

DEHN TWISTS AND HEEGAARD FLOER HOMOLOGY

Inaugural-Dissertation

zur Erlangung des Doktorgrades

der Mathematisch-Naturwissenschaftlichen Fakultät

der Universität zu Köln

vorgelegt von

BIJAN SAHAMIE

aus Oberhausen

2009

Berichterstatter: Prof. Hansjörg Geiges, Ph.D. (Cantab)
Prof. Dr. George Marinescu

Tag der mündlichen Prüfung: 16.10.2009

Zusammenfassung

Wir leiten eine Beschreibung der Hut-Version der Heegaard Floer Homologie her im Falle, dass das zugehörige Heegaard Diagramm durch einen Dehn Twist modifiziert wurde. Als Resultat dieser Beschreibung erhalten wir eine neue exakte Sequenz in der Hut-Version der Heegaard Floer Homologie. Um den in der Beschreibung und den Sequenzen auftauchenden Moduln eine geeignete geometrische Interpretation zu geben, verallgemeinern wir die Knotenhomologie $\widehat{\text{HFK}}$ auf homologisch nicht-triviale Knoten und schwächen die Zulässigkeitsbedingungen in ihrer Definition ab. Als Teil der gewonnenen exakten Sequenzen erhalten wir eine Abbildung von der wir zeigen, dass sie nicht von den Wahlen abhängt, die für ihre Definition notwendig sind, sondern nur vom Kobordismus abhängt, der durch den Dehn Twist induziert wird. Mit dieser Abbildung leiten wir eine Transformationsregel her, welche die Invariante für Legendre-Knoten und die Kontaktklasse miteinander verbindet. Wir geben drei Anwendungen dieser Beziehung. Zuletzt beschäftigen wir uns mit der Beziehung der neu gewonnenen exakten Sequenz und dem bekannten exakten Chirurgiedreieck in der Knotenhomologie. Mit einer geeigneten Modifikation ihres Konstruktionsprozesses sind wir in der Lage eine starke Beziehung zu den neu gewonnenen exakten Sequenzen herzuleiten mit dem Ergebnis, dass wir einen Zusammenhang herstellen zwischen dem Zählen holomorpher Dreiecke in zweifach-punktierten Heegaard-Trippeln und dem Zählen holomorpher Scheiben in punktierten Heegaard Diagrammen.

Abstract

We derive a representation of the hat-version of Heegaard Floer homology in case we change the associated Heegaard diagram with a Dehn Twist. Result of this description is a new exact sequence in the hat-version of Heegaard Floer homology. To give the involved modules a suitable geometric interpretation, we generalize the knot Floer homology $\widehat{\text{HFK}}$ to homologically non-trivial knots and relax the admissibility conditions used in their definition. As part of the exact sequence we obtain a map, which we show not to depend on the choices made in its definition, but on the cobordism induced by the Dehn Twist. With this map we derive a naturality property between the invariant of Legendrian knots $\widehat{\mathcal{L}}$ and the contact element and give three applications. Finally, we investigate the relationship between the newly defined exact sequences and the well-known surgery exact triangle in knot Floer homology. With a suitable

modification of the construction process of the surgery exact triangle we derive a strong relationship to the newly defined exact sequences. This, finally, results in a relationship between counting holomorphic triangles in doubly-pointed Heegaard triple diagrams and counting holomorphic discs in pointed Heegaard diagrams.

Contents

1	Introduction	7
2	Introduction to HF Theory	11
2.1	Introduction to $\widehat{\text{HF}}$ as a Model for Heegaard Floer Theory	11
2.1.1	Heegaard Diagrams	11
2.1.2	Introduction to $\widehat{\text{HF}}$ — Topology and Analysis	12
2.1.3	The Structure of the Moduli Spaces	23
2.1.4	Choice of Almost Complex Structure	30
2.1.5	Dependence on the Choice of Orientation Systems	31
2.2	The Homologies HF^∞ , HF^+ , HF^-	33
2.3	Topological Invariance	34
2.3.1	Stabilizations/Destabilizations	36
2.3.2	Independence of the Choice of Almost Complex Structures	37
2.3.3	Isotopy Invariance	40
2.3.4	Handle slide Invariance	42
2.4	Knot Floer Homologies	49
2.4.1	Refinements	54
2.5	Maps Induced By Cobordisms	55
2.6	The Surgery Exact Triangle	57

2.7	The Contact Element and $\widehat{\mathcal{L}}$	61
2.7.1	Contact Structures	61
2.7.2	Open Books	63
2.7.3	Open Books, Contact Structures and Heegaard Diagrams	64
2.7.4	The Contact Class	68
2.7.5	The Invariant $\widehat{\mathcal{L}}$	73
3	Dehn Twists in $\widehat{\text{HF}}$ Homology	75
3.1	Algebraic Preliminaries	75
3.2	Two New Exact Sequences in Heegaard Floer Homology	77
3.2.1	Positive Dehn Twists	77
3.2.2	Negative Dehn Twists	86
3.3	Invariance	87
3.4	Implications to Contact Geometry	102
3.4.1	Stabilizations of Legendrian Knots and Open Books	106
3.5	Applications – Vanishing Results of the Contact Element	116
4	Holomorphic Discs and Surgery Exact Triangles	123
4.1	Surgery Exact Triangle and Dehn Twist Sequence	125
4.2	Chain Maps and Holomorphic Discs	137
4.2.1	General Definition	137
4.2.2	Properties	139
	Bibliography	144

Chapter 1

Introduction

At the beginning of the new millennium Ozsváth and Szábo defined a Floer-type homology theory called Heegaard Floer homology (in the following HFT), assigning to a Spin^c -3-manifold (Y, s) a bunch of homologies, which are all connected with each other by exact sequences (see [40], [39]). As all Floer homologies it has its origins in the work of Gromov (see [19]), who brought holomorphic curves into the realm of symplectic geometry, and the work of Floer, who was the first to transfer the Morse homological scheme to the symplectic category (see [10],[11], [12], [13] and [14]). From that time many flavors of Floer homologies arose like for instance Seiberg-Witten Floer homology (see [23]). The motivation for the development of HFT was to give a more topological description of Seiberg-Witten theory (see [43]). Those two theories are conjecturally equivalent and there were some efforts made to bring those two theories together, with some success, as Taubes just recently showed in [50] the Seiberg Witten Floer homology to be isomorphic to embedded contact homology, and coming from the other side, Lipshitz giving the cylindrical reformulation of HFT (see [24]). It developed to a highly active research field with many applications and contributions in knot theory but also in contact geometry. Besides the applications, the theory itself was brought forward with recent extension of HFT to bordered manifolds. And there are two flavors of the bordered invariant, a topological and a geometric version: The Sutured Floer homology of András Juhasz (see [22]), which we interpret as a geometric degeneration of the topological theory, and the topological theory given by Robert Lipshitz, Peter Ozsváth and Dylan Thurston in [26].

Contact geometry in turn is among the important research fields of modern geometry. First of all, contact geometry developed a rich theory, which makes it a valuable field

of its own right. But besides its intrinsic value, contact geometry contributed to low-dimensional topology very fruitfully as elegant contact geometric proofs arose from it for delicate geometric theorems. Examples to mention would be Cerf's famous proof of $\Gamma_4 = 0$ (cf. [16]) or Geiges' elegant contact geometric proof of the Whitney-Graustein theorem (see [15]).

To a contact manifold (Y, ξ) one can associate an isotopy invariant $c(\xi)$ of ξ , the contact element, which is a class in the HFT $\widehat{\text{HF}}(-Y)$ of $-Y$. Furthermore if we additionally fix a Legendrian knot L we may associate a Legendrian isotopy invariant $\widehat{\mathcal{L}}(L)$ of the Legendrian knot in the associated knot Floer homology $\widehat{\text{HFK}}(-Y, L)$ of the pair $(-Y, L)$. Paolo Lisca and András Stipsicz showed in a series of papers (see [28], [29], [30], [31] and with Ghiggini [17]) that there are examples of families of contact structures where conventional topological techniques fail to detect tightness, the contact element however does. The contact element in the hands of Lisca and Stipsicz has turned out to be a very powerful tool in generating tight contact structures.

The theme of this thesis may be located exactly between the two fields of HFT and contact topology. The original question we tried to answer was if the contact element, in case it is non trivial, is always primitive, or if there are cases where it is not a primitive element. The most natural approach for tackling this problem is the one used in this thesis. Let (P, ϕ) be an open book decomposition adapted to the contact structure ξ . How does the Heegaard Floer homology of $(P, D_\delta \circ \phi)$ look like, where $\delta \subset P$ is a homologically essential embedded closed curve in P and D_δ denotes a Dehn Twist along δ ? This question is closely related to the first one since Dehn Twists of the given type can be translated into contact surgeries, which in turn can be used to generate every contact manifold. We were not able to answer the question concerning the primitiveness of the contact element. However, we discovered some new theory which will be the focus of this thesis.

What is the contribution of this thesis?

Chapter 2 is an introduction to Heegaard Floer homology with some emphasis on the hat-theory. We are aware of the existence of introductory articles to this subject but we tried to give an introduction without sweeping important details under the carpet. We do not want to discredit the existing literature; the existing literature is very well written. But we provide a different focus, and we believe that there is a lack of literature with this kind of point of view. We are indeed convinced that this chapter can help graduate students or researchers, especially those outside of Columbia, Princeton or other places with a local expert on this subject, to understand the material. This introduction was never meant to be complete or to give an overview of the given theory. We focus more

on giving the foundations and hope that after reading this first chapter the reader has developed intuition enough to understand the research literature without getting lost.

In Chapter 3 we derive a new representation of $\widehat{\text{HF}}(P, D_\delta \circ \phi)$ (Propositions 3.2.1 and 3.2.5). A consequence of this representation are the exact sequences given in Corollaries 3.2.2 and 3.2.6. These exact sequences have interesting implications. The most important contact geometric implication is Proposition 3.4.1. We set up a naturality property between the isotopy invariant of Legendrian knots and the contact element and give three applications (Proposition 3.5.1, Proposition 3.5.3 and Theorem 3.5.4). There are some problems occurring we would like to mention:

- (a) The representation of $\widehat{\text{HF}}(P, D_\delta \circ \phi)$ given in Propositions 3.2.1 and 3.2.5 describes this group as a mapping cone of two complexes which happen to be the knot Floer homologies in case the induced pair of base points (w, z) induces a null-homologous knot. However, in most situations this will not be the case. We need a geometric interpretation of these modules.
- (b) The diagram describing one of these modules does not in general fulfill the weak admissibility conditions. These are important ingredients in the compactification of the moduli spaces involved in the definition of the differentials.

Both problems (a) and (b) require a *generalization* of the given HFT which we provide in this thesis (see §2.4). However, we have to remark that the given theory already inherits all ingredients to set up the generalizations. So we cannot really say we generalized the theory but we made the observation that the given theory is not restricted to the cases where Ozsváth and Szabó define it. The knot Floer homology seems to have some interesting properties when homologically non-trivial knots come into play. There is a knot class K in $\mathbb{S}^2 \times \mathbb{S}^1$ whose associated knot Floer homology vanishes. This fact is central in the proof of Theorem 3.5.4. This is the first example we know with this property.

In Chapter 3 we investigate the relationship between the sequences given in Corollaries 3.2.2 and 3.2.6 and the well-known surgery exact triangle in knot Floer homology. We see that with a slight modification of the construction process of the surgery exact triangle we are able to define a surgery exact triangle in the knot Floer homology involving the cobordism maps \widehat{F} . Indeed this sequence and the one defined in Chapter 3 stay in a strong relationship which we outline in Theorem 4.1.5. In consequence we see that the sequences given in Chapter 3 admit refinements with respect to Spin^c -structures and can be defined with coherent orientations. This consequence is summarized very

briefly in Corollary 4.1.8. Secondly we learn that there is a relationship between counting holomorphic triangles in Heegaard triple diagrams and counting holomorphic discs in Heegaard diagrams: The sequences from Chapter 3 are induced by short exact sequences of chain complexes. The induced connecting morphism f can be defined by counting holomorphic discs with suitable boundary conditions. The relationship in Chapter 4 relates this map to the cobordism maps in knot Floer homology. These cobordism maps are defined by a count of holomorphic triangles with suitable boundary conditions. Finally we derive properties of the connecting morphisms f .

Acknowledgments

First of all, my thanks go to my advisor Hansjörg Geiges for pointing my interest on Heegaard Floer theory, funding my work, his support concerning my stay at Columbia University and his constructive comments, which helped making the exposition clearer at numerous spots. My sincere gratitude goes to András Stipsicz for giving me an introduction to Heegaard Floer homology, our fruitful and enlightening discussions, his constructive comments, and for pointing my interest on [20], which – in some way – builds the foundation of my work. Furthermore, I thank Peter Ozsváth for our conversations during my stay at Columbia, which were both fun and inspiring. One of these conversations motivated what is done in the first part of chapter 4.

On the non-mathematical side I have to thank my wife Cornelia who always tolerated and understood my mental absence during intense phases of my work. And at the institute there was Eva Nowak who always shared her complaints with me and who always had an open ear to mine, sometimes with coffee and cookies.

Finally, thanks go to the DAAD for the financial support concerning my stay at Columbia University.

Chapter 2

Introduction to HF Theory

2.1 Introduction to $\widehat{\text{HF}}$ as a Model for Heegaard Floer Theory

2.1.1 Heegaard Diagrams

From the Geometric Topologists' point of view 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. [18]). 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 shortly discuss, what is to be understood as Heegaard diagram throughout this thesis.

A short 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 $Y \setminus \Sigma$ decomposes into two handlebodies H_0 and H_1 . We fix a handle decomposition of $\overline{Y \setminus H_1}$ relative to this splitting surface Σ , i.e. there are 2-handles $h_{1,i}^2$, $i = 1, \dots, g$,

and a 3-handle h_1^3 such that (cf. [18])

$$Y \setminus H_1 \cong (\Sigma \times [0, 1]) \cup_{\partial} (h_{1,1}^2 \cup_{\partial} \dots \cup_{\partial} h_{1,g}^2 \cup_{\partial} h_1^3). \quad (2.1.1)$$

We can rebuild Y from this by gluing in two handles $h_{0,i}^2$, $i = 1, \dots, g$, and a 3-handle h_0^3 . Hence, Y can be written as

$$Y \cong (h_0^3 \cup_{\partial} h_{0,1}^2 \cup_{\partial} \dots \cup_{\partial} h_{0,g}^2) \cup_{\partial} (\Sigma \times [0, 1]) \cup_{\partial} (h_{1,1}^2 \cup_{\partial} \dots \cup_{\partial} h_{1,g}^2 \cup_{\partial} h_1^3). \quad (2.1.2)$$

Collecting the data from this decomposition we obtain a triple (Σ, α, β) where Σ is the splitting surface of genus g , $\alpha = \{\alpha_1, \dots, \alpha_g\}$ are the images of the attaching circles of the $h_{0,i}^2$ interpreted as sitting in Σ and $\beta = \{\beta_1, \dots, \beta_g\}$ the images of the attaching circles of the 2-handles $h_{1,i}^2$ 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 $h_{0,i}^2$. 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 $h_{0,i}^1$ over $h_{0,j}^1$ means, in the dual picture, that $h_{0,j}^2$ is slid over $h_{0,i}^2$.

2.1.2 Introduction to $\widehat{\text{HF}}$ — Topology and Analysis

Given a closed, oriented 3-manifold Y , we fix a Heegaard diagram (Σ, α, β) of Y as defined in §2.1.1. We can associate to it the triple $(\text{Sym}^g(\Sigma), \mathbb{T}_{\alpha}, \mathbb{T}_{\beta})$ which we will explain now:

By $\text{Sym}^g(\Sigma)$ 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.

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

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

$$\pi: \Sigma^{\times g} \longrightarrow \text{Sym}^g(\Sigma)$$

the projection map.

The attaching circles α and β define submanifolds

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

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

The chain complex

Define $\widehat{\text{CF}}(\Sigma, \alpha, \beta)$ as the free \mathbb{Z} -module (or \mathbb{Z}_2 -module) generated by the intersection points $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$ inside $\text{Sym}^g(\Sigma)$.

Definition 2.1.1. A map ϕ of the 2-disc \mathbb{D}^2 (regarded as the unit 2-disc in \mathbb{C}) into the symmetric product $\text{Sym}^g(\Sigma)$ is said to **connect** two points $x, y \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ if

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

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

In case $g \leq 2$ we have to define the object $\pi_2(x, y)$ 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 [40].

Fixing a point $z \in \Sigma \setminus (\alpha \cup \beta)$, we can construct a differential

$$\widehat{\partial}_z: \widehat{\text{CF}}(\Sigma, \alpha, \beta) \longrightarrow \widehat{\text{CF}}(\Sigma, \alpha, \beta)$$

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

$$\widehat{\partial}_z x = \sum_{y \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} \widehat{\partial}_z x \Big|_y \cdot y$$

of all intersection points $y \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$. The definition of the coefficients will occupy the remainder of this paragraph. The idea resembles other Floer homology theories.

The goal is to define $\widehat{\partial_z x} \Big|_y$ as a signed count of holomorphic Whitney discs connecting x and 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 §2.1.3. 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 disc, 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

$$\frac{\partial \phi}{\partial s}(s, t) + \mathcal{J}_s \left(\frac{\partial \phi}{\partial t}(s, t) \right) = 0. \quad (2.1.3)$$

Looking into (2.1.3) it is easy to see that a holomorphic Whitney disc ϕ can be reparametrized by a constant shift in \mathbb{R} -direction without violating (2.1.3).

Definition 2.1.2. Given two points $x, y \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, we denote by $\mathcal{M}_{\mathcal{J}_s}(x, y)$ the set of holomorphic Whitney discs connecting x and y . We call this set **moduli space of holomorphic Whitney discs** connecting x and y . Given a homotopy class $[\phi] \in \pi_2(x, y)$, 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. §2.1.3 or see [40]) 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

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

The so-called **signed count** of 0-dimensional components of $\widehat{\mathcal{M}}(x, y)$ 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 $\widehat{\mathcal{M}}(x, y)$. 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 $\widehat{\mathcal{M}}(x, y)$ correspond to the 1-dimensional components of $\mathcal{M}(x, y)$. 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 $\widehat{\mathcal{M}}(x, y)$, a sign. The signed count will respect the signs attached. There is a technical condition called **coherence** (see [40] or cf. §2.1.3) one has to impose on the orientations. This technical condition ensures that the morphism $\widehat{\partial}_z$ is a differential.

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

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

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

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

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

Theorem 2.1.3 (see [40]). *The assignment $\widehat{\partial}_z$ is well-defined.*

Theorem 2.1.4 (see [40]). *The morphism $\widehat{\partial}_z$ is a differential.*

We will give sketches of the proofs of the last two theorems later in §2.1.3. At the moment we do not know enough about Whitney discs and the symmetric product to prove it.

Definition 2.1.5. We denote by $\widehat{\text{CF}}(\Sigma, \alpha, \beta, z)$ the chain complex given by the data $(\widehat{\text{CF}}(\Sigma, \alpha, \beta), \partial_z)$. Denote by $\widehat{\text{HF}}(Y)$ the induced homology theory $H_*(\widehat{\text{CF}}(\Sigma, \alpha, \beta), \partial_z)$.

The notation should indicate that the homology theory does not depend on the data chosen. 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 2.1.2). Nevertheless the orientations are not written down. We guess there are two reasons: The first would be that most of the time it is not really important which system 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 would be 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 2.1.31.

On Holomorphic Discs in the Symmetric Product

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

Definition 2.1.6. For a Whitney disc ϕ we denote by $\mu(\phi)$ the **formal dimension** of \mathcal{M}_ϕ . 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}_ϕ , although even in case \mathcal{M}_ϕ is not a manifold the number $\mu(\phi)$ is defined (cf. §2.1.3). Just to give some intuition, note that the moduli spaces are the zero-set of a section 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. 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 zero set of the section is a manifold, i.e. 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 2.1.7. *In case $g(\Sigma) > 2$ the 2nd homotopy group $\pi_2(\text{Sym}^g(\Sigma))$ 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 discs.*

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

$$\mathbb{S}^2 \rightarrow \text{Sym}^g(\Sigma), \quad y \mapsto \{(y, \eta(y), y, \dots, y)\}$$

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 [40].

Lemma 2.1.8 (see [40]). *There is an exact sequence*

$$0 \rightarrow \pi_2(\text{Sym}^g(\Sigma)) \rightarrow \pi_2(x, x) \rightarrow \ker(n_z) \rightarrow 0.$$

The map n_z provides a splitting for the sequence.

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

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

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

Lemma 2.1.9. *The Kernel of n_z interpreted as a map on $\pi_2(x, x)$ is isomorphic to $H^1(Y; \mathbb{Z})$.*

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

$$*: \pi_2(x, x) \times \pi_2(x, y) \longrightarrow \pi_2(x, y),$$

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

$$\begin{array}{ccc} \pi_2(x, y) & \xrightarrow{\cong} & \pi_2(x, x) \cong \mathbb{Z} \oplus H^1(Y; \mathbb{Z}) \\ & \searrow & \swarrow \\ & \{*\} & \end{array} \quad (2.1.4)$$

as principal bundles over a one-point space, which is another way of saying that the concatenation action endows $\pi_2(x, y)$ with a group structure after fixing a unit element in $\pi_2(x, y)$. To address the well-definedness of $\widehat{\partial}_z$ we have to show that the sum in the definition of $\widehat{\partial}_z$ 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 2.1.22), hence their signed count is finite. Assuming this property we are able to show well-definedness of $\widehat{\partial}_z$ in case Y is a homology sphere.

Proof of Theorem 2.1.3 for $b_1(Y) = 0$. Observe that

$$\widehat{\mathcal{M}}(x, y)_{n_z=0}^0 = \bigsqcup_{\phi \in H(x, y, 1)} \widehat{\mathcal{M}}_\phi, \quad (2.1.5)$$

where $H(x, y, 1) \subset \pi_2(x, y)$ is the subset of homotopy classes admitting holomorphic representatives with $\mu(\phi) = 1$ and $n_z = 0$. We have to show that $H(x, y, 1)$ is a finite set. Since $b_1(Y) = 0$ the cohomology $H^1(Y; \mathbb{Z})$ vanishes. By our preliminary discussion, given a reference disc $\phi_0 \in \pi_2(x, y)$, any $\phi_{xy} \in \pi_2(x, y)$ can be written

as a concatenation $\phi_{xy} = \phi * \phi_0$, where ϕ is an element in $\pi_2(x, x)$. Since we are looking for discs with index one we have to find all $\phi \in \pi_2(x, x)$ satisfying the property $\mu(\phi) = 1 - \mu(\phi_0)$. Recall that Y is a homology sphere and thus $\pi_2(x, x) \cong \mathbb{Z} \otimes \{S\}$. Hence, the disc 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 disc satisfying the property $\mu = 1$ and $n_z = 0$. \square

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 we have to go into the theory a bit more.

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

Lemma 2.1.10 (see [40]). *The group $\pi_1(\text{Sym}^g(\Sigma))$ is abelian.*

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

$$(\gamma_1, \dots, \gamma_g): \mathbb{S}^1 \rightarrow \Sigma^{\times g}.$$

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

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

Lemma 2.1.11 (see [40]). *The map Φ induces an isomorphism*

$$\Phi_*: H_1(\text{Sym}^g(\Sigma)) \rightarrow H_1(\Sigma; \mathbb{Z}).$$

By surgery theory (see [18], p. 111) we know that

$$\frac{H_1(\Sigma; \mathbb{Z})}{[\alpha_1], \dots, [\alpha_g], [\beta_1], \dots, [\beta_g]} \cong H_1(Y; \mathbb{Z}) \quad (2.1.6)$$

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

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

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

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

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

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

with the following property.

Lemma 2.1.12. *If $\epsilon(x, y)$ is non-zero the set $\pi_2(x, y)$ is empty.*

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

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

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

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

$$s_z: \mathbb{T}_\alpha \cap \mathbb{T}_\beta \longrightarrow \text{Spin}_3^c(Y) \cong H^2(Y; \mathbb{Z}), \quad (2.1.7)$$

such that $\text{PD}(\epsilon(x, y)) = s_z(x) - s_z(y)$. We point the reader interested in the definition of s_z to [40]. Thus, fixing a Spin^c -structure s , the \mathbb{Z} -module (or \mathbb{Z}_2 -module) $\widehat{\text{CF}}(\Sigma, \alpha, \beta, z; s)$ generated by $(s_z)^{-1}(s)$ defines a subcomplex of $\widehat{\text{CF}}(\Sigma, \alpha, \beta, z)$. The associated homology is denoted by $\widehat{\text{HF}}(Y, s)$, and it is a submodule of $\widehat{\text{HF}}(Y)$. Especially note that

$$\widehat{\text{HF}}(Y) = \bigoplus_{s \in \text{Spin}_3^c(Y)} \widehat{\text{HF}}(Y, s).$$

Since $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$ 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 $\Sigma \setminus \{\alpha \cup \beta\}$ there will be a difference between the Spin^c -structure associated to an intersection point. For details we point to [40].

