THE MANY FACES OF THE CASSON INVARIANT

COLBY RILEY

Abstract. Since being introduced by Andrew Casson, his namesake Invariant has received considerable interest as being a both practically computable and theoretically interesting tool for investigating homology 3 spheres. Since then, it has proven a springboard for new directions in low-dimensional topology. With this comes different perspectives on the invariant, each having their own benefits and drawbacks. In this survey paper, we explore these perspectives and some of their distinct applications.

Contents

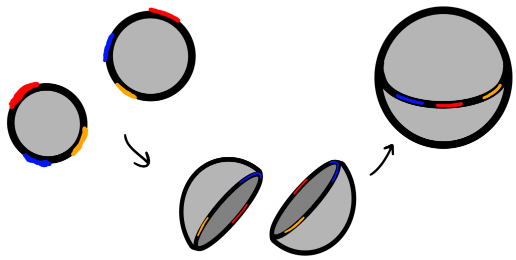

Figure 1. Gluing together two disks along their boundaries. Heegaard splittings do this one dimension higher

1. Introduction

The primary notion of equality between oriented three manifolds is that of orientation preserving homeomorphism. More precisely, if M, N are closed 3 manifolds, then an orientation is a choice of fundamental class $[M], [N]$ for each of them. Then if $f : M \to N$ is a homeomorphism, it is orientation presrving if $f_*[M] = [N]$, and it is orientation reversing if $f_*[M] = -[N]$. If $M = (M, [M])$ is an oriented manifold, then we define $\bar{M} = (M, -[M])$ to be the manifold with opposite orientation.

We will consider oriented 3 manifolds up to orientation preserving homeomorphism. General closed three manifolds have many homological obstructions, and since homology is easy to calculate, we will not care much about these. Instead, we focus on various forms of homology 3 spheres: that is, closed 3 manifolds M with $H_*(M; \mathbb{Z}) = H_*(S^3; \mathbb{Z})$. We may at times also consider rational homology spheres, with the same definition as regular (integral) homology spheres but with $\mathbb{Z}$ replaced by $\mathbb{Q}$. Such manifolds are automatically orientable. To understand homology 3 spheres, we try to find (hopefully simple) invariants of oriented homology 3 spheres. The Casson invariant is a perfect example of this, being a $\mathbb{Z}$-valued invariant of oriented homology three spheres which is both understandable theoretically and easy to compute. Each section of this paper introduces a new definition of the Casson invariant to understand this framework better.

1.1. Heegaard splittings.

Consider the standard embedding of the solid torus into space $D^2 \times S^1 \hookrightarrow \mathbb{R}^3 \cup \{\infty\} = S^3$. If we delete the interior of the solid torus from S^3, we get another compact 3 manifold with boundary, consisting of the outside of the torus. This exterior is also homeomorphic to a solid torus. Thus, we can consider S^3 to be built by 'gluing together' two solid tori. Another way to see this is to equate D^4 with $D^2 \times D^2$ and take the boundary, resulting in

$$S^3 = \partial D^4 = (\partial D^2 \times D^2) \cup (D^2 \times \partial D^2) = S^1 \times D^2 \cup D^2 \times S^1$$

In general, if X, Y are two compact manifolds with boundaries ∂X, ∂Y, and if $f : \partial X \to \partial Y$ is a homeomorphism, then we can define a new manifold $X \cup_f Y$ to be the space

$$X \cup_f Y := (X \sqcup Y)/(x \sim f(x))$$

where $x \in \partial X$. Then $X \cup_f Y$ is another compact manifold, this time without boundary.

If X and Y are oriented manifolds, we can give $X \cup_f Y$ an orientation by requiring that f is an orientation *reversing* homeomorphism of the boundaries (where the orientation of ∂X is induced by the orientation of X). This motivates the following definition:

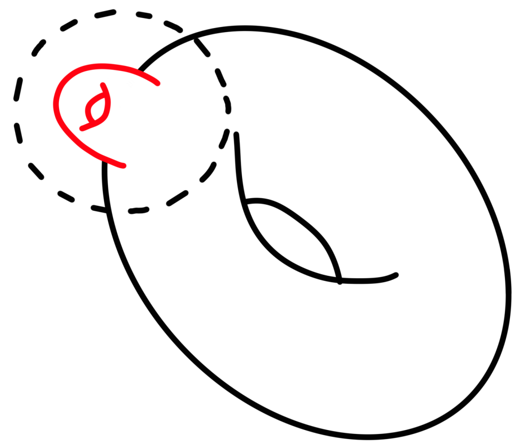

FIGURE 2. Stabilizing a Heegaard splitting.

Definition. *A Heegaard splitting of a closed, oriented 3 manifold M is a triple (M, H_1, H_2) consisting of two genus g handlebodies $H_1, H_2 \subset M$ such that $H_1 \cup H_2 = M$ and $H_1 \cap H_2 = \partial H_1 = \partial H_2$.*

Considering the inclusions $i_1, i_2 : H_1, H_2 \hookrightarrow M$ the Heegaard splitting views M as $H_1 \cup_{i_1 \circ i_2^{-1}} H_2$. A helpful theorem is that in fact, every three manifold can be obtained by such a gluing.

Theorem 1 ([9]). *Every closed, oriented 3 manifold admits a Heegaard splitting.*

Proof. There are two proofs. The first proof is to work in the Piecewise Linear category. Take a triangulation of the three manifold and thicken the one skeleton. The vertices become copies of the ball B^3 and edges become handles $D^2 \times I$. The result is a handlebody of some genus; dually, the manifold with this skeleton removed will also be a handlebody. The second proof works in the smooth category. The index 0 points give a ball B^3, and the index 1 points provide 1-handle attachements. Dually, the index 3 and 2 points also give balls and 1-handles, and together these two handlebodies glue together to give the entire manifold. □

Two Heegaard splittings $(M, H_1, H_2), (M, J_1, J_2)$ are *equivalent* if there is an orientation preserving diffeomorphism $M \to M$ which sends H_i to J_i. For most purposes, equivalent Heegaard splittings can be considered the same. One more relevant construction: given a genus g Heegaard splitting (M, H_1, H_2), we can construct a genus $g+1$ Heegaard splitting called the stabilization, as follows: take a point $p \in \partial H_2$, and consider a neighborhood $p \in N_p \subset M$ diffeomorphic to B^3, small enough such that under that diffeomorphism, H_1 is the top half of B^3 and H_2 is the bottom half, as in figure 2. Then glue a new handle B to H_1 within N_p, such that it bounds a disk, again as in figure 2. Then with $H_1' = H_1 \cup B$ and $H_2' = \overline{H_2 \setminus B}$, then $M = H_1' \cup H_2'$ is a new Heegaard decomposition with genuses $g+1$.

The point of all this is:

Theorem 2 ([9]). *All Heegaard splittings of a 3 manifold M are stably equivalent, in the sense that for any two decompositions (M, H_1, H_2) and (M, J_1, J_2), after some number of stabilizations to each of them forming H_i', J_i', then $(M, H_1', H_2') \cong (M, J_1', J_2')$.*

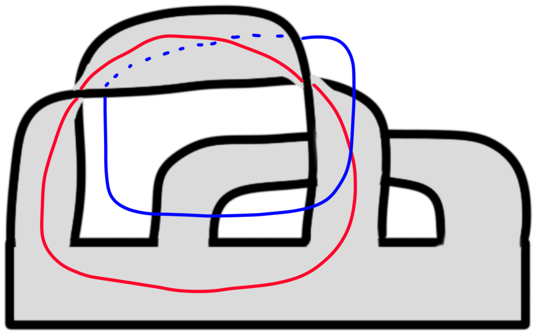

FIGURE 3. The red curve lies on the surface while its blue push off stays above the surface; the linking number is 1.

When one wants to prove a fact about 3 manifolds by using facts about Heegaard splittings, the strategy is to (1) prove that fact given a Heegaard splitting, (2) prove that equivalent Heegaard splittings preserve the fact, and lastly (3) that it remains true under stabilization. The power in this is that stabilization is such a simple construction, so the proof strategy is versatile.

1.2. **Alexander polynomial and Seifert surfaces.** Given a 3 manifold M, a knot is a (smooth) embedding $S^1 \hookrightarrow M$. A collection of disjoint knots is called a link. Knots and links are usually considered *up to isotopy*, which is a smooth map $F : [0,1] \times M \to M$ such that at every point t, $F_t : M \to M$ is a homeomorphism. Then two links are ambient isotopic, or just isotopic, if F carries the images of one link onto the other. The standard embedding $S^1 \hookrightarrow S^3$ is called the unknot.

Given a link $\mathcal{L} \subset \Sigma$ a homology sphere, a Seifert surface for $\mathcal{L}$ is a (connected, compact) oriented surface $F \subset \Sigma$ with $\partial F = \mathcal{L}$ (with the correct induced orientation). It is not hard to show that any link in any homology sphere has an associated Seifert surface.

If $k_1, k_2 \subset \Sigma$ are two knots and F is a Seifert surface for k_1, then after an isotopy of $k_1 \cup k_2$ we can assume that k_2 meets F transversely at each intersection. Whenever they meet, k_2 is heading either in the positive normal direction or the negative: weight the former as $+1$ and the latter as -1. Then define $\mathrm{lk}(k_1, k_2)$ to be the sum of these weightings over all intersections. It is evident that $\mathrm{lk}(k_1, k_2) = \mathrm{lk}(k_2, k_1)$.

For F a Seifert surface of $\mathcal{L}$, we take a list of simple closed curves $x_1, ..., x_n$ generating a basis of $H_1(F;\mathbb{Z})$, and also fix their orientations to coincide with that of F. Then for each x_i, we form the *positive push off* x_i^+, which is the same curve but pushed off of F slightly in the positive normal direction, see figure 3.

The associated Seifert matrix S to F is then the matrix with coefficients $S_{ij} = \mathrm{lk}(x_i, x_j^+)$, and the Alexander polynomial for $\mathcal{L}$ is the polynomial

$$\Delta_{\mathcal{L} \subset \Sigma}(t) = \det(t^{1/2} S - t^{-1/2} S^t)$$

where S^t is the transpose. If the ambient space is implied, especially if said space is S^3, then the notation may be simply $\Delta_{\mathcal{L}}$. The Alexander polynomial is an invariant of the oriented link. If two knots have linking number zero, then there is always a way to make sure that the Seifert surface for k_1 does not intersect k_2, but it is not always possible to find Seifert surfaces for k_1 and k_2 that have empty intersection. Latter such pairs are called *boundary links.*

Definition. *If $k : S^1 \to S^3$ is a knot, then the mirror image k^* of k is k composed with the diffeomorphism $\sigma : (x_1, x_2, x_3, x_4) \mapsto (x_1, x_2, x_3, -x_4)$. The definition extends to links in the obvious way.*

The mirror image of a knot is not always isotopic to the knot itself, but the Alexander polynomial is a bad way to try and detect this.

Lemma 1. $\Delta_k(t) = \Delta_{k^*}(t)$

Proof. The Seifert surface for k^* can also be found by taking the mirror image. This reverses the orientations of everything involved, including the positive push-off, so that in particular $\mathrm{lk}(\sigma(x_i), \sigma(x_j)^+) = \mathrm{lk}(\sigma(x_i), \sigma(x_j^+)) = -\mathrm{lk}(x_i, x_j^+)$. Thus $S_k = -S_{k^*}$. Since the homology of any oriented surface is always even-rank, the negative sign disappears in the determinant. □

There is also a way to define the Alexander polynomial for a knot in a rational homology sphere. Its definition is somewhat more general (see [12]); its most important property is that it reduces to the standard Alexander polynomial on integer homology spheres.

1.3. **Dehn surgery.** A different, usually more practical method of studying 3 manifolds is through Dehn surgery. Let M be a 3 manifold. Any knot k comes equipped with a tubular neighborhood $N(k)$ which is homeomorphic to $D^2 \times S^1$. We can equip this neighborhood with the induced orientation from M. Similarly, we can let $E(k) := \overline{M \setminus N(k)}$ be the knot exterior: it also comes with an induced orientation. Clearly the identity map serves as a function $\mathbb{1} : \partial N(k) \to \partial E(k)$ which is orientation reversing, meaning that we can take the viewpoint $\Sigma = N(k) \cup_{\mathbb{1}} E(k)$. Under the identification of $N(k)$ with $D^2 \times S^1$, we can view this instead as

$$M = E(k) \cup_f (D^2 \times S^1)$$

Dehn surgery on k in Σ is the process of replacing f with some other homeomorphism of the boundaries. In fact, every such f is determined by where it sends one closed curve. When gluing in $D^2 \times S^1$, the image of $S^1 \times \{0\}$ determines the entire rest of the homeomorphism (up to isotopy), by a two step process: first, once $S^1 \times \{0\}$ is glued in, the image of a small collared neighborhood $S^1 \times [-\epsilon, \epsilon]$ is uniquely determined. Then the rest of what we need to glue in, $D^2 \times S^1 \setminus (S^1 \times (-\epsilon, \epsilon))$ is homeomorphic to B^3, and so is uniquely determined (up to isoptopy) by where its boundary ∂B^3 is glued in–but this B^3 shares its boundary with $S^1 \times [-\epsilon, \epsilon]$ which we have already glued in, so it is uniquely determined. See figure 4.

On the other side, if a homology class $c \in H_1(\partial E(k); \mathbb{Z})$ is primitive—i.e., is not a multiple of any other homology class—then it can be represented by the image of a simple, closed, non-serparating curve (also called c) on $\partial E(k)$. Thus up to isotopy, the possible surgeries are determined by the primitive elements $c \in H_1(\partial E(k); \mathbb{Z})$. Let h_c be the homeomorphism of the boundaries which sends $S^1 \times \{0\}$ to c. Then if M was our original 3 manifold, then we denote $E(k) \cup_{h_c} D^2 \times S^1$ by $M + c \cdot k$, or sometimes M_c when the knot is understood.

In the case that Σ is a homology 3 sphere, we can do a bit better. In this case, there are two special curves in $\partial E(k)$, called the meridian and longitude, defined as follows: by Mayer-Vietoris, $H_1(E(k); \mathbb{Z}) \cong \mathbb{Z}$. Let m be a class generating $E(k)$, and pick a representative of m which is a simple, closed, non-separating curve in $\partial E(k)$. Then m is called a *meridian* of k. There is also, up to isotopy, a unique closed, simple, non-separating curve in $\partial E(k)$ which is null-homologous in $E(k)$ (and generates $H_1(N(k); \mathbb{Z})$). We will call this a *canonical longitude* or just *longitude* of k. We pick orientations of m and l such that $\mathrm{lk}(m, l) = 1$.

Now, every curve c in the previous discussion of surgery can be uniquely expressed as $c = p \cdot m + q \cdot l \in H_1(\partial E(k); \mathbb{Z})$. Thus we can refer to surgery on a knot k in a homology sphere by just

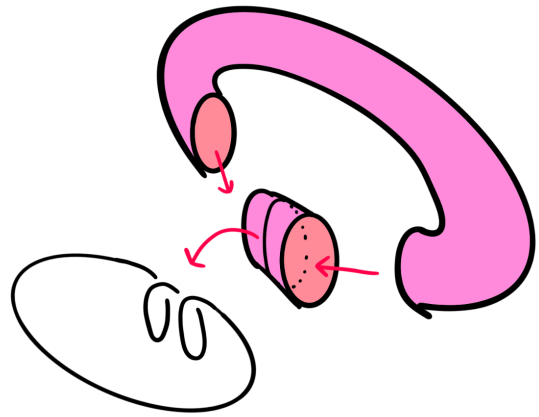

FIGURE 4. A two step process

specifying k and the pair of integers (p, q). Since c is primitive, we can assume that p and q are relatively prime, and since it turns out in addition that $(-p, -q)$ surgery and (p, q) surgery are the same (by an orientation preserving homeomorphism of $D^2 \times S^1$), in fact we can just think of (p, q) as a fraction p/q, including potentially ∞ and 0. We will write such a surgery as $\Sigma + \frac{p}{q} \cdot k$.

