KNOTS IN S^3 AND BORDERED HEEGAARD FLOER HOMOLOGY

ATANAS ATANASOV

Contents

Acknowledgements

I would like to thank my advisor Prof. Robert Lipshitz. If he had not been as devoted in his effort to teach me, this would not have been possible.

1. Introduction

Two of the key tools in 3-dimensional topology are hyperbolic geometry and gauge theory. At present, almost no relations are known between the two topics. In this paper, we start to explore whether topological field theory properties of one gauge-theoretic invariant, Heegaard Floer homology, imply a relationship with convergence of geometric structures.

This paper is split in two parts. Firstly, Sections 2 and 3 could be used as a brief introduction to bordered Heegaard Floer homology, and surgery on knots. Secondly, Section 4 describes an invariant for knots in S^3, and carries out a computation which is central to understanding it any further. We proceed to describe briefly the content of the following three sections.

Section 2. All of the material which appears here is classical, and of broad interest in the realm of low-dimensional topology. The first three subsections build the machinery of Heegaard diagrams which is used to provide a combinatorial description of 3-manifolds. The latter three subsections go carefully through the construction of link surgery.

Section 3. Figuratively speaking, this part of our exposition could be treated as a crash course in (a specific case of) bordered Heegaard Floer homology (see [7] for a detailed account of the theory). For our purposes, it is important that a bird's-eve view of the Heegaard Floer theories is presented (Section 3.1), and the relation between the non-bordered and bordered cases becomes clear (Theorem 3.2). The following two subsections aim to build the D type complex in the simple case of torus boundary. Finally, we work through an example (Section 3.4) in order to clarify the constructions we made earlier.

Section 4. Our aim in this section is to describe the constriction of an invariant for knots in S^3. This idea is inspired by W. Thurston's Hyperbolic Surgery Theorem (quoted as 4.1), which suggests a sense of stabilization for manifolds obtained by large surgeries on a fixed knot K. Section 4.2 is devoted to transferring this contraption to the realm of bordered Heegaard Floer homology. In the last subsection, we carry out an explicit computation central to any eventual further development.

1.1. Notation and conventions.

(1) The n-dimensional closed disk will be denoted by D^n.
(2) The 2-torus will be denoted by $\mathbb{T} = S^1 \times S^1$.
(3) Unless otherwise stated, all vector spaces will be over the field of two elements $\mathbb{F}_2$.
(4) Unless otherwise stated, all manifolds are assumed to be smooth and oriented.

2. Preliminaries

2.1. **Bordered manifolds.** Bordered manifolds are a refinement of manifolds with boundary. They carry an extra bit of information which allows us to work with the boundary without ambiguity. We give a motivating example before proceeding to the formal definition. Let M_1 and M_2 be manifolds with diffeomorphic boundaries. The information such a diffeomorphism exists is insufficient to glue them, hence we require an explicit such map $f\colon \partial M_1 \to \partial M_2$. Bordered manifolds attempt to fix this issue by carrying parametrizations of their boundaries. Let n be a positive integer, and suppose N is a $(n-1)$-manifold. A *bordered manifold with boundary* N is a pair (M, f), where M is an n-manifold with boundary, and $f\colon N \to \partial M$ is a diffeomorphism. We will often refer to such an object by M alone, and if necessary, we may provide the parametrization f separately. At first sight the choice of an arbitrary boundary manifold N might seem uninforming, but in most cases we will have a set of models to choose from. For example, we will be primarily interested in compact oriented bordered 3-manifolds in which case the boundary is a compact oriented surface, and diffeomorphism classes of these are classified by the genus. Returning to our motivating example, if (M_1, f_1) and (M_2, f_2) are bordered manifolds with boundaries N and $-N$, then there is a natural way to glue then, namely via $f_2 \circ f_1^{-1}\colon \partial M_1 \to \partial M_2$. Since we are working with oriented manifolds, we have to take this into account when performing a gluing in order to obtain another oriented manifold. This is the reason we asked for two copies of N with opposite orientations in the previous construction.

Our discussion leads us to define the *bordism categories* $\mathbf{Bord}_n$ for $n \geq 1$. The objects of $\mathbf{Bord}_n$ are $(n-1)$-dimensional manifolds. The morphisms between two objects N^- and N^+ are 5-tuples $(M, \partial M^-, \partial M^+, f^-, f^+)$ such that $\partial M = \partial M^- \sqcup \partial M^+$ is a decomposition of manifolds, and $f^-\colon N^- \to \partial M^-$, $f^+\colon N^+ \to \partial M^+$ are diffeomorphisms. Where possible, we will refer to such a morphism simply by M. The parametrizations of each piece of the boundary facilitate compositions in the obvious manner.

2.2. **Heegaard diagrams.** As we already mentioned, 3-manifolds will play an important role in our discussion. The aim of this section is to develop *Heegaard diagrams* which serve as combinatorial descriptions for a subclass of these objects. We will first focus on closed oriented 3-manifolds, and later expand to the bordered case.

Definition 2.1. A *Heegaard diagram* $\mathcal{H}$ is a triple $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$ consisting of

- a closed orientable surface Σ of some genus g,
- two g-tuples $\boldsymbol{\alpha} = \{\alpha_1, \dots, \alpha_g\}$ and $\boldsymbol{\beta} = \{\beta_1, \dots, \beta_g\}$ of pairwise-disjoint loops in Σ.

We further require the α_i and β_j to intersect transversally, and $\Sigma \setminus \boldsymbol{\alpha}$, $\Sigma \setminus \boldsymbol{\beta}$ to be connected. The integer g is called the *genus of* $\mathcal{H}$. A *pointed Heegaard diagram* is a quadruple $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, z)$ where the first three components satisfy the above conditions, and z is a point in $\Sigma \setminus (\boldsymbol{\alpha} \cup \boldsymbol{\beta})$. Note the abuse of notation here – $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ refer both to the set of loops and the union of their images in Σ.

The condition that $\Sigma \setminus \boldsymbol{\alpha}$ is connected is equivalent to the homological linear independence of $[\alpha_i] \in H_1(\Sigma)$; this could be inferred from the Mayer-Vietoris sequence of the decomposition $\Sigma = (\Sigma \setminus \boldsymbol{\alpha}) \cup \boldsymbol{\alpha}$. An analogous statement could be made regarding $\Sigma \setminus \boldsymbol{\beta}$ and $[\beta_i] \in H_1(\Sigma)$.

There is a canonical way to construct a 3-manifold $M_{\mathcal{H}}$ from a Heegaard diagram $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$. There are two basic observations which guide the formal definition. Firstly, the idea is to "fill" the orientable surface Σ on both of its "sides"; each of these fills amounts to constructing a 3-manifold whose boundary is Σ. Secondly, such a 3-manifold is uniquely determined by a set of curves which bound disks in it.

Let us focus on the $\boldsymbol{\alpha}$-curves for the sake of specificity. There are several ways to execute our last remark; we proceed to outline one such procedure. Let us start by thickening Σ to a manifold M_Σ whose boundary is diffeomorphic to two copies of Σ. Pick one of these, in other words, that is an embedding $\Sigma \hookrightarrow \partial M_\Sigma$. We can now identify Σ with its image, and respectively the α_i with their images in ∂M_Σ. It is always possible to pick tubular neighbourhoods of these curves in Σ, and attach thickened disks along these neighbourhoods. The boundary of the resulting manifold is one unmodified copy of Σ and a 2-sphere on the side we performed the handle attachments. Filling in the latter with a copy of D^3 yields the necessary manifold $M_{\boldsymbol{\alpha}}$ – with boundary Σ such that all α_i bound disks in it. An analogous construction with the $\boldsymbol{\beta}$-curves yields a manifold $M_{\boldsymbol{\beta}}$. Gluing $M_{\boldsymbol{\alpha}}$ and $M_{\boldsymbol{\beta}}$ along their boundaries is unambiguous, and gives the closed 3-manifold $M_{\mathcal{H}}$ we originally mentioned.

Here is another way to carry out the previous construction. The main observation is the group of diffeomorphisms of a surface Σ of genus g acts transitively on the set of g-tuples of homologically linearly independent curves. In other words, for any two sets of curves such as $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$, there is a diffeomorphism mapping one to the other. We can make a specific choice of a manifold M_0 with boundary Σ, and a set of g curves $\boldsymbol{\gamma}_0$ in it bounding disks such as the ones shown on Figure 1. Choose two diffeomorphisms $f_{\boldsymbol{\alpha}}, f_{\boldsymbol{\beta}}\colon \Sigma \to \Sigma$ satisfying $f_{\boldsymbol{\alpha}}(\boldsymbol{\gamma}_0) = \boldsymbol{\alpha}$, $f_{\boldsymbol{\beta}}(\boldsymbol{\gamma}_0) = \boldsymbol{\beta}$. The manifold $M_0 \cup_{f_{\boldsymbol{\beta}}^{-1} \circ f_{\boldsymbol{\alpha}}} M_0$, obtained by gluing two copies of M_0 via the map $f_{\boldsymbol{\beta}}^{-1} \circ f_{\boldsymbol{\alpha}}$, satisfies analogous properties as $M_{\mathcal{H}}$ in the previous paragraph.

So far we have seen how to construct a closed 3-manifold $M_{\mathcal{H}}$ corresponding to a Heegaard diagram $\mathcal{H}$. It is only natural to ask whether all such manifolds could be obtained in this manner. One way to reverse the process by using Morse theory (see [9], [11]). Let M be a closed 3-manifold, and pick a self-indexing Morse function f with a unique index points of index 0 and 3 (such a

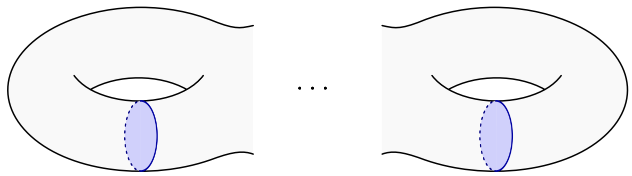

Figure 1: The manifold M_0 with the curves $\boldsymbol{\gamma}_0$ on the boundary

choice is always possible; see [11, Corollary 2.36]). We take $\Sigma = f^{-1}(3/2)$, $\boldsymbol{\alpha}$ the intersection of the ascending disks from the index 1 critical points with Σ, and $\boldsymbol{\beta}$ the intersection of the descending disks from the index 2 critical points with Σ. The Heegaard diagram $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$ yields a manifold $M_{\mathcal{H}}$ diffeomorphic to M (full details are given in [6, Section 2.2]). The Morse theory interpretation we presented is also useful to answer another relevant question: when do two Heegaard diagrams correspond to diffeomorphic 3-manifolds? It is not hard to see the following do not change the resulting manifold.

