ANNALES DE LA FACULTÉ DES SCIENCES DE TOULOUSE Mathématiques

Bijan Sahamie

Introduction to the basics of Heegaard Floer homology

Tome XXII, n° 2 (2013), p. 269-336.

<http://afst.cedram.org/item?id=AFST_2013_6_22_2_269_0>

Annales de la Faculté des Sciences de Toulouse Vol. XXII, n° 2, 2013
pp. 269–336

Introduction to the basics of Heegaard Floer homology

Bijan Sahamie[1]

Abstract. — This paper provides an introduction to the basics of Heegaard Floer homology with some emphasis on the hat theory and to the contact geometric invariants in the theory. The exposition is designed to be comprehensible to people without any prior knowledge of the subject.

Résumé. — Nous présentons une introduction aux éléments de la théorie d'Heegaard Floer et aux invariants qui proviennent d'une structure de contact. La présentation ne présuppose aucune connaissance du sujet.

Contents

(*) Reçu le 09/03/2012, accepté le 22/05/2012

(1) Mathematisches Institut der LMU München, Theresienstrasse 39, 80333 München
sahamie@math.lmu.de

Article proposé par Jean-Pierre Otal.

1. Introduction

Heegaard Floer homology was introduced by Peter Ozsváth and Zoltan Szabó at the beginning of the new millennium. Since then it developed very rapidly due to its various contributions to low-dimensional topology, particularly knot theory and contact geometry. The present paper is designed to give an introduction to **the basics** of Heegaard Floer theory with some emphasis on the hat theory. We try to provide all details necessary to communicate a complete and comprehensible picture. We would like to remark that there already are introductory articles to this subject (see [21], [22] and [23]). The difference between the existing articles and the present article is threefold: First of all we present a lot more details. We hope that these details will provide a complete picture of the basics of the theory. Our goal is to focus on those only which are relevant for the understanding of Heegaard Floer homology. Secondly, our exposition is not designed to present any applications and, in fact, we do not present any. Explaining applications to the reader would lead us too far away from the basics and would force us to make some compromise to the exposition. We felt that going into advanced elements would be disturbing to the goal of this paper. And thirdly, we have a slight contact geometric focus.

We think that the reader will profit the most from this paper when reading it completely rather than selecting a few elements: We start with a low-paced exposition and gain velocity as we move on. In this way we circumvent the creation of too many redundancies and it enables us to focus on the important facts at each stage of the paper. We expect the reader to have some knowledge about algebraic topology and surgery theory. As standard references we suggest [1] and [7].

In §2 and §3 we start with Heegaard diagrams and introduce everything necessary to construct the homology theory. We included a complete discussion of the invariance of Heegaard Floer theory (cf. §4) for two reasons. Firstly, the isomorphisms defined for showing invariance appear very frequently in the research literature. Secondly, the proof is based on constructions which can be called *the standard constructions of the theory.* Those who are impatient may just read §4.4 and skip the rest of §4. However, the remainder of the article refers to details of §4 several times. The following section, i.e. §5, is devoted to the knot theoretic variant of Heegaard Floer theory, called knot Floer homology. In §6 and §7 we outline how to assign to a 4-dimensional cobordism a map between the Floer homologies of the boundary components and derive the surgery exact triangle. This triangle is one of the most important tools, particularly for the contact geometric applications. Finally, the article focuses on preparing the reader to be able to understand the contact geometric applications as given – for instance – by Lisca and Stipsicz.

We are aware of the fact that there is a lot of material missing in this article. However, the presented theory provides a solid groundwork for understanding of what we omitted. We would like to outline at least some of the missing material. First of all the homology groups as well as the cobordism maps refine with respect to Spin^c-structures. We indicate this fact in §2 but do not outline any details. The standard reference is the article [17] of Ozsváth and Szabó. However, we suggest the reader first to familiarize with Spin^c-structures, especially with their interpretation as homology classes of vector fields (cf. [28]). Furthermore, there is an absolute $\mathbb{Q}$-grading on these homology groups (see [19]) and in case of knot Floer homologies for homologically trivial knots an additional $\mathbb{Z}$-grading (see [14]). Both gradings carry topological information and may appear as a help in explicit calculations, especially in combination with the surgery exact triangles. The knot Floer homologies admit additional exact sequences besides the surgery exact sequence. An example is the skein exact sequence (see [14] and [24]). For contact geometric applications the adjunction inequalities play a central role as they give a criterion for the vanishing of cobordism maps (see [18] or cf. [25]). Going a bit further, there are other flavors of Heegaard Floer homology: András Juhasz defined the so-called Sutured Floer homology of sutured manifolds (see [12]) and Ozsváth, Lipshitz and Thurston defined a variant of Heegaard Floer homology for manifolds with parameterized boundary (see [9]).

2. Introduction to the hat-version of Heegaard Floer homology

2.1. Heegaard Diagrams

One of the major results of Morse theory is the development of surgery and handle decompositions. Morse theory captures the manifold's topology in terms of a decomposition of it into topologically easy-to-understand pieces called **handles** (cf. [7]). In the case of closed 3-manifolds the handle decomposition can be assumed to be very symmetric. This symmetry allows us to describe the manifold's diffeomorphism type by a small amount of data. Heegaard diagrams are omnipresent in low-dimensional topology. Unfortunately there is no convention what precisely to call a Heegaard diagram; the definition of this notion underlies slight variations in different sources. Since Heegaard Floer Homology intentionally uses a non-efficient version of Heegaard diagrams, i.e. we fix more information than needed to describe the manifold's type, we briefly discuss, what is to be understood as Heegaard diagram throughout this article.

A brief summary of what we will discuss would be that we fix the data describing a handle decomposition relative to a splitting surface. Let Y be a closed oriented 3-manifold and $\Sigma \subset Y$ a **splitting surface**, i.e. a surface of genus g such that $Y \backslash \Sigma$ decomposes into two handlebodies H_0 and H_1. We fix a handle decomposition of $\overline{Y \backslash H_1}$ relative to this splitting surface Σ, i.e. there are 2-handles $h^2_{1,i}$, $i = 1, \ldots, g$, and a 3-handle h^3_1 such that (cf. [7])

$$Y \backslash H_1 \cong (\Sigma \times [0,1]) \cup_\partial (h^2_{1,1} \cup_\partial \ldots \cup_\partial h^2_{1,g} \cup_\partial h^3_1). \tag{2.1}$$

We can rebuild Y from this by gluing in 2-handles $h^2_{0,i}$, $i = 1, \ldots, g$, and a 3-handle h^3_0. Hence, Y can be written as

$$Y \cong (h^3_0 \cup_\partial h^2_{0,1} \cup_\partial \ldots \cup_\partial h^2_{0,g}) \cup_\partial (\Sigma \times [0,1]) \cup_\partial (h^2_{1,1} \cup_\partial \ldots \cup_\partial h^2_{1,g} \cup_\partial h^3_1). \tag{2.2}$$

Collecting the data from this decomposition we obtain a triple $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$ where Σ is the splitting surface of genus g, $\boldsymbol{\alpha} = \{\alpha_1, \ldots, \alpha_g\}$ are the images of the attaching circles of the $h^2_{0,i}$ interpreted as sitting in Σ and $\boldsymbol{\beta} = \{\beta_1, \ldots, \beta_g\}$ the images of the attaching circles of the 2-handles $h^2_{1,i}$ interpreted as sitting in Σ. This will be called a **Heegaard diagram of** Y. Observe that these data determine a Heegaard decomposition *in the classical sense* by dualizing the $h^2_{0,i}$. Dualizing a k-handle $D^k \times D^{3-k}$ means to reinterpret this object as $D^{3-k} \times D^k$. Both objects are diffeomorphic but observe that the former is a k-handle and the latter a $(3-k)$-handle. Observe that the α-curves are the co-cores of the 1-handles in the dualized picture, and that sliding $h^1_{0,i}$ over $h^1_{0,j}$ means, in the dual picture, that $h^2_{0,j}$ is slid over $h^2_{0,i}$.

2.2. Topology and Analytic Background

Given a closed oriented 3-manifold Y, we fix a Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$ of Y as defined in §2.1. We can associate to it the triple $(\mathrm{Sym}^g(\Sigma), \mathbb{T}_{\boldsymbol{\alpha}}, \mathbb{T}_{\boldsymbol{\beta}})$ which we will explain now:

By $\mathrm{Sym}^g(\Sigma)$ we denote the **g-fold symmetric product** of Σ, defined by taking the quotient under the canonical action of S_g on $\Sigma^{\times g}$, i.e.

$$\mathrm{Sym}^g(\Sigma) = \Sigma^{\times g}/S_g.$$

Although the action of S_g has fixed points, the symmetric product is a manifold. The local model is given by $\mathrm{Sym}^g(\mathbb{C})$ which itself can be identified with the set of normalized polynomials of degree g. An isomorphism is given by sending a point $[(p_1, \ldots, p_g)]$ to the normalized polynomial uniquely determined by the zero set $\{p_1, \ldots, p_g\}$. Denote by

$$\pi\colon \Sigma^{\times g} \longrightarrow \mathrm{Sym}^g(\Sigma)$$

the projection map. The attaching circles $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ define submanifolds

$$\mathbb{T}_{\boldsymbol{\alpha}} = \alpha_1 \times \ldots \times \alpha_g \quad \text{and} \quad \mathbb{T}_{\boldsymbol{\beta}} = \beta_1 \times \ldots \times \beta_g$$

in $\Sigma^{\times g}$. Obviously, the projection π embeds these into the symmetric product. In the following we will denote by $\mathbb{T}_{\boldsymbol{\alpha}}$ and $\mathbb{T}_{\boldsymbol{\beta}}$ the manifolds embedded into the symmetric product.

2.2.1. The chain complex

Let us start with a definition.

Definition 2.1. — *A map ϕ of the 2-disk D^2 (regarded as the unit 2-disk in $\mathbb{C}$) into the symmetric product $\mathrm{Sym}^g(\Sigma)$ is said to **connect** two points $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{T}_{\boldsymbol{\alpha}} \cap \mathbb{T}_{\boldsymbol{\beta}}$ if*

$$\begin{aligned} \phi(i) &= \boldsymbol{x}, \\ \phi(-i) &= \boldsymbol{y}, \\ \phi(\partial D^2 \cap \{z \in \mathbb{C} \,|\, Re(z) < 0\}) &\subset \mathbb{T}_{\boldsymbol{\alpha}}, \\ \phi(\partial D^2 \cap \{z \in \mathbb{C} \,|\, Re(z) > 0\}) &\subset \mathbb{T}_{\boldsymbol{\beta}}. \end{aligned}$$

*Continuous maps of the 2-disk into the symmetric product $\mathrm{Sym}^g(\Sigma)$ that connect two intersection points $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{T}_{\boldsymbol{\alpha}} \cap \mathbb{T}_{\boldsymbol{\beta}}$ are called **Whitney disks**. The set of homotopy classes of Whitney disks connecting $\boldsymbol{x}$ and $\boldsymbol{y}$ is denoted by $\pi_2(\boldsymbol{x}, \boldsymbol{y})$ in case $g > 2$.*

In case $g \leqslant 2$ we have to define the object $\pi_2(\boldsymbol{x},\boldsymbol{y})$ slightly different. However, we can always assume, without loss of generality, that $g > 2$ and, thus, we will omit discussing this case at all. We point the interested reader to [17, §2.4].

We fix a point $z \in \Sigma\backslash(\boldsymbol{\alpha}\cup\boldsymbol{\beta})$ and denote by $\mathcal{H}$ the quadruple $(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},z)$. This is called a pointed Heegaard diagram. In the remainder of this article we will use the letter $\mathcal{H}$ both for Heegaard diagrams and pointed Heegaard diagrams. Given a pointed Heegaard diagram, we define $\widehat{\mathrm{CF}}(\mathcal{H})$ as the free $\mathbb{Z}$-module (or $\mathbb{Z}_2$-module) generated by the intersection points $\mathbb{T}_{\boldsymbol{\alpha}} \cap \mathbb{T}_{\boldsymbol{\beta}}$ inside $\mathrm{Sym}^g(\Sigma)$. In case of $\mathbb{Z}_2$-coefficients, think of $\widehat{\mathrm{CF}}(\mathcal{H})$ as a $\mathbb{Z}_2$-vector space equipped with a canonical basis given by the elements $\mathbb{T}_{\boldsymbol{\alpha}} \cap \mathbb{T}_{\boldsymbol{\beta}}$. On this module, we can construct a differential

$$\widehat{\partial}_{\mathcal{H}}\colon \widehat{\mathrm{CF}}(\mathcal{H}) \longrightarrow \widehat{\mathrm{CF}}(\mathcal{H})$$

by defining it on the generators of $\widehat{\mathrm{CF}}(\mathcal{H})$. Given a point $\boldsymbol{x} \in \mathbb{T}_{\boldsymbol{\alpha}} \cap \mathbb{T}_{\boldsymbol{\beta}}$, we define $\widehat{\partial}_{\mathcal{H}}\boldsymbol{x}$ to be a linear combination

$$\widehat{\partial}_{\mathcal{H}}(\boldsymbol{x}) = \sum_{\boldsymbol{y}\in\mathbb{T}_{\boldsymbol{\alpha}}\cap\mathbb{T}_{\boldsymbol{\beta}}} \widehat{\partial}_{\mathcal{H}}\boldsymbol{x}\Big|_{\boldsymbol{y}} \cdot \boldsymbol{y}.$$

The definition of the coefficients will occupy the remainder of this section. The idea resembles other Floer homology theories. The goal is to define $\widehat{\partial}_{\mathcal{H}}\boldsymbol{x}\Big|_{\boldsymbol{y}}$ as a signed count of holomorphic Whitney disks connecting $\boldsymbol{x}$ and $\boldsymbol{y}$ which are rigid up to reparametrization. First we have to introduce almost complex structures into this picture. A more detailed discussion of these will be given in §2.3. For the moment it will be sufficient to say that we choose a generic path $(\mathcal{J}_s)_{s\in[0,1]}$ of almost complex structures on the symmetric product. Identifying the unit disk, after taking out the points $\pm i$, in $\mathbb{C}$ with $[0,1]\times\mathbb{R}$ we define ϕ to be **holomorphic** if it satisfies for all $(s,t) \in [0,1]\times\mathbb{R}$ the equation

$$\frac{\partial\phi}{\partial s}(s,t) + \mathcal{J}_s\Big(\frac{\partial\phi}{\partial t}(s,t)\Big) = 0. \tag{2.3}$$

Looking into (2.3) it is easy to see that a holomorphic Whitney disk ϕ can be reparametrized by a constant shift in $\mathbb{R}$-direction without violating (2.3).

Definition 2.2. — *Given two points $\boldsymbol{x},\boldsymbol{y} \in \mathbb{T}_{\boldsymbol{\alpha}} \cap \mathbb{T}_{\boldsymbol{\beta}}$, we denote by $\mathcal{M}_{\mathcal{J}_s}(\boldsymbol{x},\boldsymbol{y})$ the set of holomorphic Whitney disks connecting $\boldsymbol{x}$ and $\boldsymbol{y}$. We call this set* **moduli space of holomorphic Whitney disks** *connecting $\boldsymbol{x}$ and $\boldsymbol{y}$. Given a homotopy class $[\phi] \in \pi_2(\boldsymbol{x},\boldsymbol{y})$, denote by $\mathcal{M}_{\mathcal{J}_s,[\phi]}$ the space of holomorphic representatives in the homotopy class of ϕ.*

In the following the generic path of almost complex structures will not be important and thus we will suppress it from the notation. Since the path is chosen generically (cf. §2.3 or see [17]) the moduli spaces are manifolds. The constant shift in $\mathbb{R}$-direction induces a free $\mathbb{R}$-action on the moduli spaces. Thus, if $\mathcal{M}_{[\phi]}$ is non-empty its dimension is greater than zero. We take the quotient of $\mathcal{M}_{[\phi]}$ under the $\mathbb{R}$-action and denote the resulting spaces by

$$\widehat{\mathcal{M}}_{[\phi]} = \mathcal{M}_{[\phi]}/\mathbb{R} \text{ and } \widehat{\mathcal{M}}(\boldsymbol{x},\boldsymbol{y}) = \mathcal{M}(\boldsymbol{x},\boldsymbol{y})/\mathbb{R}.$$

The so-called **signed count** of 0-dimensional components of $\widehat{\mathcal{M}}(\boldsymbol{x},\boldsymbol{y})$ means in case of $\mathbb{Z}_2$-coefficients simply to count mod 2. In case of $\mathbb{Z}$-coefficients we have to introduce **coherent orientations** on the moduli spaces. We will roughly sketch this process in the following.

Obviously, in case of $\mathbb{Z}$-coefficients we cannot simply count the 0-dimensional components of $\widehat{\mathcal{M}}(\boldsymbol{x},\boldsymbol{y})$. The defined morphism would not be a differential. To circumvent this problem we have to introduce signs appropriately attached to each component. The 0-dimensional components of $\widehat{\mathcal{M}}(\boldsymbol{x},\boldsymbol{y})$ correspond to the 1-dimensional components of $\mathcal{M}(\boldsymbol{x},\boldsymbol{y})$. Each of these components carries a canonical orientation induced by the free $\mathbb{R}$-action given by constant shifts. We introduce orientations on these components. Comparing the artificial orientations with the canonical shifting orientation we can associate to each component, i.e. each element in $\widehat{\mathcal{M}}(\boldsymbol{x},\boldsymbol{y})$, a sign. The signed count will respect the signs attached. There is a technical condition called **coherence** (see [17, Definition 3.11] or cf. §2.3) one has to impose on the orientations. This technical condition ensures that the morphism $\widehat{\partial}_{\mathcal{H}}$ is a differential.

The chosen point $z \in \Sigma\backslash(\boldsymbol{\alpha}\cup\boldsymbol{\beta})$ will be part of the definition. The path $(\mathcal{J}_s)_{s\in[0,1]}$ is chosen in such a way that

$$V_z = \{z\} \times \mathrm{Sym}^{g-1}(\Sigma) \hookrightarrow \mathrm{Sym}^g(\Sigma)$$

is a complex submanifold. For a Whitney disk (or its homotopy class) ϕ define $n_z(\phi)$ as the intersection number of ϕ with the submanifold V_z. We define

$$\widehat{\partial}_{\mathcal{H}}\boldsymbol{x}\Big|_{\boldsymbol{y}} = \#\widehat{\mathcal{M}}(\boldsymbol{x},\boldsymbol{y})^0_{n_z=0},$$

i.e. the signed count of the 0-dimensional components of the unparametrized moduli spaces of holomorphic Whitney disks connecting $\boldsymbol{x}$ and $\boldsymbol{y}$ with the property that their intersection number n_z is trivial.

THEOREM 2.3 (see §4 of [17]). — *The assignment $\widehat{\partial}_{\mathcal{H}}$ is well-defined.*

Theorem 2.4 (see Theorem 4.15 of [17]). — *The morphism $\widehat{\partial}_{\mathcal{H}}$ is a differential.*

We will give sketches of the proofs of the last two theorems later in §2.3. At the moment we do not know enough about Whitney disks and the symmetric product to prove it.

Definition 2.5. — *Denote by $\widehat{\mathrm{HF}}(\mathcal{H})$ the homology theory $H_*(\widehat{\mathrm{CF}}(\mathcal{H}),\widehat{\partial}_{\mathcal{H}})$. If the pointed Heegaard diagram $\mathcal{H}$ is the Heegaard diagram of a manifold Y, then we also write $\widehat{\mathrm{HF}}(Y)$ for $\widehat{\mathrm{HF}}(\mathcal{H})$.*

The notation should indicate that the homology theory does not depend on the chosen data. It is a topological invariant of the manifold Y, although this is not the whole story. The theory depends on the choice of coherent system of orientations. For a manifold Y there are $2^{b_1(Y)}$ numbers of non-equivalent systems of coherent orientations. The resulting homologies can differ (see Example 2.3). Nevertheless the orientations are not written down. We guess there are two reasons: The first is that most of the time it is not really important which system of coherent orientations is chosen. All reasonable constructions will work for every coherent orientation system and in case there is a specific choice needed this will be explicitly stated. The second reason is that it is possible to give a *convention* for the choice of coherent orientation systems. Since we have not developed the mathematics to state the convention precisely we point the reader to Theorem 2.31.

2.2.2. On Holomorphic disks in the Symmetric Product

In order to be able to discuss a first example we briefly introduce some properties of the symmetric product.

Definition 2.6. — *For a Whitney disk ϕ we denote by $\mu(\phi)$ the* **formal dimension** *of $\mathcal{M}_\phi$. We also call $\mu(\phi)$ the* **Maslov index** *of ϕ.*

For the readers that have not heard anything about Floer homology at all, just think of $\mu(\phi)$ as the dimension of the space $\mathcal{M}_\phi$, although even in case $\mathcal{M}_\phi$ is not a manifold the number $\mu(\phi)$ is defined (cf. §2.3). Just to give some intuition, note that the moduli spaces are the zero-set of a section, S say, in a Banach bundle one associates to the given setup. The linearization of this section at the zero set is a Fredholm operator. Those operators carry a property called Fredholm index (cf. Definition 2.21). The number μ is the Fredholm index of that operator. Even if the moduli spaces are no manifolds this number is defined. It is called formal dimension or **expected dimension** since in case the section S intersects the zero-section of the Banach-bundle transversely (and hence the moduli spaces are manifolds) the

Fredholm index μ equals the dimension of the moduli spaces. So, negative indices are possible and make sense in some situations. One can think of negative indices as the number of missing degrees of freedom to give a manifold.

LEMMA 2.7. — *In case $g(\Sigma) > 2$ the 2nd homotopy group $\pi_2(\mathrm{Sym}^g(\Sigma))$ is isomorphic to $\mathbb{Z}$. It is generated by an element S with $\mu(S) = 2$ and $n_z(S) = 1$, where n_z is defined the same way as it was defined for Whitney disks.*

Let $\eta\colon \Sigma \longrightarrow \Sigma$ be an involution such that Σ/η is a sphere. The map

$$\mathbb{S}^2 \longrightarrow \mathrm{Sym}^g(\Sigma),\ \ y \longmapsto \{(y, \eta(y), y, \dots, y)\}$$

is a representative of S. Using this representative it is easy to see that $n_z(S) = 1$. It is a property of μ as an index that it behaves additive under concatenation. Indeed the intersection number n_z behaves additive, too. To develop some intuition for the holomorphic spheres in the symmetric product we state the following result from [17].

LEMMA 2.8 (see Proposition 2.15 of [17]). — *There is an exact sequence*

$$0 \longrightarrow \pi_2(\mathrm{Sym}^g(\Sigma)) \longrightarrow \pi_2(\boldsymbol{x}, \boldsymbol{x}) \longrightarrow \ker(n_z) \longrightarrow 0.$$

The map n_z provides a splitting for the sequence.

Observe that we can interpret a Withney disk in $\pi_2(\boldsymbol{x}, \boldsymbol{x})$ as a family of paths in $\mathrm{Sym}^g(\Sigma)$ based at the constant path $\boldsymbol{x}$. We can also interpret an element in $\pi_2(\mathrm{Sym}^g(\Sigma))$ as a family of paths in $\mathrm{Sym}^g(\Sigma)$ based at the constant path $\boldsymbol{x}$. Interpreted in this way there is a natural map from $\pi_2(\mathrm{Sym}^g(\Sigma))$ into $\pi_2(\boldsymbol{x}, \boldsymbol{x})$. The map n_z provides a splitting for the sequence as it may be used to define the map

$$\pi_2(\boldsymbol{x}, \boldsymbol{x}) \longrightarrow \pi_2(\mathrm{Sym}^g(\Sigma))$$

sending a Whitney disk ϕ to $n_z(\phi) \cdot S$. This obviously defines a splitting for the sequence (cf. Lemma 2.7).

LEMMA 2.9 (cf. Proposition 2.15 of [17]). — *The kernel of n_z interpreted as a map on $\pi_2(\boldsymbol{x}, \boldsymbol{x})$ is isomorphic to $H^1(Y; \mathbb{Z})$.*

With the help of concatenation we are able to define an action

$$*\colon \pi_2(\boldsymbol{x}, \boldsymbol{x}) \times \pi_2(\boldsymbol{x}, \boldsymbol{y}) \longrightarrow \pi_2(\boldsymbol{x}, \boldsymbol{y}),$$

which is obviously free and transitive. Thus, we have an identification

$$\begin{tikzcd} \pi_2(\boldsymbol{x},\boldsymbol{y}) \arrow[rr, "\cong"] \arrow[dr] & & \pi_2(\boldsymbol{x},\boldsymbol{x}) \arrow[dl] \\ & \{*\} & \end{tikzcd} \tag{2.4}$$

as principal bundles over a one-point space, which is another way of saying that the concatenation action endows $\pi_2(\boldsymbol{x},\boldsymbol{y})$ with a group structure after fixing a unit element in $\pi_2(\boldsymbol{x},\boldsymbol{y})$. To address the well-definedness of $\widehat{\partial}_{\mathcal{H}}$ we have to show that the sum in the definition of $\widehat{\partial}_{\mathcal{H}}$ is finite. For the moment let us assume that for a generic choice of path $(\mathcal{J}_s)_{s\in[0,1]}$ the moduli spaces $\widehat{\mathcal{M}}_\phi$ with $\mu(\phi)=1$ are compact manifolds (cf. Theorem 2.22), hence their signed count is finite. Assuming this property we are able to show well-definedness of $\widehat{\partial}_{\mathcal{H}}$ in case Y is a homology sphere.

Proof of Theorem 2.3 for $b_1(Y)=0$. — Observe that

$$\widehat{\mathcal{M}}(\boldsymbol{x},\boldsymbol{y})^0_{n_z=0} = \bigsqcup_{\phi\in H(\boldsymbol{x},\boldsymbol{y},1)} \widehat{\mathcal{M}}_\phi, \tag{2.5}$$

where $H(\boldsymbol{x},\boldsymbol{y},1)\subset\pi_2(\boldsymbol{x},\boldsymbol{y})$ is the subset of homotopy classes admitting holomorphic representatives with $\mu(\phi)=1$ and $n_z=0$. We have to show that $H(\boldsymbol{x},\boldsymbol{y},1)$ is a finite set. Since $b_1(Y)=0$ the cohomology $H^1(Y;\mathbb{Z})$ vanishes. By our preliminary discussion, given a reference disk $\phi_0\in\pi_2(\boldsymbol{x},\boldsymbol{y})$, every $\phi_{\boldsymbol{x}\boldsymbol{y}}\in\pi_2(\boldsymbol{x},\boldsymbol{y})$ can be written as a concatenation $\phi_{\boldsymbol{x}\boldsymbol{y}}=\phi*\phi_0$, where ϕ is an element in $\pi_2(\boldsymbol{x},\boldsymbol{x})$. Since we are looking for disks with index one we have to find all $\phi\in\pi_2(\boldsymbol{x},\boldsymbol{x})$ satisfying the property $\mu(\phi)=1-\mu(\phi_0)$. Recall that Y is a homology sphere and thus $\pi_2(\boldsymbol{x},\boldsymbol{x})\cong\mathbb{Z}\langle S\rangle$. Hence, the disk ϕ is described by an integer $k\in\mathbb{Z}$, i.e. $\phi=k\cdot S$. The property $\mu(S)=2$ tells us that

$$1-\mu(\phi_0)=\mu(\phi)=\mu(k\cdot S)=k\cdot\mu(S)=2k.$$

There is at most one $k\in\mathbb{Z}$ satisfying this equation, so there is at most one homotopy class of Whitney disks satisfying the property $\mu=1$ and $n_z=0$. □

In case Y has non-trivial first cohomology we need an additional condition to make the proof work. The given argument obviously breaks down in this case. To fix this we impose a topological/algebraic condition on the Heegaard diagram. Before we can define these *admissibility* properties (see Definition 2.17) we have to go into the theory a bit more.