As a simple example, if k is the unknot, then $S^3 + \frac{0}{1} \cdot k \cong S^1 \times S^2$. ∞ surgery on any knot will likewise result in the original manifold.

Other than just doing surgery on one knot, we can do surgeries on a link. A framed link is a link $\mathcal{L} \subset \Sigma$ and, for each component k of the link, a primitive homology class. Then by choosing the knot neighborhoods $N(k)$ to be small enough, we can do surgery on each knot separately at once. The resulting manifold will be notated as $\Sigma + \mathcal{L}$ or $\Sigma_{\mathcal{L}}$. If Σ is a homology sphere, then we can specify a fraction p/q instead of the homology classes, consistent with the above paragraphs. If all the framings are integers, then the surgery is an integer surgery. This leads us to the following:

Theorem 3 (Lickorish–Wallace theorem)**.** *Every closed, orientable three manifold can be obtained by integer surgery on a framed link in S^3.*

1.4. **Kirby calculus.** It is a natural question to ask when two (integer-framed) links $\mathcal{L}_1, \mathcal{L}_2$ give rise to the same 3-manifold after performing surgery. The answer to this is Kirby Calculus. On the one hand, given any three manifold we can connect sum it with S^3 to get back the same manifold. In terms of surgery, this is the same as adding a copy of a ± 1 framed unknot to the link which is a boundary link with your given link $\mathcal{L}$. This is because ± 1 surgery on the unknot results in S^3 again. Adding or deleting such a ± 1 framed unknot is called the 'first Kirby move'.

A more complicated move is called handle sliding. Take two knots k_1, k_2, with integer framings n_1, n_2. Let k_2' be a longitude of $E(k_2)$ (so that $\text{lk}(k_2, k_2') = n_2$). Then let $k_\# = k_1 \# k_2'$ the connected sum of the two knots. Then the handle slide of k_1 over k_2 consists of replacing the link $k_1 \cup k_2$ with the link $k_\# \cup k_2$, as in figure 5. The framing of k_2 remains the same, and the framing of $k_\#$ becomes $n_1 + n_2 + 2\text{lk}(k_1, k_2)$. This is the 'second Kirby move', also called a handle slide.

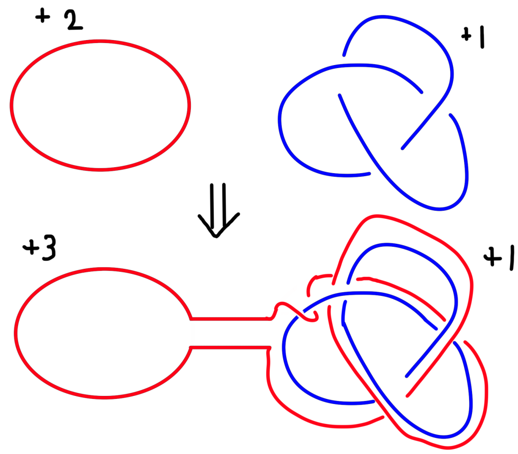

FIGURE 5. Notice that, in order to maintain a 0 linking number with the knot itself (so that it is really null-homologous in E(☘)), the longitude of the trefoil includes an extra 'twist.'

Theorem 4 (Kirby's Theorem, [4]). *Applying Kirby moves to a link does not change the resulting manifold. If surgery on $\mathcal{L}_1, \mathcal{L}_2$ give the same manifold, then the two links are related by a series of Kirby moves.*

Given a link $\mathcal{L}$, we can create a 'linking matrix' consisting of the entries $L_{ij} = \mathrm{lk}(k_i, k_j)$ ranging over all components k_i. Interpret the linking number of k_i with itself to be the framing of k_i. The resulting matrix L is an integral, symmetric matrix. The effect of the first Kirby move is to add a new row and column to L with the only new non-zero entry being a ± 1 on the last diagonal. The effect of the second Kirby move on two components k_i and k_j is to add the jth row to the ith row, then the jth column to the ith column. This operation is just the symmetrized version of the standard elementary row and column operations. In this way, many questions about the properties of three manifolds can be reduced to asking questions about the properties of matrices under the equivalence of these two moves.

Lemma 2 (Saveliev). *Let Σ be a homology sphere. Then there is a link $k_1 \cup \ldots \cup k_n$ in S^3 such that*

(1) $\Sigma = S^3 + \epsilon_1 \cdot k_1 + \ldots + \epsilon_n \cdot k_n$.
(2) $\epsilon_i = \pm 1$ *for all* i.
(3) $\mathrm{lk}(k_i, k_j) = 0$ *for all* $i \neq j$.

Proof. The lemma is equivalent to stating that any such Σ has a presentation with a diagonal matrix with only ± 1s on the diagonal. Start with any linking matrix for Σ. It turns out that the linking matrix represents the same quadratic form as the intersection form of an appropriate manifold with boundary Σ, which for Σ a homology sphere must be invertible over $\mathbb{Z}$ by the obvious long exact sequence. Thus in particular its determinant is ± 1. We may also assume that the linking matrix is indefinite by appending ± 1s as needed. It is a linear algebra fact that we may diagonalize such a matrix by elementary (integer) row and column operations: since $\mathcal{L}$ was symmetric, we may assume these operations are also symmetric, and so realizable by the second Kirby move. Then since the determinant of the matrix was ± 1, these diagonalized entries must be ± 1s. $\square$

We will call framed links of this form *preferred links.* These preferred links have the extra benefit that every intermediate step, $\Sigma_i = \Sigma + \epsilon_1 k_1 + ... + \epsilon_i \cdot k_i$, is also a homology sphere. Furthermore, since the linking numbers are zero between each knot, it is not hard to find a Seifert surface for k_i which is disjoint from all other k_js to calculate the Alexander polynomial.

Lemma 3. *If $\mathcal{L}$ is a preferred link in S^3 with surgery Σ, then framed surgery on $\mathcal{L}^*$ with coefficients negated gives Σ with opposite coefficients.*

Proof. Define $\phi : \Sigma + \mathcal{L} \to \Sigma + \mathcal{L}^*$ as follows: if $x \in E(\mathcal{L})$, then send $x \mapsto x^*$ the image of x under the reflection. If x is part of a surgered component of $\Sigma + \mathcal{L}$, then $x \in D^2 \times S^1$ for some solid torus. Then send x to the equivalent solid torus of the mirror imaged knot, but under a reflection of the solid torus across the longitude. Since the surgery coefficients were reversed, this map glues together continuously, and its continuous inverse is exactly the same by symmetry. Thus the two manifolds are homeomorphic. The fact that ϕ is orientation reversing is then immediate from examining ϕ around a point $x \in E(\mathcal{L})$; around this point, it is just the reflection map with local degree -1. □

Remark. We will often denote Σ with the opposite orientation as $\overline{\Sigma}$. Sometimes, this may conflict the notation for the topological closure of a space, but it should be clear in context.

Lemma 4. *Let M be a rational homology sphere. Then there is a sequence of knots $(k_i)_i^m$ with integral framings n_i such that for each j, $S^3 + n_1 k_1 + ... + n_j k_j$ is a rational homology sphere, and $S^3 + n_1 k_1 + ... n_m k_m = M$.*

Proof. Start with any integer surgery description of M, with some integrally framed sequence of knots $k_1, ..., k_m$. Let $\mathcal{L}_i$ be the link associated to the collection of the first i knots, let L_i be the linking matrix of $\mathcal{L}_i$. For the same reason that the determinant of the linking matrix of an integer homology sphere had to be ± 1, the determinant of the linking matrix for a rational homology sphere must be non-zero. Thus our problem is just to apply Kirby moves to make some L' with $\det(L'_i) \neq 0$ for all i.

For any set of integers $b_1, ..., b_m$, we let $B = (b_i \delta_{ij})_{i,j}$ be the diagonal matrix with b_i along the diagonal. Using the first Kirby move, expand $\mathcal{L}$ to be

$$\begin{pmatrix} \mathcal{L} & \\ & \mathbb{1}_m \end{pmatrix}$$

Then by using the second Kirby move and adding b_i times the b_{i+m}th row / column to the ith row / column, we get another linking matrix

$$\mathcal{L}' = \begin{pmatrix} \mathcal{L} + B^2 & B \\ B & \mathbb{1}_m \end{pmatrix}$$

The determinants $\det(\mathcal{L}'_k)$ are each polynomials in $b_1, ..., b_m$. For $k \leq m$, we can choose $b_1, ..., b_m$ to be large enough that $\mathcal{L} + B^2$ is a positive definite quadratic form, and thus $\mathbb{Z} \ni \det(\mathcal{L}'_k) \neq 0$.

For $m + 1 \leq k \leq 2m$, we could choose all the b_i to be zero, resulting in a determinant $\det(\mathcal{L}'_k)$ which is just the same as $\det(\mathcal{L}) \neq 0$. Thus for any of the k, there is some $(b_1, ..., b_m)$ such that $\mathbb{Z} \ni \det(\mathcal{L}'_k) \neq 0$. Since the determinants are all integer polynomials, some basic algebra shows that there must be some $(b_1, ..., b_m)$ that is non-zero on all the determinants simultaneously. This $(b_1, ..., b_m)$ gives us our result. □

2. Axioms

The simplest definition of the Casson invariant is an axiomatic one, where we specify that there exists an invariant of homology 3 spheres satisfying given properties. Often times, this black box method is all one needs in practice to perform calculations with the invariant, in particular in order to simply differentiate between oriented smooth manifolds.

Definition ([9])**.** *A Casson invariant is an invariant λ of oriented homology 3 spheres satisfying the following properties:*

(1) $\lambda(S^3) = 0$ *and* λ *is surjective.*

(2) For a knot $k \subset \Sigma$, *and* $m \in \mathbb{Z}$, *the difference*

$$\lambda\left(\Sigma + \frac{1}{m+1} \cdot k\right) - \lambda\left(\Sigma + \frac{1}{m} \cdot k\right)$$

Is independent of m.

We define the difference to be $\lambda'_\Sigma(k)$ or just $\lambda'(k)$. The third axiom is a kind of independence condition. If $k, l \subset \Sigma$ are both knots, and within Σ, $\operatorname{lk}(k,l) = 0$, then for any $m, n \in \mathbb{Z}$,

$$\begin{aligned}\lambda\left(\Sigma + \frac{1}{m+1} \cdot k + \frac{1}{n+1} \cdot l\right) &- \lambda\left(\Sigma + \frac{1}{m} \cdot k + \frac{1}{n+1} \cdot l\right) \\ &- \lambda\left(\Sigma + \frac{1}{m+1} \cdot k + \frac{1}{n} \cdot l\right) \\ &+ \lambda\left(\Sigma + \frac{1}{m} \cdot k + \frac{1}{n} \cdot l\right)\end{aligned}$$

is also independent of m and n, since it is equal to both

$$\lambda'_{\Sigma + \frac{1}{n+1} l}(k) - \lambda'_{\Sigma + \frac{1}{n} l}(k)$$

and

$$\lambda'_{\Sigma + \frac{1}{m+1} k}(l) - \lambda'_{\Sigma + \frac{1}{m} k}(l)$$

which are independent of m and n respectively by axiom 2. Thus, we can define this difference as $\lambda''(l,k)$. Then we declare

(3) For k, l *bounding disjoint Seifert surfaces in* Σ, $\lambda''(k,l) = 0$.

These axioms for the Casson invariant essentially states that there is an invariant of homology spheres which satisfies a very precise and well behaved surgery formula that depends only on the knots in question. Notable, this definition does not establish existence of the invariant, which must be proved via one of the other definitions. However, assuming existence, we obtain a surgery formula which enables efficient computation of the invariant in most circumstances.

Theorem 5. *The Casson invariant, if it exists, satisfies, for any knot* $k \subset \Sigma$, *the surgery formula:*

$$\lambda'_\Sigma(k) = \frac{1}{2}\Delta''_{k\subset\Sigma}(1) \cdot \lambda'(\text{trefoil})$$

and $\lambda'(\text{trefoil}) = \pm 1$.

Proof. Start with any knot k and pick a presentation of the knot. At each crossing, we can swap the crossing from being over/under to under/over or vice versa. After applying enough of these, the knot will be unknotted. The operation of swapping a crossing can be realized by a Kirby move on the diagram as in the first image in figure 6, so we place disjoint disks D_i at every crossing with boundaries c_i that are boundary links in both S^3 and $S^3 + k$. For any given such c and k, let k_c be the knot k after applying the crossing change at c.

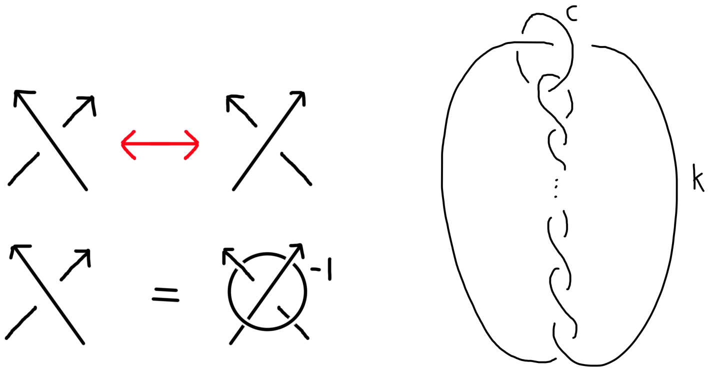

FIGURE 6. Swapping a crossing and the result after swapping enough crossings

We want to look at $\lambda'(k_c) - \lambda'(k) = \lambda'_{S^3+c}(k) - \lambda'(k) = \lambda''(k, c)$. This difference is the same as that of $\lambda'(k_{c'}) - \lambda'(k_c) = \lambda''_{S^3+c'}(k, c)$ for c' one of the other links: it can be verified by using the definition of λ'' and utilizing the fact that $\lambda''(c, c') = \lambda''_{\Sigma+k}(c, c') = 0$ because of our choice of them being boundary links, and axiom (3).

By induction, then, the change in λ' from introducing the twist c does not actually depend on what we do at any of the other twists, so we might as well do all of them first. This leaves us with a knot k' as in figure 6, with n full twists.

Call $k_n := k'_c$ and notice that $k_1 = \text{trefoil}$. Finally, using the same technique as above, we get that $\lambda'(k_n) - \lambda'(k_{n-1}) = \lambda'(k_1) = \lambda'(\text{trefoil})$. By using the skein relations for the Alexander polynomial, the same relationship holds for $\Delta''_{k_n}(1) - \Delta''_{k_{n-1}}(1) = \Delta''_{\text{trefoil}}(1)$. Together, this means that the change from $\lambda(k)$ to $\lambda(k_c)$ is the same as that for the Alexander polynomial, and in particular that $\lambda'(k)$ is proportional to $\lambda'(\text{trefoil})$. By property (1) this means that $\lambda'(\text{trefoil}) = \pm 1$.

This proof goes through the same for $\Delta''_k(1)$ using the skein relations, and since $\Delta''_{\text{trefoil}}(1) = 2$, this means that $\lambda'(k) = \frac{1}{2}\Delta''_k(1)\lambda'(\text{trefoil})$. □

To demonstrate the power of the surgery formula, we can now use lemma 2 to prove several properties.

Theorem 6. *Assuming the Casson Invariant exists,*

(1) It is unique up to the choice of $\lambda'(\text{trefoil})$.
(2) $\lambda(\overline{\Sigma}) = -\lambda(\Sigma)$
(3) $\lambda(\Sigma_1 \# \Sigma_2) = \lambda(\Sigma_1) + \lambda(\Sigma_2)$

Proof.