Example 2.1.1. The Heegaard diagram given by the data $(T^2, \{\mu\}, \{\lambda\})$ (cf. §2.1.1) is the 3-sphere. To make use of Lemma 2.1.7 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 . The chain complex $\widehat{\text{CF}}(D, z)$ equals one copy of \mathbb{Z} since it is generated by one single intersection point which we denote by x . We claim that $\widehat{\partial}_z x = 0$. Denote by $[\phi]$ a homotopy class of Whitney discs connecting x with itself. This is a holomorphic sphere which can be seen with Lemma 2.1.8, Lemma 2.1.9 and the fact that $H^1(\mathbb{S}^3) = 0$. By Lemma 2.1.7 the set $\pi_2(\text{Sym}^g(\Sigma))$ 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 $\mathcal{M}(x, x)_{n_z=0}^1$, i.e. the space of holomorphic Whitney discs connecting x with itself, with $\mu = 1$ and $n_z = 0$, is empty. Hence

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

A Low-Dimensional Model for Whitney Discs

The exact sequence in Lemma 2.1.8 combined with Lemma 2.1.9 and (2.1.4) gives an interpretation of Whitney discs as homology classes. Given a disc ϕ , we define its associated homology class by $\mathcal{H}(\phi)$, i.e.

$$0 \longrightarrow \pi_2(\text{Sym}^g(\Sigma)) \longrightarrow \pi_2(x, x) \xrightarrow{\mathcal{H}} H_2(Y; \mathbb{Z}) \longrightarrow 0. \quad (2.1.8)$$

In the following we intend to give a description of the map \mathcal{H} . Given a Whitney disc ϕ , we can lift this disc to a map $\tilde{\phi}$ by pulling back the branched covering π (cf. diagram (2.1.9)).

$$\begin{array}{ccccc} F/S_{g-1} = \widehat{\mathbb{D}} & \xrightarrow{\bar{\phi}} & \Sigma \times \text{Sym}^{g-1}(\Sigma) & \longrightarrow & \Sigma \\ \uparrow & & \uparrow & & \\ \phi^* \Sigma^{\times g} = F & \xrightarrow{\tilde{\phi}} & \Sigma^{\times g} & & \\ \downarrow & & \pi \downarrow & & \\ \mathbb{D}^2 & \xrightarrow{\phi} & \text{Sym}^g(\Sigma) & & \end{array} \quad (2.1.9)$$

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 $\bar{\phi}$ pictured in (2.1.9). Composing it with the projection onto the surface Σ we define a map

$$\hat{\phi}: \widehat{\mathbb{D}} \longrightarrow \Sigma.$$

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

Definition 2.1.13. Denote by $\mathcal{D}_1, \dots, \mathcal{D}_m$ the closures of the components of the complement of the attaching circles $\Sigma \setminus \{\alpha \cup \beta\}$. 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 $\lambda_1, \dots, \lambda_m \in \mathbb{Z}$.

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

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

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

$$\hat{\phi}(\widehat{\mathbb{D}}) = \mathcal{D}(\phi)$$

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

Lemma 2.1.14. *Two Whitney discs $\phi_1, \phi_2 \in \pi_2(x, x)$ are homotopic if and only if their domains are equal.*

Proof. Given two discs ϕ_1, ϕ_2 whose domains are equal, by definition $\mathcal{H}(\phi_1) = \mathcal{H}(\phi_2)$. By (2.1.8) 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 discs 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 discs is very inconvenient. However, not all domains are carried by holomorphic discs. Obviously, the equality $[\mathcal{D}(\phi)] = \widehat{\phi}_*[\widehat{\mathbb{D}}]$ connects the boundary conditions imposed on Whitney discs 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 2.1.11). Suppose we have fixed two intersections $x = \{x_1, \dots, x_g\}$ and $y = \{y_1, \dots, y_g\}$ connected by a Whitney disc ϕ . The boundary $\partial(\phi(\mathbb{D}^2))$ defines a connecting curve γ_{xy} . It is easy to see that

$$\text{im}(\widehat{\phi}|_{\partial\widehat{\mathbb{D}}}) = \Phi(\gamma_{xy}) = \gamma_1 + \dots + \gamma_g.$$

Restricting the γ_i to the α -curves we get a chain connecting the set x_1, \dots, x_g with y_1, \dots, y_g , and restricting the γ_i to the β -curves we get a chain connecting the set y_1, \dots, y_g with x_1, \dots, x_g . This means each boundary component of $\widehat{\mathbb{D}}$ consists of a set of arcs alternating through α -curves and β -curves.

Definition 2.1.15. 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 disc is called **periodic** if its associated domain is a periodic domain. The subgroup of periodic classes in $\pi_2(x, x)$ is denoted by Π_x .

Theorem 2.1.16 (see [40]). *For a Spin^c -structure s and a periodic class $\phi \in \Pi_x$ 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 disc with a topological property. Note that, because of the additivity of the expected dimension μ , the homology groups can be endowed with a relative grading defined by

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

where ϕ is an arbitrary element of $\pi_2(x, y)$. In the case of homology spheres this defines a relative \mathbb{Z} -grading because by Theorem 2.1.16 the expected dimension vanishes for all periodic discs. In case of non-trivial homology they just vanish modulo $\delta(s)$, where

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

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

Definition 2.1.17. A pointed Heegaard diagram $(\Sigma, \alpha, \beta, z)$ is called **weakly admissible** for the Spin^c -structure s if for every non-trivial periodic domain \mathcal{D} such that $\langle c_1(s), \mathcal{H}(\mathcal{D}) \rangle = 0$ the domain has positive and negative coefficients.

With this technical condition imposed the $\widehat{\partial}_z$ is a well-defined map on the subcomplex $\widehat{\text{CF}}(\Sigma, \alpha, \beta, s)$. From admissibility it follows that for every $x, y \in (s_z)^{-1}(s)$ and $j, k \in \mathbb{Z}$ there exists just a finite number of $\phi \in \pi_2(x, y)$ with $\mu(\phi) = j$, $n_z(\phi) = k$ and $\mathcal{D}(\phi) \geq 0$. The last condition means that all coefficients in the associated domain are greater or equal to zero.

Proof of Theorem 2.1.3 for $b_1(Y) \neq 0$. Recall that holomorphic discs 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. §2.1.3) includes that all the $\{z_i\} \times \text{Sym}^{g-1}(\Sigma)$ are complex submanifolds. Thus, holomorphic Whitney discs always satisfy $\mathcal{D}(\phi) \geq 0$. \square

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

Theorem 2.1.18 (see [40]). *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 $\phi: S \rightarrow \Sigma$ 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.*

2.1.3 The Structure of the Moduli Spaces

The material in this paragraph is presented without any details. The exposition pictures the bird's eye view of the material. Recall from the last paragraphs 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 (j, η) be a Kähler structure on the Heegaard surface Σ , i.e. η is a symplectic form and j an almost-complex structure that tames η . Let z_1, \dots, z_m be points, one in each component of $\Sigma \setminus \{\alpha \cup \beta\}$. Denote by V an open neighborhood in $\text{Sym}^g(\Sigma)$ of

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

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

Definition 2.1.19. An almost complex structure J on $\text{Sym}^g(\Sigma)$ is called (j, η, V) -**nearly symmetric** if J agrees with $\text{sym}^g(j)$ over V and if J tames $\pi_*(\eta^{\times g})$ over \overline{V}^c . The set of (j, η, V) -nearly symmetric almost-complex structures will be denoted by $\mathcal{J}(j, \eta, V)$.

The almost complex structure $\text{sym}^g(j)$ on $\text{Sym}^g(\Sigma)$ is the natural almost complex structure induced by the structure j . Important for us is that the structure J agrees with $\text{sym}^g(j)$ on V . This makes the $\{z_i\} \times \text{Sym}^{g-1}(\Sigma)$ complex submanifolds with respect to J . This is necessary to guarantee positive intersections with Whitney discs. Without this property the proof of Theorem 2.1.3 would break down in the case the manifold has non-trivial topology.

We are interested in holomorphic Whitney discs, i.e. discs in the symmetric product which are solutions of (2.1.3). Denote by the $\partial_{\mathcal{J}_s}$ the Cauchy-Riemann type operator defined by equation (2.1.3). Define $\mathcal{B}(x, y)$ as the space of Whitney discs connecting x and y such that the discs converge to x and y exponentially with respect to some Sobolev space norm in a neighborhood of i and $-i$ (see [40]). With these assumptions the solution $\partial_{\mathcal{J}_s}\phi$ lies in a space of L^p -sections

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

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

Theorem 2.1.20. *The bundle $\mathcal{L} \longrightarrow \mathcal{B}(x, y)$ is a Banach bundle.*

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

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

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 2.1.21. A map f between Banach manifolds is called **Fredholm** if for every point p the differential $T_p f$ is a Fredholm operator, i.e. has finite-dimensional kernel and cokernel. The difference $\dim \ker T_p f - \dim \text{coker } T_p f$ 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 2.1.22. (see [40]) *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 $\mathcal{M}_{\mathcal{J}_s}(x, y)$ are smooth manifolds for all $x, y \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$.*

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

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

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, ϕ) , where \mathcal{J}_s is a path of (j, η, V) -nearly symmetric almost complex structures and ϕ a Whitney disc. 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 one-dimensional components to be finite. This means we require the spaces $\widehat{\mathcal{M}}(x, y)_{n_z=0}^0$ to be compact. One ingredient of the compactness is the admissibility property introduced in Definition 2.1.17. In (2.1.5) we observed that

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

where $H(x, y, 1)$ is the set of homotopy classes of Whitney discs with $n_z = 0$ and expected dimension $\mu = 1$. Admissibility guarantees that $H(x, y, 1)$ 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 discs. The existence of this energy bound shows that the moduli spaces $\widehat{\mathcal{M}}(x, y)$ admit a compactification by adding solutions to the space.

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 discs holomorphically. Given two Whitney discs $\phi_1 \in \pi_2(x, y)$ and $\phi_2 \in \pi_2(y, w)$, 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 2.1.23. We call the pair (ϕ_2, ϕ_1) a **broken** holomorphic Whitney disc.

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 $t \rightarrow \infty$ converges to the broken disc (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 disc, 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 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 [33]).

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

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

But are these the only boundary components? If this is the case, by adding broken discs 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 discs in a 0-dimensional moduli space converging to a broken disc. But despite that there is a second reason for us to wish broken flow lines to compactify the moduli spaces. The map $\widehat{\partial}_z$ should be a boundary operator. Calculating $\widehat{\partial}_z \circ \widehat{\partial}_z$ we see that the coefficients in the resulting equation equal the number of boundary components corresponding to broken discs 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 2.1 illustrates the geometric picture behind that phenomenon. Bubbling is some kind of breaking phenomenon but the components here are discs 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 disc 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 §2.1.2. Thus, $\phi = \phi_1 * k \cdot S$ where $n_z(S) = 1$. In consequence $n_z(\phi)$ is non-zero, contradicting the assumptions.




Figure 2.1: Bubbling of spheres.

Definition 2.1.24. For a point $x \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ an α -**degenerate** disc is a holomorphic disc $\phi: [0, \infty) \times \mathbb{R} \rightarrow \text{Sym}^g(\Sigma)$ with the following boundary conditions $\phi(\{0\} \times \mathbb{R}) \subset \mathbb{T}_\alpha$ and $\phi(p) \rightarrow x$ as $x \rightarrow \infty$.

Given a degenerate disc 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 disconnect 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 of Theorem 2.1.4 with \mathbb{Z}_2 -coefficients. Fix an intersection $x \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$. We compute

$$\begin{aligned} \widehat{\partial}_z x &= \widehat{\partial}_z \left(\sum_{y \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} \# \widehat{\mathcal{M}}(x, y)_{n_z=0}^0 \cdot y \right) \\ &= \sum_{y, w \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} \# \widehat{\mathcal{M}}(x, y)_{n_z=0}^0 \# \widehat{\mathcal{M}}(y, w)_{n_z=0}^0 \cdot w. \end{aligned}$$

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

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

Gluing gives an injection

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

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

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

which shows 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

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

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}_ϕ . 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

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

Thus, a section of the determinant line bundle defines an orientation of \mathcal{M}_ϕ . 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 $o(\phi_1)$ and $o(\phi_2)$ be sections of the determinant line bundles of the associated moduli spaces, then obviously we need that under the identification given above we have

$$o(\phi_1) \wedge o(\phi_2) = o(\phi_2 * \phi_1). \tag{2.1.10}$$

In consequence, a **coherent system of orientations** is a section $o(\phi)$ of the determinant line bundle $\det(\phi)$ for each homotopy class of Whitney discs ϕ connecting two intersection points such that equation (2.1.10) 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 giving a construction for

coherent systems of orientations. Namely, fix a Spin^c -structure s and let $\{x_0, \dots, x_l\}$ be the points representing s , i.e. $(s_z)^{-1}(s) = \{x_0, \dots, x_l\}$. Let ϕ_1, \dots, ϕ_q be a set of periodic classes in $\pi_2(x_0, x_0)$ representing a basis for $H^1(Y; \mathbb{Z})$, denote by θ_i an element of $\pi_2(x_0, x_i)$. A coherent system of orientations is constructed by choosing sections over all chosen discs, i.e. $o(\phi_i)$, $i = 1, \dots, q$ and $o(\theta_j)$, $j = 1, \dots, l$. Namely, for each homotopy class $\phi \in \pi_2(x_i, x_j)$ we have a representation (cf. Lemma 2.1.8, Lemma 2.1.9 and (2.1.4))

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

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

To give a proof of Theorem 2.1.4 in case of \mathbb{Z} -coefficients we have to translate orientations on the 0-dimensional components of the moduli spaces $\widehat{\mathcal{M}}_{J_s}(x, y)$ of connecting Whitney discs into signs. For ϕ with $\mu(\phi) = 1$ the translation action naturally induces an orientation on \mathcal{M}_ϕ . 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 2.1.4 for \mathbb{Z} -coefficients. We stay in the notation of the earlier proof. With the coherent system of orientations introduced we made the map

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

orientation preserving. Hence, we see that $c(x, w)$ equals

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

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

More General Theories

There are variants of Heegaard Floer homology which do not force the condition $n_z = 0$. To make the compactification work in that case we have to take care of boundary degenerations and spheres bubbling off. Both can be shown to be controlled in the sense that the proof of Theorem 2.1.4 for the general theories works the same way with some slight additions due to bubbling and degenerations. The thesis mainly

focuses on the $\widehat{\text{HF}}$ -theory, so we mostly exclude these matters from our exposition. Note just that we get rid of bubbling by a proper choice of almost complex structure. By choosing j on Σ appropriately there is a contractible open neighborhood of $\text{sym}^g(j)$ in $\mathcal{J}(j, \eta, V)$ for which all spheres miss the intersections $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$. Moreover, for a generic choice of path $(\mathcal{J}_s)_{s \in [0,1]}$ inside this neighborhood the signed count of degenerate discs is zero. With this information it is easy to modify the given proof for the general theories. We leave this to the interested reader or point him to [40].

2.1.4 Choice of Almost Complex Structure

Let Σ be endowed with a complex structure j and let $U \subset \Sigma$ be a subset diffeomorphic to a disc.

Theorem 2.1.25 (Riemann mapping theorem). *There is a 3-dimensional family of holomorphic identifications of U with the unit disc $\mathbb{D} \subset \mathbb{C}$.*

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

Corollary 2.1.26. *Let $\phi: \mathbb{D}^2 \rightarrow \text{Sym}^g(\Sigma)$ be a holomorphic disc with $\mathcal{D}(\phi)$ isomorphic to a disc. Then the moduli space $\widehat{\mathcal{M}}_\phi$ contains a unique element.*

There are several ways to achieve this special situation. 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 2.1.27. *Let $\phi \in \pi_2(x, y)$ be an α -injective homotopy class and j a complex structure on Σ . For generic perturbations of the α -curves the moduli space $\mathcal{M}_{\text{sym}^g(j), \phi}$ 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 structure in such a way that homotopy classes of Whitney discs with disc-shaped domains just admits a unique element. This is exactly what can be achieved in general to make the $\widehat{\text{HF}}$ -theory combinatorial. For a class of Heegaard diagrams called **nice** diagrams all moduli spaces with $\mu = 1$ just admits one single element. In addition we have an exact description of how these domains look like. In \mathbb{Z}_2 -coefficients with nice diagrams this results in a method of calculating the differential $\widehat{\partial}_z$ by counting the number of domains that fit into the scheme. This is successfully done for instance for the $\widehat{\text{HF}}$ -theory in [47].

Definition 2.1.28 (see [47]). A pointed Heegaard diagram $(\Sigma, \alpha, \beta, z)$ is called **nice** if any region not containing z is either a bigon or a square.

Definition 2.1.29 (see [47]). A homotopy class is called an empty embedded $2n$ -gon if it is topologically an embedded disc with $2n$ vertices at its boundary, it does not contain any 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 all homotopy classes $\phi \in H(x, y, 1)$ with $\mu(\phi) = 1$ that admit holomorphic representatives are empty embedded bigons or empty embedded squares. Furthermore, for a generic choice of 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. Thus, the theory can be computed combinatorially. We note the following property.

Theorem 2.1.30 (see [47]). *Every 3-manifold admits a nice Heegaard diagram.*

2.1.5 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 o and o' be two orientation systems. We measure their difference

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

by saying that, given a periodic class $\phi \in \pi_2(x, x)$, we define $\delta(\phi) = 0$ if $o(\phi)$ and $o'(\phi)$ coincide, i.e. define equivalent sections, and $\delta(\phi) = 1$, if $o(\phi)$ and $o'(\phi)$ 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 2.1.2. 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

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

Hence we can interpret Spin^c -structures as homomorphisms $\mathbb{Z} \longrightarrow \mathbb{Z}$. For a number $q \in \mathbb{Z}$ define s_q to be the Spin^c -structure whose associated characteristic class, which

we also call s_q , is given by $s_q(1) = q$. 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

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

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 2.2. By boundary orientations $\mathbb{Z} \langle \langle \mathcal{D}_1 - \mathcal{D}_2 \rangle \rangle$ are all possible periodic domains.




Figure 2.2: An admissible Heegaard diagram for $\mathbb{S}^2 \times \mathbb{S}^1$.

Figure 2.2 shows that the chain module is generated by the points x and y . A straightforward computation gives $\epsilon(x, y) = 0$ (see §2.1.2 for a definition) and, hence, both intersections belong to the same Spin^c -structure we will denote by s_0 . Thus, the chain complex $\widehat{\text{CF}}(\Sigma, \alpha, \beta; s_0)$ equals $\mathbb{Z} \otimes \{x, y\}$. The regions \mathcal{D}_1 and \mathcal{D}_2 are both disc-shaped and hence α -injective. Thus, the Riemann mapping theorem (see §2.1.4) gives

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

These two discs 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. Hence, we may choose the signs on ϕ_1 and ϕ_2 arbitrarily. Thus, there are two equivalence classes of orientation systems. We define o^* to be the system of orientations where the signs differ and o_0 where they are equal. Thus, we get two different homology theories

$$\begin{aligned} \widehat{\text{HF}}(\mathbb{S}^2 \times \mathbb{S}^1, s_0; o^*) &= \mathbb{Z} \oplus \mathbb{Z} \\ \widehat{\text{HF}}(\mathbb{S}^2 \times \mathbb{S}^1, s_0; o_0) &= \mathbb{Z}_2. \end{aligned}$$

However, there is a special choice of coherent orientation systems. We point the reader to §2.2 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 these Heegaard Floer group. The resulting group is denoted by $\underline{\text{HF}}^\infty$. We point the reader to [40] for a precise definition. As a matter of completeness we cite:

Theorem 2.1.31 (see [39], Theorem 10.12). *Let Y be a closed oriented 3-manifold. Then there is a unique equivalence class of orientation system such that for each torsion Spin^c -structure s_0 there is an isomorphism*

$$\underline{\text{HF}}^\infty(Y, s_0) \cong \mathbb{Z}[U, U^{-1}]$$

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

2.2 The Homologies HF^∞ , HF^+ , HF^-

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

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

where μ^{-1} are the homotopy classes in $\pi_2(x, y)$ 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 §2.1 breaks down.

Definition 2.2.1. A Heegaard diagram $(\Sigma, \alpha, \beta, z)$ is called **strongly admissible** for the Spin^c -structure s if for every non-trivial periodic domain \mathcal{D} such that $\langle c_1(s), H(\mathcal{D}) \rangle = 2n \geq 0$ 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 discs contribute to the moduli space $\mathcal{M}_{\mathcal{J}, s}(x, y)$ (cf. §2.1).

Theorem 2.2.2. *The map ∂^- is a differential.*

As mentioned in §2.1, in this more general case we have to take a look at bubbling and degenerate discs. The proof follows the same lines as the proof of Theorem 2.1.4. With the remarks made in §2.1 it is easy to modify the given proof to a proof of Theorem 2.2.2 (see [40]). We define

$$\mathrm{CF}^\infty(\Sigma, \alpha, \beta; s) = \mathrm{CF}^-(\Sigma, \alpha, \beta; s) \otimes_{\mathbb{Z}[U^{-1}]} \mathbb{Z}[U, U^{-1}]$$

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 $\mathrm{CF}^+(\Sigma, \alpha, \beta, s)$. Finally we get back to $\widehat{\mathrm{CF}}$ by

$$\widehat{\mathrm{CF}}(\Sigma, \alpha, \beta; s) = \frac{U \cdot \mathrm{CF}^-(\Sigma, \alpha, \beta; s)}{\mathrm{CF}^-(\Sigma, \alpha, \beta; s)}.$$

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:

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

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 2.2.3. *There are two long exact sequences*

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

where s is a Spin^c -structure of Y .

The explicit description illustrated above can be derived directly from the definition of the complexes. In this thesis we will mainly focus on the $\widehat{\mathrm{HF}}$ -theory so we leave this to the interested reader (see also [40]).

2.3 Topological Invariance

Given two Heegaard diagrams (Σ, α, β) and $(\Sigma', \alpha', \beta')$ of a manifold Y , they 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. The second reason to mention is our usage of the isomorphisms we will construct later in this thesis. We will deal with the \widehat{HF} -case and point the reader to [40] 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 discs 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.

2.3.1 Stabilizations/Destabilizations

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




Figure 2.3: An admissible Heegaard diagram for $\mathbb{S}^2 \times \mathbb{S}^1 \# \mathbb{S}^2 \times \mathbb{S}^1$.

Example 2.3.1. We fix admissible Heegaard diagrams $(T_i^2, \alpha_i, \beta_i)$ $i = 1, 2$ for $\mathbb{S}^2 \times \mathbb{S}^1$ as in Example 2.1.2. 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

$$\mathcal{J}_s^i|_D = \text{sym}(j_i).$$

Figure 2.3 pictures the Heegaard diagram we get for the connected sum. Denote by T a small connected sum tube inside $\Sigma = T_1^2 \# T_2^2$. By construction the induced almost complex structure equals

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

All intersection points belong to the same Spin^c -structure s_0 . For suitable Spin^c -structures s_1, s_2 on $\mathbb{S}^2 \times \mathbb{S}^1$ we have that $s_0 = s_1 \# s_2$ and

$$\widehat{\text{CF}}(\Sigma, \alpha, \beta, s_1 \# s_2) = \mathbb{Z} \otimes \{(x_i, y_j) \mid i, j \in \{1, 2\}\} \cong \widehat{\text{CF}}(T_1^2, s_1) \otimes \widehat{\text{CF}}(T_2^2, s_2).$$

The condition $n_z = 0$ implies that for every holomorphic disc $\phi: \mathbb{D}^2 \rightarrow \text{Sym}^g(\Sigma)$ the low-dimensional model (cf. §2.1) $\widehat{\phi}: \widehat{\mathbb{D}} \rightarrow \Sigma$ stays away from the tube T . Consequently we can split up $\widehat{\mathbb{D}}$ into

$$\widehat{\mathbb{D}} = \widehat{\mathbb{D}}_1 \sqcup \widehat{\mathbb{D}}_2,$$

where $\widehat{\mathbb{D}}_i$ are the components containing the preimage $(\widehat{\phi})^{-1}(T_i^2 \setminus D)$. Restriction to these components determines maps $\widehat{\phi}_i: \widehat{\mathbb{D}}_i \rightarrow T_i^2$ inducing Whitney discs ϕ_i in the symmetric product $\text{Sym}^1(T^2)$. Thus, the moduli spaces split:

$$\begin{aligned} \mathcal{M}_{(\mathcal{J}^1 \# \mathcal{J}^2)_s}((x_i, y_k), (x_j, y_l))_{n_z=0} &\xrightarrow{\cong} \mathcal{M}_{\mathcal{J}_s^1}(x_i, x_j)_{n_z=0} \times \mathcal{M}_{\mathcal{J}_s^2}(y_k, y_l)_{n_z=0} \\ \phi &\longmapsto (\phi_1, \phi_2). \end{aligned}$$

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 $a_i \in \widehat{\text{CF}}(T_i^2, s_i)$, $i = 1, 2$ 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

$$\widehat{\text{HF}}(\mathbb{S}^2 \times \mathbb{S}^1 \# \mathbb{S}^2 \times \mathbb{S}^1, s_1 \# s_2; o_1 \otimes o_2) \cong \widehat{\text{HF}}(\mathbb{S}^2 \times \mathbb{S}^1, s_1; o_1) \otimes \widehat{\text{HF}}(\mathbb{S}^2 \times \mathbb{S}^1, s_2; o_2).$$

The same line of arguments shows the general statement.

Theorem 2.3.1 (see [39]). *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.*

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

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 $\widehat{\text{HF}}(\mathbb{S}^3) = \mathbb{Z}$. By the classification of finitely generated abelian groups and the behavior of the tensor product, invariance follows.

2.3.2 Independence of the Choice of Almost Complex Structures

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

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

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

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

where $H(x, y, 0) \subset \pi_2(x, y)$ are the homotopy classes with expected dimension $\mu = 0$ and intersection number $n_z = 0$. There is an energy bound for all holomorphic Whitney discs which is independent of the particular Whitney disc or its homotopy class (see [40]). Thus, the moduli spaces are Gromov-compact manifolds, i.e. can be compactified by adding solutions coming from broken discs, bubbling of spheres and boundary degenerations (cf. §2.1.3). Since we stuck to the $\widehat{\text{HF}}$ -theory we impose the condition $n_z = 0$ which circumvents bubbling of spheres and boundary degenerations (see §2.1.3).