There is an obstruction to finding Whitney disks connecting two given intersection points $\boldsymbol{x},\boldsymbol{y}$. The two points $\boldsymbol{x}$ and $\boldsymbol{y}$ can certainly be connected via paths inside $\mathbb{T}_{\boldsymbol{\alpha}}$ and $\mathbb{T}_{\boldsymbol{\beta}}$. Fix two paths $a\colon I \longrightarrow \mathbb{T}_{\boldsymbol{\alpha}}$ and $b\colon I \longrightarrow \mathbb{T}_{\boldsymbol{\beta}}$ such that $-\partial b = \partial a = \boldsymbol{y} - \boldsymbol{x}$. This is the same as saying we fix a closed curve $\gamma_{\boldsymbol{xy}}$ based at $\boldsymbol{x}$, going to $\boldsymbol{y}$ along $\mathbb{T}_{\boldsymbol{\alpha}}$, and moving back to $\boldsymbol{x}$ along $\mathbb{T}_{\boldsymbol{\beta}}$. Obviously $\gamma_{\boldsymbol{xy}} = b + a$. Is it possible to extend the curve $\gamma_{\boldsymbol{xy}}$, after possibly homotoping it a bit, to a disk? If so, this would be a Whitney disk. Thus, finding an obstruction can be reformulated as: Is $[\gamma_{\boldsymbol{xy}}] = 0 \in \pi_1(\mathrm{Sym}^g(\Sigma))$?

Lemma 2.10 (see Lemma 2.6 of [17]). — *The group $\pi_1(\mathrm{Sym}^g(\Sigma))$ is abelian.*

Given a closed curve $\gamma \subset \mathrm{Sym}^g(\Sigma)$ in general position (i.e. not meeting the diagonal of $\mathrm{Sym}^g(\Sigma)$), we can lift this curve to

$$(\gamma_1, \ldots, \gamma_g)\colon \mathbb{S}^1 \longrightarrow \Sigma^{\times g}.$$

Projection onto each factor Σ defines a 1-cycle. We define

$$\Phi(\gamma) = \gamma_1 + \ldots + \gamma_g.$$

Lemma 2.11 (see Lemma 2.6 of [17]). — *The map Φ induces an isomorphism*

$$\Phi_*\colon H_1(\mathrm{Sym}^g(\Sigma)) \longrightarrow H_1(\Sigma; \mathbb{Z}).$$

By surgery theory (see [7, p. 111]) we know that

$$\frac{H_1(\Sigma; \mathbb{Z})}{[\alpha_1], \ldots, [\alpha_g], [\beta_1], \ldots, [\beta_g]} \cong H_1(Y; \mathbb{Z}) \tag{2.6}$$

The curve $\gamma_{\boldsymbol{xy}}$ is homotopically trivial in the symmetric product if and only if $\Phi_*([\gamma_{\boldsymbol{xy}}])$ is trivial. If we pick different curves a and b to define another curve $\eta_{\boldsymbol{xy}}$, the difference

$$\Phi(\gamma_{\boldsymbol{xy}}) - \Phi(\eta_{\boldsymbol{xy}})$$

is a sum of α-and β-curves. Thus, interpreted as a cycle in $H_1(Y; \mathbb{Z})$, the class

$$[\Phi(\gamma_{\boldsymbol{xy}})] \in H_1(Y; \mathbb{Z})$$

does not depend on the choices made in its definition. We get a map

$$\begin{array}{rccc} \epsilon\colon & (\mathbb{T}_{\boldsymbol{\alpha}} \cap \mathbb{T}_{\boldsymbol{\beta}})^{\times 2} & \longrightarrow & H_1(Y; \mathbb{Z}) \\ & (\boldsymbol{x}, \boldsymbol{y}) & \longmapsto & [\Phi(\gamma_{\boldsymbol{xy}})]_{H_1(Y;\mathbb{Z})} \end{array}$$

with the following property.

Lemma 2.12. — *If $\epsilon(\boldsymbol{x}, \boldsymbol{y})$ is non-zero the set $\pi_2(\boldsymbol{x}, \boldsymbol{y})$ is empty.*

Proof. — Suppose there is a connecting disk ϕ then with $\gamma_{\boldsymbol{xy}} = \partial(\phi(D^2))$ we have

$$\epsilon(\boldsymbol{x},\boldsymbol{y}) = [\Phi(\gamma_{\boldsymbol{xy}})]_{H_1(Y;\mathbb{Z})} = \frac{\Phi_*\big([\gamma_{\boldsymbol{xy}}]_{H_1(\mathrm{Sym}^g(\Sigma))}\big)}{[\alpha_1],\ldots,[\alpha_g],[\beta_1],\ldots,\beta_g]} = 0$$

since $[\gamma_{\boldsymbol{xy}}]_{\pi_1(\mathrm{Sym}^g(\Sigma))} = 0$. $\square$

As a consequence we can split up the chain complex $\widehat{\mathrm{CF}}(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},z)$ into subcomplexes. It is important to notice that there is a map

$$\mathfrak{s}_z\colon \mathbb{T}_{\boldsymbol{\alpha}} \cap \mathbb{T}_{\boldsymbol{\beta}} \longrightarrow \mathrm{Spin}^c_3(Y) \cong H^2(Y;\mathbb{Z}), \tag{2.7}$$

such that $\mathrm{PD}(\epsilon(\boldsymbol{x},\boldsymbol{y})) = \mathfrak{s}_z(\boldsymbol{x}) - \mathfrak{s}_z(\boldsymbol{y})$. For a definition of the map $\mathfrak{s}_z$ we point the reader to [17, §2.6]. Thus, fixing a Spin^c-structure $\mathfrak{s}$, the $\mathbb{Z}$-module (or $\mathbb{Z}_2$-module) $\widehat{\mathrm{CF}}(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},z;\mathfrak{s})$ generated by $(\mathfrak{s}_z)^{-1}(\mathfrak{s})$ defines a subcomplex of $\widehat{\mathrm{CF}}(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},z)$. The associated homology is denoted by $\widehat{\mathrm{HF}}(Y,\mathfrak{s})$, and it is a submodule of $\widehat{\mathrm{HF}}(Y)$. Especially note that

$$\widehat{\mathrm{HF}}(Y) = \bigoplus_{\mathfrak{s}\in\mathrm{Spin}^c_3(Y)} \widehat{\mathrm{HF}}(Y,\mathfrak{s}).$$

Since $\mathbb{T}_{\boldsymbol{\alpha}} \cap \mathbb{T}_{\boldsymbol{\beta}}$ consists of finitely many points, there are just finitely many groups in this splitting which are non-zero. In general this splitting will depend on the choice of base-point. If z is chosen in a different component of $\Sigma\backslash\{\boldsymbol{\alpha}\cup\boldsymbol{\beta}\}$ there will be a difference between the Spin^c-structure associated to an intersection point. For details we point to [17, Lemma 2.19].

Example 2.1. — The Heegaard diagram given by the data $(T^2,\{\mu\},\{\lambda\})$ (cf. §2.1) is the 3-sphere. To make use of Lemma 2.7 we add two stabilizations to get a Heegaard surface of genus 3, i.e.

$$D = (T^2\#T^2\#T^2, \{\mu_1,\mu_2,\mu_3\},\{\lambda_1,\lambda_2,\lambda_3\}),$$

where μ_i are meridians of the tori, and λ_i are longitudes. The complement of the attaching curves is connected. Thus, we can arbitrarily choose the base point z. Denote by $\mathcal{H}$ the associated pointed Heegaard diagram, The chain complex $\widehat{\mathrm{CF}}(\mathcal{H})$ equals one copy of $\mathbb{Z}$ since it is generated by one single intersection point which we denote by $\boldsymbol{x}$. We claim that $\widehat{\partial}_{\mathcal{H}}\boldsymbol{x} = 0$. Denote by $[\phi]$ a homotopy class of Whitney disks connecting $\boldsymbol{x}$ with itself. This is a holomorphic sphere which can be seen with Lemma 2.8, Lemma 2.9 and the fact that $H^1(\mathbb{S}^3) = 0$. By Lemma 2.7 the set $\pi_2(\mathrm{Sym}^g(\Sigma))$ is generated by S with the property $n_z(S) = 1$. The additivity of n_z under concatenation shows that $[\phi]$ is a trivial holomorphic sphere and $\mu([\phi]) = 0$.

Thus, the space $\mathcal{M}(\boldsymbol{x},\boldsymbol{x})^1_{n_z=0}$, i.e. the space of holomorphic Whitney disks connecting $\boldsymbol{x}$ with itself, with $\mu = 1$ and $n_z = 0$, is empty. Hence,

$$\widehat{\mathrm{HF}}(\mathbb{S}^3) \cong \mathbb{Z}.$$

2.2.3. A Low-Dimensional Model for Whitney disks

The exact sequence in Lemma 2.8 combined with Lemma 2.9 and (2.4) gives an interpretation of Whitney disks as homology classes. Given a disk ϕ, we define its associated homology class by $\mathcal{H}(\phi)$, i.e.

$$0 \longrightarrow \pi_2(\mathrm{Sym}^g(\Sigma)) \longrightarrow \pi_2(\boldsymbol{x},\boldsymbol{x}) \xrightarrow{\mathcal{H}} H_2(Y;\mathbb{Z}) \longrightarrow 0. \tag{2.8}$$

In the following we intend to give a description of the map $\mathcal{H}$. Given a Whitney disk ϕ, we can lift this disk to a map $\widetilde{\phi}$ by pulling back the branched covering π given in Diagram (2.9).

$$\begin{tikzcd}
F/S_{g-1} = \widehat{D}^2 \arrow[r, "\overline{\phi}"] \arrow[rrr, bend left, "\phi"] & \Sigma & \mathrm{Sym}^{g-1}(\Sigma) \arrow[r] & \Sigma \\
\phi^*\Sigma^{\times g} = F \arrow[u] \arrow[d] \arrow[rr, "\widetilde{\phi}"] & & \Sigma^{\times g} \arrow[u] \arrow[d, "\pi"] & \\
D^2 \arrow[rr, "\phi"] & & \mathrm{Sym}^g(\Sigma) &
\end{tikzcd} \tag{2.9}$$

Let $S_{g-1} \subset S_g$ be the subgroup of permutations fixing the first component. Modding out S_{g-1} we obtain the map $\overline{\phi}$ pictured in (2.9). Composing it with the projection onto the surface Σ we define a map

$$\widehat{\phi}\colon \widehat{D^2} \longrightarrow \Sigma.$$

The image of this map $\widehat{\phi}$ defines what is called a domain.

Definition 2.13. — *Denote by $\mathcal{D}_1, \ldots, \mathcal{D}_m$ the closures of the components of the complement of the attaching circles $\Sigma\backslash\{\boldsymbol{\alpha} \cup \boldsymbol{\beta}\}$. Fix one point z_i in each component. A* ***domain*** *is a linear combination*

$$\mathcal{A} = \sum_{i=1}^{m} \lambda_i \cdot \mathcal{D}_i$$

with $\lambda_1, \ldots, \lambda_m \in \mathbb{Z}$.

For a Whitney disk ϕ we define its **associated domain** by

$$\mathcal{D}(\phi) = \sum_{i=1}^{m} n_{z_i}(\phi) \cdot \mathcal{D}_i.$$

The map $\widehat{\phi}$ and $\mathcal{D}(\phi)$ are related by the equation

$$\widehat{\phi}(\widehat{D^2}) = \mathcal{D}(\phi)$$

as chains in Σ relative to the set $\boldsymbol{\alpha} \cup \boldsymbol{\beta}$. We define $\mathcal{H}(\phi)$ as the associated homology class of $\widehat{\phi}_*[\widehat{D^2}]$ in $H_2(Y;\mathbb{Z})$. The correspondence is given by closing up the boundary components by using the core disks of the 2-handles represented by the α-curves and the β-curves.

Lemma 2.14. — *Two Whitney disks $\phi_1, \phi_2 \in \pi_2(\boldsymbol{x}, \boldsymbol{x})$ are homotopic if and only if their domains are equal.*

Proof. — Given two disks ϕ_1, ϕ_2 whose domains are equal, by definition $\mathcal{H}(\phi_1) = \mathcal{H}(\phi_2)$. By (2.8) they can only differ by a holomorphic sphere, i.e. $\phi_1 = \phi_2 + k \cdot S$. The equality $\mathcal{D}(\phi_1) = \mathcal{D}(\phi_2)$ implies that $n_z(\phi_1) = n_z(\phi_2)$. The equation

$$0 = n_z(\phi_2) - n_z(\phi_1) = n_z(\phi_2) - n_z(\phi_2 + k \cdot S) = 2k$$

forces k to vanish. □

The interpretation of Whitney disks as domains is very useful in computations, as it provides a low-dimensional model. The symmetric product is $2g$-dimensional, thus an investigation of holomorphic disks is very inconvenient. However, not all domains are carried by holomorphic disks. Obviously, the equality $[\mathcal{D}(\phi)] = \widehat{\phi}_*[\widehat{D^2}]$ connects the boundary conditions imposed on Whitney disks to boundary conditions of the domains. It is not hard to observe that the definition of $\widehat{\phi}$ follows the same lines as the construction of the isomorphism Φ_* of homology groups discussed earlier (cf. Lemma 2.11). Suppose we have fixed two intersections $\boldsymbol{x} = \{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_g\}$ and $\boldsymbol{y} = \{\boldsymbol{y}_1, \ldots, \boldsymbol{y}_g\}$ connected by a Whitney disk ϕ. The boundary $\partial(\phi(D^2))$ defines a connecting curve $\gamma_{\boldsymbol{xy}}$. It is easy to see that

$$\mathrm{im}(\widehat{\phi}\Big|_{\partial\widehat{D^2}}) = \Phi(\gamma_{\boldsymbol{xy}}) = \gamma_1 + \ldots + \gamma_g.$$

Restricting the γ_i to the α-curves we get a chain connecting the set $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_g$ with $\boldsymbol{y}_1, \ldots, \boldsymbol{y}_g$, and restricting the γ_i to the β-curves we get a chain connecting the set $\boldsymbol{y}_1, \ldots, \boldsymbol{y}_g$ with $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_g$. This means each boundary component of $\widehat{D^2}$ consists of a set of arcs alternating through α-curves and β-curves.

Definition 2.15. — *A domain is called* **periodic** *if its boundary is a sum of α-and β-curves and $n_z(\mathcal{D}) = 0$, i.e. the multiplicity of $\mathcal{D}$ at the domain $\mathcal{D}_z$ containing z vanishes.*

Of course a Whitney disk is called **periodic** if its associated domain is a periodic domain. The subgroup of periodic classes in $\pi_2(\boldsymbol{x}, \boldsymbol{x})$ is denoted by $\Pi_{\boldsymbol{x}}$.

Theorem 2.16 (see Theorem 4.9 of [17]). — *For a* Spin^c*-structure s and a periodic class $\phi \in \Pi_{\boldsymbol{x}}$ we have the equality $\mu(\phi) = \langle c_1(s), \mathcal{H}(\phi)\rangle$.*

This is a deep result connecting the expected dimension of a periodic disk with a topological property. Note that, because of the additivity of the expected dimension μ, the homology groups $\widehat{\mathrm{HF}}(Y, \mathfrak{s})$ can be endowed with a relative grading defined by

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

where ϕ is an arbitrary element of $\pi_2(\boldsymbol{x}, \boldsymbol{y})$. In the case of homology spheres this defines a relative $\mathbb{Z}$-grading because by Theorem 2.16 the expected dimension vanishes for all periodic disks. In case of non-trivial homology they just vanish modulo $\delta(s)$, where

$$\delta(\mathfrak{s}) = \gcd_{A \in H_2(Y;\mathbb{Z})} \langle c_1(\mathfrak{s}), A\rangle ,$$

i.e. it defines a relative $\mathbb{Z}_{\delta(\mathfrak{s})}$-grading.

Example 2.2. — Here, we reconsider the situation of Example 2.1. To see that $\widehat{\partial}_{\mathcal{H}}\boldsymbol{x} = 0$, observe that the differential $\widehat{\partial}_{\mathcal{H}}$ decreases the grading by 1: Suppose there is a holomorphic disk ϕ of index $\mu(\phi) = 1$ and $n_z(\phi) = 0$ that connectes $\boldsymbol{x}$ with some point $\boldsymbol{y}$, then

$$\mathrm{gr}(\boldsymbol{x}, \boldsymbol{x}) = \mu(\phi) - n_z(\phi) = 1 - 0 = 1.$$

So, in our case $\widehat{\partial}_{\mathcal{H}}\boldsymbol{x} = 0$ since $\mathrm{gr}(\boldsymbol{x}, \boldsymbol{x}) = 0$ and $\boldsymbol{x}$ is the only generator.

Definition 2.17. — *A pointed Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, z)$ is called* **weakly admissible** *for the* Spin^c*-structure $\mathfrak{s}$ if for every non-trivial periodic domain $\mathcal{D}$ such that $\langle c_1(\mathfrak{s}), \mathcal{H}(\mathcal{D})\rangle = 0$ the domain has positive and negative coefficients.*

With this technical condition imposed $\widehat{\partial}_{\mathcal{H}}$ is a well-defined map on the subcomplex $\widehat{\mathrm{CF}}(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s})$. From admissibility it follows that for every $\boldsymbol{x}, \boldsymbol{y} \in (\mathfrak{s}_z)^{-1}(\mathfrak{s})$ and $j, k \in \mathbb{Z}$ there exists just a finite number of $\phi \in \pi_2(\boldsymbol{x}, \boldsymbol{y})$

with $\mu(\phi) = j$, $n_z(\phi) = k$ and $\mathcal{D}(\phi) \geqslant 0$. The last condition means that all coefficients in the associated domain are greater or equal to zero.

Proof of Theorem 2.3 for $b_1(Y) \neq 0$. — Recall that holomorphic disks are either contained in a complex submanifold C or they intersect C always transversely and always positive. The definition of the path $(\mathcal{J}_s)_{s\in[0,1]}$ (cf. §2.3) includes that all the $\{z_i\} \times \mathrm{Sym}^{g-1}(\Sigma)$ are complex submanifolds. Thus, holomorphic Whitney disks always satisfy $\mathcal{D}(\phi) \geqslant 0$. □

We close this section with a statement that appears to be useful for developing intuition for Whitney disks. It helps imagining the strong connection between the disks and their associated domains.

Theorem 2.18 (see Lemma 2.17 of [17]). — *Consider a domain $\mathcal{D}$ whose coefficients are all greater than or equal to zero. There exists an oriented 2-manifold S with boundary and a map $\phi\colon S \longrightarrow \Sigma$ with $\phi(S) = \mathcal{D}$ with the property that ϕ is nowhere orientation-reversing and the restriction of ϕ to each boundary component of S is a diffeomorphism onto its image.*

2.3. The Structure of the Moduli Spaces

The material in this section is presented without any details. The exposition pictures the bird's eye view of the material. Recall from the last sections that we have to choose a path of almost complex structures appropriately to define Heegaard Floer theory. So, a discussion of these structures is inevitable. However, a lot of improvements have been made the last years and we intend to mention some of them.

Let $(\mathfrak{j}, \eta)$ be a Kähler structure on the Heegaard surface Σ, i.e. η is a symplectic form and $\mathfrak{j}$ an almost-complex structure that tames η. Let $z_1, \ldots, z_m$ be points, one in each component of $\Sigma\backslash\{\boldsymbol{\alpha} \cup \boldsymbol{\beta}\}$. Denote by V an open neighborhood in $\mathrm{Sym}^g(\Sigma)$ of

$$D \cup \big(\bigcup_{i=1}^{m} \{z_i\} \times \mathrm{Sym}^{g-1}(\Sigma)\big),$$

where D is the diagonal in $\mathrm{Sym}^g(\Sigma)$.

Definition 2.19. — *An almost complex structure $\mathcal{J}$ on $\mathrm{Sym}^g(\Sigma)$ is called $(\mathfrak{j}, \eta, V)$-**nearly symmetric** if $\mathcal{J}$ agrees with $sym^g(\mathfrak{j})$ over V and if $\mathcal{J}$ tames $\pi_*(\eta^{\times g})$ over $\overline{V}^c$. The set of $(\mathfrak{j}, \eta, V)$-nearly symmetric almost-complex structures will be denoted by $\mathcal{J}(\mathfrak{j}, \eta, V)$.*

The almost complex structure $sym^g(\mathfrak{j})$ on $\mathrm{Sym}^g(\Sigma)$ is the natural almost complex structure induced by the structure $\mathfrak{j}$. Important for us is that the structure $\mathcal{J}$ agrees with $sym^g(\mathrm{j})$ on V. This makes the $\{z_i\} \times \mathrm{Sym}^{g-1}(\Sigma)$ complex submanifolds with respect to $\mathcal{J}$. This is necessary to guarantee positive intersections with Whitney disks. Without this property the proof of Theorem 2.3 would break down in the case the manifold has non-trivial topology.

We are interested in holomorphic Whitney disks, i.e. disks in the symmetric product which are solutions of (2.3). Denote by $\partial_{\mathcal{J}_s}$ the Cauchy-Riemann operator defined by equation (2.3) (see [17, p. 1050]). Define $\mathcal{B}(\boldsymbol{x}, \boldsymbol{y})$ as the space of Whitney disks connecting $\boldsymbol{x}$ and $\boldsymbol{y}$ such that the disks converge to $\boldsymbol{x}$ and $\boldsymbol{y}$ exponentially with respect to some Sobolev space norm in a neighborhood of i and $-i$ (see [17, §3.2], especially [17, p. 1049]). With these assumptions the solution $\partial_{\mathcal{J}_s}\phi$ lies in a space of L^p-sections

$$L^p([0,1] \times \mathbb{R}, \phi^*(T\mathrm{Sym}^g(\Sigma))).$$

These fit together to form a bundle $\mathcal{L}$ over the base $\mathcal{B}(\boldsymbol{x}, \boldsymbol{y})$.

Theorem 2.20. — *The bundle $\mathcal{L} \longrightarrow \mathcal{B}(\boldsymbol{x}, \boldsymbol{y})$ is a Banach bundle.*

By construction the operator $\partial_{\mathcal{J}_s}$ is a section of that Banach bundle. Let us define $\mathcal{B}_0 \hookrightarrow \mathcal{B}(\boldsymbol{x}, \boldsymbol{y})$ as the zero section, then, obviously,

$$\mathcal{M}_{\mathcal{J}_s}(\boldsymbol{x}, \boldsymbol{y}) = (\partial_{\mathcal{J}_s})^{-1}(\mathcal{B}_0).$$

Recall from the Differential Topology of finite-dimensional manifolds that if a smooth map intersects a submanifold transversely, then its preimage is a manifold. There is an analogous result in the infinite-dimensional theory. The generalization to infinite dimensions requires an additional property to be imposed on the map. We will now define this property.

Definition 2.21. — *A map f between Banach manifolds is called* **Fredholm** *if for every point p the differential T_pf is a Fredholm operator, i.e. has finite-dimensional kernel and cokernel. The difference* $\dim\ker T_pf - \dim\operatorname{coker} T_pf$ *is called the* **Fredholm index** *of f at p.*

Fortunately the operator $\partial_{\mathcal{J}_s}$ is an elliptic operator, and hence it is Fredholm for a generic choice of path $(\mathcal{J}_s)_{s\in[0,1]}$ of almost complex structures.

Theorem 2.22 *(see Theorem 3.18 of [17]). — For a dense set of paths $(\mathcal{J}_s)_{s\in[0,1]}$ of (j,η,V)-nearly symmetric almost complex structures the moduli spaces $\mathcal{M}_{\mathcal{J}_s}(\boldsymbol{x}, \boldsymbol{y})$ are smooth manifolds for all $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{T}_{\boldsymbol{\alpha}} \cap \mathbb{T}_{\boldsymbol{\beta}}$.*

The idea is similar to the standard Floer homological proof. One realizes these paths as regular values of the Fredholm projection

$$\pi\colon \mathcal{M} \longrightarrow \Omega(\mathcal{J}(j,\eta,V)),$$

where $\Omega(\mathcal{J}(j,\eta,V))$ denotes the space of paths in $\mathcal{J}(j,\eta,V)$ and $\mathcal{M}$ is the unparametrized moduli space consisting of pairs $(\mathcal{J}_s,\phi)$, where $\mathcal{J}_s$ is a path of (j,η,V)-nearly symmetric almost complex structures and ϕ a Whitney disk. By the Sard-Smale theorem the set of regular values is an open and dense set of $\mathcal{J}(j,\eta,V)$.

Besides the smoothness of the moduli spaces we need the number of 1-dimensional components to be finite. This means we require the spaces $\widehat{\mathcal{M}}(\boldsymbol{x},\boldsymbol{y})^0_{n_z=0}$ to be compact. One ingredient of the compactness is the admissibility property introduced in Definition 2.17. In (2.5) we observed that

$$\widehat{\mathcal{M}}(\boldsymbol{x},\boldsymbol{y})^0_{n_z=0} = \bigsqcup_{\phi\in H(\boldsymbol{x},\boldsymbol{y},1)} \widehat{\mathcal{M}}_\phi,$$

where $H(\boldsymbol{x},\boldsymbol{y},1)$ is the set of homotopy classes of Whitney disks with $n_z=0$ and expected dimension $\mu=1$. Admissibility guarantees that $H(\boldsymbol{x},\boldsymbol{y},1)$ is a finite set. Thus, compactness follows from the compactness of the $\widehat{\mathcal{M}}_\phi$. The compactness proof follows similar lines as the Floer homological approach. It follows from the existence of an *energy bound* independent of the homotopy class of Whitney disks (see [17, §3.4]). The existence of this energy bound shows that the moduli spaces $\widehat{\mathcal{M}}(\boldsymbol{x},\boldsymbol{y})$ admit a compactification by adding solutions to the space in a controlled way.