(1) Let Σ be a homology sphere. By lemma 2, we can assume $\Sigma = S^3 + \sum_{i=1}^n \epsilon_i \cdot k_i$ for $\text{lk}(k_i, k_j) = 0$ for $i \neq j$, and $\epsilon_i = \pm 1$. We induct on the size of n. First, let $n = 1$. If $\epsilon_1 = 1$, then the surgery formula gives us that

$$\frac{1}{2}\Delta''_k(1)\lambda'(\text{trefoil}) = \lambda'(k) = \lambda(S^3 + \frac{1}{\epsilon_1} \cdot k_1) - \lambda(S^3 + \frac{1}{\epsilon_1 - 1} \cdot k_1)$$

$$= \lambda(S^3 + \epsilon_1 \cdot k_1) - \lambda(S^3 + \frac{1}{0} \cdot k_1) = \lambda(\Sigma) - \lambda(S^3) = \lambda(\Sigma)$$

Similarly, if $\epsilon_1 = -1$, then

$$\begin{aligned}\frac{1}{2}\Delta''_k(1)\lambda'(\trefoil) = \lambda'(k) &= \lambda(S^3 + \frac{1}{0}\cdot k_1) - \lambda(S^3 + \frac{1}{\epsilon_1}\cdot k_1)\\ &= \lambda(S^3) - \lambda(\Sigma) = -\lambda(\Sigma)\end{aligned}$$

So the invariant is unique up to the choice of $\lambda'(\trefoil)$, which by theorem 5 is ± 1. Now assuming the statement has been proven for $n-1$, let $\Sigma = S^3 + \sum_{i=1}^n \epsilon_i \cdot k_i$. Then let also $\Sigma_j = S^3 + \sum_{i=1}^j \epsilon_i \cdot k_i$. The statement about linking numbers guarantees that Σ_j is a homology sphere, and that $\Sigma_j + \epsilon_j \cdot k_j = \Sigma_{j+1}$. Then again, if $\epsilon_n = 1$, then

$$\begin{aligned}\lambda(\Sigma) - \lambda(\Sigma_{n-1}) &= \lambda(\Sigma_{n-1} + \frac{1}{1}\cdot k_n) - \lambda(\Sigma_{n-1} + \frac{1}{0}\cdot k_n)\\ &= \frac{1}{2}\Delta''_{k_n\subset\Sigma_{n-1}}(1)\lambda'(\trefoil)\end{aligned}$$

And the story is the same with $\epsilon_n = -1$. Thus by adding $\lambda(\Sigma_0)$ to both sides, we conclude that $\lambda(\Sigma)$ is uniquely determined. In particular, we get by the same induction the stronger statement that

$$\lambda(\Sigma) = \left(\sum_{i=1}^n \frac{\epsilon_i}{2}\lambda''_{k_i\subset\Sigma_i}(1)\right)\lambda'(\trefoil)$$

(2) With the notation as before, by lemma 3 we can reverse the orientation of Σ by the surgery $\overline{\Sigma} = -\epsilon_1 k_1^* \cup ... \cup -\epsilon_n k_n^*$. Then also by lemma 1 (noticing that since the linking numbers are 0, the Alexander polynomials can be computed without special reference to the ambient surgered-manifold), the above equation implies the result.

(3) Let $\Sigma = S^3 + \sum_i \epsilon_i k_i$ for $1 \leq i \leq n$ and $\Sigma' = S^3 + \sum_j \epsilon_j l_j$ for $n \leq j \leq m$ Then the connect sum can be made by performing the surgery on each of the two links separately, with no interaction between the two of them. That is, we can consider $\cup_i k_i$, $\cup_j l_j$ to both be in S^3, but that the two links are separated by some copy of S^2. Then also the Alexander polynomial of a given knot l_j is the same regardless of the surgery performed outside the copy of S^2 since any Seifert surface can be perturbed to be fully inside of the S^2.

Thus we get by the above formula that

$$\begin{aligned}\lambda(\Sigma\#\Sigma') &= \left(\sum_{i=1}^n \frac{\epsilon_i}{2}\lambda''_{k_i\subset\Sigma_i}(1)\right)\lambda'(\trefoil) + \left(\sum_{j=n}^m \frac{\epsilon_j}{2}\lambda''_{l_i\subset\Sigma_i}(1)\right)\lambda'(\trefoil)\\ &= \left(\sum_{i=1}^n \frac{\epsilon_i}{2}\lambda''_{k_i\subset\Sigma_i}(1)\right)\lambda'(\trefoil) + \left(\sum_{j=n}^m \frac{\epsilon_j}{2}\lambda''_{l_i\subset\Sigma'_j}(1)\right)\lambda'(\trefoil)\\ &= \lambda(\Sigma) + \lambda(\Sigma')\end{aligned}$$

□

2.1. **Applications: chirality.** A particularly interesting question about oriented three manifolds is to ask when there is an orientation reversing homeomorphism from M to itself, that is, that $M \cong \bar{M}$. For example, considering $S^3 \subset \mathbb{R}^4$, the map

$$(x_1, x_2, x_3, x_4) \mapsto (x_1, x_2, x_3, -x_4)$$

is an isomorphism $S^3 \to \bar{S}^3$. We call these manifolds *achiral*. The Casson invariant makes it easy to find chiral 3 manifolds.

Theorem 7. *The Poincare homology sphere is chiral, and there are infinitely many chiral 3 manifolds in $\mathcal{M}$.*

Proof. If M is a homology 3 sphere, and $M \cong \bar{M}$, then on the one hand, $\lambda(M) = \lambda(\bar{M})$. On the other hand, by theorem 6, $\lambda(\bar{M}) = -\lambda(M)$, so in fact $\lambda(M) = -\lambda(M) \implies \lambda(M) = 0$. The Poincare homology sphere $\Sigma(2,3,5)$ is defined as -1 surgery on [trefoil]. By the surgery formula, $\lambda(\Sigma(2,3,5)) = \pm 1$, so the manifold cannot be achiral. Then we can find infinitely many different such manifolds by considering the nth connect sum $\#^n \Sigma(2,3,5)$, which, via the additivity of the Casson invariant gives a Casson invariant of $\pm n$. $\square$

3. Representations

The original, and probably more enlightening, definition of the Casson invariant is as a tool for counting representations of the fundamental group into $SU(2)$. One of the consequences of the Poincare Conjecture is that for homology spheres, the fundamental group is often a key differentiator between these spaces. Although studying the fundamental group is too hard, group theory tells us that studying representations of a group is usually tractable. This section will mostly follow [9]. We will usually use M to denote any smooth, compact manifold (possibly with boundary) of dimension ≤ 3. H will typically be a handlebody, and F a surface (often the boundary of a handlebody). F_0 will be F with a small disc removed (a deformation retract of F minus a point).

Definition. $R(M) := \text{Hom}(\pi_1 M, SU(2))$ *with the compact open topology.*

Since any such manifold has a finitely presented fundamental group (with a generating set of size, say, m), this Hom is exactly determined by where it sends its generators, with the relations acting as polynomial relations inside of $SU(2)^m$. In this way, $R(M)$ is a subvariety of $SU(2)^m$ and is best thought of as such.

Example 1. *For H genus g, then $R(H) \cong SU(2)^g$. The orientation of $R(H)$ is determined by the orientation of $SU(2)$ and an orientation of $H_1(H;\mathbb{Z})$. If F_0 genus g, $R(F_0) \cong SU(2)^{2g}$.*

Proof. The fundamental group of H is free on g generators $x_1, ..., x_g$. Then a given homomorphism ϕ is determined exactly by where it sends each $\phi(x_i)$, so we have an isomorphism

$$R(M) \cong SU(2)^g$$

Given by $\phi \mapsto (\phi(x_1), ..., \phi(x_g))$. The orientation comes simply enough after this. The generators $\{x_1, ..., x_g\}$ are also generators of the homology, and an ordering of the generators of the homology determines an ordering of the output in ϕ. Then F_0 is the same, since its fundamental group is the free product on $2g$ generators. $\square$

Definition. *A representation $\alpha \in R(M)$ is called* reducible *if, as an action on $\mathbb{C}^2$, there is a nontrivial invariant subspace. Otherwise it is* irreducible. *Let $R^{irr}(M) \subset R(M)$ be the subset of irreducible representations of $R(M)$.*

This is the same definition of reducibility as in standard representation theory. Since the conjugacy classes of $SU(2)$ are copies of S^1, a reducible representation is one which factors as a 1 dimensional representation into $U(1) = S^1$. In particular, this means that such representations are abelian. To demonstrate that we do not actually lose much, we see that

Lemma 5. *For Σ a homology sphere, $R^{irr}(\Sigma) = R(\Sigma) \setminus \theta$ (the trivial representation).*

Proof. As previously mentioned, reducible representations are abelian. But this means that they factor through $\pi_1 M/[\pi_1 M, \pi_1 M]$ the commutator; since Σ is a homology sphere, this is trivial. Thus every such representation is trivial. □

There is an action of $SO(3)$ on the space of irreducible representations by conjugation: for $g \in SO(3)$, $(g \cdot \alpha)(x) = g\alpha(x)g^{-1}$. By irreducibility, this action is free, and so the quotient by this action retains most desirable topological properties.

Definition. $\mathcal{R}(M) := R^{irr}(M)/SO(3)$*, where the quotient is by the SO(3) action defined above.*

Now we may finally see some examples.

Example 2. *For H a genus g handlebody, $\mathcal{R}(H)$ is a smooth, oriented open manifold of dimension $3g-3$. $\mathcal{R}(F_0)$ is similar, with dimension $6g-3$.*

Proof. $R(H) = SU(2)^g$ as discussed before, and so is in particular a smooth, oriented manifold of dimension $3g$. The set of reducible representations is closed, since it is an algebraic subset of $R(H)$ formed by adding all the commutator relations (since reducible representations are exactly those which factor through a copy of $U(1)$). Thus $R^{irr}(H)$ is an open submanifold of $SU(2)^g$. Then since the $SO(3)$ action on $R^{irr}(H)$ is free, the quotient will also be a smooth, open manifold, and the dimension will be 3 less since $\dim SO(3) = 3$.

In particular, if $g = 1$, i.e., the handlebody is the torus, then all the representations will be reducible so $\mathcal{R}(H)$ will be empty. □

Example 3 ([9])**.** *For F a genus g surface, $\mathcal{R}(F)$ is a smooth, oriented, open manifold of dimension $6g-6$. Its orientation is induced by that of $\mathcal{R}(F_0)$.*

Lemma 6. *The association $M \mapsto \mathcal{R}(M)$ is a contravariant functor.*

Proof. π_1 is a covariant functor, Hom is a contravariant functor in its first argument, and the process of taking the irreducible subspace and quotienting by the $SO(3)$ action and clearly preserves functoriality. □

One might hope that for Σ a homology sphere, that $\mathcal{R}(\Sigma)$ would be always be a finite collection of points, but this is sometimes too much to ask for. Intersection theory tells us that when we want to count something, it is fruitful to consider it as a subspace of a larger space. Thus, we can take a Heegaard splitting $\Sigma = H_1 \cup_f H_2$, with $\partial H_1 := F$, and consider the diagram of inclusions

$$\begin{array}{ccc} \Sigma & \xleftarrow{i_1} & H_1 \\ {\scriptstyle i_2}\uparrow & & \uparrow{\scriptstyle j_1} \\ H_2 & \xleftarrow{j_2} & F \end{array}$$

Of course, one may also add the inclusion of F_0 into F. Since these inclusions each induce surjections on the fundamental group, by functoriality, this gives us several similar commutative diagrams of injections

$$\begin{array}{ccc} R(\Sigma) & \hookrightarrow & R(H_1) \\ \downarrow & & \downarrow \\ R(H_2) & \hookrightarrow & R(F) \end{array} \qquad \begin{array}{ccc} R^{irr}(\Sigma) & \hookrightarrow & R^{irr}(H_1) \\ \downarrow & & \downarrow \\ R^{irr}(H_2) & \hookrightarrow & R^{irr}(F) \end{array} \qquad \begin{array}{ccc} \mathcal{R}(\Sigma) & \hookrightarrow & \mathcal{R}(H_1) \\ \downarrow & & \downarrow \\ \mathcal{R}(H_2) & \hookrightarrow & \mathcal{R}(F) \end{array}$$

Which lets us consider $\mathcal{R}(\Sigma)$ as the intersection $\mathcal{R}(H_1) \cap \mathcal{R}(H_2) \subset \mathcal{R}(F)$ of oriented submanifolds inside an oriented manifold. The orientation of $\mathcal{R}(H_2)$ and $\mathcal{R}(F)$ are both determined by the choice of orientations for $H_1(F;\mathbb{Z})$; F comes with an orientation induced from being the boundary of H_1 (whose orientation is induced from Σ). Then orient $H_1(F;\mathbb{Z})$ by the natural symplectic basis corresponding to $H^1(F;\mathbb{Z})$ and Poincare Duality. Then given any orientation of $H_1(H_1;\mathbb{Z})$, orient $H_1(H_2;\mathbb{Z})$ so that their direct sum $H_1(H_1;\mathbb{Z}) \oplus H_1(H_2;\mathbb{Z}) = H_1(F;\mathbb{Z})$ is orientation preserving. Thus there are really two choices we make when we orient everything: (1) our choice of which handlebody to be called H_1, and (2) our choice of orientation of $H_1(H_1;\mathbb{Z})$.

By example 2, $\dim \mathcal{R}(H_1) = \dim \mathcal{R}(H_2) = 3g-3$, which in intersection theory we would expect generically to have a 0 dimensional intersection in a $6g-6$ dimensional surface. This is reassuring, since we want $\mathcal{R}(\Sigma)$ to be a finite collection of points. The obstruction to this being the case is exactly the (lack of) transversality of the intersection.

3.1. **Transversality.** The transversality of intersections is related to the cohomology of these spaces. If π is a finitely presented discrete group, it is a fact that the Zariski tangent space (the algebro-geometric tangent space) at a point $\alpha \in \operatorname{Hom}(\pi, SU(2))$, can be identified with cocycles in group cohomology: in fact, $T_\alpha \operatorname{Hom}(\pi, SU(2)) \cong Z^1_\alpha(\pi; \mathfrak{su}(2))$. Then dividing by coboundaries is the equivalent of modding out by the $SO(3)$ action, so that

$$T_\alpha(\operatorname{Hom}(\pi, SU(2))/SO(3)) \cong H^1_\alpha(\pi; \mathfrak{su}(2))$$

Where $H^1_\alpha(\pi; \mathfrak{su}(2))$ is the group cohomology with π acting on M via the adjoint: $x \cdot u = \operatorname{Ad}_{\alpha(x)} u$ for $u \in \mathfrak{su}(2)$. This correspondence of the tangent spaces with the group cohomology is not difficult, but also beyond the scope of the current paper (see [9]).

As an example, if θ is the trivial representation, then the cocycle condition is $\zeta(xy) = \zeta(x) + \operatorname{Ad}_{\theta(x)} \zeta(y) = \zeta(x) + \zeta(y)$, so that it must be an actual homomorphism, and in particular $\zeta(xy) = \zeta(yx)$, so ζ descends to the abelizanization of π. Thus, if $\pi = \pi_1 M$, then by the universal coefficient theorem,

$$H^1_\theta(\pi_1 M; \mathfrak{su}(2)) = \operatorname{Hom}(\pi_1 M; \mathfrak{su}(2)) = \operatorname{Hom}(H_1(M); \mathfrak{su}(2)) = H^1(M; \mathfrak{su}(2))$$

In particular, this gives us

Lemma 7. *For Σ a homology sphere, the intersection $R(H_1) \cap R(H_2) \subset R(F)$ is transversal at the trivial representation θ.*

Proof. To show that they are transversal inside $R(F)$, it is easier to consider them in the larger space $R(F_0)$, where F_0 is F with a small disk deleted from it. Since $\pi_1 F_0$ is a free group on $2g$ generators, $R(F_0) \cong SU(2)^{2g}$ is an oriented compact surface as well (c.f. example 2), and is generally easier to work with. Now the tangent spaces being transverse at θ is the same as saying that $j_1^* T_\theta R(H_1) + j_2^* T_\theta R(H_2) = T_\theta R(F_0)$.