Isotopies: replace some α_i by an isotopic curve α_i' through isotopies which leave α_j fixed for $j \neq i$; or, analogously for the $\boldsymbol{\beta}$-curves;

Handleslides: replace some α_i by α_i' for which there exists α_j such that $\alpha_i \cup \alpha_i' \cup \alpha_j$ bound a pair of pants disjoint from the remaining α_k where $k \neq i, j$; or, analogously for the $\boldsymbol{\beta}$-curves;

(De)stabilizations: a stabilization replaces $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$ with $(\Sigma', \boldsymbol{\alpha}', \boldsymbol{\beta}')$, where $\Sigma' = \Sigma \# \mathbb{T}$, $\boldsymbol{\alpha}' = \boldsymbol{\alpha} \cup \{\alpha_{g+1}\}$, $\boldsymbol{\beta}' = \boldsymbol{\beta} \cup \{\beta_{g+1}\}$ such that $\alpha_{g+1}, \beta_{g+1}$ are a pair of curves in $\mathbb{T}$ meeting transversally in a single point; a destabilization is the reverse procedure.

In fact, these three moves are sufficient (see [17] for proof).

Proposition 2.2. *Two Heegaard diagrams correspond to diffeomorphic 3-manifolds if and only if they are related by a sequence of isotopies, handleslides, and (de)stabilizations.*

Figure 2 presents two examples of genus 1 Heegaard diagrams. It is not hard to see that $\mathcal{H}_1$ corresponds to the manifold S^3. This is the usual decomposition of S^3 as the union of two solid tori – a tubular neighbourhood of an unknot in S^3 and its exterior. The diagram $\mathcal{H}_2$ stands for $S^1 \times S^2$ seen as two copies of a solid torus glued along the identity map between their boundaries.

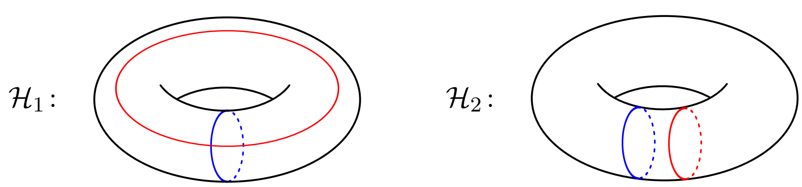

Figure 2: Two Heegaard diagrams: the $\boldsymbol{\alpha}$- and $\boldsymbol{\beta}$-curves are shown respectively in red and blue.

Remark 2.3. The discussion of this section can be carried out using PL (piecewise-linear) 3-manifolds as in [15], [16], and [4]. It turns out the two points of view are equivalent though this is specific to 3-manifolds (see [10], [2]).

2.3. **Bordered Heegaard diagrams.** We have developed two significant concepts so far – bordered manifolds, and Heegaard diagrams. The majority of Section 2.2 was devoted to describing how Heegaard diagrams determine closed 3-manifolds. It is only logical to ask whether we could augment Definition 2.1 to yield *bordered Heegaard diagrams*, objects which will correspond to bordered oriented compact 3-manifolds. In what follows, we recall without proof the relevant concepts as in [7, Section 4.1], [6, Section 2.2]. Some details may be omitted since the development closely resembles that of Section 2.2.

Definition 2.4. A *bordered Heegaard diagram* $\mathcal{H}$ is a triple $(\overline{\Sigma}, \overline{\boldsymbol{\alpha}}, \boldsymbol{\beta})$ consisting of

- a compact, orientable surface $\overline{\Sigma}$ of some genus g with one boundary component,
- a g-tuple $\boldsymbol{\beta} = \{\beta_1, \dots, \beta_g\}$ of pairwise-disjoint loops in the interior of $\overline{\Sigma}$,
- a $(g+k)$-tuple $\overline{\boldsymbol{\alpha}}$ of pairwise-disjoint curves in $\overline{\Sigma}$, split into a $(g-k)$-tuple $\overline{\boldsymbol{\alpha}}^c = \{\alpha_1^c, \dots, \alpha_{g-k}^c\}$ of loops in the interior of $\overline{\Sigma}$, and a $2k$-tuple $\overline{\boldsymbol{\alpha}}^a = \{\alpha_1^a, \dots, \alpha_{2k}^a\}$ of arcs in $\overline{\Sigma}$ with boundary in $\partial\overline{\Sigma}$ (and transverse to $\partial\overline{\Sigma}$).

We also require the $\alpha_i^{a,c}$ and β_j to intersect transversally, and $\overline{\Sigma} \setminus \overline{\boldsymbol{\alpha}}$, $\overline{\Sigma} \setminus \boldsymbol{\beta}$ to be connected. The integer g is called the *genus of* $\mathcal{H}$. A *pointed bordered Heegaard diagram* is a quadruple $(\overline{\Sigma}, \overline{\boldsymbol{\alpha}}, \boldsymbol{\beta}, z)$ where the first three components satisfy the above conditions, and z is a point in $\overline{\Sigma} \setminus (\overline{\boldsymbol{\alpha}} \cup \boldsymbol{\beta})$.

Again, the condition $\overline{\Sigma} \setminus \boldsymbol{\beta}$ is connected is equivalent to the homological linear independence of $[\beta_i] \in H_1(\overline{\Sigma})$. Analogously, $\overline{\Sigma} \setminus \overline{\boldsymbol{\alpha}}$ is connected if and only if $[\alpha_i^{c,a}] \in H_1(\overline{\Sigma}, \partial\overline{\Sigma})$ are homologically linearly independent.

Let us fix a bordered Heegaard diagram $\mathcal{H}$ and discuss how to construct a bordered 3-manifold $M_{\mathcal{H}}$. The double of $\mathcal{H}$, denoted $2\mathcal{H}$, is a Heegaard diagram obtained by gluing two copies of $\mathcal{H}$ along their common boundary. There is a natural $\mathbb{Z}/2$-action on $2\mathcal{H}$ flipping from one copy of $\mathcal{H}$ to the other. This ascends to a $\mathbb{Z}/2$-action on the manifold $M_{2\mathcal{H}}$ whose fixed set is a null-homologous surface F splitting $M_{2\mathcal{H}}$ into two parts. We define $M_{\mathcal{H}}$ to be a fundamental domain for this action with boundary F.

There is an alternative interpretation of F which demonstrates that $M_{\mathcal{H}}$ is not only a manifold with boundary, but is also bordered. Keeping track of Euler characteristics, one can deduce that $\overline{\Sigma} \setminus \overline{\boldsymbol{\alpha}}$ is a planar surface with $2(g-k)+1$ boundary components, only one of which meets $\partial\overline{\Sigma}$. Let N be the union of a collar neighbourhood of $\partial\overline{\Sigma}$ and a tubular neighbourhood of $\overline{\boldsymbol{\alpha}}^a$. Its boundary ∂N consists of two connected components, one of which contains $\partial\overline{\Sigma}$. Filling each of these components to $N \cup D^2 \cup D^2$, we obtain a surface which could be identified with F as above. There is a canonical (up to isotopy) identification of the obtained surface with a model surface of genus k, hence $M_{\mathcal{H}}$ is bordered.

As in the non-bordered case in the previous section, this construction is reversible – every bordered 3-manifold M can be expressed as the induced manifold $M_{\mathcal{H}}$ for some bordered Heegaard diagram $\mathcal{H}$. Once again, filling the complete details of the argument would require some analysis and Morse theory, so we will be content with only an outline here.

Let M be a 3-manifold with a single boundary component, and consider a Riemannian metric $\widetilde{g}$ on M, and a self-indexing Morse function $f\colon M \to \mathbb{R}$.

Definition 2.5. We say that $\widetilde{g}$ and f are *boundary compatible* with M if

- the boundary of M is geodesic,
- $\nabla f|_{\partial M}$ is tangent to ∂M,
- the unique index 0 and 2 critical points of $f|_{\partial M}$ are respectively the unique index 0 and 3 index critical points of f on M,
- the index 1 critical points of $f|_{\partial M}$ are also index 1 critical points of f on M.

We claim for every Riemannian 3-manifold with a single boundary component $(M, \widetilde{g})$, there exists a boundary compatible f. To find one, start with f with the desired properties on ∂M, then extend to a collar neighbourhood of ∂M. We proceed to extend arbitrarily to the interior of M, and then perturb to make $(\widetilde{g}, f)$ a Morse-Smale pair, and f self-indexing. Then, we can take

$\overline{\Sigma} = f^{-1}(3/2)$, $\overline{\boldsymbol{\alpha}}$ the intersection of the ascending discs of the index 1 critical points of f with $\overline{\Sigma}$, and $\boldsymbol{\beta}$ the intersection of the descending disks of the index 2 critical points with $\overline{\Sigma}$. The bordered Heegaard diagram $\mathcal{H} = (\overline{\Sigma}, \overline{\boldsymbol{\alpha}}, \boldsymbol{\beta})$ corresponds to M. This is analogous to the construction from the previous section. The only difference is we had to impose several technical conditions on f which allow us to carry through despite the presence of a boundary component.

The previous paragraph disregarded the presence of a bordered structure. We note the given procedure could extend a compatible f defined on the boundary ∂M to one on the entire manifold M. This allows us to construct a bordered Heegaard diagram for any bordered 3-manifold. Once again, there is a result analogous to Proposition 2.2.