Without giving the precise definition we would like to give some intuition of what happens at the boundaries. First of all there is an operation called **gluing** making it possible to concatenate Whitney disks holomorphically. Given two Whitney disks $\phi_1\in\pi_2(\boldsymbol{x},\boldsymbol{y})$ and $\phi_2\in\pi_2(\boldsymbol{y},\boldsymbol{q})$, gluing describes an operation to generate a family of holomorphic solutions $\phi_2\#_t\phi_1$ in the homotopy class $\phi_2*\phi_1$.

Definition 2.23. — *We call the pair (ϕ_2,ϕ_1) a* **broken** *holomorphic Whitney disk.*[1]

Moreover, one can think of this solution $\phi_2\#_t\phi_1$ as sitting in a small neighborhood of the boundary of the moduli space of the homotopy class $\phi_2*\phi_1$, i.e. the family of holomorphic solutions as $t\to\infty$ converges to the broken disk (ϕ_2,ϕ_1). There is a special notion of convergence used here. The limiting objects can be described intuitively in the following way: Think of the disk, after removing the points $\pm i$, as a strip $\mathbb{R}\times[0,1]$. Choose a properly embedded arc or an embedded $\mathbb{S}^1$ in $\mathbb{R}\times[0,1]$. Collapse the curve or the $\mathbb{S}^1$ to a point. The resulting object is a potential limiting object. The objects at

(1) This might be a sloppy and informal definition but appropriate for our intuitive approach.

the limits of sequences can be derived by applying several knot shrinkings and arc shrinkings simultaneously where we have to keep in mind that the arcs and knots have to be chosen such that they do not intersect (for a detailed treatment see [11, §4]).

We see that every broken disk corresponds to a boundary component of the compactified moduli space, i.e. there is an injection

$$f_{\text{glue}}\colon \mathcal{M}_{\phi_2} \times \mathcal{M}_{\phi_1} \hookrightarrow \partial \mathcal{M}_{\phi_2 * \phi_1}.$$

But are these the only boundary components? If this is the case, by adding broken disks to the space we would compactify it. This would result in the finiteness of the 0-dimensional spaces $\widehat{\mathcal{M}}_\phi$. A compactification by adding broken flow lines means that the 0-dimensional components are compact in the usual sense. A simple dimension count contradicts the existence of a family of disks in a 0-dimensional moduli space converging to a broken disk. But despite that there is a second reason for us to wish broken flow lines to compactify the moduli spaces. The map $\widehat{\partial}_\mathcal{H}$ should be a boundary operator. Calculating $\widehat{\partial}_\mathcal{H} \circ \widehat{\partial}_\mathcal{H}$ we see that the coefficients in the resulting equation equal the number of boundary components corresponding to broken disks at the ends of the 1-dimensional moduli spaces. If the gluing map is a bijection the broken ends generate all boundary components. Hence, the coefficients vanish mod 2.

There are two further phenomena we have to notice. Besides breaking there might be **spheres bubbling off**. This description can be taken literally to some point. Figure 1 illustrates the geometric picture behind that phenomenon. Bubbling is some kind of breaking phenomenon but the components here are disks and spheres. We do not need to take care of spheres bubbling off at all. Suppose that the boundary of the moduli space associated to the homotopy class ϕ we have breaking into a disk ϕ_1 and a sphere S_1, i.e. $\phi = \phi_1 * S_1$. Recall that the spheres in the symmetric product are generated by S, described in §2.2. Thus, $\phi = \phi_1 * k \cdot S$ where $n_z(S) = 1$. In consequence $n_z(\phi)$ is non-zero, contradicting the assumptions.

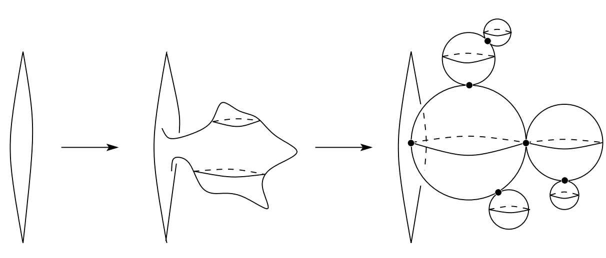

Figure 1. — Bubbling of spheres.

Definition 2.24. — *For a point $\boldsymbol{x}\in\mathbb{T}_{\boldsymbol{\alpha}}\cap\mathbb{T}_{\boldsymbol{\beta}}$ an* α-**degenerate** *disk is a holomorphic disk $\phi\colon[0,\infty)\times\mathbb{R}\longrightarrow\mathrm{Sym}^g(\Sigma)$ such that $\phi(\{0\}\times\mathbb{R})\subset\mathbb{T}_{\boldsymbol{\alpha}}$ and $\phi(p)\to\boldsymbol{x}$ as $\boldsymbol{x}\to\infty$.*

Given a degenerate disk ψ, the associated domain $\mathcal{D}(\psi)$ equals a sphere with holes, i.e. $\mathcal{D}(\psi)$ equals a surface in Σ with boundary the α-curves. Since the α-curves do not diskonnect Σ, the domain covers the whole surface. Thus, $n_z(\psi)$ is non-zero, showing that degenerations are ruled out by assuming that $n_z=0$.

Proof. — [Proof of Theorem 2.4 with $\mathbb{Z}_2$-coefficients] Fix an intersection $\boldsymbol{x}\in\mathbb{T}_{\boldsymbol{\alpha}}\cap\mathbb{T}_{\boldsymbol{\beta}}$. We compute

$$\begin{aligned}\widehat{\partial}^2_{\mathcal{H}}\boldsymbol{x} &= \widehat{\partial}_{\mathcal{H}}\Big(\sum_{\boldsymbol{y}\in\mathbb{T}_{\boldsymbol{\alpha}}\cap\mathbb{T}_{\boldsymbol{\beta}}}\#\widehat{\mathcal{M}}(\boldsymbol{x},\boldsymbol{y})^0_{n_z=0}\cdot\boldsymbol{y}\Big)\\ &= \sum_{\boldsymbol{y},\boldsymbol{q}\in\mathbb{T}_{\boldsymbol{\alpha}}\cap\mathbb{T}_{\boldsymbol{\beta}}}\#\widehat{\mathcal{M}}(\boldsymbol{x},\boldsymbol{y})^0_{n_z=0}\#\widehat{\mathcal{M}}(\boldsymbol{y},\boldsymbol{q})^0_{n_z=0}\cdot\boldsymbol{q}.\end{aligned}$$

We have to show that the coefficient in front of $\boldsymbol{q}$, denoted by $c(\boldsymbol{x},\boldsymbol{q})$ vanishes. Observe that the coefficient precisely equals the number of components (mod 2) in

$$\widehat{\mathcal{M}}(\boldsymbol{x},\boldsymbol{y})^0_{n_z=0}\times\widehat{\mathcal{M}}(\boldsymbol{y},\boldsymbol{q})^0_{n_z=0}.$$

Gluing gives an injection

$$\widehat{\mathcal{M}}(\boldsymbol{x},\boldsymbol{y})^0_{n_z=0}\times\widehat{\mathcal{M}}(\boldsymbol{y},\boldsymbol{q})^0_{n_z=0}\hookrightarrow\partial\widehat{\mathcal{M}}(\boldsymbol{x},\boldsymbol{q})^1_{n_z=0}.$$

By the compactification theorem the gluing map is a bijection, since bubbling and degenerations do not appear due to the condition $n_z=0$. Thus, (mod 2) we have

$$\begin{aligned}c(\boldsymbol{x},\boldsymbol{q}) &= \#(\widehat{\mathcal{M}}(\boldsymbol{x},\boldsymbol{y})^0_{n_z=0}\times\widehat{\mathcal{M}}(\boldsymbol{y},\boldsymbol{q})^0_{n_z=0})\\ &= \partial\widehat{\mathcal{M}}(\boldsymbol{x},\boldsymbol{q})^1_{n_z=0}\\ &= 0,\end{aligned}$$

which proves the theorem. □

Obviously, the proof breaks down in $\mathbb{Z}$-coefficients. We need the mod 2 count of ends. There is a way to fix the proof. The goal is to make the map

$$f_{\mathrm{glue}}\colon\mathcal{M}_{\phi_2}\times\mathcal{M}_{\phi_1}\hookrightarrow\partial\mathcal{M}_{\phi_2*\phi_1}$$

orientation preserving. For this to make sense we need the moduli spaces to be oriented. An orientation is given by choosing a section of the **determinant line bundle** over the moduli spaces. The determinant line bundle is defined as the bundle $\det([\phi]) \longrightarrow \mathcal{M}_\phi$ given by putting together the spaces

$$\det(\psi) = \bigwedge^{\max} \ker(D_\psi \partial_{\mathcal{J}_s}) \otimes \bigwedge^{\max} \ker((D_\psi \partial_{\mathcal{J}_s})^*),$$

where ψ is an element of $\mathcal{M}_\phi$. If we achieve transversality for $\partial_{\mathcal{J}_s}$, i.e. it has transverse intersection with the zero section $\mathcal{B}_0 \hookrightarrow \mathcal{L}$ then

$$\begin{array}{rcccc} \det(\psi) & = & \bigwedge^{\max} \ker(D_\psi \partial_{\mathcal{J}_s}) & \otimes & \mathbb{R}^* \\ & = & \bigwedge^{\max} T_\psi \mathcal{M}_\phi & \otimes & \mathbb{R}^*. \end{array}$$

Thus, a section of the determinant line bundle defines an orientation of $\mathcal{M}_\phi$. These have to be chosen in a coherent fashion to make f_{glue} orientation preserving. The gluing construction gives a natural identification

$$\det(\phi_1) \wedge \det(\phi_2) \xrightarrow{\cong} \det(\phi_2 \#_t \phi_1).$$

Since these are all line bundles, this identification makes it possible to identify sections of $\det([\phi_1]) \wedge \det([\phi_2])$ with sections of $\det([\phi_2 * \phi_1])$. With this isomorphism at hand we are able to define a coherence condition. Namely, let $\mathfrak{o}(\phi_1)$ and $\mathfrak{o}(\phi_2)$ be sections of the determinant line bundles of the associated moduli spaces, then we need that under the identification given above we have

$$\mathfrak{o}(\phi_1) \wedge \mathfrak{o}(\phi_2) = \mathfrak{o}(\phi_2 * \phi_1). \tag{2.10}$$

In consequence, a **coherent system of orientations** is a section $\mathfrak{o}(\phi)$ of the determinant line bundle $\det(\phi)$ for each homotopy class of Whitney disks ϕ connecting two intersection points such that equation (2.10) holds for each pair for which concatenation makes sense. It is not clear if these systems exist in general. By construction with respect to these coherent systems of orientations the map f_{glue} is orientation preserving.

In the case of Heegaard Floer theory there is an easy way to give a construction of coherent systems of orientations. Namely, fix a Spin^c-structure s and let $\{\boldsymbol{x}_0, \ldots, \boldsymbol{x}_l\}$ be the points representing s, i.e. $(\mathfrak{s}_z)^{-1}(\mathfrak{s}) = \{\boldsymbol{x}_0, \ldots, \boldsymbol{x}_l\}$. Let $\phi_1, \ldots, \phi_q$ be a set of periodic classes in $\pi_2(\boldsymbol{x}_0, \boldsymbol{x}_0)$ representing a basis for $H^1(Y; \mathbb{Z})$ denote by θ_i an element of $\pi_2(\boldsymbol{x}_0, \boldsymbol{x}_i)$. A coherent system of orientations is constructed by choosing sections over all chosen disks, i.e. $\mathfrak{o}(\phi_i)$, $i = 1, \ldots, q$ and $\mathfrak{o}(\theta_j)$, $j = 1, \ldots, l$. Namely, for each homotopy class $\phi \in \pi_2(\boldsymbol{x}_i, \boldsymbol{x}_j)$ we have a presentation (cf. Lemma 2.8, Lemma 2.9 and (2.4))

$$\phi = a_1 \phi_1 + \ldots + a_q \phi_q + \theta_j - \theta_i$$

inducing an orientation $\mathfrak{o}(\phi)$. This definition defines a coherent system.

To give a proof of Theorem 2.4 in case of $\mathbb{Z}$-coefficients, we have to translate orientations on the 0-dimensional components of the moduli spaces $\widehat{\mathcal{M}}_{\mathcal{J}_s}(\boldsymbol{x},\boldsymbol{y})$ of connecting Whitney disks into signs. For ϕ with $\mu(\phi)=1$ the translation action naturally induces an orientation on $\mathcal{M}_\phi$. Comparing this orientation with the coherent orientation induces a sign. We define the **signed count** as the count of the elements by taking into account the signs induced by the comparison of the action orientation with the coherent orientation.

Proof of Theorem 2.4 for $\mathbb{Z}$-coefficients. — We stay in the notation of the earlier proof. With the coherent system of orientations introduced we made the map

$$f_{\text{glue}}\colon \widehat{\mathcal{M}}(\boldsymbol{x},\boldsymbol{y})^0_{n_z=0}\times\widehat{\mathcal{M}}(\boldsymbol{y},\boldsymbol{q})^0_{n_z=0}\hookrightarrow\partial\widehat{\mathcal{M}}(\boldsymbol{x},\boldsymbol{q})^1_{n_z=0}$$

orientation preserving. Hence, we see that $c(\boldsymbol{x},\boldsymbol{q})$ equals

$$\#(\widehat{\mathcal{M}}(\boldsymbol{x},\boldsymbol{y})^0_{n_z=0}\times\widehat{\mathcal{M}}(\boldsymbol{y},\boldsymbol{q})^0_{n_z=0})$$

which in turn equals the oriented count of boundary components of $\partial\widehat{\mathcal{M}}(\boldsymbol{x},\boldsymbol{q})^1_{n_z=0}$. Since the space is 1-dimensional, this count vanishes. □

2.3.1. Other Heegaard Floer Theories

There are variants of Heegaard Floer homology which do not require the condition $n_z=0$. To make the compactification work in that case we have to take care of boundary degenerations and bubbling of spheres. Both can be shown to be controlled in the following way: we get rid of bubbling by a proper choice of almost complex structure. Namely, by choosing $\mathfrak{j}$ on Σ appropriately there is a contractible open neighborhood of $sym^g(\mathfrak{j})$ in $\mathcal{J}(\mathfrak{j},\eta,V)$ for which all spheres miss the intersections $\mathbb{T}_{\boldsymbol{\alpha}}\cap\mathbb{T}_{\boldsymbol{\beta}}$. Moreover, for a generic choice of path $(\mathcal{J}_s)_{s\in[0,1]}$ inside this neighborhood the signed count of degenerate disks is zero (see [17, Theorem 3.15]). With this information it is easy to modify the given proof for the other theories. We leave this to the interested reader or point him to [17, p. 1066].

2.4. Choice of Almost Complex Structure

Let Σ be endowed with a complex structure $\mathfrak{j}$ and let $U\subset\Sigma$ be a subset diffeomorphic to a disk.

Theorem 2.25 (Riemann mapping theorem). — *There is a 3-dimensional connected family of holomorphic identifications of U with the unit disk $D^2\subset\mathbb{C}$.*

Consequently, suppose that all moduli spaces are compact manifolds for the path of almost complex structures $(\mathcal{J}_s)_{s\in[0,1]}$ where $\mathcal{J}_s = sym^g(\mathfrak{j})$ for all $s \in [0,1]$. In this case we conclude from the Riemann mapping theorem the following corollary.

COROLLARY 2.26 *Let* $\phi\colon D^2 \longrightarrow \mathrm{Sym}^g(\Sigma)$ *be a holomorphic disk with* $\mathcal{D}(\phi)$ *isomorphic to a disk. Then the moduli space* $\widehat{\mathcal{M}}_\phi$ *contains a unique element.*

We call a domain $\mathcal{D}(\phi)$ α**-injective** if all its multiplicities are 0 or 1 and its interior is disjoint from the α-circles. We then say that the homotopy class ϕ is α**-injective**.

THEOREM 2.27. — *Let* $\phi \in \pi_2(\boldsymbol{x}, \boldsymbol{y})$ *be an* α*-injective homotopy class and* $\mathfrak{j}$ *a complex structure on* Σ*. For generic perturbations of the* α*-curves the moduli space* $\mathcal{M}_{sym^g(\mathfrak{j}),\phi}$ *is a smooth manifold.*

In explicit calculations it will be nice to have all homotopy classes carrying holomorphic representatives to be α-injective. In this case we can choose the path of almost complex structures in such a way that homotopy classes of Whitney disks with disk-shaped domains just admit a unique element. This is exactly what can be achieved in general to make the $\widehat{\mathrm{HF}}$-theory combinatorial. For a class of Heegaard diagrams called **nice diagrams** all moduli spaces with $\mu = 1$ just admit one single element. In addition, we have a precise description of how these domains look like. In $\mathbb{Z}_2$-coefficients with nice diagrams this results in a method to calculate the differential $\widehat{\partial}_{\mathcal{H}}$ by counting the number of domains that *are of a certain shape*. For details we point the reader to [27].

DEFINITION 2.28 (see [27]). — *A pointed Heegaard diagram* $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, z)$ *is called* **nice** *if every region not containing* z *is either a bigon or a square.*

DEFINITION 2.29 (see [27]). — *A homotopy class is called an empty embedded* $2n$**-gon** *if it is topologically an embedded disk with* $2n$ *vertices at its boundary, it does not contain any* $\boldsymbol{x}_i$ *or* y_i *in its interior, and for each vertex* v*, the average of the coefficients of the four regions around* v *is* $1/4$.

For a nice Heegaard diagram one can show that elements $\phi \in H(\boldsymbol{x}, \boldsymbol{y}, 1)$ with $\mu(\phi) = 1$ that admit holomorphic representatives are empty embedded bigons or empty embedded squares. Furthermore, for a generic choice of $\mathfrak{j}$ on Σ the moduli spaces are regular under a generic perturbation of the α-curves and β-curves. The moduli space $\widehat{\mathcal{M}}_\phi$ contains one single element. Hence, to compute the Heegaard Floer homology in this situation we have to scan the Heegaard diagram for domains of a certain type. No holomorphic information is needed to compute the theory. We note the following property.

Theorem 2.30 (see [27]). — *Every 3-manifold admits a nice Heegaard diagram.*

2.5. Dependence on the Choice of Orientation Systems

From their definition it is easy to reorder the orientation systems into equivalence classes. The elements in these classes give rise to isomorphic homologies. Let $\mathfrak{o}$ and $\mathfrak{o}'$ be two orientation systems. We measure their difference

$$\delta\colon H^1(Y;\mathbb{Z}) \longrightarrow \mathbb{Z}_2$$

by saying that, given a periodic class $\phi \in \pi_2(\boldsymbol{x},\boldsymbol{x})$, we define $\delta(\phi) = 0$ if $\mathfrak{o}(\phi)$ and $\mathfrak{o}'(\phi)$ coincide, i.e. define equivalent sections, and $\delta(\phi) = 1$, if $\mathfrak{o}(\phi)$ and $\mathfrak{o}'(\phi)$ define non-equivalent sections. Thus, two systems are equivalent if $\delta = 0$. Obviously, there are $2^{b_1(Y)}$ different equivalence classes of orientation systems. In general the Heegaard Floer homologies will depend on choices of equivalence classes of orientation systems. As an illustration we will discuss an example.

Example 2.3. — The manifold $\mathbb{S}^2 \times \mathbb{S}^1$ admits a Heegaard splitting of genus one, namely (T^2, α, β, z) where α and β are two distinct meridians of T^2.

Unfortunately this is not an admissible diagram. By the universal coefficient theorem

$$H^2(\mathbb{S}^2 \times \mathbb{S}^1;\mathbb{Z}) \cong \mathrm{Hom}(H_2(\mathbb{S}^2 \times \mathbb{S}^1;\mathbb{Z}),\mathbb{Z}) \cong \mathrm{Hom}(\mathbb{Z},\mathbb{Z}).$$

Hence we can interpret Spin^c-structures as homomorphisms $\mathbb{Z} \longrightarrow \mathbb{Z}$. For a number $q \in \mathbb{Z}$ define $\mathfrak{s}_q$ to be the Spin^c-structure whose associated characteristic class, which we also call $\mathfrak{s}_q$, is given by $\mathfrak{s}_q(1) = q$. The two curves α and β cut the torus into two components, where z is placed in one of them. Denote the other component with $\mathcal{D}$. It is easy to see that the homology class $\mathcal{H}(\mathcal{D})$ is a generator of $H_2(\mathbb{S}^2 \times \mathbb{S}^1;\mathbb{Z})$. Thus, we have

$$\langle c_1(\mathfrak{s}_q), \mathcal{H}(\lambda\cdot\mathcal{D})\rangle = \langle 2\cdot\mathfrak{s}_q, \mathcal{H}(\lambda\cdot\mathcal{D})\rangle = 2\cdot\mathfrak{s}_q(\lambda\cdot 1) = 2\lambda q.$$

This clearly contradicts the weak admissibility condition. We fix this problem by perturbing the β-curve slightly to give a Heegaard diagram as illustrated in Figure 2. By boundary orientations $\mathbb{Z}\,\langle(\mathcal{D}_1 - \mathcal{D}_2)\rangle$ are all possible periodic domains.

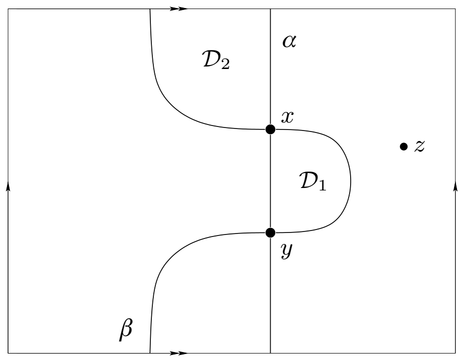

Figure 2. — An admissible Heegaard diagram for $\mathbb{S}^2 \times \mathbb{S}^1$.

Figure 2 shows that the chain module is generated by the points $\boldsymbol{x}$ and $\boldsymbol{y}$. A straightforward computation gives $\epsilon(\boldsymbol{x}, \boldsymbol{y}) = 0$ (see §2.2 for a definition) and, hence, both intersections belong to the same Spin^c-structure we will denote by $\mathfrak{s}_0$. Thus, the chain complex $\widehat{\mathrm{CF}}(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}; \mathfrak{s}_0)$ equals $\mathbb{Z}\langle\{\boldsymbol{x}, \boldsymbol{y}\}\rangle$. The regions $\mathcal{D}_1$ and $\mathcal{D}_2$ are both disk-shaped and hence α-injective. Thus, the Riemann mapping theorem (see §2.4) gives

$$\#\widehat{\mathcal{M}}_{\phi_1} = 1 \text{ and } \#\widehat{\mathcal{M}}_{\phi_2} = 1.$$

These two disks differ by the periodic domain generating $H^1(\mathbb{S}^2 \times \mathbb{S}^1; \mathbb{Z})$. Thus, we are free to choose the orientation on this generator (cf. §2.3). Hence, we may choose the signs on ϕ_1 and ϕ_2 arbitrarily. Thus, there are two equivalence classes of orientation systems. We define $\mathfrak{o}_0$ to be the system of orientations where the signs differ and $\mathfrak{o}_1$ where they are equal. Thus, we get two different homology theories

$$\begin{aligned} \widehat{\mathrm{HF}}(\mathbb{S}^2 \times \mathbb{S}^1, \mathfrak{s}_0; \mathfrak{o}_0) &= \mathbb{Z} \oplus \mathbb{Z} \\ \widehat{\mathrm{HF}}(\mathbb{S}^2 \times \mathbb{S}^1, \mathfrak{s}_0; \mathfrak{o}_1) &= \mathbb{Z}_2. \end{aligned}$$

However, there is a special choice of coherent orientation systems. We point the reader to §3 for a definition of HF^∞. Additionally, instead of using $\mathbb{Z}$-coefficients, we can use the ring $\mathbb{Z}[H_1(Y)]$ as coefficients for defining this Heegaard Floer group. The resulting group is denoted by $\underline{\mathrm{HF}^\infty}$. We point the reader to [17, §11.0.1] for a precise definition or to [18, §8]. As a matter or completeness we cite:

THEOREM 2.31 (see Theorem 10.12 of [18]). — *Let Y be a closed oriented 3-manifold. Then there is a unique equivalence class of orientation* *systems such that for each torsion* Spin^c*-structure* $\mathfrak{s}_0$ *there is an isomorphism*

$$\underline{\mathrm{HF}^{\infty}}(Y, \mathfrak{s}_0) \cong \mathbb{Z}[U, U^{-1}]$$

as $\mathbb{Z}[U, U^{-1}] \otimes_{\mathbb{Z}} \mathbb{Z}[H^1(Y;\mathbb{Z})]$*-modules.*

3. Other versions of Heegaard Floer homology

Given a pointed Heegaard diagram $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, z)$, we define $\mathrm{CF}^-(\mathcal{H};\mathfrak{s})$ as the free $\mathbb{Z}[U^{-1}]$-module generated by the points of intersection $(\mathfrak{s}_z)^{-1}(\mathfrak{s}) \subset \mathbb{T}_{\boldsymbol{\alpha}} \cap \mathbb{T}_{\boldsymbol{\beta}}$. For an intersection $\boldsymbol{x}$ we define

$$\partial^-_{\mathcal{H}} \boldsymbol{x} = \sum_{\boldsymbol{y} \in (\mathfrak{s}_z)^{-1}(\mathfrak{s})} \sum_{\phi \in \mu^{-1}(1)} \#\widehat{\mathcal{M}}_\phi \cdot U^{-n_z(\phi)} \boldsymbol{y},$$

where $\mu^{-1}(1)$ are the homotopy classes in $\pi_2(\boldsymbol{x}, \boldsymbol{y})$ with expected dimension equal to one. Note that in this theory we do not restrict to classes with $n_z = 0$. This means even with weak admissibility imposed on the Heegaard diagram the proof of well-definedness as it was done in §2 breaks down.

Definition 3.1. — *A Heegaard diagram* $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, z)$ *is called* ***strongly admissible*** *for the* Spin^c*-structure* $\mathfrak{s}$ *if for every non-trivial periodic domain* $\mathcal{D}$ *such that* $\langle c_1(\mathfrak{s}), H(\mathcal{D})\rangle = 2n \geqslant 0$ *the domain* $\mathcal{D}$ *has some coefficient greater than* n.

Imposing strong admissibility on the Heegaard diagram we can prove well-definedness by showing that only finitely many homotopy classes of Whitney disks contribute to the moduli space $\mathcal{M}_{\mathcal{J}_s}(\boldsymbol{x}, \boldsymbol{y})$ (cf. §2).