By the correspondence above (and noting that $H^1(F_0; \mathfrak{su}(2)) = H^1(F; \mathfrak{su}(2))$), it is the same as asking that $i_1^* + i_2^* : H^1(H_1; \mathfrak{su}(2)) \oplus H^1(H_2; \mathfrak{su}(2)) \to H^1(F; \mathfrak{su}(2))$ is an isomorphism. But these maps are part of the Mayer-Vietoris sequence, and since Σ is a homology sphere, exactness gives us an isomorphism. □

Lemma 8. *For (Σ, H_1, H_2) a handlebody decomposition of a homology sphere, then the intersection $\mathcal{R}(H_1) \cap \mathcal{R}(H_2) \subset \mathcal{R}(F)$ is transversal at α iff $H^1_\alpha(\pi_1\Sigma; \mathfrak{su}(2)) = 0$.*

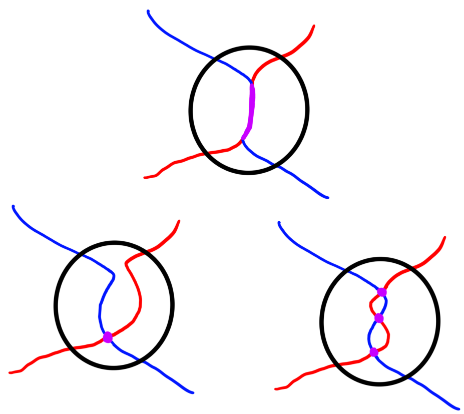

FIGURE 7. Two different ways of compactly resolving an intersection

Proof. By the associations of the tangent spaces to group cohomology, this is a purely homological question. In fact, the Mayer Vietoris sequence gives us

$$0 \longrightarrow H^1_\alpha(\pi_1\Sigma, \mathfrak{su}(2)) \longrightarrow H^1_\alpha(\pi_1 H_1, \mathfrak{su}(2)) \oplus H^1_\alpha(\pi_1 H_1, \mathfrak{su}(2)) \xrightarrow{j_1^*+j_2^*} H^1_\alpha(\pi_1 F, \mathfrak{su}(2))$$

An intersection is transversal iff the map $i_1+i_2 : T_\alpha\mathcal{R}(H_1)\oplus T_\alpha\mathcal{R}(H_2) \to T_\alpha\mathcal{R}(F)$ is surjective. Thus by the correspondence with group cohomology, $j_1^* + j_2^*$ must be surjective. By dimension counting the tangent spaces ($3g-3+3g-3 = 6g-6$) it is surjective iff it is injective. By exactness this will hold iff $H^1_\alpha(\pi_1\Sigma; \mathfrak{su}(2)) = 0$. □

This last lemma is particularly interesting because it says that the transversality of the intersection does not depend on the Handlebody decomposition we choose. Thus we can separate homology spheres into 2 categories: those with all transverse intersections and those with some non-transverse intersections - we call the former nondegenerate and the latter degenerate.

3.2. **Completing the definition.**

Lemma 9. $\mathcal{R}(\Sigma)$ *is compact.*

Proof. Since $R(M) \subset SU(2)^k$ is a subvariety (so closed) of a compact space, it is also compact. Next, by lemma 5, $R^{irr}(\Sigma) = R(\Sigma)\setminus\{\theta\}$, so to show that $R^{irr}(\Sigma)$ is compact, we need to show that θ is an isolated point of $R(\Sigma)$. By functoriality of the above diagrams, we know that $R(\Sigma) = R(H_1)\cap R(H_2)$. By lemma 7, we know that the two spaces intersect transversely at θ. By dimension counting, the intersection is generically 0 dimensional, and thus transverse intersections are isolated points: because of this, $R^{irr}(\Sigma) = R(\Sigma) \setminus \{\theta\}$ is still compact, since all we did was remove a single isolated point.

Lastly, once we know that $R^{irr}(\Sigma)$ is compact, the $SO(3)$ quotient is of course also compact. □

At this point, if $\mathcal{R}(H_1)$ really intersected $\mathcal{R}(H_2)$ transversely, then it would follow that $\mathcal{R}(\Sigma)$ was a compact manifold of dimension 0, and thus a finite collection of points. This is not necessarily the case. Instead, we choose a relatively compact neighborhood of $\mathcal{R}(\Sigma)$, and perturb $\mathcal{R}(H_2)$ with support inside of this neighborhood so that the intersection is transverse.

Let $\tilde{\mathcal{R}}(H_2)$ be this perturbation to $\mathcal{R}(H_2)$ that makes the intersection transverse: then the intersection is still compact, but now it is a manifold. By examples 2 and 3, it is the intersection of two manifolds of dimension $3g-3$ inside a space of dimension $6g-6$, so the resulting intersection is 0 dimensional and so a finite collection of points. Finally, we may count them. Since there was no canonical way to choose $\tilde{\mathcal{R}}(H_2)$, we must count them with sign since, as in figure 7, it could be the case that canceling pairs are added. But when they are counted with sign, we finally get:

Definition. *Let $x \in \mathcal{R}(H_1) \cap \tilde{\mathcal{R}}(H_2)$. Then since the intersection is transverse, $T_x\mathcal{R}(H_1) \oplus T_x\tilde{\mathcal{R}}(H_2) \cong T_x\mathcal{R}(F)$. If this identification is orientation preserving, let $\mu(x) = 0$, and otherwise $\mu(x) = 1$. Then*

$$\lambda(\Sigma, H_1, H_2) := \frac{(-1)^g}{2} \sum_{x \in \mathcal{R}(H_1) \cap \tilde{\mathcal{R}}(H_2)} (-1)^{\mu(x)}$$

Remark. This definition a priori depends on the choices we made earlier about orientation; specifically, the choice of orientation for $H_1(H_1;\mathbb{Z})$. However, if we swap the orientation of this homology group, then the orientation of the other group $H_1(H_2;\mathbb{Z})$ must swap as well by our specification of the orientation of their direct sum. Thus the actual orientation of $T_x\mathcal{R}(H_1) \oplus T_x\tilde{\mathcal{R}}(H_2)$ does not change. It is similar to show that the invariant does not actually depend on which handlebody we called H_1.

Finally we have defined the Casson invariant for handlebody decompositions. It takes some more effort to show that this definition satisfies the same properties as the previous definition. For a full proof of these properties, see [9]. However, we will prove some of them to give a taste of what it is like to work with the invariant. In particular, we will show that λ does not actually depend on the handlebody decomposition.

Theorem 8. *The Casson invariant $\lambda(\Sigma, H_1, H_2)$ is independent of the handlebody decomposition, and so is actually an invariant of the underlying oriented manifold.*

Proof. Fix two handlebody decompositions (Σ, H_1, H_2) and (Σ, H_1', H_2'). If they are equivalent, then there is an orientation preserving homeomorphism f between the two, and by chasing through the functoriality of $\mathcal{R}$, it is clear that they give the same invariant. Thus by theorem 2, it suffices to show that the invariant does not change under stabilization; let H_1', H_2' be the stabilization of H_1, H_2. Denote the genus of H_1, H_2 as g, and let $F = H_1 \cap H_2$, $F' = H_1' \cap H_2'$ be the genus g ($g+1$, respectively) boundary surface. The representation space $R(F)$ is somewhat complicated, so we excise a small disc from F to create a new F_0 whose representation space is a better behaved ambient space.

The idea of the proof is that the extra elements of the fundamental group for H_1' and H_2' do not interact with each other. Thus, when we take the intersections in the representation spaces, they should not contribute anything to $\mathcal{R}(\Sigma)$.

$\pi_1(F')$ has two new natural generators, in addition to those of $\pi_1 F'$. Let a be the generator which is also a new generator of $\pi_1(H_1')$ and b the new generator of $\pi_1(H_2')$. Since $\pi_1 H_1' = \langle a \rangle * \pi_1 H_1$, and similarly with $\pi_1 H_2'$, we get that their representation spaces are $R(H_k') = SU(2) \times R(H_k)$. Similarly, $\pi_1(F_0') = \langle a \rangle * \langle b \rangle * \pi_1(F_0)$, so when viewing the inclusions of $F_0' \hookrightarrow H_1', H_2'$, we get identifications

$$R(H_1') = SU(2) \times \{1\} \times R(H_1) \subset SU(2) \times SU(2) \times \mathcal{R}(F_0)$$

$$R(H_2') = \{1\} \times SU(2) \times R(H_2) \subset SU(2) \times SU(2) \times \mathcal{R}(F_0)$$

So $R(H_1')\cap R(H_2') = \{1\}\times\{1\}\times(R(H_1)\cap R(H_2)) \subset R(F_0)$. Next, taking out the reducible elements, we get that

$$R^{irr}(H_1')\cap R^{irr}(H_2') = \{1\}\times\{1\}\times(R^{irr}(H_1)\cap R^{irr}(H_2))$$

And by the definition of the $SO(3)$ action, it acts component-wise on the product factors, so that the quotient becomes

$$\mathcal{R}(H_1')\cap\mathcal{R}(H_2') = \{1\}\times\{1\}\times(\mathcal{R}(H_1)\cap\mathcal{R}(H_2))$$

That is to say, they are pointwise equal! Although we may have to perturb both $\mathcal{R}(H_2')$ and $\mathcal{R}(H_2)$ to make the intersections transversal, they can be perturbed so that this intersection result is still true.

Now assuming that the intersections are transversal, we lastly need to deal with potential sign issues in the count. Remembering that our definition included a factor of $(-1)^g$, this means we want to show that the sign count of the points in $\mathcal{R}(H_1')\cap\mathcal{R}(H_2')$ is the opposite of that of $\mathcal{R}(H_1)\cap\mathcal{R}(H_2)$.

Fix a point $\alpha\in\mathcal{R}(H_1)\cap\mathcal{R}(H_2)$. First we compare the orientations of $T_\alpha\mathcal{R}(M_1')\oplus T_\alpha\mathcal{R}(M_2')$ with that of $T_\alpha\mathcal{R}(M_1)\oplus T_\alpha\mathcal{R}(M_2)$. However it is that $\mathcal{R}(M_k)$ is oriented, the natural orientation of $T_\alpha\mathcal{R}(M_k')$ is that of $\mathfrak{su}(2)\oplus T_\alpha\mathcal{R}(M_k)$. Thus, we get the orientation preserving identification

$$\begin{aligned}T_\alpha\mathcal{R}(H_1')\oplus T_\alpha\mathcal{R}(H_2') &= \mathfrak{su}(2)\oplus T_\alpha\mathcal{R}(M_1)\oplus\mathfrak{su}(2)\oplus T_\alpha\mathcal{R}(M_2)\\ &= (-1)^{g-1}\mathfrak{su}(2)\oplus\mathfrak{su}(2)\oplus T_\alpha\mathcal{R}(M_1)\oplus T_\alpha\mathcal{R}(M_2)\end{aligned}$$

The $(-1)^{g-1}$ term appearing because we pull the $\mathfrak{su}(2)$ term through $\dim T_\alpha\mathcal{R}(M_1) = 3g-3\equiv g-1 \pmod 2$ terms.

Next, we must compare this orientation of $T_\alpha\mathcal{R}(F')$ with $T_\alpha\mathcal{R}(F')$. These orientations both come from the orientations of $\mathcal{R}(F_0')$ and $\mathcal{R}(F_0)$. These orientations come from a symplectic basis for the two, and thus the two new elements of homology of $H_1(F_0';\mathbb{Z})$ are next two each other in the orientation; that is,

$$T\mathcal{R}(F_0') = \mathfrak{su}(2)\oplus\mathfrak{su}(2)\oplus T\mathcal{R}(F_0)$$

and the same will be true of F', F:

$$T\mathcal{R}(F') = \mathfrak{su}(2)\oplus\mathfrak{su}(2)\oplus T\mathcal{R}(F)$$

The orientation of $H_1(F';\mathbb{Z})$ is naturally the symplectic basis: but the orientations of $H_1(H_k;\mathbb{Z})$ were by a product. These orientations actually differ: the first looks like $\mathbb{R}\oplus\mathbb{R}\oplus H_1(F;\mathbb{Z})$, while the second looks like $\mathbb{R}\oplus H_1(H_1;\mathbb{Z})\oplus\mathbb{R}\oplus H_1(H_2;\mathbb{Z})$. These only agree after we pull the second $\mathbb{R}$ through the first homology, picking up a factor of $(-1)^g$ when we want to make that replacement. Putting these together, we get that

$$\begin{aligned}T_\alpha\mathcal{R}(H_1')\oplus T_\alpha\mathcal{R}(H_2') &= (-1)^{g-1}\mathfrak{su}(2)\oplus\mathfrak{su}(2)\oplus T_\alpha\mathcal{R}(H_1)\oplus T_\alpha\mathcal{R}(H_2)\\ &= \mu(\alpha)(-1)^{g-1}\mathfrak{su}(2)\oplus\mathfrak{su}(2)\oplus T_\alpha\mathcal{R}(F)\\ &= \mu(\alpha)(-1)^{g-1}(-1)^g T_\alpha\mathcal{R}(H_1)\oplus T_\alpha\mathcal{R}(H_2)\\ &= -\mu(\alpha)T_\alpha\mathcal{R}(H_1)\oplus T_\alpha\mathcal{R}(H_2)\end{aligned}$$

As desired. □

So the Casson invariant is well defined for homology 3 spheres, and we may write it simply as $\lambda(\Sigma)$. Although it is easy, let's do a simple calculation first.

Example 4. $\lambda(S^3) = 0$.

Proof. We prove this example twice with two different handlebody decompositions to demonstrate theorem 8.

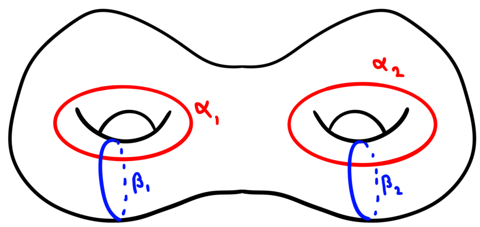

FIGURE 8. Genus 2 standard Heegaard splitting of S^3

(1) S^3 has a handlebody decomposition as the union of two genus 0 handles: its northern and southern hemispheres H_1, H_2, with intersection S^2. All such surfaces are simply connected, so $R(S^3) = R(H_1) = R(H_2) = R(S^2) = \{\theta\}$. So also $R^{irr}(S^3) = R^{irr}(H_1) = R^{irr}(H_2) = R^{irr}(S^2) = \emptyset$, and the same is true of $\mathcal{R}$. Then $\mathcal{R}(H_1)$ trivially intersects $\mathcal{R}(H_2)$ transversely with intersection empty, so there is nothing to count and

$$\lambda(S^3) = \lambda(S^3, H_1, H_2) = \frac{(-1)^0}{2} \sum_{x \in \emptyset} (-1)^{\mu(x)} = 0$$

(2) S^3 also has a handlebody decomposition formed by the standard embedding of the solid genus 2 handlebody. Call this H_1, and the outer one $H_2 = S^3 \setminus (\operatorname{int} H_2)$. Their intersection is the standard embedding of the genus 2 surface F. Then consider the intersection $\mathcal{R}(H_1) \cap \mathcal{R}(H_2)$. Let's take a closer look at the fundamental groups involved: considering the loops in figure 8.

The α loops generate $\pi_1 H_1$, the β loops generate $\pi_1 H_2$, and both together generate $\pi_1 F$. Inside of $\pi_1 F$, the only relation is the product of commutators, and this is the only relation between any α_i and a β_i. If $\phi \in \mathcal{R}(H_1) \cap \mathcal{R}(H_2)$ then there is a representative $\phi_0 \in R(H_1) \cap R(H_2)$. Then the fact that $\phi_0 \in R(H_1) \subset R(F)$ means that $\phi_0 = \phi_0|_{H_1} \implies \phi_0(\beta_i) = 0$ for any β_i. By mirrored reasoning, $\phi_0(\alpha_i) = 0$, thus actually $\phi_0 = \theta$, a contradiction. Thus $\mathcal{R}(H_1) \cap \mathcal{R}(H_2)$ is again empty and so trivially transverse, and we can conclude as before.

□

As a taste for how this definition feels different from the axiomatic one, we will prove parts 2 and 3 of theorem 6 from the representation counting perspective.

Lemma 10. *$\lambda(\overline{\Sigma}) = -\lambda(\Sigma)$ and $\lambda(\Sigma_1 \# \Sigma_2) = \lambda(\Sigma_1) + \lambda(\Sigma_2)$.*

Proof. For simplicity's sake, we prove both cases assuming that Σ is non-degenerate. Also, fix an orientation of $SU(2)$ without further reference.