Proposition 2.6. *Two bordered Heegaard diagrams correspond to isomorphic bordered 3-manifolds if and only if they are related by a sequence of isotopies, handleslides, and (de)stabilizations. The only difference is that handleslides for $\overline{\boldsymbol{\alpha}}$ may only occur for $\overline{\boldsymbol{\alpha}}^a$ over $\overline{\boldsymbol{\alpha}}^c$.*

We refer to [7, Section 4.1] for a proof.

Remark 2.7. In many of the discussions which follow, one is typically interested in bordered pointed Heegaard diagrams. We will provide pointed examples until the meaning of the extra marking becomes clear later.

We have already seen the double of a bordered Heegaard diagram denoted $2\mathcal{H}$. This is a specific case of a more general construction: consider two bordered diagrams $\mathcal{H}_1 = (\overline{\Sigma}_1, \overline{\boldsymbol{\alpha}}_1, \boldsymbol{\beta}_1)$, $\mathcal{H}_2 = (\overline{\Sigma}_2, \overline{\boldsymbol{\alpha}}_2, \boldsymbol{\beta}_2)$, and an orientation reversing diffeomorphism $f\colon \partial\overline{\Sigma}_1 \to \partial\overline{\Sigma_2}$, which pairs arcs from $\overline{\boldsymbol{\alpha}}_1^a$ and $\overline{\boldsymbol{\alpha}}_2^a$ via their endpoints. If $\mathcal{H}_1$ and $\mathcal{H}_2$ are pointed, we further require that f maps one marking to the other. Under these assumptions, one can form the diagram $\mathcal{H} = \mathcal{H}_1 \cup_f \mathcal{H}_2$ by gluing $\mathcal{H}_1$ and $\mathcal{H}_2$ along f. Going through the constructions of the associated bordered 3-manifolds, f naturally induces a diffeomorphism $\widetilde{f}\colon \partial M_{\mathcal{H}_1} \to \partial M_{\mathcal{H}_2}$ such that $M_{\mathcal{H}} = M_{\mathcal{H}_1} \cup_{\widetilde{f}} M_{\mathcal{H}_2}$.

To some extent, we can also interpret connect sums of bordered 3-manifolds in terms of their Heegaard diagrams. Before doing so however, we make several general remarks. In the construction of the connect sum of two manifolds, one deletes an open ball from each, and then identifies the remaining pieces via an orientation preserving diffeomorphism of their boundaries. It is a subtle but important point that for connected oriented manifolds without boundary the isomorphism type of the end result does not depend on the open balls we choose, hence the unambiguous notation $-\#-$. A lengthly discussion of this topic can be found in [4, Chapter 3].

Even if our manifolds are connected and oriented, allowing boundary changes matters slightly. We have the choice of either removing coordinate balls from the interior of the manifolds, or half-balls which touch their boundaries (mixing the two options does not yield a sensible construction). In the former case, referred to as *connect sum*, the isomorphism type does not depend on the positions of the balls we choose. In the latter case, referred to as *boundary connect sum*, the isomorphism type depends only on the boundary components touching the half-balls we select. For example, suppose M_1 and M_2 are two manifolds each of which has a single boundary component. The connect sum yields a manifold with two boundary components, whereas the boundary connect sum yields a manifold with a single boundary component which is well-defined up to isomorphism.

This analogy does not extend as freely to the bordered case. Let $\mathcal{H}_1$ and $\mathcal{H}_2$ be bordered Heegaard diagrams corresponding to the manifolds M_1 and M_2 respectively. Each of these has a single boundary component. We are interested in taking a boundary connect sum, so the result has a single boundary component too. The isomorphism type of the result however depends on the position of the half-balls we choose. In a similar fashion, one can take an appropriate boundary connect sum of the Heegaard diagrams $\mathcal{H}_1$ and $\mathcal{H}_2$ by removing half-disks near the boundary not touching any of the $\overline{\boldsymbol{\alpha}}_{1,2}$, $\boldsymbol{\beta}_{1,2}$, or the markings $z_{1,2}$ if present. It is not hard to convince oneself that the isomorphism class of $\mathcal{H}_1\#\mathcal{H}_2$ depends only on the position of the half-balls relative to the arc endpoints $\partial\overline{\boldsymbol{\alpha}}_{1,2}$. It is possible to interpret $M_{\mathcal{H}_1\#\mathcal{H}_2}$ as a certain boundary connect sum $M_1\#M_2$, and vice versa, but we will not go into more details with this discussion.

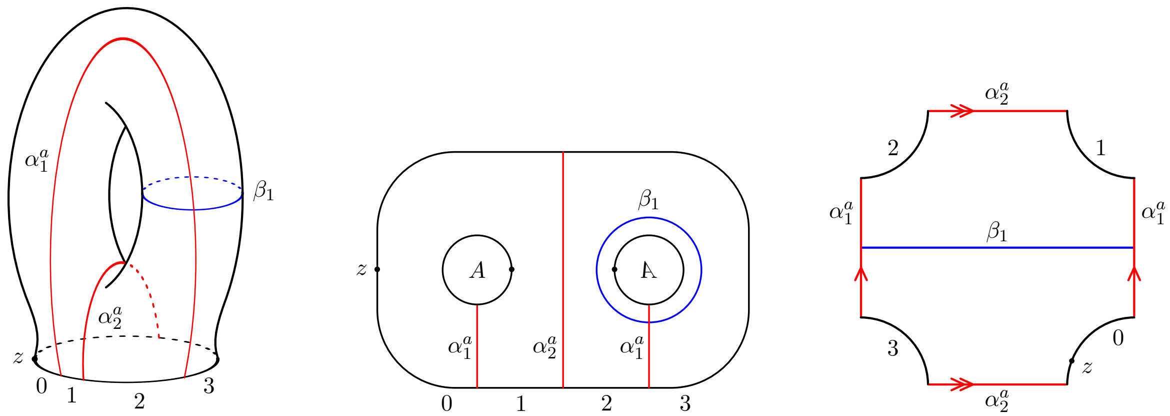

Presentation 1 Presentation 2 Presentation 3

Figure 3: Three ways to visualize the same pointed bordered Heegaard diagram.

Before closing this section, we provide a simple example of the Heegaard diagram for a genus one handlebody (topologically a solid torus). The first presentation in Figure 3 matches our construction the closest. We have shown a punctured genus one surface $\overline{\Sigma}$, two $\overline{\boldsymbol{\alpha}}$-curves α_1^a, α_2^a in red, and one $\boldsymbol{\beta}$-curve β_1 in blue. The marking on $\partial\overline{\Sigma}$ has been denoted as the point z. For reasons which will become clear, we are also interested in the connected components of $\overline{\Sigma}\setminus\overline{\boldsymbol{\alpha}}$ which have been marked 0 to 3 in the counterclockwise direction, 0 being the one containing z.

Presentation 2 could be read as a recipe for obtaining the former one. The two regions marked A and $\reflectbox{$A$}$ instruct us to remove their interiors, and glue the boundaries (alternatively glue a tube between them). The preferred gluing orientation is indicated by both letters inside, and the solid dot marking on the boundary. It is important to note the arc α_1^a is shown in two fragments but these are joined after the prescribed identification.

Finally, presentation 3 is obtained from the first one by cutting along the arcs α_1^a and α_2^a. Arrows along the cut regions serve as directions to reconstruct the original object, though these may be omitted for the sake of saving space. If the bordered Heegaard diagram in question has $2k$ arcs, then this presentation would be a polygon of $4k$ edges with chipped vertices. In practice, this is only feasible for $k = 1$ as the number of gluings becomes too large to imagine easily. In conclusion, the third presentation is preferred for the case $k = 1$, and the second for $k > 1$.

2.4. **Knots and links.** The aim of this section, among other things, is to provide rigorous answers to the following two questions:

- What is a knot?
- What is an equivalence of knots?

Without further ado, we address the first one.

Definition 2.8. A *knot* K in a manifold M is an embedded submanifold diffeomorphic to S^1. A *link with n components* L in M is an embedded submanifold diffeomorphic to $\bigsqcup_{i=1}^n S^1$.

We will be working with knots in 2- and 3-manifolds, though the definition makes sense in general. It is often more useful to think of knots or links as maps into M, hence the following equivalent formulation.

Definition 2.9. A *knot* K in a manifold M is an embedding $K\colon S^1 \hookrightarrow M$. Similarly, a *link with n components* L is an embedding $L\colon \bigsqcup_{i=1}^n S^1 \hookrightarrow M$.

At this point, we can answer our second question. We will specialize to knots though all formulations carry through using links as well.

Definition 2.10. Let $K, K'\colon S^1 \hookrightarrow M$ be two knots in M. In increasing strength, we present three notions of knot isomorphisms.

- **(Map) equivalence:** There exists a diffeomorphism $h\colon M \to M$ such that $h \circ K = K'$.
- **Oriented equivalence:** Same as above, but we require h to be orientation preserving.
- **(Ambient) isotopy:** There exists a homotopy h_t for $0 \le t \le 1$ such that all $h_t\colon M \to M$ are diffeomorphisms, h_0 is the identity map on M, and $h_1 = h$ satisfies $h \circ K = K'$.

In practice, we will most often be interested in knots up to ambient isotopy.

2.5. **Facts about tori and solid tori.** In this section, we recall (without proof) several basic results which lead to the rigorous formulation of knot surgery as presented in Section 2.6. Our exposition follows several sections from [16, Chapter 2].

We start by answering two questions:

- What are the isotopy types of knots in $\mathbb{T}$?
- What is the mapping class group of $\mathbb{T}$?

Recall the torus is $\mathbb{T} = S^1 \times S^1$, so $\pi_1(\mathbb{T}) = \pi_1(S^1) \times \pi_1(S^1) \cong \mathbb{Z}^2$. Without loss of generality, we identify these two groups, so we will denote elements of $\pi_1(\mathbb{T})$ by pairs of integers (a, b). The former question is handled by the following result (see [16, Chapter 2.C] for a proof).

Theorem 2.11. *A map $S^1 \to \mathbb{T}$ is homotopic to a knot (an embedding) if and only if its class $(a, b) \in \pi_1(\mathbb{T})$ satisfies $a = b = 0$, or $\gcd(a, b) = 1$. Conversely, two knots $K, K'\colon S^1 \hookrightarrow \mathbb{T}$ are isotopic if and only if $[K] = \pm[K']$ as elements of $\pi_1(\mathbb{T})$.*

Before answering the second question, we have to clarify one of the terms.

