CHAPTER 6: FLOER THEORIES

DANNY CALEGARI

ABSTRACT. These are notes on Floer theories on 3-manifolds, which are being transformed into Chapter 6 of a book on 3-Manifolds. These notes follow a course given at the University of Chicago in Winter 2020.

CONTENTS

1. CLASSICAL INVARIANTS

1.1. Homology, linking, surgery.

1.1.1. *Alexander duality.* If X is a compact locally contractible subspace of the sphere S^n (for example, if X is a submanifold), Alexander duality is an isomorphism $\tilde{H}_q(S^n - X) \cong \tilde{H}^{n-q-1}(X)$.

Suppose $n = 3$ and X is a knot K. Then $H_1(S^3 - K) = H^1(K) = \mathbb{Z}$. It's convenient to replace $S^3 - K$ by the homotopy equivalent manifold $M := S^3 - N(K)$ where $N(K)$ is an open neighborhood of K. Then Lefschetz duality gives a chain of isomorphisms from $H_1(K) = H^1(M) = H_2(M, \partial M)$. An orientation for K gives a preferred generator for $H^1(M)$, represented by a homotopy class of map $M \to S^1$. If we make this map smooth, the preimage of a regular value is an oriented proper surface F in M whose boundary wraps once around the knot K, and by throwing away (necessarily homologically trivial) closed components if necessary, we can arrange for the surface to be connected. Such a surface — oriented, embedded, connected and with boundary equal to K — is called a *Seifert surface* for the knot. In § 1.2.6 we shall give an algorithm (Seifert's algorithm) to construct a Seifert surface from a knot projection.

If K' is any oriented knot in M the algebraic intersection of homology classes $[K'] \cap [F]$ is called the *linking number* of K' with K, and denoted $\mathrm{lk}(K', K)$. Linking number with

Date: November 8, 2023.

K realizes the chain of isomorphisms $H_1(M) = H^1(K) = \mathrm{Hom}(H_1(K);\mathbb{Z}) = \mathbb{Z}$. Linking number is symmetric, and changes sign when either orientation is reversed.

If $n = 3$ and X is a genus g surface F with one or zero boundary components, then $H_1(S^3 - F) = H^1(F) = \mathbb{Z}^{2g}$. Alexander duality is compatible with restriction to subspaces, so the pairing is given as before by linking number: if $a_1 \cdots a_{2g}$ is a basis for $H_1(F)$, there is a dual basis $\alpha_1 \cdots \alpha_{2g}$ for $H_1(S^3 - F)$ given by the formula $\mathrm{lk}(\alpha_i, a_j) = \delta_{ij}$.

Alexander duality works just as above in any manifold with the homology of S^n. A 3-manifold Y with the same homology as S^3 is called a *homology 3-sphere*. Thus, (for example) a knot in Y has a Seifert surface, oriented knots in Y have a well-defined linking number, and so on.

1.1.2. *Homology of the boundary.* If M is a compact manifold obtained from S^3 by removing a tubular neighborhood of a knot, then ∂M is a torus, and the inclusion $\partial M \to M$ induces a map on homology $H_1(\partial M) \to H_1(M)$ whose kernel is the primitive $\mathbb{Z}$ subspace generated by the class of the longitude; in particular, it has half dimension in $H_1(\partial M)$.

This is a special case of a more general fact:

Proposition 1.1. *Let M be a compact oriented 3-manifold with boundary ∂M. Then the kernel of $H_1(\partial M) \to H_1(M)$ is a Lagrangian subspace of $H_1(\partial M)$ with respect to the intersection pairing; in particular, it is half dimensional in $H_1(\partial M)$.*

Proof. This is an algebraic fact, and can be proved easily from Lefschetz duality and the long exact sequence, but it is illuminating to give a geometric proof. Let L denote the kernel of $H_1(\partial M) \to H_1(M)$. We show first that L is isotropic: the intersection pairing is trivial on L.

Let α, β be elements in L. To be in the kernel of the map on homology is to bound oriented immersed surfaces F_α, F_β in M. If we make these surfaces transverse, their intersection is a 1-manifold, consisting of closed loops, and intervals that run between pairs of intersection points of α with β of opposite sign. Thus $[\alpha] \cap [\beta] = 0$.

To show that L is half dimensional (i.e. Lagrangian), let α_i be a collection of loops in M that are a basis for the image of $H_1(\partial M)$. Each α_i is homologous to some β_i in ∂M, and the homology is realized by a surface F_i. Let G_i be a collection of proper surfaces in M representing classes in $H_2(M, \partial M)$ dual to the α_i, so that $[\alpha_i] \cap [G_j] = \delta_{ij}$. Then by making F_i and G_j transverse, we see as before that $[\beta_i] \cap [\partial G_j] = \delta_{ij}$. By definition the classes $[\partial G_j]$ are in L, so this gives a pairing between L and the image of $H_1(\partial M)$ in $H_1(M)$ and we are done. □

1.1.3. *Dehn surgery.* If Y is an oriented 3-manifold and K is a knot in Y, we may build a new 3-manifold by removing a solid torus neighborhood $N(K)$ of K and gluing it back by some diffeomorphism of the boundary. The result depends only on the (unoriented) isotopy class of the loop $\gamma \subset \partial N(K)$ (called a *slope*) that bounds a disk in the new solid torus, and is called the result of *Dehn surgery along K with slope γ*.

Oriented essential simple closed curves on $\partial N(K)$ up to isotopy are in bijection with primitive homology classes in $H_1(\partial N(K))$. If we choose a basis m, l for homology, we have $[\gamma] = p[m] + q[l]$ for coprime (p, q). Thus slopes are parameterized by $p/q \in \mathbb{Q} \cup \infty$, once we have chosen the basis (m, l).

The usual convention is to choose m to be a *meridian* for K; i.e. an essential loop bounding a disk in $N(K)$. The curve l is called a *longitude*, and intersects m in one point. Since Y is oriented, so is $\partial N(K)$ and we can choose orientations on m and l so that $[m] \cap [l] = 1$.

There are $\mathbb{Z}$ possibilities for l, and in general there is no canonical way to make a choice. However, if Y is a homology sphere we can choose l to be the boundary of a Seifert surface F for K. Thus if Y is a homology sphere, there is a canonical parameterization of slopes on $\partial N(K)$ by $\mathbb{Q} \cup \infty$, and we can speak unambiguously about *(p/q)-Dehn surgery along K*, denoted $Y_K(p/q)$. If Y is a homology sphere, then $H_1(Y_K(p/q)) = \mathbb{Z}/p\mathbb{Z}$.

Example 1.2. If K is the unknot, then $S^3_K(p/q)$ is the Lens space $L(p,q)$, which is $S^2 \times S^1$ if $p = 0$ and S^3 if $p = \pm 1$.

Example 1.3. If K is the right-handed trefoil, then 1 surgery on K is called the Poincaré homology sphere. The fundamental group of this sphere is a group of order 120 called the binary icosahedral group, and as we shall see in § 2.8.1 is the preimage in the group $\mathrm{SU}(2)$ of the group of orientation-preserving symmetries of the regular icosahedron, thought of as a subgroup of $\mathrm{SO}(3)$.

Theorem 1.4 (Dehn surgery presentation)**.** *Every oriented 3-manifold Y is obtained by integer Dehn surgery on some link L in S^3.*

Proof. Choose a Heegaard splitting of Y. There is a (unique) Heegaard splitting $H_1 \cup_\Sigma H_2$of S^3 of the same genus, and Y is obtained from this splitting by cutting along Σ and regluing H_2 by some mapping class φ of Σ. Express φ as a product of Dehn twists or their inverses $\varphi = \tau_1^{\pm 1} \cdots \tau_n^{\pm 1}$ where τ_i is a Dehn twist along a curve $\gamma_i \subset \Sigma$.

Parameterize a collar neighborhood of Σ as $\Sigma \times [0,1]$ and push γ_i to the level surface $\Sigma \times i/n$. The γ_i become in this way the components of an n-component link L in S^3. Let l' be the framing of γ_i coming from its embedding in Σ. Then doing $m \pm l'$ surgery on γ_i has the effect at the level of handlebodies of changing the gluing by the Dehn twist $\tau_i^{\pm 1}$. Note if γ_i is separating in Σ then $l' = l$, a longitude for γ_i, so that we are doing ± 1 surgery. Otherwise l' differs from l by some multiple of the meridian, and all we know is that the surgery coefficient is an integer. □

1.1.4. *Integer surgery and handlebodies.* Suppose K is a knot in an oriented 3-manifold Y. A *framing* for K is a trivialization of the normal bundle; equivalently, a framing is an identification of a neighborhood $N(K)$ with $D^2 \times S^1$. The data of a framing is the same thing as a section of $\partial N(K)$ thought of as an oriented S^1 bundle over K; i.e. a choice of a longitude l. If Y is a homology 3-sphere, there is a canonical longitude (the one bounding a Seifert surface for K), so if we orient K the framings are canonically parameterized by $\mathbb{Z}$, the coefficient of m when we express the framing longitude in terms of the canonical m, l basis. Otherwise, the set of framings is in bijection with $\mathbb{Z}$ but not canonically.

Suppose $Y = \partial W$ for some 4-manifold W. Let L be a link in Y, and for each component L_i choose some framing $l_i \in \partial N(L_i)$. We may build a new 4-manifold W' by attaching 2-handles to W along L. That is, for each i we glue a $D^2 \times D^2$ along $S^1 \times D^2$ by identifying the solid torus with $N(L_i)$ via the given framing. Evidently $\partial W'$ is equal to the result of Dehn surgery on Y along L with slopes equal to the framing.

Now, let L be an n-component link in S^3, and for each component L_i choose an integer n_i. Let Y be the result of n_i surgery along each component L_i. If we think of $S^3 = \partial B^4$ then we have exhibited $Y = \partial W$ where W is obtained from B^4 by attaching 2-handles as above. Since by Theorem 1.4 every oriented 3-manifold is obtained by integer surgery on a link in S^3, it follows that every oriented 3-manifold bounds a simply-connected 4-manifold.

Attaching a 2-handle to B^4 is the same as wedging with a 2-sphere at the level of homotopy. Thus $H_1(W) = 0$ and $H_2(W) = \mathbb{Z}^n$, where there is a bijection between the $\mathbb{Z}$ factors and the components L_i, and an orientation of L_i determines a generator of the associated $\mathbb{Z}$ in $H_2(W)$.

There is a fundamental relationship between the intersection pairing of $H_2(W)$ and linking between the components of L.

Lemma 1.5 (Intersection pairing is linking matrix)**.** *With notation as above, let A denote the (symmetric) matrix whose ij entry equals $\operatorname{lk}(L_i, L_j)$ and whose ii entry is n_i. Then A is the intersection matrix for $H_2(W)$ with the given basis.*

Proof. For each oriented L_i choose a Seifert surface F_i in S^3, and let $F_i' \subset W$ be obtained by attaching the core D^2 of the 2-handle. Then the $[F_i']$ are a basis for $H_2(W)$. By pushing F_i slightly into B^4 we see that $[F_i'] \cap [F_j']$ is equal to the signed intersection of L_i with F_j, which is precisely $\operatorname{lk}(L_i, L_j)$.

The surface F_i and the core D^2 on F_i' can both be pushed off themselves, and the difference on their boundary is given by the difference of the given framing of L_i and the canonical longitude framing. Thus $[F_i'] \cap [F_i'] = n_i$. $\square$

In particular, if Y is a homology sphere, then Y bounds a simply-connected 4-manifold W with definite intersection pairing.

1.1.5. *Kirby moves.* Suppose we have a presentation of Y as integer surgery on a link L in S^3. We may obtain new presentations by combinations of the following two moves:

(1) **(stabilization):** add an isolated unknot with framing ± 1; or
(2) **(handle slide):** band connect sum L_i to a framed pushoff of L_j to produce a new component L_i' and replace L_i by L_i' with framing $n_i + n_j + 2\operatorname{lk}(L_i, L_j)$.

These moves are sometimes abbreviated to K1 and K2 for short.

If we think of the surgery diagram as a handlebody description of a 4-manifold W that Y bounds, then the first move connect sums W with $\mathbb{CP}^2$ or $\overline{\mathbb{CP}}^2$ and the second move slides the attaching circle associated to the ith handle over the jth handle. Thus neither operation changes the homeomorphism type of the result of surgery.

At the level of the linking matrix, K1 sums with the matrix (± 1), while K2 adds the jth row and column to the ith row and column.

One corollary is the following:

Theorem 1.6 (Homology sphere presentation)**.** *Let Y be an integer homology sphere. Then Y is obtained by Dehn surgery on a link L where the components are pairwise unlinked and the coefficients are all ± 1.*

Proof. Start with any integral surgery presentation. Apply K1 twice to add two isolated unknots with framings 1 and -1. Now the linking matrix A is odd, indefinite and unimodular, and is therefore isomorphic over $\mathbb{Z}$ to a diagonal matrix with ± 1s. Achieve this diagonalization by a sequence of K2 moves. $\square$

Let L be a link as promised by Theorem 1.6. Doing the surgeries one by one gives a sequence of homology 3-spheres

$$S^3 = Y_0 \to Y_1 \to Y_2 \to \cdots \to Y_n = Y$$

where each Y_{i+1} is obtained from Y_i by doing ± 1 surgery on some knot $K_{i+1} \subset Y_i$. Here Y_i is obtained from S^3 by performing the surgeries on components L_1 through L_i, and K_{i+1} is the image of the component L_{i+1} in Y_i. Said another way: ± 1 surgery on a knot in a homology sphere produces a new homology sphere, and any homology 3-sphere may be obtained from S^3 by finitely many such operations.

A more elaborate procedure shows:

Theorem 1.7 (Even presentation)**.** *Let Y be any oriented 3-manifold. Then Y is obtained by integral Dehn surgery on a link L with all coefficients even.*

1.2. **The Alexander Polynomial.** Alexander [2] introduced a polynomial invariant of knots in the 3-sphere.

1.2.1. *Alexander module.* Let K be a knot in S^3, and let M be S^3 minus an open neighborhood $N(K)$ of K. By Mayer–Vietoris, M has the homology of a solid torus; i.e. $H_1(M) = \mathbb{Z}$ and $H_2 = H_3 = 0$. It follows that there is a unique $\mathbb{Z}$ cover $\hat{M}$.

Let t denote the generator of the deck group $\mathbb{Z}$ of the covering $\pi : \hat{M} \to M$. Then the deck group action makes the homology $H_1(\hat{M})$ into a *module* over the ring $\mathbb{Z}[t, t^{-1}]$, called the *Alexander module* of K.

1.2.2. *Seifert matrix and a presentation for the Alexander module.* Let's examine the structure of this module more closely. Let F be a Seifert surface for K, and let N be obtained from M by cutting along F. Then N is a sutured manifold with a single annulus suture whose core is a longitude for K, and whose boundary consists of two copies $F^{\pm}$ of F. We can think of the cover $\hat{M}$ as being obtained from $\mathbb{Z}$ copies of N stacked end to end, where F_n^+ is glued to F_{n-1}^- compatibly with their identification to F. We identify N with one of these copies, so that every other copy is $t^n N$ for some nonzero n.

Since F is connected, by Mayer-Vietoris $H_1(\hat{M})$ is generated as a module by $H_1(N)$. Further, the map $\hat{M} \to M$ induces zero on H_1, so for every loop α in N there is a surface $f : S \to M$ which bounds $\pi(\alpha)$. The surface S has a cyclic cover $\hat{S}$ that bounds α togther with its translates by t^n. The restriction of $\hat{S}$ to N witnesses a homology from α to the difference of two classes in F^+ and F^- respectively. Thus in fact $H_1(\hat{M})$ is generated as a module by $H_1(F)$.

To obtain a presentation for this module, we need to understand homologies in N between classes in F^+ and F^-. We use Alexander duality to obtain a linking pairing between $H_1(N)$ and $H_1(R)$: if $a_1 \cdots a_{2g}$ is a basis for $H_1(F)$ there is a dual basis $\alpha_1 \cdots \alpha_{2g}$ of $H_1(N)$ for which $\mathrm{lk}(\alpha_i, a_j) = \delta_{ij}$. The inclusion of $F^{\pm}$ into N induces two maps $H_1(F) \to H_1(N)$: geometrically, a class a_i may be pushed off F to the positive or negative side to obtain loops a_i^+ and a_i^- in N representing homology classes $\sum \mathrm{lk}(a_i^+, a_j)\alpha_j$ and $\sum \mathrm{lk}(a_i^-, a_j)\alpha_j$ in $H_1(N)$ respectively.

On the other hand, as thought of as homology classes in $\hat{M}$, the class a_i in R^- and the same class a_i in F^+ differ by the action of the generator t of the deck group. Thus we obtain relations

$$\sum \mathrm{lk}(a_i^+, a_j)\alpha_j = t \sum \mathrm{lk}(a_i^-, a_j)\alpha_j$$

By Mayer–Vietoris all relations in $H_1(\hat{M})$ arise in this way, from equalities between classes under the inclusions $H_1(F^\pm) \to H_1(N)$. Finally, observe that $\mathrm{lk}(a_i^-, a_j) = \mathrm{lk}(a_i, a_j^+) = -\mathrm{lk}(a_j^+, a_i)$ so if we define the *Seifert matrix* V to have ij entry $V_{ij} := \mathrm{lk}(a_i^+, a_j)$ we obtain a $\mathbb{Z}[t, t^{-1}]$-module presentation

$$H_1(\hat{M}) = H_1(F)/(tV - V^T)$$

1.2.3. *The Alexander Polynomial.* Now it's an algebraic fact that if M is a module over a commutative ring A with unit which has an $s \times r$ presentation matrix P, the ideal generated by all $r \times r$ minors of A is an invariant of M called the *order ideal*. If P is a square matrix, the order ideal is principal and generated by $\det(P)$. Any other generator will differ by multiplication by a monomial $\pm t^n$, so we may normalize the generator (if we like, and following Alexander) to have no powers of t^{-1}, and to have a positive constant term.

This evidently applies in our case, so the polynomial $\det(tV - V^T)$ (suitably normalized) is an invariant of K, called the *Alexander polynomial*, and denoted $\Delta_K(t)$ or just $\Delta(t)$ if K is understood.

From the construction we can derive several useful properties of this polynomial

Theorem 1.8 (Alexander polynomial). *Let K be a knot in S^3, and $\Delta(t)$ its Alexander polynomial. Let* $\mathrm{genus}(K)$ *denote the least genus of a Seifert surface for K.*

(1) *homology:* $\Delta(1) = \pm 1$;
(2) *symmetry:* $\Delta(t^{-1}) = t^{-2g}\Delta(t)$ *up to units;*
(3) *genus:* $\deg(\Delta(t)) \le 2\,\mathrm{genus}(K)$;
(4) *fibered: if K is a fibered knot, $\Delta(t)$ is monic, and* $\deg(\Delta(t)) = 2\,\mathrm{genus}(K)$.

Proof. We prove each bullet in turn.

(1): The Seifert form V and the intersection form ι on $H_1(R)$ are evidently related by $\iota = V - V^T$. Since ι is a nondegenerate symplectic form, this implies that $\Delta(1) = \det(\iota) = 1$ up to units.