(1) Pick a genus g Heegaard splitting $\Sigma = H_1 \cup_F H_2$. The orientation of Σ gives an induced orientation to H_1, H_2 and F. Then $\bar{\Sigma}$ clearly has a Heegaard splittiing consisting of $(\bar{H}_1) \cup_{\bar{F}} (\bar{H}_2)$. Every point of $\alpha \in \mathcal{R}(\Sigma)$ has an associated point (the same point!) in $\mathcal{R}(\bar{\Sigma})$. Thus it suffices to prove that $\bar{\mu}(\alpha) \cong \mu(\alpha) + 1 \pmod 2$, where $\bar{\mu}$ is the orientation in $\mathcal{R}(\bar{\Sigma})$ and μ is the orientation in $\mathcal{R}(\Sigma)$.

The sign of $\mu(\alpha)$ is positive or negative corresponding to whether the orientations of the spaces $T_\alpha\mathcal{R}(H_1) \oplus T_\alpha\mathcal{R}(H_2)$ and $T_\alpha\mathcal{R}(F)$ agree. Thus, we want to show that if these orientations do agree, then the orientations of $T_\alpha\mathcal{R}(\bar{H}_1) \oplus T_\alpha\mathcal{R}(\bar{H}_2)$ and $T_\alpha\mathcal{R}(\bar{F})$ disagree. By symmetry, the other case is then the same.

First, we see what happens to the orientations of $H^1(F;\mathbb{R}), H^1(H_1;\mathbb{R})$, and $H^1(H_2;\mathbb{R})$. Changing the orientation of F changes the orientation of $H^1(F)$ by a factor of $(-1)^g$. Since the orientation of $H^1(H_1)$ is such that the product over the oriented generators is the fundamental class, changing the orientation of H_1 changes the orientation of $H^1(H_1)$ by a factor of $(-1)^g$ as well. Since we choose the bases so that $H^1(H_1) \oplus H^1(H_2) = H^1(F)$ is orientation preserving, this means that the orientation of $H^1(H_2)$ does not change.

Next, we study how the orientation of $\mathcal{R}$ changes based on the orientation of H^1. The orientation of $\mathcal{R}(F)$ changes by a factor of $(-1)^{g+1}$, so we just focus on $H^1(H_k)$. $\mathcal{R}(M_k)$ inherits the orientation from $SU(2)^g/SO(3)$ induced by the identification of $\mathrm{Hom}(\pi_1 H_k, SU(2)) \cong SU(2)^g$. Changing the orientation of $H^1(H_k)$ changes the orientation of $\mathrm{Hom}(\pi_1 H_k, SU(2))$ by the same factor, since both correspond to the same generating basis vectors. Thus, changing the orientation of $H^1(H_k)$ by $(-1)^g$ changes the orientation of $\mathcal{R}(H_k)$ by $(-1)^g$, and preserving the orientation of one preserves the orientation of the other.

Finally, we can put it all together. The orientation of $T_\alpha\mathcal{R}(-F)$ is $(-1)^{g+1}$ the orientation of $T_\alpha\mathcal{R}(F)$. Since the orientation of $H^1(-H_1)$ differs by a factor of $(-1)^g$, so does the orientation of $T_\alpha\mathcal{R}(-H_1)$ does as well. Since the orientation of $H^1(-H_2)$ does not change, neither does the orientation of $T_\alpha\mathcal{R}(-H_2)$. All in all, the orientation of $T_\alpha\mathcal{R}(-H_1) \oplus T_\alpha\mathcal{R}(-H_2)$ differs by a factor of $(-1)^g$, whereas the orientation of $T_\alpha\mathcal{R}(-F)$ differs by a factor of $(-1)^{g+1}$, so they will always disagree given that the originals agreed. And that suffices for the proof.

(2) It is certainly true that $R(\Sigma_1\#\Sigma_2) \cong R(\Sigma_1) \times R(\Sigma_2)$. Then since $R(\Sigma_i) = R^{irr}(\Sigma_i) \cup \{\theta\}$,

$$R^{irr}(\Sigma_1\#\Sigma_2) \cup \{\theta_1,\theta_2\} = (R^{irr}(\Sigma_1) \cup \theta_1) \times (R^{irr}(\Sigma_2) \cup \theta_2)$$

$$= [R^{irr}(\Sigma_1) \times R^{irr}(\Sigma_2)] \cup [\theta_1 \times R^{irr}(\Sigma_2)] \cup [R^{irr}(\Sigma_1) \times \theta_2] \cup (\theta_1,\theta_2)$$

Thus

$$R^{irr}(\Sigma_1\#\Sigma_2) = [R^{irr}(\Sigma_1) \times R^{irr}(\Sigma_2)] \sqcup [\theta_1 \times R^{irr}(\Sigma_2)] \sqcup [R^{irr}(\Sigma_1) \times \theta_2]$$

The SO(3) action clearly distributes across the product identifications, so we get

$$\mathcal{R}(\Sigma_1\#\Sigma_2) = [\mathcal{R}(\Sigma_1) \times \mathcal{R}(\Sigma_2)] \sqcup [\theta_1 \times \mathcal{R}(\Sigma_2)] \sqcup [\mathcal{R}(\Sigma_1) \times \theta_2]$$

It is clear that the second and third parts of the union give when counted exactly $\lambda(\Sigma_1)$ and $\lambda(\Sigma_2)$, respectively, and that they add. The last thing we need to do is examine $\mathcal{R}(\Sigma_1) \times \mathcal{R}(\Sigma_2)$. We want to show that the signed count of it is exactly 0. Unfortunately, this part of $\mathcal{R}(\Sigma_1\#\Sigma_2)$ *is* exactly the degenerate part. Elements of $\mathcal{R}(\Sigma_1) \times \mathcal{R}(\Sigma_2)$ are equivalence classes of products (α,β) with neither α nor β reducible. By its definition, the conjugation action distributes, so in particular if (α,β_0) is conjugate to (α,β) by some g then actually $g = 1$. Thus for every α and β, the set $\{(\alpha, g\beta g^{-1}) : g \in SO(3)\}$ consists of disjoint conjugacy classes and as a subspace of $\mathcal{R}(M_1\#M_2)$ is homeomorphic to a copy of $SO(3)$ by the obvious map. So the $\mathcal{R}(\Sigma_1)\#\mathcal{R}(\Sigma_2)$ bit consists of (finitely many) disjoint unions of copies of $SO(3)$.

Since the degenerate parts of our manifold are fortunately submanifolds, we can consider a special perturbation of the intersection: one that retains the Euler characteristic of the submanifold. After such a perturbation, the SO(3) copies break into 0-dimensional components who Euler characteristics must sum (still counting sign) to $\chi(SO(3)) = 0$. The euler characteristic of a point is just 1, so in fact this is the same as saying that the signed count of the points in $\mathcal{R}(\Sigma_1) \times \mathcal{R}(\Sigma_2)$, after the perturbation, is just the sum over all the Euler characteristics of the $SO(3)$ copies, which is 0. Thus this part contributes nothing to the sum, and our addition formula is complete.

For more information about this kind of degeneracy, called *Morse-Bott nondegeneracy*, see [8].

□

In our proof of example 4, it is clear that we got 'lucky' because the generators in the handlebody decomposition of S^3 did not interact with each other. the computation for the Poincare homology sphere $\Sigma(2,3,5)$ is more difficult and warrants study.

Example 5. $\lambda(\Sigma(2,3,5)) = -1$.

Proof. First, we figure out what $\mathcal{R}(\Sigma(2,3,5))$ looks like. Then we examine transversality so that we can count the points without perturbations. Last, we find the orientations.

$\pi_1\Sigma(2,3,5)$ has the presentation

$$\pi_1(\Sigma(2,3,5)) = \langle x_1, x_2, x_3, h | [h, x_k] = 1, x_1^2 = h, x_2^3 = h^{-1}, x_3^5 = h^{-1}, x_1x_2x_3 = 1\rangle$$

Let α be a representation with $\alpha(h) \neq \pm 1$. Then it lies inside of a unique $U(1)$. Since $h \in Z(\pi_1)$, all other elements would also have to be inside of that $U(1)$, thus the whole representation is reducible.

Now assume that $\alpha(h) = 1$. Then $\alpha(x_1)^2 = 1 \implies \alpha(x_1) = \pm 1$. But then $\alpha(x_2) = \pm\alpha(x_3)^{-1}$, and so in particular they commute with each other. But then every element commutes with every other element, so the representation is reducible.

Now assume that $\alpha(h) = -1$. Then $\alpha(x_1)^2 = -1$, so after composing with some conjugation, $\alpha(x_1) = \pm i$. Both of these choices are conjugate to each-other by the matrix $\begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}$ (In general, $e^{i\theta}$ is conjugate to $e^{-i\theta}$ by this matrix).

Similarly, the power condition says that $\alpha(x_2)$ and $\alpha(x_3)$ will be in the conjugacy classes of primitive 3rd, 5th roots of -1, respectively, with positive imaginary part. But there is only one such option for x_2, and two such options for x_3. Thus, the conjugacy class of $\alpha(x_1), \alpha(x_2)$ are fixed, and there at most two choices for conjugacy classes of $\alpha(x_3)$. One can manually check that this means that there are at most 2 elements of $\mathcal{R}(\Sigma(2,3,5))$.

Using MatLab, we realize these both of these representations by picking:

$$\alpha(x_1) = \begin{pmatrix} i & 0 \\ 0 & -i \end{pmatrix}, \alpha(x_2) = \begin{pmatrix} \frac{1}{2} - \cos(\frac{2\pi}{5})i & \cos(\frac{\pi}{5}) \\ -\cos(\frac{\pi}{5}) & \frac{1}{2} + \cos(\frac{2\pi}{5})i \end{pmatrix}, \alpha(x_3) = \begin{pmatrix} \cos(\frac{\pi}{5}) & \frac{1}{2} + \cos(\frac{2\pi}{5})i \\ -\frac{1}{2} + \cos(\frac{2\pi}{5})i & \cos(\frac{\pi}{5}) \end{pmatrix}$$

And

$$\alpha(x_1) = \begin{pmatrix} i & 0 \\ 0 & -i \end{pmatrix}, \alpha(x_2) = \begin{pmatrix} \frac{1}{2} - \cos(\frac{\pi}{5})i & \cos(\frac{2\pi}{5}) \\ -\cos(\frac{2\pi}{5}) & \frac{1}{2} + \cos(\frac{\pi}{5})i \end{pmatrix}, \alpha(x_3) = \begin{pmatrix} \cos(\frac{\pi}{5}) - \frac{1}{2}i & -\cos(\frac{2\pi}{5})i \\ \cos(\frac{2\pi}{5}) & \frac{1}{2} + \cos(\frac{2\pi}{5})i \end{pmatrix}$$

These two representations satisfy the required relations that $\alpha(x_1)^2 = \alpha(x_2)^3 = \alpha(x_3)^5 = -1$, and $\alpha(x_1)\alpha(x_2)\alpha(x_3) = 1$. They can be shown to not be conjugate to each-other by the following reason: suppose they were. Then since $\alpha(x_1) = i$ in both cases, the element which conjugates them must (as a quaternion) have no j or k components, and so is of the form $\begin{pmatrix} e^{i\theta} & 1 \\ e^{-i\theta} & 0 \end{pmatrix}$. But it is clear that no such elements conjugate $\alpha(x_2)$ and $\alpha(x_3)$ in the required manner. Thus, $R^{irr}(\Sigma(2,3,5))$ consists of two conjugacy classes, and $\mathcal{R}(\Sigma(2,3,5)) = R^{irr}(\Sigma(2,3,5))/SO(3)$ (modding out by exactly conjugation) consists of 2 points.

Next, we want to make sure that these two points actually intersect transversely. By lemma 8, we can do this without reference to the handlebody decomposition—just by looking at $H^1_\alpha(\pi_1\Sigma(2,3,5);\mathfrak{su}(2))$. Fix one of the above irreducible representations α. Let $\zeta : \pi_1\Sigma(2,3,5) \to \mathfrak{su}(2)$ be a cocycle, so that $\zeta(xy) = \zeta(x) + \mathrm{Ad}_{\alpha(x)}(y)$. First off, since h is in the center of $\pi_1\Sigma(2,3,5)$, for any $g \in \pi_1\Sigma(2,3,5)$,

$$\zeta(h) + \zeta(g) = \zeta(h) + \mathrm{Ad}_{\alpha(h)}\,\zeta(g) = \zeta(hg) = \zeta(gh) = \zeta(g) + \mathrm{Ad}_{\alpha(g)}\,\zeta(h)$$

$$\implies \mathrm{Ad}_{\alpha(g)}\,\zeta(h) = \zeta(h)$$

For all g. But since α is irreducible, the only point which is fixed under every adjoint action is 0, so $\zeta(h) = 0$. The relations $x_1^2 = x_2^3 = x_3^5 = h$ then give rise to the requirements

$$(1 + \mathrm{Ad}_{\alpha(x_1)})\zeta(x_1) = 0$$

$$(1 + \mathrm{Ad}_{\alpha(x_2)} + \mathrm{Ad}^2_{\alpha(x_2)})\zeta(x_2) = 0$$

$$(1 + \mathrm{Ad}_{\alpha(x_3)} + \mathrm{Ad}^2_{\alpha(x_3)} + \mathrm{Ad}^3_{\alpha(x_3)} + \mathrm{Ad}^4_{\alpha(x_3)})\zeta(x_3) = 0$$

As well as $x_1x_2x_3 = 1$ giving the requirement

$$0 = \zeta(x_1x_2x_3) = \zeta(x_1) + \mathrm{Ad}_{\alpha(x_1)}\,\zeta(x_2) + \mathrm{Ad}_{\alpha(x_1)}\,\mathrm{Ad}_{\alpha(x_2)}\,\zeta(x_3)$$

Each of these equations we can consider as operators on $\mathfrak{su}(2)$ which give linear constraints on the choices of $\zeta(x_k)$. Initially, we had $\mathbb{R}^3$ choices for each $\zeta(x_k)$ ($\mathbb{R}^9$ total): each of the first three conditions is a requirement that $\zeta(z_k)$ is in the kernel of the relevant operator, which can be calculated to be 2 dimensional. Thus we get 1 constraint from each of the first 3 equations. The last equation can also be viewed as an operator from the set of (thus far valid) choices of $\zeta(x_k)$ into $\mathbb{R}^3$. Since α is irreducible, this operator must be surjective and thus we get an $\mathbb{R}^3$ of constraints from this. Thus the dimension of the cocycles is $9 - 3 - 3 = 3$.

The coboundaries are the combined span of Image$(1 - \mathrm{Ad}_{\alpha(g)})$ ranging over all $g \in \pi_1(\Sigma(2,3,5))$. Since the α are irreducible, these combined spans must be all of $\mathbb{R}^3$: if they were not, then there would be some nontrivial kernel and thus element $x \in \mathfrak{su}(2)$ which was fixed by every $\mathrm{Ad}_{\alpha(g)}$, violating irreducibility. Thus the space of coboundaries is 3 dimensional. So the cohomology is $Z^1_\alpha(\pi_1\Sigma(2,3,5);\mathfrak{su}(2))/B^1_\alpha(\pi_1\Sigma(2,3,5);\mathfrak{su}(2)) = \mathbb{R}^3/\mathbb{R}^3 = 0$. Thus by lemma 8, the intersections are transversal.

Lastly, we discuss orientations, following [6] In fact, at both points α, the intersection gives a sign of -1. In order to do this, we need an explicit Heegaard splitting. An exercise in Kirby calculus shows that surgery on the LHS of figure 9 also gives $\Sigma(2,3,5)$, and the RHS demonstrates a genus 2 Heegaard splitting of the homology sphere. The handlebody which is formed from the -2 and 3 surgeries as depicted is H_1, and the compliment is H_2.