Definition 2.12. For any smooth manifold M, its self-diffeomorphisms form a group $\mathrm{Aut}(M)$ with binary operation composition of maps. The identity element is the identity map $\mathrm{id}_M\colon M \to M$. The diffeomorphisms smoothly isotopic (homotopic through diffeomorphisms) to id_M form a normal subgroup $\mathrm{Aut}_0(M)$. The *mapping class group of* M is the quotient $\mathrm{Mod}(M) = \mathrm{Aut}(M)/\mathrm{Aut}_0(M)$.

Remark 2.13. The mapping class group can be constructed in the category of topological spaces and continuous maps, but we will focus our attention to the smooth case. One can also work with pointed topological spaces or pointed smooth manifolds, and define the *based mapping class group.*

There is an action of $\mathrm{Aut}(M)$ on M by evaluation. If we assume that M is connected, this action becomes transitive. In order to avoid some technical details, we will make this assumption for the remainder of this section. For the sake of mathematical correctness, we fix a basepoint in $x_0 \in M$, and denote $\pi_1(M) = \pi_1(M, x_0)$. Any diffeomorphism $f\colon M \to M$ induces an isomorphism $f_*\colon \pi_1(M) \to \pi_1(M)$ which is well-defined up to conjugation, that is, $[f_*]$ is an element of the group of outer automorphisms $\mathrm{Out}(\pi_1(M)) = \mathrm{Aut}(\pi_1(M))/\mathrm{Inn}(\pi_1(M))$. Therefore, we have a homomorphism $\widetilde{\Phi}_M\colon \mathrm{Aut}(M) \to \mathrm{Out}(\pi_1(M))$ given by $f \mapsto [f_*]$. It turns out $\mathrm{Aut}_0(M)$ is contained in the kernel of $\widetilde{\Phi}_M$, hence there is an induced homomorphism $\Phi_M\colon \operatorname{Mod}(M) \to \operatorname{Out}(\pi_1(M))$. In the case of $M = \mathbb{T}$, we know that $\pi_1(\mathbb{T}) \cong \mathbb{Z}^2$ is abelian, so

$$\operatorname{Out}(\pi_1(\mathbb{T})) \cong \operatorname{Aut}(\pi_1(\mathbb{T})) \cong \operatorname{Aut}(\mathbb{Z}^2) = GL(2, \mathbb{Z}).$$

Theorem 2.14. *The map $\Phi_{\mathbb{T}}\colon \operatorname{Mod}(\mathbb{T}) \to GL(2,\mathbb{Z})$ is an isomorphism.*

Given an element $A \in GL(2,\mathbb{Z})$, it is easy to write a diffeomorphism $f_A\colon \mathbb{T} \to \mathbb{T}$ such that $\Phi_{\mathbb{T}}([f_A]) = A$. First, we treat A as an element of $GL(2,\mathbb{R})$, hence there is an induced isomorphism $\widetilde{f}_A\colon \mathbb{R}^2 \to \mathbb{R}^2$. Recalling that A has integer entries, it follows that $\widetilde{f}_A$ restricts to a bijection $\mathbb{Z}^2 \to \mathbb{Z}^2$, therefore there is an induced isomorphism of Lie groups $f_A\colon \mathbb{R}^2/\mathbb{Z}^2 \to \mathbb{R}^2/\mathbb{Z}^2$. Noting that $\mathbb{R}^2/\mathbb{Z}^2$ and $\mathbb{T}$ are diffeomorphic completes the necessary construction.

A pair of knots (K, K') in $\mathbb{T}$ is called *complimentary* if their homology classes span $H_1(\mathbb{T})$, and they intersect transversally in a single point. We could have equivalently required the homotopy classes to span $\pi_1(\mathbb{T})$. When talking about the homology or homotopy class of a knot, there is a choice of orientation involved, yet that is unambiguous when we treat the element as a generator. The statement below follows from the Theorems 2.11, 2.14, and the remark above.

Corollary 2.15. *For any two pairs of complimentary knots, there exists a diffeomorphism $\mathbb{T} \to \mathbb{T}$ sending one pair to the other.*

This corollary is often useful when combined with the fact that for any knot K in $\mathbb{T}$, there exist a complimentary one K'.

"Instead of just the frosting, now consider the whole doughnut." [16]

Definition 2.16. A *solid torus* is a space P diffeomorphic to $S^1 \times D^2$. A *framing of P* is an explicit diffeomorphism $h\colon S^1 \times D^2 \to P$.

Let us fix a solid torus P.

Proposition 2.17. *For an embedded non-contractible curve μ in ∂P, the following are equivalent:*

(1) μ is homologically trivial in P,
(2) μ is homotopically trivial in P,
(3) μ bounds a disk in P,
(4) for some framing $h\colon S^1 \times D^2 \to P$, we have $\mu = h(\{1\} \times \partial D^2)$.

A curve μ satisfying any of the given conditions is called a meridian *of P.*

Proposition 2.18. *For an embedded non-contractible curve λ in ∂P, the following are equivalent:*

(1) λ represents a generator of $H_1(P) \cong \pi_1(P) \cong \mathbb{Z}$,
(2) λ intersects some meridian μ of P transversally in a single point,
(3) for some framing $h\colon S^1 \times D^2 \to P$, we have $\lambda = h(S^1 \times \{1\})$.

A curve λ satisfying any of the given conditions is called a longitude *of P.*

One is typically interested in the isotopy type of meridians and longitudes in ∂P. The following result sheds some light on this question.

Proposition 2.19. *Any two meridians in P are isotopic to each other. Any two longitudes of P are equivalent via a diffeomorphism of P, through they need not be isotopic.*

Consider an embedding of P in S^3, and let $Q = S^3 \setminus \operatorname{int}(P)$. It turns out that the homology of Q is isomorphic to the homology of a solid torus, yet they may not be diffeomorphic. The Mayer-Vietoris sequence of the decomposition $S^3 = P \cup Q$ implies a meridian of P always represents a generator of $H_1(Q) \cong \mathbb{Z}$. This homology group is even more interesting for the following fact: there is a unique isotopy type of a longitude which represents the zero cycle in $H_1(Q)$. Furthermore, there is also a unique (up to isotopy) framing $h\colon S^1 \times D^2 \to P$ such that $h(S^1 \times \{1\})$ is the zero cycle in $H_1(Q)$. This is called the *preferred framing.* In conclusion, the isotopy type of the meridian μ is a feature inherent to P. There are however infinitely many longitudes. If one chooses an embedding $P \hookrightarrow S^3$, then there is a *preferred longitude* λ.

Consider a knot $K\colon S^1 \to M$. There is an essentially unique choice of a tubular neighbourhood P for K in M. We define *the meridian of* K to be the meridian for the solid torus P. Similarly, *the longitude of* K is defined as a longitude for P. If $M = S^3$, then *the (preferred) longitude of* K (note the definite article) is the preferred longitude of P.

2.6. **Surgery on links in 3-manifolds.** Let M be a 3-manifold, perhaps with boundary, or even bordered, and let $L = \bigsqcup_{i=1}^n L_i$ be a link with n components in M. Furthermore, suppose we are given:

(1) disjoint closed tubular neighbourhoods N_i of the L_i in the interior $\mathrm{int}(M) \subset M$,
(2) a non-contractible simple closed curve η_i in each boundary ∂N_i.

By Corollary 2.15, there always exists a diffeomorphism $h_i\colon \partial N_i \to \partial N_i$ which send the meridian curve μ_i of L_i to η_i. Let h be the union of all h_i, that is, a map $h\colon \bigsqcup_{i=1}^n \partial N_i \to \bigsqcup_{i=1}^n \partial N_i$.

Definition 2.20. The *(Dehn) surgery on* M *along* L *with instructions (1) and (2)* is the manifold M' given by

$$M' = \left(M \setminus \bigcup_{i=1}^n \mathrm{int}(N_i) \right) \cup_h \left(\bigcup_{i=1}^n N_i \right).$$

Since regular neighbourhoods are essentially unique, M' does not depend on the choice of N_i. Furthermore, it is not hard to see the diffeomorphism class of M' depends only on the homotopy type of the curves η_i. As a matter of fact, the meridians μ_i are also well-defined only up to homotopy.

We can now restrict our attention to the case $M = S^3$. Consider an oriented link L with components L_i in S^2. For each L_i there are essentially unique choices for a meridian μ_i, and a preferred longitude λ_i. We can decompose each curve ν_i in terms of the basis $[\mu_i], [\lambda_i]$ for $\pi_1(\mathbb{T})$, say

$$h_*([\mu_i]) = [\eta_i] = a_i[\lambda_i] + b_i[\mu_i],$$

where $a_i, b_i \in \mathbb{Z}$. Depending on the orientation of η_i, there is a sign ambiguity in the expression above. Theorem 2.11 implies $\gcd(a_i, b_i) = 1$ (we drop the case $a_i = b_i = 0$ since η_i is not contractible). Therefore, we can dismiss the sign ambiguity without losing any information by defining

$$r_i = \frac{b_i}{a_i}$$

if $a_i \neq 0$, and $r_i = \infty$ if $a_i = 0$, $b_i = \pm 1$. The number $r_i \in \widetilde{\mathbb{Q}} = \mathbb{Q} \cup \{\infty\}$ is called the *surgery coefficient* associated to the component L_i. It is not hard to see r_i is also independent of the orientation of L_i. The only subtle assumption we are making is that S^3 is endowed with a fixed orientation – reversing that would reverse the signs of all surgery coefficients.

Remark 2.21. Let L be a link in S^3 with rational numbers (in $\widetilde{\mathbb{Q}}$) assigned to each of its components. We can form the link L' by dropping all components marked with ∞, and copying the rest of the surgery coefficients to L'. It turns our that performing surgeries on L and L' yields the same manifold. Therefore, we may drop ∞ as a surgery coefficient.