(2): We compute

$$\Delta(t^{-1}) = \det(t^{-1}V - V^T) = t^{-2g}\det(V - tV^T) = t^{-2g}\det((tV - V^T)^T) = t^{-2g}\Delta(t)$$

(3): Obvious

(4): If K is fibered then N is a product and $\Delta(t)$ is the characteristic polynomial of the monodromy matrix. $\square$

Because of bullets (1) and (2) it follows that there is a unique normalized Alexander Laurent polynomial satisfying $\Delta(1) = 1$ and $\Delta(t^{-1}) = \Delta(t)$. In terms of a Seifert matrix with a symplectic basis this is given by $\Delta_K(t) = \det(t^{1/2}V - t^{-1/2}V^T)$ where V is the Seifert form. We call this the *symmetric normalization* if we need to be specific. Note that for this normalization, $\Delta''(1)$ is even.

1.2.4. *Examples.* The trivial knot has $\Delta = 1$. The trefoil (of either handedness) has $\Delta = t - 1 + t^{-1}$. The figure 8 knot has $\Delta = -t + 3 - t^{-1}$. These two knots are both fibered, with fiber of genus 1 and monodromy $\left(\begin{smallmatrix} 0 & 1 \\ -1 & 1 \end{smallmatrix}\right)$ and $\left(\begin{smallmatrix} 2 & 1 \\ 1 & 1 \end{smallmatrix}\right)$ respectively, and Δ is the characteristic polynomial of the mondromy matrices in each case.

The knot 5_2 in Rolfsen's tables has $\Delta = 2t - 3 + 2t^{-1}$. Since this is not monic, the knot is not fibered.

1.2.5. *Slice knots.* A knot K in S^3 is *slice* (one also says *topologically slice*) if it bounds an embedded tame (i.e. locally flat) disk in B^4. It is *smoothly slice* if the disk can be taken to be smooth. Smoothly slice implies topologically slice but not vice versa. Note that every knot K bounds an embedded disk in B^4, simply by coning to 0; but such a knot is not locally flat at 0 unless K is the unknot.

The Alexander polynomial can be used to show that certain knots are not slice.

Theorem 1.9 (Fox–Milnor)**.** *Suppose K is slice. Then $\Delta_k(t) = p(t)p(t^{-1})$ for some polynomial p.*

Proof. We give a proof that works for smoothly slice knots.

Let F be a Seifert surface for K of genus g. If K is smoothly slice, we can complete F to a closed embedded surface F' of the same genus in B^4. By Alexander duality and general position, F' bounds a 3-manifold M in B^4. Linking number of knots in S^3 is equal to algebraic intersection number of surfaces they bound in B^4. Since $F' = \partial M$ it follows from Proposition 1.1 that the kernel of $H_1(F') \to H_1(M)$ is a Lagrangian subspace W of $H_1(F')$; in particular it has dimension g. If α, β are in W, we can span them by surfaces S, T in W. Then the pushoff α^+ bounds a surface S^+ obtained by pushing off M, and therefore is disjoint from T.

It follows that by choosing a suitable basis, the Seifert matrix V has a block diagonal form

$$V = \begin{pmatrix} 0 & A \\ B & C \end{pmatrix}$$

so that $\Delta = \det(t^{1/2}A - t^{-1/2}B^T)\det(t^{1/2}B - t^{-1/2}A^T) = p(t)p(t^{-1})$ for some p.

If K is merely topologically slice the same argument works, but one needs to know topological transversality in dimension 4; see e.g. the book by Freedman–Quinn [8] Chapter 9. □

Example 1.10. Let K be the figure 8 knot. Then $\Delta_K(-1) = 5$. Since 5 is not square, K is not slice.

1.2.6. *Seifert's Algorithm and Link Polynomial.* Seifert gave an algorithm to produce a Seifert surface F for an oriented knot K from a knot projection. The algorithm is as follows: first, resolve each of the crossings compatibly with the orientation to produce a family of oriented circles with disjoint embedded projections to the plane. Bound each circle with a disk, oriented compatibly, and connect up the disks with a twisted band for each crossing. The result is a Seifert surface; see Figure 1.

If L is a link, then for any orientation of the components and a suitable projection, the same algorithm produces a Seifert surface F compatible with the given orientation (note that a Seifert surface must be *connected*; thus one needs to choose a projection in which

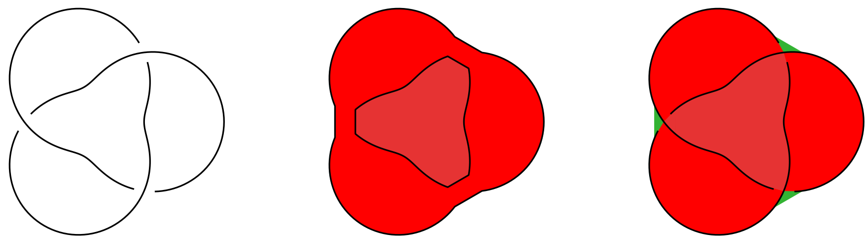

FIGURE 1. Seifert's algorithm: resolve crossings compatible with orientation, bound oriented loops by oriented disks, and attach a twisted band for each crossing.

the various components cross themselves). Such a surface gives rise to a Seifert linking matrix V, and we can define $\Delta_L(t) = \det(t^{1/2}V - t^{-1/2}V^T)$, the Alexander polynomial of the *oriented* link L.

If L has more than one component then the intersection form on $H_1(F)$ is degenerate, so $\Delta_L(1) = 0$. Note that with this normalization, Δ_L will have terms with half-integral powers of t when L has an even number of components.

1.2.7. *Skein relation.* Let L_+ and L_- be two oriented links whose projections are related only by changing one specific crossing. Let L_0 be the oriented link obtained by resolving the crossing compatibly with the orientations.

Seifert's algorithm relates Seifert surfaces for the three oriented links in a simple way; analyzing the effect on the linking matrix, one obtains a relation between the Alexander polynomials of the links, called the *skein relation*, and first discovered by Alexander. For the symmetric normalization this has the form

$$\Delta_{L_+} - \Delta_{L_-} = (t^{1/2} - t^{-1/2})\Delta_{L_0}$$

In Alexander's normalization the right hand side is $(t-1)\Delta_{L_0}$. This skein relation gives an inductive algorithm to compute Δ.

To prove the Skein relation, simply apply Seifert's algorithm to obtain three Seifert surfaces F_+, F_- and F_0 for the three links in question. Then $F_\pm$ are each obtained from F_0 by adding a band, and the two bands differ only by a twist. The core of the band together with any arc in F_0 between the endpoints gives the new generator a for $H_1(F_\pm)$; for any $b \in H_1(F_0)$ the linking numbers $\mathrm{lk}(a^+, b)$ are the same for F_+ as for F_-, while the values of $\mathrm{lk}(a^+, a)$ for F_+ and F_- differ by 1. Thus the Seifert forms $V_\pm$ for $F_\pm$ are obtained from V by adding the same row and column, differing only by 1 in the upper left entry, from which the Skein relation follows immediately.

1.2.8. *Reducible representations.* de Rham showed that non-abelian reducible representations of $\pi_1(S^3 - K)$ into $\mathrm{SL}(2, \mathbb{C})$ correspond to the roots of $\Delta_K(t)$. The number of representations associated to each root is determined by a refinement of Δ coming from the ideals generated by the minors of a presentation matrix of various sizes. One clean statement of de Rham's theorem is the following:

Theorem 1.11 (de Rham)**.** *Let K be a knot in S^3, let M be the complement of a tubular neighborhood of K, and let $\mu \in \pi_1(M)$ represent a meridian of K. Then there is a nonabelian upper-triangular representation $\rho : \pi_1(M) \to \mathrm{SL}(2,\mathbb{C})$ for which the upper off-diagonal entry of $\rho(\mu)$ is m if and only if m^2 is a root of $\Delta_K(t)$.*

This is proved by thinking of $H_1(\hat{M})$ as the abelianization of the commutator subgroup of $\pi_1(M)$, and recovering its structure via 2-step solvable representations to $\mathrm{SL}(2,\mathbb{C})$.

1.2.9. *Knots in Homology Spheres.* The only properties of S^3 used in the definition of the Alexander polynomial are those coming from Alexander duality. Thus for any homology 3-sphere Y and any knot K in Y there is an Alexander polynomial $\Delta_K(t)$ with symmetric normalization equal to $\det(t^{1/2}V - t^{-1/2}V^T)$ for Seifert matrix V defined using linking numbers exactly as above, and this polynomial satisfies all the properties of Theorem 1.8.

2. The Casson Invariant

In 1985 Casson gave a series of lectures introducing a new invariant of homology 3-spheres. The invariant is defined from a Heegaard decomposition, and behaves in a predictable way under Dehn surgery, in which regards it bears a close family resemblance to the theory of Heegaard Floer Homology that we shall take up in § 4.

A basic reference for the Casson invariant is Akbulut–McCarthy [1] or Saveliev [18].

2.1. **Spin structures and Parallelizability.** In each dimension n the group $\mathrm{Spin}(n)$ is the connected double cover of $\mathrm{SO}(n)$. For $n = 3$ we have $\mathrm{Spin}(3) = \mathrm{SU}(2)$. Any oriented $\mathbb{R}^n$ bundle E over a reasonable space X can by given a fiberwise metric, reducing its structure group to $\mathrm{SO}(n)$. A *Spin structure* on E is a lift of its structure group from $\mathrm{SO}(n)$ to $\mathrm{Spin}(n)$. A Spin structure on an oriented manifold is (by abuse of notation) a Spin structure on its tangent bundle.

Since $\pi_1(\mathrm{SO}(n)) = \mathbb{Z}/2\mathbb{Z})$ (when $n \geq 3$), we have $\pi_2(B\mathrm{SO}(n)) = \mathbb{Z}/2\mathbb{Z}$ and therefore also $H^2(B\mathrm{SO}(n);\mathbb{Z}/2\mathbb{Z}) = \mathbb{Z}/2\mathbb{Z}$. The pullback of the generator of this $\mathbb{Z}/2\mathbb{Z}$ under a classifying map $X \to B\mathrm{SO}(n)$ associated to a principal $\mathrm{SO}(n)$ bundle $E \to X$ is called the *second Stiefel-Whitney class* $w_2(E) \in H^2(X;\mathbb{Z}/2\mathbb{Z})$. The oriented double covering $\mathrm{Spin}(n) \to \mathrm{SO}(n)$ induces a map on classifying spaces $B\mathrm{Spin}(n) \to B\mathrm{SO}(n)$ which is an isomorphism on π_i for $i \neq 2$, so $w_2(E)$ precisely measures the obstruction to the existence of a lift.

Hence: the obstruction to the existence of a Spin structure on an oriented bundle $E \to X$ is the second Stiefel-Whitney class $w_2(E) \in H^2(X;\mathbb{Z}/2\mathbb{Z})$, and if such a structure exists, the set of structures up to isomorphism is in bijection with $H^1(X;\mathbb{Z}/2\mathbb{Z})$.

Proposition 2.1 (Oriented 3-Manifolds are spin)**.** *Let M be an oriented 3-manifold. Then $w_2(M) := w_2(TM) = 0$. In other words, M admits a Spin structure.*

The following proof is from Kirby [11] Ch. VII Thm. 1:

Proof. It suffices to prove this when M is compact, since nontriviality of a bundle may be detected on some compact piece. Further, we may assume M is closed, since we may reduce to this case by doubling.

Let $X \subset M$ denote the Poincaré dual to w_2. If w_2 is nonzero, X is a 1-manifold, which may be taken to be connected. Since any nonzero class is primitive in homology (with $\mathbb{Z}/2\mathbb{Z}$ coefficients) we may find a closed embedded surface S (possibly non-orientable) which intersects X transversely in a single point x. By the definition of X, there is a Spin structure s on $M - X$ which does not extend over X.

Now, $TM|S = TS \oplus \nu$ where ν is the normal bundle. Since M is orientable, if S is orientable ν is trivial, and otherwise ν is non-orientable, precisely around non-orientable loops on S. Thus $TM|S$ is trivial, and therefore Spin.

The set of Spin structures on $TM|S$ is parameterized by $H^1(S;\mathbb{Z}/2\mathbb{Z})$, which also parameterizes the set of Spin structures on $H^1(S-x;\mathbb{Z}/2\mathbb{Z})$. Now, s restricts to a spin structure on $TM|(S-x)$. Since restriction $H^1(S;\mathbb{Z}/2\mathbb{Z}) \to H^1(S-x;\mathbb{Z}/2\mathbb{Z})$ is an isomorphism, it follows that s extends to a Spin structure on $TM|S$; in other words, it extends over x, contrary to the definition of X. This contradiction shows that $w_2 = 0$ after all. □

Corollary 2.2 (Oriented 3-Manifolds are parallelizable)**.** *Every oriented 3-manifold M is parallelizable. I.e. TM is isomorphic to a trivial bundle $M \times \mathbb{R}^3$.*

Proof. Since M is oriented and spin, TM admits a Spin structure, classified up to isomorphism by the homotopy class of a map from M to $B\mathrm{Spin}(3)$. Since $\mathrm{Spin}(3) = S^3$, the classifying space $B\mathrm{Spin}(3)$ has $\pi_4 = \mathbb{Z}$ and $\pi_i = 0$ for $i < 4$. Thus every map from a 3-manifold to $B\mathrm{Spin}(3)$ is homotopically trivial. □

Here's how to explicitly construct a trivialization. Since the dimension of M is odd, the Euler characteristic is zero by duality. Thus M admits a nowhere zero vector field v, and $v^\perp$ is an oriented $\mathbb{R}^2$ bundle whose Euler class (i.e. first Chern class) satisfies $c_1(v^\perp) = w_2 = 0$ mod 2. For any knot X we can trivialize TM along a neighborhood NX so that v is a constant section. Relative to this trivialization, any other section w of $TM|NX$ is a map from NX to S^2 in such a way that v is the constant map to the north pole. Define a map $\phi : NX \to S^2$ so that on each $D^2 \times$ point it maps ∂D to the north pole, and so that the induced map from $D^2/\partial D^2 \to S^2$ has degree n. Let w be the vector field on NX associated with ϕ, and observe that w agrees with v on ∂NX so we can extend it by $w = v$ on $M - NX$. Now $c_1(w^\perp) = c_1(v^\perp) + 2nX^D$ where X^D is the Poincaré dual of X so after finitely many operations of this kind we can find a nowhere zero vector field v for which $c_1(v^\perp) = 0$. Thus $v^\perp$ is a trivial $\mathbb{R}^2$ bundle, and admits a nowhere zero section u, and once we have two nowhere zero linearly independent sections v and u we can find a third.

2.2. **Rokhlin's Theorem.**

Now suppose W is a smooth oriented 4-manifold. A class A in $H_2(W)$ is represented by a smooth oriented embedded surface S. There is a decomposition $TW|S = TS \oplus \nu$ where ν is the normal bundle. Evidently $w_2(TW)[S] = w_2(TS)[S] + w_2(\nu)[S]$ mod 2. Since S is oriented, $\chi(S)$ is even, so $w_2(TS)[S] = 0$, and $w_2(A)$ is equal to the self-intersection number of S mod 2.

If W is simply-connected, or more generally has no 2-torsion in its homology, then $H_2(W;\mathbb{Z}/2\mathbb{Z}) = H_2(W) \otimes \mathbb{Z}/2\mathbb{Z}$ and therefore W is spin if and only if the intersection form on (integer) homology is *even*.

The signature σ of an even unimodular symmetric form (i.e. the difference of the number of positive and negative eigenvalues) is always a multiple of 8. However it turns out that not all such forms are realized as the intersection form on the homology of a smooth 4-manifold. Famously, one has Rokhlin's theorem:

Theorem 2.3 (Rokhlin)**.** *Let W be a smooth, closed, oriented spin 4-manifold. Then the signature $\sigma(W)$ is divisible by 16.*

One way to prove this is to use the *Dirac operator*. This is an elliptic differential operator mapping between sections of a certain bundle (the bundle of *spinors*) that can be defined from a spin structure. The index of this operator can be computed from the Atiyah–Singer index theorem, and is equal to $-p_1(W)/24$. The index is a difference of dimensions of complex vector spaces, and is therefore an integer. However, in dimension 4, it turns out that the kernel and cokernel of the Dirac operator are *quaternionic* vector spaces, so that their complex dimensions are even. Hence p_1 is divisible by 48, and since $\sigma(W) = p_1(W)/3$ Rokhlin's theorem follows.

2.3. **The Rokhlin Invariant.** Every homology 3-sphere M is spin in a unique way, and bounds a smooth spin 4-manifold W. Because M is a homology sphere, the intersection pairing on $H_2(W)$ is nondegenerate, and because W is spin, it is even. Thus for algebraic reasons the signature is divisible by 8. The *Rokhlin invariant* of M, denoted $\mu(M)$, is $\sigma(W)/8 \bmod 2$. This is well-defined, because if W' is another smooth spin 4-manifold bounding M, then $W \cup -W'$ is a closed smooth spin 4-manifold, and therefore has $\sigma(W \cup -W')$ divisible by 16, by Rokhlin's Theorem 2.3. But $\sigma(W \cup -W') = \sigma(W) - \sigma(W')$; the claim follows.

Example 2.4. The three-sphere bounds B^4 so $\mu(S^3) = 0$. If M is the Poincaré homology 3-sphere then M bounds a 4-manifold with intersection form equal to the E_8 lattice, thus $\mu(M) = 1$.

2.4. **Arf Invariant and behavior under surgery.** Let W be a smooth closed oriented 4-manifold, simply-connected for simplicity. A *characteristic surface* is an oriented embedded surface F Poincaré dual to w_2; i.e. such that $[F] \cap x = x \cap x \bmod 2$ for every homology class x. Because F is characteristic, $\sigma(W)$ is congruent to $[F] \cap [F] \bmod 8$.

Associated to the pair (W, F) is a quadratic form on $H_1(F; \mathbb{Z}/2\mathbb{Z})$. A quadratic form on a $\mathbb{Z}/2\mathbb{Z}$-vector space has a mod 2 invariant, called the *Arf invariant*, and we obtain in this way a mod 2 invariant $\mathrm{Arf}(W, F)$.

Rokhlin's theorem generalizes to the following formula:

$$\mathrm{Arf}(W, F) = \frac{\sigma(W) - [F] \cap [F]}{8} \bmod 2$$

If M is an integral homology sphere embedded in W and separating F into $F' \cup D^2$ then $F \cap M$ is a knot $K \subset M$, and the quadratic form on $H_1(F; \mathbb{Z}/2\mathbb{Z})$ is isomorphic to the mod 2 reduction of the Seifert form of K. In this case the Arf invariant may be calculated from the Seifert matrix, and is equal to $\Delta_K''(1)$ where the Alexander polynomial is normalized so that $\Delta_K(t^{-1}) = \Delta_K(t)$ (i.e. it is a Laurent polynomial symmetric about the constant term) and $\Delta_K(1) = 1$.

Using this property one obtains a surgery formula for the Rokhlin invariant: if K is a knot in an integer homology sphere M, and Δ_K is the (normalized) Alexander polynomial of K, Then

$$\mu(M_K(1/n)) = \mu(M) + (n/2)\Delta_K''(1) \bmod 2$$

2.5. **Representation Varieties.** The *Casson Invariant* of a homology 3-sphere is an integer lift of the Rokhlin invariant. It can be defined informally as half the number of conjugacy classes of representations from $\pi_1(M)$ into $\mathrm{SU}(2)$, counted with multiplicity and with sign.

For a finitely generated group π, we define $R(\pi) := \mathrm{Hom}(\pi, \mathrm{SU}(2))$. This is a real algebraic variety, and contains an open subvariety $R^*(\pi) := \mathrm{Hom}_{\mathrm{irr}}(\pi, \mathrm{SU}(2))$ of irreducible representations. Note that a representation is reducible if and only if it is conjugate to a diagonal one.