The orientations of the representation spaces start by orienting the homologies of the handlebodies, which starts with the fundamental group of the boundary surface F, which is symplectically oriented by the curves (m_1, l_1, m_2, l_2). The orientation of the handlebodies should be thought of with respect to their inclusions of $H_1(F;\mathbb{Z})$ into them, i.e., thought of in terms of the generators m_k, l_k. After finding this, one finds local parameterizations of the two spaces $\mathcal{R}(H_1), \mathcal{R}(H_2)$ around each point α and compare them to a parameterization of $\mathcal{R}(F)$. Lescop calculates in [6] that the

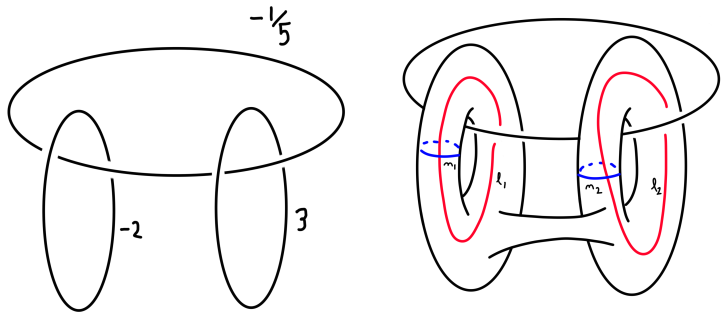

FIGURE 9. A genus 2 Heegaard splitting of $\Sigma(2,3,5)$

determinant of the diffeomorphism going between them is -1 regardless of α, and so each point contributes a -1. Thus we have proven $\lambda(\Sigma(2,3,5)) = \frac{1}{2}(-1-1) = -1$. □

Corollary. *The Casson invariant is not completely determined by the fundamental group.*

Proof. By Van Kampen's theorem, it is easy to see that $\Sigma(2,3,5)\#\Sigma(2,3,5)$ and $\Sigma(2,3,5)\#\overline{\Sigma(2,3,5)}$ have the same fundamental group. But

$$\begin{aligned}\lambda(\Sigma(2,3,5)\#\Sigma(2,3,5)) &= 2\lambda(\Sigma(2,3,5)) = -2\\ &\neq 0 = \lambda(\Sigma(2,3,5)) - \lambda(\Sigma(2,3,5))\\ &= \lambda(\Sigma(2,3,5)\#\overline{\Sigma(2,3,5)})\end{aligned}$$

where we used Theorem 6. □

Thus the Casson invariant, despite its initial definition as counting representations of the fundamental group, actually ends up containing slightly different information, the differences being mostly related to orientation.

3.3. **Applications: property P.** Before the Poincare conjecture was proven, one potential method of finding a counterexample was to simply apply Dehn surgery on a knot in S^3 and hope that (a) the fundamental group of the resulting 3 manifold was trivial, and (b) that some other 3 manifold invariant was nontrivial. A knot is said to have property P if no surgery on it could possibly give such a counter-example, i.e., that no nontrivial surgery on the knot results in a manifold with nontrivial fundamental group. The statement of property P is every knot except the unknot has property P. Kronheimer and Mrowka proved this in general in 2003 in [5]; the tools of the Casson invariant do not get us all the way to the proof, but they make significant progress.

It has been known for some time that a nontrivial surgery on a nontrivial knot never itself gives S^3 back, so another way of proving property P would be to prove the Poincare conjecture. Perelman proved the Poincare conjecture in 2003, in fact before Kronheimer and Mrowka's paper.

Theorem 9. *Let k be a knot. If there exists some nontrivial p, q such that p/q surgery along k on S^3 has $\pi_1(S^3 + p/q \cdot k) = 0$, then $\Delta_k''(1) = 0$.*

Proof. First, we remember that $1/0$ surgery is the identity and so is trivial. We now assume that $q \neq 0$. If $p \neq 1$, then the resulting manifold will have homology $\mathbb{Z}/p\mathbb{Z}$ (by Mayer Vietoris) so will have a nontrivial fundamental group. Thus we can also assume that $p = 1$. The remaining cases are all of the form $1/q$. Then by our computation of the Casson invariant with $q \neq 0$ and $\Delta_k''(1) \neq 0$, we know that

$$\lambda(S^3 + \frac{1}{q} \cdot k) = \frac{q}{2}\Delta_k''(1) \neq 0$$

Since a nontrivial Casson invariant implies a nontrivial homomorphism of the fundamental group into SU(2), it must be in particular that the fundamental group is non-trivial, which suffices for the proof. □

4. Combinatorics

The Casson-Walker invariant, a generalization of the Casson invariant, serves two functions. Firstly, it extends the Casson invariant from an invariant of integer homology 3 spheres to one of rational homology 3 spheres, opening up the properties of the invariant to a much larger class of manifolds; of particular relevance is Dehn surgery with coefficients other than $\frac{1}{q}$.

The invariant may be defined in two different ways, both shown by Walker in [12]. The first is a construction similar to the SU(2) representation definition of the Casson invariant. We will not go over the details of this, but the main innovation is a way to count the reducible representations, which now make some real contribution. The second definition is purely combinatorial. Unlike the axiomatic definition of the Casson invariant, this definition actually establishes existence of the invariant, and retroactively the existence of the Casson invariant as well, using purely combinatorial means. A disadvantage of the combinatorial definition is that it lacks an immediate topological interpretation which the first one gives.

Definition. *Let $((\cdot)) : \mathbb{R} \to \mathbb{R}$ be the sawtooth function*

$$((x)) = \begin{cases} x - \lfloor x \rfloor - 1/2 & x \in \mathbb{R} \setminus \mathbb{Z} \\ \frac{1}{2} & x \in \mathbb{Z} \end{cases}$$

Then define the Dedekind sum $s : \mathbb{Z} \times (\mathbb{Z} \setminus 0) \to \mathbb{R}$ by

$$s(q,p) = \operatorname{sign}(p) \cdot \sum_{k=1}^{|p|} \left(\left(\frac{k}{p}\right)\right)\left(\left(\frac{kq}{p}\right)\right)$$

Let $\langle \cdot, \cdot \rangle$ be the bilinear form on $H^1(T^2; \mathbb{Z})$ defined by Poincare duality. In the case of the torus, we remark that in particular this form is alternating and in a suitable basis is of the form

$$\begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}$$

Definition. *Let $a, b, l \in H_1(T^2; \mathbb{Z})$ so that a and b are represented by simple closed curves and $\langle a, l \rangle, \langle b, l \rangle \neq 0$. Then choose a basis $x, y \in H_1(T^2; \mathbb{Z})$ so that $l = dy$ for some $d \in \mathbb{Z}$ and $\langle x, y \rangle = 1$. Then define*

$$\tau(a, b; l) = -s(\langle x, a \rangle, \langle y, a \rangle) + s(\langle x, b \rangle, \langle y, b \rangle) + \frac{d^2 - 1}{12} \frac{\langle a, b \rangle}{\langle a, l \rangle \langle b, l \rangle}$$

This definition is independent of the choice of x, y; this can be verified by picking a new basis, expanding it out in terms of the old one, and using a couple identities related to s.

Definition. *The Casson Walker invariant is an invariant λ_W of oriented rational homology spheres with the following combinatorial description:*

(1) $\lambda_W(S^3) = 0$.

(2) *If $k \subset \Sigma$ is a knot in a rational homology sphere, l is a longitude of $E(k)$, $a, b \in H_1(\partial E(k); \mathbb{Z})$ are primitive, and $\langle a, l\rangle, \langle b, l\rangle \neq 0$, then*

$$\lambda_W(\Sigma + b \cdot k) = \lambda_W(\Sigma + a \cdot k) + \tau(a, b; l) + \frac{1}{2}\frac{\langle a, b\rangle}{\langle a, l\rangle\langle b, l\rangle}\Delta''_{k\subset\Sigma}(1)$$

For this to be a well defined invariant, we must show two things: (1) that every rational homology sphere can be described by a sequence of surgeries satisfying the above restrictions, and (2) that out of every such sequence arises the same number. We will do this at the end, but first we shall get some practice in computing it.

Theorem 10. *Let k be a knot in an integral homology sphere Σ. Then for $q, p \neq 0$, $\lambda_W(\Sigma + \frac{p}{q} \cdot k) = \lambda_W(\Sigma) - s(q, p)/2 + \frac{q}{2p}\Delta''_k(1)$.*

Proof. Since $q, p \neq 0$, the resulting surgery is indeed a rational homology sphere. Let $E(k)$ be the knot exterior as usual, and identify $N(k)$ with $D^2 \times S^1$. Now, let m be the meridian and l the longitude of $\partial E(k)$, such that m generates $H^1(E(k))$ and vanishes in $H^1(E(k))$, while l generates $H^1(E(k))$ and vanishes in $H^1(E(k))$. Orient m, l so that $\langle m, l\rangle = 1$.

Then we can set $a = pm + ql$, $b = m$. These are both primitive because (p, q) is coprime. The p/q surgery is defined by sending $\partial D^2 \times \{0\}$ to a, b, respectively.

Since $q \neq 0$, we know that $\langle a, l\rangle = p$, $\langle b, l\rangle = 1$; neither of these are 0, so a, b, l satisfies the surgery formula conditions. We can now pick a basis x, y by setting $x = m + l$, $y = l$. We get that $y = 1 \cdot l$ so that $d = 1$, and that $\langle x, y\rangle = 1$. Thus

$$\begin{aligned}\tau(a, b; l) &= -s(\langle x, a\rangle, \langle y, a\rangle) + s(\langle x, b\rangle, \langle y, b\rangle) + \frac{d^2 - 1}{12}\frac{\langle a, b\rangle}{\langle a, l\rangle\langle b, l\rangle} \\ &= -s(\langle m + l, pm + ql\rangle, \langle l, pm + ql\rangle) + s(\langle m + l, m\rangle, \langle l, m\rangle) \\ &= -s(q - p, -p) + s(-1, -1)\end{aligned}$$

Since $s(-1, -1) = 0$, and $s(q - p, -p) = -s(q, p)$, we get that $\tau(a, b; l) = s(q, p)$. Now notice that $K_b = \Sigma$, since it is $1/0$ surgery. Thus we have that

$$\lambda_W(K_b) = \lambda_W(K_a) + \frac{1}{2}\tau(a, b; l) + \frac{\langle a, b\rangle}{\langle a, l\rangle\langle b, l\rangle}\frac{1}{2}\Delta''_k(1) = \frac{s(q, p)}{2} - \frac{q}{2p}\Delta''_k(1)$$

$$\implies \lambda_W(\Sigma + \frac{p}{q} \cdot k) = \lambda_W(\Sigma) - \frac{1}{2}s(q, p) + \frac{q}{2p}\Delta''_k(1)$$

□

Remark. In particular, since lens spaces $L(p, q)$ are made by doing $-p/q$ surgery on the unknot in S^3, this tells us that $\lambda_W(L(p, q)) = -s(q, p)/2$.

Corollary. *If Σ is an integer homology sphere, then $\lambda(\Sigma) = \lambda_W(\Sigma)$.*

Proof. By lemma 2, every homology sphere can be obtained by a sequence of integer surgeries on knots with coefficients ± 1, with each intermediate surgery a homology sphere (and starting with S^3). By definition, $\lambda_W(S^3) = \lambda(S^3) = 0$. Thus it suffices by induction to prove that if $\lambda(\Sigma) = \lambda_W(\Sigma)$, then $\lambda(\Sigma \pm k) = \lambda_W(\Sigma \pm k)$. By the above theorem,

$$\lambda_W(\Sigma \pm k) = \lambda_W(\Sigma) - \frac{1}{2}s(\pm 1, 1) + \frac{\pm 1}{2(1)}\Delta''_k(1)$$

$$= \lambda(\Sigma) \pm \frac{1}{2}\Delta''_k(1) = \lambda(\Sigma \pm k)$$

Where we use the easy calculation that $s(\pm 1, 1) = 0$. □

Now we should actually prove existence.

Definition. *A permissible surgery sequence (PSS), denoted $\mathcal{N}$, of length n is, for each $1 \leq i \leq n$, the following data: a (closed) rational homology sphere N^i, a knot $k_i \subset N^i$, and two primitive elements $m_i, s_i \in H_1(E(k); \mathbb{Z})$ such that $N^i_{m_i} \cong N^i$ and $N^i_{s_i} \cong N^{i+1}$ (for $i < n$). Furthermore, $N^1 \cong S^3$. We say that $\mathcal{N}$ represents N^n. For each i, we let l_i be the canonical longitude of k_i.*

For a PSS $\mathcal{N}$, we say that

$$\lambda(\mathcal{N}) := \sum_{i=1}^{n} \left(\tau(m_i, s_i; l_i) + \frac{1}{2} \frac{\langle m_i, s_i \rangle}{\langle m_i, l_i \rangle \langle s_i, l_i \rangle} \Delta''_{k_i \subset N^i}(1) \right)$$

Then the existence of the Casson-Walker invariant becomes the statement that every rational homology sphere is represented by some PSS, and that any two PSS's that represent the same rational homology sphere have the same invariant $\lambda(\mathcal{N})$.

By viewing each knot as living in the original S^3, each PSS gives rise to a framed link $\mathcal{L}$, along with an ordering of the links components $(k_i)_i$. Such links are called ordered framed links (OFLs). If the framing of each component is an integer, they are called integral OFLs.

Theorem 11. *Every rational homology sphere can be reached by a sequence of such surgeries.*

Proof. This is a corollary of lemma 4: pick some sequence of knots k_i, link $\mathcal{L}$, and partial links $\mathcal{L}_i$ as in the lemma. Let N^i be the surgery on $\mathcal{L}_{i-1}$, and the primitive elements m_i, s_i be those represented by $1/0$ surgery and n_i surgery, respectively. Then by the lemma, each N^i is a rational homology sphere and N^n is surgery on $\mathcal{L}$; and we are done. □

Remark. In the language of OFLs, this tells us that every rational homology sphere is given by surgery on an integral OFL $\mathcal{L}$, such that the link $\mathcal{L}_{k-1}$ is also a rational homology sphere. Call such integral OFLs permissible. Then we can rephrase the theorem as stating that every rational homology sphere is represented by an integral, permissible OFL.

We can then also go the other way, saying that a PSS is integral if its corresponding OFL is integral.

Theorem 12. *Any two permissible, integral surgery sequences leading to the same rational homology sphere give the same value of λ_W.*

Proof. A slight variation of Kirby's theorem gives that any two permissible, integral OFLs are related by a sequence of four moves:

(1) isotopy of the link.
(2) Add or remove an unlinked ± 1 unknot anywhere within the sequence, as in the first Kirby move.
(3) slide k_i over k_j for $i > j$ as in the second Kirby move.
(4) Swap the ordering of k_i, k_{i+1}, as long as the result still gives a PSS.

One then needs to show that the Casson-Walker invariant does not change when applying any of the above moves. Isotopy is clear: none of the values of the homology classes or surgered results change. Next, we consider the first Kirby move. On the OFL side, we are slotting in a ± 1 unlinked unknot k in between positions i and $i + 1$. The actual sign doesn't matter, so let's assume its $+1$ surgery. Surgery on this does not change the manifold, so the resulting sequence is still a permissible, integral OFL. The $+1$ surgery sends the meridian m to $m + l$, with longitude l. Thus in the language of the definition of the surgery formula, $a = m$, $b = m+l$, and we can check that $\langle a, l\rangle = \langle m, l\rangle = 1 \neq 0$ and similarly with $\langle b, l\rangle$. Since the manifold does not change before or after the surgery, the contribution of each other component of the PSS does not change in the calculation of λ. The contribution of this component is

$$\tau(m, m+l; l) + \frac{1}{2}\frac{\langle m, m+l\rangle}{\langle m, l\rangle\langle m+l, l\rangle}\Delta''_{k\subset N^i}(1)$$

The first term is 0 by calculating out τ: indeed, we get that the d in the definition of τ is just 1, and the Dedekind functions vanish by basic properties thereof. The second term is also 0 because, being an unknot in N^i, its Alexander polynomial is trivial.