Theorem 3.2 (Theorem 4.15 of [17]). — *The map* $\partial^-_{\mathcal{H}}$ *is a differential.*

As mentioned in §2, in this case we have to take a look at bubbling and degenerate disks. The proof follows the same lines as the proof of Theorem 2.4. With the remarks made in §2 it is easy to modify the given proof to a proof of Theorem 3.2 (see [17, p. 1066]). We define

$$\mathrm{CF}^{\infty}(\mathcal{H};\mathfrak{s}) = \mathrm{CF}^-(\mathcal{H};\mathfrak{s}) \otimes_{\mathbb{Z}[U^{-1}]} \mathbb{Z}[U, U^{-1}]$$

and denote by ∂^∞ the induced differential. From the definition we get an inclusion of $\mathrm{CF}^- \hookrightarrow \mathrm{CF}^\infty$ whose cokernel is defined as $\mathrm{CF}^+(\mathcal{H};\mathfrak{s})$. Finally we get back to $\widehat{\mathrm{CF}}$ by

$$\widehat{\mathrm{CF}}(\mathcal{H};\mathfrak{s}) = \frac{U \cdot \mathrm{CF}^-(\mathcal{H};\mathfrak{s})}{\mathrm{CF}^-(\mathcal{H};\mathfrak{s})}.$$

The associated homology theories are denoted by HF^∞, HF^- and $\widehat{\mathrm{HF}}$. There are two long exact sequences which can be derived easily from the definition of the Heegaard Floer homologies. To give an intuitive picture look at the following illustration:

$$\begin{array}{lcccccccccc}
\mathrm{CF}^\infty & = & \dots & U^{-3} & U^{-2} & U^{-1} & U^0 & U^1 & U^2 & U^3 & \dots \\
\mathrm{CF}^- & = & \dots & U^{-3} & U^{-2} & U^{-1} & & & & & \\
\widehat{\mathrm{CF}} & = & & & & & U^0 & & & & \\
\mathrm{CF}^+ & = & & & & & U^0 & U^1 & U^2 & U^3 & \dots
\end{array}$$

We see why the condition of weak admissibility is not strong enough to give a well-defined differential on CF^∞ or CF^-. However, weak admissibility is enough to make the differential on CF^+ well-defined, since the complex is bounded from below with respect to the obvious filtration given by the U-variable.

Lemma 3.3. — *There are two long exact sequences*

$$\begin{array}{lccccccc}
\dots \longrightarrow & \mathrm{HF}^-(Y;\mathfrak{s}) & \longrightarrow & \mathrm{HF}^\infty(Y;\mathfrak{s}) & \longrightarrow & \mathrm{HF}^+(Y;\mathfrak{s}) & \longrightarrow & \dots \\
\dots \longrightarrow & \widehat{\mathrm{HF}}(Y;\mathfrak{s}) & \longrightarrow & \mathrm{HF}^+(Y;\mathfrak{s}) & \longrightarrow & \mathrm{HF}^+(Y;\mathfrak{s}) & \longrightarrow & \dots
\end{array}$$

where $\mathfrak{s}$ is a Spin^c*-structure of Y.*

The explicit description illustrated above can be derived directly from the definition of the complexes. We leave this to the interested reader (see also [17, Lemma 4.4 and p. 1066]).

4. Topological Invariance

Let Y be a closed, oriented 3-manifold and suppose we are given two Heegaard diagrams $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$ and $\mathcal{H}' = (\Sigma', \boldsymbol{\alpha}', \boldsymbol{\beta}')$ which both represent Y. These diagrams are equivalent after a finite sequence of isotopies of the attaching circles, handle slides of the α-curves and β-curves and stabilizations/destabilizations. Two Heegaard diagrams are equivalent if there is a diffeomorphism of the Heegaard surface interchanging the attaching circles. Obviously, equivalent Heegaard diagrams define isomorphic Heegaard Floer theories. To show that Heegaard Floer theory is a topological invariant of the manifold Y we have to see that each of the moves, i.e. isotopies, handle slides and stabilization/destabilizations yield isomorphic theories. We will briefly sketch the topological invariance. This has two reasons: First of all the invariance proof uses ideas that are standard in Floer homology theories and hence appear frequently. The ideas provided from the invariance proof happen to be the standard techniques for proving exactness of sequences, proving invariance properties, and proving the existence of morphisms between Floer homologies. Thus, knowing the invariance proof, at least at the level of ideas, is crucial for an understanding of most of the papers published in this field. We will deal with the $\widehat{\mathrm{HF}}$-case and and point the reader to [17] for a general treatment.

The invariance proof contains several steps. We start showing invariance under the choice of path of admissible almost complex structures. Isotopies of the attaching circles are split up into two separate classes: Isotopies that generate/cancel intersection points and those which do not change the chain module. The invariance under the latter Heegaard moves immediately follows from the independence of the choice of almost complex structures. Such an isotopy is carried by an ambient isotopy inducing an isotopy of the symmetric product. We perturb the almost complex structure and thus interpret the isotopy as a perturbation of the almost complex structure. The former Heegaard moves have to be dealt with separately. We mimic the generation/cancellation of intersection points with a Hamiltonian isotopy and with it explicitly construct an isomorphism of the respective homologies by counting disks with dynamic boundary conditions. Stabilizations/ destabilizations is the easiest part to deal with: it follows from the behavior of the Heegaard Floer theory under connected sums. Finally, handle slide invariance will require us to define what can be regarded as the Heegaard Floer homological version of the pair-of-pants product in Floer homologies. This product has two nice applications. The first is the invariance under handle slides and the second is the association of maps to cobordisms giving the theory the structure of a topological field theory.

4.1. Stabilizations/Destabilizations

We determine the groups $\widehat{\mathrm{HF}}(\mathbb{S}^2\times\mathbb{S}^1\#\mathbb{S}^2\times\mathbb{S}^1)$ as a model calculation for how the groups behave under connected sums.

Example 4.1. — We fix admissible Heegaard diagrams (T^2_i,α_i,β_i) $i=1,2$ for $\mathbb{S}^2\times\mathbb{S}^1$ as in Example 2.3. To perform the connected sum of $\mathbb{S}^2\times\mathbb{S}^1$ with itself we choose 3-balls such that their intersection D with the Heegaard surface fulfills the property

$$\mathcal{J}^i_s\big|_{D^2}=\mathrm{sym}(\mathfrak{j}_i).$$

Figure 3 pictures the Heegaard diagram we get for the connected sum. Denote by T a small connected sum tube inside $\Sigma=T^2_1\#T^2_2$. By construction the induced almost complex structure equals

$$(\mathcal{J}^1\#\mathcal{J}^2)_s\big|_{T\times\Sigma}=sym^2(\mathfrak{j}^1\#\mathfrak{j}^2).$$

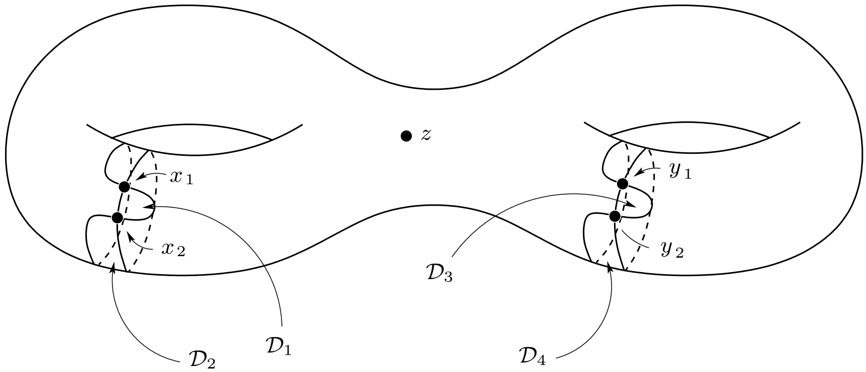

Figure 3. — An admissible Heegaard diagram for $\mathbb{S}^2 \times \mathbb{S}^1 \# \mathbb{S}^2 \times \mathbb{S}^1$.

All intersection points belong to the same Spin^c-structure s_0. For suitable Spin^c-structures $\mathfrak{s}_1$, $\mathfrak{s}_2$ on $\mathbb{S}^2 \times \mathbb{S}^1$ we have that $\mathfrak{s}_0 = \mathfrak{s}_1 \# \mathfrak{s}_2$ and

$$\widehat{\mathrm{CF}}(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s}_1 \# \mathfrak{s}_2) = \mathbb{Z} \otimes \{(\boldsymbol{x}_i, \boldsymbol{y}_j) \,|\, i, j \in \{1,2\}\} \cong \widehat{\mathrm{CF}}(T_1^2, \mathfrak{s}_1) \otimes \widehat{\mathrm{CF}}(T_2^2, \mathfrak{s}_2).$$

The condition $n_z = 0$ implies that for every holomorphic disk ϕ: $D^2 \longrightarrow \mathrm{Sym}^g(\Sigma)$ the low-dimensional model (cf. §2) $\widehat{\phi}$: $\widehat{D^2} \longrightarrow \Sigma$ stays away from the tube T. Consequently we can split up $\widehat{D^2}$ into

$$\widehat{D^2} = \widehat{D^2}_1 \sqcup \widehat{D^2}_2,$$

where $\widehat{D^2}_i$ are the components containing the preimage $(\widehat{\phi})^{-1}(T_i^2 \backslash D)$. Restriction to these components determines maps $\widehat{\phi}_i$: $\widehat{D^2}_i \longrightarrow T_i^2$ which induce Whitney disks ϕ_i in the symmetric product $\mathrm{Sym}^1(T^2)$. Thus, the moduli spaces split:

$$\begin{aligned} \mathcal{M}_{(\mathcal{J}^1 \# \mathcal{J}^2)_s}((\boldsymbol{x}_i, \boldsymbol{y}_k), (\boldsymbol{x}_j, \boldsymbol{y}_l))_{n_z=0} &\overset{\cong}{\longrightarrow} \mathcal{M}_{\mathcal{J}_s^1}(\boldsymbol{x}_i, \boldsymbol{x}_j)_{n_z=0} \times \mathcal{M}_{\mathcal{J}_s^2}(\boldsymbol{y}_k, \boldsymbol{y}_l)_{n_z=0} \\ \phi &\longmapsto (\phi_1, \phi_2). \end{aligned}$$

For moduli spaces with expected dimension $\mu = 1$, a dimension count forces one of the factors to be constant. So, the differential splits, too, i.e. for $a_i \in \widehat{\mathrm{CF}}(T_i^2, \mathfrak{s}_i)$, $i = 1, 2$, we see that

$$\widehat{\partial}_{(\mathcal{J}^1 \# \mathcal{J}^2)_s}(a_1 \otimes a_2) = \widehat{\partial}_{\mathcal{J}_s^1}(a_1) \otimes a_2 + a_1 \otimes \widehat{\partial}_{\mathcal{J}_s^2}(a_2).$$

And, consequently,

$$\widehat{\mathrm{HF}}(\mathbb{S}^2 \times \mathbb{S}^1 \# \mathbb{S}^2 \times \mathbb{S}^1, \mathfrak{s}_1 \# \mathfrak{s}_2; \mathfrak{o}_1 \wedge \mathfrak{o}_2) \cong \widehat{\mathrm{HF}}(\mathbb{S}^2 \times \mathbb{S}^1, \mathfrak{s}_1; \mathfrak{o}_1) \otimes \widehat{\mathrm{HF}}(\mathbb{S}^2 \times \mathbb{S}^1, \mathfrak{s}_2; \mathfrak{o}_2).$$

The same line of arguments shows the general statement.

Theorem 4.1 (see Theorem 1.5 of [18]). — *For closed oriented* 3*-manifolds Y_i, $i=1,2$ the Heegaard Floer homology of the connected sum $Y_1\#Y_2$ equals the tensor product of the Heegaard Floer homologies of the factors, i.e.*

$$\widehat{\mathrm{HF}}(Y_1\#Y_2)=H_*(\widehat{\mathrm{CF}}(Y_1)\otimes\widehat{\mathrm{CF}}(Y_2)),$$

where the chain complex on the right carries the natural induced boundary.

Stabilizing a Heegaard diagram of Y means, on the manifold level, to do a connected sum with $\mathbb{S}^3$. We know that $\widehat{\mathrm{HF}}(\mathbb{S}^3)=\mathbb{Z}$. By the classification of finitely generated abelian groups and the behavior of the tensor product, invariance follows.

4.2. Independence of the Choice of Almost Complex Structures

Suppose we are given a 1-dimensional family of paths of $(\mathfrak{j},\eta,V)$-nearly symmetric almost complex structures $(\mathcal{J}_{s,t})$. Given a Whitney disk ϕ, we define $\mathcal{M}_{\mathcal{J}_{s,t},\phi}$ as the moduli space of Whitney disks in the homotopy class of ϕ which satisfy the equation

$$\frac{\partial\phi}{\partial s}(s,t)+\mathcal{J}_{s,t}\Big(\frac{\partial\phi}{\partial t}(s,t)\Big)=0.$$

Observe that there is no free translation action on the moduli spaces as on the moduli spaces we focused on while discussing the differential $\widehat{\partial}_{\mathcal{H}}$. We define a map $\widehat{\Phi}_{\mathcal{M}_{\mathcal{J}_{s,t}}}$ between the theories $(\widehat{\mathrm{CF}}(\mathcal{H}),\widehat{\partial}_{\mathcal{J}_{s,i}})$ for $i=0,1$ by defining for $\boldsymbol{x}\in\mathbb{T}_{\boldsymbol{\alpha}}\cap\mathbb{T}_{\boldsymbol{\beta}}$

$$\widehat{\Phi}_{\mathcal{J}_{s,t}}(\boldsymbol{x})=\sum_{\boldsymbol{y}\in\mathbb{T}_{\boldsymbol{\alpha}}\cap\mathbb{T}_{\boldsymbol{\beta}}}\ \sum_{\phi\in H(\boldsymbol{x},\boldsymbol{y},0)}\#\mathcal{M}_{\mathcal{J}_{s,t},\phi}\cdot\boldsymbol{y},$$

where $H(\boldsymbol{x},\boldsymbol{y},0)\subset\pi_2(\boldsymbol{x},\boldsymbol{y})$ are the homotopy classes with expected dimension $\mu=0$ and intersection number $n_z=0$. There is an energy bound for all holomorphic Whitney disks which is independent of the particular Whitney disk or its homotopy class (see [17, p. 1052–1053]). Thus, the moduli spaces are Gromov-compact manifolds, i.e. can be compactified by adding solutions coming from broken disks, bubbling of spheres and boundary degenerations (cf. §2.3). Since we consider the $\widehat{\mathrm{HF}}$-theory we impose the condition $n_z=0$ which circumvents bubbling of spheres and boundary degenerations (see §2.3).

To check that $\widehat{\Phi}_{\mathcal{J}_{s,t}}$ is a chain map, we compute

$$\begin{aligned}\widehat{\partial}_{J_{s,1}} \circ \widehat{\Phi}_{J_{s,t}}(\boldsymbol{x}) - \widehat{\Phi}_{J_{s,t}} \circ \widehat{\partial}_{J_{s,0}}(\boldsymbol{x}) &= \sum_{\substack{\boldsymbol{y},\boldsymbol{q} \\ \phi\in H(\boldsymbol{x},\boldsymbol{y},0),\psi\in H(\boldsymbol{y},z,1)}} \#\mathcal{M}_{\mathcal{J}_{s,t}}(\phi)\#\widehat{\mathcal{M}}_{\mathcal{J}_{s,1}}(\psi)\boldsymbol{q} \\ &\quad - \sum_{\substack{\boldsymbol{y},\boldsymbol{q} \\ \phi\in H(\boldsymbol{x},\boldsymbol{y},1),\psi\in H(\boldsymbol{y},z,0)}} \#\widehat{\mathcal{M}}_{\mathcal{J}_{s,0}}(\phi)\#\mathcal{M}_{\mathcal{J}_{s,t}}(\psi)\boldsymbol{q} \\ &= \sum_z c(\boldsymbol{x},\boldsymbol{q}) \cdot \boldsymbol{q}.\end{aligned}$$

The coefficient $c(\boldsymbol{x},\boldsymbol{q})$ is given by

$$\sum_{\boldsymbol{y},I} \left(\#\mathcal{M}_{\mathcal{J}_{s,t},\phi} \cdot \#\widehat{\mathcal{M}}_{\mathcal{J}_{s,1},\psi} - \#\widehat{\mathcal{M}}_{\mathcal{J}_{s,0},\widetilde{\psi}} \cdot \#\mathcal{M}_{\mathcal{J}_{s,t},\widetilde{\phi}}\right), \tag{4.1}$$

where I consists of pairs

$$(\phi,\widetilde{\phi}) \in H(\boldsymbol{x},\boldsymbol{y},0) \times H(\boldsymbol{y},\boldsymbol{q},0) \text{ and} (\psi,\widetilde{\psi}) \in H(\boldsymbol{x},\boldsymbol{y},1) \times H(\boldsymbol{y},\boldsymbol{q},1).$$

Looking at the ends of the moduli spaces $\mathcal{M}_{\mathcal{J}_{s,t}}(\eta)$ for $\eta \in H(\boldsymbol{x},\boldsymbol{q},1)$, the gluing construction (cf. §2.3) together with the compactification argument mentioned earlier provides the following ends:

$$\Big(\bigsqcup_{\eta=\psi*\phi} (\mathcal{M}_{\mathcal{J}_{s,t}}(\phi) \times \widehat{\mathcal{M}}_{\mathcal{J}_{s,1}}(\psi))\Big) \sqcup \Big(\bigsqcup_{\eta=\widetilde{\psi}*\widetilde{\phi}} (\widehat{\mathcal{M}}_{\mathcal{J}_{s,0}}(\widetilde{\psi}) \times \mathcal{M}_{\mathcal{J}_{s,t}}(\widetilde{\phi}))\Big), \tag{4.2}$$

where the expected dimensions of ϕ and $\widetilde{\phi}$ are 1 and of ψ and $\widetilde{\psi}$ they are 0. A signed count of (4.2) precisely reproduces (4.1) and hence $c(\boldsymbol{x},\boldsymbol{q}) = 0$ – at least in $\mathbb{Z}_2$-coefficients. To make this work in general, i.e. with coherent orientations, observe that we have the following condition imposed on the sections:

$$\mathfrak{o}_{s,t}(\phi) \wedge \mathfrak{o}_1(\psi) = -\mathfrak{o}_0(\widetilde{\psi}) \wedge \mathfrak{o}_{s,t}(\widetilde{\phi}).$$

We get an identification of orientation systems, ξ say, such that $\widehat{\Phi}_{\mathcal{J}_{s,t}}$ is a chain map between

$$(\widehat{\mathrm{CF}}(\mathcal{H}), \widehat{\partial}^{\mathfrak{o}}_{\mathcal{J}_{s,0}}) \longrightarrow (\widehat{\mathrm{CF}}(\mathcal{H}), \widehat{\partial}^{\xi(\mathfrak{o})}_{\mathcal{J}_{s,1}}).$$

We reverse the direction of the isotopy and define a map $\widehat{\Phi}_{\mathcal{J}_{s,1-t}}$. The compositions

$$\widehat{\Phi}_{\mathcal{J}_{s,1-t}} \circ \widehat{\Phi}_{\mathcal{J}_{s,t}} \text{ and } \widehat{\Phi}_{\mathcal{J}_{s,t}} \circ \widehat{\Phi}_{\mathcal{J}_{s,1-t}}$$

are both chain homotopic to the identity. In the following we will discuss the chain homotopy equivalence for the map $\widehat{\Phi}_{\mathcal{J}_{s,t}} \circ \widehat{\Phi}_{\mathcal{J}_{s,1-t}}$.

Define a path $\mathcal{J}_{s,t}(\tau)$ such that $\mathcal{J}_{s,t}(0) = \mathcal{J}_{s,t} * \mathcal{J}_{s,1-t}$ and $\mathcal{J}_{s,t}(1) = \mathcal{J}_{s,0}$. The existence of this path follows from the fact that we choose the paths inside a contractible set (cf. §2.3 or see [17]). Define the moduli space

$$\mathcal{M}_{\mathcal{J}_{s,t}(\tau),\phi} = \bigcup_{\tau\in[0,1]} \mathcal{M}_{\mathcal{J}_{s,t}(\tau),\phi}.$$

THEOREM 4.2. — *For a generic choice of $\mathcal{J}_{s,t}(\tau)$ the moduli space $\mathcal{M}_{\mathcal{J}_{s,t}(\tau),\phi}$ is a compact manifold of dimension $\mu(\phi)+1$.*

There are two types of boundary components: the one type of boundary component coming from variations of the Whitney disk ϕ which are breaking, bubbling or degenerations and the other type of ends coming from variations of the almost complex structure.

We define a map

$$\widehat{H}_{\mathcal{J}_{s,t}(\tau)}(\boldsymbol{x}) = \sum_{\boldsymbol{y}\in\mathbb{T}_{\boldsymbol{\alpha}}\cap\mathbb{T}_{\boldsymbol{\beta}}} \sum_{\phi\in H(\boldsymbol{x},\boldsymbol{y},-1)} \#\mathcal{M}_{\mathcal{J}_{s,t}(\tau),\phi}\cdot\boldsymbol{y},$$

where $H(\boldsymbol{x},\boldsymbol{y},-1) \subset \pi_2(\boldsymbol{x},\boldsymbol{y})$ are the homotopy classes ϕ with $n_z(\phi)=0$ and expected dimension $\mu(\phi)=-1$. According to Theorem 4.2, the manifold $\mathcal{M}_{\mathcal{J}_{s,t}(\tau),\phi}$ is 0-dimensional. We claim that $\widehat{H}$ is a chain homotopy between $\widehat{\Phi}_{\mathcal{J}_{s,t}} \circ \widehat{\Phi}_{\mathcal{J}_{s,1-t}}$ and the identity. By definition, the equation

$$\widehat{\Phi}_{\mathcal{J}_{s,t}} \circ \widehat{\Phi}_{\mathcal{J}_{s,1-t}} - \mathrm{id} - (\widehat{\partial}_{\mathcal{J}_{s,0}} \circ \widehat{H}_{\mathcal{J}_{s,t}(\tau)} + \widehat{H}_{\mathcal{J}_{s,t}(\tau)} \circ \widehat{\partial}_{\mathcal{J}_{s,1}}) = 0 \tag{4.3}$$

has to hold. Look at the ends of $\mathcal{M}_{\mathcal{J}_{s,t,\tau}}(\psi)$ for $\mu(\psi)=0$. This is a 1-dimensional space, and there are the ends

$$\Big(\bigsqcup_{\psi=\eta*\phi} \widehat{\mathcal{M}}_{\mathcal{J}_{s,0},\eta} \times \mathcal{M}_{\mathcal{J}_{s,t}(\tau),\phi}\Big) \sqcup \Big(\bigsqcup_{\psi=\widetilde{\eta}*\widetilde{\phi}} \mathcal{M}_{\mathcal{J}_{s,t}(\tau),\widetilde{\eta}} \times \widehat{\mathcal{M}}_{\mathcal{J}_{s,1},\widetilde{\phi}}\Big)$$

coming from variations of the Whitney disk, and the ends

$$\mathcal{M}_{\mathcal{J}_{s,t}(0),\psi} \sqcup \mathcal{M}_{\mathcal{J}_{s,t}(1),\psi}$$

coming from variations of the almost complex structure.

These all together precisely produce the coefficients in equation (4.3). Thus, the Floer homology is independent of the choice of $(\mathfrak{j},\eta,V)$-nearly symmetric path. Variations of η and V just change the contractible neighborhood $\mathcal{U}$ around $\xi^g_{sym}(\mathfrak{j})$ containing the admissible almost complex structures. So, the theory is independent of these choices, too. A $\mathfrak{j}'$-nearly symmetric path can be approximated by a $\mathfrak{j}$-symmetric path given that $\mathfrak{j}'$ is close to $\mathfrak{j}$. The set of complex structures on a surface Σ is connected, so step by step one can move from a $\mathfrak{j}$-symmetric path to any $\mathfrak{j}'$-symmetric path.

4.3. Isotopy Invariance

Every isotopy of an arbitrary attaching circle can be divided into two classes: creation/anhillation of pairs of intersection points and isotopies not affecting transversality. An isotopy of an α-circle of the latter type induces an isotopy of $\mathbb{T}_{\boldsymbol{\alpha}}$ in the symmetric product. Compactness of the $\mathbb{T}_{\boldsymbol{\alpha}}$ tells us that there is an ambient isotopy ϕ_t carrying the isotopy. With this isotopy we perturb the admissible path of almost complex structures as

$$\widetilde{\mathcal{J}}_s = (\phi_1^{-1})_* \circ \mathcal{J}_s \circ (\phi_1)_*$$

giving rise to a path of admissible almost complex structures. The diffeomorphism ϕ_1 induces an identification of the chain modules. The moduli spaces defined by $\mathcal{J}_s$ and $\widetilde{\mathcal{J}}_s$ are isomorphic. Hence, if we denote by $\mathcal{H}$ the diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, z)$ and by $\mathcal{H}'$ the diagram $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\beta}, z)$, we get

$$H_*(\widehat{\mathrm{CF}}(\mathcal{H}), \widehat{\partial}_{\mathcal{H}}^{\mathcal{J}_s}) = H_*(\widehat{\mathrm{CF}}(\mathcal{H}'), \widehat{\partial}_{\mathcal{H}'}^{\widetilde{\mathcal{J}}_s}) = H_*(\widehat{\mathrm{CF}}(\mathcal{H}'), \widehat{\partial}_{\mathcal{H}'}^{\mathcal{J}_s}), \tag{4.4}$$

where the last equality follows from the considerations in §4.2. This chain of equalities shows that the discussed isotopies can be interpreted as variations of the almost complex structure.

The creation/cancellation of pairs of intersection points is done with an exact Hamiltonian isotopy supported in a small neighborhood of two attaching circles. We cannot use the methods from §4.2 to create an isomorphism between the associated Floer homologies. At a certain point the isotopy violates transversality as the attaching tori do not intersect transversely. Thus, the arguments of §4.2 for the right equality in (4.4) break down.

Consider an exact Hamiltonian isotopy ψ_t of an α-curve generating a canceling pair of intersections with a β-curve. We will just sketch the approach used in this context, since the ideas are similar to the ideas introduced in §4.2.