The group $\mathrm{SU}(2)$ acts on $R(\pi)$ by conjugation. This action factors through $\pm\mathrm{Id}$ so descends to an action of $\mathrm{SO}(3)$. This action of $\mathrm{SO}(3)$ is free on $R^*(\pi)$.

At a nontrivial reducible representation the stabilizer is a circle subgroup, and at the trivial representation the stabilizer is all of $\mathrm{SO}(3)$. Denote by $X(\pi)$ the quotient $X(\pi) := R(\pi)/\mathrm{SU}(2)$ and also denote $X^*(\pi) := R^*(\pi)/\mathrm{SU}(2)$.

2.6. **Group Cohomology.** We can interpret the tangent spaces to $R(\pi)$ and $X(\pi)$ in terms of group cohomology. Fix a finitely presented group π and a representation $\rho : \pi \to \mathrm{SU}(2)$. Let ρ_t be a smooth deformation of ρ. There is a map $u : \pi \to \mathfrak{su}(2)$ defined by

$$u(g) = \frac{d}{dt}\Big|_{t=0} \rho_t(g)\rho(g^{-1})$$

Differentiating the equation $\rho_t(gh) = \rho_t(g)\rho_t(h)$ at $t = 0$ gives the relation

$$u(gh) = u(g) + \mathrm{ad}(\rho(g))u(h)$$

In other words, u is a 1-cocycle with values in $\mathfrak{su}(2)$. We denote the space of 1-cocycles by $Z^1_\rho(\pi, \mathfrak{su}(2))$, where the subscript ρ indicates the π-module structure on $\mathfrak{su}(2)$.

Suppose we choose a presentation $\pi = \langle S \mid L\rangle$. In this way we can think of π as a quotient of a free group F_S. Points in $R(F_S)$ are parameterized by $\mathrm{SU}(2)^{|S|}$ and $R(\pi)$ is the preimage of $(\mathrm{id})^{|L|}$ under the evaluation map $\Phi : \mathrm{SU}(2)^{|S|} \to \mathrm{SU}(2)^{|L|}$.

The *Zariski tangent space* to $R(\pi)$ at ρ is the kernel of the map $d\Phi$ at ρ, and is evidently isomorphic to $Z^1_\rho(\pi, \mathfrak{su}(2))$. If $\rho \in \Phi^{-1}(\mathrm{id})^{|L|}$ and Φ is a submersion at ρ, this kernel may be identified with the ordinary (smooth) tangent space $T_\rho R(\pi)$. Informally, we call such a ρ a *smooth point* of $R(\pi)$.

For π free of rank n, we have $R(\pi) = \mathrm{SU}(2)^n$ and evidently every point is smooth, so $T_\rho R(\pi) = Z^1_\rho(\pi, \mathfrak{su}(2))$.

For π a surface group, we have the presentation

$$\pi = \langle a_1, b_1, \cdots, a_g, b_g \mid \prod_i [a_i, b_i]\rangle$$

In this case we have the following:

Lemma 2.5 (Smooth submanifold). *Let $\Phi : \mathrm{SU}(2)^{2g} \to \mathrm{SU}(2)$ be defined by*

$$\Phi(A_1, B_1, \cdots, A_g, B_g) = \prod_i [A_i, B_i]$$

If $\rho \in \mathrm{SU}(2)^{2g} = R(F_{2g})$ is irreducible, then $d\Phi|_\rho$ surjects onto $\mathfrak{su}(2)$.

This can be proved by a direct calculation. Thus for π a surface group, the tangent space to $R(\pi)$ is smooth at every irreducible representation.

If $\xi \in \mathfrak{su}(2)$, conjugating by $\exp(t\xi)$ produces a family of deformations with $u(g) = \mathrm{ad}(\rho(g))\xi - \xi$; i.e. u is a 1-coboundary, an element of $B^1_\rho(\pi, \mathfrak{su}(2))$. This exactly parameterizes the tangent vectors to the conjugation action of $\mathrm{SU}(2)$. Since this action is free (mod the discrete center $\pm\mathrm{Id}$) at an irreducible representation, it follows that at a smooth point ρ of $R^*(\pi)$ we may identify $T_\rho X^*(\pi)$ with $H^1_\rho(\pi, \mathfrak{su}(2))$. In particular, this identification is valid throughout X^* for π a free or surface group.

2.7. **Heegaard Splittings.** Now, let M be a homology 3-sphere, and fix a Heegaard splitting of M; i.e. a decomposition into handlebodies H_1, H_2 of genus g glued along their common boundary surface Σ. Let Σ_* be the result of removing a small disk from Σ.

By abuse of notation, for a topological space Y we denote $R(Y) := R(\pi_1(Y))$ and analogously $R^*(Y)$, $X(Y)$, $X^*(Y)$. The diagram of inclusions

$$\Sigma_* \to \Sigma \to H_i \to M$$

induces diagrams of inclusions the other way

$$R(M) \to R(H_i) \to R(\Sigma) \to R(\Sigma_*)$$

Since each inclusion of spaces is surjective on π_1, the maps on representation varieties are inclusions.

Since $\pi_1(\Sigma^*)$ and $\pi_1(H_i)$ are free, each $R(H_i)$ is just $\mathrm{SU}(2)^g$ and $R(\Sigma_*)$ is $\mathrm{SU}(2)^{2g}$. In particular, the $R(H_i)$ are submanifolds of complementary dimension in $R(\Sigma_*)$, and we can compute their intersection number (which is well-defined up to sign, depending on an orientation). For an arbitrary three manifold, the intersection number is 0 if $\beta_1(M) > 0$ and otherwise it has absolute value equal to $|H_1(M;\mathbb{Z})|$. Thus if M is a homology sphere, the intersection number is ± 1 and we may choose an orientation on $R(\Sigma_*)$ for which the intersection is 1.

By Lemma 2.5, the subspace $R^*(\Sigma)$ is a smooth submanifold of $R^*(\Sigma_*)$ of codimension 3, and since the conjugation action of $\mathrm{SU}(2)$ at an irreducible representation is free mod the (discrete) center, it follows that $X^*(\Sigma_*)$ and $X^*(\Sigma)$ are smooth oriented manifolds of dimension $6g-3$ and $6g-6$ respectively, and the $X^*(H_i)$ are smooth oriented submanifolds of $X^*(\Sigma)$ of dimension $3g-3$.

Now, although $X^*(\Sigma)$ is noncompact, it may be compactified by adding back the conjugacy classes of reducible representations. Since M is a homology sphere, every reducible representation of $\pi_1(M)$ is trivial. Furthermore the trivial representation is isolated in $R(M) = R(H_1) \cap R(H_2)$, or else a nontrivial tangent vector to $R(M)$ at the identity would give a nontrivial homomorphism from $H^1(M)$ to the Lie algebra $\mathfrak{su}(2)$. Thus the intersection of the $X^*(H_i)$ in $X^*(\Sigma)$ is compact, and there is a well-defined algebraic intersection $X^*(H_1) \cap X^*(H_2)$ that can be computed by perturbing $X^*(H_1)$ (say) by a compactly supported perturbation so that it's transverse to $X^*(H_2)$, and then computing intersection numbers in the obvious way.

The *Casson Invariant* $\lambda(M)$ is (up to multiplication by some universal constant ± 1 to be determined in § 2.8.1) defined to be $(-1)^g/2$ times the algebraic intersection number of the $X^*(H_i)$ in $X^*(\Sigma)$.

2.8. **Invariance under stabilization.** To show this is well-defined one must show that it doesn't depend on the choice of Heegaard splitting. By the Reidemeister–Singer theorem, any two Heegaard splittings of a given 3-manifold have a common stabilization. So we just need to check that the invariant doesn't change under stabilization.

Let $M = H_1' \cup_{\Sigma'} H_2'$ be the result of stabilization. Σ_*' is obtained from Σ_* by boundary connect summing a once-punctured torus, whose meridian and longitude curves α, β bound disks in H_2' and H_1' respectively. Thus $R(\Sigma_*')$ is equal to $R(\Sigma_*) \times \mathrm{SU}(2)^2$, where the two extra $\mathrm{SU}(2)$ factors parameterize the images under a representation of α and β respectively, and similarly we have $X^*(\Sigma_*') = X^*(\Sigma_*) \times \mathrm{SU}(2)^2$. With these coordinates, $X^*(H_1') = X^*(H_1) \times \mathrm{SU}(2) \times 1$ and $X^*(H_2') = X^*(H_2) \times 1 \times \mathrm{SU}(2)$, and if we perturb $X^*(H_1)$ (say) to be transverse to $X^*(H_2)$, the product of the perturbation with $\mathrm{SU}(2) \times 1$ is transverse to $X^*(H_2')$. Thus after perturbation, there is a bijection between the points of $X^*(H_1) \cap X^*(H_2)$ and $X^*(H_1') \cap X^*(H_2')$, and one can check that the sign of the intersection is reversed. It follows that our formula for $\lambda(M)$ is invariant under stabilization, and is therefore an invariant of M.

2.8.1. *Example: the Poincaré homology sphere.* We shall now compute $\lambda(M)$ for M the Poincaré homology sphere. It will turn out that $\lambda(M) = \pm 1$. By convention, we normalize the definition of the Casson invariant by multiplying by a universal constant so that $\lambda(M) = 1$.

Let's denote $\pi_1(M) = \Gamma$. If K is the right-hand trefoil, Seifert van-Kampen gives a presentation

$$\pi_1(S^3 - K) = \langle x, y \mid xyx = yxy \rangle$$

Doing $+1$ surgery adds a relation

$$\Gamma = \langle x, y \mid xyx = yxy, yx^2y = x^3 \rangle$$

Making the substitution $z = xy$ reduces this to

$$\Gamma = \langle x, z \mid zx = x^{-1}z^2, z^{-1}zxz = x^3 \rangle = \langle x, z \mid (zx)^2 = z^3 = x^5 \rangle$$

Evidently Γ has a trivial abelianization. Furthermore, the element $I = z^3 = x^5$ is central, so we deduce that Γ is a central extension of the alternating group A_5:

$$A_5 := \langle x, z \mid (zx)^2 = z^3 = x^5 = 1 \rangle$$

which has $H^2(A_5; \mathbb{Z}/2\mathbb{Z}) = \mathbb{Z}/2\mathbb{Z}$, and exhibits Γ as the nontrivial central $\mathbb{Z}/2\mathbb{Z}$ extension of A_5.

Now, A_5 may be realized as a subgroup of $\mathrm{SO}(3)$ as follows. The elements z and x both normally generate, so their image in any nontrivial representation must have order 3 and 5 respectively. There is only one conjugacy class in $\mathrm{SO}(3)$ of order 3, namely a rotation of order $2\pi/3$ fixing the north and south pole $\pm p$. There are two nontrivial conjugacy classes of elements of order 5, namely rotations of order $2\pi/5$ and $4\pi/5$ fixing antipodal points $\pm q$ where we may take q to lie on the Greenwich meridian. For each fixed choice of $2\pi/5$ or $4\pi/5$ the product xz is a rotation through an angle α that depends only on the latitude of q, and there is a unique latitude for which $\alpha = \pi$.

We thus obtain precisely two nontrivial representations of A_5 in $\mathrm{SO}(3)$ up to conjugacy, and their preimages in $\mathrm{SU}(2)$ give the two nontrivial conjugacy classes of representations of Γ to $\mathrm{SU}(2)$. One may check that the character varieties are transverse at these representations, and that the intersection numbers have the same sign. Thus $\lambda(M) = \pm 1$ as claimed.

2.9. **A symplectic structure on $X^*(\Sigma)$.** Proposition 1.1 dualizes to show that for any compact oriented M, the image $H^1(M;\mathbb{R}) \to H^1(\partial M;\mathbb{R})$ is a Lagrangian subspace with respect to the cup product pairing. Because $\mathbb{R}$ is abelian, we may think of $H^1(M;\mathbb{R})$ as the character variety of homomorphisms from $\pi_1(M)$ to $\mathbb{R}$ up to conjugacy (conjugation acts trivially).

It turns out that character varieties of surface groups are symplectic in great generality. We follow the discussion in Goldman [9]. Recall that for Σ a closed oriented surface of genus g that the tangent space $T_\rho X^*(\Sigma)$ is identified with $H^1_\rho(\Sigma, \mathfrak{su}(2))$.

There is a nondegenerate form (the Killing form) on $\mathfrak{su}(2)$ given by $\langle A, B\rangle := \mathrm{tr}(AB)$. Composing with cup product defines a pairing

$$H^1_\rho(\Sigma,\mathfrak{su}(2)) \times H^1_\rho(\Sigma,\mathfrak{su}(2)) \xrightarrow{\cup} H^2_\rho(\Sigma,\mathfrak{su}(2)\otimes\mathfrak{su}(2)) \xrightarrow{\langle\cdot,\cdot\rangle} H^2(\Sigma,\mathbb{R}) \xrightarrow{\int} \mathbb{R}$$

where the map $H^2(\Sigma,\mathbb{R}) \to \mathbb{R}$ is obtained by integrating over a fundamental class. Poincaré duality implies that this pairing is nondegenerate; i.e. it defines a nondegenerate 2-form ω on $X^*(\Sigma)$. It turns out that this 2-form is closed, and therefore defines a symplectic structure on $X^*(\Sigma)$.

If M is any compact oriented 3-manifold with $\partial M = \Sigma$ then (evidently) the class of Σ becomes trivial in $H_2(M)$. Dually, this implies that the symplectic form vanishes identically on the image of $X^*(M)$ in $X^*(\Sigma)$. For $M = H_i$, a handlebody of genus g, the map $X^*(H_i) \to X^*(\Sigma)$ is an inclusion of half dimension, and the image is therefore a Lagrangian submanifold.

2.10. **Behavior under surgery.** To actually calculate λ in practice it turns out to be very useful to understand how it varies under surgery. If M is a homology 3-sphere and K is a knot, there is a well-defined meridian and longitude (up to sign), and for any integer n the result of $1/n$ surgery on K produces a new homology 3-sphere that we denote $M_K(1/n)$.

Theorem 2.6 (Casson invariant surgery formula)**.** *Let K be a knot in an integer homology sphere M. Then there is a formula*

$$\lambda'(M,K) := \lambda(M_K(1/(n+1))) - \lambda(M_K(1/n)) = \frac{1}{2}\Delta''_K(1)$$

where Δ_K is the Alexander polynomial of K, normalized to satisfy $\Delta(t) = \Delta(t^{-1})$ and $\Delta(1) = 1$.

In particular, $\lambda(M_K(1/n)) = \lambda(M) + (n/2)\Delta''_K(1)$.

By Theorem 1.6, any homology sphere may be obtained from S^3 by a sequence of ± 1 surgeries, and therefore by the surgery formula in Theorem 2.6, the value of λ is completely determined from the Alexander polynomials of the surgery knots. It follows that λ is the unique invariant of homology 3-spheres satisfying $\lambda(S^3) = 0$ and satisfying the surgery formula.

By the surgery formula for the Rokhlin invariant, we deduce:

Corollary 2.7. *The mod 2 reduction of the Casson invariant $\lambda(M)$ is equal to the Rokhlin invariant $\mu(M)$.*

2.11. **Proof of Theorem 2.6, part I.** Theorem 2.6 will be proved in a sequence of steps. First we prove a preliminary combinatorial lemma.

Lemma 2.8 (Seifert surface gives Heegaard splitting)**.** *Let K be a knot in a homology sphere M. Then there is a Seifert surface F for K so that the complement of $H_1 := F\times[0,1]$ is a handlebody H_2.*

Proof. Choose any Seifert surface F' for K. A neighborhood N' of F' is a handlebody, with boundary two copies $F'_\pm$ of F' connected along an annulus with core K.

The complement $M - N'$ is not necessarily an open handlebody, but if we triangulate it, the complement of the 1-skeleton is. A neighborhood of this 1-skeleton can be isotoped to the neighborhood of a homotopic graph $\Gamma \subset M - F'$ attached to F' at a single vertex on the positive side (say), and then a neighborhood N of $F' \cup \Gamma$ is a handlebody with handlebody complement.

The problem is that N is now no longer a thickened Seifert surface for K; it is a thickened Seifert surface with a "thickened graph" Γ attached on the positive side. To correct this we drill out Γ together with an unknotted arc α that runs across N' from the vertex where Γ attaches to the other side. This adds a new thickened copy of Γ to the complement, which therefore remains a handlebody, while N minus a neighborhood of $\Gamma\cup\alpha$ is homeomorphic to a product $F \times [0,1]$ where F is a new Seifert surface for K with the desired properties. ∎

Lemma 2.9 (λ' is well-defined)**.** *Let K be a knot in a homology sphere M. Then for any n the difference $\lambda(M_K(1/(n+1))) - \lambda(M_K(1/n))$ does not depend on n and is an invariant of the pair (M,K), denoted $\lambda'(M,K)$.*

Proof. By Lemma 2.8 the knot K has a Seifert surface F for which $H_1 := F \times [0,1]$ is a handlebody in a Heegaard splitting $M = H_1 \cup_\Sigma H_2$. The Heegaard surface Σ is made from two copies of F that we denote $F^\pm$ joined by an annulus whose core is K. With this setup, $(1/n)$ surgery on K is accomplished by cutting and regluing H_2 by the nth power τ^n of a Dehn twist on Σ along K.

The knot K decomposes Σ into $F^\pm$, and a pair of representations $\rho^\pm : \pi_1(F^\pm) \to \mathrm{SU}(2)$ glue together to give a representation $\rho : \pi_1(\Sigma) \to \mathrm{SU}(2)$ if and only if $\rho^+(K) = \rho^-(K)$ (here we are implicitly choosing a basepoint for π_1 on K, and thinking of K itself as an element of $\pi_1(F^\pm)$). The Dehn twist τ acts on representations by

$$\tau^* : (\rho^+,\rho^-) \to (\rho^+, \rho(K)\rho^-\rho(K)^{-1})$$

i.e. it acts by conjugating ρ^- by the common element $\rho(K)$.

Every element α in $\mathrm{SU}(2)$ except for $\pm\mathrm{Id}$ has distinct 1-dimensional eigenspaces, with eigenvalues $e^{i\theta}, e^{-i\theta}$ for some unique $\theta \in (-\pi,\pi)$. For $t \in [0,1]$ let α_t denote the element with the same eigenspaces and with eigenvalues $e^{it\theta}, e^{-it\theta}$. Then $(t,\alpha) \to \alpha_t$ defines a deformation retraction of $\mathrm{SU}(2) - \{-\mathrm{Id}\}$ to Id. Note that α_t commutes with α.

Let $R_-(\Sigma)$ be the subspace of $R(\Sigma)$ where $\rho(K) = -\mathrm{Id}$ and let $X_-(\Sigma) \subset X^*(\Sigma)$ be its image. Note that τ^* fixes $R_-(\Sigma)$. The self-maps Id and τ^* of $R(\Sigma) - R_-(\Sigma)$ are isotopic, by an isotopy that drags each representation ρ along the path $t \to (\rho^+, \rho(K)_t\rho^-\rho(K)_t^{-1})$ for

$t \in [0,1]$, where $\rho(K)_t$ is the 1-parameter family associated to $\rho(K)$ as above. This isotopy is compatible with the conjugation action, and descends to an isotopy of $X^*(\Sigma) - X_-(\Sigma)$.