The second Kirby move may appear to change PSS at first. However, because we are sliding the knot k_i over k_j for $i > j$, there is no effect on the surgered spheres for $k \leq i$, as the only knot which changes is k_i. Then by the definition of the handle slide, after the surgery by k_j, the image of the post-slide knot k'_i becomes identical to the pre-slide knot k_i. Thus the actual PSS does not change except by an isotopy, with no changes to the knots, closed curves, or surgered manifolds in the sequence.

The proof of invariance under the final move can be found in [12]. The main idea is just to carefully compute how the Alexander polynomial changes under the surgery. If it is the k_i and k_{i+1} knots which have their orders exchanged, then we can consider the manifold E which is the combined knot exterior of both k_i and k_{i+1}. Then there is an inclusion $i_* : H_1(\partial E(k_i); \mathbb{Q}) \to H_1(E; \mathbb{Q})$. The generator of $H_1(\partial E(k_i); \mathbb{Q})$ which is also a generator of $H_1(E(k_i); \mathbb{Q})$ does not vanish, so the map is at least rank 1. Whether it is an isomorphism depends on the image of the generator which is null-homologous in $H_1(E(k_i); \mathbb{Q})$, but potentially not null-homologous in $H_1(E; \mathbb{Q})$. Then the proof of the result relies on carefully examining the effect on the Alexander polynomial under the two cases of i_* being rank one or an isomorphism. It is complicated by the fact that the rational homology sphere definition of the Alexander polynomial must be used, which is harder to compute than the analogous definition for integral homology spheres.

□

4.1. **Applications: cosmetic surgery.** Cosmetic surgery is a problem asking how unique Dehn surgery is. Ideally, we would like Dehn surgery to be as close to a 1-to-1 description of closed 3 manifolds as possible, so uniqueness of Dehn surgery would be a striking result. Kirby calculus answers the question of when surgery on two different knots or links gives the same manifold. A different direction is to ask when Dehn surgery on the same knot with two different coefficients can give the same manifold. For the unknot, this can indeed happen, but conjecturally this cannot happen with any other knots. This is called the cosmetic surgery conjecture. It comes in two flavors:

Conjecture 1 (Purely cosmetic surgery). *Let $k \subset S^3$ be a non-trivial knot. Then if r, s are two surgery coefficients and $S^3 + r \cdot k$ is orientation preserving homeomorphic to $S^3 + s \cdot k$, then $r = s$.*

Conjecture 2 (Chirally cosmetic surgery). *Let $k \subset S^3$ be a knot which is not isotopic to its mirror image, and is not a $(2, n)$-torus knot. Then if r, s are two surgery coefficients and $S^3 + r \cdot k$ is orientation reversing homeomorphic to $S^3 + s \cdot k$, in fact $r = s$.*

There has been much progress related to these two conjectures, mostly of the form of obstructions to admitting such cosmetic surgeries. The Casson-Walker invariant allows us to give some such obstructions.

Theorem 13. *Let k be a knot, and let $m \neq n \in \mathbb{Z}$. If p/q surgery and p/q' surgery on k result in orientation preserving diffeomorphic manifolds, then $(q - q')\Delta''_k(1) = p(s(q, p) - s(q', p'))$.*

Proof. We may assume that $q \neq q'$ and $q', q \neq 0$. By the surgery formula,

$$\lambda_W(S^3 + \frac{p}{q} \cdot k) = \lambda_W(S^3 + \frac{p}{q'} \cdot k) \implies -\frac{1}{2}s(q,p) + \frac{q}{2p} \cdot \Delta_k''(1) = -\frac{1}{2}s(q',p) + \frac{q'}{2p} \cdot \Delta_k''(1)$$
$$\implies (q - q')\Delta_k''(1) = p(s(q,p) - s(q',p))$$

□

Corollary. *If k admits purely cosmetic integer surgery via the pair $\pm m$, then $\Delta_k''(1) = \frac{(m-1)(m-2)}{12}$.*

Proof. In this case, $q = 1$, $q' = -1$, and $p = m$, Then $s(q,p) - s(q',p) = s(1,m) - s(-1,m) = 2s(1,m)$, and an explicit formula for $s(1,m)$ concludes. □

Remark. If k is instead a null-homologous knot in a rational homology sphere Σ (such that Dehn-surgery by coefficients still makes sense), then the above proof carries out verbatim, since the only extra term added is a $\Delta_W(\Sigma)$ term on both sides, which cancels.

Since $q = q' \pmod p$ implies that $s(q,p) = s(q',p)$, this gives us a plethora of new values where $\Delta_k''(1)$ must vanish. It turns out that for p prime, $s(q,p) = s(q',p)$ iff $q = q' \pmod p$ or $qq' = 1 \pmod p$.

Theorem 14 ([11])**.** *If there is a chirally cosmetic surgery $p/q, p/q'$ with $p, q, q' \neq 0$ on a knot k, then*

$$(q + q')\Delta_k''(1) = p(s(q',p) + s(q,p))$$

Proof. If they are orientation *reversing* homeomorphic, then

$$\lambda_W(S^3 + \frac{p}{q} \cdot k) = -\lambda_W(S^3 + \frac{p}{q'} \cdot k)$$
$$\implies -\frac{1}{2}s(q,p) + \frac{q}{2p}\Delta_k''(1) = -(-\frac{1}{2}s(q',p) + \frac{q'}{2p}\Delta_k''(1))$$
$$\implies (q + q')\Delta_k''(1) = p(s(q',p) + s(q,p))$$

□

We can extract more information on Δ_k by performing topological operations on the knot. The general idea is as follows: many processes that alter 3 manifolds in proscribed ways cooperate nicely with surgery. In particular, if $\Sigma \to \Sigma'$ is some process, and $\Sigma = \Sigma_0 + n \cdot k$, then we hope that the process 'preserves' this surgery description, in the sense that $\Sigma' = \Sigma_0' + n' \cdot k'$, for k' a new knot. In the most favorable cases, k' is related, but not identical to, k. Thus, we can extract information about k and Σ by looking at the properties of k'. Let's look at this more specifically.

The specific case we will cover is that of double coverings. If $\Sigma = S^3 \pm 2 \cdot k$, then $H_1(\Sigma; \mathbb{Z}) \cong \mathbb{Z}/2$. Thus the commutator $[\pi_1\Sigma, \pi_1\Sigma] \subset \pi_1\Sigma$ is index 2, and this gives rise to a double covering $\tilde{\Sigma}$ over Σ by covering space theory. While it doesn't make sense to create an index 2 cover of S^3, we can create a *branched cover* of S^3 over k.

Since $H_1(E(k); \mathbb{Z}) \cong \mathbb{Z}$, fundamental group π_1 has a unique index 2 subgroup containing the commutator group; call this subgroup $H \subset \pi_1$; thus there is a unique index 2 covering space $\tilde{E}(k)$. It is in particular finite index and is thus also a compact 3 manifold with boundary. Said boundary must be homeomorphic to $S^1 \times S^1$ because the boundary is itself a two sheeted covering space for the boundary $\partial E(k) \cong S^1 \times S^1$, and an two sheeted covering space of a torus is a torus. There is a unique way to glue in a $D^2 \times S^1$ such that the result extends the cover to a branched cover $p : X \to S^3$. The branching set of the cover are $k \subset S^3$ the way this can be thought of is that the knot exterior can be as thin as we want it to be (creating an honest covering space no matter how thin), so the eventual branched cover should only be branched at k itself. Then $\tilde{k} := p^{-1}(k)$, the branching set in X, is also a null homologous knot, and the result becomes:

Lemma 11 ([1], [3])**.** *In the notation of above ($\Sigma = S^3 \pm 2k$), $\tilde{\Sigma} = X \pm \tilde{k}$. Furthermore, $\Delta_{\tilde{k}}(t) = \Delta_k(t^{1/2})\Delta_k(t^{-1/2})$.*

Theorem 15. *Suppose $k \subset S^3$ has a Seifert surface with genus 2, and suppose that k admits a ± 2 cosmetic surgery. Then $\Delta_k(t) = 1$ identically.*

Proof. If F is a genus 2 Seifert surface for k, then the corresponding Seifert matrix is 4×4, which means the Alexander polynomial has non-zero coefficients of degrees at most t^2. By symmetry of the Alexander polynomial of a knot, we can write $\Delta_k(t)$ as

$$\Delta_k(t) = at^2 + bt + c + bt^{-1} + at^{-2}$$

Since k admits a ± 2 cosmetic surgery, by corollary 4.1, $\Delta_k''(1) = 0$. By the above lemma, $\tilde{k}$ admits a ± 1 cosmetic surgery, and so we similarly get $\Delta_{\tilde{k}}''(1) = 0$ by the remark after 4.1. Writing out the Alexander polynomial for $\Delta_{\tilde{k}}(t)$ based on the above lemma, we have

$$\Delta_{\tilde{k}}(t)\Delta_k(t^{1/2})\Delta_k(t^{-1/2}) = a^2t^2 + (2ac - b^2)t + (2a^2 - 2b^2 + c^2) + (2ac - b^2)t^{-1} + a^2t^{-2}$$

Thus we have the two relations:

$$0 = \Delta_k''(1) = [2a + 2bt^{-3} + 6at^{-4}]|_{t=1} = 2a + 2b + 6a \implies b = -4a$$

$$0 = \Delta_{\tilde{k}}''(1) = 2(a^2) + 2(2ac - b^2) + 6(a^2) \implies 4a^2 + 2ac = b^2$$

Putting these together gives us

$$4a^2 + 2ac = 16a^2 \implies a(12a - c) = 0$$

Thus $a = 0$ or $12a = c$. For the last restriction, we remember that $|\Delta_k(1)| = 1$, so that

$$2a + 2(-4a) + c = \pm 1 \implies c = 6a \pm 1$$

Thus it must be that $a = b = 0$, and so $\Delta_k(t) = \pm 1$, which up to normalization is just $\Delta_k(t) = 1$. □

Remark. The above theorem has more relevance than it seems. It turns out that *all* non-trivial knots with ± 2 purely cosmetic surgery have genus 2 Seifert surfaces. In fact, [1] proves that all non-trivial knots with cosmetic surgery must have surgeries of the form ± 2, so in fact every knot with cosmetic surgery has trivial Alexander polynomial.

5. Gauge theory

When constructing the Casson invariant, it may seem as if we are throwing away a certain amount of information. For example, in the case of $\Sigma(2, 3, 5)$, no pertubation was needed to get $\mathcal{R}^{irr}(H_1)$, $\mathcal{R}^{irr}(H_2)$ into general position, so perhaps there is a way to retain not just the signed count of the points, but to relate the points to each-other in a well defined way. Because in some manifolds a pertubation is needed, the signed count was the most we could do.

Gauge theory provides an equivalent way of investigating the structure of $\mathrm{Hom}(\pi_1\Sigma, SU(2))/SO(3)$ which preserves more of its details. Although this Instanton Floer Homology is too much to discuss in detail, we can at least see the beginnings of the tools and its power.

Let M be a 3 manifold, and let $P = M \times SU(2)$. The projection $\pi : M \times SU(2) \to M$ makes P into a fiber bundle over M. There is more structure here: there is a natural group action $SU(2) \times P \to SU(2)$ given by $g \cdot (x, h) = (x, hg)$ (this is called a principle bundle). As with many fiber bundles, we are concerned with the lifting of paths. Let $\gamma : [0,1] \to M$, and pick a specific point $(\gamma(0), y) \in \pi^{-1}(\gamma(0))$. Then we may ask for a lift $\tilde{\gamma}$ making the following diagram commute:

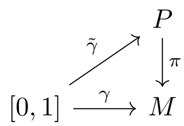

and such that $\tilde{\gamma}(0) = (\gamma(0), y)$. Consider the collection of functions which take in a path and an initial lift, and output a full lift. Firstly, there is the obvious 'trivial' lift: $\tilde{\gamma}(t) = (\gamma(t), y)$. But there are many more. To restrict the options, we require that these functions vary smoothly and respect the group action, in a sense that will be made precise. It turns out that these types of functions appear naturally as a consequence of the notion of connections.

Definition. *Let $\Omega^k(M; \mathfrak{su}(2)) = \Gamma(\bigwedge^k(T^*M) \otimes \mathfrak{su}(2))$ be the space of global sections of the kth exterior power of the cotangent bundle with coefficients in $\mathfrak{su}(2)$. $\mathscr{A} := \Omega^1(M; \mathfrak{su}(2))$ is called the space of connections, and the elements of $\mathscr{A}$ are connections.*

A connection A is thus a smoothly varying collection of linear maps $A_p : T_pM \to \mathfrak{su}(2)$. We can extend A to be defined over P and take in any element (v, u) of $T(M \times SU(2)) \cong TM \times \mathfrak{su}(2)$ by the rule

$$A_{(p,g)}(v, u) = \operatorname{Ad}_{g^{-1}} A_p(v)$$

This definition has the property that for any $g, h \in SU(2)$,

$$\operatorname{Ad}_h A_{(p,gh)} = \operatorname{Ad}_h \operatorname{Ad}_{h^{-1}g^{-1}} A_p = \operatorname{Ad}_{hh^{-1}} \operatorname{Ad}_{g^{-1}} A_p = A_{(p,g)}$$

Connections are the answer to our path lifting criteria. Given a connection A, we can require all lifts $\tilde{\gamma}$ of γ to satisfy $A_{\tilde{\gamma}(t)}(\tilde{\gamma}'(t)) = 0$ for all t. In fact, given a connection A, there always exists such a lift and it is unique, by existence and uniqueness of ODE's. So we can make the following definition:

Definition. *The holonomy of A is the map*

$$\operatorname{hol}_A : \gamma \mapsto \tilde{\gamma}(1)$$

Which inputs a path $\gamma : [0,1] \to M$ with base-point x, and whose output concerns the unique lift of γ along A, outputting in $SU(2)$.

hol_A is a morphism of groupoids $\Omega M \to SU(2)$, where ΩM (not to be confused with $\Omega^k M$) is the loop space of A at the point x. The image of this morphism is called the holonomy group of A.

Definition. *A connection is called flat if, for any loops γ, η in M which are based homotopic, $\operatorname{hol}_A(\gamma) = \operatorname{hol}_A(\eta)$.*

Given a flat connection and some choice of basepoint, hol_A thus defines a representation $\pi_1 M \to SU(2)$. If $\mathscr{F} \subset \mathscr{A}$ is the space of flat connections, the map $A \mapsto \operatorname{hol}_A$ so gives a map $\mathscr{F} \to \operatorname{Hom}(\pi_1 M, SU(2)) = R(M)$. This space is at the very least non-empty: the 'trivial' connection θ, which lifts the paths according to the trivial lift. Now one may start to see where we are going.

Theorem 16. *There is a bijective correspondence $\mathscr{F} \leftrightsquigarrow R(M)$.*

Proof. We already have a map hol going one way. To go the other, we follow Floer in [2]. $\pi_1 M$ acts on the universal cover $\tilde{M}$, and given a representation $\phi \in R(M)$, the action extends to an action on $\tilde{M} \times \mathbb{C}^2$. π_1 acts on the first factor via deck transformations, and on the second factor via its representation into $SU(2)$, which then acts on $\mathbb{C}^2$ in the usual way.

The quotient of this action, $(\tilde{M} \times \mathbb{C}^2)/\pi_1 M$, is an $SU(2)$ bundle over M. Since all $SU(2)$ bundles over a 3 manifold are trivial (the classifying space $BSU(2)$ is 3-connected), this bundle is a product $M \times SU(2)$, with an inherited connection from the trivial connection on $\tilde{M} \times \mathbb{C}^2$ (defined for this product just as it is defined for M above). These can be checked to be inverses, and thus in fact the correspondence between $R(M)$ and $\mathscr{F}$ is a bijection. □

Next, we establish the equivalents of the $SO(3)$ action on $R(M)$.