To check that $\widehat{\Phi}$ is a chain map, we compute

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

The coefficient $c(x, z)$ is given by

$$\sum_{y,I} (\#\mathcal{M}_{\mathcal{J}_{s,t},\phi} \cdot \#\widehat{\mathcal{M}}_{\mathcal{J}_{s,1},\psi} - \#\widehat{\mathcal{M}}_{\mathcal{J}_{s,0},\tilde{\psi}} \cdot \#\mathcal{M}_{\mathcal{J}_{s,t},\tilde{\phi}}), \quad (2.3.1)$$

where I consists of pairs

$$(\phi, \tilde{\phi}) \in H(x, y, 0) \times H(y, z, 0) \text{ and } (\psi, \tilde{\psi}) \in H(x, y, 1) \times H(y, z, 1).$$

Looking at the ends of the moduli spaces $\mathcal{M}_{J_{s,t}}(\eta)$ for an $\eta \in H(x, z, 1)$, the gluing construction (cf. §2.1.3) together with the compactification argument mentioned earlier provides the following ends:

$$\left(\bigsqcup_{\eta=\psi*\phi} (\mathcal{M}_{J_{s,t}}(\phi) \times \widehat{\mathcal{M}}_{J_{s,1}}(\psi)) \right) \sqcup \left(\bigsqcup_{\eta=\tilde{\psi}*\tilde{\phi}} (\widehat{\mathcal{M}}_{J_{s,0}}(\tilde{\psi}) \times \mathcal{M}_{J_{s,t}}(\tilde{\phi})) \right), \quad (2.3.2)$$

where the expected dimensions of ϕ and $\tilde{\phi}$ are 1 and of ψ and $\tilde{\psi}$ they are 0. A signed count of (2.3.2) precisely reproduces (2.3.1) and hence $c(x, z) = 0$ – 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:

$$o_{s,t}(\phi) \wedge o_1(\psi) = -o_0(\tilde{\psi}) \wedge o_{s,t}(\tilde{\phi}).$$

For an arbitrary coherent orientation system $o_{s,t}$ we get an identification of orientation systems, $\xi_{o_{s,t}}$ say, such that Φ is a chain map between

$$(\widehat{\text{CF}}(\Sigma, \alpha, \beta, z), \widehat{\partial}_{\mathcal{J}_{s,0}}^o) \longrightarrow (\widehat{\text{CF}}(\Sigma, \alpha, \beta, z), \widehat{\partial}_{\mathcal{J}_{s,1}}^{\xi_{o_{s,t}}}).$$

Observe that we can choose the coherent system $o_{s,t}$ arbitrarily. This will only affect the identification $\xi_{o_{s,t}}$.

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

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

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. §2.1.3 or see [40]). 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 2.3.2. *Let $\mathcal{J}_{(t_1, \dots, t_n)}$ be an n -parameter family of generic almost complex structures and ϕ a homotopy class of Whitney discs with expected dimension $\mu(\phi)$. Then \mathcal{M} , defined as the union of $\mathcal{M}_{\mathcal{J}_{(t_1, \dots, t_n)}, \phi}$ over all $\mathcal{J}_{(t_1, \dots, t_n)}$ in the family, is a manifold of dimension $\mu(\phi) + n$.*

There are two types of boundary components: the one type of boundary component coming from variations of the Whitney disc 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

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

where $H(x, y, -1) \subset \pi_2(x, y)$ are the homotopy classes ϕ with $n_z(\phi) = 0$ and expected dimension $\mu(\phi) = -1$. According to Theorem 2.3.2, 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

$$\widehat{\Phi}_{\mathcal{J}_{s,t}} \circ \widehat{\Phi}_{\mathcal{J}_{s,1-t}} - \text{id} - (\widehat{\partial}_{\mathcal{J}_{s,0}} \circ \widehat{H}_{\mathcal{J}_{s,t}(\tau)} + \widehat{H}_{\mathcal{J}_{s,t}(\tau)} \circ \widehat{\partial}_{\mathcal{J}_{s,1}}) = 0 \quad (2.3.3)$$

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

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

coming from variations of the Whitney disc, 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 (2.3.3). Thus, the Floer homology is independent of the choice of (j, η, V) -nearly symmetric path. Variations of η and V just change the contractible neighborhood \mathcal{U} around $\xi_{sym}^g(j)$ containing the admissible almost complex structures. So, the theory is independent of these choices, too. A j' -nearly symmetric path can be approximated by a j -symmetric path given that j' is close to j . The set of complex structures on a surface Σ is connected, so step by step one can move from a j -symmetric path to any j' -symmetric path.

2.3.3 Isotopy Invariance

Every isotopy of an attaching circle can be divided into two classes: creation/annihilation 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}_α in the symmetric product. Compactness of the \mathbb{T}_α 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

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

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 $\tilde{\mathcal{J}}_s$ are isomorphic. Hence

$$H_*(\widehat{\text{CF}}(\Sigma, \alpha, \beta), \widehat{\partial}_z^{\mathcal{J}_s}) = H_*(\widehat{\text{CF}}(\Sigma, \alpha', \beta), \widehat{\partial}_z^{\tilde{\mathcal{J}}_s}) = H_*(\widehat{\text{CF}}(\Sigma, \alpha', \beta), \widehat{\partial}_z^{\mathcal{J}_s}), \quad (2.3.4)$$

where the last equality follows from the considerations in §2.3.2. This chain of equalities shows that the isotopies discussed 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 §2.3.2 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 §2.3.2 for the right equality in (2.3.4) 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 §2.3.2.

Define $\pi_2^t(x, y)$ as the set of Whitney discs ϕ with dynamic boundary conditions in the following sense:

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

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

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

where $H_t(x, y, 0) \subset \pi_2^t(x, y)$ are the homotopy classes with expected dimension $\mu = 0$ and $n_z = 0$. Using the low-dimensional model introduced in §2.1, Ozvath and Szabo prove the following property.

Theorem 2.3.3 (see [40], §7.3). *There exists a t -independent energy bound for holomorphic Whitney discs independent of its homotopy class.*

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

Theorem 2.3.4. *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 §2.3.2. We leave the proof to the interested reader.

2.3.4 Handle slide Invariance

The Pair-of-Pants Product

In this paragraph 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 §2.1 before continuing.

Definition 2.3.5. A set of data $(\Sigma, \alpha, \beta, \gamma)$, where Σ is a surface of genus g and α, β, γ three sets of attaching circles, is called a **Heegaard triple diagram**.

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

$$\widehat{f}_{\alpha\beta\gamma}: \widehat{\text{CF}}(\Sigma, \alpha, \beta, z) \otimes \widehat{\text{CF}}(\Sigma, \beta, \gamma, z) \longrightarrow \widehat{\text{CF}}(\Sigma, \alpha, \gamma, z)$$

by counting holomorphic triangles with suitable boundary conditions: A **Whitney triangle** is a map $\phi: \Delta \longrightarrow \text{Sym}^g(\Sigma)$ with boundary conditions as illustrated in Figure 2.4. We call the respective boundary segments its α -, β - and γ -boundary. The boundary points, as should be clear from the picture, are $x \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, $w \in \mathbb{T}_\alpha \cap \mathbb{T}_\gamma$ and $y \in \mathbb{T}_\beta \cap \mathbb{T}_\gamma$. The set of homotopy classes of Whitney discs connecting x, w and y is denoted by $\pi_2(x, y, w)$.

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

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




Figure 2.4: A Whitney triangle and its boundary conditions.

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

$$0 \longrightarrow \pi_2(\text{Sym}^g(\Sigma)) \longrightarrow \pi_2(x, y, w) \longrightarrow \ker(n_z) \longrightarrow 0, \quad (2.3.5)$$

where n_z provides a splitting for the sequence. We define

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

where U_α , U_β and U_γ are the handlebodies determined by the 2–handles associated to the attaching circles α , β and γ , and e_α , e_β and e_γ are the edges of the triangle Δ . The manifold $X_{\alpha\beta\gamma}$ is 4-dimensional with boundary

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

Lemma 2.3.6. *The kernel of n_z equals $H_2(X_{\alpha\beta\gamma}; \mathbb{Z})$*

Combining (2.3.5) with Lemma 2.3.6 we get an exact sequence

$$0 \longrightarrow \pi_2(\text{Sym}^g(\Sigma)) \longrightarrow \pi_2(x, y, w) \xrightarrow{\mathcal{H}} H_2(X_{\alpha\beta\gamma}; \mathbb{Z}) \longrightarrow 0, \quad (2.3.6)$$

where \mathcal{H} is defined similarly as for discs (cf. §2.1.2). Of course there is a low-dimensional model for triangles and the discussion we have done for discs carries over verbatim for triangles. The condition $n_z = 0$ makes the product $f_{\alpha\beta\gamma}$ well-defined in case $H_2(X_{\alpha\beta\gamma}; \mathbb{Z})$ is trivial. Analogous to our discussion for Whitney discs and the differential, we have to include a condition controlling the periodic triangles, i.e. the triangles associated to elements in $H_2(X_{\alpha\beta\gamma}; \mathbb{Z})$. 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 2.3.7. A pointed triple diagram $(\Sigma, \alpha, \beta, \gamma, z)$ 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 discs to triangles. One can show that for given $j, k \in \mathbb{Z}$ there exist just a finite number of Whitney triangles $\phi \in \pi_2(x, y, w)$ with $\mu(\phi) = j$, $n_z(\phi) = k$ and $\mathcal{D}(\phi) \geq 0$.

For a given homotopy class $\psi \in \pi_2(x, y, w)$ with $\mu(\psi) = 1$ we compute the ends by shrinking a properly embedded arc to a point (see the description of convergence in §2.1.3). There are three different ways to do this in a triangle. Each time we get a concatenation of a disc with a triangle. By boundary orientations we see that each of these boundary components contributes to one of the terms in the following sum

$$\widehat{f}_{\alpha\beta\gamma} \circ (\widehat{\partial}_{\alpha\beta}(x) \otimes y) + \widehat{f}_{\alpha\beta\gamma} \circ (x \otimes \widehat{\partial}_{\beta\gamma}(y)) - \widehat{\partial}_{\alpha\gamma} \circ \widehat{f}_{\alpha\beta\gamma}(x \otimes y). \quad (2.3.7)$$

Conversely, the coefficient at any of these terms is given by a product of signed counts of moduli spaces of discs and moduli spaces of triangles and hence – by gluing – comes from one of these contributions. The sum in (2.3.7) vanishes, showing that $\widehat{f}_{\alpha\beta\gamma}$ descends to a pairing $\widehat{f}_{\alpha\beta\gamma}^*$ between the Floer homologies.

Holomorphic rectangles

Recall that the set of biholomorphisms of the unit disc 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 biholomorphisms of the unit disc 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 2.3.8. Given another set of attaching circles δ defining a map $\widehat{f}_{\alpha\gamma\delta}$, the following equality holds:

$$\widehat{f}_{\alpha\beta\gamma}^*(\widehat{f}_{\alpha\gamma\delta}^*(\cdot \otimes \cdot) \otimes \cdot) - \widehat{f}_{\alpha\beta\delta}^*(\cdot \otimes \widehat{f}_{\beta\gamma\delta}^*(\cdot \otimes \cdot)) = 0. \quad (2.3.8)$$

This property is called **associativity**.




Figure 2.5: Ends of the moduli space of holomorphic rectangles.

If we count holomorphic Whitney rectangles with boundary conditions in α , β , γ and δ and with $\mu = 1$ (see Definition 2.1.6) the ends of the associated moduli space will look like pictured in Figure 2.5. 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 disc by variation of the rectangle. By Figure 2.5 an appropriate count of holomorphic rectangles will be a natural candidate for a chain homotopy proving equation (2.3.8). Define a pairing

$$H: \widehat{\text{CF}}(\Sigma, \alpha, \beta, z) \otimes \widehat{\text{CF}}(\Sigma, \beta, \gamma, z) \otimes \widehat{\text{CF}}(\Sigma, \gamma, \delta, z) \longrightarrow \widehat{\text{CF}}(\Sigma, \alpha, \delta, z)$$

by counting holomorphic Whitney rectangles with boundary components as indicated in Figure 2.6 and $\mu = 0$. By counting ends of the moduli space of holomorphic




Figure 2.6: The boundary conditions of rectangles for the definition of H .

rectangles with $\mu = 1$ we have six contributing ends. These ends are pictured in Figure 2.5. The four ends coming from breaking contribute to

$$\widehat{\partial} \circ H(\cdot \otimes \cdot \otimes \cdot) + H \circ \widehat{\partial}(\cdot \otimes \cdot \otimes \cdot). \quad (2.3.9)$$

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

$$\widehat{f}_{\alpha\beta\gamma}(\widehat{f}_{\alpha\gamma\delta}(\cdot \otimes \cdot) \otimes \cdot) - \widehat{f}_{\alpha\beta\delta}(\cdot \otimes \widehat{f}_{\beta\gamma\delta}(\cdot \otimes \cdot)). \quad (2.3.10)$$

We see that the sum of (2.3.9) and (2.3.10) vanishes, showing that H is a chain homotopy proving associativity.

Special Case – Handle Slides

Handle slides provide special Heegaard triple diagrams. Let $(\Sigma, \alpha, \beta, z)$ be an admissible pointed Heegaard diagram and define $(\Sigma, \alpha, \gamma, z)$ 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 $Y_{\beta\gamma}$ 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 2.1.2 and Theorem 2.3.1,

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

where the last identification is done using the $\bigwedge^*(H_1/\text{Tor})$ -module structure (see [40]). 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 x_i^+ and x_i^- . Which one is denoted how is determined by the following criterion: there is a disc-shaped domain connecting x_i^+ with x_i^- with boundary in β_i and γ_i . The point

$$x^+ = \{x_1^+, \dots, x_g^+\}$$

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 [40].

Plugging in the generator we define a map

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

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 $\widehat{F}_{\alpha\gamma\delta}$ be the associated map. Then the associativity given in (2.3.8) translates into

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

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. The proof is carried out as a model for proofs that will be discussed in the remainder of this thesis.

Lemma 2.3.9. *Given the map $\widehat{f}_{\alpha\gamma\delta}$, we have*

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

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

Proof. The complement of the β -circles in Σ is a sphere with holes. We have a precise description of how the sets γ and δ look like relative to β . 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 $(\Sigma, \beta, \gamma, \delta)$ will look like given in Figure 2.7. 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 $\{\widehat{\Theta}_{\beta\gamma}, \widehat{\Theta}_{\gamma\delta}\}$, 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 disc 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 $\widehat{F}_{\beta\gamma\delta} (\widehat{\Theta}_{\beta\gamma} \otimes \widehat{\Theta}_{\gamma\delta})$. 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 $\widehat{\phi}: \widehat{\mathbb{D}} \rightarrow \Sigma$ from a surface with boundary $\widehat{\mathbb{D}}$ whose associated domain equals the sum of \mathcal{D}_1 -domains. The map $\widehat{\phi}$ is a biholomorphism and thus $\widehat{\mathbb{D}}$ is a disjoint union of triangles. The




Figure 2.7: The Heegaard surface cut open along the β -curves.

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 $\widehat{\mathbb{D}} \rightarrow \mathbb{D}^2$. Since $\widehat{\mathbb{D}}$ is a union of g discs, this covering is unique, too (up to equivalence) and thus the associated moduli space is a one-point space. \square

Lemma 2.3.9 and (2.3.4) combine to give the composition law

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

We call a holomorphic triangle **small** if it is supported within the thin strips of isotopy between β and δ .

Lemma 2.3.10 (see [40], Lemma 9.10). *Let $F: A \rightarrow B$ 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 $l(x) < F_0(x)$. 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 **area filtration** (cf. [40]), which is indeed bounded from below. In this situation the top-dimensional generator $\widehat{\Theta}_{\beta\delta}$ is generated by a single intersection point $x^+ \in \mathbb{T}_\beta \cap \mathbb{T}_\delta$. The map $\widehat{F}_{\alpha\beta\delta}$ is induced by

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

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 β and δ . The map f_0 is filtration preserving and l , if the δ -curves are close enough to the β -curves, strictly decreasing. By Lemma 2.3.10 the map $\widehat{F}_{\alpha\beta\delta}$ is an isomorphism between the associated Heegaard Floer homologies.

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

Theorem 2.3.11. *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.

2.4 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 $\overline{Y \setminus \nu K}$ relative to its boundary. Note that in the literature the knot Floer variants are **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. But the generalization comes at a price. 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 2.4.1. A Heegaard diagram (Σ, α, β) 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 disc of the 2-handle, represented by β_1 , once and is disjoint from the others (after possibly isotoping the knot K).

Lemma 2.4.2. *Every pair (Y, K) admits a Heegaard diagram subordinate to K .*

Proof. By surgery theory (see [18], p. 104) we know that there is a handle decomposition of $Y \setminus \nu K$, i.e.

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

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

$$Y \cong h^3 \cup_{\partial} h_1^{2*} \cup_{\partial} (T^2 \times I) \cup_{\partial} h_2^1 \cup_{\partial} \dots \cup_{\partial} h_g^1 \cup_{\partial} h_1^2 \cup_{\partial} \dots \cup_{\partial} h_g^2 \cup_{\partial} h^3. \quad (2.4.1)$$

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 (Σ, α, β) where $\alpha = \alpha_1^* \cup \{\alpha_2, \dots, \alpha_g\}$ are the co-cores of the 1-handles and $\beta = \{\beta_1, \dots, \beta_g\}$ are the attaching circles of the 2-handles. \square

Having fixed such a Heegaard diagram (Σ, α, β) 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 $(\Sigma, \alpha, \beta, w, z)$. Conversely, a doubly-pointed Heegaard diagram uniquely determines a topological knot class: Connect z with w in the complement of the attaching circles α and $\beta \setminus \beta_1$ with an arc δ that crosses β_1 once. Connect w with z in the complement of β using an arc γ . The union $\delta \cup \gamma$ represents the knot class K represents. The orientation on K is given by orienting δ such that $\partial \delta = z - w$. If we use a different path $\tilde{\gamma}$ in the complement of β , we observe that $\tilde{\gamma}$ is isotopic to γ (in Y): Since $\Sigma \setminus \beta$ is a sphere with holes an isotopy can move γ across the holes by doing handle slides. Isotope the knot along the core discs 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 $\widehat{\text{CFK}}(Y, K)$ is the free \mathbb{Z}_2 -module (or \mathbb{Z} -module) generated by the intersections $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$. The boundary operator $\widehat{\partial}^w$, for $x \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, is defined by

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

where $H(x, y, 1) \subset \pi_2(x, y)$ are the homotopy classes with $\mu = 1$ and $n_z = n_w = 0$. We denote by $\widehat{\text{HFK}}(Y, K)$ the associated homology theory $H_*(\widehat{\text{CFK}}(Y, K), \widehat{\partial}^w)$. 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 2.4.3. ([38]) *Let $(\Sigma, \alpha, \beta, z, w)$ and $(\Sigma', \alpha', \beta', z', w')$ 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 β_2, \dots, β_g . These isotopies may not cross I .*
- (m_3) *Handle slides of β_1 over the β_2, \dots, β_g and isotopies.*
- (m_4) *Stabilizations/destabilizations.*

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

Proof. By Theorem 4.2.12 of [18] 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 (2.4.1). 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_i^{2*} over the h_i^2 . This corresponds to move (m_3). \square

Proposition 2.4.4. *Let $K \subset Y$ be an arbitrary knot. The knot Floer homology group $\widehat{\text{HF}}\text{K}(Y, K)$ is a topological invariant of the knot type of K in Y . These homology groups split with respect to $\text{Spin}^c(Y)$.*

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 §2.3. Each of these maps is defined by counting holomorphic discs with punctures, whose properties are shown by defining maps by counting holomorphic discs with punctures.

Isotopies/Almost Complex Structure. Denote by 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 discs with punctures (for a precise definition look into §2.3.2 and §2.3.3) 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 2.1.6). 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 J , respecting w means, that we perturb J through nearly symmetric almost complex structures such that V (cf. Definition 2.1.19) also contains $\{w\} \times \text{Sym}^{g-1}(\Sigma)$. 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}^τ , $\tau \in [0, 1]$, and combining them to

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

The spaces \mathcal{M}_{-1}^τ are defined like done in §2.3.2 and §2.3.3. 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. §2.3.2 and §2.3.3). 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

$$\widehat{f}_{\alpha\beta\gamma}: \widehat{\text{CFK}}(\Sigma, \alpha, \beta, w, z) \otimes \widehat{\text{CFK}}(\Sigma, \beta, \gamma, w, z) \longrightarrow \widehat{\text{CFK}}(\Sigma, \alpha, \gamma, w, z)$$

induced by a doubly-pointed Heegaard triple diagram $(\Sigma, \alpha, \beta, \gamma, w, z)$. We have to see, that in case the triple is induced by a handle slide, the knot Floer homology $\widehat{\text{HF}}(\Sigma, \beta, \gamma, w, z)$ carries a top-dimensional generator $\widehat{\Theta}_{\beta\gamma}$, 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 $\Sigma \setminus \{\beta \cup \gamma\}$. Hence, we have an identification

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

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. 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.

Admissibility

The admissibility condition given in Definition 2.1.17 suffices to give a well-defined theory. However, since we have an additional point w in play, we can relax the admissibility condition.

Definition 2.4.5. We call a doubly-pointed Heegaard diagram $(\Sigma, \alpha, \beta, w, z)$ **extremely weakly admissible** for the Spin^c -structure s if for every non-trivial periodic domain, with $n_w = 0$ and $\langle c_1(s), \mathcal{H}(\mathcal{D}) \rangle = 0$, the domain has both positive and negative coefficients.

With a straightforward adaptation of the proof of well-definedness in the case of $\widehat{\partial}_z$ we get the following result (see [40], Lemma 4.17, cf. Definition 2.1.17 and cf. proof of Theorem 2.1.3).

Theorem 2.4.6. *Let $(\Sigma, \alpha, \beta, w, z)$ be an extremely weakly admissible Heegaard diagram. Then $\widehat{\partial}^w$ is well-defined and a differential.* \square

Note that Ozsváth and Szabó impose weak admissibility of the Heegaard diagram $(\Sigma, \alpha, \beta, z)$. The introduction of our relaxed condition is done since we will find setups in this thesis where it is convenient to relax the admissibility condition like introduced.

Other knot Floer homologies

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

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

where $H(x, y, 1) \subset \pi_2(x, y)$ 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 2.4.5. Using this construction, and continuing like in §2.2, we define variants we denote by HFK^∞ and HFK^+ . The groups are naturally connected by exact sequences analogous to those presented in Lemma 2.2.3.

2.4.1 Refinements

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

$$\text{Spin}^c(Y_0(K)) = \text{Spin}^c(Y) \times \mathbb{Z}. \quad (2.4.2)$$

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 (2.4.2). Given a Spin^c -structure t on $Y_0(K)$, we associate to it the pair (s, k) , where s is the restriction of 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 t and restrict this vector field to $Y \setminus \nu K$. Observe that we may regard $Y \setminus \nu K$ 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 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 t we can associate a link L_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_t can be written as a sum

$$L_t = k \cdot \mu_0 + \dots,$$

and thus we can compute k with

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

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 disc in $Y_0(K)$.

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

$$\widehat{\text{CFK}}(Y, K, s) = \bigoplus_{t \in H_s} \widehat{\text{CFK}}(Y, K, t),$$

where $H_s \subset \text{Spin}^c(Y_0(K))$ are the elements extending $s \in \text{Spin}^c(Y)$. We have to show that $\widehat{\partial}^w$ preserves this splitting. We point the interested reader to [38].

2.5 Maps Induced By Cobordisms

The pairing introduced in §2.3.4 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. Proposition 4.2.13 in [18]). 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 [44].

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 $(\Sigma, \alpha, \beta, \gamma)$ where $\gamma_i, i \geq 2$, 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 2.5.1. *The cobordism $X_{\alpha\beta\gamma} \cup_{\partial} (\#^{g-1}D^3 \times \mathbb{S}^1)$ is diffeomorphic to the cobordism W_K given by the framed surgery along K .*

We define

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

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

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

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 2.3.9 to show a composition law reading

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

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 $L = K_1 \sqcup \cdots \sqcup K_m$, observe that we can obviously define a map

$$\widehat{F}_L: \widehat{\text{HF}}(Y) \longrightarrow \widehat{\text{HF}}(Y_L),$$

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 2.3.9, will suffice to show that the map \widehat{F}_L associated to multiple attachments is a composition

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

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.

2.6 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 talk about framings as an integer $n \in \mathbb{Z}$. This identification will be done whenever it makes sense.

There is a long exact sequence

$$\cdots \xrightarrow{\partial_*} \widehat{\text{HF}}(Y) \xrightarrow{\widehat{F}_1} \widehat{\text{HF}}(Y_K^n) \xrightarrow{\widehat{F}_2} \widehat{\text{HF}}(Y_K^{n+\mu}) \xrightarrow{\partial_*} \cdots, \quad (2.6.1)$$

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 [39])

$$0 \longrightarrow \widehat{\text{CF}}(Y) \xrightarrow{\widetilde{F}_1} \widehat{\text{CF}}(Y_K^n) \xrightarrow{\widehat{F}_2} \widehat{\text{CF}}(Y_K^{n+\mu}) \longrightarrow 0. \quad (2.6.2)$$

The map ∂_* in (2.6.1) denotes the induced coboundary. This enables us to prove the existence of the surgery exact triangle.

Theorem 2.6.1. *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*

$$\begin{array}{ccc} \widehat{\text{HF}}(Y) & \xrightarrow{\widehat{F}_1} & \widehat{\text{HF}}(Y_K^n) \\ & \searrow \widehat{F}_3 & \swarrow \widehat{F}_2 \\ & \widehat{\text{HF}}(Y_K^{n+\mu}) & \end{array}$$

which is called **surgery exact triangle**.




Figure 2.8: The topological situation in the exact triangle.

Proof. Observe that the topological situation is very symmetric. The long exact sequence (2.6.1) corresponds to the topological situation pictured in Figure 2.8. Each arrow in Figure 2.8 corresponds to an exact sequence of type (2.6.1). With the identifications given, we can concatenate the three sequences to give the surgery exact sequence of Theorem 2.6.1. \square

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 paragraph. 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

$$\widehat{\text{CF}}(Y) \xrightarrow{\widehat{f}_1} \widehat{\text{CF}}(Y_K^n) \xrightarrow{\widehat{f}_2} \widehat{\text{CF}}(Y_K^{n+\mu}) \quad (2.6.3)$$

is null-chain homotopic. Let $(\Sigma, \alpha, \beta, z)$ 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_K^n can be described by $(\Sigma, \alpha, \gamma, z)$ where $\gamma_i, i \geq 2$ are isotopic push-offs of the β_i such that β_i and γ_i meet in two cancelling intersections transversely. 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 δ where $\delta_i, i \geq 2$ 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, (Σ, α, δ) is a Heegaard diagram of $Y_K^{n+\mu}$. By associativity (2.3.8), the composition $\widehat{f}_2 \circ \widehat{f}_1$ is chain homotopic to

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

where the chain homotopy H is given by counting holomorphic rectangles with suitable boundary conditions (cf. §2.3.4). To compute $\widehat{f}_{\beta\gamma\delta}(\widehat{\Theta}_{\beta\gamma} \otimes \widehat{\Theta}_{\gamma\delta})$ we use a model calculation. Figure 2.9 illustrates the Heegaard triple diagram.