Define $\pi_2^{\Psi_t}(\boldsymbol{x}, \boldsymbol{y})$ as the set of Whitney disks ϕ with dynamic boundary conditions in the following sense:

$$\begin{array}{rcl} \phi(i) & = & \boldsymbol{x}, \\ \phi(-i) & = & \boldsymbol{y}, \\ \phi(0+it) & \in & \Psi_t(\mathbb{T}_{\boldsymbol{\alpha}}) \\ \phi(1+it) & \subset & \mathbb{T}_{\boldsymbol{\beta}} \end{array}$$

for all $t \in \mathbb{R}$. Spoken geometrically, we follow the isotopy with the α-boundary of the Whitney disk. Correspondingly, we define the moduli spaces of $\mathcal{J}_s$-holomorphic Whitney disks with dynamic boundary conditions as $\mathcal{M}^{\Psi_t}(\boldsymbol{x},\boldsymbol{y})$. For $\boldsymbol{x}\in\mathbb{T}_{\boldsymbol{\alpha}}\cap\mathbb{T}_{\boldsymbol{\beta}}$ define

$$\widehat{\Gamma}_{\Psi_t}(\boldsymbol{x})=\sum_{\boldsymbol{y}\in\mathbb{T}_{\boldsymbol{\alpha}}\cap\mathbb{T}_{\boldsymbol{\beta}}}\;\sum_{\phi\in H_t(\boldsymbol{x},\boldsymbol{y},0)}\#\mathcal{M}^{\Psi_t}_{\mathcal{J}_s,\phi}\cdot\boldsymbol{y},$$

where $H_t(\boldsymbol{x},\boldsymbol{y},0)\subset\pi_2^{\Psi_t}(\boldsymbol{x},\boldsymbol{y})$ are the homotopy classes with expected dimension $\mu=0$ and $n_z=0$. Using the low-dimensional model introduced in §2, Ozsváth and Szabó prove the following property.

Theorem 4.3 (see [17], §7.3). — *There exists a t-independent energy bound for holomorphic Whitney disks independent of its homotopy class.*

The existence of this energy bound shows that there are Gromov compactifications of the moduli spaces of Whitney disks with dynamic boundary conditions.

Theorem 4.4. — *The map $\widehat{\Gamma}_{\Psi_t}$ is a chain map. Using the inverse isotopy we define $\widehat{\Gamma}_{\Psi_{1-t}}$ such that the compositions $\widehat{\Gamma}_{\Psi_t}\circ\widehat{\Gamma}_{\Psi_{1-t}}$ and $\widehat{\Gamma}_{\Psi_{1-t}}\circ\widehat{\Gamma}_{\Psi_t}$ are chain homotopic to the identity.*

The proof follows the same lines as in §4.2. We leave the proof to the interested reader.

4.4. Handle slide Invariance

4.4.1. The Pair-of-Pants Product

In this section we will introduce the Heegaard Floer incarnation of the pair-of-pants product and with it associate to cobordisms maps between the Floer homologies of their boundary components. In case the cobordisms are induced by handle slides the associated maps are isomorphisms on the level of homology. The maps we will introduce will count holomorphic triangles in the symmetric product with appropriate boundary conditions. We have to discuss well-definedness of the maps and that they are chain maps. To do that we have to follow similar lines as it was done for the differential. Because of the strong parallels we will shorten the discussion here. We strongly advise the reader to first read §2 before continuing.

Definition 4.5. — *A set of data $(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},\gamma)$, where Σ is a surface of genus g and $\boldsymbol{\alpha}$, $\boldsymbol{\beta}$, $\boldsymbol{\gamma}$ three sets of attaching circles, is called a* **Heegaard triple diagram**.

We denote the 3-manifolds determined be taking pairs of these attaching circles as $Y_{\boldsymbol{\alpha\beta}}$, $Y_{\boldsymbol{\beta\gamma}}$ and $Y_{\boldsymbol{\alpha\gamma}}$. We fix a point $z \in \Sigma\backslash\{\boldsymbol{\alpha} \cup \boldsymbol{\beta} \cup \boldsymbol{\gamma}\}$ and define a product

$$\widehat{f}_{\boldsymbol{\alpha},\boldsymbol{\beta\gamma}}\colon \widehat{\mathrm{CF}}(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, z) \otimes \widehat{\mathrm{CF}}(\Sigma, \boldsymbol{\beta}, \boldsymbol{\gamma}, z) \longrightarrow \widehat{\mathrm{CF}}(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\gamma}, z)$$

by counting holomorphic triangles with suitable boundary conditions: A **Whitney triangle** is a map $\phi\colon \Delta \longrightarrow \mathrm{Sym}^g(\Sigma)$ with boundary conditions as illustrated in Figure 4. We call the respective boundary segments its **α-, β- and γ-boundary**. The boundary points, as should be clear from the picture, are $\boldsymbol{x} \in \mathbb{T}_{\boldsymbol{\alpha}} \cap \mathbb{T}_{\boldsymbol{\beta}}$, $\boldsymbol{q} \in \mathbb{T}_{\boldsymbol{\alpha}} \cap \mathbb{T}_{\gamma}$ and $\boldsymbol{y} \in \mathbb{T}_{\boldsymbol{\beta}} \cap \mathbb{T}_{\gamma}$. The set of homotopy classes of Whitney disks connecting $\boldsymbol{x}$, $\boldsymbol{q}$ and $\boldsymbol{y}$ is denoted by $\pi_2(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{q})$.

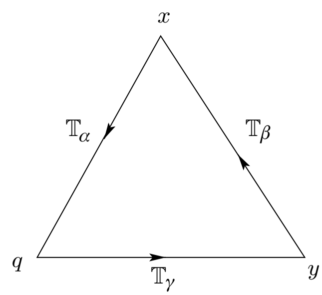

Figure 4. — A Whitney triangle and its boundary conditions.

Denote by $\mathcal{M}_\phi^\Delta$ the moduli space of holomorphic triangles in the homotopy class of ϕ. Analogous to the case of disks we denote by $\mu(\phi)$ its expected/formal dimension. For $\boldsymbol{x} \in \mathbb{T}_{\boldsymbol{\alpha}} \cap \mathbb{T}_{\boldsymbol{\beta}}$ define

$$\widehat{f}_{\boldsymbol{\alpha},\boldsymbol{\beta\gamma}}(\boldsymbol{x} \otimes \boldsymbol{y}) = \sum_{w \in \mathbb{T}_{\boldsymbol{\alpha}} \cap \mathbb{T}_{\boldsymbol{\gamma}}} \sum_{\phi \in H(\boldsymbol{x},\boldsymbol{y},\boldsymbol{q},0)} \#\mathcal{M}_\phi^\Delta \cdot \boldsymbol{q},$$

where $H(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{q}, 0) \subset \pi_2(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{q})$ is the subset with $\mu = 0$ and $n_z = 0$. The set of homotopy classes of Whitney disks fits into an exact sequence

$$0 \longrightarrow \pi_2(\mathrm{Sym}^g(\Sigma)) \longrightarrow \pi_2(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{q}) \longrightarrow \ker(n_z) \longrightarrow 0, \qquad (4.5)$$

where n_z provides a splitting for the sequence. We define

$$X_{\boldsymbol{\alpha\beta\gamma}} = \frac{(\Delta \times \Sigma) \cup e_{\boldsymbol{\alpha}} \times U_{\boldsymbol{\alpha}} \cup e_{\boldsymbol{\beta}} \times U_{\boldsymbol{\beta}} \cup e_{\gamma} \times U_{\gamma}}{(e_{\boldsymbol{\alpha}} \times \Sigma) \sim (e_{\boldsymbol{\alpha}} \times \partial U_{\boldsymbol{\alpha}}), (e_{\boldsymbol{\beta}} \times \Sigma) \sim (e_{\boldsymbol{\beta}} \times \partial U_{\boldsymbol{\beta}}), (e_{\gamma} \times \Sigma) \sim (e_{\gamma} \times \partial U_{\gamma})},$$

where $U_{\boldsymbol{\alpha}}$, $U_{\boldsymbol{\beta}}$ and U_γ are the handlebodies determined by the 2$-$handles associated to the attaching circles $\boldsymbol{\alpha}$, $\boldsymbol{\beta}$ and $\boldsymbol{\gamma}$, and $e_{\boldsymbol{\alpha}}$, $e_{\boldsymbol{\beta}}$ and e_γ are the edges of the triangle Δ. The manifold $X_{\boldsymbol{\alpha\beta\gamma}}$ is 4-dimensional with boundary

$$\partial X_{\boldsymbol{\alpha\beta\gamma}} = Y_{\boldsymbol{\alpha\beta}} \sqcup Y_{\boldsymbol{\beta\gamma}} \sqcup -Y_{\boldsymbol{\alpha\gamma}}.$$

Lemma 4.6 (pp. 1094–1095 of [17]). — *The kernel of n_z equals $H_2(X_{\boldsymbol{\alpha\beta\gamma}};\mathbb{Z})$*

Combining (4.5) with Lemma 4.6 we get an exact sequence

$$0 \longrightarrow \pi_2(\mathrm{Sym}^g(\Sigma)) \longrightarrow \pi_2(\boldsymbol{x},\boldsymbol{y},\boldsymbol{q}) \xrightarrow{\mathcal{H}} H_2(X_{\boldsymbol{\alpha\beta\gamma}};\mathbb{Z}) \longrightarrow 0, \tag{4.6}$$

where $\mathcal{H}$ is defined similarly as for disks (cf. §2.2). Of course there is a low-dimensional model for triangles and the discussion we have done for disks carries over verbatim for triangles. The condition $n_z = 0$ makes the product $f_{\boldsymbol{\alpha};\boldsymbol{\beta\gamma}}$ well-defined in case $H_2(X_{\boldsymbol{\alpha\beta\gamma}};\mathbb{Z})$ is trivial. Analogous to our discussion for Whitney disks and the differential, we have to include a condition controlling the periodic triangles, i.e. the triangles associated to elements in $H_2(X_{\boldsymbol{\alpha\beta\gamma}};\mathbb{Z})$. A domain $\mathcal{D}$ of a triangle is called **triply-periodic** if its boundary consists of a sum of α-,β- and γ-curves such that $n_z = 0$.

Definition 4.7. — *A pointed triple diagram $(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},\boldsymbol{\gamma},z)$ is called* **weakly admissible** *if all triply-periodic domains $\mathcal{D}$ which can be written as a sum of doubly-periodic domains have both positive and negative coefficients.*

This condition is the natural transfer of weak-admissibility from disks to triangles. One can show that for given $j,k \in \mathbb{Z}$ there exist just a finite number of Whitney triangles $\phi \in \pi_2(\boldsymbol{x},\boldsymbol{y},\boldsymbol{q})$ with $\mu(\phi) = j$, $n_z(\phi) = k$ and $\mathcal{D}(\phi) \geqslant 0$.

For a given homotopy class $\psi \in \pi_2(\boldsymbol{x},\boldsymbol{y},\boldsymbol{q})$ with $\mu(\psi) = 1$ we compute the ends by shrinking a properly embedded arc to a point (see the description of convergence in §2.3). There are three different ways to do this in a triangle. Each time we get a concatenation of a disk with a triangle. By boundary orientations we see that each of these boundary components contributes to one of the terms in the following sum

$$\widehat{f}_{\boldsymbol{\alpha},\boldsymbol{\beta\gamma}} \circ (\widehat{\partial}_{\boldsymbol{\alpha\beta}}(\boldsymbol{x}) \otimes \boldsymbol{y}) + \widehat{f}_{\boldsymbol{\alpha},\boldsymbol{\beta\gamma}} \circ (\boldsymbol{x} \otimes \widehat{\partial}_{\boldsymbol{\beta\gamma}}(\boldsymbol{y})) - \widehat{\partial}_{\boldsymbol{\alpha\gamma}} \circ \widehat{f}_{\boldsymbol{\alpha},\boldsymbol{\beta\gamma}}(\boldsymbol{x} \otimes \boldsymbol{y}). \tag{4.7}$$

Conversely, the coefficient at any of these terms is given by a product of signed counts of moduli spaces of disks and moduli spaces of triangles and hence – by gluing – comes from one of these contributions. The sum in (4.7) vanishes, showing that $\widehat{f}_{\boldsymbol{\alpha},\boldsymbol{\beta\gamma}}$ descends to a pairing $\widehat{f}^{*}_{\boldsymbol{\alpha},\boldsymbol{\beta\gamma}}$ between the Floer homologies.

4.4.2. Holomorphic rectangles

Recall that the set of biholomorphisms of the unit disk is a 3-dimensional connected family. If we additionally fix a point we decrease the dimension of that family by one. A better way to formulate this is to say that the set of biholomorphishms of the unit disk with one fixed point is a 2-dimensional family. Fixing two further points reduces to a 0-dimensional set. If we additionally fix a fourth point the rectangle together with these four points uniquely defines a conformal structure. Variation of the fourth point means a variation of the conformal structure. Indeed one can show that there is a uniformization of a holomorphic rectangle, i.e. a rectangle with fixed conformal structure, which we denote by $\square$,

$$\square \longrightarrow [0,l] \times [0,h],$$

where the ratio l/h uniquely determines the conformal structure. With this uniformization we see that $\mathcal{M}(\square) \cong \mathbb{R}$. The uniformization is area-preserving and converging to one of the ends of $\mathcal{M}(\square)$ means to stretch the rectangle infinitely until it breaks at the end into a concatenation of two triangles.

Theorem 4.8 (see Lemma 9.6 of [17]). — *Given another set of attaching circles* $\boldsymbol{\delta}$ *defining a map* $\widehat{f}_{\boldsymbol{\alpha},\boldsymbol{\gamma\delta}}$*, the following equality holds:*

$$\widehat{f}^{*}_{\boldsymbol{\alpha},\boldsymbol{\beta\gamma}}(\widehat{f}^{*}_{\boldsymbol{\alpha},\boldsymbol{\gamma\delta}}(\,\cdot\otimes\cdot\,)\otimes\cdot\,) - \widehat{f}^{*}_{\boldsymbol{\alpha},\boldsymbol{\beta\delta}}(\,\cdot\otimes\widehat{f}^{*}_{\boldsymbol{\beta},\boldsymbol{\gamma\delta}}(\,\cdot\otimes\cdot\,)) = 0. \tag{4.8}$$

This property is called **associativity**.

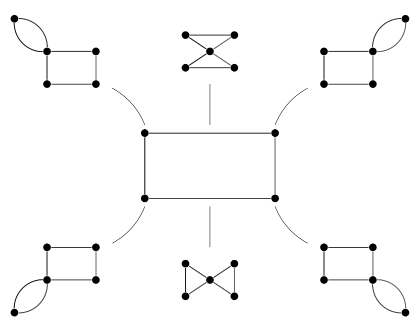

Figure 5. — Ends of the moduli space of holomorphic rectangles.

If we count holomorphic Whitney rectangles with boundary conditions in $\boldsymbol{\alpha}$, $\boldsymbol{\beta}$, $\boldsymbol{\gamma}$ and $\boldsymbol{\delta}$ and with $\mu = 1$ (see Definition 2.6) the ends of the associated moduli space will look like pictured in Figure 5. Note that we are talking about holomorphicity with respect to an arbitrary conformal structure on the rectangle. There will be two types of ends. We will have a degeneration into a concatenation of triangles by variation of the conformal structure on the rectangle and breaking into a concatenation of a rectangle with a disk by variation of the rectangle. By Figure 5 an appropriate count of holomorphic rectangles will be a natural candidate for a chain homotopy proving equation (4.8). Define a pairing

$$H\colon \widehat{\mathrm{CF}}(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},z)\otimes\widehat{\mathrm{CF}}(\Sigma,\boldsymbol{\beta},\boldsymbol{\gamma},z)\otimes\widehat{\mathrm{CF}}(\Sigma,\boldsymbol{\gamma},\boldsymbol{\delta},z)\longrightarrow\widehat{\mathrm{CF}}(\Sigma,\boldsymbol{\alpha},\boldsymbol{\delta},z)$$

by counting holomorphic Whitney rectangles with boundary components as indicated in Figure 6 and $\mu=0$. By counting ends of the moduli space of holomorphic rectangles with $\mu=1$ we have six contributing ends. These ends are pictured in Figure 5. The four ends coming from breaking contribute to

$$\widehat{\partial}\circ H(\,\cdot\otimes\cdot\otimes\cdot\,)+H\circ\widehat{\partial}(\,\cdot\otimes\cdot\otimes\cdot\,). \tag{4.9}$$

In addition there are two ends coming from degenerations of the conformal structure on the rectangle. These give rise to

$$\widehat{f}_{\boldsymbol{\alpha},\boldsymbol{\beta}\boldsymbol{\gamma}}(\widehat{f}_{\boldsymbol{\alpha},\boldsymbol{\gamma}\boldsymbol{\delta}}(\,\cdot\otimes\cdot\,)\otimes\cdot\,)-\widehat{f}_{\boldsymbol{\alpha},\boldsymbol{\beta}\boldsymbol{\delta}}(\,\cdot\otimes\widehat{f}_{\boldsymbol{\beta},\boldsymbol{\gamma}\boldsymbol{\delta}}(\,\cdot\otimes\cdot\,)). \tag{4.10}$$

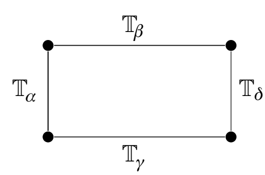

Figure 6. — The boundary conditions of rectangles for the definition of H.

We see that the sum of (4.9) and (4.10) vanishes, showing that H is a chain homotopy proving associativity.

4.4.3. Special Case – Handle Slides

Handle slides provide special Heegaard triple diagrams. Let $(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},z)$ be an admissible pointed Heegaard diagram and define $(\Sigma,\boldsymbol{\alpha},\boldsymbol{\gamma},z)$ by handle sliding β_1 over β_2. We push the γ_i off the β_i to make them intersect transversely in two cancelling points. This defines a triple diagram, and obviously $Y_{\boldsymbol{\beta}\boldsymbol{\gamma}}$ equals the connected sum $\#^g(\mathbb{S}^2\times\mathbb{S}^1)$.

A very important observation is that the Heegaard Floer groups of connected sums of $\mathbb{S}^2\times\mathbb{S}^1$ admit a top-dimensional generator. By Example 2.3 and Theorem 4.1,

$$\widehat{\mathrm{HF}}(\#^{g-1}(\mathbb{S}^2\times\mathbb{S}^1),\mathfrak{o}_0)\cong\mathbb{Z}^{2g-2}\cong H_*(T^g;\mathbb{Z}),$$

where the last identification is done using the $\bigwedge^*(H_1/\mathrm{Tor})$-module structure on the Heegaard Floer homology groups (see [17, Lemma 9.1] and [17, Proposition 4.17]). We claim that the behavior of the Heegaard Floer groups under connected sums can be carried over to the module structure, and thus it remains to show the assertion for the case $g = 1$. But this is not hard to see.

Each pair (β_i, γ_i) has two intersections $\boldsymbol{x}_i^+$ and $\boldsymbol{x}_i^-$. Which one is denoted how is determined by the following criterion: there is a disk-shaped domain connecting $\boldsymbol{x}_i^+$ with $\boldsymbol{x}_i^-$ with boundary in β_i and γ_i. The point

$$\boldsymbol{x}^+ = \{\boldsymbol{x}_1^+, \dots, \boldsymbol{x}_g^+\}$$

is a cycle whose associated homology class is the top-dimensional generator we denote by $\widehat{\Theta}_{\beta\gamma}$. For a detailed treatment of the top-dimensional generator we point the reader to [17, §9.1].

Plugging in the generator we define a map

$$\widehat{F}_{\boldsymbol{\alpha},\boldsymbol{\beta\gamma}} = \widehat{f}^*_{\boldsymbol{\alpha},\boldsymbol{\beta\gamma}}(\,\cdot \otimes \widehat{\Theta}_{\boldsymbol{\beta\gamma}})\colon \widehat{\mathrm{HF}}(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, z) \longrightarrow \widehat{\mathrm{HF}}(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\gamma}, z)$$

between the associated Heegaard Floer groups. Our intention is to show that this is an isomorphism.

We can slide the γ_1 back over γ_2 to give another set of attaching circles we denote by δ. Of course we make the curves intersecting all other sets of attaching circles transversely and introduce pairs of intersections points of the δ-curves with the γ-and β-curves.

Let $\widehat{F}_{\boldsymbol{\alpha},\boldsymbol{\gamma\delta}}$ be the associated map. Then the associativity given in (4.8) translates into

$$\widehat{f}^*_{\boldsymbol{\alpha},\boldsymbol{\beta\gamma}}(\widehat{f}^*_{\boldsymbol{\alpha},\boldsymbol{\gamma\delta}}(\,\cdot \otimes \widehat{\Theta}_{\boldsymbol{\gamma\delta}}) \otimes \widehat{\Theta}_{\boldsymbol{\beta\gamma}}) - \widehat{f}^*_{\boldsymbol{\alpha},\boldsymbol{\beta\delta}}(\,\cdot \otimes \widehat{f}^*_{\boldsymbol{\beta},\boldsymbol{\gamma\delta}}(\widehat{\Theta}_{\boldsymbol{\beta\gamma}} \otimes \widehat{\Theta}_{\boldsymbol{\gamma\delta}})) = 0.$$

The proof of the following lemma will be done in detail. It is the first explicit calculation using the low-dimensional model in a non-trivial manner.

Lemma 4.9 (see Lemma 9.7 of [17]). — *Given the map $\widehat{f}_{\boldsymbol{\alpha},\boldsymbol{\gamma\delta}}$, we have*

$$\widehat{f}_{\boldsymbol{\beta},\boldsymbol{\gamma\delta}}(\widehat{\Theta}_{\boldsymbol{\beta\gamma}} \otimes \widehat{\Theta}_{\boldsymbol{\gamma\delta}}) = \widehat{\Theta}_{\boldsymbol{\beta\delta}}.$$

Hence, we have $\widehat{F}_{\boldsymbol{\beta},\boldsymbol{\gamma\delta}}(\widehat{\Theta}_{\boldsymbol{\beta\gamma}}) = \widehat{\Theta}_{\boldsymbol{\beta\delta}}$.

Proof. — The complement of the β-circles in Σ is a sphere with holes. We have a precise description of how the sets $\boldsymbol{\gamma}$ and $\boldsymbol{\delta}$ look like relative to $\boldsymbol{\beta}$. The Heegaard surface cut open along the β-curves can be identified with a sphere with holes by using an appropriate diffeomorphism. Doing so, the diagram $(\Sigma, \boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{\delta})$ will look like given in Figure 7. In each component we have to have a close look at the domains $\mathcal{D}_1$, $\mathcal{D}_2$ and $\mathcal{D}_3$. To improve the illustration in the picture we have separated them. There are exactly two domains contributing to holomorphic triangles with boundary points in $\{\widehat{\Theta}_{\boldsymbol{\beta\gamma}}, \widehat{\Theta}_{\boldsymbol{\gamma\delta}}\}$, namely $\mathcal{D}_1$ and $\mathcal{D}_3$. The domain $\mathcal{D}_3$ can be written as a sum of $\mathcal{D}_1$ and $\mathcal{D}_2$, the former carrying $\mu = 0$, the latter carrying $\mu = 1$. Consequently, every homotopy class of triangles using $\mathcal{D}_3$-domains can be written as a concatenation of a triangle with a disk with the expected dimensions greater than or equal to those mentioned. Consequently, the expected dimension of the triangle using a $\mathcal{D}_3$-domain is strictly bigger than zero and thus does not contribute to $\widehat{F}_{\boldsymbol{\beta},\boldsymbol{\gamma\delta}}(\widehat{\Theta}_{\boldsymbol{\beta\gamma}} \otimes \widehat{\Theta}_{\boldsymbol{\gamma\delta}})$. All holomorphic triangles relevant to us have domains which are a sum of $\mathcal{D}_1$-domains. Taking boundary conditions into account we see that we need a $\mathcal{D}_1$-domain in each component. Thus, there is a unique homotopy class of triangles interesting to us. By the Riemann mapping theorem there is a unique holomorphic map $\widehat{\phi}\colon \widehat{D^2} \longrightarrow \Sigma$ from a surface with boundary whose associated domain equals the sum of $\mathcal{D}_1$-domains. The map $\widehat{\phi}$ is a biholomorphism and, thus, $\widehat{D^2}$ is a disjoint union of triangles. The uniqueness of $\widehat{\phi}$ tells us that the number of elements in the associated moduli space equals the number of non-equivalent g-fold branched coverings $\widehat{D^2} \longrightarrow D^2$. Since $\widehat{D^2}$ is a union of g disks, this covering is unique, too (up to equivalence) and, thus, the associated moduli space is a one-point space. $\square$

Lemma 4.9 and (4.4) combine to give the composition law

$$\widehat{F}_{\boldsymbol{\alpha},\boldsymbol{\beta\delta}} = \widehat{F}_{\boldsymbol{\alpha},\boldsymbol{\gamma\delta}} \circ \widehat{F}_{\boldsymbol{\alpha},\boldsymbol{\beta\gamma}}.$$

We call a holomorphic triangle **small** if it is supported within the thin strips of isotopy between $\boldsymbol{\beta}$ and $\boldsymbol{\delta}$.

LEMMA 4.10 (see Lemma 9.10 of [17]). — *Let $F\colon A \longrightarrow B$ be a map of filtered groups such that F can be decomposed into $F_0 + l$, where F_0 is a filtration-preserving isomorphism and $l(\boldsymbol{x}) < F_0(\boldsymbol{x})$. Then, if the filtration on B is bounded from below, the map F is an isomorphism of groups.*

There are two important observations to make. The first is that we can equip the chain complexes with a filtration, called the **energy filtration** (cf. [17, p. 1122]), which is indeed bounded from below. In this situation the top-dimensional generator $\widehat{\Theta}_{\boldsymbol{\beta\delta}}$ is generated by a single intersection point $\boldsymbol{x}^+ \in \mathbb{T}_{\boldsymbol{\beta}} \cap \mathbb{T}_{\delta}$. The map $\widehat{F}_{\boldsymbol{\alpha},\boldsymbol{\beta\delta}}$ is induced by

$$\widehat{f}_{\boldsymbol{\alpha},\boldsymbol{\beta\delta}}(\,\cdot \otimes \boldsymbol{x}^+),$$

which in turn can be decomposed into a sum of f_0 and l, where f_0 counts small holomorphic triangles and l those triangles whose support is not contained in the thin strips of isotopy between $\boldsymbol{\beta}$ and $\boldsymbol{\delta}$. The map f_0 is filtration preserving and l, if the δ-curves are close enough to the β-curves, strictly decreasing. By Lemma 4.10 the map $\widehat{F}_{\boldsymbol{\alpha},\boldsymbol{\beta}\boldsymbol{\delta}}$ is an isomorphism between the associated Heegaard Floer homologies.