Definition. *The Gauge group $\mathscr{G}$ of $M \times SU(2)$ is the group of principal bundle automorphisms of $M \times SU(2)$ (i.e., fiber preserving homeomorphisms which are invariant under the group action).*

By the $SU(2)$ invariance, any such automorphism of $M \times SU(2)$ is exactly determined by where it sends $(M, 1)$. Thus we can identify $\mathscr{G}$ with $C^\infty(M, SU(2))$ when needed. In particular, the exterior derivative of $g \in \mathscr{G}$ is an $\mathfrak{su}(2)$-valued 1-form. Then $\mathscr{G}$ acts on $\mathscr{A}$ by

$$g \cdot A = g^{-1} dg + g^{-1} A g$$

It can be checked that $g \cdot A$ satisfies the same invariance properties as the original connection; in particular, if A is flat, then so is $g \cdot A$. The action of $\mathscr{G}$ on $\mathscr{A}$ corresponds to the action of $SU(2)$ on $R(M)$, which can be checked by looking at how $\mathscr{G}$ affects holonomy. Specifically, we can identify the $\operatorname{stab}(A)$ with the centralizer of the holonomy group in $SU(2)$. We say that a connection is irreducible if $\operatorname{stab}(A) = \{\pm 1\} = Z(SU(2))$. If A is flat, then irreducibility of A is exactly surjectivity of $\operatorname{hol}_A \in R(M)$ because the nontrivial subgroups of $SU(2)$ are abelian. This gives a correspondence also between flat, irreducible connections ($:= \mathscr{F}^*$) and elements of $R^{irr}(M)$. We can form the quotient $\mathscr{R}^* := \mathscr{F}^*/\mathscr{G}$, which corresponds to $\mathcal{R}(M)$. To summarize:

Theorem 17. *There is a correspondence of the following form:*

$$\begin{array}{ccc} R(M) & \leftrightsquigarrow & \mathscr{F} \\ R^{irr}(M) & \leftrightsquigarrow & \mathscr{F}^* \\ \mathcal{R}(M) & \leftrightsquigarrow & \mathscr{R}^* \end{array}$$

As well as $\operatorname{Hom}(\pi_1 M, SU(2))/SO(3) \leftrightsquigarrow \mathscr{R} := \mathscr{F}/\mathscr{G}$.

The spaces $\mathscr{R}$, $\mathscr{R}^*$ live inside of the larger spaces $\mathscr{B} := \mathscr{A}/\mathscr{G}$ and $\mathscr{B}^*$, the subset of $\mathscr{B}$ formed by the orbits of irreducible connections. This latter space is an infinite dimensional Banach manifold, and this is the extra structure which we may use to gain more information about $\mathcal{R}(M)$.

Remark. Recall that in section 3.1, we found out that the transversality of $\mathcal{R}(H_1) \cap \mathcal{R}(H_2)$ was actually a property of Σ itself, not a particular decomposition, and that this gave us a convenient way to discuss degenerate versus nondegenerate homology three spheres. In the spirit of the correspondence theorem, we extend this without any more mention to $\mathscr{R}^*$.

5.1. **Spectral flow.** A 1-form still has a notion of wedge product, even when it has $\mathfrak{su}(2)$ coefficients: we just need to specify how the wedge acts on the coefficients. In this case, it is easy since we can use the lie bracket. For example, in local coordinates we may have

$$\left(\begin{pmatrix} i & 0 \\ 0 & -i \end{pmatrix} dx \right) \wedge \left(\begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} dy \right) = \begin{pmatrix} 0 & 2i \\ 2i & 0 \end{pmatrix} dx \wedge dy$$

Similarly, the differential $d : \Omega^k(M; \mathfrak{su}(2)) \to \Omega^{k+1}(M; \mathfrak{su}(2))$ is defined by ignoring the $\mathfrak{su}(2)$ component of the tensor entirely and acting only on the underlying form.

Definition. *For a connection $A \in \Omega^1(M; \mathfrak{su}(2))$ The covariant derivative is a linear map $\Omega^k(M;\mathfrak{su}(2)) \to \Omega^{k+1}(M;\mathfrak{su}(2))$ defined by*

$$d_A u = du + A \wedge u - u \wedge A$$

Now choose a Riemannian metric on M. Hodge theory and the Hodge star operator $\star$ tells us that the dual operator $d_A^* = -\star d_A \star$ is an adjoint to d_A.

Definition. *For A a connection, define the linear operator*

$$K_A : (\Omega^0 \oplus \Omega^1)(M;\mathfrak{su}(2)) \to (\Omega^0 \oplus \Omega^1)(M;\mathfrak{su}(2))$$

by the block matrix

$$\begin{pmatrix} 0 & d_A^* \\ d_A & \star d_A \end{pmatrix}$$

The association $A \mapsto K_A$ is continuous; this can be checked just by tracing through the definitions.

Theorem 18 (Lemma 2.4 of [10]). *On the L^2 completion of $\Omega^0 \oplus \Omega^1$ to the appropriate Hilbert space, K_A is a(n unbounded) self adjoint Fredholm operator. In particular, it has a discrete, real spectrum $\sigma(K_A) \subset \mathbb{R}$ with finite multiplicity at each eigenvalue.*

The hope would be to assign to each $A \in \mathscr{R}^*$ a number μ which counted the number of points on the spectrum. Since the operators K_A are unbounded, though, the spectrum may be infinite and we cannot count them directly. Instead, we could try to *compare* the spectrums, and find a relative difference of sorts. To this end, we define spectral flow.

Spectral flow has its roots in functional analysis. An elegant, functional analytic definition of spectral flow can be found in [7]. We will here give a different, non-rigorous definition which can be made rigorous. Given a curve $\gamma : [0,1] \to \mathscr{A}$, $t \mapsto A_t$, we get an associated curve of operators K_{A_t} by continuity. This then facilitates a continuous change in the spectrums $\sigma(K_{A_t})$ as well. By specifying a given point $p \in \sigma(K_{A_0})$, we get a curve that links p to some other point in $\sigma(K_{A_1})$.

We can then measure how many of the eigenvalues (with multiplicity) change from negative to positive over the course of this path, and define $\mathrm{sf}(A_0, A_1)$ to be the number of sign changes of the spectrum of K_{A_t} from negative to positive for any generic path $\gamma : t \mapsto A_t$.

Slightly more formally: let $\Gamma = \{(t,y) | y \in \sigma(K_{A_t})\}$ be the graph of the spectrums over t. Since the spectrums are discrete, we can pick a $\delta > 0$ such that there is no $p \in (\sigma(K_{A_0}) \cup \sigma(K_{A_1})) \cap ([-\delta,\delta] \setminus \{0\})$. Then for each $p \in \sigma(K_{A_0})$ (counting multiplicity), the 'continuous path' of spectra gives rise to an actual continuous path starting at p (the restriction of the spectral flow to the specific value p), and we can count the signed intersection number of this curve with the line from $(-\delta, 0)$ to $(1, \delta)$, see figure 10. This intersection number is either $-1, 0$, or 1. Then we can add this intersection number across all points p to get the spectral flow: we label this $\mathrm{sf}(A_0, A_1)$.

We summarize the key facts about spectral flow in the following lemma:

Lemma 12 (Lemma 2.5, 2.6 of [10]). *For two connections A, B, the spectral flow $sf(A,B)$ is always finite. It depends on the choice of curve connecting A and B only up to homotopy rel. endpoints. Mod 8, it is well-defined regardless of curve chosen and is gauge-invariant, and thus descends to a well-defined function on $\mathscr{R}$.*

Finally, we can define, for any $\alpha \in \mathscr{R}^*$, $\mu(\alpha) = \mathrm{sf}(\theta, \alpha)$ for θ the trivial connection, giving a function $\mathscr{R}^* \to \mathbb{Z}/8$. If $\mathscr{R}^*$ is non-degenerate, then there are only finitely many points in $\mathscr{R}^*$, so

Definition (The Casson invariant, again). *If $\mathscr{R}^*$ is non-degenerate, then define*

$$\lambda(\Sigma) = \frac{1}{2} \sum_{\alpha \in \mathcal{R}(\Sigma)} (-1)^{\mu(\alpha)}$$

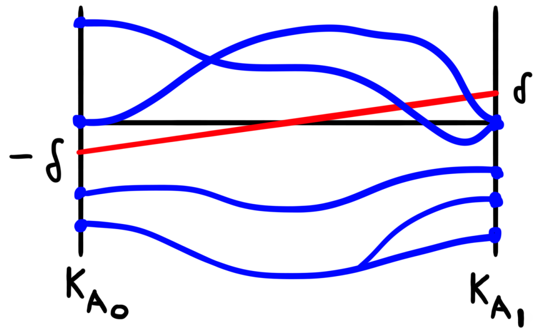

FIGURE 10. $\mathrm{sf}(A_0, A_1) = -2$

That this definition is well-defined (not depending on the various choices we have made, such as that of a Riemmanian metric) and equal to the original Casson Invariant is the subject of the paper [10]. The key insight is that the above definitions all extend to the structure of Heegaard splittings and their handlebodies, and that the intersection of the handlebodies gives rise to another Fredholm operator which can be compared to the spectral flow.

5.2. **Morse theory.** When $\mathcal{R}(\Sigma)$ is degenerate, there is no intersection we can perturb to make things finite again. We could instead try to perturb the function K_A to some new function K'_A. We would still need to know which points to sum this new function over; we can solve this issue by having K'_A be the Hessian of some other operator, and perturbing this other operator to have finitely many critical points. Then we can let K'_A be the Hessian of the perturbed operator and sum over the critical points.

The curvature of a connection $A \in \Omega^1(M; \mathfrak{su}(2))$ is the 2-form $F_A = dA + A \wedge A$. As an example, if

$$A = \begin{pmatrix} i & 0 \\ 0 & -i \end{pmatrix} dx + \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} dy$$

Then

$$F_A = dA + A \wedge A = A \wedge A = \left[\begin{pmatrix} i & 0 \\ 0 & -i \end{pmatrix}, \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} \right] dx \wedge dy + \left[\begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}, \begin{pmatrix} i & 0 \\ 0 & -i \end{pmatrix} \right] dy \wedge dx$$
$$= \begin{pmatrix} 0 & 4i \\ 4i & 0 \end{pmatrix} dx \wedge dy$$

Theorem 19. *A connection is flat iff $F_A = 0$.*

Let θ be the 'trivial' connection in $\mathscr{A}$ induced by the product structure, and pick a path $[0,1] \to \mathscr{A}$ which begins at θ and ends at a connection A. This path determines a connection on the trivial $SU(2)$ bundle over $[0,1] \times M$ (over which connections are defined exactly analogously as over M itself, and all the same theory carries over). Then, calling this connection A_γ, we define

$$\mathrm{cs}(A) := \int_{[0,1] \times M} \mathrm{tr}(F_{A_\gamma} \wedge F_{A_\gamma})$$

Where the 'trace' is the $\mathfrak{su}(2)$-lie algebra pairing (ignoring the differential form component), i.e., $\mathrm{tr}((u \otimes \omega) \wedge (v \otimes \sigma)) = -\frac{1}{2}\mathrm{tr}(uv)\omega \wedge \sigma$. It turns out that this function is independent of the path

γ; it is also true that for $g \in \mathscr{G}$, $\mathrm{cs}(g \cdot A) - \mathrm{cs}(A) \in \mathbb{Z}$, so that the function descends to a map $\mathrm{cs} : \mathscr{B} \to \mathbb{R}/\mathbb{Z}$.

Remember that $\mathscr{B}^*$ is an infinite dimensional Banach manifold, and so there is a notion of gradient of functions. It turns out that the gradient $\nabla\,\mathrm{cs}$ can be computed to be

$$(\nabla\,\mathrm{cs})(A) = -\frac{1}{4\pi^2} F_A$$

Thus, the 'critical points' of cs on $\mathscr{B}^*$ are exactly the flat, irreducible connections, i.e., $\mathscr{R}^*$! Since constant multiples don't matter when computing kernels, index, or spectrum, we can ignore the $-\frac{1}{4\pi^2}$ factor. The hessian of cs is in turn $\star d_A : \ker d_A^* \to \ker d_A^*$. If A is flat and irreducible, then the kernel of K_A is the same as the kernel of $\star d_A$, so that whenever one is nondegenerate, so is the other. This is why it makes sense to define the Casson invariant using K_A for nondegenerate spaces. For degenerate spaces it now suffices to take a perturbation of the Chern-Simons functional which has finitely many critical points. One can then define the adapted K'_A on these critical points and continue as before. See [8] for an in-depth explanation.

5.3. **"Applications": Instanton Floer.** Instanton Floer homology extends the gauge theoretical definition of the Casson invariant up to its logical conclusion: the spectral flow gave a number valid not just mod 2, but modulo 8. Assuming that Σ is nondegenerate, so that the set of points $x \in \mathcal{R}(\Sigma)$ is finite, we can assign to each $k \in \mathbb{Z}/8$ the free abelian group generated by the set $\{x \in \mathcal{R}(\Sigma) | \mu(x) \cong k \mod 8\}$.

Denote each of these abelian groups as $IC_k(\Sigma)$, with $IC(\Sigma) = \bigoplus_{k=0}^{8} IC_k$ the a graded abelian group. There exists boundary operators $\partial : IC_k \to IC_{k-1}$ (with $\partial^2 = 0$, of course) making this into a chain complex. The associated homology, $I_*(\Sigma)$, is the Instanton Floer Homology of Σ. The boundary maps are somewhat complicated to define, and even more complicated to compute. So what does Instanton Floer have to offer that the Casson invariant doesn't? For one, it retains more of the original information about the number of points in $\mathcal{R}(\Sigma)$. If the boundary operators all vanish - which frequently occurs - then the homology contains perfect information about $\mathcal{R}(\Sigma)$. Furthermore, IC is functorial with respect to cobordisms. This makes it a powerful way of examining the theoretical structure of homology spheres. Just as with the Casson invariant, for degenerate manifolds there is a way of perturbing the functions so that everything becomes finite.

Today, Instanton Floer is used extensively. Its applications are wide ranging; it strengthens of the Casson invariant in almost every area. For example, Kronheimer and Mrowka in their proof of property P in [5], and Daemi et al. use it in their progress on cosmetic surgery in [1]. Notice that we used the Casson invariant to make more limited progress in these areas. For more information about Instanton Floer homology, see [2] and [8].

6. Acknowledgements

This paper would not be possible without my advisor, Zoltán Szabó; his help was invaluable at all stages of the process, and I cannot thank him enough. I would also like to thank Ollie Thakar for his ever valuable back-of-the-napkin explanations of gauge theory, as well as all of my friends who incessantly supported me throughout the semester. This paper was written in fulfillment of the junior independent work requirement at Princeton University.

References

[1] Aliakbar Daemi, Mike Miller Eismeier, and Tye Lidman. Filtered instanton homology and cosmetic surgery, 2025.

[2] Andreas Floer. An instanton-invariant for 3-manifolds. *Communications in Mathematical Physics*, 118(2):215–240, 1988.

[3] Christopher Herald, Paul Kirk, and Charles Livingston. Metabelian representations, twisted alexander polynomials, knot slicing, and mutation. *Mathematische Zeitschrift*, 265:925–949, 08 2010.
[4] Robion Kirby. A calculus for framed links in s3. *Inventiones mathematicae*, 45:35–56, 1978.
[5] P.B. Kronheimer and T.S. Mrowka. Witten's conjecture and property p. *Geometry & Topology*, 8:295–310, 2004.
[6] Christine Lescop. *Un calcul élémentaire de l'invariant de Casson des sphères d'homologie entière fibrées de Seifert à trois fibres exceptionnelles*, volume 38. L'Enseignement mathématique, 1992.
[7] John Phillips. Self-adjoint fredholm operators and spectral flow. *Canadian Mathematical Bulletin*, 39(4):460–467, 1996.
[8] Nikolai Saveliev. *Invariants for Homology 3-Spheres*, volume 140. 01 2002.
[9] Nikolai Saveliev. *An Introduction to the Casson Invariant*. De Gruyter, Berlin, Boston, 2012.
[10] Clifford Henry Taubes. Casson's invariant and gauge theory. *J. Diff. Geom.*, 31(2):547–599, 1990.
[11] Konstantinos Varvarezos. Heegaard floer homology and chirally cosmetic surgeries. *Communications in Analysis and Geometry*, 2021.
[12] Kevin Walker. *An Extension of Casson's Invariant. (AM-126)*. Princeton University Press, 1992.