Applying this isotopy to $X^*(H_1) - X_-(\Sigma)$ we see that $X^*(H_1)$ and $\tau^* X^*(H_1)$ are homologous, modulo a cycle δ supported in a neighborhood of $X_-(\Sigma)$ and therefore the quantity $\lambda(M_k(1/(n+1))) - \lambda(M_k(1/n))$ is equal to half the algebraic intersection of δ with $(\tau^{n+1})^* X^*(H_2)$ up to sign (one needs to check that the isotopy preserves transversality of $R(H_1)$ and $R(H_2)$ near the trivial representation). But δ is supported near $X_-(\Sigma)$ which is fixed pointwise by τ, so this algebraic intersection is independent of n, as claimed. □

In fact, one can give an explicit description of the difference cycle δ. The Seifert surface Σ is made from two copies of the Seifert surface F, and either inclusion of F into H_1 is a homotopy equivalence. A representation ρ of $\pi_1(H_1)$ is therefore the same thing as a representation (ρ^+, ρ^-) of $\pi_1(\Sigma)$ for which $\rho^+ = \rho^-$.

If we choose standard free generators $\alpha_1, \beta_1, \cdots, \alpha_h, \beta_h$ for $\pi_1(H_1) = \pi_1(F)$, then K is the representative of $\prod_i [\alpha_i, \beta_i]$, and in Lemma 2.5 we defined the map $B : R(F) \to \mathrm{SU}(2)$ by $B(\rho) = \rho(K)$, and asserted that the differential of B is surjective at irreducible representations. In particular, $-\mathrm{Id}$ is a regular value, so that $X^*(H_1)$ is transverse to $X_-(\Sigma)$.

Let ρ^+ be any representation of $\pi_1(H_1)$ with $\rho^+(K) = -\mathrm{Id}$. Then ρ^+ can be perturbed to a nearby representation with $\rho^+_\alpha(K) = \alpha$ for any α near $-\mathrm{Id}$ in $\mathrm{SU}(2)$. The track of ρ^+_α under the isotopy conjugates ρ^- by the α_t in $\mathrm{SU}(2)$, and the union over all α close to $-\mathrm{Id}$ together fills up $\mathrm{SU}(2)$. Thus the cycle δ is parameterized by pairs (ρ, g) up to conjugacy, where $\rho : \pi_1(H_1) \to \mathrm{SU}(2)$ has $\rho(K) = -\mathrm{Id}$.

Here is another way to say this. Restriction to subsurfaces defines a map $p : X_-(\Sigma) \to X_-(F) \times X_-(F)$ and the preimage of the diagonal $p^{-1}(\Delta)$ is the set of pairs of representations (ρ^+, ρ^-) that differ by conjugation, up to *simultaneous* conjugacy. From an element of $p^{-1}(\Delta)$ we can recover (ρ^+, g) in δ by $\rho^- = g\rho^+ g^{-1}$ up to the ambiguity of multiplying g by $-\mathrm{Id}$. In other words, homologically speaking, $\delta = 2p^{-1}(\Delta)$. For the reader who is uncomfortable with pulling back homology classes, note that we can express this dually in terms of cohomology: the Poincaré dual of δ is equal to 2 times the pullback under p^* of the Poincaré dual of Δ. Note that since Δ is an integral cycle, this implies that λ' (and therefore also λ) is an integer.

Now, if M is a homology sphere, and K_1, K_2 are knots in M with $\mathrm{lk}(K_1, K_2) = 0$ then for any integers m, n the result of $(1/m, 1/n)$ surgery on K_1, K_2 is also a homology sphere. We denote the result by $M(1/m, 1/n)$.

For any integers m, n define

$$\begin{aligned}\lambda''(M, K_1, K_2) := \lambda(M(1/(m+1), 1/(n+1))) - \lambda(M(1/m, 1/(n+1))) \\ - \lambda(M(1/(m+1), 1/n)) + \lambda(M(1/m, 1/n))\end{aligned}$$

Since the right hand is equal both to $\lambda'(M_{K_2}(1/(n+1)), K_1)$ and to $\lambda'(M_{K_1}(1/(m+1)), K_2)$ its value is in fact independent of both m and n, and depends only on K_1, K_2.

A 2-component link $L = K_1 \cup K_2$ in a homology sphere is a *boundary link* if the K_i bound disjoint Seifert surfaces. This implies that $\mathrm{lk}(K_1, K_2) = 0$, but the converse is false.

Lemma 2.10 ($\lambda'' = 0$ for boundary links)**.** *Let M be a homology sphere, and let $L := K_1 \cup K_2$ be a boundary link in M. Then $\lambda''(M, K_1, K_2) = 0$.*

Proof. We first claim that we can find a Seifert surface F for K_1 so that the complement of $H_1 := F \times [-1, 1]$ is a handlebody H_2, and such that K_2 is a separating curve on $F \times 1$. To see this, first choose disjoint Seifert surfaces F_1, F_2 for K_1, K_2. Then thicken F_2 to a handlebody with boundary G so that K_2 is a separating curve on G, and tube F_1 to G to produce a new Seifert surface F' for K_1 that contains K_2 as a separating curve. Now proceed as in the proof of Lemma 2.8 to obtain the desired F.

Let δ be the difference cycle associated to K_1, so that $\lambda'(M, K_1)$ is half the intersection of δ with $X^*(H_2)$ up to sign. If τ denotes the Dehn twist along K_2, then $\lambda''(M, K_1, K_2)$ is half the intersection of $X^*(H_2)$ with $\delta - \tau^*(\delta)$ up to sign. So it suffices to show that τ^* acts trivially on the homology of $X_-(\Sigma)$.

From our description of the Poincaré dual of δ as 2 times the pullback of the Poincaré dual of Δ it follows that the homology class of δ depends only on the action of τ on the (co)-homology of $X_-(F)$. Now, it can be shown directly that the action of the mapping class group $\mathrm{Mod}(F)$ on the homology of $X_-(F)$ factors through its action on the homology of F; a detailed exposition of this fact can be found in Akbulut-McCarthy [1] Theorem VI.2.4. Since K_2 is separating in $F \times 1$, the Dehn twist τ acts trivially on $H_1(F \times 1)$, and we are done. □

2.12. **Proof of Theorem 2.6, part II.** We now show that Theorem 2.6 follows formally from Lemma 2.9 and Lemma 2.10.

If $K \cup K'$ is a boundary link in a homology sphere M, Lemma 2.10 says that $\lambda'(K', M) = \lambda'(K', M_K(1))$. We claim that the same is true for the Alexander polynomial of K:

Lemma 2.11. *Let $K \cup K'$ be a boundary link in a homology sphere M. Then the Alexander polynomials of K' as computed in M or in $M_K(1)$ are equal.*

Proof. Let F and F' be disjoint Seifert surfaces for K and K'. If α is any loop on F' then α is disjoint from F, so that $\mathrm{lk}(\alpha, K) = 0$. This implies that we can find a Seifert surface F_α for α disjoint from K. But for any β the matrix entry $\mathrm{lk}(\beta^+, \alpha)$ of the Seifert form of K' is the algebraic intersection of β^+ with F_α, and since both are disjoint from K the intersection is the same whether computed in M or $M_K(1)$. □

Now if M is a homology sphere, we can obtain S^3 by a sequence of ± 1 surgeries on a succession of knots $K_1, \cdots, K_m$. If K is a knot in M we can move each of K and K_1 individually by isotopy until $K \cup K_1$ is a boundary link. It follows by induction and Lemma 2.11 that to prove Theorem 2.6 it suffices to prove it for knots in S^3.

Since $\lambda'(K)$ and $(1/2)\Delta''_K(1)$ both vanish on the unknot, it suffices to show that these quantities both change in the same way under a crossing change for a knot in S^3. Orient K, and let c be a small unknotted circle linking a crossing in such a way that $\mathrm{lk}(K, c) = 0$. Performing ± 1 surgery on c (as appropriate) takes S^3 to itself, and takes K to a knot K_c obtained from K by changing the crossing. We want to understand the difference $\lambda''(S^3, K, c) = \lambda'(S^3, K) - \lambda'(S^3, K_c)$.

A pair of linking circles c, c' for K is called a *crossing pair* if they bound disjoint disks D, D' that each intersect K in two points, in such a way that these two pairs of points are unlinked in K. Note that this implies that c, c' are a boundary link in $S^3 - K$ and hence in $S^3_K(\pm 1)$, since we can tube D, D' with cylindrical neighborhoods of disjoint arcs on K to make disjoint genus one Seifert surfaces for c, c' in the complement of K. By Lemma 2.10 it follows that $\lambda''(S^3, K, c) = \lambda''(S^3, K_{c'}, c)$.

Let $K_{c'}$ be the result of the crossing change at c', and $K_{cc'}$ the result of the crossing change at both c and c'.

Lemma 2.12. *With notation as above, there is an identity*

$$\Delta''_K(1) - \Delta''_{K_c}(1) = \Delta''_{K'_c}(1) - \Delta''_{K_{cc'}}(1)$$

Proof. Denote by K_0 and $K_{0c'}$ the results of resolving the crossings of K and $K_{c'}$ at c, and by K_{00} the result of resolving the crossings of K at both c and c'. Note that the condition on c, c' implies that K_{00} has 3 components; the other two links each have two components.

The Skein formula for the normalized Alexander polynomial gives formulae

$$\Delta_{K_c} - \Delta_K = \pm(t^{1/2} - t^{-1/2})\Delta_{K_0}, \quad \Delta_{K_{cc'}} - \Delta_{K_{c'}} = \pm(t^{1/2} - t^{-1/2})\Delta_{K_{0c'}},$$
$$\Delta_{K_0} - \Delta_{K_{0c'}} = \pm(t^{1/2} - t^{-1/2})\Delta_{K_{00}}$$

depending on the sign of the crossing change, and therefore the difference of the first two left hand sides is equal to $\pm(t^{1/2} - t^{-1/2})^2\Delta_{K_{00}}$. Since K_{00} is a link with more than one component, its Alexander polynomial vanishes at 1, so this expression has a zero of at least 3rd order at 1, and therefore its second derivative vanishes at 1. □

Thus the change in both $\lambda'(S^3, K)$ and $(1/2)\Delta''_K(1)$ under a crossing change at c is the same for K as for any knot obtained from K by a crossing change in any c' for which c, c' is a crossing pair.

Start with any knot K and any crossing c. The disk D cuts K into two strands. By suitable choices of c' we may move each strand arbitrarily through itself but not through the other strand. After finitely many such moves, we may reduce to a knot K' of a very simple form, as illustrated in Figure 2.

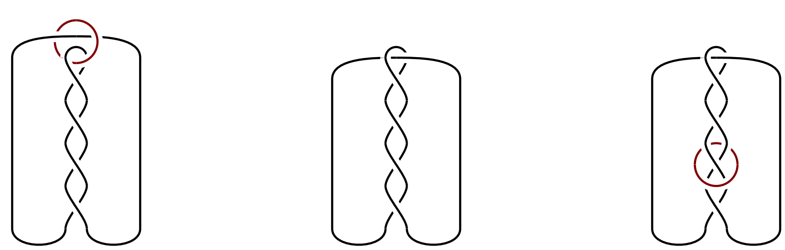

FIGURE 2. The unknot K' and linking circle c; the knot $K'_c = K(n)$ for $n = 2$; and the knot $K(n)$ with linking circle c'.

In other words, we may arrange for K' to be an unknot, and for K'_c to be of the form $K(n)$ where the two strands of K on either side of D have linking number n. Thus we always have an identity of the form

$$\lambda''(S^3, K, c) = \lambda'(K'_c) - \lambda'(K') = \lambda'(K(n))$$

On the other hand, if c' is an unknot as in the figure, twisting across a disk D' bounded by c' takes $K(n)$ to $K(n+1)$. So

$$\lambda''(S^3, K(n), c) = \lambda'(K(n+1)) - \lambda'(K(n))$$

is independent of n, and we must simply check that $\lambda'(K(n)) = (1/2)\Delta''_{K(n)}(1)$ for $n = 1$. Now, $K(1)$ is the trefoil, and $(1/2)\Delta''(1) = 1$. On the other hand, $+1$ surgery on the right handed trefoil gives rise to the Poincaré homology sphere M, and as we computed in § 2.8.1 one also has $\lambda(M) = 1$. This completes the proof of Theorem 2.6.

3. Instanton homology

Shortly after Casson introduced his invariant, Floer [5, 6] building on work of Taubes [20] managed to interpret it as the Euler characteristic of a homology theory. In fact Floer gave *two* homology theories with this property; the equivalence between these two theories is still open, and known as the Atiyah-Floer conjecture. Both are defined via Morse theory on certain infinite dimensional spaces. Before introducing these theories we recall the relationship between Morse theory and homology on finite dimensional manifolds.

3.1. Morse Theory and homology.

Let M be a compact smooth n-manifold. A smooth function $f : M \to \mathbb{R}$ is *Morse* if the critical points $df = 0$ are nondegenerate; i.e. near each such point p there are smooth local coordinates x_i vanishing at p such that

$$f(x) = f(p) - x_1^2 - \cdots - x_i^2 + x_{i+1}^2 + \cdots + x_n^2$$

for some i, the *index* of the critical point.

It is often convenient to assume f is *self-indexing*: i.e. that $f(p) = i$ for every critical point of index i. Every compact smooth manifold admits many self-indexing Morse functions.

For each $t \in \mathbb{R}$ define $M_t := f^{-1}(-\infty, t]$. Then M_t is empty for $t < 0$, and is equal to all of M for $t > n$. Furthermore, M_t is a smooth manifold with boundary $\partial M_t := f^{-1}(t)$ whenever $t \neq 0, 1, \cdots, n$. If we choose a Riemannian metric on M, we can define the gradient vector field $\mathrm{grad}(f)$. The flowlines of $\mathrm{grad}(f)$ provide a diffeomorphism from M_s to M_t whenever $i < s < t < i+1$ for some integer i, and if $i-1 < s < i < t < i+1$ we obtain M_t from M_s by attaching i-handles, one for each critical point of index i. Thus at the level of homotopy, M has the homotopy type of a CW complex with one i-cell for each critical point of index i.

The critical points freely generate the cellular chain groups for this structure. To compute the homology of M from f we need to be able to see differentials between chain groups in adjacent dimension. For each pair of critical points p, q of indexes i and j we can consider the space $F(p, q)$ of flowlines of $\mathrm{grad}(f)$ that are asymptotic to p and q in the positive and negative direction respectively. For generic f, the space $F(p, q)$ has the structure of the interior of a compact oriented manifold with corners of dimension $i - j - 1$. The space $F(p, q)$ can be compactified by adding products $F(p, r_1) \times F(r_1, r_2) \times \cdots F(r_k, q)$ for intermediate collections of critical points $r_1, \cdots, r_k$.

Thus when $i - j = 1$, the space $F(p, q)$ is a finite set of points, and when $i = j = 2$ the space $F(p, q)$ is a finite union of circles, and open intervals compactified by points of the form $F(p, r) \times F(r, q)$ where the index of r is $i - 1 = j + 1$.

The manifolds $F(p,q)$ can be oriented by thinking of them as intersections of the (oriented) manifolds of *all* flowlines asymptotic to p (resp. q) in the future (resp. past). So: for $i-j=1$ the space $F(p,q)$ is a finite set of *signed* points, and we can count this space with sign to get an integer $n(p,q)$.

Now: think of an index i critical point p as a generator in the (cellular) chain group in dimension i. We define

$$\partial p = \sum_{\text{index}(q)=i-1} n(p,q)q$$

Then

$$\partial\partial p = \sum_{\text{index}(q)=i-2}\ \sum_{\text{index}(r)=i-1} n(p,r)n(r,q)q$$

But for each q, the sum $\sum n(p,r)n(r,q)$ is equal to the number of boundary points of the 1-manifold $F(p,q)$, counted with sign, which is evidently zero. Thus ∂ is the differential of a chain complex, and the homology of this complex is $H_*(M)$.

3.2. **Gauge theory.** $\mathbb{C}^2$ bundles over 3-manifolds are classified by c_1; thus a $\mathbb{C}^2$ bundle E over an integral homology sphere M is trivial, and we can pick a trivialization.

The trivialization identifies sections of E with smooth $\mathbb{C}^2$-valued functions on M. Exterior d acting on such functions can be thought of as an $\mathfrak{su}(2)$-connection on E and every other $\mathfrak{su}(2)$-connection differs from this by a unique 1-form with coefficients in $\mathfrak{su}(2)$; in other words, we may identify the space $\mathcal{A}$ of $\mathfrak{su}(2)$ connections on E with $\Omega^1(M,\mathfrak{su}(2))$, and the notation $A\in\mathcal{A}$ means both the connection, and the associated matrix of 1-forms.

We let d_A denote the covariant derivative associated to a connection A. It acts on sections of E by $d_A\sigma = d\sigma + A\sigma$, and on $\mathfrak{su}(2)$-valued functions (i.e. $\Omega^0(M,\mathfrak{su}(2))$) by essentially the same formula $d_AB = dB + A\cdot B$ where now $A\cdot B$ denotes the adjoint action; i.e. $d_AB = dB + AB - BA$. We extend this action to all of $\Omega^*(M,\mathfrak{su}(2))$ by the Leibniz rule; thus in general $d_AB = dB + A\wedge B - (-1)^{\deg(B)}B\wedge A$.

Any two trivializations of E are related by an element of $\mathcal{G} := C^\infty(M,\mathrm{SU}(2))$, also called the *gauge group*. It acts on connections by $g\cdot A := gAg^{-1} + gdg^{-1}$. The tangent space to $\mathcal{G}$ at the identity is $\Omega^0(M,\mathfrak{su}(2))$, and under the action, a tangent vector B pushes forward to a tangent vector $-d_AB \in T_A\mathcal{A} = \mathcal{A}$.

We denote the quotient by this action $\mathcal{B} = \mathcal{A}/\mathcal{G}$. Let $\mathcal{A}^*$ denote the space of irreducible connections (those for which the holonomy is an irreducible subgroup of $\mathrm{SU}(2)$), and define $\mathcal{B}^* := \mathcal{A}^*/\mathcal{G}$. The group $\mathcal{G}$ acts freely (mod $\pm\mathrm{Id}$) on $\mathcal{A}^*$, and the quotient $\mathcal{B}^*$ is a smooth, infinite dimensional manifold whose tangent space at each point may be identified with $\Omega^1(M,\mathfrak{su}(2))/d_A\Omega^0(M,\mathfrak{su}(2))$.

3.3. **Curvature.** The *curvature* of a connection A is $F_A := dA + A\wedge A$, an algebraic operator on sections of E and on $\Omega^*(M,\mathfrak{su}(2))$. As operators, $d_Ad_A = F_A$ since

$$\begin{aligned} d_Ad_AB &= d_A(dB + A\cdot B) \\ &= ddB + A\cdot dB + dA\cdot B - A\cdot dB + A\wedge A\cdot B \\ &= F_A\cdot B \end{aligned}$$

The (differential) *Bianchi identity* is the identity $d_AF_A = 0$. The gauge group acts on curvature by $g\cdot F_A = F_{g\cdot A} = gF_Ag^{-1}$, as it should because F_A is a tensor.

A connection is *flat* if $F_A = 0$. At a flat connection, $\Omega^*(M, \mathfrak{su}(2))$ becomes a complex with respect to d_A, with de Rham homology $H^*_{dR}(M, \mathfrak{su}(2))$.

For any $a, A \in \mathcal{A}$ we have $F_{A+ta} = F_A + td_A a + t^2 a \wedge a$. Therefore the tangent space to the space of flat connections at A is $\ker d_A$, and we may identify $H^1_{dR}(M, \mathfrak{su}(2))$ with the tangent space to the the space of flat connections modulo gauge equivalence.

