence vacuously characteristic in a rank 1 lattice); the corresponding spin^c structure on $H_1(\partial)$ is given by $[\partial D]$.

Exercise 1.5.3

Show that there can be boundary spin^c structures which don't extend over X. Show that there can be boundary spin^c structures with non-unique extensions. Bonus: for $X_n(K)$, compute $\partial : H_2(X_2(K), \partial) \to H_1(\partial)$.

Proof : Working purely on the level of homology, for the first problem, we just want the boundary to have more H_1 than the manifold itself has H_2. Set $X = S^1 \times D^3$, where $H_1(Y) = \mathbb{Z}$, and inspect the long exact sequence.

For the second problem, set $X = (\mathbb{CP}^2)^\circ$ where M° denotes the punctured manifold $M \setminus \mathrm{int}(B^4)$; $\partial X = Y = S^3$, and once again, inspection of the long exact sequence provides the result. ■

A Taxonomy of Floer H_* Theories

There are roughly three main types of Floer homology:

For the purposes of this class, symplectic manifolds don't exist.

Instanton (I_*) due to Floer

Monopole (HM) due to Kronheimer-Mrowka

Heegaard (HF) due to Ozsváth-Szabó

Fundamentally, all these homology theories come from some kind of functorial association taking $(Y^3, \mathfrak{s})$ to an infinite dimensional Lie group G equipped with a Morse function $f : G \to \mathbb{R}$, and an infinite-dimensional analogue of Morse homology in that setting. The various Floer theories give us chain complexes C_* over a field $\mathbb{F}$ or a polynomial ring $\mathbb{F}[U]$ for some formal variable U (more than one formal variable is irrelevant to us in this class), generically denoted R. The Floer theories also capture essentially the same data as various gauge theoretic invariants: monopole Floer homology is equivalent to Seiberg-Witten theory, instanton Floer to Donaldson theory, and Heegaard-Floer (conjecturally) to Seiberg-Witten theory again.

The chain complexes C_* will be finitely generated, graded, and the chain homotopy type of C_* will be an invariant of $(Y, \mathfrak{s})$. spin^c-cobordisms $(W, \mathfrak{s}) : (Y_1, \mathfrak{s}_1) \to (Y_2, \mathfrak{s}_2)$ will induce an R-linear map $F_{(W,\mathfrak{s})} : C_*(Y_1, \mathfrak{s}_1) \to$

$C_*(Y_2, \mathfrak{s}_2)$ (i.e, our chain complexes are TQFTs). $H_*(W)$ and $\mathfrak{s}$ govern the grading shifts in the chain complex, so the classical topology determines the grading shift (via the Atiyah-Singer index theorem).

We can extract invariants of closed spin^c 4-manifolds from Floer homologies as cobordisms between the empty manifold. For reasons we will see, these will always vanish when $b_2^+ - b_1 \equiv 0 \pmod 2$.

I only know one incarnation of this parity problem for invariants of smooth structures on 4-manifolds, the version for Seiberg-Witten invariants; there, it arises from having our moduli space sitting inside of $K(\mathbb{Z}, 2) = \mathbb{CP}^\infty$, and we get our invariants by evaluating the homology class of the moduli space against $[\mathbb{CP}^1]^k = [\mathbb{CP}^k]$ which is an even-dimensional cocyle (at least this is how it works in the simply connected case). It's very bizarre to me that this fact that seems very specific to the setup of Seiberg-Witten theory is apparently a core limitation of all gauge/Floer-theoretic invariants of 4-manifolds.

For what follows, we will first develop the theory for 4-manifolds with boundary, then for closed 4-manifolds, and, finally, we will discuss some gluing formulae.

Heegaard Floer

There are two flavors of Heegaard-Floer chain complexes of interest to us: $\widehat{\mathrm{CF}}$ (over $\mathbb{F}$), and CF^- (over $\mathbb{F}[U]$). For our purposes, $\mathbb{F} = \mathbb{F}_2$; much of what we will discuss is also known over $\mathbb{Z}$, but we leave those extensions to the Floer theorists. The Floer chain complexes are $\mathbb{Q}$-graded for torsion spin^c structures. Note that $\mathrm{CF}^- \neq \widehat{\mathrm{CF}} \otimes_{\mathbb{F}} \mathbb{F}[U]$.

We will write CF° and $\mathrm{HF}^\circ(Y, \mathfrak{s}) = H_*(\mathrm{CF}^\circ(Y, \mathfrak{s}))$ when we are talking about either flavor. Also note that if the spin^c structure is not defined, then $\mathrm{HF}(Y) = \bigoplus_{\mathfrak{s} \in \mathrm{Spin}^c(Y)} \mathrm{HF}^\circ(Y, \mathfrak{s})$ by convention.

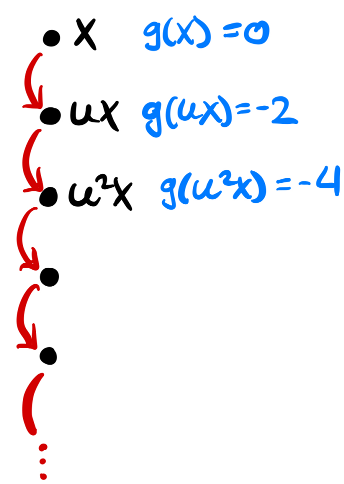

Figure 1.13: The HF^- tower for S^3, whose top grading is 0 and rank is 1 (this is all of the salient information but this tower serves as a good example for more complicated towers to come).