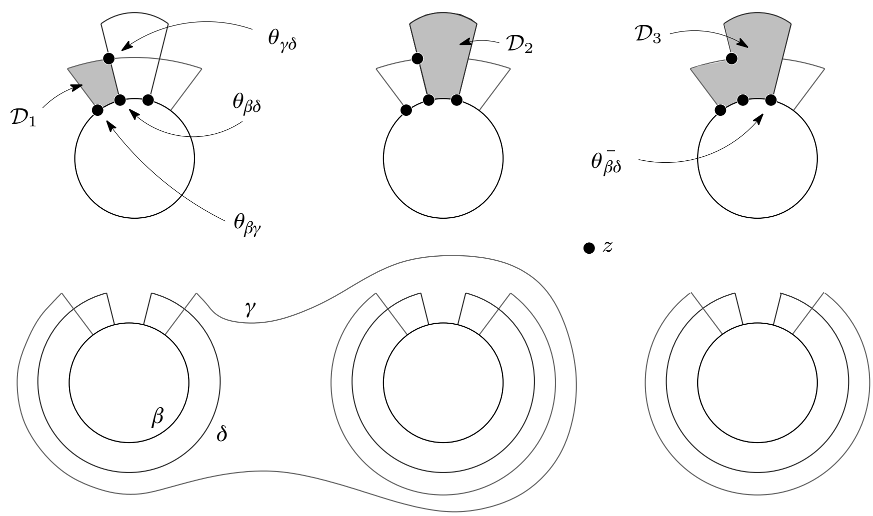

Figure 7. — The Heegaard surface cut open along the β-curves.

To conclude topological invariance we have to see that the following claim is true.

Theorem 4.11. — *Two pointed admissible Heegaard diagrams associated to a 3-manifold are equivalent after a finite sequence of Heegaard moves, each of them connecting two admissible Heegaard diagrams, which can be done in the complement of the base-point z.*

The only situation where the point z seems to be an obstacle arises when trying to isotope an attaching circle, α_1 say, over the base-point z. But observe that cutting the α-circles out of Σ we get a sphere with holes. We can isotope α_1 freely and pass the holes by handle slides. Thus, the requirement not to pass z is not an obstruction at all. Instead of passing z we can go *the other way around the surface* by isotopies and handle slides.

5. Knot Floer Homologies

Knot Floer homology is a variant of the Heegaard Floer homology of a manifold. Recall that the Heegaard diagrams used in Heegaard Floer theory come from handle decompositions relative to a splitting surface. Given a knot $K \subset Y$, we can restrict to a subclass of Heegaard diagrams by requiring the handle decomposition to come from a handle decomposition of $\overline{Y \backslash \nu K}$ relative to its boundary. Note that at first in the literature the knot Floer variants were defined for homologically trivial knots only. However, the definition can be carried over nearly one-to-one to give a well-defined topological invariant for arbitrary knot classes. However, in the homologically trivial case it is possible to subdivide the groups in a special manner giving rise to a refined invariant, which cannot be defined in the non-trivial case.

Given a knot $K \subset Y$, we can specify a certain subclass of Heegaard diagrams.

Definition 5.1. — *A Heegaard diagram $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$ is said to be* **subordinate** *to the knot K if K is isotopic to a knot lying in Σ and K intersects β_1 once, transversely and is disjoint from the other β-circles.*

Since K intersects β_1 once and is disjoint from the other β-curves we know that K intersects the core disk of the 2-handle, represented by β_1, once and is disjoint from the others (after possibly isotoping the knot K).

Lemma 5.2. — *Every pair (Y, K) admits a Heegaard diagram subordinate to K.*

Proof. — By surgery theory (see [7, p. 104]) we know that there is a handle decomposition of $Y \backslash \nu K$, i.e.

$$Y \backslash \nu K \cong (T^2 \times [0,1]) \cup_\partial h_2^1 \cup_\partial \dots h_g^1 \cup_\partial h_1^2 \cup_\partial \dots \cup_\partial h_g^2 \cup_\partial h^3$$

We close up the boundary $T^2 \times \{0\}$ with an additional 2-handle h_1^{2*} and a 3-handle h^3 to obtain

$$Y \cong h^3 \cup_\partial h_1^{2*} \cup_\partial (T^2 \times I) \cup_\partial h_2^1 \cup_\partial \dots h_g^1 \cup_\partial h_1^2 \cup_\partial \dots \cup_\partial h_g^2 \cup_\partial h^3. \quad (5.1)$$

We may interpret $h^3 \cup_\partial h_1^{2*} \cup_\partial (T^2 \times [0,1])$ as a 0-handle h^0 and a 1-handle h_1^{1*}. Hence, we obtain the following decomposition of Y:

$$h^0 \cup_\partial h_1^{1*} \cup_\partial h_2^1 \cup_\partial \dots \cup_\partial h_g^1 \cup_\partial h_1^2 \cup_\partial \dots \cup_\partial h_g^2 \cup_\partial h^3.$$

We get a Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$ where $\boldsymbol{\alpha} = \alpha_1^* \cup \{\alpha_2, \dots, \alpha_g\}$ are the co-cores of the 1-handles and $\boldsymbol{\beta} = \{\beta_1, \dots, \beta_g\}$ are the attaching circles of the 2-handles. $\square$

Having fixed such a Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$ we can encode the knot K in a pair of points. After isotoping K onto Σ, we fix a small interval I in K containing the intersection point $K \cap \beta_1$. This interval should be chosen small enough such that I does not contain any other intersections of K with other attaching curves. The boundary ∂I of I determines two points in Σ that lie in the complement of the attaching circles, i.e. $\partial I = z - w$, where the orientation of I is given by the knot orientation. This leads to a doubly-pointed Heegaard diagram $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, w, z)$. Conversely, a doubly-pointed Heegaard diagram uniquely determines a topological knot class: Connect z with w in the complement of the attaching circles $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}\backslash\beta_1$ with an arc δ that crosses β_1 once. Connect w with z in the complement of $\boldsymbol{\beta}$ using an arc γ. The union $\delta \cup \gamma$ is represents the knot klass K represents. The orientation on K is given by orienting δ such that $\partial\delta = z - w$. If we use a different path $\widetilde{\gamma}$ in the complement of $\boldsymbol{\beta}$, we observe that $\widetilde{\gamma}$ is isotopic to γ (in Y): Since $\Sigma\backslash\boldsymbol{\beta}$ is a sphere with holes an isotopy can move γ across the holes by doing handle slides. Isotope the knot along the core disks of the 2-handles to cross the holes of the sphere. Indeed, the knot class does not depend on the specific choice of δ-curve.

The knot chain complex $\widehat{\mathrm{CFK}}(\mathcal{H})$ is the free $\mathbb{Z}_2$-module (or $\mathbb{Z}$-module) generated by the intersections $\mathbb{T}_{\boldsymbol{\alpha}} \cap \mathbb{T}_{\boldsymbol{\beta}}$. The boundary operator $\widehat{\partial}^w_{\mathcal{H}}$, for $\boldsymbol{x} \in \mathbb{T}_{\boldsymbol{\alpha}} \cap \mathbb{T}_{\boldsymbol{\beta}}$, is defined by

$$\partial^{\bullet;\bullet}_{\mathcal{H}}(\boldsymbol{x}) = \sum_{\boldsymbol{y} \in \mathbb{T}_{\boldsymbol{\alpha}} \cap \mathbb{T}_{\boldsymbol{\beta}}} \sum_{\phi \in H(\boldsymbol{x}, \boldsymbol{y}, 1)} \#\widehat{\mathcal{M}}_\phi \cdot \boldsymbol{y},$$

where $H(\boldsymbol{x}, \boldsymbol{y}, 1) \subset \pi_2(\boldsymbol{x}, \boldsymbol{y})$ are the homotopy classes with $\mu = 1$ and $n_z = n_w = 0$. We denote by $\widehat{\mathrm{HFK}}(Y, K)$ the associated homology theory $H_*(\widehat{\mathrm{CFK}}(\mathcal{H}), \partial^{\bullet;\bullet}_{\mathcal{H}})$. The crucial observation for showing invariance is that two Heegaard diagrams subordinate to a given knot can be connected by moves that *respect the knot complement.*

Lemma 5.3 *([14]). — Let* $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, z, w)$ *and* $(\Sigma', \boldsymbol{\alpha}', \boldsymbol{\beta}', z', w')$ *be two Heegaard diagrams subordinate to a given knot* $K \subset Y$*. Let* I *denote the interval inside* K *connecting* z *with* w*, interpreted as sitting in* Σ*. Then these two diagrams are isomorphic after a sequence of the following moves:*

(m_1) Handle slides and isotopies among the α*-curves. These isotopies may not cross* I*.*

(m_2) Handle slides and isotopies among the $\beta_2, \dots, \beta_g$*. These isotopies may not cross* I*.*

(m_3) Handle slides of β_1 *over the* $\beta_2, \dots, \beta_g$ *and isotopies.*

(m_4) Stabilizations/destabilizations.

For the convenience of the reader we include a short proof of this lemma.

Proof. — By [7, Theorem 4.2.12] we can transform two relative handle decompositions into each other by isotopies, handle slides and handle creation/annihilation of the handles written at the right of $T^2 \times [0,1]$ in (5.1). Observe that the 1-handles may be isotoped along the boundary $T^2 \times \{1\}$. Thus, we can transform two Heegaard diagrams into each other by handle slides, isotopies, creation/annihilation of the 2-handles h_i^2 and we may slide the h_i^1 over h_j^1 and over h_1^{1*} (the latter corresponds to h_i^1 sliding over the boundary $T^2 \times \{1\} \subset T^2 \times I$ by an isotopy). But we are not allowed to move h_1^{1*} off the 0-handle. In this case we would lose the relative handle decomposition. In terms of Heegaard diagrams we see that these moves exactly translate into the moves given in (m_1) to (m_4). Just note that sliding the h_i^1 over h_1^{1*}, in the dual picture, looks like sliding h_1^{2*} over the h_i^2. This corresponds to the move (m_3). $\square$

Proposition 5.4. — *Let $K \subset Y$ be an arbitrary knot. The knot Floer homology group $\widehat{\mathrm{HFK}}(Y,K)$ is a topological invariant of the knot type of K in Y. These homology groups split with respect to* $\mathrm{Spin}^c(Y)$.

Proof. — Given one of the moves (m_1) to (m_4), the associated Heegaard Floer homologies are isomorphic, which is shown using one of the isomorphisms given in §4. Each of these maps is defined by counting holomorphic disks with punctures, whose properties are shown by defining maps by counting holomorphic disks with punctures.

Isotopies/Almost Complex Structure. Denote by $\mathcal{J}$ the path of almost complex structures used in the definition of the Heegaard Floer homologies. Let M be an isotopy or perturbation of J. Let $\widehat{\Phi}$ be the isomorphism induced by M. We split the isomorphism up into

$$\widehat{\Phi} = \widehat{\Phi}^w + \widehat{\Phi}^{\neq},$$

where $\widehat{\Phi}^w$ is defined by counting holomorphic disks with punctures (for a precise definition look into §4.2 and §4.3) that fulfill $n_w = 0$. Let us denote with $\mathcal{M}_0$ the associated moduli space used to define the map $\widehat{\Phi}$. The index indicates the value of the index μ. The chain map property of $\widehat{\Phi}$ was shown by counting ends of $\mathcal{M}_1$ which contains the same objects we needed to define $\widehat{\Phi}$ but now with the index fulfilling $\mu = 1$ (see Definition 2.6). We restrict our attention to $\mathcal{M}_0^w$ and $\mathcal{M}_1^w$, the superscript w indicates that we look at the holomorphic elements in $\mathcal{M}_0$ (or $\mathcal{M}_1$ respectively) with intersection number $n_w = 0$: The additivity of the intersection number n_w and the positivity of intersections guarantees that the ends of $\mathcal{M}_1^w$ lie within the space $\mathcal{M}_0^w$ provided that M respects the point w. If M is an isotopy, respecting w means that no attaching circle crosses the point w. If M is a perturbation of $\mathcal{J}$, respecting w means, that we perturb $\mathcal{J}$ through nearly symmetric almost complex structures such that V (cf. Definition 2.19) also contains $\{w\} \times \mathrm{Sym}^{g-1}(\Sigma)$. Hence, we have the equality

$$(\partial\mathcal{M}_1)^w = \partial\mathcal{M}_1^w.$$

Thus, $\widehat{\Phi}^w$ has to be a chain map between the respective knot Floer homologies. To show that $\widehat{\Phi}$ is an isomorphism, we invert the move M we have done and construct the associated morphism $\widehat{\Psi}$. To show that $\widehat{\Psi}$ is the inverse, we construct a chain homotopy equivalence between $\widehat{\Psi} \circ \widehat{\Phi}$ and the identity (or between $\widehat{\Phi} \circ \widehat{\Psi}$ and the identity) by counting elements of $\mathcal{M}_0^{ch}$ which are defined by constructing a family of moduli spaces $\mathcal{M}_{-1}^{\tau}$, $\tau \in [0,1]$, and combining them to

$$\mathcal{M}_0^{ch} := \bigsqcup_{\tau \in [0,1]} \mathcal{M}_{-1}^{\tau}.$$

The spaces $\mathcal{M}_{-1}^{\tau}$ are defined like done in §4.2 and §4.3. We show the chain homotopy equation by counting ends of $\mathcal{M}_1^{ch}$. Restricting our attention to $\mathcal{M}^{ch,w}$, this space consists of the union of spaces $\mathcal{M}_{-1}^{\tau,w}$, $\tau \in [0,1]$ (cf. §4.2 and §4.3). We obtain the equality

$$(\partial\mathcal{M}_0^{ch})^w = \partial\mathcal{M}_0^{ch,w}.$$

And hence we see that $\widehat{\Phi}^w$ is an isomorphism.

Handle slides. In case of the knot Floer homology we are able to define a pairing

$$\widehat{f}_{\boldsymbol{\alpha},\boldsymbol{\beta\gamma}}\colon \widehat{\mathrm{CFK}}(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},w,z) \otimes \widehat{\mathrm{CFK}}(\Sigma,\beta,\gamma,w,z) \longrightarrow \widehat{\mathrm{CFK}}(\Sigma,\alpha,\gamma,w,z)$$

induced by a doubly-pointed Heegaard triple diagram $(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},\boldsymbol{\gamma},w,z)$. We have to see, that in case the triple is induced by a handle slide, the knot Floer homology $\widehat{\mathrm{HFK}}(\Sigma,\boldsymbol{\beta},\boldsymbol{\gamma},w,z)$ carries a top-dimensional generator $\widehat{\Theta}_{\boldsymbol{\beta\gamma}}$, analogous to the discussion for the Heegaard Floer homologies, with similar properties (recall the composition law). It is easy to observe that, in case of a handle slide, the points w and z lie in the same component of $\Sigma\backslash\{\boldsymbol{\beta}\cup\boldsymbol{\gamma}\}$. Hence, we have an identification

$$\widehat{\mathrm{HFK}}(\Sigma,\boldsymbol{\beta},\boldsymbol{\gamma},w,z) = \widehat{\mathrm{HF}}(\#^g(\mathbb{S}^2 \times \mathbb{S}^1)).$$

Counting triangles with $n_w = 0$, the positivity of intersections and the additivity of the intersection number n_w guarantees that the discussion carries over verbatim and gives invariance here. $\square$

Remark 5.1. — If a handle were slid over β_1, we would leave the class of subordinate Heegaard diagrams. Recall that subordinate Heegaard diagrams come from relative handle decompositions.

5.0.4. Admissibility

The admissibility condition given in Definition 2.17 suffices to give a well-defined theory. However, since we have an additional point w in the game, we can relax the admissibility condition.

Definition 5.5. — *We call a doubly-pointed Heegaard diagram* $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, w, z)$ **extremely weakly admissible** *for the* Spin^c*-structure* $\mathfrak{s}$ *if for every non-trivial periodic domain, with* $n_w = 0$ *and* $\langle c_1(\mathfrak{s}), \mathcal{H}(\mathcal{D})\rangle = 0$*, the domain has both positive and negative coefficients.*

With a straightforward adaptation of the proof of well-definedness in the case of $\widehat{\partial}_{\mathcal{H}}$ we get the following result (see [17, Lemma 4.17], cf. Definition 2.17 and cf. proof of Theorem 2.3).

Theorem 5.6. — *Let* $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, w, z)$ *be an extremely weakly admissible Heegaard diagram. Then* $\widehat{\partial}^w_{\mathcal{H}}$ *is well-defined and a differential.* □

Note that Ozsváth and Szabó impose the weak admissibility condition of the Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, z)$. The introduction of our relaxed condition is done since there are setups (see [26]) where it is convenient to relax the admissibility condition like introduced.

5.0.5 Other knot Floer homologies

By permitting variations of n_z in the differential we define the homology $\mathrm{HFK}^{\bullet,-}$: Let $\mathrm{CFK}^{\bullet,-}(Y,K)$ be the $\mathbb{Z}[U^{-1}]$-module (or $\mathbb{Z}_2[U^{-}1]$-module) generated by the intersection points $\mathbb{T}_{\boldsymbol{\alpha}} \cap \mathbb{T}_{\boldsymbol{\beta}}$. A differential $\partial^{\bullet,-}_{\mathcal{H}}$ is defined by

$$\partial^{\bullet,-}_{\mathcal{H}}(\boldsymbol{x}) = \sum_{\boldsymbol{y} \in \mathbb{T}_{\boldsymbol{\alpha}} \cap \mathbb{T}_{\boldsymbol{\beta}}} \sum_{\phi \in H(\boldsymbol{x}, \boldsymbol{y}, 1)} \#\widehat{\mathcal{M}}_\phi \cdot \boldsymbol{y},$$

where $H(\boldsymbol{x}, \boldsymbol{y}, 1) \subset \pi_2(\boldsymbol{x}, \boldsymbol{y})$ are the homotopy classes with $n_w = 0$ (possibly $n_z \neq 0$) and $\mu = 1$. To make this a well-defined map we may impose the strong admissibility condition on the underlying Heegaard diagram or relax it like it was done for weak admissibility in Definition 5.5. Using this construction, and continuing like in §3, we define variants we denote by $\mathrm{HFK}^{\bullet,\infty}$ and $\mathrm{HFK}^{\bullet,+}$. The groups are naturally connected by exact sequences analogous to those presented in Lemma 3.3.

5.1. Refinements

If the knot K is null-homologous, we get, using a Mayer-Vietoris computation, that

$$\mathrm{Spin}^c(Y_0(K)) = \mathrm{Spin}^c(Y) \times \mathbb{Z}. \tag{5.2}$$

Alternatively, by interpretation of Spin^c-structures as homology classes of vector fields, i.e. homotopy classes over the 2-skeleton of Y, we can prove this result and see that there is a very geometric realization of the correspondence (5.2). Given a Spin^c-structure $\mathfrak{t}$ on $Y_0(K)$, we associate to it the pair $(\mathfrak{s}, k)$, where $\mathfrak{s}$ is the restriction of $\mathfrak{t}$ on Y and k an integer we will define in a moment. Beforehand, we would like to say in what way the phrase *restriction of t onto Y* makes sense. Pick a vector field v in the homology class of $\mathfrak{t}$ and restrict this vector field to $Y \backslash \nu K$. Observe that we may regard $Y \backslash \nu K$ as a submanifold of $Y_0(K)$. The restricted vector field may be interpreted as sitting on Y. We extend v to the tubular neighborhood νK of K in Y, which determines a Spin^c-structure $\mathfrak{s}$ on Y. However, the induced Spin^c-structure does not depend on the special choice of extension of v on νK, since K is homologically trivial.

To a Spin^c-structure $\mathfrak{t}$ we can associate a link $L_\mathfrak{t}$ and its homology class determines the Spin^c-structure. Denote by μ_0 a meridian of K in Y, interpreted as sitting in $Y_0(K)$. Then $L_\mathfrak{t}$ can be written as a sum

$$L_\mathfrak{t} = k \cdot \mu_0 + \dots,$$

and thus we can compute k with

$$k = lk^Y(L, \lambda) = \#^Y(L, F) = \#^{Y_0(K)}(L, \widehat{F}) = \left\langle \frac{1}{2} c_1(\mathfrak{t}), [\widehat{F}] \right\rangle,$$

where λ is a push-off of K in Y and $\widehat{F}$ is obtained by taking a Seifert surface F of K in Y and capping it off with a disk in $Y_0(K)$.

We can try to separate intersection points $\mathbb{T}_{\boldsymbol{\alpha}} \cap \mathbb{T}_{\boldsymbol{\beta}}$ with respect to Spin^c-structures of $Y_0(K)$. This defines a refined invariant $\widehat{\mathrm{CFK}}(Y, K, \mathfrak{t})$, for every $\mathfrak{t} \in \mathrm{Spin}^c(Y_0(K))$, and we have

$$\widehat{\mathrm{CFK}}(Y, K, \mathfrak{s}) = \bigoplus_{\mathfrak{t} \in H_\mathfrak{s}} \widehat{\mathrm{CFK}}(Y, K, \mathfrak{t}),$$

where $H_\mathfrak{s} \subset \mathrm{Spin}^c(Y_0(K))$ are the elements extending $\mathfrak{s} \in \mathrm{Spin}^c(Y)$. We have to show that the differential preserves this splitting. We point the interested reader to [14].

6. Maps Induced By Cobordisms

The pairing introduced in §4.4 can be used to associate maps to cobordisms. In general, every cobordism between two connected 3-manifolds Y and Y' can be decomposed into 1-handles, 2-handles and 3-handles (cf. [7, Proposition 4.2.13]). All cobordisms appearing through our work will be induced by surgeries on a 3-manifold. A surgery corresponds to a 2-handle attachment to the trivial cobordism $Y \times I$. For this reason we will not discuss 1-handles and 3-handles. We will give the construction for cobordisms obtained by attachments of one single 2-handle. For a definition of the general, very similar construction, we point the interested reader to [19].

Given a framed knot $K \subset Y$, we fix an admissible Heegaard diagram subordinate to K. Without loss of generality, we can choose the diagram such that $\beta_1 = \mu$ is a meridian of the first torus component of Σ. The framing of K is given, by pushing K off itself onto the Heegaard surface. The resulting knot on Σ is determined by $\lambda + n \cdot \mu$, for a suitable $n \in \mathbb{Z}$. With this done, we can represent the surgery by the Heegaard triple diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma})$ where γ_i, $i \geqslant 2$, are isotopic push-offs of the β_i, perturbed, such that γ_i intersects β_i in a pair of cancelling intersection points. The curve γ_1 equals $\lambda + n \cdot \mu$.

Proposition 6.1. — *The cobordism $X_{\boldsymbol{\alpha\beta\gamma}} \cup_\partial (\#^{g-1} D^3 \times \mathbb{S}^1)$ is diffeomorphic to the cobordism W_K given by the framed surgery along K.*

We define

$$\widehat{F}_{W_K} = \widehat{f}^*_{\boldsymbol{\alpha};\boldsymbol{\beta\gamma}}$$

as the map induced by the cobordism W_K. Of course, for this to make sense, we have to show that $\widehat{F}_{W_K}$ does not depend on the choices made in its definition. This is shown by the following recipe: Suppose we are given maps $\widehat{F}_1$ and $\widehat{F}_2$, induced by two sets of data that can be connected via a Heegaard move. Then these maps fit into a commutative box

$$\begin{array}{ccc} \widehat{\mathrm{HF}} & \xrightarrow{\widehat{F}_1} & \widehat{\mathrm{HF}} \\ \cong \downarrow & & \downarrow \cong \\ \widehat{\mathrm{HF}} & \xrightarrow{\widehat{F}_2} & \widehat{\mathrm{HF}} \end{array}$$

where the associated Heegaard Floer homologies are connected by the isomorphism induced by the move done to connect the diagrams. If we did a handle slide, we use associativity together with a conservation property analogous to Lemma 4.9 to show a composition law reading

$$\widehat{F}_{\boldsymbol{\alpha},\boldsymbol{\gamma\gamma'}} \circ \widehat{F}_{\boldsymbol{\alpha},\boldsymbol{\beta\gamma}} = \widehat{F}_{\boldsymbol{\alpha},\boldsymbol{\beta\gamma'}}.$$

In a similar vein one covers handle slides among the α-circles. Invariance under isotopies and changes of almost complex structures is shown by proving that the isomorphisms induced by these moves make the corresponding diagram commute.

Given a framed link $L = K_1 \sqcup \ldots \sqcup K_m$, observe that we can obviously define a map

$$\widehat{F}_L\colon \widehat{\mathrm{HF}}(Y) \longrightarrow \widehat{\mathrm{HF}}(Y_L),$$

where Y_L is the manifold obtained by surgery along L in Y, in the same way we did for a single attachment. We claim that associativity, together with a conservation law like given in Lemma 4.9, will suffice to show that the map $\widehat{F}_L$ associated to multiple attachments is a composition

$$\widehat{F}_L = \widehat{F}_{K_m} \circ \ldots \circ \widehat{F}_{K_1}$$

of the maps $\widehat{F}_{K_i}$ associated to the single attachments along the K_i. The associativity will prove that the maps in this chain *commute*. Although we have to be careful by saying *they commute*. The maps, as we change the order of the attachments, are defined differently and, thus, differ depending on the attachment order.

There is a procedure for defining maps associated to 1-handle attachments and 3-handle attachments. Their construction is not very enlightening, and the cobordisms appearing in our discussions will mostly be induced by surgeries.

7. The Surgery Exact Triangle

Denote by K a knot in Y and let n be a framing of that knot. We will briefly recall the notion of framings to fix the notation. Given a tubular neighborhood $\nu K \hookrightarrow Y$ of K, we fix a meridian μ of the boundary $\partial\nu K$. A framing is given by a push-off n of K, sitting on $\partial\nu K$, such that $\#(\mu, n) = 1$. The pair μ, λ determines a basis for $H_1(\partial\nu K; \mathbb{Z})$. Any other framing λ' can be written as $\lambda' = m \cdot \mu + \lambda$ for an integer $m \in \mathbb{Z}$ and, vice versa, any of these linear combinations determines a framing on K. Thus, when writing n as a framing for K it makes sense to talk about the framing $n + \mu$. If the knot is homologically trivial, it bounds a Seifert surface which naturally induces a framing on the knot called **the Seifert framing**. This serves as a canonical framing and having fixed this framing we can think of framings as integers $n \in \mathbb{Z}$. This identification will be done whenever it makes sense.