Figure 2.9: Heegaard triple diagram for computation of $\widehat{f}_{\beta\gamma\delta}(\widehat{\Theta}_{\beta\gamma} \otimes \widehat{\Theta}_{\gamma\delta})$.

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 2.3.9). Hence, in \mathbb{Z}_2 -coefficients

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

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 2.8, makes it possible to carry over the proof given here.

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

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

denote the null-homotopy of $\widehat{f}_2 \circ \widehat{f}_1$ (cf. §2.3.4). Define the chain complex $A_{\widehat{f}_1, \widehat{f}_2}$ to be given by the module $A = \widehat{\text{CF}}(Y) \oplus \widehat{\text{CF}}(Y_K^n) \oplus \widehat{\text{CF}}(Y_K^{n+\mu})$ 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 2.6.2. *The sequence (2.6.3) is exact on the homological level at $\widehat{\text{CF}}(Y_K^n)$ if $H_*(A_{\widehat{f}_1, \widehat{f}_2}) = 0$.*

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

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

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

Definition 2.6.3. For a chain map $f: A \longrightarrow B$ 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

$$\partial_f = \begin{pmatrix} \partial_A & 0 \\ f & \partial_B \end{pmatrix}$$

The mapping cone is a chain complex (cf. Lemma 3.1.1).

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

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

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

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

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 2.6.4 ([42], Lemma 4.2). *Let $\{A_i\}_{i \in \mathbb{Z}}$ be a collection of modules and let*

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

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 $H_i: A_i \longrightarrow A_{i+2}$. The maps

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

should induce isomorphisms between the associated homologies. Then the maps $(H_i, f_{i+1}): M(f_i) \longrightarrow A_{i+2}$ induce isomorphisms on the homological level.

If we can show that the sequence

$$\dots \xrightarrow{\widehat{f}_3} \widehat{\text{CF}}(Y) \xrightarrow{\widehat{f}_1} \widehat{\text{CF}}(Y_K^n) \xrightarrow{\widehat{f}_2} \widehat{\text{CF}}(Y_K^{n+\mu}) \xrightarrow{\widehat{f}_3} \dots$$

satisfies the assumptions of Lemma 2.6.4, 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 2.6.2, we conclude that $\text{im}(\widehat{F}_i) = \ker(\widehat{F}_{i+1})$. Hence, Theorem 2.6.1 follows.

2.7 The Contact Element and $\widehat{\mathcal{L}}$

2.7.1 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

$$\alpha \wedge d\alpha \neq 0. \quad (2.7.1)$$

Those 1-forms satisfying (2.7.1) 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 $\lambda: Y \longrightarrow \mathbb{R}^+$, we can change the contact form to $\lambda\alpha$ without affecting the contact condition (2.7.1):

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

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

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

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 $T_p Y \setminus \xi_p$. Thus, R_α is a section of TY/ξ .

Definition 2.7.1. Two contact manifolds (Y, ξ) and (Y', ξ') are called **contactomorphic** if there is a diffeomorphism $\phi: Y \rightarrow Y'$ 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 standard model for these objects.

Definition 2.7.2. The pair $(\mathbb{R}^3, \xi_{std})$, where ξ_{std} is the contact structure given by the kernel of the 1-form $dz - ydx$, 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 2.7.3 (Gray Stability, cf. [16]). *Each smooth homotopy of contact structures $(\xi_t)_{t \in [0,1]}$ 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 induced 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 2.7.4. 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 [8]. 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 2.7.5 (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 connect an arbitrary contact manifold with an arbitrary overtwisted contact manifold. For a definition of overtwistedness we point the reader to [16]. 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 (see Figure 3.19).

2.7.2 Open Books

For a detailed treatment of open books we point the reader to [9].

Definition 2.7.6. An **open book** on a closed, oriented 3-manifold Y is a pair (B, π) defining a fibration

$$P \hookrightarrow Y \setminus B \xrightarrow{\pi} \mathbb{S}^1,$$

where P is an oriented surface with boundary $\partial P = B$. For every component B_i of B there is a neighborhood $\iota: D^2 \times \mathbb{S}^1 \hookrightarrow \nu B_i \subset Y$ such that the core $C = \{0\} \times \mathbb{S}^1$ is mapped onto B_i under ι and π commutes with the projection $(D^2 \times \mathbb{S}^1) \setminus C \longrightarrow \mathbb{S}^1$ 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 $\phi: P \longrightarrow P$ 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

$$(P \times [0, 1]) / (p, 1) \sim (\phi(p), 0) \tag{2.7.2}$$

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 full torus $D^2 \times \mathbb{S}^1$: we glue a meridional disc $D^2 \times \{\star\}$ 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

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

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 $\phi: P \longrightarrow P$ such that

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

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

Two open books are called equivalent if they are diffeomorphic. The set of equivalence classes of open books is denoted by $\text{OB}(Y)$. An abstract open book defines an open book up to diffeomorphism. With the construction given above we define a map

$$\Psi: \text{ABS}(Y) \longrightarrow \text{OB}(Y)$$

and its 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.

2.7.3 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 $H_1 \cup H_2$ where $H_i = \pi^{-1}([i/2, (i+1)/2])$, $i = 0, 1$. Any curve γ in Y running from H_1 to H_2 , when projected onto \mathbb{S}^1 , has to intersect $\{1/2, 1\}$ at some point. Thus, the curve γ is forced 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 2.7.8. A system $a = \{a_1, \dots, a_n\}$ of disjoint, properly embedded curves on P is called **cut system** if $P \setminus \{a_1, \dots, a_n\}$ is topologically a disc.

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

Interpreting the curve a_i as sitting on $P_{1/2}$ and \bar{a}_i as sitting inside P_1 , we can combine them to $\alpha_i = a_i \cup_{\partial} \bar{a}_i$, $i = 1, \dots, n$. The disc $a_i \times [0, 1/2]$ can be embedded into $P \times [0, 1]$ and by going over to the quotient (2.7.2) we obtain a disc in H_0 with boundary α_i . This means we can interpret the set $\{\alpha_1, \dots, \alpha_n\}$ as a set of attaching circles for the handlebody H_0 . 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 $P_1 \rightarrow P_0$. These two maps combine to a map $\partial H_0 \rightarrow \partial H_1$. Define b_i , $i = 1, \dots, n$, as small push-offs of the α_i that intersect these transversely. Then we define $\beta_i = b_i \cup \overline{\phi(b_i)}$, $i = 1, \dots, n$. Thus the following lemma is immediate.

Lemma 2.7.9. *The triple (Σ, α, β) 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_{γ}^{\pm} the positive/negative Dehn twist along γ .

Definition 2.7.10. The abstract open book $(P', D_{\gamma}^{\pm} \circ \phi)$ 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 2.7.11. *Stabilizations preserve the underlying 3-manifold, i.e. the manifolds $Y(P', \phi')$ and $Y(P, \phi)$ are isomorphic.*

A priori, it is not clear that stabilizations preserve the associated 3-manifold. A proof of this lemma can be found in [9]. But in the following we will discuss an alternative proof. Our proof uses a construction introduced by Lisca, Ozsvá, Stipsicz and Szabó (see [27], Alternative proof of Theorem 2.11).

Lemma 2.7.12 ([27]). *There is a cut system $\{a_1, \dots, a_n\}$ on (P, ϕ) that is disjoint from $\gamma \cap P$.*

Proof. Denote by γ' the arc $\gamma \cap P$. If $P \setminus \gamma'$ 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 $P \setminus \gamma'$ disconnects 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' . \square

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, \dots, a_{n+1} is a cut system of P' . Choose the b_i , $i = 1, \dots, n + 1$, as small isotopic push-offs of the a_i . Then, for $i = 1, \dots, n$, we have

$$\begin{aligned} \phi'(b_i) &= \phi \circ D_\gamma^\pm(b_i) = \phi(b_i) \\ \phi'(b_{n+1}) &= D_\gamma^\pm \circ \phi(b_{n+1}) = D_\gamma^\pm(b_{n+1}). \end{aligned}$$

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 2.7.11. 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

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

\square

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 in the remainder of this thesis. For a detailed treatment we point the reader to [9]. Every contact structure is supported by an open book decomposition.

Theorem 2.7.13 (cf. [9]). *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. §2.6). 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 [9].

Lemma 2.7.14 (cf. [27]). *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, ξ_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 ξ_L^\pm .*

2.7.4 The Contact Class

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 paragraph, 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 2.10: We push the b_i off the a_i by following with ∂b_i the positive boundary orientation of ∂P .




Figure 2.10: 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 x_i the unique intersection point between a_i and b_i . Define

$$EH(P, \phi, \{a_1, \dots, a_{2g}\}) = \{x_1, \dots, x_{2g}\}.$$

By construction of the Heegaard diagram EH is a cycle in the Heegaard Floer homology associated to the data $(-\Sigma, \alpha, \beta, z)$.

Lemma 2.7.15 (see [39]). *The Heegaard Floer cohomology $\widehat{HF}^*(Y)$ is isomorphic to $\widehat{HF}(-Y)$.*

The Heegaard diagram $(-\Sigma, \alpha, \beta)$ 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 discs at their α -boundary and β -boundary, i.e. we will be interested in Whitney discs in (Σ, β, α) . 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 disc emanating from x_i .

Theorem 2.7.16. *The class $EH(P, \phi, \{a_1, \dots, a_{2g}\})$ 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 $P \setminus \{a_1, \dots, a_{2g}\}$ 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 2.7.17. *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 2.7.18. Let a Heegaard diagram (Σ, α, β) and a homologically essential, simple, closed curve δ on Σ be given. The Heegaard diagram (Σ, α, β) is called **δ -adapted** if the following conditions hold.

1. It is induced by an open book and the pair α, β is induced by a cut system (cf. §2.7.3) for this open book.
2. The curve δ intersects β_1 once and does not intersect any other of the $\beta_i, i \geq 2$.

We can always find δ -adapted Heegaard diagrams. This is already stated in [20] and [27] but not proved. We wish to give a proof because this specific choice is crucial throughout this thesis

Lemma 2.7.19. *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, \dots, 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 disconnects 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 disconnect the page. If they disconnected, 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 $\Sigma \setminus \{a_1, \dots, a_n\}$ is a disc shows that every element in $H_1(P, \partial P)$ can be written as a linear combination of the curves a_1, \dots, 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 $\{a_1, \dots, a_n\}$ is a cut system, then $\{\phi(a_1), \dots, \phi(a_n)\}$ is a cut system, too. It suffices to show that there is a cut system $\{a_1, \dots, a_n\}$ such that δ intersects a_i once if and only if $i = 1$.




Figure 2.11: Possible choice of curve γ .

We start by taking a band sum of δ with a small arc γ as shown in Figure 2.11. 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 disc 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 $(P, \partial P)$

$$\begin{array}{ccccccccc} H_2(P) & \longrightarrow & H_2(P, \partial P) & \xrightarrow{\partial_*} & H_1(\partial P) & \longrightarrow & H_1(P) & \xrightarrow{\iota_*} & H_1(P, \partial P) & \longrightarrow & 0 \\ \parallel & & \cong & & \cong & & & & & & \\ 0 & \longrightarrow & \mathbb{Z}\langle [P] \rangle & \xrightarrow{\partial_*} & \mathbb{Z}\langle [\partial P] \rangle & \longrightarrow & H_1(P) & \xrightarrow{\iota_*} & H_1(P, \partial P) & \longrightarrow & 0 \end{array}$$

we see that ∂_* is surjective since $\partial_*[P] = [\partial P]$. Hence, exactness of the sequence implies that the inclusion $\iota: P \rightarrow (P, \partial P)$ 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. $\iota_*[\delta] = [a_2]$. 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 primitive, too. To see, that we can extend to a cut system such that δ is disjoint from a_3, \dots, 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, \dots, S_4 say, in this boundary. We extend a_2 to a cut system a_2, \dots, a_n of P' and arrange the boundary points of the curves a_3, \dots, a_n to be disjoint from S_1, \dots, S_4 . The set a_1, \dots, 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 2.7.20. *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 disc. 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,

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

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

Recall that a positive/negative Giroux stabilization of an open book (P, ϕ) is defined as the open book $(P', D_\gamma^\pm \circ \phi)$ 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 2.7.10). Using the proofs of Lemma 2.7.11 and Lemma 2.7.12, 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 $\alpha = \{\alpha_1, \dots, \alpha_n\}$ the associated attaching circles. We define a map

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

by assigning to $x \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ the element $\Phi(x) = (x, q)$ where q is the unique intersection point $\gamma \cap a_{n+1}$. This is an isomorphism by reasons similar to those given in Example 2.3.1.

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 2.7.21 ([41]). *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*

$$\widehat{F}_{-W}: \widehat{\text{HF}}(-Y) \longrightarrow \widehat{\text{HF}}(-Y')$$

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

Proof. Let an open book (P, ϕ) adapted to (Y, ξ, L) be given. By Lemma 2.7.14, a $(+1)$ -contact surgery acts on the monodromy as a composition with a negative Dehn

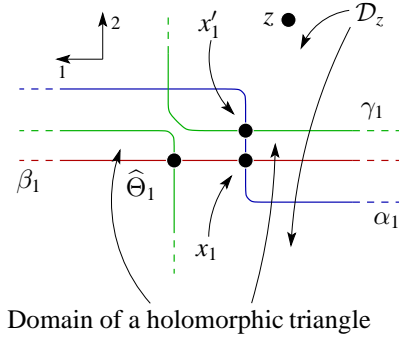


Figure 2.12: Significant part of the Heegaard triple diagram.

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 2.12. The contact element $c(Y, \xi)$ is represented by the point $\{x_1, \dots, x_n\}$. Obviously, there is only one domain which carries a holomorphic triangle. It is the small holomorphic triangle connecting x_1 and x'_1 (cf. §2.3.4). Thus, there is only one domain with positive coefficients, with $n_z = 0$, connecting the points $\{x_1, \dots, x_n\}$ with $\{x'_1, \dots, x'_n\}$. By considerations similar to those given at the end of the proof of Lemma 2.3.9, we see that the associated moduli space is a one-point space. Hence, the result follows. \square

2.7.5 The Invariant $\widehat{\mathcal{L}}$

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 Lisca, Ozsváth, Stipsicz and Szabó and was defined in [27]. 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. §2.7.3). Choose a cut system that induces an L -adapted Heegaard diagram (cf. §3.2.1, Definition 2.7.18 and Lemma 2.7.19). Figure 2.13 illustrates the positioning of a point w in the Heegaard diagram induced by the open book. Similar to the case of the contact element those intersection points $\alpha_i \cap \beta_i$ who sit on $P \times \{1/2\}$ determine one specific generator of $\widehat{\text{CF}}(-Y)$. This element may be interpreted as




Figure 2.13: Positioning of the point w depending on the knot orientation.

sitting in $\widehat{\text{CFK}}(-Y, L)$, and it is a cycle there, too. The induced element in the knot Floer homology is denoted by $\widehat{\mathcal{L}}(L)$.

Remark. 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.

Chapter 3

Dehn Twists in $\widehat{\text{HF}}$ Homology

3.1 Algebraic Preliminaries

We outline some algebraic tools used in the next sections. We present this material for the sake of completeness.

Lemma 3.1.1. *Suppose we are given two complexes (C, ∂_C) and (D, ∂_D) and a morphism $f: D \rightarrow C$ of complexes. Then $(C \oplus D, \partial^f)$ is a chain complex where $\partial^f := \partial_C + f - \partial_D$, i.e.*

$$\partial^f = \begin{pmatrix} \partial_C & f \\ 0 & -\partial_D \end{pmatrix}.$$

Proof. For $(p, q) \in C \oplus D$ we calculate

$$\begin{aligned} (\partial^f)^2(p, q) &= \partial^f(\partial_C p + f(q), -\partial_D q) \\ &= (\partial_C^2 p + \partial_C f(q) + f(-\partial_D q), \partial_D^2 q) \\ &= 0, \end{aligned}$$

where the last equality holds, since ∂_C and ∂_D are differentials and f is a chain map. \square

A nice, immediate consequence of this construction is the following Lemma.

Lemma 3.1.2. *There is a long exact sequence*

$$\dots \xrightarrow{-f_*} H_*(C, \partial_C) \xrightarrow{\Gamma_1} H_*(C \oplus D, \partial^f) \xrightarrow{\Gamma_2} H_*(D, -\partial_D) \xrightarrow{-f_*} \dots,$$

where f_* is the map in homology induced by f , and Γ_1 and Γ_2 are given as follows:

- Γ_1 is induced by the map

$$\gamma_1: (C, \partial_C) \longrightarrow (C \oplus D, \partial^f), x \longmapsto x \oplus 0;$$

- Γ_2 is induced by the map

$$\gamma_2: (C \oplus D, \partial^f) \longrightarrow (D, -\partial_D), x \oplus y \longmapsto -y.$$

Proof. We first have to see that γ_1 and γ_2 are chain maps. Given an element $c \in C$, observe that

$$\gamma_1(\partial_C c) = \partial_C c = \partial^f c = \partial^f \gamma_1(c).$$

Furthermore, we see that

$$\gamma_2(\partial^f(c \oplus 0)) = \gamma_2(\partial_C c) = 0 = \gamma_2(c \oplus 0) = -\partial_D(\gamma_2(c \oplus 0)).$$

We continue with an element $d \in D$:

$$\gamma_2(\partial^f(0 \oplus d)) = \gamma_2(f(d) - \partial_D(d)) = \partial_D(d) = -\partial_D(\gamma_2(0 \oplus d)).$$

Thus, both γ_1 and γ_2 are chain maps. Finally, γ_1 and γ_2 obviously fit into the short exact sequence

$$0 \longrightarrow (C, \partial_C) \xrightarrow{\gamma_1} (C \oplus D, \partial^f) \xrightarrow{\gamma_2} (D, -\partial_D) \longrightarrow 0$$

of chain complexes. Hence, by standard results in Algebraic Topology (see [1]) this short exact sequence induces a long exact sequence

$$\dots \xrightarrow{\partial_*} H_*(C, \partial_C) \xrightarrow{\Gamma_1} H_*(C \oplus D, \partial^f) \xrightarrow{\Gamma_2} H_*(D, -\partial_D) \xrightarrow{\partial_*} \dots$$

It remains to show that the connecting homomorphism ∂_* equals $-f_*$. Recall that for $d \in \ker(\partial_D)$ the morphism ∂_* is defined by

$$\partial_*[d] = [\gamma_1^{-1}(\partial^f(\gamma_2^{-1}(d)))].$$

Of course, γ_1 and γ_2 are not necessarily invertible. However, we take the preimages as given in the equation, and, by standard algebraic topology, all the elements in the preimage will belong to the same equivalence class. Observe:

$$\begin{aligned} \partial_*[d] &= [\gamma_1^{-1}(\partial^f(\gamma_2^{-1}(d)))] \\ &= [\gamma_1^{-1}(\partial^f(0 \oplus -d))] \\ &= [\gamma_1^{-1}(-f(d))] \\ &= -[f(d)] \\ &= -f_*[d] \end{aligned}$$

□

Of course, the whole construction works if f goes the other way, i.e. $f: C \longrightarrow D$. In this case we form the complex $C \oplus D$ with the differential

$$\partial_f = \begin{pmatrix} \partial_C & 0 \\ f & -\partial_D \end{pmatrix}.$$

In an analogous manner we obtain a long exact sequence

$$\dots \xrightarrow{-f_*} H_*(D, -\partial_D) \xrightarrow{\Gamma_1} H_*(C \oplus D, \partial_f) \xrightarrow{\Gamma_2} H_*(C, \partial_C) \xrightarrow{-f_*} \dots$$

3.2 Two New Exact Sequences in Heegaard Floer Homology

3.2.1 Positive Dehn Twists

Let an open book (P, ϕ) and a homologically essential closed curve δ in P be given. We first ask how a Dehn twist along δ would change the associated Heegaard Floer homology. There is a specific choice of attaching circles that are – in a sense – adapted to the closed curve δ . Figure 3.1 depicts a small neighborhood of the point $\delta \cap \beta_1$ in the Heegaard diagram induced by the open book decomposition. The page at the right side of the boundary pictured in Figure 3.1 is $P \times \{1/2\}$. The dotted line indicates the neighborhood of ∂P where the monodromy ϕ is the identity. The proof of Lemma 2.7.19 shows that we can arrange a neighborhood of $\delta \cap \beta_1$ to look like in Figure 3.1, i.e. it is possible to arrange the curve δ and the attaching circles like indicated in Figure 3.1 due to the arguments given in the proof of Lemma 2.7.19.

With respect to the surface orientation given in Figure 3.1 this is the appropriate setup for performing a positive Dehn twist along δ : Denote by β' the β -curves after performing the Dehn twist. Obviously, $\beta' = \{\beta'_1, \beta_2, \dots, \beta_{2g}\}$. Observe that

$$\mathbb{T}_\alpha \cap \mathbb{T}_{\beta'} = \mathbb{T}_\alpha \cap \mathbb{T}_\beta \sqcup \mathbb{T}_\alpha \cap \mathbb{T}_\delta, \quad (3.2.1)$$

where \mathbb{T}_δ is given by the set $\delta = \{\delta, \beta_2, \dots, \beta_{2g}\}$ (by abuse of notation since δ also denotes the curve on P but what is meant will be clear from the context). The set of curves δ may be interpreted as a set of attaching circles. In the following we will call the arc $\beta'_1 \cap \beta_1$ the **β -part of β'_1** and the arc $\beta'_1 \cap \delta$ the **δ -part of β'_1** . Figure 3.2 depicts the situation before and after the Dehn twist.

The main observation is that there can be no holomorphic disc in (Σ, α, β') that connects a $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$ -intersection of $\mathbb{T}_\alpha \cap \mathbb{T}_{\beta'}$ with a $\mathbb{T}_\alpha \cap \mathbb{T}_\delta$ -intersection of $\mathbb{T}_\alpha \cap \mathbb{T}_{\beta'}$.




Figure 3.1: A small neighborhood of $\delta \cap \beta_1$ in the Heegaard surface $\Sigma = P \times \{1/2\} \cup (-P) \times \{0\}$.




Figure 3.2: Before and after the positive Dehn twist.

Suppose there is a disc ϕ starting at $x \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ and going to $y \in \mathbb{T}_\alpha \cap \mathbb{T}_{\delta'}$ along its α -boundary. Then, at the β -boundary, the disc ϕ has to run from y to x along the β' -curves. Since $\delta \cap \beta_1$ contains only one point, namely the intersection that can be seen in Figures 3.1 and 3.2, the disc has to run through either D_* or D_{**} (since $n_z(\phi) = 0$ we cannot use the D_z -region). But since we are moving from the δ -part of β'_1 to the β -part of β'_1 , we see that $n_*(\phi) < 0$ or $n_{**}(\phi) < 0$, in contradiction to holomorphicity. So, there are just three choices for the β -boundary of a holomorphic disc.

1. It starts at the δ -part of β'_1 and stays there.
2. It starts at the β -part of β'_1 and stays there.

3. It starts at the β -part of β'_1 and runs to the δ -part of β'_1 and stays there.

This immediately shows that

$$\widehat{\text{HF}}(Y^\delta) = H_*(\widehat{\text{CF}}(\alpha, \beta) \oplus \widehat{\text{CF}}(\alpha, \delta), \partial),$$

where ∂ is of the form

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

If we perform a negative Dehn twist along δ in the situation indicated in Figure 3.1, we would connect \mathcal{D}_* with \mathcal{D}_{**} and keep separate \mathcal{D}_w and \mathcal{D}_z . Observe that we would have, a priori, no control of holomorphic discs like in the case of positive Dehn twists. To get back into business, in case of negative Dehn twists, we have to first isotope δ inside the page of the open book appropriately (see §3.2.2).

Proposition 3.2.1. *Let (Σ, α, β) be a δ -adapted Heegaard diagram of Y and denote by Y^δ the manifold obtained from Y by composing the gluing map, given by the attaching curves α, β , with a positive Dehn twist along δ as indicated in Figure 3.2. Then the following holds:*

$$\widehat{\text{HF}}(Y^\delta) \cong H_*(\widehat{\text{CF}}(\alpha, \beta) \oplus \widehat{\text{CF}}(\alpha, \delta), \partial^f),$$

where ∂^f is of the form

$$\begin{pmatrix} \widehat{\partial}_{\alpha\beta}^w & f \\ 0 & \widehat{\partial}_{\alpha\delta}^w \end{pmatrix}$$

with f a chain map between $(\widehat{\text{CF}}(\alpha, \delta), \widehat{\partial}_{\alpha\delta}^w)$ and $(\widehat{\text{CF}}(\alpha, \beta), \widehat{\partial}_{\alpha\beta}^w)$.

Proof. There is a natural identification of intersection points

$$\mathbb{T}_\alpha \cap \mathbb{T}_{\beta'} \longleftrightarrow \mathbb{T}_\alpha \cap \mathbb{T}_\beta \sqcup \mathbb{T}_\alpha \cap \mathbb{T}_\delta,$$

i.e. we get an isomorphism

$$\epsilon: \widehat{\text{CF}}(\alpha, \beta') \xrightarrow{\cong} \widehat{\text{CF}}(\alpha, \beta) \oplus \widehat{\text{CF}}(\alpha, \delta)$$

of modules. Pick an intersection point $x \in \mathbb{T}_\alpha \cap \mathbb{T}_{\beta'}$ such that $\epsilon(x) \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$. Looking at the boundary

$$\widehat{\partial}^\delta x = \sum_y \sum_\phi \# \widehat{\mathcal{M}}_\phi \cdot y \tag{3.2.2}$$

we want to see that the moduli space of holomorphic discs connecting x with an intersection $y \in \epsilon^{-1}(\mathbb{T}_\alpha \cap \mathbb{T}_\delta)$ is empty: Assume this were not the case. This means

there were a holomorphic disc ϕ connecting x with an element $y = (y_1, \dots, y_n) \in \epsilon^{-1}(\mathbb{T}_\alpha \cap \mathbb{T}_\delta)$. Observe that y_1 is a point in $\delta \cap \alpha_1$. Hence, $\mathcal{D}(\phi)$ includes \mathcal{D}_* or \mathcal{D}_{**} since these are the only domains giving a connection between $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$ and $\mathbb{T}_\alpha \cap \mathbb{T}_\delta$. Boundary orientations force the coefficient of ϕ at \mathcal{D}_* or \mathcal{D}_{**} to be negative. Since holomorphic maps are orientation preserving, this cannot be the case. So, the point x can be connected to points in $\epsilon^{-1}(\mathbb{T}_\alpha \cap \mathbb{T}_\beta)$ only.