In conclusion, any link in S^3 with a rational number assigned to each of its components specifies unambiguously surgery instructions, hence a closed oriented 3-manifold. In fact, the converse is also true (see [5, Chapter 12] for a proof). If K is a knot in S^3, and $r \in \widetilde{\mathbb{Q}}$, we will denote by $S^3_r(K)$ the surgery on K with coefficient r.

Theorem 2.22 (Lickorish-Wallace)**.** *Any closed oriented 3-manifold may be obtained by a surgery on a link in S^3. In fact, all surgery coefficients can be chosen to be ± 1.*

Before closing our discussion of surgery on 3-manifolds, it is important to make a remark about its relation to 4-manifolds. Consider a 4-manifold N with boundary. Attaching *2-handle* to N is the operation of gluing a copy of $D^2 \times D^2$ along $S^1 \times D^2$ to ∂N. It is not hard to see that a 2-handle attachment changes the boundary of N by a surgery on a knot. Similarly, attaching several 2-handles has the effect of a surgery on a link on the boundary. Hence the following result.

Corollary 2.23. *Any closed oriented 3-manifold is the boundary of some oriented 4-manifold.*

Equally interesting is the following operation: given a closed oriented 3-manifold M, we may thicken it to $M \times I$, and perform 2-handle attachments on one side of its boundary. The resulting 4-manifold is a cobordism between M and the result of the surgery corresponding to the handle attachments. It follows from Theorem 2.22 that any two closed oriented 3-manifolds may be connected via a bordism. In other words, the graph associated to the category $\mathbf{Bord}_4$ is connected.

3. Bordered Heegaard Floer homology

Heegaard Floer homology is a collective title for several related topological invariants for 3-manifolds which were defined and developed by P. Ozsvàth and Z. Szabò in a series of papers (the first of these is [13]). In [14], these ideas were further extended to functor-like objects emanating from $\mathbf{Bord}_4$ (there are a few details about gluing of spin^c structures which require careful handling). Similar methods were adapted to provide invariants for knots in 3-manifolds as explained in [12]. More recent work extends the theory to bordered objects [6], [7].

3.1. The conceptual picture.

The aim of this section is to provide a brief high-level picture of the bordered Heegaard Floer package without elaborating on any details (see [7] for a thorough treatment).

Let M be a compact oriented 3-manifold. The Heegaard Floer invariants $HF^-(M, \mathfrak{s})$, $HF^\infty(M, \mathfrak{s})$, $HF^+(M, \mathfrak{s})$, $\widehat{HF}(M, \mathfrak{s})$ are defined by taking homology of the complexes $CF^-(M, \mathfrak{s})$, $CF^\infty(M, \mathfrak{s})$, $CF^+(M, \mathfrak{s})$, $\widehat{CF}(M, \mathfrak{s})$ coming from the data of a Heegaard diagram for M. The following result is a statement of independence from several choices made along the way, one of them being a Heegaard diagram.

Theorem 3.1 (result 11.1 in [13])**.** *The relatively $\mathbb{Z}/n$-graded modules $HF^-(M, \mathfrak{s})$, $HF^\infty(M, \mathfrak{s})$, $HF^+(M, \mathfrak{s})$, and $\widehat{HF}(M, \mathfrak{s})$ are topological invariants of the underlying 3-manifold M and its* spin^c *structure $\mathfrak{s}$. The integer $n \geq 0$ is given by the divisibility of the first Chern class of $\mathfrak{s}$.*

An exposition of the bordered theory should start with an assignment of a differential graded algebra $\mathcal{A}(F)$ to each oriented surface F. In Section 3.2, we will give an explicit construction of this algebra in the case $F = \mathbb{T}$. For each manifold M with parametrized boundary F, there are two invariants – type A denoted by $\widehat{CFA}(M)$, and type D denoted by $\widehat{CFD}(M)$. The former is a right ($\mathcal{A}_\infty$-) $\mathcal{A}(F)$-module, and the latter is a left $\mathcal{A}(F)$-module. The following result draws a connection between the bordered and non-bordered theories.

Theorem 3.2 (Pairing Theorem, result 1.3 in [7]). *Let M_1 and M_2 be two compact oriented bordered 3-manifolds with boundaries $\partial M_1 = F = -\partial M_2$, and let M be the closed 3-manifold obtained by gluing M_1 and M_2 along F. Then $\widehat{CF}(M)$ is quasi-isomorphic to the $\mathcal{A}_\infty$ tensor product of $\widehat{CFA}(M_1)$ and $\widehat{CFD}(M_2)$. In particular,*

$$\widehat{HF}(M) \cong H_* \left(\widehat{CFA}(M_1) \widetilde{\otimes}_{\mathcal{A}(F)} \widehat{CFD}(M_2) \right).$$

There are several points worth clarifying.

Remark 3.3 (On tensor products). The $\mathcal{A}_\infty$ tensor product used above is a fancy extension of the typical version which attempts to capture additional derived information. Fleshing out the construction would require us to go into details about $\mathcal{A}_\infty$ structures and homological algebra. This is not necessary for our needs, hence, without loss of generality, one may imagine $\widetilde{\otimes}$ to be $\otimes$ throughout.

Remark 3.4 (On gradings). Even in the non-bordered case (for example, Theorem 3.1), one is already using a relative $\mathbb{Z}$-grading. Transitioning to the bordered theory, one has to replace $\mathbb{Z}$ by a non-commutative group which depends on F (see [7] for complete details).

3.2. **The torus algebra.** We start by recalling the structure of the *torus algebra* $\mathcal{A}(\mathbb{T}) = \mathcal{A}(\mathbb{T}, 0)$ as in [7, Section 10.1]. There is a subalgebra of idempotents $\mathcal{I}(\mathbb{T}) \subset \mathcal{A}(\mathbb{T})$ generated by ι_0 and ι_1, and the unit is $1 = \iota_0 + \iota_1$. More precisely, $\mathcal{I}(\mathbb{T}) = \mathbb{F}_2\langle \iota_0, \iota_2 \rangle$ as a vector space, and the generators satisfy relations

$$\iota_0^2 = \iota_0, \quad \iota_1^2 = \iota_1, \quad \iota_0\iota_1 = \iota_1\iota_0 = 0.$$

The algebra $\mathcal{A}(\mathbb{T})$ is generated by elements $\rho_1, \rho_2, \rho_3, \rho_{12}, \rho_{23}, \rho_{123}$ over $\mathcal{I}(\mathbb{T})$ satisfying the relations

$$\rho_1\rho_2 = \rho_{12}, \quad \rho_2\rho_3 = \rho_{23}, \quad \rho_{12}\rho_3 = \rho_{123}, \quad \rho_1\rho_{23} = \rho_{123},$$

and all other products of generators are zero. The compatibility conditions with the idempotents are

$$\begin{aligned} \rho_1 &= \iota_0\rho_1\iota_1, & \rho_2 &= \iota_1\rho_2\iota_0, & \rho_3 &= \iota_0\rho_3\iota_1, \\ \rho_{12} &= \iota_0\rho_{12}\iota_0, & \rho_{23} &= \iota_1\rho_{23}\iota_1, & \rho_{123} &= \iota_0\rho_{123}\iota_1. \end{aligned}$$

We will adopt the notation $\mathcal{A}_j = \mathcal{A}(\mathbb{T})\iota_j$ for $j = 0, 1$.

3.3. **A definition.** An understanding of the type D modules in the torus boundary case suffices for our purposes. We will briefly outline the construction of $\widehat{CFD}$ in this specific case, and then work through an example in the following section.

Let M be a compact oriented bordered 3-manifold with torus boundary, and $\mathcal{H} = (\overline{\Sigma}, \boldsymbol{\alpha}, \boldsymbol{\beta}, z)$ a pointed bordered Heegaard diagram for it. By construction, $\widehat{CFD}(M) = \widehat{CFD}(\mathcal{H})$ is well-defined up to quasi-isomorphism. We will focus on the case in which $\mathcal{H}$ has genus one, when M is a solid torus, and $\overline{\Sigma}$ is a punctured torus as in Figure 3. More specifically, we imagine cutting $\mathcal{H}$ along the arcs α_1^a, α_2^a as in Figure 3, Presentation 3. We will not make further restraining assumptions regarding the curve β_1. Consider the intersection points $\boldsymbol{\alpha} \cap \boldsymbol{\beta}$ which partition as $(\alpha_1^a \cap \beta_1) \sqcup (\alpha_2^a \cap \beta_1)$. Let

$$\alpha_1^a \cap \beta_1 = \{a_1, \ldots, a_s\} \qquad \text{and} \qquad \alpha_2^a \cap \beta_1 = \{b_1, \ldots, b_t\}$$

for two integers $s, t \geq 0$. As a left $\mathcal{A}(\mathbb{T})$-module, we set

$$\widehat{CFD}(\mathcal{H}) = \mathcal{A}_0\langle a_1, \ldots, a_s\rangle \oplus \mathcal{A}_1\langle b_1, \ldots, b_t\rangle.$$

To flesh out the differential $\partial\colon \widehat{CFD}(\mathcal{H}) \to \widehat{CFD}(\mathcal{H})$, it suffices to define it on the generators $\boldsymbol{\alpha} \cap \boldsymbol{\beta}$, and extend by linearity. For any $x \in \boldsymbol{\alpha} \cap \boldsymbol{\beta}$, we have $\partial x = \sum_{y \in \boldsymbol{\alpha} \cap \boldsymbol{\beta}} c_{x,y} y$ for some $c_{x,y} \in \mathcal{A}$. The coefficient $c_{x,y}$ is a weighted count of the holomorphic disks joining x and y.

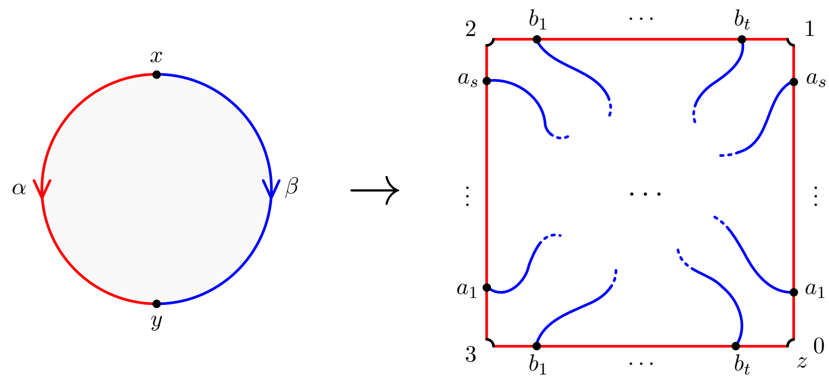

Figure 4: A holomorphic disk, pictorially.