A connection is flat if and only if holonomy transport is invariant under homotopy. In this case one obtains a representation $\rho : \pi_1(M) \to \mathrm{SU}(2)$ by parallel transporting a fiber around loops in the base. The gauge group acts on this representation by conjugating the fiber, which gives a conjugate representation. In other words: flat connections mod $\mathcal{G}$ are in natural bijection with SU(2) representations mod conjugacy. At the level of tangent spaces, this is reflected in the (de Rham) isomorphism of cohomology groups $H^1_{dR}(M, \mathfrak{su}(2)) = H^1_\rho(\pi_1(M), \mathfrak{su}(2))$.

If we write M as a Heegaard splitting $M = H_1 \cup_\Sigma H_2$ then $H^1_\rho(\pi_1(M), \mathfrak{su}(2)) = 0$ if the $X^*(H_i)$ intersect transversely at ρ. Informally we call such a ρ a *smooth point* on $X^*(M)$.

3.4. **Chern–Simons functional.** The *Chern–Simons functional* is a function $\mathbf{cs} : \mathcal{A} \to \mathbb{R}$, defined by

$$\mathbf{cs}(A) := \frac{1}{4\pi} \int_M \mathrm{tr}(A \wedge dA + \frac{2}{3} A \wedge A \wedge A)$$

Let A be a connection, and let $a \in T_A\mathcal{A} = \mathcal{A}$ be a tangent vector at A. Then

$$\begin{aligned} d\mathbf{cs}(a) &= \frac{1}{4\pi} \int_M \mathrm{tr}(a \wedge dA + A \wedge da + 2a \wedge A \wedge A) \\ &= \frac{1}{4\pi} \int_M 2\,\mathrm{tr}(a \wedge (dA + A \wedge A)) + \frac{1}{4\pi} \int_M d\,\mathrm{tr}(a \wedge A) \\ &= \frac{1}{2\pi} \int_M \mathrm{tr}(a \wedge F_A) \end{aligned}$$

In particular, A is a critical point for $\mathbf{cs}$ if and only if $F_A = 0$; i.e. if and only if A is a flat connection.

Recall that the tangents to the action of the gauge group are of the form $-d_A B$ for $B \in \Omega^0(M, \mathfrak{su}(2))$. For vectors of this form,

$$\begin{aligned} d\mathbf{cs}(-d_A B) &= \frac{1}{4\pi} \int_M \mathrm{tr}(-dB \wedge F_A - AB \wedge F_A + BA \wedge F_A) \\ &= \frac{1}{4\pi} \int_M \mathrm{tr}(BdF_A + B(A \wedge F_A - F_A \wedge A)) + \frac{1}{4\pi} \int_M d\,\mathrm{tr}(-BF_A) \\ &= \frac{1}{4\pi} \int_M \mathrm{tr}(Bd_A F_A) = 0 \text{ by the Bianchi identity} \end{aligned}$$

Thus $\mathbf{cs}$ is invariant under the connected component of the identity of $\mathcal{G}$. It is not quite invariant under the whole gauge group; there is an identification $\pi_0(\mathcal{G}) = \mathbb{Z}$ where the map is given by thinking of an element $g \in \mathcal{G}$ as a map between closed oriented 3-manifolds $M \to \mathrm{SU}(2)$, which has a degree $\deg(g) \in \mathbb{Z}$. The function $\mathbf{cs}$ transforms under the full gauge group by $\mathbf{cs}(g \cdot A) = \mathbf{cs}(A) + 2\pi \deg(g)$. Thus, $\mathbf{cs}$ is well-defined as a function on $\mathcal{B}$ with values in the circle $\mathbb{R}/2\pi\mathbb{Z}$, and with critical points precisely at gauge equivalence classes of flat connections.

Now suppose we fix a metric on M. Hodge star acts on forms on M, and we can extend it trivially to the $\mathfrak{su}(2)$ factor to define

$$* : \Omega^p(M, \mathfrak{su}(2)) \to \Omega^{3-p}(M, \mathfrak{su}(2))$$

Thus we obtain an inner product on $\mathcal{A}$ by

$$\langle A, B\rangle := \int_M \mathrm{tr}(A \wedge *B)$$

If we think of this inner product as a kind of formal Riemannian metric on $\mathcal{A}$, the 'gradient vector field' $\mathrm{grad}(\mathbf{cs})$ is (up to a constant) just $*F_A$. Since the gauge group acts on $\Omega^*(M, \mathfrak{su}(2))$ just by conjugation on the $\mathfrak{su}(2)$ factor, and since trace is invariant under conjugation, this metric is invariant under the gauge group and $\mathrm{grad}(\mathbf{cs})$ descends to a flow on $\mathcal{B}^*$.

3.5. **Instanton Homology.** Instanton Homology $I_*(M)$ is the homology theory obtained by thinking of $\mathbf{cs}$ as a Morse function on the space $\mathcal{B}^*$ of irreducible connections up to gauge equivalence. Thus — morally speaking — the chain groups are freely generated by conjugacy classes of irreducible flat connections, and the differentials are computed by counting flowlines of $\mathrm{grad}(\mathbf{cs})$ that run between critical points of adjacent index.

The most serious obstacle to making this idea meaningful is that the critical points of $\mathbf{cs}$ do not have a well-defined index. Let's suppose for simplicity that A is an irreducible flat connection whose class $[A]$ in $\mathcal{B}^*$ is a critical point for $\mathbf{cs}$, and let's determine the Hessian of $\mathbf{cs}$ at $[A]$.

Define $A_{s,t} := A + sa + tb$ and compute

$$\begin{aligned}\mathrm{Hess}(\mathbf{cs})(a, b) &= \frac{\partial^2}{\partial s \partial t}\Big|_{s=t=0} \frac{1}{4\pi} \int_M \mathrm{tr}(A_{s,t} \wedge dA_{s,t} + \frac{2}{3} A_{s,t} \wedge A_{s,t} \wedge A_{s,t}) \\ &= \frac{1}{2\pi} \int_M \mathrm{tr}(a \wedge d_A b)\end{aligned}$$

Using the inner product and the identification $T_A\mathcal{A} = \mathcal{A}$ we can therefore identify the Hessian (up to a constant) with the operator $*d_A$. Thus at a flat connection $[A] \in \mathcal{B}^*$ the Hessian is degenerate on $T_{[A]}\mathcal{B}^*$ precisely on the quotient space $H^1_{dR}(M, \mathfrak{su}(2))$, which is zero when the holonomy representation is a smooth point of $X^*(M)$.

The operator $*d_A$ is formally self-adjoint on $\Omega^1/d_A\Omega^0$ and has discrete spectrum, with eigenspaces of finite multiplicity. However, there are infinitely many eigenvalues of either sign, and no straightforward way to define an index.

Nevertheless it is possible to make sense of the *difference* of the index of two critical points, or at least to understand when this difference is equal to 1. By analogy with the finite dimensional story, and at least when critical points are nondegenerate, we expect that two equivalence classes of flat connections A, B have an index differing by 1 if the dimension of the space of flowlines of $\mathrm{grad}(\mathbf{cs})$ joining B to A is zero dimensional.

3.6. **Self-duality equation.** A flowline of $\mathrm{grad}(\mathbf{cs})$ joining B to A is an equivalence class of map $a : \mathbb{R} \to \mathcal{A}$ satisfying $da/dt = *F_{a(t)}$ with $a(t)$ converging to A or B as t goes to $\pm\infty$, up to the ambiguity of the action of the gauge group, and translation $t \to t + \text{constant}$.

Such a flowline has a natural interpretation in terms of gauge theory on the 4-manifold $M \times \mathbb{R}$ with the product metric. Let t be the coordinate in the $\mathbb{R}$ direction. Let $\overline{E}$ be a trivial $\mathbb{C}^2$ bundle over $M \times \mathbb{R}$, and let $\bar{a}$ be an $\mathfrak{su}(2)$-connection on this bundle. Suppose we write $\bar{a} = f dt + a_M$ where f is a function taking values in $\mathfrak{su}(2)$, and a_M has no component in the $\mathbb{R}$ direction. Let $g : M \times \mathbb{R} \to \mathrm{SU}(2)$ solve the equation $(\partial g/\partial t)g^{-1} = -gfg^{-1}$. Then by applying g as a gauge transformation we can replace $\bar{a}$ by a gauge-equivalent connection with $f = 0$.

The *self-duality equation* for an $\mathfrak{su}(2)$-connection A on a Riemannian 4-manifold W is the equation $*F_A = F_A$. For $W = M \times \mathbb{R}$ and $\bar{a} = a_M$ as above, then if $\alpha = \partial\bar{a}/\partial t$, we have on each fixed t_0 slice the equation $d\bar{a} = dt \wedge \alpha + d(\bar{a}(t_0))$ and therefore $F_{\bar{a}} = dt \wedge \alpha + F_{\bar{a}(t_0)}$ where the second term can be thought of as the curvature of the connection $\bar{a}(t_0)$ on the 3-manifold $M \times t_0$. For any form β on $M \times \mathbb{R}$ with no component in the t direction, we let $*^3$ denote the result of applying (3-dimensional) Hodge star to the restriction of β to each $M \times t_0$ slice. Then $*\beta = dt \wedge *^3\beta$ and therefore

$$*F_{\bar{a}} = *^3\alpha + dt \wedge *^3 F_{\bar{a}(t_0)}$$

Thus $\bar{a}$ solves the self-duality equation if and only if $\partial\bar{a}/\partial t = *^3 F_{\bar{a}(t_0)}$, which is precisely the statement that $\bar{a}(t)$ is a flowline of $\mathrm{grad}(cs)$.

Donaldson [4] famously showed used the space of solutions to the self-duality equation to obtain constraints on the topology of smooth 4-manifolds. On a noncompact 4-manifold it is important to add the constraint that the curvature $F_{\bar{a}}$ has finite energy (as measured by the L^2 norm). This implies that the connection $\bar{a}$ is asymptotically flat, and in fact must converge (up to gauge equivalence) to flat connections A and B at infinity.

In words: flowlines of $\mathrm{grad}(\mathbf{cs})$ joining equivalence classes of flat connections A and B correspond to finite energy solutions to the self-duality equations on $M \times \mathbb{R}$ up to gauge equivalence and translation. Gauge equivalence classes of solutions are called *instantons*, and their moduli space $\mathcal{M}$ decomposes into spaces $\mathcal{M}(A, B)$ for A, B as above. If A and B are nondegenerate and nontrivial and a further technical transversality condition holds, then $\mathcal{M}(A, B)$ is a smooth manifold on which $\mathbb{R}$ acts freely by translations.

It turns out that different components of $\mathcal{M}(A, B)$ can have different dimensions — but that these dimensions agree mod 8(!) and therefore we can define the *relative index* of A and B to be this dimension (counted mod 8).

Still assuming nondegeneracy, we can define a chain complex freely generated by nontrivial flat connections up to gauge equivalence, so that the coefficient of B in the differential ∂A is the number of one-dimensional components of $\mathcal{M}(A, B)$, counted with sign. The homology of this complex is $I_*(M)$. It does not depend on the choice of metric on M, and its Euler characteristic is equal to $2\lambda(M)$.

Making sense of this construction without assuming nondegeneracy is highly technically involved, and is carried out in the papers of Floer and Taubes cited above.

3.7. **The meaning of 8.** Where does 8 come from? The simplest nontrivial example of an $\mathfrak{su}(2)$ connection with self-dual curvature is as follows. If we think of $\mathbb{R}^8$ as the quaternionic plane, there is a 'Hopf map' from the unit sphere S^7 to the quaternionic projective line S^4 with fibers the unit quaternions $S^3 = \mathrm{SU}(2)$. In other words, we can think of S^7 as a principal $\mathrm{SU}(2)$ bundle over S^4 with $c_2 = 1$.

If we give S^7 its standard Riemannian metric, the orthogonal complements to the fibers gives a connection on S^7, thought of as an SU(2) bundle over S^4, and the curvature of this connection is self-dual.

The self-dual equations are conformally invariant, so we can push this solution around by the conformal group $\mathrm{SO}(5,1)$ of S^4. The stabilizer of a solution is the compact subgroup SO(5), and therefore we get a 5-dimensional space of solutions parameterized by $\mathrm{SO}(5,1)/\mathrm{SO}(5)$, which may be identified with 5-dimensional hyperbolic space (topologically an open 5-ball). This space may be compactified by gluing back S^4 itself; these extra points parameterize limits of instantons whose curvature is concentrated closer and closer to a point.

Now, on any smooth closed 4-manifold W with a principal SU(2) bundle E, and for any instanton, Taubes [19] shows how to 'insert' one of these 'limit' instantons near any point and perturb the result to get an honest new self-dual connection. The insertion changes the bundle E to E' with $c_2(E') = c_2(E)+1$, and depends on 8 parameters — 4 parameters for the point of insertion, 1 parameter for a 'scale' factor (from the conformal invariance), and an extra 3 parameters for an element of SU(2) measuring the difference in gauge at the gluing. This 'explains' why $\dim \mathcal{M}(E') = \dim \mathcal{M}(E)+8$ (a formal justification follows from the Atiyah-Singer index theorem which gives a formula $\dim \mathcal{M}(E) = 8c_2 - 3(1-b_1+b_2^+)$).

If $W = M \times \mathbb{R}$ the Chern class c_2 vanishes, and the components of $\mathcal{M}(A,B)$ of different dimensions correspond to paths of connections in $\mathcal{B}^*$ in different homotopy classes rel. endpoints.

4. Heegaard Floer Homology

4.1. **Lagrangian Intersection Homology.** Floer's second construction [6] of a homology theory whose Euler characteristic is 2λ makes use of the symplectic structure on $X^*(\Sigma)$, and is really a homology theory for pairs of Lagrangians in a symplectic manifold in general.

Let P,ω denote a symplectic manifold, and let L_1, L_2 denote Lagrangian submanifolds. Thus if the dimension of P is $2n$, the L_i are smooth submanifolds of dimension n. Let Ω denote the space of smooth maps $z : I \to P$ from L_1 to L_2. In any sufficiently small neighborhood $z_0 \in U \subset \Omega$ of a point z_0 we can define a function $a : U \to \mathbb{R}$ as follows. For any other $z \in U$ we can join z_0 to z by a path z_t in Ω, which sweeps out a rectangle $F \subset P$ with left and right edges on L_1 and L_2 respectively. Then set $a(z) := \int_F \omega$. This is well-defined on a sufficiently small neighborhood U since if z_t' is another path sweeping out another rectangle F', then $F \cup F$ is a cylinder whose ends can be capped off by small disks D_i in L_i to make a sphere S, and

$$\int_F \omega - \int_{F'} \omega = \int_S \omega = 0$$

because ω is closed, and S is null-homotopic for sufficiently small U.

In general the indeterminacy in a is generated by the periods of ω on cylinders interpolating between L_1 and L_2. Such cylinders are given by intersections of conjugacy classes $\pi_1(L_1) \cap g\pi_1(L_2)g^{-1}$ up to the action of $\pi_2(P)$. Thus (for example) if $H_1(L_1) \cap H_1(L_2) = 0$ and $\pi_2(P)$ is trivial, we may define a globally.

In any case, the differential da is well-defined. The tangent space $T_z\Omega$ is just the space of vector fields ξ along z, and

$$da(\xi) = \int_0^1 \omega(\dot{z}(t), \xi)dt$$

Since ω is nondegenerate, the critical points of a are the constant maps; i.e. the points of $L_1 \cap L_2$.

4.1.1. *Holomorphic Whitney disks.* To define a metric on Ω and hence make sense of $\text{grad}(a)$ and gradient flowlines, we choose an almost complex structure — i.e. an endomorphism J of the tangent bundle squaring to $-\text{Id}$ and preserving the symplectic form — for which $g := \omega(\cdot, J\cdot)$ is positive definite (i.e. it defines a metric on P). Then we get a metric on Ω by integrating: for ξ_1, ξ_2 vector fields along z,

$$\langle \xi_1, \xi_2 \rangle := \int_0^1 \omega(\xi_1, J\xi_2)$$

Thus

$$da(\xi) = \int_0^1 \omega(\dot{z}, \xi) = \int_0^1 \omega(J\dot{z}, J\xi) = \langle J\dot{z}, \xi \rangle$$

so that we may write (formally) $\text{grad}(a) = J\dot{z}$. Up to sign, 'trajectories' of $\text{grad}(a)$ are maps $u : I \times I \to P$ where (if we give the I factors coordinates s and t)

$$\frac{\partial u}{\partial s} + J\frac{\partial u}{\partial t} = 0$$

which is simply the statement that u is a holomorphic map with respect to the (almost) complex structure.

If $x, y \in L_1 \cap L_2$ then flowlines from x to y are the holomorphic maps of the unit disk $u : D \to P$ sending $-i$ to x and i to y, and such that the arc of ∂D with positive (resp. negative) real part maps to L_1 (resp. L_2). We call such a map a *holomorphic (Whitney) disk*. The disk admits a real 1-dimensional family of holomorphic automorphisms fixing $-i$ and i, and acting as a translation of the hyperbolic geodesic joining these points; thus the space $\mathcal{M}(x, y)$ of flowlines between x and y admits a free $\mathbb{R}$ action, and has dimension at least 1 if it is nonempty.

Then (modulo technical difficulties! — see § 5.3) one defines a homology theory whose chain groups are freely generated by intersections $L_1 \cap L_2$, and whose differentials count 1-dimensional components of $\mathcal{M}(x, y)$ with a certain sign (or with $\mathbb{Z}/2\mathbb{Z}$ coefficients, ignoring sign). Specializing to the case that $P = X^*(\Sigma)$ and $L_i = X^*(H_i)$ one recovers 2λ as the Euler characteristic of the theory.

4.1.2. *Maslov Index.* As with the case of Instanton homology, the moduli space $\mathcal{M}(x, y)$ can have components of different dimension in general. This dimension is a relative index called the *Maslov Index*, and depends on a choice of a path in Ω (i.e. a homotopy class of Whitney disk u) joining x to y. Somewhat informally, we denote the homotopy classes of such maps by $\pi_2(x, y)$.

Associated to the real vector space $\mathbb{R}^n$ one has the cotangent bundle $T^*\mathbb{R}^n$ and the complexification $\mathbb{C}^n$. The former has a natural symplectic structure, and the latter a natural complex structure. The standard Euclidean inner product on $\mathbb{R}^n$ extends to the standard Hermitian inner product on $\mathbb{C}^n$, and we may choose an identification of $\mathbb{C}^n$ with $T^*\mathbb{R}^n$ as real vector spaces compatibly with these metrics. Under this identification, $\mathrm{Sp}(2n,\mathbb{R})\cap \mathrm{O}(2n) = \mathrm{U}(n)$. The group $\mathrm{U}(n)$ acts transitively on the set $\Lambda(n)$ of Lagrangian subspaces of $\mathbb{R}^{2n}$, and the stabilizers are conjugates of the subgroup $\mathrm{O}(n)$. Thus we have an identification $\Lambda(n) = \mathrm{U}(n)/\mathrm{O}(n)$. Taking the square of the (complex) determinant defines a fibration $\det^2 : \mathrm{U}(n)/\mathrm{O}(n) \to S^1$. If we pull back the (oriented) generator of $H^1(S^1)$ we get a generator μ of $H^1(\Lambda(n)) = \mathbb{Z}$.

If V is a linear Lagrangian subspace of $\mathbb{R}^{2n}$, Arnol'd [3] defines the *train* $T(V)$ to be the set of Lagrangian subspaces whose intersection with V is nontrivial. Each train is Poincaré dual to μ, and the complement $\Lambda(n) - T(V)$ is simply-connected. Thus, if V_1, V_2 are any two linear Lagrangian subspaces of $\mathbb{R}^{2n}$, there is a canonical homotopy class of path in $\Lambda(n)$ from V_1 to V_2 that does not cross $T(V)$.