There is a long exact sequence

$$\ldots \xrightarrow{\partial_*} \widehat{\mathrm{HF}}(Y) \xrightarrow{\widehat{F}_1} \widehat{\mathrm{HF}}(Y_K^n) \xrightarrow{\widehat{F}_2} \widehat{\mathrm{HF}}(Y_K^{n+\mu}) \xrightarrow{\partial_*} \ldots, \tag{7.1}$$

where $\widehat{F}_i$ denote the maps associated to the cobordisms induced by the surgeries. The map $\widehat{F}_2$ is induced by a surgery along a meridian of K with framing -1. The exactness of the sequence is proved by showing that $\widehat{F}_1$ – on the chain level – can be perturbed within its chain homotopy class to fit into a short exact sequence of chain complexes and chain maps (see [18, Proposition 9.7])

$$0 \longrightarrow \widehat{\mathrm{CF}}(Y) \xrightarrow{\widetilde{\widehat{F}}_1} \widehat{\mathrm{CF}}(Y_K^n) \xrightarrow{\widehat{F}_2} \widehat{\mathrm{CF}}(Y_K^{n+\mu}) \xrightarrow{\partial_*} 0, \tag{7.2}$$

The map ∂_* in (7.1) denotes the induced coboundary. This enables us to prove the existence of the surgery exact triangle.

THEOREM 7.1. — *In the situation described above, let ν denote a meridian of μ and $\widehat{F}_3$ the map induced by surgery along ν with framing -1. There is a long exact sequence*

$$\begin{tikzcd} \widehat{\mathrm{HF}}(Y) \arrow[r, "\widehat{F}_1"] & \widehat{\mathrm{HF}}(Y_K^n) \arrow[r, "\widehat{F}_2"] & \widehat{\mathrm{HF}}(Y_K^{n+\theta}) \arrow[ll, bend left, "\widehat{F}_3"] \end{tikzcd}$$

which is called **surgery exact triangle***.*

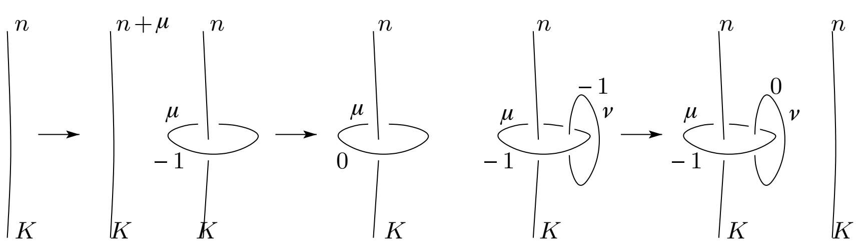

Figure 8. — The topological situation in the exact triangle.

Proof. — Observe that the topological situation is very symmetric. The long exact sequence (7.1) corresponds to the topological situation pictured in Figure 8. Each arrow in Figure 8 corresponds to an exact sequence of type (7.1). With the identifications given, we can concatenate the three sequences to give the surgery exact sequence of Theorem 7.1. □

A second proof, one more appealing to our aesthetic sense, although only valid for $\mathbb{Z}_2$-coefficients, was also developed by Ozsváth and Szabó. We will discuss the proof in the remainder of this section. It contains a very interesting algebraic approach for showing exactness of a sequence.

The composition $\widehat{f}_2 \circ \widehat{f}_1$ in the sequence

$$\widehat{\mathrm{CF}}(Y) \xrightarrow{\widehat{f}_1} \widehat{\mathrm{CF}}(Y_K^n) \xrightarrow{\widehat{f}_2} \widehat{\mathrm{CF}}(Y_K^{n+\mu}) \tag{7.3}$$

is null-chain homotopic. Let $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, z)$ be a Heegaard diagram subordinate to the knot $K \subset Y$. We can choose the data such that β_1 is a meridian of the first torus component of Σ. A Heegaard diagram of Y^n_K can be described by $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\gamma}, z)$ where γ_i, $i \geqslant 2$, are isotopic push-offs of the β_i such that β_i and γ_i meet transversely in two canceling intersections. The curve γ_1 equals $n \cdot \beta_1 + \lambda$ where λ is the longitude of the first torus component of Σ determining the framing on K. We define a fourth set of attaching circles $\boldsymbol{\delta}$ where δ_i, $i \geqslant 2$ are push-offs of the γ_i which meet the γ_i and δ_i in two cancelling intersections. The curve δ_1 equals $(n+1)\beta_1 + \lambda$. Thus, $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\delta})$ is a Heegaard diagram of $Y^{n+\mu}_K$. By associativity (4.8), the composition $\widehat{f}_2 \circ \widehat{f}_1$ is chain homotopic to

$$\widehat{f}_{\boldsymbol{\alpha};\boldsymbol{\beta}\boldsymbol{\delta}}(\,\cdot \otimes \widehat{f}_{\boldsymbol{\beta};\boldsymbol{\gamma}\boldsymbol{\delta}}(\widehat{\Theta}_{\boldsymbol{\beta}\boldsymbol{\gamma}} \otimes \widehat{\Theta}_{\boldsymbol{\gamma}\boldsymbol{\delta}})),$$

where the chain homotopy H is given by counting holomorphic rectangles with suitable boundary conditions (cf. §4.4). To compute $\widehat{f}_{\boldsymbol{\beta};\boldsymbol{\gamma}\boldsymbol{\delta}}(\widehat{\Theta}_{\boldsymbol{\beta}\boldsymbol{\gamma}} \otimes \widehat{\Theta}_{\boldsymbol{\gamma}\boldsymbol{\delta}})$ we use a model calculation. Figure 9 illustrates the Heegaard triple diagram.

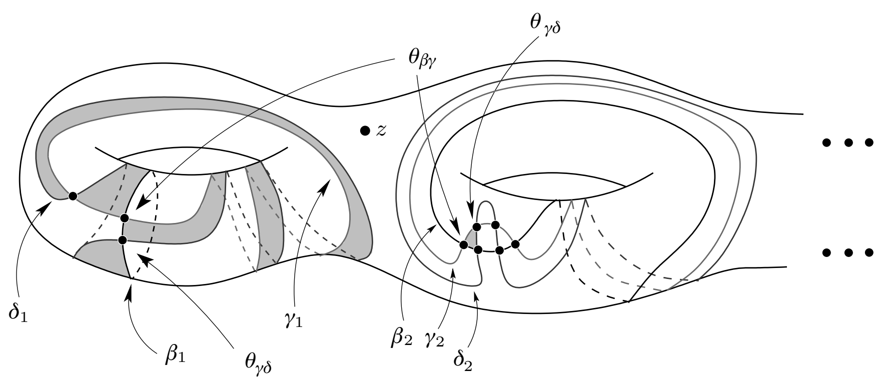

Figure 9. — Heegaard triple diagram for the computation of $\widehat{f}_{\boldsymbol{\beta},\boldsymbol{\gamma}\boldsymbol{\delta}}(\widehat{\Theta}_{\boldsymbol{\beta}\boldsymbol{\gamma}} \otimes \widehat{\Theta}_{\boldsymbol{\gamma}\boldsymbol{\delta}})$.

There are exactly two homotopy classes of Whitney triangles we have to count. Each domain associated to the homotopy classes is given by a disjoint union of triangles. Thus, the moduli spaces associated to these homotopy classes each carry one single element (cf. Lemma 4.9). Hence, in $\mathbb{Z}_2$-coefficients

$$\widehat{f}_{\boldsymbol{\beta};\boldsymbol{\gamma}\boldsymbol{\delta}}(\widehat{\Theta}_{\boldsymbol{\beta}\boldsymbol{\gamma}} \otimes \widehat{\Theta}_{\boldsymbol{\gamma}\boldsymbol{\delta}}) = 2 \cdot \widehat{\Theta}_{\boldsymbol{\beta}\boldsymbol{\delta}} = 0.$$

In general we have to see that we can choose the signs of the associated elements differently. But observe that the domains of both homotopy classes contributing in our signed count differ by a triply-periodic domain. We can choose the signs on these elements differently.

This discussion carries over verbatim for any of the maps in the surgery exact sequence. The symmetry of the situation, as indicated in Figure 8, makes it possible to carry over the proof given here.

There is an algebraic trick to show exactness on the homological level. Let

$$H\colon \widehat{\mathrm{CF}}(Y) \longrightarrow \widehat{\mathrm{CF}}(Y_K^{n+\mu})$$

denote the null-homotopy of $\widehat{f}_2 \circ \widehat{f}_1$ (cf. §4.4). Define the chain complex $A_{\widehat{f}_1,\widehat{f}_2}$ to be given by the module $A = \widehat{\mathrm{CF}}(Y) \oplus \widehat{\mathrm{CF}}(Y_K^n) \oplus \widehat{\mathrm{CF}}(Y_K^{n+\mu})$ with the differential

$$\partial = \begin{pmatrix} \widehat{\partial}_Y & 0 & 0 \\ \widehat{f}_1 & \widehat{\partial}_{Y_K^n} & 0 \\ H & \widehat{f}_2 & \widehat{\partial}_{Y_K^{n+\mu}} \end{pmatrix}.$$

Lemma 7.2. — *The sequence (7.3) is exact on the homological level at $\widehat{\mathrm{CF}}(Y_K^n)$ if the homology $H_*(A_{\widehat{f}_1,\widehat{f}_2})$ is trivial.*

Proof. — Suppose we are given an element $b \in \widehat{\mathrm{CF}}(Y_K^n) \cap \ker(\widehat{f}_2)$ with $\widehat{\partial}_{Y_K^n} b = 0$. Since $H_*(A_{\widehat{f}_1,\widehat{f}_2}, \partial)$ is trivial there is an element $(x, y, q) \in A$ such that $(0, b, 0) = \partial(x, y, q)$. Thus, we have

$$b = \widehat{f}_1(x) + \widehat{\partial}_{Y_K^n}(y)$$

proving, that $[b] \in \mathrm{im}(\widehat{F}_1)$. □

Definition 7.3. — *For a chain map $f\colon A \longrightarrow B$ between $\mathbb{Z}_2$-vector spaces we define its* **mapping cone** *to be the chain complex $M(f)$, given by the module $A \oplus B$ with differential*

$$\partial_f = (\partial_A \quad 0 f \quad \partial_B .)$$

The mapping cone is a chain complex.

From the definition of mapping cones there is a short exact sequence of chain complexes

$$0 \longrightarrow \widehat{\mathrm{CF}}(Y_K^{n+\mu}) \xrightarrow{\widehat{f}_1} A_{\widehat{f}_1,\widehat{f}_2} \xrightarrow{\widehat{f}_2} M(\widehat{f}_1) \longrightarrow 0$$

inducing a long exact sequence between the associated homologies. The connecting morphism of this long exact sequence is induced by

$$(H, \widehat{f}_2)\colon M(\widehat{f}_1) \longrightarrow \widehat{\mathrm{CF}}(Y_K^{n+\mu}).$$

The triviality of $H_*(A_{\widehat{f}_1,\widehat{f}_2},\partial)$ is the same as saying that $(H,\widehat{f}_2)_*$ is an isomorphism.

LEMMA 7.4 (LEMMA 4.2 OF [20],) *Let $\{A_i\}_{i\in\mathbb{Z}}$ be a collection of modules and let*

$$\{f_i\colon A_i \longrightarrow A_{i+1}\}_{i\in\mathbb{Z}}$$

be a collection of chain maps such that $f_{i+1}\circ f_i$, $i\in\mathbb{Z}$ is chain homotopically trivial by a chain homotopy $H_i\colon A_i \longrightarrow A_{i+2}$. The maps

$$\psi_i = f_{i+2}\circ H_i + H_{i+1}\circ f_i\colon A_i \longrightarrow A_{i+3}$$

should induce isomorphisms between the associated homologies. Then the maps

$$(H_i, f_{i+1})\colon M(f_i) \longrightarrow A_{i+2}$$

induce isomorphisms on the homological level.

If we can show that the sequence

$$\begin{tikzcd} \widehat{\mathrm{CF}}(Y) \arrow[r, "\widehat{f}_1"] & \widehat{\mathrm{CF}}(Y^n_K) \arrow[r, "\widehat{f}_2"] & \widehat{\mathrm{CF}}(Y^{n+\theta}_K) \arrow[ll, bend left, "\widehat{f}_3"] \end{tikzcd}$$

satisfies the assumptions of Lemma 7.4, then for every pair $\widehat{f}_i$ and $\widehat{f}_{i+1}$, the associated map $(H,\widehat{f}_{i+1})_*$ is an isomorphism. With the arguments from above, i.e. analogous to Lemma 7.2, we conclude that $\mathrm{im}(\widehat{F}_i) = \ker(\widehat{F}_{i+1})$. Hence, Theorem 7.1 follows.

8. The Contact Element and Legendrian knot Invariant

8.1. Contact Structures

A 3-dimensional contact manifold is a pair (Y,ξ) where Y is a 3-dimensional manifold and $\xi\subset TY$ a hyperplane bundle that can be written as the kernel of a 1-form α with the property

$$\alpha\wedge d\alpha \neq 0. \tag{8.1}$$

Differential 1-forms satisfying (8.1) are called **contact forms**. Given a contact manifold (Y,ξ), the associated contact form is not unique. Suppose α is a contact form of ξ then, given a non-vanishing function $\lambda\colon Y \longrightarrow \mathbb{R}^+$, we can change the contact form to $\lambda\alpha$ without affecting the contact condition (8.1):

$$\lambda\alpha\wedge d(\lambda\alpha) = \lambda\alpha\wedge d\lambda\wedge\alpha + \lambda^2\alpha\wedge d\alpha = \lambda^2\alpha\wedge d\alpha \neq 0.$$

The existence of a contact form implies that the normal direction TY/ξ is trivial. We define a section R_α by

$$\alpha(R_\alpha) \neq 0 \text{ and } \quad \iota_{R_\alpha} d\alpha = 0.$$

This vector field is called **Reeb field** of the contact form α. The contact condition implies that $d\alpha$ is a non-degenerate form on ξ. Thus, $\iota_{R_\alpha} d\alpha = 0$ implies that for each point $p \in Y$ the vector $(R_\alpha)_p$ is an element of $T_pY \backslash \xi_p$. Thus, R_α is a section of TY/ξ.

Definition 8.1. — *Two contact manifolds* (Y, ξ) *and* (Y', ξ') *are called* **contactomorphic** *if there is a diffeomorphism* $\phi\colon Y \longrightarrow Y'$ *preserving the contact structures, i.e. such that* $T\phi(\xi) = \xi'$*. The map* ϕ *is a* **contactomorphism**.

It is a remarkable property of contact manifolds that there is a unique local model for these objects.

Definition 8.2. — *The pair* $(\mathbb{R}^3, \xi_{std})$*, where* ξ_{std} *is the contact structure given by the kernel of the* 1*-form* $dz - y\,dx$*, is called* **standard contact space**.

Every contact manifold is locally contactomorphic to the standard contact space. This is known as **Darboux's theorem**. As a consequence we will not be able to derive contact invariants by purely local arguments, in contrast to differential geometry where for instance curvature is a constraint to the existing local model.

Theorem 8.3 (Gray Stability, cf. Theorem 2.2.2 of [6]). — *Each smooth homotopy* $(\xi_t)_{t \in [0,1]}$ *of contact structures is induced by an ambient isotopy* ϕ_t*, i.e. the condition* $T\phi_t(\xi_0) = \xi_t$ *applies for all* $t \in [0, 1]$.

An isotopy which is induced by a homotopy of contact structures is called **contact isotopy**. So, a homotopy of contact structures can be interpreted as an isotopy and, vice versa, an isotopy induces a homotopy of contact structures. As in the case of vector fields, we have a natural connection to isotopies, i.e. objects whose existence and form will be closely related to the manifold's topology.

A **contact vector field** X is a vector field whose local flow preserves the contact structure. An embedded surface $\Sigma \hookrightarrow Y$ is called **convex** if there is a neighborhood of Σ in Y in which a contact vector field exists that is transverse to Σ. The existence of a contact vector field immediately implies that there is a neighborhood $\Sigma \times \mathbb{R} \hookrightarrow Y$ of Σ in which the contact structure is invariant in $\mathbb{R}$-direction. Thus, convex surfaces are the objects along which we glue contact manifolds together.

Definition 8.4. — *A knot $K \subset Y$ is called* **Legendrian** *if it is tangent to the contact structure.*

The contact condition implies that, on a 3-dimensional contact manifold (Y, ξ), only 1-dimensional submanifolds, i.e. knots and links, can be tangent to ξ. Every Legendrian knot admits a tubular neighborhood with a convex surface as boundary. Hence, it is possible to mimic surgical constructions to define the contact geometric analogue of surgery theory, called **contact surgery**. Contact surgery in arbitrary dimensions was introduced by Eliashberg in [4]. His construction, in dimension 3, corresponds to (-1)-contact surgeries. For 3-dimensional contact manifolds Ding and Geiges gave in [2] a definition of contact-r-surgeries (cf. also [3]) for arbitrary $r \in \mathbb{Q} > 0$. It is nowadays one of the most significant tools for 3-dimensional contact geometry. Its importance relies in the following theorem.

Theorem 8.5 (see [3]). — *Given a contact manifold (Y, ξ), there is a link $\mathbb{L} = \mathbb{L}^+ \sqcup \mathbb{L}^-$ in $\mathbb{S}^3$ such that contact-$(+1)$-surgery along the link $\mathbb{L}^+$ and contact-(-1)-surgery along $\mathbb{L}^-$ in $(\mathbb{S}^3, \xi_{std})$ yields (Y, ξ).*

Moreover, if we choose cleverly, we can accomplish $\mathbb{L}^+$ to have just one component. Using (-1)-contact surgeries only, we can transform an arbitrary overtwisted contact manifold into an arbitrary (not necessarily overtwisted) contact manifold. For a definition of overtwistedness we point the reader to [6, Definition 4.5.1]. Thus, starting with a knot K so that $(+1)$-contact surgery along K yields an overtwisted contact manifold (Y', ξ'), for any contact manifold (Y, ξ), we can find a link $\mathbb{L}^-$, such that (-1)-contact surgery along $\mathbb{L}^-$ in (Y', ξ') yields (Y, ξ). An example for such a knot K is the Legendrian shark, i.e. the Legendrian realization of the unknot with $tb = -1$ and $rot = 0$.

8.2. Open Books

For a detailed treatment of open books we point the reader to [5].

Definition 8.6. — *An* **open book** *on a closed, oriented 3-manifold Y is a pair (B, π) defining a fibration*

$$P \hookrightarrow Y \backslash B \pi \longrightarrow \mathbb{S}^1,$$

where P is an oriented surface with boundary $\partial P = B$. For every component B_i of B there is a neighborhood $\iota\colon D^2 \times \mathbb{S}^1 \hookrightarrow \nu B_i \subset Y$ such that the core $C = \{0\} \times \mathbb{S}^1$ is mapped onto B_i under ι and π commutes with the projection $(D^2 \times \mathbb{S}^1) \backslash C \longrightarrow \mathbb{S}^1$ given by $(r \cdot \exp(it), \exp(is)) \longmapsto \exp(it)$. The submanifold B is called **binding** *and P the* **page of the open book**.

An **abstract open book** is a pair (P,ϕ) consisting of an oriented genus-g surface P with boundary and a homeomorphism $\phi\colon P \longrightarrow P$ that is the identity near the boundary of P. The surface P is called **page** and ϕ the **monodromy**. Given an abstract open book (P,ϕ), we may associate to it a 3-manifold. Let $c_1,\dots,c_k$ denote the boundary components of P. Observe that

$$(P\times[0,1])/(p,1)\sim(\phi(p),0) \tag{8.2}$$

is a 3-manifold. Its boundary is given by the tori

$$((c_i\times[0,1])/(p,1)\sim(p,0))\cong c_i\times\mathbb{S}^1.$$

Fill in each of the holes with a solid torus $D^2\times\mathbb{S}^1$: we glue a meridional disk $D^2\times\{\star\}$ onto $\{\star\}\times\mathbb{S}^1\subset c_i\times\mathbb{S}^1$. In this way we define a closed oriented 3-manifold $Y(P,\phi)$. Denote by B the union of the cores of the tori $D^2\times\mathbb{S}^1$. The set B is called **binding**. By definition of abstract open books we obtain an open book structure

$$P\hookrightarrow Y(P,\phi)\backslash B\longrightarrow\mathbb{S}^1$$

on $Y(P,\phi)$. Conversely, given an open book by cutting a small tubular neighborhood νB out of Y, we obtain a P-bundle over $\mathbb{S}^1$. Thus, there is a homeomorphism $\phi\colon P\longrightarrow P$ such that

$$Y\backslash\nu B\cong(P\times[0,1])/(p,1)\sim(\phi(p),0).$$

Inside the standard neighborhood νB, as given in the definition, the homeomorphism ϕ is the identity. So, the pair (P,ϕ) defines an abstract open book.

Definition 8.7. — *Two abstract open books (P,ϕ) and (P,ϕ') are called* **equivalent** *if there is a homeomorphism $h\colon P\longrightarrow P$ which is the identity near the boundary such that $\phi\circ h=\phi'\circ h$. We denote by $\mathrm{ABS}(Y)$ the set of abstract open books (P,ϕ) with $Y(P,\phi)=Y$, up to equivalence.*

The set of equivalence classes of open books is denoted by $\mathrm{OB}(Y)$. An abstract open book defines an open book up to diffeomorphism. The discussion from above provides us with a map

$$\Psi\colon \mathrm{ABS}(Y)\longrightarrow\mathrm{OB}(Y)$$

and an inverse. Thus, to some point, open books and abstract open books are the same objects. Sometimes, it is more convenient to deal with abstract open books rather than open books themselves.

8.3. Open Books, Contact Structures and Heegaard Diagrams

Given an open book (B,π) or an abstract open book (P,ϕ), define a surface Σ by gluing together two pages at their boundary

$$\Sigma = P_{1/2} \cup_\partial P_1.$$

The manifold Y equals the union $H_0 \cup H_1$ where $H_i = \pi^{-1}([i/2,(i+1)/2])$, $i=0,1$. Any curve γ in Y running from H_0 to H_1, when projected onto $\mathbb{S}^1$, has to intersect $\{1/2,1\}$ at some point. Thus, the curve γ has to intersect Σ. The submanifolds H_i are handlebodies of genus $g(\Sigma)$ and

$$Y = H_0 \cup_\partial H_1$$

is a Heegaard decomposition of Y.

Definition 8.8. — *A system $a=\{a_1,\dots,a_n\}$ of disjoint properly embedded arcs on P is called* **cut system** *if $P\backslash\{a_1,\dots,a_n\}$ is topologically a disk.*

A system of arcs is a cut system if and only if it defines a basis for the first homology of $(P,\partial P)$.

We interpret the curve a_i as sitting on $P_{1/2}$ and $\overline{a_i}$, i.e. the curve a_i with reversed orientation, as sitting inside P_1. These two can be combined to $\alpha_i = a_i \cup_\partial \overline{a_i}$, $i=1,\dots,n$, which all sit in Σ. Referring to the relation between open books and abstract open books discussed in §8.2, observe that

$$H_1 = \pi^{-1}([1/2,1]) = (P\times[1/2,1])/{\sim}$$

where $\sim$ identifies points $(p,0)$ with $(\phi(p),1)$ for $p\in P$ and points (p,t) with (p,t') for $p\in\partial P$ and $t,t'\in[1/2,1]$. Thus $a_i\times[1/2,1]$ determines a disk in H_1 whose boundary is α_i. This means we can interpret the set $\{\alpha_1,\dots,\alpha_n\}$ as a set of attaching circles for the handlebody H_1. The gluing of the two handlebodies H_0 and H_1 is given by the pair (id,ϕ) where id is the identity on $P_{1/2}$ and ϕ the monodromy, interpreted as a map $P_1 \longrightarrow P_0$. These two maps combine to a map $\partial H_1 \longrightarrow \partial H_0$. Define b_i, $i=1,\dots,n$, as small push-offs of the a_i that intersect these transversely in a single point (see Figure 10). Then by the gluing of the two handlebodies H_0 and H_1 the α-curves define a Heegaard diagram with β-curves given by $\beta_i = b_i \cup \overline{\phi(b_i)}$, $i=1,\dots,n$. Thus the following lemma is immediate.

Lemma 8.9. — *The triple $(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta})$ is a Heegaard diagram of Y.* □

Given an abstract open book (P,ϕ), define P' by attaching a 1-handle to P, i.e. $P' = P\cup h^1$. Choose a knot γ in P' that intersects the co-core of h^1

once, transversely. The monodromy ϕ can be extended as the identity over h^1, and, thus, may be interpreted as a homeomorphism of P'. We denote by $D^{\pm}_{\gamma}$ the positive/negative Dehn twist along γ.

Definition 8.10. — *The abstract open book* $(P', D^{\pm}_{\gamma} \circ \phi)$ *is called a* **positive/negative Giroux stabilization** *of* (P, ϕ).

We will see that open books, up to positive Giroux stabilizations, correspond one-to-one to isotopy classes of contact structures.

Lemma 8.11. — *Stabilizations preserve the underlying* 3*-manifold, i.e. the manifolds* $Y(P', \phi')$ *and* $Y(P, \phi)$ *are diffeomorphic.*

A priori, it is not clear that stabilizations preserve the associated 3-manifold. A proof of this lemma can be found in [5]. But in the following we will discuss an alternative proof. Our proof uses a construction introduced by Lisca, Ozsváth, Stipsicz and Szabó (see the alternative proof of Theorem 2.11 in [10, p. 1320]).

Lemma 8.12 (see p. 1321 of [10]). — *There is a cut system* $\{a_1, \ldots, a_n\}$ *on* (P, ϕ) *that is disjoint from* $\gamma \cap P$.

Proof. — Denote by γ' the arc $\gamma \cap P$. If $P \backslash \gamma'$ is connected, we choose a_1 to be a push-off of γ' and then extend it to a cut system of P. This is possible since $H_1(P, \partial P)$ is torsion free and $[a_1]$ a primitive element in it. If $P \backslash \gamma'$ diskonnects into the components P_1 and P_2, then we may choose cut systems on P_i, $i = 1, 2$, arbitrarily. The union of these cut systems will be a cut system of P and disjoint from γ'. □

The given cut system on P can be extended to a cut system on P'. We can choose a_{n+1} as the co-core of h^1. The set of curves $a_1, \ldots, a_{n+1}$ is a cut system of P'. Choose the b_i, $i = 1, \ldots, n+1$, as small isotopic push-offs of the a_i. Then, for $i = 1, \ldots, n$, we have

$$\begin{array}{rclcl} \phi'(b_i) & = & \phi \circ D^{\pm}_{\gamma}(b_i) & = & \phi(b_i) \\ \phi'(b_{n+1}) & = & D^{\pm}_{\gamma} \circ \phi(b_{n+1}) & = & D^{\pm}_{\gamma}(b_{n+1}). \end{array}$$

Consequently, $\phi'(b_{n+1})$ looks like γ outside the handle h^1. The curve β_{n+1} has to be disjoint from all α_i, $i < n+1$.