At this point, we elaborate on the specific meaning of holomorphic disks for the genus 1 case, and describe the associated weights. For the sake of specificity, we realize the 2-disk as $D^2 = \{z \in \mathbb{C} \mid |z| \leq 1\}$. At its heart, a holomorphic disk from x to y is a map $\varphi\colon D^2 \to \overline{\Sigma}$ which satisfies several properties as suggested by Figure 4:

(1) $\varphi(i) = x$,
(2) $\varphi(-i) = y$,
(3) $\varphi(\partial D^2 \cap \{z \mid \operatorname{Re}(z) < 0\}) \subset \boldsymbol{\alpha} \cup (\partial\overline{\Sigma} \setminus \{z\})$,
(4) $\varphi(\partial D^2 \cap \{z \mid \operatorname{Re}(z) > 0\}) \subset \boldsymbol{\beta}$,
(5) φ is holomorphic.

Items (1) and (2) merely say that the image of φ should join x and y. Condition (3) requires the left arc between i and $-i$ (red in Figure 4) to map along the union of the $\boldsymbol{\alpha}$-arcs, and the punctured boundary of $\overline{\Sigma}$, that is $\partial\overline{\Sigma} \setminus \{z\}$. Condition (4) is a restatement of (3) for the right arc which should map to the $\boldsymbol{\beta}$-curve (equivalently, we could have added $\partial\overline{\Sigma} \setminus \{z\}$). Finally, (5) forces φ to be orientation preserving, and imposes a type of rigidity requirement. Recall that we split $\partial\overline{\Sigma} \setminus \{z\}$ into four regions labelled 0 to 3 in Figure 4. Suppose the right arc in going from i to $-i$ touches regions with labels $i_1, \ldots, i_k$ in that order for some $k \geq 0$. Then the weight of φ is given by $\prod_{j=1}^{k} \rho_{i_j} \in \mathcal{A}(\mathbb{T})$, whereas $c_{x,y}$ is the sum over the weights of all holomorphic disks joining x and y.

3.4. **An example.** Now that we have described all ingredients that go into the D type module $\widehat{CFD}(\mathcal{H})$, we proceed to work through an example. Consider the bordered pointed Heegaard diagram $\mathcal{H}$ shown in Figure 5 (a). The associated D type module is $\widehat{CFD}(\mathcal{H}) = \mathcal{A}_0 a \oplus \mathcal{A}_1 b_1 \oplus \mathcal{A}_1 b_2$. It remains to compute its differential ∂ by counting holomorphic disks. There is a unique disk from a to b_1 shown in (b) whose weight is ρ_3. Considering a and b_1, there are two disks shown in (b) and (c). The former one has weight ρ_1, and the latter $\rho_3\rho_2\rho_3 = 0$. To read properly the disk shown in (c), we should recall that the red edges are identified in pairs, and note that the darker shade indicates an overlap. Finally, there is a unique disk joining b_1 and b_2 shown in (e) with weight $\rho_2\rho_3 = \rho_{23}$. A moment of thought would show there are no other holomorphic disks present. Putting this information together, we obtain

$$\partial a = \rho_3 b_1 + \rho_1 b_2, \qquad \partial b_1 = \rho_{23} b_2, \qquad \partial b_2 = 0.$$

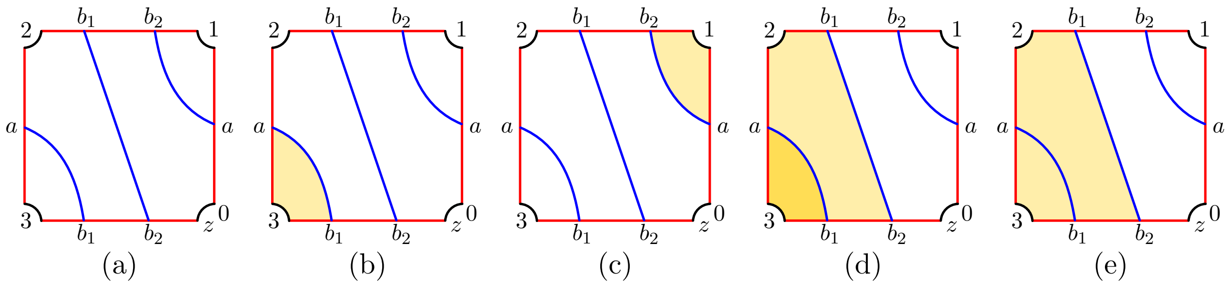

Figure 5: A bordered pointed Heegaard diagram $\mathcal{H}$, and all relevant holomorphic disks.

Note that using the notation from Section 4.2, we have $\mathcal{H} = \mathcal{H}_{-2}$. One can specialize the general computation to our results here.

4. Knots in S^3

Our aim in this section is to outline a potentially interesting invariant for knots in S^3. Directions for future investigation are presented in Section 5. Our ideas draw inspiration from a result by W. Thurston (Theorem 4.1), and we will start by describing this.

4.1. **Hyperbolic surgery and Gromov-Hausdorff convergence.** Let M be a complete Riemannian manifold. We call M *hyperbolic* if it has constant sectional curvature -1. As a consequence of Mostow's Rigidity Theorem [1, Chapter C], if a smooth 3-manifold has a hyperbolic structure, it is essentially unique (generally, this holds in dimension $n \geq 3$). Therefore, it makes sense to call a 3-manifold hyperbolic if it is possible to endow it with a hyperbolic structure but without a reference to a specific such structure. Recall that Riemannian manifolds, in particular hyperbolic ones, are naturally metric spaces. Therefore, one can endow a suitable class of hyperbolic 3-manifolds with the pointed Gromov-Hausdorff metric [3]. Convergence in this metric space is referred to as *Gromov-Hausdorff convergence.*

A knot K in S^3 is called *hyperbolic* if its compliment $S^3 \setminus K$ is a hyperbolic 3-manifold. It turns our that for a hyperbolic knot K, all but finitely many of the manifolds $S^3_r(K)$ for $r \in \mathbb{Q}$ are hyperbolic. The following is a very intriguing result (see [1, Chapter E] for a proof).

Theorem 4.1 (W. Thurston's Hyperbolic Surgery Theorem)**.** *Let K be a hyperbolic knot in S^3, and consider a sequence of rational numbers $p_i/q_i \in \mathbb{Q}$ for $i \geq 0$. If $p_i^2 + q_i^2 \to \infty$, then the sequence of manifolds $S^3_{p_i/q_i}(K)$ converge to $S^3 \setminus K$ in the Gromov-Hausdorff sense.*

Remark 4.2. The result above could be suitably formulated in greater generality – it holds for a link L in a 3-manifold M provided $M \setminus L$ is hyperbolic.

Informally, the collection of surgeries with "large" coefficients encodes topological information about $K \subset S^3$. In what follows, we will attempt to extract such information via the Heegaard Floer package.

4.2. **A bordered interpretation of surgery.** Let K be a knot in S^3. Fix a tubular neighbourhood M_∞ of K, and let N be the closure of the complement $S^3 \setminus M_\infty$. We can treat both N and M_∞ as bordered 3-manifolds with boundary $\mathbb{T}$ such that the meridian and preferred longitude of the boundary torus $\mathbb{T}$ are respectively a meridian and preferred longitude for K. Fix a bordered Heegaard diagram $\mathcal{H}_N$ for N, and similarly $\mathcal{H}_\infty$ for M_∞. Without loss of generality, we may assume $\mathcal{H}_\infty$ is as shown in the left side of Figure 6. Similarly, for any $r \in \mathbb{Q}$, we can construct a diagram $\mathcal{H}_r$ in which the curve β_1 has slope r relative to the coordinate system induced by the arcs α_1^a and

α_2^a. These observations are interesting since for any $r \in \mathbb{Q}$, we have $S_r^3(K) = N \cup_{\mathbb{T}} M_r$, where M_r

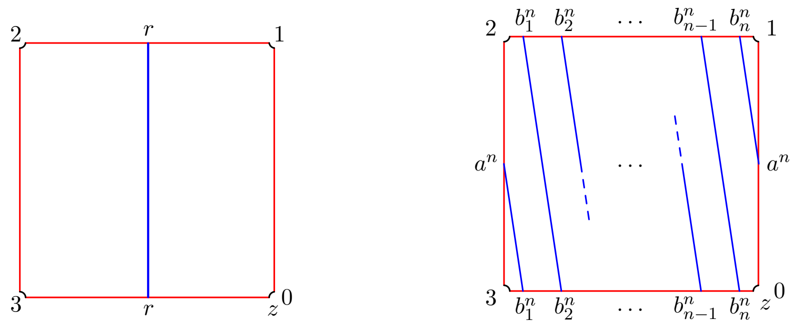

Figure 6: The bordered Heegaard diagrams $\mathcal{H}_\infty$, and $\mathcal{H}_{-n}$ for $n \geq 0$.

corresponds to the bordered Heegaard diagram $\mathcal{H}_r$. The Pairing Theorem (3.2) then implies that $\widehat{CF}(S^3(K))$ is quasi-isomorphic to $\widehat{CFA}(\mathcal{H}_N)\widetilde{\otimes}_{\mathcal{A}(\mathbb{T})}\widehat{CFD}(\mathcal{H}_r)$.