Explicitly, we can find a unitary matrix $U \in \mathrm{U}(n)$ with $U(V_1) = V_2$, and by multiplying by a suitable element of $\mathrm{O}(n)$ if necessary we can assume the eigenvalues of U are of the form $e^{i\theta_j}$ for n numbers $\theta_j \in [0,\pi)$. Thus normalized, we can think of U as the endpoint of a path of unitary matrices U_t with the same eigenspaces, and eigenvalues $e^{it\theta_j}$.

We are now in a position to define the Maslov index, following Viterbo [22]. Let $u : D \to P$ be a Whitney disk joining x to y. Since D is contractible, the pullback u^*TP has a (symplectic) trivialization as $D \times \mathbb{R}^{2n}$. The circle $S^1 = \partial D$ factorizes as the union of two arcs $\alpha_1 \cup \alpha_2$ where $u : \alpha_i \to L_i$, and where we orient the arcs so that each $u(\alpha_i)$ runs from x to y. Using the trivialization of u^*TP we can think of each $\alpha_i^*TL_i$ as a path in $\Lambda(n)$. Join the endpoints of $\alpha_1^*TL_1$ to the endpoints of $\alpha_2^*TL_2$ by the procedure above to get a quadrilateral $\gamma_u : S^1 \to \Lambda(n)$. The *Maslov Index*, is equal to $\mu(\gamma_u)$ where μ is the oriented generator of $H^1(\Lambda(n))$. By abuse of notation we denote this $\mu(u)$. Evidently $\mu(u)$ depends only on the class of u in $\pi_2(x,y)$.

Here is an equivalent procedure. Since the inclusion of α_1 into D is a homotopy equivalence, we may choose a symplectic trivialization of our bundle which restricts to a trivialization over α_1 for which $TL_1|\alpha_1$ is constant, and TL_2 is the same (perpendicular) subspace at x and y. Then relative to this trivialization, $TL_2|\alpha_2$ is a *closed* loop in $\Lambda(n)$, whose winding number is $\mu(u)$.

If classes $u, u' \in \pi_2(x,y)$ have homotopic boundary values, they glue together (up to homotopy) to make a map $\phi := u \cup -u' : S^2 \to P$.

Proposition 4.1. *With notation as above, there is a formula*

$$\mu(u) - \mu(u') = 2c_1(P)[\phi]$$

Proof. The pullback ϕ^*TP is a symplectic bundle over S^2, which is classified by c_1. Thus the difference of $\mu(u)$ and $\mu(u')$ is proportional to $c_1(P)[\phi]$, where $[\phi]$ denotes the image of ϕ under the Hurewicz map. To compute c_1 we look at the winding number of det on a clutching function for the trivializations of ϕ^*TP over u and $-u$. Since the fibration $\Lambda(n) \to S^1$ is given by the square of the determinant, the constant of proportionality is 2. □

Viterbo [22] shows that the index so defined is the formal dimension of the moduli space $\mathcal{M}(x,y)$ of pseudo-holomorphic curves in the homotopy class of u. We shall give a justification for this fact in § 5.2. Further work is necessary to show that for suitable perturbations of the complex structure (see § 5.3) these moduli spaces are smooth manifolds of the desired dimension.

4.2. **Definition of Heegaard Floer Homology.** Heegaard Floer Homology as defined by Ozsváth–Szabó [13, 14] is a version of Lagrangian Intersection Homology in the sense of § 4.1 adapted to a symplectic manifold (P, ω) and a pair of transverse Lagrangian submanifolds L_1, L_2 associated to a decorated Heegaard splitting of a 3-manifold. It is very similar to the setup for the Casson invariant, in that P is (very nearly) the character variety of $\pi_1(\Sigma)$, except that we take representations into the abelian group $\mathrm{U}(1)$ instead of $\mathrm{SU}(2)$.

If we give Σ the structure of a Riemann surface, then the Jacobian variety $J\Sigma$ is topologically equivalent to $\mathrm{Hom}(\pi_1(\Sigma), \mathrm{U}(1))$. For Heegaard Floer Homology one works not with $J\Sigma$, but with a birationally equivalent space $S^g\Sigma$.

4.2.1. *Heegaard diagrams.* Let H be a genus g handlebody. An *attaching set* α is a collection of g essential simple closed curves α_i in $\Sigma := \partial H$ satisfying

(1) the α_i are disjoint;
(2) the complement $\Sigma - \alpha$ is connected; and
(3) the α_i bound disjoint embedded disks D_i in H.

The handlebody H can be recovered from the pair (Σ, α) by thickening Σ, attaching 2-handles along the α_i, and capping off the spherical boundary that results with a 3-ball. We should therefore think of an attaching set as a presentation for H.

If we fix a handlebody, any two attaching sets α, α' presenting H are related up to isotopy by a finite sequence of *handle slides*: choose a pair α_i, α_j and an embedded arc $\delta \subset \Sigma$ running from α_i to α_j and interior disjoint from α, and slide α_i along δ and push it over the disk D_j. This replaces α_i by α_i', where the triple $\alpha_i, \alpha_j, \alpha_i'$ bound an embedded pair of pants in the complement of the other α_k.

Let M be a 3-manifold. A *Heegaard diagram* for M is a surface Σ together with two attaching sets α, β presenting handlebodies H_1, H_2 so that $H_1 \cup_\Sigma H_2$ is a Heegaard splitting of M. Usually we insist that α and β intersect in general position.

Any two Heegaard splittings of M have a common stabilization. At the level of Heegaard diagrams, stabilization is achieved by taking a torus T with basis curves m, l, connect summing Σ to T, and adding m to the collection α and l to the collection β.

Proposition 4.2 (Heegaard diagram)**.** *Any two Heegaard diagrams of the same 3-manifold are related by a sequence of one of three basic moves:*

(1) isotopy of one of α or β;
(2) handle slide of one of α or β; and
(3) stabilization or its inverse.

Geometrically, a genus g Heegaard diagram for M corresponds to a self-indexing Morse function f on M with one minimum, one maximum, and Σ a level set between the index 1 and 2 critical points. The index 1 points are all in H_1, the index 2 points are in H_2, the α are the intersection of the ascending manifolds of the index 1 points with Σ, and the β are the intersection of the descending manifolds of the index 2 points with Σ. Proposition 4.2 is thus a restatement of some standard facts from Morse theory.

4.2.2. *Pointed Heegaard diagrams.* A *pointed* Heegaard diagram is a Heegaard diagram (Σ, α, β) together with the choice of a basepoint $z \in \Sigma - (\alpha \cup \beta)$. A *pointed isotopy* of α or β is one in which the curves stay disjoint from the basepoint. A *pointed handle slide* is a handle slide for which the pair of pants region cobounding the sliding circles is disjoint from the basepoint.

It is not hard to show that any two pointed Heegaard diagrams representing the same 3-manifold are related by pointed isotopy, pointed handle slides and stabilization.

4.2.3. *Symmetric Products.* When we defined the Casson invariant, we associated three spaces $X^*(\Sigma), X^*(H_1), X^*(H_2)$ to a Heegaard diagram, and derived the invariant from the configuration of the $X^*(H_i)$ in $X^*(\Sigma)$.

Heegaard Floer Homology is defined analogously. If Σ has genus g, the spaces that take the analog of the character varieties X^* are, roughly speaking, configuration spaces of g unordered points.

Given a smooth surface Σ of genus g, define the *symmetric product* $S^g\Sigma$ to be the quotient of the g-fold product Σ^g by the full symmetry group of the factors. This is a smooth manifold; in fact, if we give Σ the structure of an algebraic curve, then $S^g\Sigma$ is a complex projective variety. As a complex projective variety, the complex structure depends on that of Σ, but as a symplectic manifold, it is independent of choices.

The symmetric product is important in algebraic geometry: $S^g\Sigma$ is birationally equivalent to the Jacobian variety $J\Sigma$ of Σ. If we fix g linearly independent holomorphic 1-forms $\omega_1, \cdots, \omega_g$ and a basepoint $z \in \Sigma$ then there is the Abel-Jacobi map $u : \Sigma \to J\Sigma$ obtained by taking each point $p \in \Sigma$ to the vector of integrals $\int_z^p \omega_i$, which is well-defined modulo the period lattice. The map u extends to symmetric powers by using the additive group structure on $J\Sigma$; i.e. for $(p_1, \cdots, p_k) \in S^k\Sigma$ define $u(p_1, \cdots, p_k) = \sum u(p_i)$.

It can be shown that at a generic point, the map $u : S^g\Sigma \to J\Sigma$ is a submersion, and is therefore a birational isomorphism. The preimage of any point in $J\Sigma$ is a complete linear system, and therefore a projective space of some dimension (generically 0). Since topologically, $J\Sigma$ is a complex torus of dimension g, it follows that $\pi_1(S^g\Sigma)$ is abelian, and equal to $H_1(S^g\Sigma) = H_1(\Sigma) = \mathbb{Z}^{2g}$ (this is easy to see directly, since $\pi_1 S^n X$ is abelian for any topological space X and any $n > 1$).

The homology of Σ^g can be computed from the Künneth formula, and the homology of $S^g\Sigma$ is the part invariant under the symmetric group. Thus $H_2(S^g\Sigma)$ has dimension $\binom{2g}{2} + 1$, freely generated by $\Lambda^2 H_1$ and the image of $H_2(\Sigma)$ included as a factor. Since $H_2(\pi_1) = \Lambda^2 H_1$ it follows that the image of π_2 under the Hurewicz map is $\mathbb{Z}$.

Example 4.3 (Low genus examples). For $g = 1$, $S^1\Sigma = J\Sigma = \Sigma$, a complex 1-torus.

For $g = 2$, $S^2\Sigma$ is a blow-up of $J\Sigma$ at one point. The exceptional curve of the blow-up is the image of $S := (y, \tau(y))$ where τ is the hyperelliptic involution of Σ. Thus $\pi_1 = \mathbb{Z}^4$ and π_2 is freely generated by S as a π_1-module. By decomposing the tangent bundle along S into the tangent and normal parts, we get the formula

$$c_1(S^2\Sigma)[S] = \chi(S) + [S] \cap [S] = 2 - 1 = 1$$

For $g = 3$, if Σ is canonical, the exceptional locus in $J\Sigma$ is a copy of Σ itself. Each point in this copy is blown up to a 2-sphere S. The blow-up locus is surjective on π_1, so $\pi_2 = \mathbb{Z}$ and is trivial as a π_1-module. We have $c_1(S^3\Sigma)[S] = 1$ as before.

In algebraic geometry, the blowup locus in $J\Sigma$ (i.e. the locus of 'special divisors') depends on Σ, and is studied by Brill–Noether theory. However, the topological picture is substantially simpler, at least in low dimensions. For $g > 2$ one has $\pi_2(S) = \mathbb{Z}$ generated by $S := (y, \tau(y), z, \cdots, z)$ where τ is a hyperelliptic involution, and satisfies $c_1(S^g\Sigma)[S] = 1$ (by essentially the same calculation as above, since the extra z coordinates give trivial summands of the tangent bundle along S). The involution τ always exists topologically, but not holomorphically unless Σ is hyperelliptic.

4.2.4. *Lagrangian tori.* If α is an attaching set of curves in Σ, we define the torus $T_\alpha \subset S^g\Sigma$ to be the (unordered) product $\alpha_1 \times \cdots \times \alpha_g$. Evidently this is a Lagrangian subspace of $S^g\Sigma$.

Now suppose $(\Sigma, \alpha, \beta, z)$ is a pointed Heegaard diagram for M. Since α and β are in general position, so are the tori T_α and T_β; i.e. they meet transversely in finitely many points.

The basepoint z determines a subspace $V_z := z \times S^{g-1}\Sigma_g$. Under the Abel-Jacobi map, V_z is the preimage of the Theta divisor Θ. It is evidently disjoint from T_α and T_β.

Now, $H_1(M)$ is the quotient of $H_1(\Sigma)$ by the subspaces spanned by α and β. Under the identification of $H_1(\Sigma)$ with $H_1(S^g\Sigma)$ these are the subspaces $H_1(T_\alpha)$ and $H_1(T_\beta)$ respectively. Thus $H_1(S^g\Sigma)/(H_1(T_\alpha) + H_1(T_\beta)) = H_1(M)$.

For any two points $x, y \in T_\alpha \cap T_\beta$ we can choose paths $a \subset T_\alpha$ and $b \subset T_\beta$ from x to y. Then $a - b$ is a 1-cycle, and the image of its homology class $\epsilon(x, y)$ in $H_1(M)$ (under the identification above) is well-defined and independent of the paths a and b. The points of $T_\alpha \cap T_\beta$ fall into equivalence classes with $a \sim b$ if $\epsilon(a, b) = 0$.

4.2.5. *Whitney disks.* Fix two points $x, y \in T_\alpha \cap T_\beta$. A *Whitney disk* for the pair x, y is a map $u : D \to S^g\Sigma$ with $u(-i) = x$ and $u(i) = y$, and such that the arc of ∂D with positive (resp. negative) real part maps under u to T_α (resp. T_β). The set of homotopy classes of Whitney disks for x, y is denoted $\pi_2(x, y)$. It is a module over $\pi_2(S^g\Sigma) = \mathbb{Z}$ when $g \geq 3$ and over $\pi_2'(S^g\Sigma) := \pi_2(S^g\Sigma)/\pi_1(S^g\Sigma) = \mathbb{Z}$ if $g \geq 2$. Note that $\pi_2(x, y)$ is empty if $\epsilon(x, y) \neq 0$.

Proposition 4.4. *If $\epsilon(x, y) = 0$ then $\pi_2(x, y)$ is not empty, and if $g \geq 3$ it's isomorphic to $\mathbb{Z} \oplus H^1(M; \mathbb{Z})$.*

Proof. A pair of arcs $\gamma_i \subset T_i$ from x to y whose union is null-homologous bounds a map from a surface to $S^g\Sigma$. Since $\pi_1(S^g\Sigma)$ is abelian, this surface can be compressed to a (Whitney) disk. Thus if $\epsilon(x, y) = 0$ then $\pi_2(x, y)$ is nonempty.

Now let $\phi_1, \phi_2 \in \pi_2(x, y)$ be two classes. If we glue the domain disks D_i of ϕ_i along $x \cup y$ we obtain a (pinched) cylinder A together with a map $\Phi : A \to S^g\Sigma$ interpolating between ϕ_1 and ϕ_2. The difference in homotopy classes ϕ_1, ϕ_2 is measured by the homotopy class of $(A, \partial^+ A, \partial^- A) \to (S^g\Sigma, T_\alpha, T_\beta)$.

The homotopy class of the boundary terms are determined by their homology, because the $\pi_1(T_i)$ are abelian. The images of $\partial^\pm A$ are loops $\ell_i \in T_i$ which define homology classes $[\ell_i] \in H_1(T_i)$ with the same image in $H_1(S^g\Sigma)$. Using $H_1(S^g\Sigma) = H_1(\Sigma)$ we identify

$$[\ell_i] \in \ker : H_1(\Sigma) \to H_1(H_i)$$

Using the following fragment of the Mayer-Vietoris sequence

$$0 \to H_2(M) \to H_1(\Sigma) \to H_1(H_1) \oplus H_1(H_2)$$

the pair of classes $[\ell_i]$ together determine a class $\beta \in H^1(M) = H_2(M)$. Every such class arises, since a homology between loops in the T_i is realized by a map of a surface to $S^g\Sigma$, and $\pi_1(S^g\Sigma)$ is abelian so any such surface may be compressed to a cylinder. Thus we obtain a surjection $\pi_2(x, y) \to H^1(M; \mathbb{Z})$.

Any two annuli A, A' with the same boundary differ by an element of $\pi_2(S^g\Sigma) = \mathbb{Z}$ if $g \geq 3$, and the action of π_2 is free, because the generator is nontrivial in homology. □

When $g = 2$ there is still a surjective map $\pi_2(x, y) \to \mathbb{Z} \oplus H^1(M; \mathbb{Z})$ where we just record the (relative) homology class of a Whitney disk.

4.2.6. Spin^c *structures.* In each dimension n the group $\mathrm{Spin}^c(n)$ is the connected double cover of $\mathrm{SO}(n) \times S^1$ that unwraps each factor. A Spin^c structure on an oriented manifold is a lift of an $\mathrm{SO}(n)$ structure on the tangent bundle to a principal $\mathrm{Spin}^c(n)$ bundle compatible with the quotient map $\mathrm{SO}(n) = \mathrm{Spin}^c(n)/S^1$.

Since every oriented 3-manifold M is parallelizable, every M admits a Spin^c structure. $\mathrm{Spin}^c(3) = \mathrm{U}(2)$, and $\mathrm{SO}(3)$ can be thought of as the quotient of $\mathrm{U}(2)$ by the center. The subgroup $\mathrm{U}(1) \oplus 1$ sits in $\mathrm{U}(2)$ as diagonal matrices of the form $\left(\begin{smallmatrix} e^{i\theta} & 0 \\ 0 & 1 \end{smallmatrix}\right)$. If M has a Spin^c-structure, then after choosing a trivialization this circle subgroup acts by positively oriented rotations around some oriented axis in the tangent space at each point. The field of oriented axes is (up to scale) a nowhere zero vector field v on M.

An automorphism of a spin structure is a map from M to $\mathrm{U}(2)$ and the stabilizer of the field of $\mathrm{U}(1) \oplus 1$ subgroups is a map from M to $\mathrm{U}(1)$. Thus the ambiguity in the map from Spin^c structures to vector fields is parameterized by maps from M to $\mathrm{U}(2)/\mathrm{U}(1) \oplus 1$ which is homeomorphic to S^3. Two nowhere zero vector fields are related via this equivalence if and only if they are homotopic outside of a ball. This is because maps from M to S^3 are classified by degree, and therefore any two such maps are homotopic outside of a small ball.

Two nowhere zero vector fields homotopic outside a ball are said to be *homologous*, and the equivalence classes are called *Euler structures*. On a 3-manifold they are in natural bijection with Spin^c classes, as we have just shown. This observation is due to Turaev [21].

Since M is oriented, a nowhere zero vector field v determines a complex line bundle $v^\perp$ (well-defined up to homotopy). Changing v by a homology does not affect the isomorphism class of $v^\perp$, because the difference in the clutching maps is determined by a homotopy class of map from $\partial B^3 = S^2$ to $\mathrm{U}(1)$, which is necessarily trivial. This complex line bundle is naturally identified with the determinant line bundle of the $\mathrm{U}(2)$ bundle, since $\det : \mathrm{U}(1) \oplus 1 \to S^1$ is an isomorphism; in particular, it is naturally associated to the Spin^c structure $\mathfrak{s}$, and we denote its first Chern class by $c_1(\mathfrak{s})$.

If we choose a trivialization of M, a nowhere zero vector field v determines a map $\phi_v : M \to S^2$ relative to this trivialization. The pullback of the generator e of $H^2(S^2)$ is a class $\phi_v(e) \in H^2(M)$ Poincaré dual to the 1-manifold $X \subset M$ where v points to the north pole (relative to the trivialization). Changing the trivialization changes the class of ϕ_v by 2-torsion, but for any two vector fields v, w the difference $\phi_v(e) - \phi_w(e) \in H^2(M)$ is well-defined independent of the trivialization. Thus the set of Spin^c structures has a free transitive $H^2(M)$ action, so that this set is in bijection with $H^2(M)$ but not canonically.