Next observe that discs ϕ appearing in the sum (3.2.2) all have the property $n_*(\phi) = n_{**}(\phi) = 0$. Indeed, suppose there were a disc ϕ with nonnegative intersection n_* or n_{**} . The β -boundary of ϕ starts at x and runs through $\partial\mathcal{D}_*$ or $\partial\mathcal{D}_{**}$. The disc ϕ is holomorphic, so, the β -boundary runs from the β -part to the δ -part of $\mathbb{T}_{\beta'}$. At the end of the β -boundary of ϕ the disc converges to a point in $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$. Thus, the β -boundary of ϕ has to come back through either \mathcal{D}_* or \mathcal{D}_{**} . The boundary orientation would force ϕ to negatively intersect $\{*\} \times \text{Sym}^{g-1}(\Sigma)$ or $\{**\} \times \text{Sym}^{g-1}(\Sigma)$. This cannot happen.




Figure 3.3: Picture of the three different boundary conditions arising in our discussion.

Denote by $[a, c]$ the small arc in β'_1 running through Figure 3.3 and define $[b, d]$ analogously. All discs arising in the sum have boundary conditions in \mathbb{T}_α and

$$\mathbb{T}_{\beta'} \setminus \{ \{ [a, c] \sqcup [b, d] \} \times \beta_2 \times \dots \times \beta_g \}.$$

Observe that $\mathbb{T}_{\beta'} \setminus \{ \{ [a, c] \sqcup [b, d] \} \times \beta_2 \times \dots \times \beta_g \}$ has two components, one lying in \mathbb{T}_β and one lying in \mathbb{T}_δ . Since the β -boundary of the disc ϕ starts in \mathbb{T}_β , it remains there all the time. Moreover, looking at discs ϕ in $(\Sigma, \alpha, \beta, z, w)$ with $n_z(\phi) = n_w(\phi) = 0$, an analogous line of arguments as above shows that the β -boundary of these discs stays away from

$$[a, b] \times \beta_2 \times \dots \times \beta_g,$$

where $[a, b]$ is the arc in β pictured in the right of Figure 3.3. Thus, the boundary conditions for discs connecting intersections $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$ are the same in $(\Sigma, \alpha, \beta', z)$ and $(\Sigma, \alpha, \beta, z, w)$. Thus, we have

$$\widehat{\partial}^\delta x = \widehat{\partial}_{\alpha\beta}^w x.$$

Now suppose that $x \in \epsilon^{-1}(\mathbb{T}_\alpha \cap \mathbb{T}_\delta)$. Then

$$\begin{aligned} \widehat{\partial}^\delta x &= \sum_y \sum_\phi \# \widehat{\mathcal{M}}_\phi \cdot y \\ &= \sum_{y \in \mathbb{T}_\alpha \cap \mathbb{T}_\delta} \sum_\phi \# \widehat{\mathcal{M}}_\phi \cdot y + \sum_{z \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} \sum_\phi \# \widehat{\mathcal{M}}_\phi \cdot z. \end{aligned}$$

With an analogous line of arguments as above we see that the first sum counts discs with $n_* = n_{**} = n_z = 0$ only. The triviality of these intersection numbers and holomorphicity implies that the discs have boundary conditions in \mathbb{T}_α and

$$\mathbb{T}_{\beta'} \setminus \{[a, c] \sqcup [b, d]\} \times \beta_2 \times \dots \times \beta_g.$$

As mentioned above this set has two components, where one of them lies in \mathbb{T}_δ . The β -boundary of ϕ starts in \mathbb{T}_δ and therefore remains there all the time. Again, we see that discs connecting intersection points $\mathbb{T}_\alpha \cap \mathbb{T}_\delta$ in $(\Sigma, \alpha, \beta', z)$ and $(\Sigma, \alpha, \delta, z, w)$ have to fulfill identical boundary conditions. Thus, the moduli spaces are isomorphic. This shows the equality

$$\widehat{\partial}^\delta x = \widehat{\partial}_{\alpha\delta}^w x + \sum_{z \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} \sum_\phi \# \widehat{\mathcal{M}}_\phi \cdot z.$$

In the right sum we only count discs where $n_* \neq 0$ or $n_{**} \neq 0$. We will denote this right sum with $f(x)$. We have to see that f defines a chain map

$$f: (\widehat{\text{CF}}(\alpha, \delta), \widehat{\partial}_{\alpha\delta}^w) \longrightarrow (\widehat{\text{CF}}(\alpha, \beta), \widehat{\partial}_{\alpha\beta}^w).$$

This can be proved in two ways: We know that $\partial^\delta = \partial_{\alpha\beta}^w + \partial_{\alpha\delta}^w + f$. Hence, f is a sum of three boundaries. The equality $0 = (\partial^\delta)^2$ implies that f is a chain map (cf. Lemma 3.1.1). The second way is to test the chain map property directly. To do so, pick a generator $y \in \mathbb{T}_\alpha \cap \mathbb{T}_{\beta'}$ lying in the preimage of $\mathbb{T}_\alpha \cap \mathbb{T}_\delta$ under ϵ . Observe that $(\widehat{\partial}_{\alpha\beta}^w \circ f - f \circ \widehat{\partial}_{\alpha\delta}^w)(x)$ equals

$$\begin{aligned} &\sum_{z \in \mathbb{T}_\alpha \cap \mathbb{T}_\delta} \left(\sum_{(y, \phi_2, \phi_1)} \# \widehat{\mathcal{M}}(\phi_2) \# \widehat{\mathcal{M}}(\phi_1) - \sum_{(y', \phi'_2, \phi'_1)} \# \widehat{\mathcal{M}}(\phi'_2) \# \widehat{\mathcal{M}}(\phi'_1) \right) \cdot z \\ &= \sum_{z \in \mathbb{T}_\alpha \cap \mathbb{T}_\delta} c(x, z) \cdot z, \end{aligned}$$

where the first sum in the definition of $c(x, z)$ goes over elements (y, ϕ_2, ϕ_1) in the set $\mathbb{T}_\alpha \cap \mathbb{T}_\beta \times \pi_2(y, z) \times \pi_2(x, y)$ with $\mu(\phi_2) = \mu(\phi_1) = 1$, and the second sum goes over $(y', \phi'_2, \phi'_1) \in \mathbb{T}_\alpha \cap \mathbb{T}_\delta \times \pi_2(y, z) \times \pi_2(x, y)$ with $\mu(\phi'_2) = \mu(\phi'_1) = 1$. Furthermore, look at the boundary of a moduli space $\widehat{\mathcal{M}}(\phi)$ connecting a point in $\mathbb{T}_\alpha \cap \mathbb{T}_\delta$ with a point in $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$ with $\mu(\phi) = 2$. Observe that we do not have to take care of boundary degenerations or spheres bubbling off since we are looking for maps with $n_z = 0$ (cf. [40]). The only phenomenon appearing at the boundary is breaking. The boundary of $\widehat{\mathcal{M}}(\phi)$ is modelled on

$$\bigsqcup_{\phi_1 * \phi_2 = \phi} \widehat{\mathcal{M}}(\phi_1) \times \widehat{\mathcal{M}}(\phi_2).$$

There are two cases. Either $n_*(\phi_1) = n_*(\phi)$ or $n_*(\phi_2) = n_*(\phi)$ (the discussion for n_{**} is analogous):

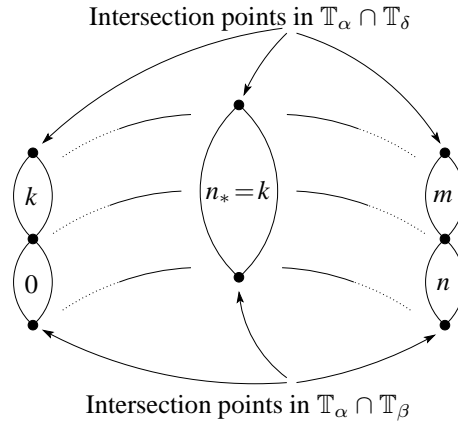


Figure 3.4: Here we figure a moduli space with $\mu = 2$ and its possible ends.

To prove this, we have to show that a given family of discs ϕ_n in $\widehat{\mathcal{M}}(\phi)$ cannot converge to a broken disc $\phi_1 * \phi_2$ with $n = n_*(\phi_1) \neq 0$ and $m = n_*(\phi_2) \neq 0$. Figure 3.4 represents a moduli space of discs with $\mu = 2$ and $n_*(\phi_n) = k$. We know that $n + m = k$, since intersection numbers behave additively under concatenation. Assume that n, m were both non-zero: Since n is non-zero, we know that ϕ_1 connects a point in $\mathbb{T}_\alpha \cap \mathbb{T}_\delta$ with one in $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$. The bottom intersection is a $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$ -intersection, since ϕ_n connects $\mathbb{T}_\alpha \cap \mathbb{T}_\delta$ with an $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$ -intersection by assumption. Hence, ϕ_2 connects a point of $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$ with a point in $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$ and runs through the domain \mathcal{D}_* . This is simply not possible because of orientation reasons. Thus, either

$n_*(\phi_1) = k$ and $n_*(\phi_2) = 0$ or $n_*(\phi_1) = 0$ and $n_*(\phi_2) = k$. This means the ends of $\widehat{\mathcal{M}}(\phi)$ precisely look like

$$\left(\bigsqcup_{\phi_2 * \phi_1 = \phi} \widehat{\mathcal{M}}(\phi_2)^* \times \widehat{\mathcal{M}}(\phi_1) \right) \sqcup \left(\bigsqcup_{\phi'_2 * \phi'_1 = \phi} \widehat{\mathcal{M}}(\phi_2) \times \widehat{\mathcal{M}}(\phi_1)^* \right),$$

where $*$ means that the associated discs have non-trivial intersection number n_* or n_{**} . Now consider the union of moduli spaces of discs connecting the point x and z with Maslov index 2. According to our discussion, the ends look like

$$\left(\bigsqcup_{(y, \phi_2, \phi_1)} \widehat{\mathcal{M}}(\phi_2) \times \widehat{\mathcal{M}}(\phi_1)^* \right) \sqcup \left(\bigsqcup_{(y', \phi'_2, \phi'_1)} \widehat{\mathcal{M}}(\phi'_2)^* \times \widehat{\mathcal{M}}(\phi'_1) \right),$$

where the first union goes over $(y, \phi_2, \phi_1) \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta \times \pi_2(y, z) \times \pi_2(x, y)$ with $\mu(\phi_2) = \mu(\phi_1) = 1$ and the second union goes over $(y', \phi'_2, \phi'_1) \in \mathbb{T}_\alpha \cap \mathbb{T}_\delta \times \pi_2(y, z) \times \pi_2(x, y)$ with $\mu(\phi'_2) = \mu(\phi'_1) = 1$. Hence, the coefficients $c(x, z)$ all vanish, proving the theorem. \square

An immediate, simple algebraic consequence (cf. §3.1) of this description is the following Corollary.

Corollary 3.2.2. *Let $K \subset Y$ be the knot determined by δ . Then there is a long exact sequence*

$$\dots \xrightarrow{\partial_*} \widehat{\text{HF}}\widehat{\text{K}}(Y, K) \xrightarrow{\Gamma_1} \widehat{\text{HF}}(Y_{-1}(K)) \xrightarrow{\Gamma_2} \widehat{\text{HF}}\widehat{\text{K}}(Y_0(K), \mu) \xrightarrow{\partial_*} \dots$$

with $\partial_* = -f_*$ where f is the map defined in the proof of Proposition 3.2.1. The knot μ denotes a meridian of K .

Proof. With Proposition 3.2.1 we see that $\widehat{\text{HF}}(Y^\delta)$ fulfills the assumptions of Lemma 3.1.1 and therefore Lemma 3.1.2 applies. Finally, we apply Proposition 2.4.4 to identify $H_*(\widehat{\text{CF}}, \widehat{\partial}^w)$ with the respective knot Floer homology. It is easy to observe that with respect to the framing induced by the open book the manifold Y^δ equals $Y_{-1}(K)$, i.e. the result of (-1) -surgery along the knot K . We obtain the sequence

$$\dots \xrightarrow{\partial_*} \widehat{\text{HF}}\widehat{\text{K}}(Y, K) \xrightarrow{\Gamma_1} \widehat{\text{HF}}(Y_{-1}(K)) \xrightarrow{\Gamma_2} \widehat{\text{HF}}\widehat{\text{K}}(Y_{\alpha\delta}, K_2) \xrightarrow{\partial_*} \dots,$$

where $(Y_{\alpha\delta}, K_2)$ is the pair given by the data $(\Sigma, \alpha, \delta, z, w)$. It is easy to see that the pair (w, z) in the diagram (Σ, α, δ) determines β_1 up to orientation, i.e. the attaching circle β_1 interpreted as a knot in $Y_{\alpha\delta}$. This attaching circle β_1 is a meridian for a tubular neighborhood μ of K in Y . Finally, we have to see that $Y_{\alpha\delta}$ equals the 0-surgery along K with respect to the framing induced by the open book. This is straightforward. \square

Corollary 3.2.3. *In the situation of Proposition 3.2.1 we define a map*

$$f: \widehat{\text{CFK}}(\Sigma, \alpha, \delta, z, w) \longrightarrow \widehat{\text{CFK}}(\Sigma, \alpha, \beta, z, w)$$

by sending an element $x \in \mathbb{T}_\alpha \cap \mathbb{T}_\delta$ to

$$f(x) = \sum_{z \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} \sum_{\phi \in H(x, y, 1)} \# \widehat{\mathcal{M}}_\phi \cdot y,$$

where $H(x, y, 1)$ are classes in $\pi_2^{\alpha\beta'}(x, y)$ with $\mu = 1$ and with the pair of intersection numbers $(n_*(\phi), n_{**}(\phi)) \neq (0, 0)$. We denote by $\pi_2^{\alpha\beta'}(x, y)$ the homotopy classes of Whitney discs associated to the diagram $(\Sigma, \alpha, \beta', z)$. The defined f is a chain map and its induced map on homology satisfies $f_* = \partial_*$ where ∂_* is the connecting morphism in the sequence given in Corollary 3.2.2. \square

A few words about admissibility: The reader may have noticed that we did not say anything about admissibility of the Heegaard diagram $(\Sigma, \alpha, \delta, z, w)$, but nonetheless talk about the knot Floer homology $\widehat{\text{HFK}}(Y_{\alpha\delta}, K_2)$ induced by this diagram. We could restrict to just saying we take the homology induced by the data. The respective boundary operator is well defined (finite sum) since $\widehat{\partial}^\delta$ is. However, we would like to remark that the diagram $(\Sigma, \alpha, \delta, z, w)$ is always admissible *in a relaxed sense*. We may relax the weak-admissibility condition imposed by Ozsváth and Szabó for the definition of knot Floer homology to the extreme weak-admissibility condition given in Definition 2.4.5. The diagram $(\Sigma, \alpha, \delta, z, w)$ is always extremely weakly-admissible: Let \mathcal{D} be a non-trivial periodic domain with $n_w(\mathcal{D}) = 0$ (see §2.4) and let s be an arbitrary Spin^c -structure such that $\langle c_s(s), \mathcal{H}(\mathcal{D}) \rangle = 0$. By definition of the boundary, $\partial\mathcal{D}$ can be written as

$$\partial\mathcal{D} = \sum_{i \geq 1} \lambda_i \alpha_i + \kappa_1 \delta + \sum_{j \geq 2} \kappa_j \beta_j.$$

Assuming that $\lambda_i \neq 0$ for a $i \geq 2$ or $\kappa_j \neq 0$ for a $j \geq 2$, we see that \mathcal{D} has both positive and negative coefficients due to the fact that $\partial\mathcal{D}$ runs through a configuration like given in Figure 2.10. Thus, let us assume that λ_i and κ_j would vanish, for all $i, j \geq 2$. The boundary of \mathcal{D} could be written as

$$\partial\mathcal{D} = \lambda_1 \alpha_1 + \kappa_1 \delta.$$

However, κ_1 has to vanish, since δ runs through $\partial\overline{\mathcal{D}}_w \cap \partial\overline{\mathcal{D}}_z$ (see Figure 3.3). Finally, we get that $\partial\mathcal{D} = \lambda_1 \alpha_1$. Examining the middle part of Figure 3.3 we see that the part

of α_1 which is at the right of δ is surrounded by the region \mathcal{D}_z . Thus, $\lambda_1 = 0$.

With help of the geometric realization of the $\bigwedge^*(H_1/Tor)$ -module structure given in [40] we can easily prove the following proposition.

Proposition 3.2.4. *The maps Γ_1 and Γ_2 from the exact sequence of Corollary 3.2.2 respect the $\bigwedge^*(H_1/Tor)$ -module structure of the Heegaard Floer groups in the following sense. Let $\gamma \subset \Sigma$ be a curve. Then the following identities hold:*

$$\begin{aligned} A_{[\gamma]_{Y\delta}}^{Y\delta}(\Gamma_1(x)) &= \Gamma_1(A_{[\gamma]_Y}^Y(x)) \\ \Gamma_2(A_{[\gamma]_{Y\delta}}^{Y\delta}(x)) &= A_{[\gamma]_{Y\alpha\delta}}^{Y\alpha\delta}(\Gamma_2(x)) \end{aligned}$$

Proof. Recall the geometric realization of the $\bigwedge^*(H_1/Tor)$ -module structure. Given a point $x \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta \subset \mathbb{T}_\alpha \cap \mathbb{T}_{\beta'}$ (cf. the proof of Proposition 3.2.1 for the appropriate identification), by definition

$$A_{[\gamma]_{Y\delta}}^{Y\delta}(x) = \sum_y \sum_{\phi \in H(x,y,1)} a(\gamma, \phi) \cdot y,$$

where $H(x, y, 1) \subset \pi_2(x, y)$ is the set of Whitney discs with $n_z = 0$ and $\mu = 1$. Furthermore,

$$a(\gamma, \phi) = \#\widehat{\mathcal{M}}_\phi \cdot \#(u(\{-1\} \times \mathbb{R}, \gamma \times \text{Sym}^{g-1}(\Sigma))_{\mathbb{T}_\alpha}.$$

where the right factor denotes the intersection number of $u(\{-1\} \times \mathbb{R})$ and $\gamma \times \text{Sym}^{g-1}(\Sigma)$ inside \mathbb{T}_α . Fixing another point $y \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, recall that these points are connected by $\widehat{\partial}_{\alpha\beta}^w$ if and only if they are connected by $\widehat{\partial}^\delta$. Moreover, there is an identification of the respective moduli spaces. Thus, fixing a disc ϕ connecting these points (in $\alpha\beta'$), we know – since $n_z(\phi) = 0$ – that ϕ connects these intersection points in the $\alpha\beta$ -diagram, too. Denoting by $[\phi]$ its class in π_2 , we see that

$$\#\widehat{\mathcal{M}}_{[\phi]}^{\alpha\beta} = \#\widehat{\mathcal{M}}_{[\phi]}^{\alpha\beta'}.$$

Moreover, the intersection number in \mathbb{T}_α used to define $a(\gamma, [\phi])$ coincides for both diagrams since ϕ is a common representative. Thus, we see that

$$a^{Y\delta}(\gamma, [\phi]) = a^Y(\gamma, [\phi]).$$

Recall that there are no connections from $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$ -intersections to a $\mathbb{T}_\alpha \cap \mathbb{T}_\delta$ -intersection in the α, β' -diagram. Hence, the first equality given in the proposition follows.

To show the second, fix a point $x \in \mathbb{T}_\alpha \cap \mathbb{T}_\delta \subset \mathbb{T}_\alpha \cap \mathbb{T}_{\beta'}$. Use the same line of arguments as above to show that the following identity holds:

$$\begin{aligned} A_{[\gamma]_Y^\delta}^{Y^\delta}(x) &= \sum_y \sum_{\phi \in H(x,y,1)} a^{Y^\delta}(\gamma, \phi) \cdot y + \sum_z \sum_{\psi \in H(x,z,1)} a^{Y^\delta}(\gamma, \psi) \cdot z \\ &= A_{[\gamma]_{Y_{\alpha\delta}}^{Y_{\alpha\delta}}}(x) + \sum_z \sum_{\psi \in H(x,z,1)} a^{Y^\delta}(\gamma, \psi) \cdot z. \end{aligned}$$

The second sum is an element in $\widehat{\text{CF}}(\Sigma, \alpha, \beta, z, w)$. Recall that Γ_2 is induced by the projection onto $\widehat{\text{CF}}(\Sigma, \alpha, \delta, z, w)$. Hence, the second sum cancels when projected under the map Γ_2 . The second equality of the proposition follows. \square

In §3.3 we will derive suitable naturality properties of the sequence to show that the maps involved in the sequences are indeed topological. We will be interested in the maps denoted by Γ_1 since these are directly related to the surgery represented by the Dehn twist.

3.2.2 Negative Dehn Twists

The approach for negative Dehn twists is pretty much the same as for positive Dehn twists. In §3.2.1 we already mentioned that the situation indicated in Figure 3.1 is not suitable for performing negative Dehn twists. Performing a negative twist, we could not make an a priori statement about what generators can be connected by holomorphic discs like we did in §3.2.1. To get back into business we just need to isotope the curve δ inside the page a bit (or equivalently isotope some of the attaching circles). Figure 3.5 indicates a possible perturbation suitable for our purposes. Comparing Figures 3.2 and 3.5 we see that we isotoped the curve δ a bit. Observe that with this perturbation done, we again can read off the behavior of holomorphic discs like in §3.2.1 (carry over the discussion of §3.2.1 to this situation). As a consequence, the following proposition can be proved. The proof of Proposition 3.2.1 carries over verbatim to a proof of Proposition 3.2.5.

Proposition 3.2.5. *Let (Σ, α, β) be a δ -adapted Heegaard diagram of Y and denote by Y^δ the manifold obtained from Y by composing the gluing map, given by the attaching curves α, β , with a negative Dehn twist along δ as hinted in Figure 3.5. Then we have*

$$\widehat{\text{HF}}(Y^\delta) \cong H_*(\widehat{\text{CF}}(\alpha, \beta) \oplus \widehat{\text{CF}}(\alpha, \delta), \partial^f),$$




Figure 3.5: Before and after a negative Dehn twist along δ .

where ∂^f is of the form

$$\begin{pmatrix} \widehat{\partial}_{\alpha\beta}^w & 0 \\ f & \widehat{\partial}_{\alpha\delta}^w \end{pmatrix}$$

with f a chain map between $(\widehat{\text{CF}}(\alpha, \delta), \widehat{\partial}_{\alpha\delta}^w)$ and $(\widehat{\text{CF}}(\alpha, \beta), \widehat{\partial}_{\alpha\beta}^w)$. \square

Corollary 3.2.6. *Let $K \subset Y$ be the knot determined by δ . Then there is a long exact sequence*

$$\dots \xrightarrow{\partial_*} \widehat{\text{HFK}}(Y_0(K), \mu) \xrightarrow{\Gamma_2} \widehat{\text{HF}}(Y_{+1}(K)) \xrightarrow{\Gamma_1} \widehat{\text{HFK}}(Y, K) \xrightarrow{\partial_*} \dots$$

with $\partial_* = -f_*$ where f is the map defined in the proof of Proposition 3.2.5. The knot μ denotes a meridian of K . Moreover, identities hold similar to those given in Proposition 3.2.4. \square

3.3 Invariance

Our goal in this paragraph is to show that the map Γ_1 in the sequences introduced are topological, i.e. just depend on the cobordism associated to the surgery represented by the Dehn twist. To do that, we have to generalize our approach a bit and try to see that everything we have done, especially the proof of Proposition 3.2.1, works without using a Heegaard diagram that is necessarily induced by an open book. Obviously, the geometric situation given in Figure 3.3 builds the foundation of the proof. To clarify the situation, look at Figure 3.6.




Figure 3.6: The important geometric configuration.

We, for the moment, stick to the notation of §3.2. We need the curve δ to intersect β_1 once, transversely and to be disjoint from the other β -circles. In addition, the top right domain at the point $\delta \cap \beta_1 \in \Sigma$ has to contain the base point z (cf. Figure 3.6). Given this configuration, the proof of Proposition 3.2.1 applies. The situation figured, does not occur exclusively when the Heegaard diagram is induced by an open book.




Figure 3.7: Preparation of the Heegaard diagram.

Given a Heegaard diagram subordinate to a knot K , we can isotope the knot K onto the Heegaard surface. The isotoped knot intersects just one β -circle once, transversely. Without loss of generality K intersects β_1 . To generate a geometric configuration like indicated in Figure 3.6, we may isotope the knot again to move the intersection $\beta_1 \cap K$ to lie next to a \mathcal{D}_z -region: Cutting the α -circles out of the Heegaard surface, we obtain a sphere with holes. The region \mathcal{D}_z is a region in this sphere. Either \mathcal{D}_z is the whole sphere with holes or not. In case it is the whole sphere all the β -circles touch the region \mathcal{D}_z and we are done. In case \mathcal{D}_z is not the whole sphere, there has to be at least one β -circle touching the boundary of \mathcal{D}_z . If β_1 touches the boundary of \mathcal{D}_z , we are done. If β_1 does not touch the boundary of \mathcal{D}_z , we obtain a configuration like indicated in part (a) of Figure 3.7. Without loss of generality we assume that β_2 touches \mathcal{D}_z . Note that it not possible for β_2 to separate \mathcal{D}_z from β_1 , since the complement of the β -circles in Σ is connected. We are allowed to slide β_1 over this β -circle (cf. part (b)

of Figure 3.7). After the handle slide there is a small arc a inside β_1 touching \mathcal{D}_z . By a small isotopy of the knot K we can move the intersection point $K \cap \beta_1$ along the new β_1 -circle until it enters the arc a (cf. part (c) of Figure 3.7).