We can specialize to the case when the surgery coefficient is a negative integer. Recall our discussion on the relation between surgery on 3-manifolds and attaching handles on 4-manifolds (Section 2.6). It follows that for each $n > 0$, there is a bordism from $S_{-n}^3(K)$ to $S_{-(n-1)}^3$. This induces a morphism of graded differential modules $\Psi_n\colon \widehat{CF}(S_{-n}^3(K)) \to \widehat{CF}(S_{-(n-1)}^3(K))$. An analogous Pairing Theorem for bordisms implies there exists a commutative diagram

$$\begin{tikzcd}
\widehat{CF}(S^3_{-n}(K)) \arrow[r, "\Psi_n"] \arrow[d, "\cong"'] & \widehat{CF}(S^3_{-(n-1)}(K)) \arrow[d, "\cong"] \\
\widehat{CFA}(\mathcal{H}_N)\widetilde{\otimes}_{\mathcal{A}(\mathbb{T})}\widehat{CFD}(\mathcal{H}_{-n}) \arrow[r, "\mathrm{id}_{\widehat{CFA}(\mathcal{H}_N)}\otimes\psi_n"] & \widehat{CFA}(\mathcal{H}_N)\widetilde{\otimes}_{\mathcal{A}(\mathbb{T})}\widehat{CFD}(\mathcal{H}_{-(n-1)})
\end{tikzcd}$$

in which both vertical arrows are quasi-isomorphisms, and $\psi_n\colon \widehat{CFD}(\mathcal{H}_{-n}) \to \widehat{CFD}(\mathcal{H}_{-(n-1)})$ is a morphism of differential left $\mathcal{A}(\mathbb{T})$-modules. The maps Ψ_n fit into an inverse system

$$\cdots \longrightarrow \widehat{CF}(S^3_{-n}) \xrightarrow{\Psi_n} \widehat{CF}(S^3_{-(n-1)}(K)) \longrightarrow \cdots \longrightarrow \widehat{CF}(S^3_{-1}(K)) \xrightarrow{\Psi_1} \widehat{CF}(S^3_0(K)),$$

whose limit $C_K = \varprojlim \widehat{CF}(S^3_{-n}(K))$ is an object of interest. A priori, it seems we need to compute $\widehat{CF}(S^3_{-n}(K))$ for all $n > 0$ in order to gain some knowledge about C_K. One can however reduce this to two calculations – that of $\widehat{CFA}(\mathcal{H}_N)$ which depends on the knot K, and a second one which is independent of K. Let us consider the inverse system

$$\cdots \longrightarrow \widehat{CFD}(\mathcal{H}_{-n}) \xrightarrow{\psi_n} \widehat{CFD}(\mathcal{H}_{-(n-1)}) \longrightarrow \cdots \longrightarrow \widehat{CFD}(\mathcal{H}_{-1}) \xrightarrow{\psi_1} \widehat{CFD}(\mathcal{H}_0)$$

with limit $C = \varprojlim \widehat{CFD}(\mathcal{H}_{-n})$. Since the module $\widehat{CFA}(\mathcal{H}_N)$ is finitely generated, some formal observations about algebra (see Appendix) imply that

$$C_K \cong \widehat{CFA}(\mathcal{H}_N)\widetilde{\otimes}_{\mathcal{A}(\mathbb{T})}C.$$

Our next goal is to compute the module C.

4.3. **The inverse limit of $\widehat{\boldsymbol{CFD}}(\mathcal{H}_{-n})$.** Let us start by describing $\widehat{CFD}(\mathcal{H}_\infty)$, and $\widehat{CFD}(\mathcal{H}_{-n})$ for $n \geq 0$ (see Figure 6 for the relevant bordered Heegaard diagrams). The module $\widehat{CFD}(\mathcal{H}_\infty)$ is generated by r satisfying $\partial r = \rho_{23} r$. We have $r\iota_1 = r$, so $\widehat{CFD}(\mathcal{H}_\infty) = \mathcal{A}_1\langle r\rangle$. Similarly, the complex $\widehat{CFD}(\mathcal{H}_{-n})$ is generated by $a^n, b_1^n, \ldots, b_n^n$ and these satisfy

$$\partial a^n = \rho_3 b_1^n + \rho_1 b_n^n, \quad \partial b_1^n = \rho_{23} b_2^n, \quad \ldots \quad \partial b_{n-1}^n = \rho_{23} b_n^n, \quad \partial b_n^n = 0, \tag{4.3}$$

when $n > 0$. The remaining case $n = 0$ is degenerate in the sense that $\partial a_0 = \rho_{12} a_0$. For all $n \geq 0$, we have $a^n \iota_0 = a^n$ and $b_i^n \iota_1 = b_i^n$, so $\widehat{CFD}(\mathcal{H}_{-n}) = \mathcal{A}_0\langle a^n\rangle \oplus \mathcal{A}_1\langle b_1^n, \ldots, b_n^n\rangle$.

The following summarizes our claim, and the remainder of this section is devoted to its proof.

Proposition 4.4. *Let $C = \varprojlim \widehat{CFD}(\mathcal{H}_{-n})$. There is an isomorphism of $\mathcal{A}(\mathbb{T})$-modules*

$$\Phi\colon C \to \mathcal{A}_0 \times \prod_{n\geq 1} \mathcal{A}_1.$$

Remark 4.5. Although $\mathcal{A}_0 \times \prod_{n\geq 1} \mathcal{A}_1$ is not a differential module naturally, the morphism Φ induced a differential structure on it. This is explicitly given by Equations 4.11.

We start by constructing the aforementioned inverse system. For each $n > 0$, there is an exact sequence

$$0 \longrightarrow \widehat{CFD}(\mathcal{H}_\infty) \xrightarrow{\varphi_n} \widehat{CFD}(\mathcal{H}_{-n}) \xrightarrow{\psi_n} \widehat{CFD}(\mathcal{H}_{-(n-1)}) \longrightarrow 0,$$

where the maps φ_n, ψ_n are given by

$$\varphi_n(r) = \sum_{i=1}^n b_i^n + \rho_2 a^n, \tag{4.6}$$

and

$$\begin{aligned} &\psi_n(a^n) = a^{n-1}, && \psi_n(b_1^n) = b_1^{n-1} + \rho_2 a^{n-1}, \\ &\psi_n(b_i^n) = b_{i-1}^{n-1} + b_i^{n-1} \text{ for } 2 \leq i \leq n-1, && \psi_n(b_n^n) = b_{n-1}^{n-1}. \end{aligned} \tag{4.7}$$

These expressions are derived by counting triangles in the certain bordered Heegaard diagrams as in [7, Section 10.2]. It is easy to verify these maps fit into a short exact sequence as already indicated. This completes the construction of the speculated inverse system.

Remark 4.8. Applying the homology functor to the short exact sequence above yields the corresponding surgery exact triangle (see [7, Section 10.2]).

By definition $C = \varprojlim \widehat{CFD}(\mathcal{H}_{-n})$ is a submodule of

$$C' = \prod_{n\geq 0} \widehat{CFD}(\mathcal{H}_{-n}) = \prod_{n\geq 0}\left(\mathcal{A}_0 a^n \oplus \bigoplus_{i=1}^n \mathcal{A}_1 b_i^n\right) = \left(\prod_{n\geq 0} \mathcal{A}_0 a^n\right) \times \left(\prod_{1\leq i\leq n} \mathcal{A}_1 b_i^n\right).$$

Therefore, any element $c \in C'$ can be written as $c = (c^n)_{n\geq 0}$, where $c^n \in \widehat{CFD}(\mathcal{H}_{-n})$. If $c^n = (c_i^n)_{0\leq i\leq n}$, then $c = (c_i^n)_{0\leq i\leq n}$, where $c_0^n \in \mathcal{A}_0$ are the coefficients of a^n, and $c_i^n \in \mathcal{A}_1$ are the coefficients of b_i^n. One can also use an infinite upper-triangular matrix

$$c = \begin{pmatrix} c_0^0 & c_0^1 & c_0^2 & c_0^3 & \cdots \\ & c_1^1 & c_1^2 & c_1^3 & \cdots \\ & & c_2^2 & c_2^3 & \cdots \\ & & & c_3^3 & \cdots \\ & & & & \ddots \end{pmatrix}$$

whose columns correspond to $c^n \in \widehat{CFD}(\mathcal{H}_{-n})$. By construction, an element $c = (c^n)_n \in C'$ lies in C if and only if $\psi_n(c^n) = c^{n-1}$ for every $n > 0$. The following diagram illustrates element c and the maps ψ_n going left between adjacent columns.

$$\begin{tikzcd}
c_0^0 & c_0^1 \arrow[l] & c_0^2 \arrow[l] & c_0^3 \arrow[l] & \cdots \arrow[l] \\
& c_1^1 \arrow[ul, "\rho_2"] & c_1^2 \arrow[l] \arrow[ul, "\rho_2"] & c_1^3 \arrow[l] \arrow[ul, "\rho_2"] & \cdots \arrow[l] \arrow[ul, "\rho_2"] \\
& & c_2^2 \arrow[ul] & c_2^3 \arrow[l] \arrow[ul] & \cdots \arrow[l] \arrow[ul] \\
& & & c_3^3 \arrow[ul] & \cdots \arrow[l] \arrow[ul] \\
& & & & \ddots \arrow[ul]
\end{tikzcd}$$

Labels along arrows indicate multiplicative coefficients; the absence thereof means the coefficient is 1. In other words, $c \in C$ if and only if

$$c_0^n = c_0^{n+1} + \rho_2 c_1^{n+1}, \quad c_i^n = c_i^{n+1} + c_{i+1}^{n+1} \tag{4.9}$$

for all $n \geq 0$ and $1 \leq i \leq n$.

Define the map

$$\Phi \colon C \to \mathcal{A}_0 \times \prod_{n \geq 1} \mathcal{A}_1$$

by $c = (c_i^n)_{n,i} \mapsto (c_n^n)_n$ which diagrammatically corresponds to filtering out the coefficients along the main diagonal. We proceed to verify this is an isomorphism. This amounts to showing the diagonal entries c_n^n determine all entries above the diagonal c_i^n for $0 \leq i < n$. Suppose we are given an arbitrary $d = (d_n) \in \mathcal{A}_0 \times \prod_{n \geq 1} \mathcal{A}_1$, and we would like to construct some $c = (c_i^n)_{n,i} \in C$ such that $c_n^n = d_n$ for all $n \geq 0$. The choices

$$\begin{aligned} c_0^n &= d_0 + \rho_2 \sum_{j=1}^{n} (-1)^j \binom{n}{j} d_j \quad \text{for} \quad n \geq 0, \\ c_i^n &= \sum_{j=0}^{n-i} (-1)^j \binom{n-i}{j} d_{i+j} \quad \text{for} \quad 1 \leq i \leq n \end{aligned} \tag{4.10}$$

satisfy Equations 4.9. Furthermore, proceeding by induction on $n-i$ it is not hard to see Equations 4.9 uniquely determine all c_i^n. This proves the map Φ is an isomorphism as claimed.