In any trivialization, we can take the field of cross products of v with the north pole to obtain a section of $v^\perp$ zero exactly where v is vertical. Thus we obtain the formula

$$c_1(\mathfrak{s}) = c_1(v^\perp) = \phi_v(e) - \phi_{-v}(e)$$

Ozsváth-Szabó associate a Spin^c class $\mathfrak{s}(x)$ to every $x \in T_\alpha \cap T_\beta$ in such a way that $\mathfrak{s}(x) = \mathfrak{s}(y)$ iff $\epsilon(x,y) = 0$. Think of a Heegaard surface as a level set of a Morse function f separating the index $0,1$ from the index $2,3$ points. The α curves are the boundaries of ascending manifolds of the 1-handles, and the β curves are the boundaries of the descending manifolds of the 2-handles. Thus a point $x \in T_\alpha \cap T_\beta$ is a g-tuple of flowlines pairing the index 1 and index 2 critical points. The marked point z likewise can be thought of as a flowline pairing the index 0 and 3 critical points. Outside these flowlines, grad$f$ is a nowhere zero vector field, and because the flowlines run between critical points of opposite parity, grad$f$ may be extended to a nonsingular vector field over neighborhoods of these flowlines, thereby determining a unique Euler structure, which is to say a Spin^c class $\mathfrak{s}(x)$.

4.2.7. *The definition of* $\widehat{HF}$.

Definition 4.5 ($\mathcal{M}$ and $\widehat{\mathcal{M}}$)**.** For $\phi \in \pi_2(x,y)$, denote by $\mathcal{M}(\phi)$ the space of holomorphic Whitney disks in the homotopy class of ϕ. The group $\mathbb{R}$ acts on $\mathcal{M}(\phi)$ by reparameterizations of the disk, and this action is free unless ϕ is the trivial element of $\pi_2(x,x)$. Denote by $\widehat{\mathcal{M}}(\phi)$ the quotient $\mathcal{M}(\phi)/\mathbb{R}$ or the empty set if ϕ is trivial.

Under a suitable transversality condition, $\mathcal{M}(\phi)$ is a manifold of dimension equal to the Maslov index $\mu(\phi)$. Unfortunately, to achieve this transversality we need to work not with the given complex structure on $S^g\Sigma$, but with a generic almost complex structure. We ignore this issue for the moment, and return to it in § **??**. Thus $\dim(\widehat{\mathcal{M}}(\phi)) = \mu(\phi) - 1$.

The moduli spaces $\mathcal{M}$ and $\widehat{\mathcal{M}}$ may be oriented. If $\mu(\phi) = 1$ then $\widehat{\mathcal{M}}(\phi)$ is compact — i.e. it consists of finitely many (oriented) points, and we denote by $c(\phi)$ the signed count of these points. When $\mu(\phi) = 2$ the $\widehat{\mathcal{M}}(\phi)$ are 1-dimensional though not typically compact, but they can be compactified by products of lower dimensional $\widehat{\mathcal{M}}(\phi')$ as one expects from Morse theory.

Let S denote the holomorphic sphere that generates $\pi_2(S^g\Sigma)$ (up to the action of π_1 if $g = 2$). Then by Proposition 4.1 we have

$$\mu(\phi + kS) = \mu(\phi) + 2kc_1(S^g\Sigma)[S] = \mu(\phi) + 2k$$

Definition 4.6. For a homotopy class $\phi \in \pi_2(x,y)$ denote by $n_z(\phi)$ the algebraic intersection number of $\phi(D)$ with V_z.

If $\mathcal{M}(\phi)$ is nonempty, ϕ has holomorphic representatives, and since V_z is also holomorphic, necessarily $n_z \ge 0$ in this case. Note that V_z will not necessarily stay holomorphic for a generic almost complex structure, but we may choose such a structure sufficiently close to the integrable structure so that $n_z \ge 0$ still holds.

Furthermore, for S the holomorphic sphere as above, $S \cap V_z = 1$ so that $n_z(\phi + kS) = n_z(\phi) + k$.

We're now ready to give the definition of (the hat version of) Heegaard Floer Homology. To a pointed Heegaard diagram $(\Sigma, \alpha, \beta, z)$, associate a free abelian group $\widehat{CF}$ generated by the points of $T_\alpha \cap T_\beta$, and let $\widehat{CF}(\mathfrak{s})$ be the summand generated by x with $\mathfrak{s}(x) = \mathfrak{s}$. Thus $\widehat{CF} = \oplus_{\mathfrak{s}} \widehat{CF}(\mathfrak{s})$. For each $\mathfrak{s}$ define a *boundary map* $\partial : \widehat{CF}(\mathfrak{s}) \to \widehat{CF}(\mathfrak{s})$ by the formula

$$\partial x := \sum_{\mathfrak{s}(y)=\mathfrak{s}(x)} \sum_{\phi \in \pi_2(x,y):\mu(\phi)=1, n_z(\phi)=0} c(\phi) y$$

Then $\partial^2 = 0$, and the homology of the resulting complex is denoted $\widehat{HF} = \oplus_{\mathfrak{s}} \widehat{HF}(\mathfrak{s})$.

If M is a rational homology sphere, there is at most one class in each $\pi_2(x, y)$ with $\mu(\phi) = 1$, so this sum is finite. When M has nontrivial H^1 one must arrange somehow (e.g. by restricting the class of admissible Heegaard diagrams) that there are only finitely many classes $\phi \in \pi_2(x, y)$ with $\mu(\phi) = 1$ and $\widehat{\mathcal{M}}(\phi)$ nonempty; we defer a discussion of this point to § **??**.

If M is a rational homology sphere there is also a relative grading on the set of $x \in T_\alpha \cap T_\beta$ with $\mathfrak{s}(x) = \mathfrak{s}$, namely

$$\mathrm{gr}(x, y) := \mu(\phi) - 2n_z(\phi)$$

Equivalently, $\mathrm{gr}(x, y)$ is the value of the Maslov index $\mu(\phi)$ for the unique class $\phi \in \pi_2(x, y)$ with $n_z(\phi) = 0$.

The homology groups $\widehat{HF}$ depend on many choices. The chain groups depend on a choice of pointed Heegaard diagram, and the differentials depend on a generic choice of almost complex structure on $S^g\Sigma$. Nevertheless, it turns out that the homology $\widehat{HF}$ does not depend on these choices. In other words:

Theorem 4.7 (Oszváth–Szabó [13])**.** *The homology groups $\widehat{HF}(\mathfrak{s})$ are independent of choices, and are therefore an invariant of M.*

We shall sketch the proof of this in the sequel.

4.2.8. *The definitions of HF^*.* A refinement of the hat version of homology considers all holomorphic disks (not just those with $n_z = 0$). Define $CF^\infty(\mathfrak{s})$ to be the free abelian group generated by pairs $[x, i]$ where $\mathfrak{s}(x) = \mathfrak{s}$ and i is an integer. Define a relative grading

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

(this makes sense if M is a rational homology sphere). Define a boundary map

$$\partial[x, i] := \sum_{\mathfrak{s}(y)=\mathfrak{s}(x)} \sum_{\phi \in \pi_2(x,y):\mu(\phi)=1} c(\phi)[y, i - n_z(\phi)]$$

Then $\partial^2 = 0$ and the homology of the complex is denoted $HF^\infty = \oplus_{\mathfrak{s}} HF^\infty(\mathfrak{s})$.

Again, the formula for ∂ is manifestly a finite sum when M is a rational homology sphere, since there is still at most one class in each $\pi_2(x, y)$ with $\mu = 1$. Nevertheless as before it turns out one can arrange for general M for the sums to be finite, so that the homology groups are defined.

The chain groups CF^∞ admit an automorphism U of degree -2

$$U : [x, i] \to [x, i - 1]$$

compatible with the boundary operator; thus HF^∞ admits the structure of a $\mathbb{Z}[U,U^{-1}]$-module. It turns out (for rational homology spheres) it is always isomorphic to $\mathbb{Z}[U,U^{-1}]$. More subtle structure comes from the filtration arising from n_z.

Define CF^- to be the subgroup of CF^∞ freely generated by $[x,i]$ with $i<0$, and let CF^+ denote the quotient $CF^+ := CF^\infty/CF^-$. Since $n_z(\phi)\ge 0$ for any ϕ with $\widehat{\mathcal{M}}(\phi)$ nonempty, it follows that the coefficient of $[y,j]$ in $\partial[x,i]$ is zero unless $i\ge j$. Thus CF^- is a subcomplex of CF^∞ and we get a short exact sequence of complexes

$$0\to CF^-\to CF^\infty\to CF^+\to 0$$

The automorphism U takes CF^- into itself, and therefore induces an endomorphism of CF^+. With this notation, there is another short exact sequence of chain groups

$$0\to \widehat{CF}\to CF^+\xrightarrow{U} CF^+\to 0$$

Theorem 4.8 (Oszváth–Szabó [13])**.** *The homology groups $HF^\pm(\mathfrak{s})$ are independent of choices, and are therefore an invariant of M.*

5. Proofs

In this section we sketch the proofs of Theorems 4.7 and 4.8. We emphasize that these are just sketches. Furthermore, at several points we restrict attention to the hat-version of HF, and to the situation where $b_1(M)=0$ where this makes the arguments materially simpler. For full details see [?] (and, at some places, Floer [6, 7]).

5.1. Gromov's compactness theorem.

Example 5.1. For $\epsilon\in\mathbb{C}$ let C_ϵ be the curve $z_1z_2=\epsilon$ in $\mathbb{C}^2$. When $\epsilon\ne 0$ the curve C_ϵ is an annulus foliated by circles $|z|=$ constant. But as $\epsilon\to 0$ the C_ϵ degenerate to C_0, a pair of transverse complex lines. This is a model for the way in which a family of Riemann surfaces in a complex manifold can degenerate to a singular Riemann surface.

Ignoring the embedding, C_0 is obtained topologically from an annulus by collapsing the meridian circle to a point. If C is a smooth surface and γ_i is a collection of disjoint simple loops, we can obtain a singular surface $\bar{C}$ by collapsing each γ_i to a point p_i. Let $\hat{C}$ be obtained from $\bar{C}-\cup p_i$ by adding one point for each end. Then $\hat{C}$ is a closed surface, and there is a map $\hat{C}\to\bar{C}$ which is 2–1 on the preimages of the p_i, and is 1–1 elsewhere. A *complex structure* on $\bar{C}$ is simply a complex structure on $\hat{C}$. Equivalently, it is a complex structure on $\bar{C}-\cup p_i$ for which the modulus of every annular end is infinite.

By abuse of notation, we refer to the *components* of $\bar{C}$, by which we mean the components of $\hat{C}$. These are Riemann surfaces with marked points (i.e. the preimages of the p_i).

Let P be a smooth manifold with an almost-complex structure J and a compatible metric g (i.e one for which J preserves lengths and has Jv perpendicular to v). A *(pseudo)-holomorphic curve* is a Riemann surface C and a smooth map $u:C\to P$ whose differential du is complex linear for all vectors $v\in TC$; i.e. $du(iv)=Jdu(v)$. A *cusp curve* is a singular Riemann surface $\bar{C}$ together with a map $u:\bar{C}\to P$ whose differential is complex linear away from the p_i.

A sequence of holomorphic curves $u_n:C_n\to P$ *converges weakly* to a cusp curve $u:\bar{C}\to P$ if

(1) the areas of the $u_n(C_n)$ converge to the area of $u(\bar{C})$; and
(2) there are families of disjoint loops $\gamma_{n,i} \subset C_n$ and diffeomorphisms $\phi_n : \bar{C} - \cup p_i \to C_n - \cup\gamma_{n,i}$ so that the maps $u_n\phi_n$ converge uniformly to u on compact subsets of $\bar{C} - \cup p_i$.

Degenerations of pseudo-holomorphic curves are controlled by:

Theorem 5.2 (Gromov's compactness Theorem [10], Thm. 1.5.B)**.** *Let P be a smooth manifold with an almost-complex structure J and a compatible metric g. Then any sequence of holomorphic curves of fixed genus and uniformly bounded area has a subsequence which converges weakly to a cusp curve.*

If P, ω is a symplectic manifold and J is an almost-complex structure compatible with ω then the area of a curve depends only on its homology class; thus pseudo-holomorphic curves in a fixed homology class all have bounded area, and we can 'compactify' the space $\mathcal{M}$ of such curves by adding cusp curves. The word 'compactify' is in quotes because cusp curves (or ordinary curves for that matter) may admit noncompact families of automorphisms, e.g. if some component has genus 0 and fewer than 3 marked points. Dividing out by such automorphisms, the quotient space $\widehat{\mathcal{M}}$ is compactified by cusp curves.

Gromov's proof of Theorem 5.2 has the following key ingredients. We treat the case when we have compatible J, g, ω for simplicity.

(1) **Minimal surfaces:** The symplectic form ω calibrates complex subspaces of the tangent space; thus every holomorphic curve is a minimal surface, so its Gauss curvature is uniformly bounded above by some constant K, and it satisfies a uniform isoperimetric inequality (by comparison with a surface of constant curvature K).
(2) **Gromov–Schwarz Lemma:** For every conformal $u : \mathbb{D} \to P$ Cauchy–Schwarz gives

$$\frac{\partial}{\partial r}\text{area}(u(\mathbb{D}_r)) = \int_{\partial\mathbb{D}_r} |du|^2 \geq \frac{1}{2\pi r}\left(\int_{\partial\mathbb{D}_r} |du|\right)^2 = \frac{1}{2\pi r}\text{length}(u(\partial\mathbb{D}_r))^2$$

Thus the isoperimetric inequality implies (by integrating) that for every conformal map $\mathbb{D} \to P$ the area of the image controls the derivative at 0, at least for areas $< 2\pi/K$ if K is positive. A bootstrap argument in the jet space gives inductive control on all higher derivatives at 0.
(3) **Extension over punctures:** By the Gromov–Schwarz Lemma, a holomorphic map $u : \mathbb{D}^* \to P$ of small area is uniformly Lipschitz in the hyperbolic metric, so concentric circles around the puncture have very small image. In a minimal surface, a very thin tube can't be very long, so these circles must converge, and u fills in over the puncture.
(4) **Extracting a limit:** Fix small constants $A < 2\pi/K$ and $\epsilon > 0$ so that every minimal surface in P intersects every ball of radius at least ϵ in a subsurface of area at least A. If $u : C \to P$ is a holomorphic curve we can find a maximal subset of points $Q \subset C$ so that the ϵ balls about the points of $u(Q)$ are disjoint. Then the cardinality of $|Q|$ is bounded above by the area of $u(C)$, and below by 3 if we take ϵ small enough. Thus $C - Q$ is conformally hyperbolic with bounded complexity, and

by Gromov–Schwarz, the norm of du in the hyperbolic metric is uniformly bounded on $C - Q$.

If $u_n : C_n \to P$ is a sequence of curves and $Q_n \subset C_n$ points as above, either the hyperbolic metrics on $C_n - Q_n$ have a convergent subsequence, or there is a subsequence for which these metrics degenerate by stretching necks centered at finitely many essential simple closed loops, in which case $C_n - Q_n$ converges to a singular $\bar{C} - Q$. Since the u_n are equicontinuous in the hyperbolic metrics, some subsequence of the $u_n|C_n - Q_n \to P$ converges on compact subsets to a holomorphic map $u|\bar{C} - Q \to P$, which extends to $u : \bar{C} \to P$ by step (3).

5.2. **Formal dimension of the Moduli space.** Let $u : C \to P$ be a holomorphic curve of genus g. A smooth variation of u (keeping the domain fixed) is tangent to a section of u^*TP, and a variation through holomorphic maps is tangent to a holomorphic section, where we think of u^*TP as a complex vector bundle, of complex dimension n if P has real dimension $2n$.

Suppose for the moment that the complex structure on P is integrable. Then $E := u^*TP$ is a holomorphic vector bundle, and the Euler characteristic $\chi(E) := \dim_{\mathbb{C}} H^0(C; E) - \dim_{\mathbb{C}} H^1(C; E)$ can be calculated by the Riemann–Roch formula $\chi(E) = \mathrm{ch}(E)\mathrm{Td}(TC)[C]$ where $\mathrm{ch}(E)$ is the Chern character

$$\mathrm{ch}(E) = \mathrm{rank}(E) + c_1(E) + \frac{1}{2}(c_1(E)^2 - 2c_2(E)) + \cdots$$

and Td is the Todd class of the (holomorphic) tangent bundle TC

$$\mathrm{Td}(TC) = 1 + \frac{1}{2}c_1(TC) + \frac{1}{12}(c_1(TC)^2 + c_2(TC)) + \cdots$$

Since C is 1 (complex) dimensional, the only relevant numbers are $\mathrm{rank}(E) = n$ and the Chern numbers $c_1(E)[C]$ and $c_1(TC)[C] = \chi(C) = 2 - 2g$.

If E were a trivial $\mathbb{C}^n$ bundle, we would have $\chi(E) = n(1 - g)$. The correction term for a nontrivial bundle over a curve is the first chern class, so

$$\chi(E) = c_1(E)[C] + n(1 - g)$$

This is the *formal dimension* of the (Zariski tangent) space of holomorphic maps from C to P at u. Of course this is a complex dimension; the real dimension is twice this number.

For an almost-complex structure we can still define an operator $\bar{\partial}_J$ on smooth (real) variations $u(t)$ of u as follows. For any smooth $u(t)$ the differential $du(t)$ is a section of the real bundle $\Omega^1(C, u(t)^*TP)$, and we denote by $\bar{\partial}u(t)$ the projection of du to $\Omega^{0,1}_J(C, u(t)^*TP)$ where the subscript J denotes the dependence on the almost-complex structure. If we fix a holomorphic connection ∇ on M, then for small t we may parallel transport $\bar{\partial}u(t)$ via the connection along nearby geodesics to a section of $\Omega^{0,1}_J(C, u(0)^*TP)$; denote the result by $\bar{\partial}_J u(t)$. Thus, the operator $\bar{\partial}_J$ vanishes precisely on (pseudo)-holomorphic variations. This operator is quasi-linear and elliptic, and its linearization has the same symbol as the usual Cauchy–Riemann operator $\bar{\partial}$. Because index is a homotopy invariant, we may compute it as in the integrable case.

To compute the formal dimension of the moduli space $\mathcal{M}$ of all genus g holomorphic curves in the homotopy class ϕ of u we want to allow the complex structure on C to vary.

To compute the formal dimension of $\widehat{\mathcal{M}}$ we must quotient by the automorphism group of C. For $g \geq 2$ the (complex) dimension of the space of complex structures on a genus g curve is $3g-3$, whereas for $g = 0, 1$ the dimension is $0, 1$ respectively. On the other hand, the (complex) dimension of $\mathrm{Aut}(C)$ for C a genus g curve is $3, 1$ when $g = 0, 1$ and is otherwise 0. Thus the difference is $3g-3$ for any g, and the formal dimension of $\widehat{\mathcal{M}}$ is

$$\dim_{\mathbb{R}} \widehat{\mathcal{M}} = 2c_1(u^*TP)[C] + 2(n-3)(1-g)$$

If $L \subset P$ is a Lagrangian submanifold, we can consider holomorphic curves with boundary on L. For compatible J, g, ω holomorphic curves are perpendicular to the boundary, so by doubling we get closed holomorphic curves, and we are interested in subspaces of the moduli spaces of the double invariant under the symmetry. This subspace has half the dimension; i.e.

$$\dim_{\mathbb{R}} \widehat{\mathcal{M}} = c_1^D + (n-3)(1-g)$$

where c_1^D is the first Chern number of the bundle over the double. If C is a disk, $g = 0$ and c_1^D is equal to the Maslov index μ of ∂C, as defined in § 4.1.2. To see this, observe that the effect of adding 1 to c_1 in the interior of C adds 2 both to c_1^D and to μ; compare Proposition 4.1. Thus $\dim_{\mathbb{R}} \mathcal{M} = \mu + n$ and $\dim_{\mathbb{R}} \widehat{\mathcal{M}} = \mu + n - 3$. If we pick a point $x \in L$ and ask for holomorphic disks with $u(i) = x$ this cuts the real dimension down by n; thus $\dim_{\mathbb{R}} \mathcal{M} = \mu$ and $\dim_{\mathbb{R}} \widehat{\mathcal{M}} = \mu - 2$.