Care has to be taken of the surgery framing. Here, we stick to surgeries or to framed knots K such that there exists a subordinate Heegaard diagram with the framing induced by the Heegaard surface coinciding with the framing of the knot. Evidence indicate that every framing can be realized in this way.

We saw that our discussion from the last paragraph can be carried over to a more general situation. We, indeed, do not need the Heegaard diagram to be induced by an open book. So far, we restricted the discussion to Heegaard diagrams induced by open books, since we are interested in applications to the contact geometric parts of the theory, which makes a discussion of this class of diagrams inevitable.

Given two Heegaard diagrams subordinate to a pair (Y, δ) , we transform the one diagram into the other by the moves introduced in Lemma 2.4.3. These moves respect the knot complement of δ . The goal is to show that each move preserves the exact sequence and the maps inherited. In the following we will call Heegaard diagrams, realizing a geometric situation as given in Figure 3.3 for a knot δ , δ -**suitable**.




Figure 3.8: Triangles that have to be counted for handle slides among the α -curves.

We begin showing invariance under handle slides among the α -curves. Although used in some papers it was never explicitly mentioned which triangles are counted for handle slides among the α -curves (cf. §2.3.4). Given a Heegaard diagram (Σ, α, β) , denote by α' the attaching circles obtained by a handle slide among the α -curves. The associated map between the respective homologies counts holomorphic triangles with boundary conditions in α, α' and β . Figure 3.8 pictures a Whitney triangle connecting a point $x \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ with a point $y \in \mathbb{T}_{\alpha'} \cap \mathbb{T}_\beta$. Observe that in this situation Θ is a top-dimensional generator of $\widehat{\text{HF}}(\alpha', \alpha)$ (note the order of the attaching circles). To not confuse the maps induced by handle slides among the α -circles with the maps

induced by handle slides among the β -circles, we introduce the following notation: let us denote by $\Gamma_{\alpha, \alpha'; \beta}$ the map induced by a handle slide among the α -circles (like indicated above) and by $\Gamma_{\alpha; \beta, \beta'}$ the map induced by a handle slide among the β -circles (like indicated in §2.3.4).

Proposition 3.3.1. *Let $(\Sigma, \alpha, \beta, z)$ be a δ -suitable Heegaard diagram and $(\Sigma, \alpha', \beta, z)$ be obtained by a handle slide of one of the α_i . Denote by*

$$\begin{aligned} \Gamma_{\alpha, \alpha'; \beta}^w &: \widehat{\text{CFK}}(\Sigma, \alpha, \beta, z, w) \longrightarrow \widehat{\text{CFK}}(\Sigma, \alpha', \beta, z, w) \\ \Gamma_{\alpha, \alpha'; \delta}^w &: \widehat{\text{CFK}}(\Sigma, \alpha, \delta, z, w) \longrightarrow \widehat{\text{CFK}}(\Sigma, \alpha', \delta, z, w) \\ \Gamma_{\alpha, \alpha'; \beta'} &: \widehat{\text{CF}}(\Sigma, \alpha, \beta', z) \longrightarrow \widehat{\text{CF}}(\Sigma, \alpha', \beta', z) \end{aligned}$$

the induced maps. These maps induce a commutative diagram with exact rows

$$\begin{array}{ccccccc} \dots & \xrightarrow{\partial_*} & \widehat{\text{HF}}(\Sigma, \alpha, \beta, z, w) & \xrightarrow{\Gamma_1} & \widehat{\text{HF}}(\Sigma, \alpha, \beta', z) & \xrightarrow{\Gamma_2} & \widehat{\text{HF}}(\Sigma, \alpha, \delta, z, w) & \xrightarrow{\partial_*} & \dots \\ & & \downarrow \Gamma_{\alpha, \alpha'; \beta}^{w,*} & & \downarrow \Gamma_{\alpha, \alpha'; \beta'}^* & & \downarrow \Gamma_{\alpha, \alpha'; \delta}^{w,*} & & \\ \dots & \xrightarrow{\partial'_*} & \widehat{\text{HF}}(\Sigma, \alpha', \beta, z, w) & \xrightarrow{\Gamma'_1} & \widehat{\text{HF}}(\Sigma, \alpha', \beta', z) & \xrightarrow{\Gamma'_2} & \widehat{\text{HF}}(\Sigma, \alpha', \delta, z, w) & \xrightarrow{\partial'_*} & \dots \end{array}$$

Proof. The proof of this proposition is quite similar to the proof of Proposition 3.2.1. To keep the exposition efficient, we do not point out all details here. Start looking at the map $\Gamma_{\alpha, \alpha'; \beta'}$. It is defined by counting triangles with boundary conditions in $\mathbb{T}_\alpha, \mathbb{T}_{\alpha'}, \mathbb{T}_{\beta'}$.




Figure 3.9: Picture of the three different boundary conditions arising in our discussion.

Figure 3.9 illustrates the boundary conditions and how they look like near the region where the Dehn twist is performed. Analogous to the discussion in the proof of

Proposition 3.2.1 the picture shows that

$$\Gamma_{\alpha, \alpha'; \beta'} = \begin{pmatrix} \Gamma_{\alpha, \alpha'; \beta}^w & \bar{\Gamma} \\ \mathbf{0} & -\Gamma_{\alpha, \alpha'; \delta}^w \end{pmatrix},$$

where $\bar{\Gamma}$ is a map defined by counting triangles that connect $\mathbb{T}_{\alpha'} \cap \mathbb{T}_{\delta}$ -intersections with $\mathbb{T}_{\alpha} \cap \mathbb{T}_{\beta}$ -intersections. This immediately shows commutativity of the first two boxes, i.e.

$$\begin{aligned} \Gamma_{\alpha, \alpha'; \beta'}^* \circ \Gamma_1 &= \Gamma'_1 \circ \Gamma_{\alpha, \alpha'; \beta}^{w,*} \\ \Gamma'_2 \circ \Gamma_{\alpha, \alpha'; \beta'}^* &= -\Gamma_{\alpha, \alpha'; \delta}^{w,*} \circ \Gamma_1. \end{aligned}$$

It remains to show that

$$\Gamma_{\alpha, \alpha'; \beta}^{w,*} \circ \partial_* = \partial'_* \circ -\Gamma_{\alpha, \alpha'; \delta}^{w,*}.$$

Recall that ∂_* equals the map f in the definition of the boundary $\widehat{\partial}^\delta$. These were defined by counting discs with $n_* \neq 0$ or $n_{**} \neq 0$. Look at the following expression

$$\Gamma_{\alpha, \alpha'; \beta}^{w,*} \circ f_* + f'_* \circ \Gamma_{\alpha, \alpha'; \delta}^{w,*}.$$

The strategy to show its vanishing is analogous to the discussion of the chain map-property of f in the proof of Proposition 3.2.1. There are two ways to see this: Recall that $\Gamma_{\alpha, \alpha'; \beta'}$ is a chain map. Hence, with the representation of $\widehat{\partial}^\delta$ given in Proposition 3.2.1, this means that

$$f'_* \circ \Gamma_{\alpha, \alpha'; \delta}^w + \Gamma_{\alpha, \alpha'; \beta}^w \circ f = \widehat{\partial}_{\alpha' \beta}^w \circ \bar{\Gamma} + \bar{\Gamma} \circ \widehat{\partial}_{\alpha' \delta}^w. \quad (3.3.1)$$

Thus,

$$\begin{aligned} 0 &= (f'_* \circ \Gamma_{\alpha, \alpha'; \delta}^w + \Gamma_{\alpha, \alpha'; \beta}^w \circ f)_* \\ &= f'_* \circ \Gamma_{\alpha, \alpha'; \delta}^{w,*} + \Gamma_{\alpha, \alpha'; \beta}^{w,*} \circ f_* \end{aligned}$$

since all maps involved are chain maps. Hence, the third box commutes, too. Alternatively, look at the ends of the moduli spaces of Whitney triangles with boundary conditions in \mathbb{T}_α , $\mathbb{T}_{\alpha'}$, $\mathbb{T}_{\beta'}$ with Maslov index 1 and non-trivial intersection number n_* or n_{**} . The ends look like given in Figure 3.10. There are three possible ends. But observe that the top end (cf. Figure 3.10) corresponds to $\Gamma(x \otimes \widehat{\partial}^{\widehat{\Theta}^+})$, which vanishes since by definition $\widehat{\partial}^{\widehat{\Theta}^+} = 0$. Hence, for our situation there are just two possible types of ends to consider (the both at the bottom of Figure 3.10). Recall that breaking is the only phenomenon that appears here (cf. proof of Proposition 3.2.1 or see [40]). Proceeding as in the proof of Proposition 3.2.1, the commutativity of the third box follows. \square

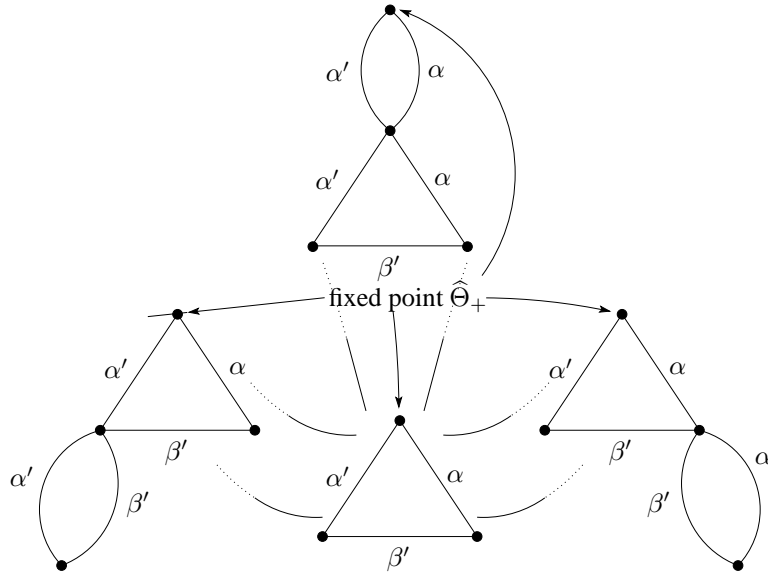


Figure 3.10: The moduli space has three possible ends. But only two of them count non-trivially, since $\widehat{\partial}\hat{\Theta}^+ = 0$.

Proposition 3.3.2. *Isotopies of the α -circles induce isomorphisms on the homologies such that all squares commute. Isotopies of the β -curves that miss the points w and z induce isomorphisms such that all squares commute.*

Proof. We realize isotopies of the attaching circles by Hamiltonian isotopies. Hence, the induced map Φ on homology is defined by counting discs with dynamic boundary conditions in the α -curves. The β -side remains untouched. Hence, by an analogous argument as in the proofs of Theorems 3.2.1 and 3.3.1 the map on homology splits into three components. The commutativity with Γ_1 and Γ_2 is then obviously true, and the only thing to show is the commutativity with the connecting homomorphism ∂_* and ∂'_* . But this again can be done by counting appropriate ends of moduli spaces or by looking into the chain map equation of Φ with respect to the representation of $\widehat{\partial}^\delta$. \square

Consider the following situation: Let $(\Sigma, \alpha, \beta, z)$ be a δ -suitable Heegaard diagram. With the discussion in §3.2.1 we obtain a long exact sequence

$$\dots \xrightarrow{\partial_*} \widehat{\text{HFK}}(\Sigma, \alpha, \beta, z, w) \xrightarrow{\Gamma_1} \widehat{\text{HF}}(\Sigma, \alpha, \beta', z) \xrightarrow{\Gamma_2} \widehat{\text{HFK}}(\Sigma, \alpha, \delta, z, w) \xrightarrow{\partial_*} \dots$$

where we define the attaching circles

$$\begin{aligned}\beta' &= \{\beta'_1, \beta_2, \dots, \beta_g\} \\ \delta &= \{\delta, \beta_2, \dots, \beta_g\}\end{aligned}$$

as it was done in §3.2.1. Define β'' by performing a handle slide among the β_i , $i \geq 2$, or by a handle slide of β'_1 over β_i . Perform the same operation on the set of attaching circles β to obtain $\tilde{\beta}$. Finally, take an isotopic push-off of δ , δ' say, that intersects δ in a cancelling pair of intersection points. Do the same with the β_i , $i \geq 2$, to get β'_i , $i \geq 2$. In this way we define another set of attaching circles δ' which is given by

$$\delta' = \{\delta', \beta'_2, \dots, \beta'_g\}.$$

Using these data we have the following result.

Proposition 3.3.3. *Let $(\Sigma, \alpha, \beta, z)$ be a δ -suitable Heegaard diagram and $(\Sigma, \alpha, \beta'', z)$ be obtained by a handle slide among the β_i , $i \geq 2$ or by a handle slide of β_1 over β_i . Denote by*

$$\begin{aligned}\Gamma_{\alpha; \beta, \tilde{\beta}}^w &: \widehat{\text{CFK}}(\Sigma, \alpha, \beta, z, w) \longrightarrow \widehat{\text{CFK}}(\Sigma, \alpha, \tilde{\beta}, z, w) \\ \Gamma_{\alpha; \delta, \delta'}^w &: \widehat{\text{CFK}}(\Sigma, \alpha, \delta, z, w) \longrightarrow \widehat{\text{CFK}}(\Sigma, \alpha, \delta', z, w) \\ \Gamma_{\alpha; \beta', \beta''} &: \widehat{\text{CF}}(\Sigma, \alpha, \beta', z) \longrightarrow \widehat{\text{CF}}(\Sigma, \alpha, \beta'', z)\end{aligned}$$

the induced maps. These maps induce a commutative diagram with exact rows

$$\begin{array}{ccccccc} \dots & \xrightarrow{\partial_*} & \widehat{\text{HFK}}(\Sigma, \alpha, \beta, z, w) & \xrightarrow{\Gamma_1} & \widehat{\text{HF}}(\Sigma, \alpha, \beta', z) & \xrightarrow{\Gamma_2} & \widehat{\text{HFK}}(\Sigma, \alpha, \delta, z, w) & \xrightarrow{\partial_*} & \dots \\ & & \Gamma_{\alpha; \beta, \tilde{\beta}}^{w,*} \downarrow & & \Gamma_{\alpha; \beta', \beta''}^* \downarrow & & \Gamma_{\alpha; \delta, \delta'}^{w,*} \downarrow & & \\ \dots & \xrightarrow{\partial'_*} & \widehat{\text{HFK}}(\Sigma, \alpha, \tilde{\beta}, z, w) & \xrightarrow{\Gamma'_1} & \widehat{\text{HF}}(\Sigma, \alpha, \beta'', z) & \xrightarrow{\Gamma'_2} & \widehat{\text{HFK}}(\Sigma, \alpha, \delta', z, w) & \xrightarrow{\partial'_*} & \dots \end{array}$$

Before going *in medias res*, we would like to explain our strategy. The idea behind all main proofs concerning the exact sequences was to show that certain holomorphic discs cannot exist. Up to this point we always used the base points w and z in the sense that we tried to see what implications can be made from the conditions $n_z = n_w = 0$. In addition, keeping in mind that holomorphic maps between manifolds of the same dimension are orientation preserving, we were able to prove everything we needed. Here, however, it is not so easy. First we would like to see that the map $\Gamma_{\alpha; \beta', \beta''}$ can be written as

$$\Gamma_{\alpha; \beta', \beta''} = \begin{pmatrix} \Gamma_1 & \bar{\Gamma} \\ 0 & \Gamma_2 \end{pmatrix}.$$

This means we would like to show that there are no triangles connecting $\alpha\beta$ -intersections of $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$ with $\alpha\delta'$ -intersections of $\mathbb{T}_\alpha \cap \mathbb{T}_{\beta''}$ (cf. Figure 3.11). This part is very similar to the proofs already given. We could try to continue in the same spirit and identify moduli spaces as we did before, but this is quite messy in this situation. The reason is that we are counting triangles, and being forced to make an intermediate stop at the point $\widehat{\Theta}$, we are able to *switch our direction* there. So, comparing the boundary conditions given in the three triple diagrams is not very convenient. Unfortunately we were not able to avoid these inconveniences completely, but could minimize them. After proving the splitting, we stick to $\Gamma_{\alpha;\beta',\beta''}$ and show that the maps $\Gamma_1, \Gamma_2, \bar{\Gamma}$ are chain maps and that all boxes in the diagram commute. This is realized by counting ends of appropriate moduli spaces of holomorphic triangles and squares. Finally, to minimize the messy task of comparing triangles in three diagrams, we just stick to Γ_1 and show that this map essentially equals $\Gamma_{\alpha;\beta,\tilde{\beta}}^w$ on the chain level. The 5-Lemma then ends the proof.




Figure 3.11: The important part of the Heegaard diagram after handle slide.

Proof. First observe that β_1' and β_1'' meet in two pairs of cancelling intersection points. Thus

$$\begin{aligned} \Gamma_{\alpha;\beta',\beta''} &= \widehat{f}_{\alpha\beta'\beta''}(\cdot \otimes \widehat{\Theta}) \\ &= \widehat{f}_{\alpha\beta'\beta''}(\cdot \otimes \{a_1, \theta_2, \dots, \theta_g\}) + \widehat{f}_{\alpha\beta'\beta''}(\cdot \otimes \{a_2, \theta_2, \dots, \theta_g\}). \end{aligned}$$

So, we are looking for triangles with intermediate intersection $\{a_1, \theta_2, \dots, \theta_g\}$ and triangles with intermediate intersection $\{a_2, \theta_2, \dots, \theta_g\}$.

Step 1– Splitting. Let $x \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ and $y \in \mathbb{T}_\alpha \cap \mathbb{T}_{\tilde{\beta}}$ be fixed. Let

$$\widehat{f}_{\alpha\beta'\beta''}(x \otimes \{a_1, \theta_2, \dots, \theta_g\}) \Big|_y$$

be the coefficient of $\widehat{f}_{\alpha\beta'\beta''}(x \cdot \otimes \{a_1, \theta_2, \dots, \theta_g\})$ at the generator y . Suppose there is a triangle starting at x and going to y along the α -boundary and then running to a_1 along its β' -boundary. From that point we have to go back to x again, following the red curve pictured in Figure 3.11. At a_1 we have two choices: we go upwards along the red curve, or we go downwards. Observe that going upwards, this would lead us to entering the \mathcal{D}_z -region at some point and force n_z to be non-zero in contradiction to our assumptions. Going downwards, we again enter the \mathcal{D}_z -region and the boundary conditions force n_z to be non-zero, again. Thus, there is no holomorphic triangle connecting x with y along a_1 . Thus

$$\widehat{f}_{\alpha\beta'\beta''}(x \otimes \{a_1, \theta_2, \dots, \theta_g\}) \Big|_y = 0.$$

The next step is to compute

$$\widehat{f}_{\alpha\beta'\beta''}(x \otimes \{a_2, \theta_2, \dots, \theta_g\}) \Big|_y.$$

Suppose there were a triangle that contributes. Going along the boundary of that triangle we would start at x and go to y along the α -boundary of the triangle and then try to go to a_2 following the pink curve in Figure 3.11. At some point we enter \mathcal{D}_z forcing n_z to be non-trivial. Hence, we have

$$\widehat{f}_{\alpha\beta'\beta''}(x \otimes \{a_2, \theta_2, \dots, \theta_g\}) \Big|_y = 0.$$

This shows that

$$\Gamma_{\alpha;\beta',\beta''} = \begin{pmatrix} \Gamma_1 & \bar{\Gamma} \\ 0 & \Gamma_2 \end{pmatrix}.$$

Step 2 – $\Gamma_1 = \Gamma_{\alpha;\beta,\tilde{\beta}}^w$. First of all it is easy to see that holomorphic triangles, contributing in $\Gamma_{\alpha;\beta,\tilde{\beta}}^w$, fulfill the property that $n_{y_1} = 0$. Hence, together with $n_w =$

$n_z = 0$ the triangles have to stay away from the regions surrounding $\beta \cap \delta$. Hence, we have

$$\Gamma_1 = \Gamma_{\alpha;\beta,\tilde{\beta}}^w + R.$$

The map R counts all holomorphic triangles not contributing to $\Gamma_{\alpha;\beta,\tilde{\beta}}^w$. Conversely, all holomorphic discs contributing to Γ_1 should be shown to fulfill $n_* = n_{**} = n_{y_1} = n_{y_2} = 0$. In this case $R = 0$ and both maps coincide on the chain level. Look at Figure 3.12: The situation for the $\alpha\beta\tilde{\beta}$ -diagram is pictured.

1. Observe that there is exactly one holomorphic triangle with $n_{**} \neq 0$. This triangle contributes to $\overline{\Gamma}$.
2. There is no holomorphic triangle contributing to Γ_1 with $n_* \neq 0$.
3. In a similar vein observe that these triangles in addition have trivial intersection with y_1 and y_2 .

Thus, we see that $R = 0$.




Figure 3.12: What happens.

Step 3 – Chain map properties and commutativity. Given points $x \in \mathbb{T}_\alpha \cap \mathbb{T}_\delta$ and $y \in \mathbb{T}_\alpha \cap \mathbb{T}_{\tilde{\beta}}$, look at the moduli space of holomorphic triangles connecting x with y , with Maslov index 1. There are, a priori, eight ends from which we just write down four. The four ends missing in Figure 3.13 are those contributing to $\Gamma(\cdot \otimes \partial\hat{\Theta})$, which

vanishes since $\partial\widehat{\Theta} = 0$. We know that $\Gamma_{\alpha;\beta',\beta''}$ is a chain map, i.e.

$$\begin{aligned} 0 &= \partial \circ \Gamma_{\alpha;\beta',\beta''} + \Gamma_{\alpha;\beta',\beta''} \circ \partial \\ &= \partial_{\alpha\tilde{\beta}}^w \circ \Gamma_1 + \Gamma_1 \circ \partial_{\alpha\beta}^w \\ &\quad + \partial_{\alpha\tilde{\beta}}^w \circ \bar{\Gamma} + f' \circ \Gamma_2 + \Gamma_1 \circ f + \bar{\Gamma} \circ \partial_{\alpha\delta}^w \\ &\quad + \partial_{\alpha\delta'}^w \circ \Gamma_2 + \Gamma_2 \circ \partial_{\alpha\delta}^w. \end{aligned}$$

The first two terms vanish since we identified Γ_1 with $\Gamma_{\alpha;\beta,\tilde{\beta}}^w$, which is a $(\partial_{\alpha\beta}^w, \partial_{\alpha\tilde{\beta}}^w)$ -chain map. The next four terms vanish since these correspond to the ends illustrated in Figure 3.13. Finally, since the whole equation is zero, the last two terms cancel each other. Thus, Γ_2 is a chain map as desired. By construction, two of three boxes in the diagram commute. We have to see that on the level of homology

$$\Gamma_1 \circ f = f' \circ \Gamma_2.$$

Recall we showed that on the chain level

$$\partial_{\alpha\tilde{\beta}}^w \circ \bar{\Gamma} + f' \circ \Gamma_2 + \Gamma_1 \circ f + \bar{\Gamma} \circ \partial_{\alpha\delta}^w = 0.$$

Hence, $\bar{\Gamma}$ is a chain homotopy between $\Gamma_1 \circ f$ and $f' \circ \Gamma_2$. □




Figure 3.13: The ends of the moduli space providing commutativity

In [27] the authors give an alternative proof for the independence of the contact element of the choice of cut system. We are especially interested in the technique they used to prove Proposition 3.3 of [27]. Recall, that given an open book (P, ϕ) , a **positive Giroux stabilization** of (P, ϕ) is the open book $(P \cup h^1, \phi \circ D_\gamma^+)$ where γ is a




Figure 3.14: Illustration of what happens while Giroux stabilizing.

closed curve in $P \cup h^1$ that intersects the co-core of h^1 once, transversely. Fixing a homologically essential, simple closed curve δ in P we call the Giroux stabilization δ -**elementary** if, after a suitable isotopy, δ intersects γ transversely in at most one point (cf. Definition 2.5. of [27]). Their invariance proof relies on the fact that, given a positive Giroux stabilization, one can choose a cut system a_1, \dots, a_n of (P, ϕ) such that the curve γ does not intersect any of the a_i . Observe that, given such a cut system for (P, ϕ) and defining a_{n+1} to be the co-core of the handle h^1 , then a_1, \dots, a_{n+1} is a cut system for the Giroux stabilized open book. Furthermore, observe that for $i \leq n$

$$\phi \circ D_\gamma^+(a_i) = \phi(a_i).$$

Figure 3.14 illustrates how $\phi \circ D_\gamma^+(\alpha_{n+1})$ looks like. Thus, all intersections between α_i and β_j for $i, j \leq n$ remain unchanged, where α_{n+1} intersects only β_{n+1} once, transversely. Furthermore, $D_\gamma^+(a_{n+1})$ is disjoint from all a_i , $i \leq n$. And, hence, β_{n+1} is disjoint from all α_i , $i \leq n$. Thus, the induced Heegaard diagram looks like a stabilized Heegaard diagram induced by the open book (P, ϕ) with cut system a_1, \dots, a_n . Denote by q the unique intersection point of α_{n+1} and β_{n+1} . Then the map

$$\Phi: \widehat{\text{CF}}(P, \phi, \{a_1, \dots, a_n\}) \longrightarrow \widehat{\text{CF}}(P \cup h^1, D_\gamma^+ \circ \phi, \{a_1, \dots, a_{n+1}\}),$$

given by sending a generator x of $\widehat{\text{CF}}(P, \phi, \{a_1, \dots, a_n\})$ to $\Phi(x) = (x, q)$, is clearly an isomorphism of chain complexes preserving the contact element.

We will, however, focus our attention on a special version of positive Giroux stabilization. Recall, that we call $(\Sigma \# T^2, \alpha', \beta')$ a stabilization of the Heegaard diagram (Σ, α, β) where we define $\alpha' = \alpha \cup \{\mu\}$ and $\beta' = \beta \cup \{\lambda\}$ with μ a meridian and λ a longitude of T^2 .