Proof of Lemma 8.11. — On the level of cobordisms the pair α_{n+1} and β_{n+1} which meet in a single point correspond to a cancelling pair of handles attached to the boundary $Y(P, \phi) \times \{1\}$ of $Y(P, \phi) \times I$. Thus, we have

$$Y(P', \phi') = \mathbb{S}^3 \# Y(P, \phi).$$

□

A contact structure ξ is **supported** by an open book (B, π) of Y if ξ is contact isotopic to a contact structure ξ' which admits a contact form α such that $d\alpha$ is a positive area form on each page $P_\theta = \pi^{-1}(\theta)$ and $\alpha > 0$ on ∂P_θ. We gave the definition as a matter of completeness, but a detailed understanding of this definition will not be interesting to us. For a detailed treatment we point the reader to [5]. Every contact structure is supported by an open book decomposition.

Theorem 8.13 (cf. [5]). — *There is a one-to-one correspondence between isotopy classes of contact structures and open book decompositions up to positive Giroux stabilization.*

Given a Legendrian knot $L \subset (Y, \xi)$, we know by definition that its tangent vector at every point of L lies in ξ. The tangent bundle of a closed oriented 3-manifold is orientable, which especially implies the triviality of $TY|_L$. The coorientability of ξ implies that $\xi|_L$ is trivial, too. By definition of Legendrian knots the tangent vector of L lies in ξ. The 2-dimensionality implies that ξ, in addition, contains a normal direction. The triviality of the tangent bundle over L implies that this normal direction determines a framing of L. This framing which is determined by the contact structure is called **contact framing**. In case of contact surgery it plays the role of the canonical 0-framing, i.e. we measure contact surgery coefficients with respect to the contact framing. Note that if L is homologically trivial, a Seifert surface determines a second framing on L. Surgery coefficients in a surgery presentation of a manifold are usually determined by measuring the surgery framing with respect to this canonical Seifert framing (cf. §7). Measuring the contact framing with respect to the Seifert framing determines a number $tb(L) \in \mathbb{Z}$ which is called the **Thurston-Bennequin invariant**. This is certainly an invariant of L under **Legendrian isotopies**, i.e. isotopies of L through Legendrian knots. By definition, the coefficients are related by

$$\text{smooth surgery coefficient} = \text{contact surgery coefficient} + tb(L).$$

It is possible to find an open book decomposition which supports ξ such that L sits on a page of the open book. Furthermore, we can arrange the page framing and the contact framing to coincide. This is the most important ingredient for applications of Heegaard Floer homology in the contact geometric world. The proof relies on the fact that it is possible to find CW-decompositions of contact manifolds which are adapted to the contact structure. These are called **contact cell decompositions**. The 1-cells in such a decomposition are Legendrian arcs. With these decompositions it is possible to directly construct an open book supporting the contact structure. Since the 1-cells are Legendrian arcs we can include a fixed Legendrian knot into the decomposition and in this way modify the open book such that the result follows. For details we point the reader to [5].

LEMMA 8.14 (cf. Proposition 2.4 of [10]). — *Let $L \subset (Y, \xi)$ be a Legendrian knot and (P, ϕ) an abstract open book supporting ξ such that L sits on a page of the underlying open book. Let $(Y_L^{\pm}, \xi_L^{\pm})$ denote the 3-manifold obtained by (± 1)-contact surgery along L. Then $(P, D_\gamma^{\mp} \circ \phi)$ is an abstract open book supporting the contact structure $\xi_L^{\pm}$.*

8.4. The Contact Element

Given a contact manifold (Y, ξ), we fix an open book decomposition (P, ϕ) which supports ξ. This open book defines a Heegaard decomposition and, with the construction stated in the last section, we are able to define a Heegaard diagram. We now put in an additional datum. The curves b_i are isotopic push-offs of the a_i. We choose them like indicated in Figure 10: We push the b_i off the a_i by following with ∂b_i the positive boundary orientation of ∂P.

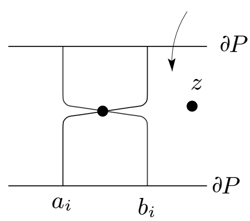

Figure 10. — Positioning of the point z and choice of b_i.

The point z is placed outside the thin strips of isotopy between the a_i and b_i. We denote by $\boldsymbol{x}_i$ the unique intersection point between a_i and b_i. Define

$$EH(P, \phi, \{a_1, \ldots, a_n\}) = \{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n\}.$$

By construction of the Heegaard diagram EH is a cycle in the Heegaard Floer homology associated to the data $(-\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, z)$. We choose the negative surface orientation since with this orientation there can be no holomorphic Whitney disks emanating from EH (cf. Figure 10).

LEMMA 8.15 (SEE PROPOSITION 2.5 OF [18]) *The Heegaard Floer cohomology $\widehat{\mathrm{HF}}^*(Y)$ is isomorphic to $\widehat{\mathrm{HF}}(-Y)$.*

The Heegaard diagram $(-\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$ is a Heegaard diagram for $-Y$ and, thus, represents the Heegaard Floer cohomology of Y. Instead of switching the surface orientation we can swap the boundary conditions of the Whitney disks at their α-boundary and β-coundary, i.e. we will be interested in Whitney disks in $(\Sigma, \boldsymbol{\beta}, \boldsymbol{\alpha})$. The element EH can be interpreted as sitting in the Heegaard Floer cohomology of Y. The push-off b_i is chosen such that there is no holomorphic disk emanating from $\boldsymbol{x}_i$.

THEOREM 8.16 (see Theorem 3.1 of [8]). — *The class $EH(P, \phi, \{a_1, \dots, a_n\})$ is independent of the choices made in its definition. Moreover, the associated cohomology class $c(Y, \xi)$ is an isotopy invariant of the contact structure ξ, up to sign. We call $c(Y, \xi)$* **contact element**.

The proof of this theorem relies on several steps we would like to sketch: An **arc slide** is a geometric move allowing us to change the cut system. Any two cut systems can be transformed into each other by a finite sequence of arc slides. Let a_1 and a_2 be two adjacent arcs. Adjacent means that in $P \backslash \{a_1, \dots, a_n\}$ one of the boundary segments associated to a_1 and a_2 are connected via one segment τ of ∂P. An arc slide of a_1 over a_2 (or vice versa) is a curve in the isotopy class of $a_1 \cup \tau \cup a_2$. We denote it by $a_1 + a_2$.

LEMMA 8.17. — *Any two cut systems can be transformed into each other with a finite number of arc slides.*

It is easy to observe that an arc slide affects the associated Heegaard diagram by two handle slides. The change under the α-circles is given by a handle slide of α_1 over α_2. But the associated β-curve moves with the α-curve, i.e. we have to additionally slide β_1 over β_2. We have to see that these handle slides preserve the contact element. To be more precise: After the first handle slide we moved out of the set of Heegaard diagrams induced by open books. Thus, we cannot see the contact element in that diagram. After the second handle slide, however, we move back into that set and, hence, see the contact element again. We have to check that the composition of the maps between the Heegaard Floer cohomologies induced by the handle slides preserves the contact element. This is a straightforward computation.

DEFINITION 8.18. — *Let a Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$ and a homologically essential, simple, closed curve δ on Σ be given. The Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$ is called* **δ-adapted** *if the following conditions hold.*

1. *It is induced by an open book and the pair $\boldsymbol{\alpha}$, $\boldsymbol{\beta}$ is induced by a cut system (cf. §8.3) for this open book.*

2. *The curve δ intersects β_1 once and does not intersect any other of the β_i, $i \geqslant 2$.*

We can always find δ-adapted Heegaard diagrams. This is already stated in [8] and [10] but not proved. For the convenience of the reader we include a proof, here.

LEMMA 8.19. — *Let (P, ϕ) be an open book and $\delta \subset P$ a homologically essential closed curve. There is a choice of cut system on P that induces a δ-adapted Heegaard diagram.*

Observe that $a_1, \ldots, a_n$ to be a cut system of a page P essentially means to be a basis of $H_1(P, \partial P)$: Suppose the curves are not linearly independent. In this case we are able to identify a surface $F \subset P$, $F \neq P$, bounding a linear combination of some of the curves a_i. But this means the cut system diskonnects the page P in contradiction to the definition. Conversely, suppose the curves in the cut system are homologically linearly independent. In this case the curves cannot diskonnect the page. If they diskonnected, we could identify a surface F in P with boundary a linear combination of some of the a_i. But this contradicts their linear independence. The fact that $\Sigma \backslash \{a_1, \ldots, a_n\}$ is a disk shows that every element in $H_1(P, \partial P)$ can be written as a linear combination of the curves $a_1, \ldots, a_n$.

Proof. — Without loss of generality, we assume that P has connected boundary: Suppose the boundary of P has two components. Choose a properly embedded arc connecting both components of ∂P. Define this curve to be the first curve a_0 in a cut system. Cutting out this curve a_0, we obtain a surface with connected boundary. The curve a_0 determines two segments S_1 and S_2 in the connected boundary. We can continue using the construction process for connected binding we state below. We just have to check the boundary points of the curves to remain outside of the segments S_1 and S_2. Given that P has more than two boundary components, we can, with this algorithm, inductively decrease the number of boundary components.

The map ϕ is an element of the mapping class group of P. Thus, if $\{a_1, \ldots, a_n\}$ is a cut system, then $\{\phi(a_1), \ldots, \phi(a_n)\}$ is a cut system, too. It suffices to show that there is a cut system $\{a_1, \ldots, a_n\}$ such that δ intersects a_i once if and only if $i = 1$.

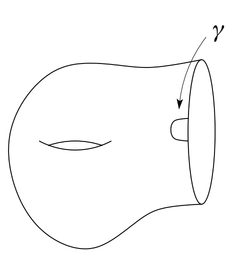

Figure 11. — Possible choice of curve γ.

We start by taking a band sum of δ with a small arc γ as shown in Figure 11. We are free to choose the arc γ. Denote the result of the band sum by a_2. The arc a_2 indeed bounds a compressing disk in the respective handlebody because its boundary lies on ∂P. Because of our prior observation it suffices to show that a_2 is a primitive class in $H_1(P, \partial P)$. Since $H_1(P, \partial P)$ is torsion free the primitiveness of a_2 implies that we can extend a_2 to a basis of $H_1(P, \partial P)$. The curves defining this basis can easily be chosen to be not closed, with their boundary lying on ∂P.

Writing down the long exact sequence of the pair $(P, \partial P)$

$$\begin{tikzcd}
H_2(P) \arrow[r] \arrow[d, phantom, "\|"] & H_2(P,\partial P) \arrow[r, "\partial_*"] \arrow[d, phantom, "\cong"] & H_1(\partial P) \arrow[r] \arrow[d, phantom, "\cong"] & H_1(P) \arrow[r, "\iota_*"] & H_1(P,\partial P) \arrow[r] & 0 \\
0 \arrow[r] & \mathbb{Z}\langle [P]\rangle \arrow[r, "\partial_*"] & \mathbb{Z}\langle [\partial P]\rangle \arrow[r] & H_1(P) \arrow[r, "\iota_*"] & H_1(P,\partial P) \arrow[r] & 0
\end{tikzcd}$$

we see that ∂_* is surjective since $\partial_*[P] = [\partial P]$. Hence, exactness of the sequence implies that the inclusion $\iota\colon P \longrightarrow (P, \partial P)$ induces an isomorphism on homology. Note that the zero at the end of the sequence appears because ∂P is assumed to be connected. Let g denote the genus of P. Of course $H_1(P;\mathbb{Z})$ is $\mathbb{Z}^{2g}$, which can be seen by a Mayer-Vietoris argument or from handle decompositions of surfaces (compute the homology using a handle decomposition). Since δ was embedded it follows from the lemma below that it is a primitive class in $H_1(P;\mathbb{Z})$. The isomorphism ι_* obviously sends δ to a_2, i.e. $\iota_*[\delta] = [\gamma]$. Thus, a_2 is primitive in $H_1(P, \partial P)$.

Cut open the surface along δ. We obtain two new boundary components, C_1 and C_2 say, which we can connect with the boundary of P with two arcs. These two arcs, in P, determine a properly embedded curve, a_1 say, whose boundary lies on ∂P. Furthermore, a_1 intersects δ in one single point, transversely. The curve a_1 is primitve, too. To see, that we can extend to a cut system such that δ is disjoint from $a_3, \ldots, a_n$, cut open the surface P along δ and a_1. We obtain a surface P' with one boundary component. The curves δ and a_1 determine 4 segments, $S_1, \ldots, S_4$ say, in this boundary. We extend a_2 to a cut system $a_2, \ldots, a_n$ of P' and arrange the boundary points of the curves $a_3, \ldots, a_n$ to be disjoint from $S_1, \ldots, S_4$. The set $a_1, \ldots, a_n$ is a cut system of P with the desired properties. $\square$

As a consequence of the proof we may arrange δ to be a push-off of a_2 outside a small neighborhood where the band sum is performed. Geometrically spoken, we cut open δ at one point and move the boundaries to ∂P to get a_2. Given a positive Giroux stabilization, we can find a special cut system which is adapted to the curve γ. It is not hard to see that there is only one homotopy class of triangles that connect the old with the new contact element and that the associated moduli space is a one-point space.

Lemma 8.20. — *An embedded circle δ in an orientable, compact surface Σ which is homologically essential is a primitive class of $H_1(\Sigma,\mathbb{Z})$.*

Proof. — Cut open the surface Σ along δ. We obtain a connected surface S with two boundary components since δ is homologically essential in Σ. We can recover the surface Σ by connecting both boundary components of S with a 1-handle and then capping off with a disk. There is a knot $K \subset S \cup h^1$ intersecting the co-core of h^1 only once and intersecting δ only once, too. To construct this knot take a union of two arcs in $S \cup h^1$ in the following way: Namely, define a as the core of h^1, i.e. as $D^1 \times \{0\} \subset D^1 \times D^1 \cong h^1$ and let b be a curve in S, connecting the two components of the attaching sphere h^1 in ∂S. We define K to be $a \cup b$. Obviously,

$$\pm 1 = \#(K,\delta) = \big\langle PD[K],[\delta]\big\rangle.$$

Since $H_1(\Sigma;\mathbb{Z})$ is torsion free, $H^1(\Sigma;\mathbb{Z}) \cong \mathrm{Hom}(H_1(\Sigma;\mathbb{Z}),\mathbb{Z})$. Thus, $[\delta]$ is primitive. $\square$

Recall that a positive/negative Giroux stabilization of an open book (P,ϕ) is defined as the open book $(P', D^{\pm}_{\gamma} \circ \phi)$ where P' is defined by attaching a 1-handle to P and γ is a embedded, simple closed curve in P' that intersects the co-core of h^1 once (see Definition 8.10). Using the proofs of Lemma 8.11 and Lemma 8.12, we see that there is a cut system $\{a_1,\dots,a_{n+1}\}$ of the stabilized open book such that γ intersects only a_{n+1} which is the co-core of h^1. Denote by $\boldsymbol{\alpha} = \{\alpha_1,\dots,\alpha_n\}$ the associated attaching circles. We define a map

$$\Phi\colon \widehat{\mathrm{CF}}(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},z) \longrightarrow \widehat{\mathrm{CF}}(\Sigma\#T^2,\boldsymbol{\alpha}\cup\{\alpha_{n+1}\},\boldsymbol{\beta}\cup\{\beta_{n+1}\},z)$$

by assigning to $\boldsymbol{x} \in \mathbb{T}_{\boldsymbol{\alpha}}\cap\mathbb{T}_{\boldsymbol{\beta}}$ the element $\Phi(\boldsymbol{x}) = (\boldsymbol{x},\boldsymbol{q})$ where $\boldsymbol{q}$ is the unique intersection point $\gamma \cap a_{n+1}$. This is an isomorphism by reasons similar to those given in Example 4.1.

With our preparations done, we can easily prove one of the most significant properties of the contact element: Its functoriality under $(+1)$-contact surgeries. We will outline the proof since it can be regarded as a model proof.

Theorem 8.21 (see Theorem 4.2 of [15]). — *Let (Y',ξ') be obtained from (Y,ξ) by $(+1)$-contact surgery along a Legendrian knot L. Denote by W the associated cobordism. Then the map*

$$\widehat{F}_{-W}\colon \widehat{\mathrm{HF}}(-Y) \longrightarrow \widehat{\mathrm{HF}}(-Y')$$

preserves the contact element, i.e. $\widehat{F}_{-W}(c(Y,\xi)) = c(Y',\xi')$.

Proof. —

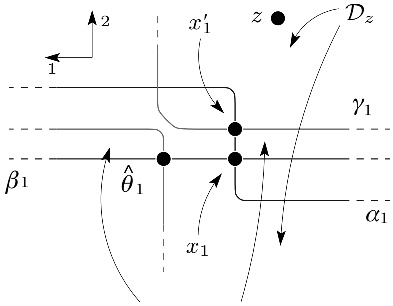

Figure 12. — Significant part of the Heegaard triple diagram.

Let an open book (P, ϕ) adapted to (Y, ξ, L) be given. By Lemma 8.14, a $(+1)$-contact surgery acts on the monodromy as a composition with a negative Dehn twist. Without loss of generality, the knot L just intersects β_1 once, transversely and is disjoint from the other β-circles. Moreover, we can arrange the associated Heegaard triple to look as indicated in Figure 12. The contact element $c(Y, \xi)$ is represented by the point $\{\boldsymbol{x}_1, \dots, \boldsymbol{x}_n\}$. Obviously, there is only one domain which carries a holomorphic triangle. It is the small holomorphic triangle connecting $\boldsymbol{x}_1$ and $\boldsymbol{x}_1'$ (cf. §4.4). Thus, there is only one domain with positive coefficients, with $n_z = 0$, connecting the points $\{\boldsymbol{x}_1, \dots, \boldsymbol{x}_n\}$ with $\{\boldsymbol{x}_1', \dots, \boldsymbol{x}_n'\}$. By considerations similar to those given at the end of the proof of Lemma 4.9, we see that the associated moduli space is a one-point space. Hence, the result follows. □

8.5. The Legendrian knot invariant

Ideas very similar to those used to define the contact element can be utilized to define an invariant of Legendrian knots we will briefly call LOSS. This invariant is due to **L**isca, **O**zsváth, **S**tipsicz and **S**zabó and was defined in [10]. It is basically the contact element but now it is interpreted as sitting in a filtered Heegaard Floer complex. The filtration is constructed with respect to a fixed Legendrian knot.

Let (Y, ξ) be a contact manifold and $L \subset Y$ a Legendrian knot. There is an open book decomposition of Y, subordinate to ξ, such that L sits on the page $P \times \{1/2\}$ of the open book (cf. §8.3). Choose a cut system that induces an L-adapted Heegaard diagram (cf. §8.4, Definition 8.18 and Lemma 8.19). Figure 13 illustrates the positioning of the point w in the Heegaard diagram induced by the open book. Similar to the case of the contact element the intersection points of $\alpha_i \cap \beta_i$ which sit on $P \times \{1/2\}$ determine a specific generator of $\widehat{\mathrm{CF}}(-Y)$. This element can be interpreted as sitting in $\widehat{\mathrm{CFK}}(-Y, L)$ and it is a cycle. The induced element in the knot Floer homology is denoted by $\widehat{\mathcal{L}}(L)$.

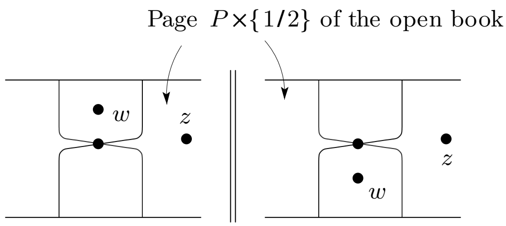

Figure 13. — Positioning of the point w depending on the knot orientation.

Remark 8.1. —

(1) Since this is an important issue we would like to recall the relation between the pair (w, z) and the knot orientation. In homology we connect z with w in the complement of the α-curves and w with z in the complement of the β-curves (oriented as is obvious from the definition). In **cohomology**, *we orient in the opposite manner, i.e. we move from z to w in the complement of the β-curves and from w to z in the complement of the α-curves.*

(2) Observe, that the definition of the invariant $\widehat{\mathcal{L}}$ as well as the contact element always comes with a specific presentation of the groups $\widehat{\mathrm{HF}}$ and $\widehat{\mathrm{HFK}}$. If we want to compare for instance invariants of two different Legendrian knots, we have to get rid of the presentation in the background. This can be done by modding out a certain mapping class group action on the homologies. We point the reader to [13].

Analogous to the properties of the contact element, the invariant $\widehat{\mathcal{L}}$ is preserved when a $(+1)$-contact surgery is performed in its complement (see [13, Corollary 3.4]): Suppose we are given a contact manifold (Y, ξ) with two Legendrian knots L and S sitting in it. Performing a $(+1)$-contact surgery along S, denote by W the associated cobordism. Furthermore, we denote by (Y_S, ξ_S) the result of the contact surgery. The cobordism $-W$ induces a map

$$\widehat{F}_{-W} \colon \widehat{\mathrm{HFK}}(-Y, L) \longrightarrow \widehat{\mathrm{HFK}}(-Y_S, L_S)$$

such that $\widehat{F}_{-W}(\widehat{\mathcal{L}}(L)) = \widehat{\mathcal{L}}(L_S)$. Here, L_S denotes the knot L in the manifold Y_S. Observe, that the cobordism maps constructed for the hat-theory can be defined the same way for knot Floer homologies.

Finally, the contact element and the invariant $\widehat{\mathcal{L}}$ are connected, too. Performing a $(+1)$-contact surgery along the knot L, denote by W the associated cobordism and by (Y_L, ξ_L) the result of the contact surgery. The cobordism $-W$ induces a map

$$\Gamma_{-W}\colon \widehat{\mathrm{HFK}}(-Y, L) \longrightarrow \widehat{\mathrm{HF}}(-Y_L^+)$$

such that $\Gamma_{-W}(\widehat{\mathcal{L}}(L)) = c(\xi_L^+)$ (see [26, Theorem 6.1]). This map is not defined by counting holomorphic triangles. It needs a specific construction we do not outline here. We point the interested reader to [26].

Bibliography

[1] Bredon (G. E.). — Geometry and Topology, Graduate Texts in Mathematics 139, Springer-Verlag (1993).

[2] Ding (F.) and Geiges (H.). — Symplectic fillability of tight contact structures on torus bundles, Algebr. Geom. Topol. 1, p. 153-172 (2001).

[3] Ding (F.) and Geiges (H.). — A Legendrian surgery presentation of contact 3-manifolds, Math. Proc. Cambridge Philos. Soc. 136, p. 583-598 (2004).

[4] Eliashberg (Y.). — Topological characterization of Stein manifolds of dimension > 2, Internat. J. Math. 1, p. 29-46.

[5] Etnyre (J. B.). — Lectures on open-book decompositions and contact structures, Amer. Math. Soc. 5, p. 103-142, (Proceedings of the Clay Mathematics Summer School) (2006).

[6] Geiges (H.). — An Introduction to Contact Topology, Cambridge Studies in Advanced Mathematics 109, Cambridge University Press (2008).

[7] Gompf (R. E.) and Stipsicz (A. I.). — 4-Manifolds and Kirby Calculus, Graduate Studies in Mathematics 20, American Mathematical Society (1999).

[8] Honda (K.), Kazez (W. H.), and Matić (G.). — On the contact class in Heegaard Floer homology, J. Diff. Geom. 83(2), p. 289-311 (2009).

[9] Lipshitz (R.), Ozsváth (P.) and Thurston (D.). — Bordered Heegaard Floer homology: Invariance and pairing, `arXiv:0810.0687`.

[10] Lisca (P.), Ozsváth (P.), Stipsicz (A. I.) and Szabó (Z.). — Heegaard Floer invariants of Legendrian knots in contact three-manifolds, J. Eur. Math. Soc (JEMS) 11(6), p. 1307-1363 (2009).

[11] McDuff (D.) and Salamon (D.). — j-Holomorphic Curves and Symplectic Topology, Colloquium Publications 52, American Mathematical Society (2004).

[12] Juhasz (A.). — Holomorphic disks and sutured manifolds, Algebr. Geom. Topol. 6, p. 1429-1457 (2006).

[13] Ozsváth (P.) and Stipsicz (A. I.). — Contact surgeries and the transverse invariant in knot Floer homology, J. Inst. Math. Jussieu bf 9(3) (2010), 601-632.

[14] Ozsváth (P.) and Szabó (Z.). — Holomorphic disks and knot invariants, Adv. Math. 186(1), p. 58-116 (2004).

[15] Ozsváth (P.) and Szabó (Z.). — Heegaard Floer homologies and contact structures, Duke Math. J. 129(1), p. 39-61 (2005).

[16] Ozsváth (P.) and Szabó (Z.). — Heegaard diagrams and holomorphic disks, Diff. faces of Geom., Int. Math. Series p. 301-348.

[17] Ozsváth (P.) and Szabó (Z.). — Holomorphic disks and topological invariants for closed three-manifolds, Ann. of Math. 159(3), p. 1027-1158 (2004).

[18] Ozsváth (P.) and Szabó (Z.). — Holomorphic disks and three-manifold invariants: Properties and applications, Ann. of Math. 159(3), p. 1159-1245 (2004).

[19] Ozsváth (P.) and Szabó (Z.). — Holomorphic triangles and invariants of smooth four-manifolds, Adv. Math. 202(2), p. 326-400 (2006).

[20] Ozsváth (P.) and Szabó (Z.). — On the Heegaard Floer homology of branched double-covers, Adv. Math. 194, p. 1-33 (2005).

[21] Ozsváth (P.) and Szabó (Z.). — Introduction to Heegaard Floer theory, Clay Math. Proc. 5, p. 3-28 (2006).

[22] Ozsváth (P.) and Szabó (Z.). — Lectures on Heegaard Floer homology, Clay Math. Proc. 5, p. 29-70 (2006).

[23] Ozsváth (P.) and Szabó (Z.). — Heegaard diagrams and holomorphic disks, Different faces of geometry p. 301-348.

[24] Ozsváth (P.) and Szabó (Z.). — On the skein exact sequence for knot Floer homology, `arXiv:0707.1165`.

[25] Ozbagci (B.) and Stipsicz (A. I.). — Surgery on Contact 3-Manifolds and Stein Surfaces, Bolyai Society Mathematical Studies 13, Springer-Verlag (2004).

[26] Sahamie (B.). — Dehn twists in Heegaard Floer homology, Algebr. Geom. Topol. 10, p. 465-524 (2010).

[27] Sarkar (S.) and Wang (J.). — An algorithm for computing some Heegaard Floer homologies, Ann. of Math., vol. 171 (2), p. 1213-1236 (2010).

[28] Turaev (V.). — Torsion invariants of Spin^c-structures on 3-manifolds, Math. Research Letters 6, p. 679-695 (1997).