For a pair of transverse Lagrangians $L_1, L_2 \subset P$ and holomorphic disks with corners passing through points of $L_1 \cap L_2$ we can obtain the correct index by doubling along the restriction to L_2 (say). Then we obtain a disk with boundary on L_1 whose index as we have just seen is μ^D (the Maslov index of the doubled free boundary). Thus $\dim_{\mathbb{R}} \mathcal{M} = \mu^D/2 = \mu$ as computed in § 4.1.2.

Suppose we have a pair of Lagrangians L_1, L_2 and points $x, y \in L_1 \cap L_2$. Let $\phi \in \pi_2(x, y)$ be a homotopy class of Whitney disk. Let $\theta \in \pi_2(P, L_1, x)$ be a homotopy class of disk with boundary on L_1 and basepoint at x, and let $S \in \pi_2(P, x)$ be a homotopy class of 2-sphere. It makes sense to consider the homotopy classes $\theta + \phi$ and $S + \phi$. By the discussion above, the indices (and therefore the formal dimensions of $\mathcal{M}$) are related by

$$\mu(\theta + \phi) = \mu(\theta) + \mu(\phi) \text{ and } \mu(S + \phi) = 2c_1[S] + \mu(\phi)$$

5.3. **Smoothness of the Moduli space.** One would like to say that for a generic choice of almost complex structure J zero is a regular value of $\bar{\partial}_J$ so that the moduli spaces $\mathcal{M}$ and $\widehat{\mathcal{M}}$ are smooth manifolds of the correct dimension. This is morally true, and follows from Sard–Smale (once we have set up the proper function spaces so that the operator is Fredholm) with a substantial caveat.

If we fix a smooth surface S, we denote by $\mathcal{T}$ the Teichmüller space of marked holomorphic structures on S, by $\mathcal{P}$ a suitable Sobolev completion of the space of smooth maps from S to P in a homology class ϕ, and by $\mathcal{J}$ the Banach space of almost-complex structures on P compatible with ω which are C^r for some fixed $r > 0$. Let $\mathcal{E}$ be the bundle over $\mathcal{T} \times \mathcal{P} \times \mathcal{J}$ whose fiber over $(u : C \to P, J)$ is (a suitable Sobolev completion of) $\Omega_J^{0,1}(C, u^*TP)$. Then $\bar{\partial}_J$ defines a section of this bundle, and the 'universal Teichmüller space' $\mathcal{T}_{\mathcal{J}}(\phi)$ is the preimage of the zero section.

One would like to argue that $\mathcal{T}_{\mathcal{J}}(\phi)$ is a smooth Banach manifold and that the projection $\mathcal{T}_{\mathcal{J}} \to \mathcal{J}$ is Fredholm, and therefore that the fibers of this map $\mathcal{T}_J(\phi) := \mathcal{T}_{\mathcal{J}} \cap \mathcal{T} \times \mathcal{P} \times J$ are smooth of the correct dimension for a Baire set of $J \in \mathcal{J}$.

In fact, this is false. The most serious issue has to do with multiple covers. For example, a genus 0 curve admits a $2d+1$ complex dimensional family of holomorphic self-maps R (rational functions) of any positive degree d, so every nonconstant holomorphic map $u : \mathbb{CP}^1 \to P$ in a homology class ϕ gives rise to a $2d+1$ dimensional family of maps of the form $uR : \mathbb{CP}^1 \to P$ in the homology class $d\phi$. But if $c_1(\phi) \le 0$ the formal dimension of $\mathcal{M}(d\phi)$ will be less than that of $\mathcal{M}(\phi)$ and thus we can never achieve transversality under such circumstances.

Away from multiply-covered curves the optimistic picture holds: $\mathcal{T}_{\mathcal{J}}$ is a smooth Banach manifold, and the slices $\mathcal{T}_J$ are generically smooth manifolds of the correct dimension.

For holomorphic disks with boundary or corners the same issue arises whenever $u : \mathbb{D} \to P$ factors as $u = fg$ for some $g : \mathbb{D} \to S$ where $f : S \to P$ is nonconstant, and g maps over every point of S with degree at least 2.

For holomorphic Whitney disks with boundary on a pair L_1, L_2 of Lagrangians this problem can be solved by considering 1-parameter families J_t of almost complex structures compatible with ω. Given J_t and a pair of points $x, y \in L_1 \cap L_2$ we consider maps $u : \mathbb{R} \times [0,1] \to P$ with $u(\cdot, j) : \mathbb{R} \to L_j$ running from x to y for $j = 0, 1$, and for which $du(iv) = J_t du(v)$ for v a tangent vector at the point (s,t). The time-dependence breaks the symmetry of a multiple cover, and lets us achieve transversality; see Floer [7]. We will essentially ignore this issue going forward.

Quotienting by the (discrete!) action of the mapping class group gives universal moduli space $\mathcal{M}_{\mathcal{J}}$; quotienting further by holomorphic reparameterizations of the domain gives $\widehat{\mathcal{M}}_{\mathcal{J}}$. When $g > 0$ the action of the mapping class group on $\mathcal{T}(S)$ is not free, so that the quotient $\mathcal{M}_{\mathcal{J}}$ and its slices $\mathcal{M}_J$ may acquire orbifold singularities even if $\mathcal{T}_{\mathcal{J}}$ and $\mathcal{T}_J$ are manifolds; however, this is not an issue for genus 0 curves (or disks) and it does not come up in our context.

5.4. **Orientations of the moduli spaces.**

5.5. **Compactification and $\partial^2 = 0$.** Let's now specialize to the case of $P = S^g\Sigma$ and a pair T_α, T_β of Lagrangians, and study the moduli spaces $\mathcal{M}(\phi)$ and $\widehat{\mathcal{M}}(\phi)$ for $\phi \in \pi_2(x,y)$ where $x, y \in T_\alpha \cap T_\beta$. These moduli spaces have dimension $\mu(\phi)$ and $\mu(\phi) - 1$ respectively, ignoring the degenerate case of $x = y$ and ϕ the constant map. The moduli spaces are smooth of the correct dimension for generic (paths of) almost-complex structures J_t. When $\mu(\phi) = 1$ so that $\dim \widehat{\mathcal{M}}(\phi) = 0$, Gromov compactness implies that $\widehat{\mathcal{M}}$ is compact; i.e. it consists of a finite set of points.

When $\mu(\phi) = 2$ we must compactify $\widehat{\mathcal{M}}$ with cusp curves. A priori there are three kinds of degeneration to consider:

(1) **sphere bubbling:** a circle in $\mathbb{D}$ pinches and bubbles off a holomorphic sphere S;
(2) **disk bubbling:** an interval in $\mathbb{D}$ not separating $i, -i$ pinches and bubbles off a holomorphic disk D with boundary on T_α (say); or
(3) **strip breaking:** an interval in $\mathbb{D}$ separating $i, -i$ pinches and $\mathbb{D}$ degenerates to a pair of holomorphic disks in $\pi_2(x,z) \times \pi_2(z,y)$ for some intermediate z.

For sufficiently large μ these degenerations may all occur, possibly multiple times.

However for $\mu(\phi) = 2$ we are in better shape. If a holomorphic sphere S pinches off we have $\phi = S + \phi'$. Since S is holomorphic, it's a positive multiple of the generator of π_2, and therefore $c_1[S] > 0$, so that $\dim \widehat{\mathcal{M}}(\phi') = \mu(\phi') - 1 < 0$, so this case will not occur.

Likewise, if a holomorphic disk D pinches off we have $\phi = \theta + \phi'$ for $\theta \in \pi_2(S^g\Sigma, T_\alpha)$ (say). Since the inclusion of T_α in $S^g\Sigma$ is injective on π_1, the loop $\partial\theta$ is homotopically trivial in T_α, so we can cap off D with a disk $E \to T_\alpha$ to make a sphere S. Since D is holomorphic and E is Lagrangian, it's still true that S is a positive multiple of the generator of π_2, so this case won't occur either.

Finally, if ϕ degenerates to $\phi' + \phi''$ for $\phi' \in \pi_2(x, z)$ and $\phi'' \in \pi_2(z, y)$ then $\mu(\phi) = \mu(\phi') + \mu(\phi'')$, and since each of these must be at least 1 for $\widehat{\mathcal{M}}$ to be nonempty, the only possibility is that $\mu(\phi') = \mu(\phi'') = 1$ and the multiplicity of the degeneration is 1.

The conclusion is that $\widehat{\mathcal{M}}(\phi)$ can be compactified to a 1-manifold with boundary, and that the boundary points correspond to products of $\widehat{\mathcal{M}}$ for classes ϕ', ϕ'' with $\mu(\phi') = \mu(\phi'') = 1$.

To prove that $\partial^2 = 0$ one must argue the converse — that a pair of disks joining x to z and z to y can be glued and then perturbed to produce a smooth disk joining x to y. This is proved by Floer [6] Prop. 4.1 by modifying Taubes' gluing construction for instantons [19].

5.6. **Dependence on J_t.** We have now shown (modulo analytic details!) that the homology groups $\widehat{HF}$ (and for similar reasons the HF^*) make sense, at least for a generic choice of path J_t of compatible almost complex structures on $S^g\Sigma$, and at least for $b_1(M) = 0$. The next step is to show that different generic choices of J_t produce isomorphic homology groups.

Suppose $J_{i,t}$ for $i = 0, 1$ are two paths of almost complex structures, giving rise to chain complexes $\widehat{CF}_i, \partial_i$ for $i = 0, 1$, and let $J_{s,t}$ be a 1-parameter family of paths interpolating between them. For convenience we extend this to $s \in \mathbb{R}$ by $J_{s,t} = J_{0,t}$ for $s \le 0$ and $J_{s,t} = J_{1,t}$ for $s \ge 1$.

We can define a chain homotopy from the complex $\widehat{CF}_0, \partial_0$ to $\widehat{CF}_1, \partial_1$ by counting holomorphic Whitney disks which are holomorphic with respect to $J_{s,t}$; i.e. that satisfy $du(iv) = J_{s,t}du(v)$ for v a tangent vector at the point (s, t). For a class $\phi \in \pi_2(x, y)$ let $\mathcal{M}_*(\phi)$ denote the space of $J_{s,t}$-holomorphic Whitney disks from x to y, and let $c_*(\phi)$ denote the signed count of points in $\mathcal{M}_*(\phi)$ when $\mu(\phi) = 0$. Note that the s-dependence of $J_{s,t}$ breaks the translational symmetry of $\mathcal{M}_*$ so that $\mathcal{M}_* = \widehat{\mathcal{M}}_*$.

Then we can define a chain map $\Phi : \widehat{CF}_0 \to \widehat{CF}_1$ by

$$\Phi(x) := \sum_y \sum_{\phi \in \pi_2(x,y) : \mu(\phi)=0, n_z(\phi)=0} c_*(\phi) y$$

To see this is a chain map, consider moduli spaces $\mathcal{M}_*(\phi)$ with $\mu(\phi) = 1$. These are 1-dimensional, and their noncompactness arises only for families in which area escapes to $\pm\infty$. In either case a suitable sequence of translates converges to a nontrivial (translation-invariant!) curve which is $J_{i,t}$ holomorphic for one of $i = 0, 1$. In other words, $\widehat{\mathcal{M}}_*(\phi)$ (which is equal to $\mathcal{M}_*(\phi)$) is compactified by products of the form $\widehat{\mathcal{M}}_0(\phi_1) \times \mathcal{M}_*(\phi_2)$ and

$\mathcal{M}_*(\phi_1)\times\widehat{\mathcal{M}}_1(\phi_2)$ for factorizations $\phi_1+\phi_2=\phi$ where each moduli space in the product is 0-dimensional. The signed count of these boundary points gives the coefficients of $\partial_1\Phi-\Phi\partial_0$ which are therefore zero.

To see that Φ induces an isomorphism in homology, define $\bar{\Phi}$ to be the chain map induced by $J_{1-s,t}$. Then the concatenation of $J_{s,t}$ and $J_{1-s,t}$ is a family of time-dependent almost complex structures from $J_{0,t}$ to itself, and by interpolating between this family and the constant family we obtain (as above) moduli spaces that induce a chain homotopy between $\bar{\Phi}\Phi$ and the identity map on $\widehat{CF}_0$.

Analogous chain homotopies can be defined between CF_i^∞. These restrict to chain homotopies between CF_i^- and induce chain homotopies between CF_i^+, in every case inducing isomorphisms on homology.

5.7. **Exact Hamiltonian isotopies.** If P is a symplectic manifold, the pairing on TP induced by the symplectic form lets us canonically identify 1-forms with vector fields. Thus a smooth function f determines a smooth vector field $X:=X_f$ by the identity $\omega(X_f,\cdot)=df$, and a family $f_t, t\in[0,1]$ of smooth functions determines a family $X_t:=X_{f_t}$ of smooth vector fields. This family defines a flow ϕ by $d\phi/dt=X_t$, called an *exact Hamiltonian isotopy*.

If L_1,L_2 are a pair of Lagrangians, it turns out under suitable circumstances that the Intersection Floer homologies of the pairs L_1,L_2 and $\phi(L_1),L_2$ are isomorphic when $\phi=\phi_1$ for an exact Hamiltonian isotopy as above.

Let $C_*(L_1,L_2)$ and $C_*(\phi(L_1),L_2)$ denote the Floer complexes generated by intersections $x\in L_1\cap L_2$ and $y\in\phi(L_1)\cap L_2$ respectively. We want to count holomorphic disks from x to y of a certain kind.

Extend the domain of ϕ to all of $\mathbb{R}$ by making it constant on $t\le 0$ and $t\ge 1$. Let $\psi:[0,1]\times\mathbb{R}\to\mathbb{D}-\pm i$ be a conformal parameterization. Let $\mathcal{M}$ denote the space of holomorphic disks $u:\mathbb{D}\to P$ with $u(-i)=x$ and $u(i)=y$, where $u\psi(0,t)\in L_2$ and $u\psi(1,t)\in\phi_t(L_1)$.

These disks fall into subsets parameterized by their homotopy class, and we can define the Maslov index μ of a homotopy class in the usual way, and the space of disks in a homotopy class is a manifold of dimension μ.

The main new technical issue is to prove Gromov compactness. Since area is not constant in a homotopy class, this is not automatic as before. Let u_0 and u_1 be two holomorphic maps in the same homotopy class, and let $U:\mathbb{D}\times[0,1]\to P$ be a homotopy between them through smooth maps with $u\psi(0,t)\in L_2$ and $u\psi(1,t)\in\phi_t(L_1)$. Since L_2 is Lagrangian, the difference in the areas of u_0 and u_1 is equal to the integral of ω over a map $F:\mathbb{R}\times[0,1]$ where $F(t,s)\in\phi_t(L_1)$.

Now, $dF(\partial_s)$ is tangent to $\phi_t(L_1)$ and $dF(\partial_t)=X_t(F)+V$ where V is tangent to $\phi_t(L_1)$. Thus

$$\int_{\mathbb{R}\times[0,1]}F^*\omega=\int\omega(X_t,dF(\partial_s))dt\,ds=\int\langle df_t,dF(\partial_s)\rangle dt\,ds=\int_{\mathbb{R}}f_t(F(t,1))-f_t(F(t,0))dt$$

and therefore the difference in areas is bounded by $\int_t\sup f_t-\inf f_t$, which is finite because $f=0$ outside $t\in[0,1]$.

It follows that the signed count of disks with $\mu = 0$ is finite, and defines a map from $C_*(L_1, L_2)$ to $C_*(\phi(L_1), L_2)$. One shows under suitable circumstances (roughly as in § 5.6) that this map is a chain map inducing an isomorphism in homology.

5.8. **Stabilization.** Invariance under stabilization of Heegaard splittings is much easier in the hat version of HF.

5.9. **Handle slides.**

6. Computation and Examples

6.1. **Knot Floer Homology.**

6.2. **Sutured Floer Homology.**

6.3. **Contact Floer Homology.**

6.4. **L-spaces.**

7. Acknowledgments

References

[1] S. Akbulut and J. McCarthy, *Casson's invariant for oriented homology 3-spheres*, Math. Notes **36**, Princeton University Press, Princeton, NJ, 1990
[2] J. Alexander, *Topological invariants of knots and links*, Trans. AMS **30** (2) (1928), 275–306
[3] V. Arnol'd, *The Sturm Theorems and Symplectic Geometry*, Func. Anal. Appl. **19** (1985), no. 4, 251–259
[4] S. Donaldson, *An application of gauge theory to the topology of 4-manifolds*, J. Diff. Geom **18** (1983), 269–316
[5] A. Floer, *An Instanton Invariant for 3-Manifolds*, Comm. Math. Phys. **118** (1988), 215–240
[6] A. Floer, *Morse theory for Lagrangian Intersections*, J. Diff. Geom. **28** (1988), 513–547
[7] A. Floer, *The Unregularized Gradient Flow of the Symplectic Action*, Comm. Pure. Appl. Math. **41** (1988), 775–813
[8] M. Freedman and F. Quinn, *Topology of 4-Manifolds*, Princeton University Press, Princeton, NJ, 1990
[9] W. Goldman, *The symplectic nature of fundamental groups of surfaces*, Adv. Math. **54** (1984), no. 2, 200–225
[10] M. Gromov, *Pseudo-holomorphic curves in symplectic manifolds*, Invent. Math. **82** (1985), 307–347
[11] R. Kirby, *The topology of 4-manifolds*, Springer LNM **1374**. Springer-Verlag, Berlin, 1989
[12] Y. Ni, *Knot Floer homology detects fibered knots*, Invent. Math. **170** (3) (2007), 577–608
[13] P. Ozsváth and Z. Szabó, *Holomorphic disks and topological invariants for closed three-manifolds*, Ann. Math. (2) **159** (2004), no. 3, 1027–1158
[14] P. Ozsváth and Z. Szabó, *Holomorphic disks and three-manifold invariants: properties and applications*, Ann. Math. (2) **159** (2004), no. 3, 1159–1245
[15] P. Ozsváth and Z. Szabó, *Holomorphic disks and knot invariants*, Adv. Math **186** (1) (2004), 58–116
[16] P. Ozsváth and Z. Szabó, *Holomorphic disks and genus bounds*, Geom. Top. **8** (2004) 311–334
[17] J. Rasmussen, *Floer homology and knot complements*, preprint, arXiv:math/0306378
[18] N. Saveliev, *Lectures on the Topology of 3-Manifolds*, Walter de Gruyter & Co., Berlin, 1999
[19] C. Taubes, *Self-dual Yang-Mills connections on non-self-dual 4-manifolds*, J. Diff. Geom. **17** (1982), no. 1, 139–170
[20] C. Taubes, *Casson's Invariant and Gauge Theory*, J. Diff. Geom. **31** (1990) 547–599
[21] V. Turaev, *Torsion invariants of* Spin^c *structures on 3-manifolds*, Math. Res. Lett. **4** (1997), 679–695
[22] C. Viterbo, *Intersection de sous-variétés lagrangiennes, fonctionnelles d'action et indice des systèmes hamiltoniens*, Bull. SMF **115** (1987) 361–390

University of Chicago, Chicago, Ill 60637 USA
E-mail address: dannyc@math.uchicago.edu