Each of the modules $\widehat{CFD}(\mathcal{H}_{-n})$ has a differential which commutes with the maps ψ_n, so C inherits one too. The isomorphism Φ induces a differential on its codomain. We proceed to give an explicit description of it. Suppose $\partial(d_n)_n = (d'_n)_n$ and $(d_n)_n = \Phi((c_i^n)_{n,i})$. It follows from Equation 4.3 that

$$\begin{aligned}
d'_0 &= \rho_{12} d_0, \\
d'_1 &= (\rho_1+\rho_3)c_0^1 = (\rho_1+\rho_3)d_0 - \rho_{12}d_1, \\
d'_n &= \rho_1 c_0^n + \rho_{23}c_{n-1}^n = \rho_1 d_0 + \rho_{12}\sum_{j=1}^{n}(-1)^j\binom{n}{j}d_j + \rho_{23}(d_{n-1}-d_n) \quad \text{for} \quad n>1.
\end{aligned} \tag{4.11}$$

This completes the proof of Proposition 4.4.

5. Final remarks

There are numerous directions one can proceed to from Section 4. The aim of this section is to outline a few of these, and describe some of the existing shortcomings.

For once, using negative integral surgery coefficients might have seemed rather unmotivated. Of course, there are infinitely many sequences of surgery coefficients we could have used to produce a direct or inverse system. The only other case we investigated is that of using positive integral coefficients – it yields a direct system, whose limit is trivial. Testing other such sequences could be one direction of potential future development.

We briefly mentioned the presence of a delicate non-commutative grading in the case of bordered Heegaard Floer homology, but this discussion was not continued in Section 4.2. The reason for this arrangement is different $\widehat{CFD}(\mathcal{H}_{-n})$ have slightly different gradings, and the maps ψ_n are not compatible with these in any understood manner. Hence, there is no sensible way to assign a grading to the inverse limit C. There are however some relations between the various gradings, and exploring these further could be another way to advance this project.

Recently (in [7, Appendix], and [8]), bordered Heegaard Floer homology has been extended to manifolds with two boundary components in which case the result is a bimodule (one has to choose a type A or D for each boundary component). In the spirit of Section 4.2, it is possible to interpret knot surgery as a succession of gluings, all intermediate pieces being bordered manifolds with two boundary components, and the beginning and end, manifolds with a single boundary component. An analogous Pairing Theorem in the bimodule setting enables us to write the module $\widehat{CFD}(\mathcal{H}_{-n})$ as the tensor product of the n-th tensor power of a "twist" bimodule, and $CFD(\mathcal{H}_\infty)$. Studying stabilization properties of increasingly large tensor powers could be another way to approach the module C we discussed.

Comparing Sections 4.1 and 4.2, we note the hyperbolic hypothesis on the knot K has been dropped out. This could be treated as an advantage of the Heegaard Floer approach in comparison to Thurston's Hyperbolic Surgery Theorem (result 4.1).

Appendix: Tensor products

Let A be a ring, M a right A-module, and N_i a collection of left A-modules indexed by $i \in I$. We are interested in the objects

$$M \otimes_A \prod_{i\in I} N_i \quad \text{and} \quad \prod_{i\in I}(M \otimes_A N_i).$$

With the present setup, these are abelian groups. If however, we require M to be an A-bimodule, then they would also inherit the structure of left A-modules. Similarly, if we require all N_i to be A-bimodules, we will have right A-modules.

There always exists a map Φ given by $m \otimes (n_i)_i \mapsto (m \otimes n_i)_i$ as shown in the diagram below.

$$\begin{tikzcd} M \otimes_A \prod_{i\in I} N_i \arrow[r, shift left, "\Phi"] & \prod_{i\in I}(M \otimes_A N_i) \arrow[l, shift left, dashed, "\Psi"] \end{tikzcd}$$

If we assume M is finitely generated, then there exists an inverse Ψ. Let $m_1, \ldots, m_k$ be a generating set for M, and consider an element $((\sum_{j=1}^k m_j a_{ij}) \otimes n_i)_i$ where $a_{ij} \in A$. Letting $n_{ij} = a_{ij} n_i$, we can rewrite this as

$$\begin{aligned}\left(\left(\sum_{j=1}^k m_j a_{ij}\right) \otimes n_i\right)_i &= \left(\sum_{j=1}^k (m_j a_{ij}) \otimes n_i\right)_i = \left(\sum_{j=1}^k m_j \otimes (a_{ij} n_i)\right)_i \\ &= \left(\sum_{j=1}^k m_j \otimes n_{ij}\right)_i = \sum_{j=1}^k (m_j \otimes n_{ij})_i,\end{aligned}$$

so it is reasonable to map this element to $\sum_{j=1}^k m_j \otimes (n_{ij})_i$ via Ψ. It is not hard to see the so defined maps Φ and Ψ are morphisms of abelian groups and inverses of each other.

Let us further assume we are given a small category $\mathcal{I}$ with set of objects I, and a functor $D\colon \mathcal{I} \to A\text{-mod}$ satisfying $D(i) = N_i$ for all $i \in I$ (here A-mod stands for the category of left A-modules). There is an induced functor $M \otimes D\colon \mathcal{I} \to \mathbf{Ab}$ given by $(M \otimes D)(i) = M \otimes N_i$ and $(M \otimes D)(f) = \mathrm{id}_M \otimes_A D(f)$ for $f \in \mathrm{Mor}_{\mathcal{I}}(i,j)$, $i, j \in I$. The limits of D and $M \otimes D$ are naturally constructed as subgroups of $\prod_i N_i$ and $\prod_i (M \otimes N_i)$. It is not hard to see the maps Φ and Ψ above restrict to isomorphisms

$$\begin{tikzcd} M \otimes_A (\lim D) \arrow[r, shift left, "\Phi"] & \lim(M \otimes D). \arrow[l, shift left, "\Psi"] \end{tikzcd}$$

As before, had the module M not been finitely generated, we would only be able to provide an inclusion Φ.

References

1. Riccardo Benedetti and Carlo Petronio, *Lectures on hyperbolic geometry*, Universitext, Springer-Verlag, Berlin, 1992. MR MR1219310 (94e:57015)
2. R. H. Bing, *An alternative proof that* 3*-manifolds can be triangulated*, Ann. of Math. (2) **69** (1959), 37–65. MR MR0100841 (20 #7269)
3. Misha Gromov, *Metric structures for Riemannian and non-Riemannian spaces*, english ed., Modern Birkhäuser Classics, Birkhäuser Boston Inc., Boston, MA, 2007, Based on the 1981 French original, With appendices by M. Katz, P. Pansu and S. Semmes, Translated from the French by Sean Michael Bates. MR MR2307192 (2007k:53049)
4. John Hempel, 3*-Manifolds*, Princeton University Press, Princeton, N. J., 1976, Ann. of Math. Studies, No. 86. MR MR0415619 (54 #3702)
5. W. B. Raymond Lickorish, *An introduction to knot theory*, Graduate Texts in Mathematics, vol. 175, Springer-Verlag, New York, 1997. MR MR1472978 (98f:57015)
6. Robert Lipshitz, *A Heegaard-Floer invariant of bordered 3-manifolds*, Ph.D. thesis, Stanford University, Palo Alto, CA, 2006.
7. Robert Lipshitz, Peter Ozsvath, and Dylan Thurston, *Bordered Heegaard Floer homology: Invariance and pairing*, Sciences-New York (2008), 1–177.
8. Robert Lipshitz, Peter S. Ozsvath, and Dylan P. Thurston, *Bimodules in bordered Heegaard Floer homology*, (2010), 145.

9. J. Milnor, *Morse theory*, Based on lecture notes by M. Spivak and R. Wells. Annals of Mathematics Studies, No. 51, Princeton University Press, Princeton, N.J., 1963. MR MR0163331 (29 #634)
10. Edwin E. Moise, *Affine structures in 3-manifolds. V. The triangulation theorem and Hauptvermutung*, Ann. of Math. (2) **56** (1952), 96–114. MR MR0048805 (14,72d)
11. Liviu I. Nicolaescu, *An invitation to Morse theory*, Universitext, Springer, New York, 2007. MR MR2298610 (2009m:58023)
12. Peter Ozsváth and Zoltán Szabó, *Holomorphic disks and knot invariants*, Adv. Math. **186** (2004), no. 1, 58–116. MR MR2065507 (2005e:57044)
13. ______, *Holomorphic disks and topological invariants for closed three-manifolds*, Ann. of Math. (2) **159** (2004), no. 3, 1027–1158. MR MR2113019 (2006b:57016)
14. ______, *Holomorphic triangles and invariants for smooth four-manifolds*, Adv. Math. **202** (2006), no. 2, 326–400. MR MR2222356 (2007i:57029)
15. V. V. Prasolov and A. B. Sossinsky, *Knots, links, braids and 3-manifolds*, Translations of Mathematical Monographs, vol. 154, American Mathematical Society, Providence, RI, 1997, An introduction to the new invariants in low-dimensional topology, Translated from the Russian manuscript by Sossinsky [Sosinskiuı]. MR MR1414898 (98i:57018)
16. Dale Rolfsen, *Knots and links*, Mathematics Lecture Series, vol. 7, Publish or Perish Inc., Houston, TX, 1990, Corrected reprint of the 1976 original. MR MR1277811 (95c:57018)
17. James Singer, *Three-dimensional manifolds and their Heegaard diagrams*, Trans. Amer. Math. Soc. **35** (1933), no. 1, 88–111. MR MR1501673