Figure 3.15: The choice of γ for a topological stabilization.

Definition 3.3.4. Let (P, ϕ) be an open book decomposition and let $(P', \phi \circ D_\gamma^+)$ be a positive Giroux stabilization. We say that the Giroux stabilization **represents a topological stabilization** if there is a cut system $\{a_1, \dots, a_n, a_{n+1}\}$ of P' with the following properties:

- (1) The set $\{a_1, \dots, a_n\}$ is a cut system for P .
- (2) Denote by (Σ, α, β) the Heegaard diagram induced by $(P, \phi, \{a_1, \dots, a_n\})$ and let $(\Sigma', \alpha', \beta')$ be the Heegaard diagram induced by $(P', \phi \circ D_\gamma^+, \{a_1, \dots, a_{n+1}\})$. The diagram $(\Sigma', \alpha', \beta')$ is a stabilization of (Σ, α, β) up to isotopy of the attaching circles.

Look into Figure 3.15. In this picture we present how to choose γ such that the positive Giroux stabilization represents a topological stabilization. Indeed, the following lemma holds.

Lemma 3.3.5. *Let (P, ϕ) be an open book decomposition and let $(P', \phi \circ D_\gamma^+)$ be a positive Giroux stabilization. The Giroux stabilization represents a topological stabilization up to isotopy of the attaching circles if and only if γ is isotopic to the black curve pictured in Figure 3.15.*

Proof. Given an open book decomposition (P, ϕ) and a positive Giroux stabilization $(P', \phi \circ D_\gamma^+)$ with γ like indicated in Figure 3.15, this stabilization clearly represents a topological stabilization up to isotopy: Recall that $P' = P \cup h^1$. Choose a cut system $\{a_1, \dots, a_n\}$ of P such that $\partial a_i, i = 1, \dots, n$, is disjoint from the region where the handle h^1 is attached on. Define a_{n+1} as the co-core of the handle h^1 . Picturing the resulting Heegaard diagrams we see that the positive Giroux stabilization represents a

topological stabilization up to isotopy.

Conversely, suppose we are given a Giroux stabilization representing a topological stabilization up to isotopy, then we have to show that γ is isotopic to the black curve, γ_s say, indicated in Figure 3.15. First note that the handle is attached on one boundary component of P . If h^1 connects two different boundary components of P , the genus of the resulting Heegaard surface would increase by 2. By assumption there is a cut system $\{a_1, \dots, a_{n+1}\}$ for P' fulfilling properties (1) and (2), given in Definition 3.3.4. As in Definition 3.3.4, denote by (Σ, α, β) and $(\Sigma', \alpha', \beta')$ the respective Heegaard diagrams. By assumption, $\Sigma' = \Sigma \# T^2$ and, after applying suitable isotopies, $\alpha_i = \alpha'_i$ and $\beta_i = \beta'_i$ for all $i = 1, \dots, n$. We have, that

$$\begin{aligned}\alpha'_{n+1} &= a_{n+1} \cup \overline{a_{n+1}} \\ \beta'_{n+1} &\sim a_{n+1} \cup \overline{\phi \circ D_\gamma^+(a_{n+1})}\end{aligned}$$

with

$$\alpha'_{n+1} \sim \mu_{T^2} \tag{3.3.2}$$

$$\beta'_{n+1} \sim \lambda_{T^2}. \tag{3.3.3}$$

By (3.3.2), we see that a_{n+1} is isotopic to the co-core of h^1 . This can be read off from Figure 3.16. Hence, we have

$$a_{n+1} \cup \overline{\phi \circ D_\gamma^+(a_{n+1})} = \beta_{n+1} \sim \lambda_{T^2} \sim a_{n+1} \cup \overline{\phi \circ D_{\gamma_s}^+(a_{n+1})}.$$

So, $\phi \circ D_\gamma^+(a_{n+1})$ is isotopic to $\phi \circ D_{\gamma_s}^+(a_{n+1})$, which is equivalent to saying that $D_\gamma(a_{n+1})$ is isotopic to $D_{\gamma_s}(a_{n+1})$. But this finally implies that γ is isotopic to γ_s . \square

Proposition 3.3.6. *Let (P, ϕ) be an open book decomposition of Y and $(P', \phi \circ D_\gamma^+)$ a positive δ -elementary Giroux stabilization representing a topological stabilization (cf. Definition 3.3.4 and look at Figure 3.15). Then there are isomorphisms ϕ_1, ϕ_2 and ϕ_3 on homology such that the following diagram commutes*

$$\begin{array}{ccccccc} \dots & \xrightarrow{\partial_*} & \widehat{\text{HFK}}(P, \phi, \delta) & \xrightarrow{\Gamma_1} & \widehat{\text{HF}}(P, D_\delta^+ \circ \phi) & \xrightarrow{\Gamma_2} & \widehat{\text{HFK}}(P, \tilde{\phi}) \xrightarrow{\partial_*} \dots \\ & & \phi_1 \downarrow \cong & & \phi_2 \downarrow \cong & & \phi_3 \downarrow \cong \\ \dots & \xrightarrow{\partial'_*} & \widehat{\text{HFK}}(P', \phi \circ D_\gamma^+, \delta) & \xrightarrow{\Gamma'_1} & \widehat{\text{HF}}(P', D_\delta^+ \circ \phi \circ D_\gamma^+, z) & \xrightarrow{\Gamma'_2} & \widehat{\text{HFK}}(P', \tilde{\phi} \circ D_\gamma^+) \xrightarrow{\partial'_*} \dots \end{array}$$

Remark. General positive Giroux stabilizations do not preserve the exact sequence. The reason is that in the general situation $\gamma \cap P$ and $\phi^{-1}(\delta)$ might intersect and cannot be separated. In the topological situation, however, the special choice of γ makes it possible to separate $\gamma \cap P$ from $\phi^{-1}(\delta)$.




Figure 3.16: The left portion pictures Σ' and the right portion the page P' and how it is obtained from P .

Proof. Denote by γ_1 the part of γ that runs through P . Since we are just doing a topological stabilization, we can attach the handle h^1 in such a way that γ_1 and $\phi^{-1}(\delta)$ are disjoint. Just choose γ like indicated in Figure 3.15. Even if $\phi^{-1}(\delta)$ intersects γ_1 , we can separate them with help of a small isotopy. By choosing a cut system $\{a_1, \dots, a_n\}$ for (P, ϕ) appropriately, we can extend this cut system to a cut system for the stabilized open book by choosing a_{n+1} like indicated in Figure 3.15. For all Heegaard diagrams in the following, we will use this cut system. Since $\phi^{-1}(\delta)$ and γ are disjoint, the associated Heegaard diagram of $(P', D_\delta^+ \circ \phi \circ D_\gamma^+)$ will look like a stabilization of the Heegaard diagram induced by the open book $(P, D_\delta^+ \circ \phi)$. The same holds for $(-P', \tilde{\phi})$ and $(-P', \tilde{\phi} \circ D_\gamma^+)$. Using the isomorphism induced by stabilizations as discussed above we can define ϕ_1 , ϕ_2 and ϕ_3 as indicated in Proposition 3.3.6. These maps are all isomorphisms and obviously commute on the chain level. \square

Theorem 3.3.7. *The map Γ_1 is topological, i.e. it just depends on the cobordism induced by the surgery.*

Proof. The cobordism induced by the Dehn twist depends only on the 3-manifold Y and the framed knot type K which the curve δ , together with its page framing, represents inside Y . This pair, on the other hand, is described by an open book decomposition adapted to δ and a δ -adapted cut system. These data determine a Heegaard diagram subordinate to the pair (Y, K) (cf. §2.4). Given another adapted open book together

with an adapted cut system, the associated Heegaard diagram is equivalent to the first after a sequence of moves which are described in Lemma 2.4.3. All of these moves are recovered via Proposition 3.3.1, Proposition 3.3.2, Proposition 3.3.3 and Proposition 3.3.6. Of course, after some point, we might leave the class of Heegaard diagrams induced by open books. But the propositions cited do not use this open book structure as discussed at the beginning of the section. \square

3.4 Implications to Contact Geometry

In this section we will focus our attention on contact manifolds (Y, ξ) . Let (P, ϕ) be an open book decomposition that is adapted to the contact structure ξ (cf. §2.7.3). Recall that the contact element and the invariant defined in [27] sit in the Heegaard Floer cohomology (cf. §2.7.4 and Lemma 2.7.15). Because of the well-known equivalence

$$\widehat{\text{HF}}^*(Y) = \widehat{\text{HF}}_*(-Y)$$

we will be interested in the behavior of $-Y$ rather than Y (cf. Lemma 2.7.15). Recall from §2.7.4 that we have two choices to extract the Heegaard Floer homology of $-Y$ from data given by a Heegaard diagram of Y . We can either switch the orientation of the Heegaard surface or switch the boundary conditions.

Let $L \subset Y$ be a Legendrian knot (cf. §2.7.1) and denote by Y_L^+ the manifold obtained by doing a $(+1)$ -contact surgery along L . There is an open book decomposition (P, ϕ) adapted to ξ such that L sits on the page $P \times \{1/2\}$ of the open book and the page framing coincides with the contact framing. A $(+1)$ -contact surgery acts on the open book like a negative Dehn twist along L , i.e. $(P, \phi \circ D_L^{-,P})$ is an adapted open book decomposition of (Y_L^+, ξ_L^+) where $D_L^{-,P}$ denotes a negative Dehn twist along L with respect to the orientation of P . Observe that L sits on the wrong page for our construction of the exact sequence. Fortunately, the identity

$$\phi \circ D_L^{-,P} = D_{\phi(L)}^{-,P} \circ \phi \tag{3.4.1}$$

holds. Thus, a surgery along L can be interpreted as a left-hand composition of the monodromy with a Dehn twist. In addition $(P, D_{\phi(L)}^{-,P} \circ \phi)$ is an adapted open book decomposition of (Y_L^+, ξ_L^+) . To see the effect on the Heegaard Floer cohomology, we have to change the surface orientation. We see that

$$-Y_L^+ = (-P, D_{\phi(L)}^{-,P} \circ \phi) = (-P, D_{\phi(L)}^{+,-P} \circ \phi). \tag{3.4.2}$$

One very important ingredient for our construction is the fact that we may choose an L -adapted Heegaard diagram where L sits on $P \times \{1/2\}$. Because of the identity (3.4.1) we need a Heegaard diagram with attaching circles adapted to $\phi(L)$ in the following sense: the curve $\phi(L)$ intersects β_1 once, transversely and is disjoint from all other β -circles. This condition is satisfied for L -adapted Heegaard diagrams since $\phi(a_i) = b_i$. This means we are able to simultaneously match all conditions for setting up the exact sequence and seeing the invariant $\widehat{\mathcal{L}}(L)$. Recall that the sequence requires the point w defining L to be in a specific domain of the Heegaard diagram. This positioning of w induces an orientation on L . On the other hand, a fixed orientation of L determines where w has to be placed. These two orientations, the one coming from the sequence and the one from the knot L itself, have to be observed carefully. We have to see whether every possible choice of orientation of L induces a positioning of w inside the Heegaard diagram that is compatible with the requirements coming from the exact sequence.

Proposition 3.4.1. *Let (Y, ξ) be a contact manifold and $L \subset Y$ an oriented Legendrian knot.*

- (i) *Let W be the cobordism induced by $(+1)$ -contact surgery along L . Then the cobordism $-W$ induces a map*

$$\Gamma_{-W}: \widehat{\text{HF}}\text{K}(-Y, L) \longrightarrow \widehat{\text{HF}}(-Y_L^+),$$

such that $\Gamma_{-W}(\widehat{\mathcal{L}}(L)) = c(Y_L^+, \xi_L^+)$.

- (ii) *If L carries a specific orientation and W denotes the cobordism induced by a (-1) -contact surgery along L . Then the cobordism $-W$ induces a map*

$$\Gamma_{-W}: \widehat{\text{HF}}(-Y_L^-) \longrightarrow \widehat{\text{HF}}\text{K}(-Y, L)$$

such that $\Gamma_{-W}(c(Y_L^-, \xi_L^-)) = 0$.

Proof. Recall that

$$\begin{aligned} -Y_L^+ &= (-P, D_{\phi(L)}^{+, -P} \circ \phi) \\ -Y_L^- &= (-P, D_{\phi(L)}^{-, -P} \circ \phi). \end{aligned}$$

We choose a cut system which is L -adapted. This means that L intersects α_1 transversely, in a single point and is disjoint from the other α -circles. Hence, $\phi(L)$ (sitting

on the other side of the Heegaard surface) intersects β_1 in a single point and is disjoint from the other β -circles. We first try to prove the results concerning the (+1)-contact surgery. After possibly isotoping the knot L slightly, we can achieve a neighborhood of $\phi(L) \cap \beta_1$ to look like the left or right part of Figure 3.17.




Figure 3.17: Setting things up for a contact (+1)-surgery.

In each part of the picture the knot L and the point w are placed in such a way that the Dehn twist associated to the (+1)-contact surgery connects the regions where the points w and z lie. Thus, each picture shows a situation in which we may apply the proof technique used for Proposition 3.2.1 (resp. Proposition 3.2.5). Observe that Figure 3.17 shows the situation for each orientation of L . Since we are doing a (+1)-contact surgery, we perform a positive Dehn twist along L with respect to the surface orientation given in Figure 3.17 (cf. Equality (3.4.2) and cf. discussion at the beginning of this paragraph). Thus, we are able to define a map

$$\Gamma^+ : \widehat{\text{HFK}}(-Y, L) \longrightarrow \widehat{\text{HF}}(-Y_L^+).$$

The situations in both pictures are designed to apply the proof technique of Proposition 3.2.1. The induced pair (w, z) determines an orientation on L . To match the induced orientation with the one of the knot L we either use the left or the right picture of Figure 3.17. By definition of Γ^+ we see that

$$\Gamma^+(\widehat{\mathcal{L}}(L)) = c(Y_L^+, \xi_L^+).$$

To cover (-1)-contact surgeries, look at Figure 3.18.

The same line of arguments as above applies to define a map

$$\Gamma^- : \widehat{\text{HF}}(-Y_L^-) \longrightarrow \widehat{\text{HFK}}(-Y, L).$$




Figure 3.18: Setting things up for a contact (-1) -surgery.

Again, recall that w is placed in the Heegaard diagram in such a way that allows us to define the map Γ^- . The pair (w, z) induces an orientation on L . The opposite orientation will be denoted by ob . What can be seen immediately from the picture is that the Dehn twist separates the contact element and the invariant $\widehat{\mathcal{L}}(L, \overline{ob})$: The arguments show that we have the following exact sequence.

$$\begin{array}{ccccccc}
 0 & \longrightarrow & \widehat{\text{CFK}}(Y_0(L), \mu) & \longrightarrow & \widehat{\text{CF}}(-Y_L^-) & \xrightarrow{\Gamma^-} & \widehat{\text{CFK}}(-Y, (L, \overline{ob})) \longrightarrow 0 \\
 & & \bullet & \longmapsto & c & & \\
 & & & & & & \bullet & \longmapsto & \widehat{\mathcal{L}}(L, \overline{ob})
 \end{array}$$

To speak in the language of the proof of Proposition 3.2.1: the element c is an $\alpha\beta$ -intersection, whereas the element $\widehat{\mathcal{L}}(L, \overline{ob})$ is an $\alpha\delta$ -intersection. By exactness, the contact element c lies in the kernel of Γ^- . \square

Definition 3.4.2. The orientation $ob(P, \phi)$ from the last proof is called the **open book orientation**.

To prove Corollary 3.4.3 we have to recall that Honda, Kazez and Matić introduced in [21] an invariant $EH(L)$ of a Legendrian knot L in the Sutured Floer homology (cf. [22]) of a contact manifold with boundary. To be more precise, given $L \subset (Y, \xi)$, they define an Legendrian isotopy invariant of L , called $EH(L)$, sitting in $\text{SFH}(-Y \setminus \nu L, \Gamma)$ where Γ are suitably chosen sutures. Furthermore, Stipsicz and Vertesi have shown in [48] that this invariant is equipped with a morphism $\text{SFH}(-Y \setminus \nu L, \Gamma) \longrightarrow \widehat{\text{HF}}(-Y, L)$ that maps $EH(L)$ to $\widehat{\mathcal{L}}(L)$. Composing this morphism with the one coming from Theorem 3.4.1 we get the following result.

Corollary 3.4.3. *There is a map*

$$\gamma: \text{SFH}(-Y \setminus \nu L, \Gamma) \longrightarrow \widehat{\text{HF}}(-Y_L^+)$$

such that $\gamma(EH(L)) = c(Y_L^+, \xi_L^+)$. □

Corollary 3.4.4. *Let L be a Legendrian knot in a contact manifold (Y, ξ) . Then $EH(L) = 0$ implies that $c(Y_L^+, \xi_L^+) = 0$.* □

It is also possible to derive these corollaries using methods coming from [48].

Proposition 3.4.5. *Let L be a Legendrian knot in a contact manifold (Y, ξ) carrying the open book orientation induced by an adapted open book (P, ϕ) . Let (P', ϕ') be the once-stabilized open book that carries the Legendrian knot $S_+(L)$ (see Proposition 3.4.11). The open book orientation $ob(P', \phi')$ coincides with the orientation included by the stabilization.*

We will give a proof of Proposition 3.4.5 in the following paragraph.

3.4.1 Stabilizations of Legendrian Knots and Open Books

Stabilizations as Legendrian Band Sums

Recall that stabilization basically means to enter a zigzag into the front projection of a Legendrian knot. If we are not in the standard contact space, we perform this operation inside a Darboux chart. Which zigzag is regarded as a positive/negative stabilization depends on the knot orientation. Positivity/Negativity is fixed by the following equations

$$\begin{aligned} tb(S_{\pm}(L)) &= tb(K) - 1 \\ rot(S_{\pm}(L)) &= rot(L) \pm 1. \end{aligned}$$

This tells us that

$$\overline{S_+(L)} = S_-(\overline{L}). \tag{3.4.3}$$

Given two Legendrian knots L and L' , we can form their **Legendrian band sum** $L\#_{L_b}L'$ in the following way: Pick a contact surgery representation of the contact manifold in such a way that the surgery link \mathbb{L} stays away from $L \cup L'$. In this way we can think of L and L' as sitting in the standard contact space and, so, can perform the band sum. We denote by L_0 and \overline{L}_0 the oriented Legendrian shark with the orientations as indicated in Figure 3.19.




Figure 3.19: The oriented Legendrian shark and its inverse.

Proposition 3.4.6. *Given a Legendrian knot L , we can realize its stabilizations as Legendrian band sums, i.e.*

$$\begin{aligned} S_+(L) &= L \#_{Lb} L_0 \\ S_-(L) &= L \#_{Lb} \overline{L_0}, \end{aligned}$$

where $\#_{Lb}$ denotes the Legendrian band-sum.

Proof. We prove the equality for positive stabilizations. The case of negative stabilizations is proved in a similar fashion. No matter what orientation the knot L carries, we will find at least one right up-cusp or one right down-cusp. In case of a right down-cusp we perform a band-sum involving this right down-cusp on L and the left up-cusp on L_0 . In case we use a right up-cusp we perform the band-sum as indicated in the left part of Figure 3.20. In Figure 3.20 we indicate the Legendrian isotopy that illustrates that we have stabilized positively. \square




Figure 3.20: The Legendrian band-sum in case of a right up-cusp and a Legendrian isotopy.

Open Books and Connected Sums

Suppose we are given open books (P_1, ϕ_1) and (P_2, ϕ_2) for manifolds (Y_1, ξ_1) and (Y_2, ξ_2) . Let B_1 be the binding of (P_1, ϕ_1) . Denote by νB_1 an equivariant tubular neighborhood of B_1 . Fix a point p on B_1 and embed a 3-ball D^3 such that it is

centered at p . Furthermore, the ball should sit inside νB_1 such that the north and south pole of D^3 equal $B_1 \cap \mathbb{S}^2$. Denote by $f_1: D^3 \rightarrow \nu B_1 \subset Y_1$ the embedding. Embed $g: D^3 \rightarrow Y_2$ in the same fashion. Compose g with a right-handed rotation r that swaps the two hemispheres of D^3 to get another embedding $f_2 = g \circ r$. Use these embeddings to perform the connected sum. By its definition, the gluing $f_2 \circ f_1^{-1}$ preserves the open book structure. Note that the rotation is needed to make the pages of the open book glue together nicely with their given orientation. Moreover, we are able to explicitly describe the resulting open book. The new page P equals $P_1 \cup_{h^1} P_2$, where h^1 is a 1-handle connecting P_1 and P_2 and the binding B equals $B_1 \# B_2$. To define the monodromy, first extend ϕ_1 and ϕ_2 as the identity along the handle and the complementary page. Then define ϕ as the composition $\phi_2 \circ \phi_1 = \phi_1 \circ \phi_2$.

Lemma 3.4.7. *The open book (P, ϕ) is an adapted open book for $(Y_1 \# Y_2, \xi_1 \# \xi_2)$.*

Proof. Observe that the given operation is a special case of the Murasugi sum. The lemma then follows from [9]. \square

Corollary 3.4.8. *Let (Y, ξ) and (Y', ξ') be contact manifolds and $L \subset Y$ a Legendrian knot. Then we have*

$$\begin{aligned} \widehat{\text{HFK}}(-Y \# Y', L) &\cong \widehat{\text{HFK}}(-Y, L) \otimes \widehat{\text{HF}}(-Y') \\ \widehat{\mathcal{L}}(Y \# Y', L) &= \widehat{\mathcal{L}}(Y, L) \otimes c(\xi') \end{aligned}$$

Proof. Let (P_1, ϕ_1) be an open book decomposition adapted to the knot L and the contact structure ξ . Denote by (P_2, ϕ_2) an open book for (Y', ξ') . We define an open book (P, ϕ) by using the open books for Y and Y' as given above. Recall, that the page P is given by joining the pages P_1 and P_2 with a 1-handle h^1 , i.e.

$$P = P_1 \cup_{h^1} P_2.$$

Denote by $f: \partial h^1 \rightarrow \partial P_1 \sqcup \partial P_2$ the attaching map. Furthermore, let $\{a_1, \dots, a_n\}$ be a cut system for P_1 and $\{a'_1, \dots, a'_m\}$ a cut system for P_2 . Choose isotopic push-offs b_i of the a_i so that a_i and b_i intersect each other in a pair x_i^+, x_i^- of intersection points. The push-offs are chosen like specified in §2.7.4 (cf. also Figure 2.10). Analogously, the curves $b'_j, j = 1, \dots, m$, are defined; denote the points of intersection by $y_j^+, y_j^-, j = 1, \dots, m$. The names are attached to the intersection points in such a way that $\{x_1^+, \dots, x_n^+\}$ represents the class $\widehat{\mathcal{L}}(Y, L)$ and that $\{y_1^+, \dots, y_m^+\}$ represents $c(\xi')$. We additionally fix base points $z_i \in P_i, i = 1, 2$, and a third one, w say, in P_1 determining the knot L . These choices induce Heegaard diagrams we denote by $(\Sigma_i, \alpha_i, \beta_i), i = 1, 2$. We require the chosen cut systems to fulfill the following two conditions:

$$(1) \quad \text{Im}(f) \cap \left(\bigcup_{i=1}^n \partial a_i \cup \bigcup_{j=1}^m \partial a'_j \right) = \emptyset$$

$$(2) \quad \text{Im}(f) \subset \partial \overline{\mathcal{D}_{z_1}} \cup \partial \overline{\mathcal{D}_{z_2}}$$

As a consequence of these two conditions and the fact that by definition $\phi|_{P_i} = \phi_i$, $i = 1, 2$ and $\phi|_{h^1} = \text{id}_{h^1}$ we see that

$$\phi(a_i) \cap a'_j = \emptyset \quad \text{and} \quad a_i \cap \phi(a'_j) = \emptyset. \quad (3.4.4)$$

The set $\{a_1, \dots, a_n\} \cup \{a'_1, \dots, a'_m\}$ is a cut system for the open book (P, ϕ) . Denote by (Σ, α, β) the induced Heegaard diagram, then with (3.4.4), we see that

$$\Sigma = \Sigma_1 \# \Sigma_2, \quad \alpha = \alpha_1 \cup \alpha_2, \quad \beta = \beta_1 \cup \beta_2$$

and the points z_i , $i = 1, 2$, lie in the regions unified by the connected sum tube. Choose a base point $z \in \Sigma$ lying in this unified region. Thus, — with the same reasoning as in the proof of [39], Proposition 6.1. — we see that

$$\widehat{\text{HFK}}(-Y \# Y, L) \cong \widehat{\text{HFK}}(-Y, L) \otimes \widehat{\text{HF}}(-Y). \quad (3.4.5)$$

By construction, the intersection point $\{x_1^+, \dots, x_n^+, y_1^+, \dots, y_m^+\}$ represents the class $\widehat{\mathcal{L}}(Y \# Y', L)$. But the isomorphism giving (3.4.5), φ say, has that property that

$$\{x_1^+, \dots, x_n^+, y_1^+, \dots, y_m^+\} \longmapsto \{x_1^+, \dots, x_n^+\} \otimes \{y_1^+, \dots, y_m^+\},$$

i.e. $\varphi(\widehat{\mathcal{L}}(-Y \# Y', L)) = \widehat{\mathcal{L}}(-Y, L) \otimes c(\xi')$. □

Lemma 3.4.9. ([9]) *If γ is a non-separating curve on a page of an open book (P, ϕ) , we can isotope the open book slightly such that γ is Legendrian and the contact framing agrees with the page framing.*

This fact follows from the Legendrian realization principle. As a consequence, we get the following corollary.