Example 1.5.4

We begin with the simplest 3-manifold, S^3; $H_1(S^3) = 1$ and therefore there is a unique spin^c structure (by convention, for integer homology 3-spheres, the spin^c structure is unique hence omitted in the notation). $\widehat{\mathrm{HF}}(S^3) = \mathbb{F}$ and $\mathrm{HF}^-(S^3) = \mathbb{F}[U]$; U has grade -2, so multiplication by U lowers the grading by 2, which is drawn as a tower, see Figure 1.13 in the margin.

Example 1.5.5

Next, we consider $S^1 \times S^2$; $H_1(S^1 \times S^2) = \mathbb{Z}$, so $\mathrm{Spin}^c(S^1 \times S^2) = 2\mathbb{Z}$. $\widehat{\mathrm{HF}}(S^1 \times S^2, \mathfrak{s}_{2i}) = 0$ for $i \neq 0$ and $\mathbb{F} \oplus \mathbb{F}$ for $i = 0$. $\mathrm{HF}^-(S^1 \times S^2, \mathfrak{s}_{2i}) = 0$ for $i \neq 0$ and $\mathbb{F}[U] \oplus \mathbb{F}[U]$ for $i = 0$. The absolute grading of the two generators for HF^- are $\pm\frac{1}{2}$, see Figure 1.14.

Example 1.5.6

Let Q be the connected sum of the right and left-handed trefoils (i.e. the square knot), and consider $S^3_{-1}(Q)$ which is a homology 3-sphere and therefore has a unique spin^c structure; its HF^- tower is depicted in Figure 1.15.

It turns out that $\widehat{\mathrm{HF}}(Y, \mathfrak{s}) = \bigoplus_i \mathbb{F}_{g_i}$ so the data of $\widehat{\mathrm{HF}}$ is the number of nonzero elements and their gradings. Similarly, $\mathrm{HF}^-(Y,\mathfrak{s}) = \bigoplus_j \mathbb{F}[U] \bigoplus_k \mathbb{F}[U]/U$ so HF^- has infinite towers and some torsion; the torsion part is denoted $\mathrm{HF}_{\mathrm{red}}$. By Figure 1.15, $\mathrm{HF}_{\mathrm{red}}(S^3_{-1}(Q)) = \langle x+y, y+z \rangle$.

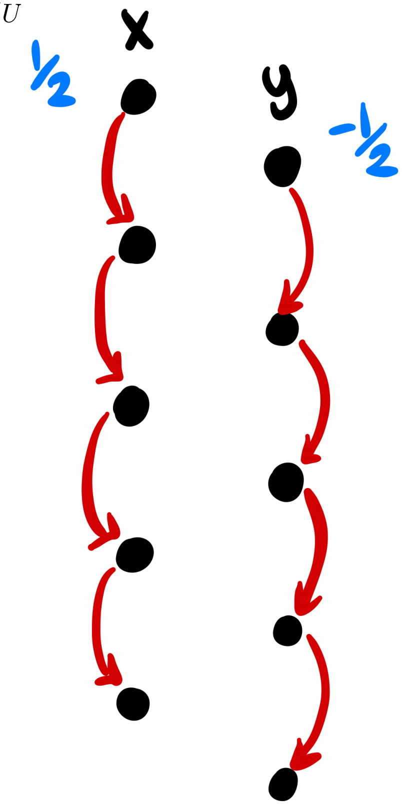

Figure 1.14: The HF^- tower for $S^1 \times S^2$ with two infinite towers and no torsion.

When $b_1(Y) = 0$, it turns out that there is a unique tower, whose grading is denoted $d(Y)$ (the d-invariant). We say that Y is an *L-space* if $\mathrm{HF}_{\mathrm{red}}(Y,\mathfrak{s})$ is trivial for all spin^c structures $\mathfrak{s}$.

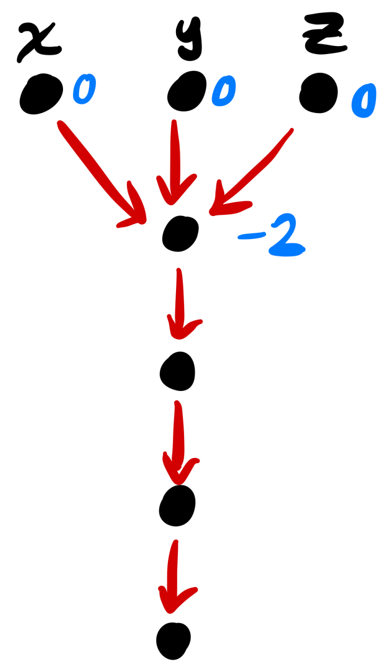

Figure 1.15: The HF^- tower for $S^3_{-1}(Q)$ with one infinite tower and nontrivial torsion.

References

Anar Akhmedov and B Doug Park. *Exotic smooth structures on small 4-manifolds*. In: *Inventiones mathematicae* 173.1 (2008), pp. 209–223.

R. Inanc Baykur and Noriyuki Hamada. *Exotic 4-manifolds with signature zero*. 2024. arXiv: 2305.10908 [math.GT]. URL: https://arxiv.org/abs/2305.10908.

Qiuyu Ren and Michael Willis. *Khovanov homology and exotic 4-manifolds*. 2024. arXiv: 2402.10452 [math.GT]. URL: https://arxiv.org/abs/2402.10452.