Corollary 3.4.10. *If the Legendrian knots $L_i \subset P_i$ sit on the ages, then, on the page P of (P, ϕ) , we will find a Legendrian knot L with the following property: There is a naturally induced contactomorphism ϕ_c such that $\phi_c(L)$ equals $L_1 \#_{Lb} L_2$ after performing a right-handed twist along the Legendrian band. Indeed, we obtain L by a band sum of L_1 and L_2 on the page P .*

Proof. Let (P_i, ϕ_i) be open books adapted to (Y_i, ξ_i, L_i) , $i = 1, 2$. On P_i there is a set of embedded, simple closed curves c_1^i, \dots, c_n^i whose associated Dehn twists generate the mapping class groups of P_i . The associated Dehn twists can be interpreted as contact surgeries along suitable Legendrian knots (cf. Theorem 2.7 in [27]). Thus, using the open book decomposition we are able to find a (maybe very inefficient) contact surgery representation of (Y_i, ξ_i) which is suitable for our purposes to perform the Legendrian band sum (cf. beginning of this section). Moreover, we can think of L_1 to pass the binding B_1 of P_1 very closely at some point: this means that there is a point p_1 in the binding, and a Darboux ball D_1 around p_1 , such that the curve intersects this Darboux ball. Suppose this is not the case, then we can isotope the Legendrian knot L_1 , which sits on P_1 , as a curve in P_1 , to pass the binding closely (as described above). The isotopy is not necessarily a Legendrian isotopy. However, by Theorem 2.7 of [27], we know that the isotoped curve determines a uniquely defined Legendrian knot, which is Legendrian isotopic to L_1 . With a slight isotopy of the open book, we can think of this new knot as sitting on P_1 . By abuse of notation, we call the new knot L_1 . After possibly isotoping the open book we can think of \mathbb{L}_1 as sitting in the complement of D_1 . We obtain a situation like indicated in the top row of Figure 3.21. Since we have the identification $(Y_1, \xi_1) \cong (\mathbb{S}^3(\mathbb{L}_1), \xi_{\mathbb{L}_1})$, the ball D_1 can be thought of as sitting in \mathbb{S}^3 . The complement of D_1 in \mathbb{S}^3 is again a ball we denote by \widetilde{D}_1 . We may make similar arrangements for L_2 : however, we would like L_2 and the associated surgery link \mathbb{L}_2 to sit inside D_1 and \widetilde{D}_1 to be the ball in which L_2 comes close to B_2 (cf. bottom row of Figure 3.21). We can form the connected sum

$$\mathbb{S}^3(\mathbb{L}_1 \sqcup \mathbb{L}_2) = \mathbb{S}^3(\mathbb{L}_1) \setminus D_1 \cup_{\partial} \mathbb{S}^2 \times [0, 1] \cup_{\partial} \mathbb{S}^3(\mathbb{L}_2) \setminus \widetilde{D}_1 \quad (3.4.6)$$

where the gluing is determined by the naturally given embeddings (cf. §4.12 in [16])

$$\iota_1: D_1 \hookrightarrow \mathbb{S}^3 \quad \text{and} \quad \iota_2: \widetilde{D}_1 \hookrightarrow \mathbb{S}^3.$$

For a detailed discussion of connected sums of contact manifolds we point the reader to [16]. The induced contact structure is the connected sum $\xi_{\mathbb{L}_1} \# \xi_{\mathbb{L}_2} = \xi_{\mathbb{L}_1 \sqcup \mathbb{L}_2}$ (cf. §4.12 of [16]). The knots L_1 and L_2 are contained in this connected sum and, here, we can perform the Legendrian band sum as defined at the beginning of this section; we can perform a band sum which looks like given in Figure 3.22. Recall that we introduced a connected sum operation such that the open books (P_i, ϕ_i) glue together to give the open book (P, ϕ) where $P = P_1 \cup_{h_1} P_2$ and ϕ is given as the composition of the two monodromies ϕ_1 and ϕ_2 . To perform the connected sum operation such that the open book structures are preserved, we have to modify the construction slightly. We modify the inclusion ι_1 by composing it with a rotation about the y -axis with angle π . Without

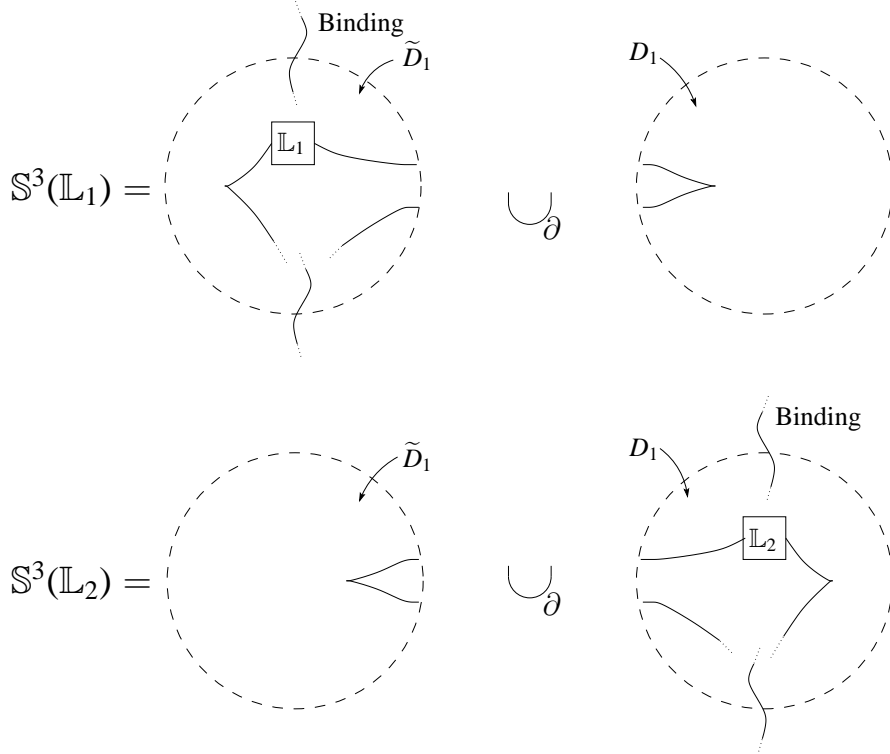


Figure 3.21: Our specific arrangement for performing the connected sum.

loss of generality we can think $L_1 \cap \partial D_1$ and $L_2 \cap \partial \tilde{D}_1$ to be identified by the gluing induced by the inclusion maps ι_1 and ι_2 . We can also assume that the rotation r swaps the two intersection points $L_1 \cap \partial D_1$. We obtain a new gluing map, f say, and get

$$Y = \mathbb{S}^3(\mathbb{L}_1) \setminus D_1 \cup_f \mathbb{S}^3(\mathbb{L}_2) \setminus \tilde{D}_1$$

with induced contact structure ξ . With this identification the knots L_1 and L_2 glue together to give a knot L . This knot L corresponds to a band sum of L_1 and L_2 on the page P (after possibly applying Proposition 3.4.9). Recall that contact structures on $\mathbb{S}^2 \times [0, 1]$ are uniquely determined, up to isotopy, by the characteristic foliations on $\mathbb{S}^2 \times \{j\}$, $j = 0, 1$ (cf. Lemma 4.12.1 and Theorem 4.9.4 of [16]). Consider the connected sum tube used in (3.4.6), and extend it with small collar neighborhoods of the boundaries of $\mathbb{S}^3(\mathbb{L}_1) \setminus D_1$ and $\mathbb{S}^3(\mathbb{L}_2) \setminus \tilde{D}_1$. The characteristic foliation $\xi_{\mathbb{L}_1 \sqcup \mathbb{L}_2}$ induces at the boundary will coincide with the characteristic foliation ξ induces on a suitably chosen tubular neighborhood of $\partial D_1 \cong \mathbb{S}^2 \times [0, 1]$ in Y . Thus, there is a contactomorphism between νD_1 and this thickened connected sum tube. Moreover,




Figure 3.22: Performing a band sum of L_1 and L_2 inside $\mathbb{S}^3(\mathbb{L}_1 \sqcup \mathbb{L}_2)$.

the contactomorphism can be extended to a contactomorphism

$$\phi_c: (Y, \xi) \longrightarrow (\mathbb{S}^3(\mathbb{L}_1 \sqcup \mathbb{L}_2), \xi_{\mathbb{L}_1 \sqcup \mathbb{L}_2})$$

which just affects the connected sum tube and fixes the rest. As one can derive with some effort, this contactomorphism basically rotates the \mathbb{S}^2 -factor once while going through the handle $\mathbb{S}^2 \times [0, 1]$. Thus, $\phi_c(L)$ looks like a band sum $L_1 \#_{Lb} L_2$ in $\mathbb{S}^3(\mathbb{L}_1 \sqcup \mathbb{L}_2)$ after twisting the band once. Figure 3.23 applies. \square

The following statement is due to Etnyre. Since there is no proof in the literature, we include a proof here for the convenience of the reader.

Proposition 3.4.11. ([9]) *Let (Y, ξ, L) be a contact manifold with Legendrian knot and (P, ϕ) and open book adapted to ξ with L on its page such that the page framing and contact framing coincide. By stabilizing the open book once we can arrange either the stabilized knot $S_+(L)$ or $S_-(L)$ to sit on the page of the stabilized open book as indicated in Figure 3.24.*

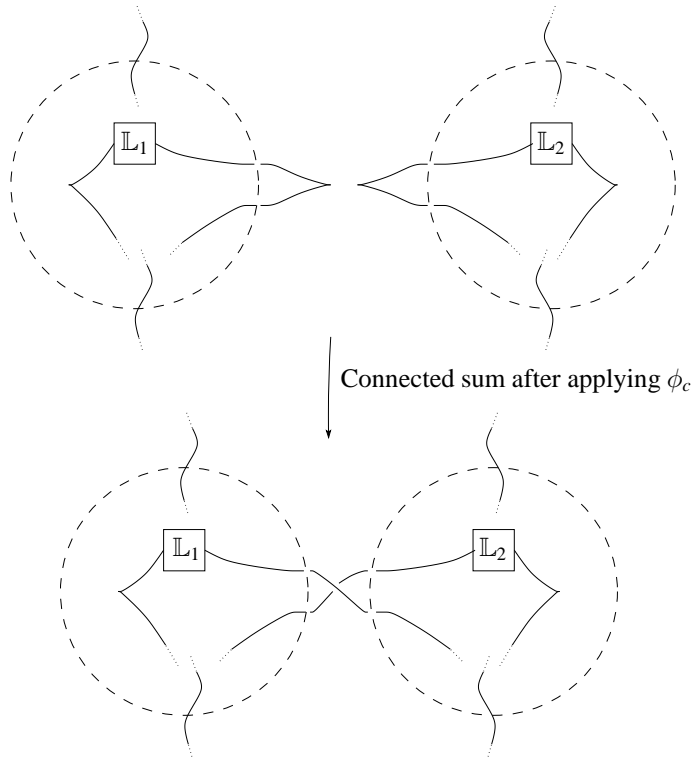


Figure 3.23: Schematic picture of the band braid after identifying (Y, ξ) with $(\mathbb{S}^3(\mathbb{L}_1 \sqcup \mathbb{L}_2), \xi_{\mathbb{L}_1 \sqcup \mathbb{L}_2})$.




Figure 3.24: The stabilized open book and a positive Legendrian stabilization.

The following result concerning the vanishing of the Legendrian invariant under positive stabilizations is due to Lisca, Ozsváth, Stipsicz and Szabó and follows from their connected sum formula given in [27]. Their proof carries over verbatim even for knots which are homologically non-trivial. Here we reprove a special case of Theorem 7.2.

of [27] using different methods.

Proposition 3.4.12 ([27], Theorem 7.2). *Given any Legendrian knot L in a contact manifold (Y, ξ) , we have $\widehat{\mathcal{L}}(S_+(L)) = 0$.*

Proof. Let (P, ϕ) be an open book decomposition adapted to (Y, ξ, L) . By Proposition 3.4.11 we know that a stabilized open book (P', ϕ') carries the stabilized knot $S_+(L)$. Furthermore, from Figures 3.24 and 3.30 we can see how the induced Heegaard diagram (adapted to capturing the contact geometric information) will look like near the base point w . This is done in Figure 3.25. We may use Proposition 3.4.5 to check that the




Figure 3.25: Parts of the Heegaard diagram induced by the open book carrying the stabilized knot.

positioning of the point w in Figure 3.25 is correct. First observe that $\widehat{\mathcal{L}}(S_+(L))$ is the homology class induced by the point

$$\{x_1, x_2, x_3, \dots, x_{2g}\}.$$

Recall that by definition of the points x_i every holomorphic disc emanating from x_i is constant. Thus, a holomorphic disc emanating from $Q := \{p, q, x_3, \dots, x_{2g}\}$ can only be non-constant at p, q . By orientation reasons and the placement of w the shaded region is the only region starting at p, q which can carry a holomorphic disc. Since it is disc-shaped, it does carry a holomorphic disc. Hence

$$\widehat{\partial}^w Q = \{x_1, x_2, x_3, \dots, x_{2g}\}$$

showing that $\widehat{\mathcal{L}}(S_+(L))$ vanishes. □




Figure 3.26: The open book necessary to carry the Legendrian unknot with $tb = -1$ and $rot = 0$.

Proof of Proposition 3.4.11. Given a triple (Y, ξ, L) , there is an open book (P, ϕ) adapted to ξ such that L sits on a page of the open book. By Proposition 3.4.6, Lemma 3.4.7 and Corollary 3.4.10 we perform a connected sum $(Y, \xi) \# (\mathbb{S}^3, \xi_{std})$ on the level of open books using the open book of $(\mathbb{S}^3, \xi_{std})$ pictured in Figure 3.26. By construction, the new open book carries the Legendrian knot L_2 pictured in Figure 3.27. In Figure 3.28 an isotopy is given, showing that L_2 corresponds to the band sum




Figure 3.27: The knot L_2 in $(Y, \xi) \# (\mathbb{S}^3, \xi_{std})$.

$L \#_{Lb} L_0$ and, thus, represents $S_{\pm}(L)$.

By Figure 3.26 what happens on the level of open books can be pictured as in Figure 3.29. \square

Proof of Proposition 3.4.5. Using Proposition 3.4.11, we have a tool to compare the open book orientation before and after the stabilization. We start with an open book adapted to the triple (Y, ξ, L) and choose an L -adapted cut system. By Proposition 3.4.11 we can generate an open book adapted to the positive stabilization by stabilizing the open book. Doing this appropriately, we may extend the cut system to an adapted cut system of the stabilized open book as indicated in Figure 3.30. Recall the rule with which the knot orientation is determined by the points (w, z) (see remark in §2.7.5). In




Figure 3.28: Legendrian isotopy showing that L_2 corresponds to the Legendrian band sum of L with the Legendrian shark L_0 .

Figure 3.30 we can now compare the open book orientation of the stabilized knot with the orientation induced by the stabilization. We see that the orientations coincide. \square

3.5 Applications – Vanishing Results of the Contact Element

In this paragraph we want to derive some applications of the theory developed in §3.2, §3.3 and §3.4. First to mention would be Proposition 3.5.1, which can also be

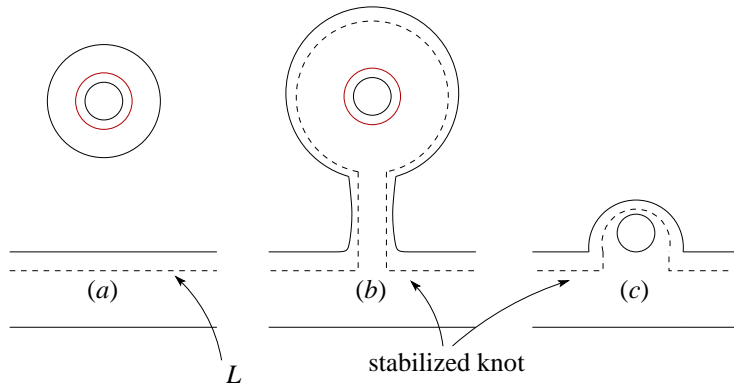


Figure 3.29: What happens during stabilization.

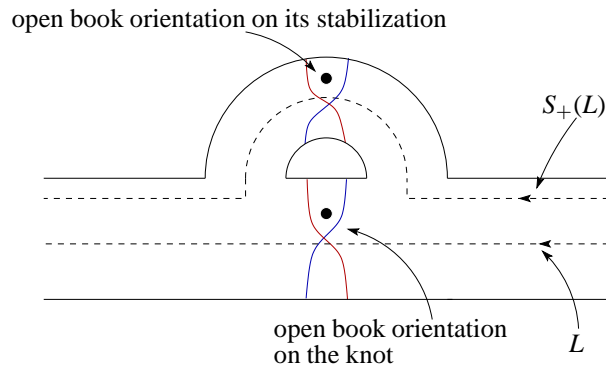


Figure 3.30: Comparing induced with open book orientation.

derived using methods developed in [29]. There, Lisca and Stipsicz show that $(+1)$ -contact surgery along stabilized Legendrian knots yield overtwisted contact manifolds, which implies the vanishing of the contact element. A second application would be Proposition 3.5.3, which is meant as a demonstration that calculating the Legendrian knot invariant and using Proposition 3.4.1 to get information about a contact element under investigation can be more convenient than using other methods, since the knot Floer homologies have additional structures we may use. A third application would be Theorem 3.5.4 which is a vanishing result of the contact element which can be easily read off from a surgery representation. This application uses the knot Floer homology for arbitrary knots and makes use of a phenomenon that seems to be special about these, namely that there are knots for which the knot Floer homology vanishes. We do not know any other example with this property.

Proposition 3.5.1. *If (Y, ξ) is obtained from (Y', ξ') by $(+1)$ -contact surgery along a Legendrian knot L which can be destabilized, the element $c(\xi)$ vanishes.*

Proof. There are two cases to cover. Give the knot L an orientation o . Suppose that

$$(L, o) = S_+(L', o').$$

Then Proposition 3.4.12 shows the vanishing of $\widehat{\mathcal{L}}(L, o)$. By Proposition 3.4.1 the element $c(\xi)$ vanishes, too. Now assume that

$$(L, o) = S_-(L', o').$$

We see that

$$(L, \bar{o}) = \overline{S_-(L', o')} = S_+(L', \bar{o}),$$

hence, $\widehat{\mathcal{L}}(L, \bar{o}) = 0$. By Proposition 3.4.1 again $c(\xi) = 0$. □

There are some immediate consequences we may derive from this theorem. The first corollary is well-known but with help of our results we are able to reprove it.

Corollary 3.5.2 (Ozsváth and Szabó). *If (Y, ξ) is overtwisted, the contact element vanishes.*

Proof. Recall that the surgery diagram given in Figure 3.31 is an overtwisted contact structure ξ' on \mathbb{S}^3 .




Figure 3.31: Surgery diagram for an overtwisted \mathbb{S}^3 in the homotopy class of ξ_{std} .

This overtwisted contact structure is homotopic to ξ_{std} as 2-plane fields (cf. [4]). By Eliashberg's classification theorem (see [7]), a connected sum of (Y, ξ) with (\mathbb{S}^3, ξ') does not change the contact manifold, i.e.

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

Denote by K the shark on the left of Figure 3.31. The manifold (Y, ξ) admits a surgery representation $\mathbb{S}^3(\mathbb{L})$ where $\mathbb{L} = K \sqcup \mathbb{L}'$. Furthermore, K and \mathbb{L}' are not linked. Denote by (Y', ξ'') the contact manifold with surgery representation $\mathbb{S}^3(\mathbb{L}')$. We obtain (Y, ξ) out of (Y', ξ'') by $(+1)$ -contact surgery along K , which can be destabilized inside Y' . Proposition 3.5.1 implies the vanishing of $c(\xi)$. \square

Remark. For a detailed discussion of the homotopy invariants of overtwisted contact structures on \mathbb{S}^3 see [5].

Another consequence is that performing a simple Lutz twist along a transverse knot kills the contact element. The resulting contact structure is clearly overtwisted. Thus, by work of Ozsváth and Szabó the contact element vanishes. But besides this approach we can show the vanishing of the contact element without referring to overtwistedness at all. In [6] a surgical description for simple Lutz twists along transverse knots is presented. This description involves $(+1)$ -contact surgeries along a Legendrian approximation L of the transverse knot and another Legendrian knot which is a stabilized version of L . Proposition 3.5.1 then implies the vanishing of the contact element.

When looking at a homologically trivial knot L , to show the vanishing of a contact element after surgery along L it can be convenient to show the vanishing of $\widehat{\mathcal{L}}(L)$ and then apply Proposition 3.4.1, because of the various gradings on the knot Floer homological level. The following proposition is meant as an illustration of this fact.

Proposition 3.5.3. *A $(+1)$ -contact surgery along the Legendrian realizations L_n given in Figure 3.32 of the Eliashberg-Chekanov twist knots E_n with $n \in -2\mathbb{N}$ all give contact manifolds with vanishing contact element.*




Figure 3.32: The Eliashberg-Chekanov twist knots E_n and Legendrian realizations L_n .

Proof. Since the invariant $\widehat{\mathcal{L}}(L_n)$ of the Legendrian realizations L_n of the knots E_n live in $\widehat{\text{HFK}}(-\mathbb{S}^3, E_n)$, and because of the correspondence

$$\widehat{\text{HFK}}(-\mathbb{S}^3, E_n) = \widehat{\text{HFK}}(\mathbb{S}^3, \overline{E_n}),$$

where $\overline{E_n}$ denotes the mirror knot, we have to compute the groups $\widehat{\text{HFK}}(\mathbb{S}^3, \overline{E_n})$. The knots are all alternating. Therefore we will stick to Theorem 1.3 of [37] for a convenient computation of the groups. We compute the Alexander-Conway polynomial using its skein relation and get

$$\Delta_{\overline{E_n}}(T) = (1 - n) + \frac{n}{2}(T^1 + T^{-1}).$$

To compute the signature of the knots E_n , we use the formula given in Theorem 6.1 of [37] and see that all these knots have signature $\sigma(\overline{E_n}) = -n - 2$. By Theorem 1.3 of [37], which describes the knot Floer homology groups of an alternating knot in terms of the coefficients of the associated Alexander-Conway polynomial, the knot Floer homology of $\overline{E_n}$ looks like

$$\widehat{\text{HFK}}_j(\mathbb{S}^3, \overline{E_n}, i) = \begin{cases} \mathbb{Z}^{-n/2}, & i = -1, j = -1 + \frac{-n-2}{2} \\ \mathbb{Z}^{|1-n|}, & i = 0, j = \frac{-n-2}{2} \\ \mathbb{Z}^{-n/2}, & i = 1, j = 1 + \frac{-n-2}{2} \\ 0, & \text{otherwise} \end{cases}.$$

According to [35], the Legendrian invariant $\widehat{\mathcal{L}}(L_n)$ lives in $\widehat{\text{HFK}}_{M(L_n)}(-\mathbb{S}^3, E_n, A(L_n))$ where $A(L_n)$ is the *Alexander grading* of L_n and $M(L_n)$ is called *Maslov grading*. These gradings are computed using the formulas (see [35])

$$\begin{aligned} 2 \cdot A(L_n) &= tb(L_n) - rot(L_n) + 1 \\ d_3(\xi_{std}) &= 2A(L_n) - M(L_n), \end{aligned}$$

where d_3 denotes the Hopf-invariant (cf. [18]). However, note that with the conventions used in Heegaard Floer theory $d_3(\xi_{std}) = 0$. With a straightforward computation we see that $tb(L_n) = -4$ and $rot(L_n) = 1$, which give the following Alexander gradings and Maslov gradings

$$\begin{aligned} A(L_n) &= -1 \\ M(L_n) &= -2. \end{aligned}$$

Consequently, we can show, by using the computed Alexander and Maslov gradings, that for every knot L_n , $n \neq 0$, the invariant $\widehat{\mathcal{L}}(L_n)$ is an element of a vanishing




Figure 3.33: The Legendrian isotopy showing that L_0 can be destabilized.

subgroup of $\widehat{\text{HFK}}(\mathbb{S}^3, \overline{E}_n)$. To show the vanishing of $\widehat{\mathcal{L}}(L_0)$ we observe that L_0 can be destabilized.

The isotopy is pictured in Figure 3.33. By Proposition 3.5.1 $c(\xi_{L_0}^+)$ vanishes, too. Using Proposition 3.4.1 the proposition follows. \square

The following theorem is a new vanishing result of the contact element, which uses the knot Floer homology for arbitrary knots. Furthermore, we make use of the fact that in $\mathbb{S}^2 \times \mathbb{S}^1$ there are homologically non-trivial knots whose associated knot Floer homology vanishes.

Theorem 3.5.4. *Let (Y, ξ) be a contact manifold given as a contact surgery along a Legendrian link in $(\mathbb{S}^3, \xi_{std})$. If the surgery diagram contains a configuration like given in Figure 3.34, the contact element $c(Y, \xi)$ vanishes.*




Figure 3.34: Configuration in a surgery diagram of (Y, ξ) killing the contact element.

Proof. We start looking at the knot Floer homology group of the pair $(\mathbb{S}^2 \times \mathbb{S}^1, G)$ where G is a specific knot representing a generator of $H_1(\mathbb{S}^2 \times \mathbb{S}^1)$: Figure 3.35 is a Heegaard diagram adapted to this specific knot G . A straightforward calculation gives $\widehat{\text{HFK}}(\mathbb{S}^2 \times \mathbb{S}^1, G) = 0$. In Figure 3.36 we see a surgery diagram of $\mathbb{S}^2 \times \mathbb{S}^1$ with the knot G in it. Returning to Figure 3.34, we can interpret K' as an ordinary knot and remove it from the surgery description. We obtain a contact manifold $(Y' \# \mathbb{S}^2 \times \mathbb{S}^1, \xi')$ and K' is a Legendrian knot in it. A $(+1)$ -contact surgery along K' will yield (Y, ξ) . Furthermore, as a topological knot, K' can be written as $K'' \# G$ where $K'' \subset Y$ and $G \subset (\mathbb{S}^2 \times \mathbb{S}^1)$ is a knot representing a generator of $H_1(\mathbb{S}^2 \times \mathbb{S}^1)$. Hence, we have (cf. [27])

$$\widehat{\text{HFK}}(Y' \# (\mathbb{S}^2 \times \mathbb{S}^1), K') = \widehat{\text{HFK}}(Y', K'') \otimes \widehat{\text{HFK}}(\mathbb{S}^2 \times \mathbb{S}^1, G) = 0.$$




Figure 3.35: Heegaard diagram adapted to G




Figure 3.36: Surgery diagram of $\mathbb{S}^2 \times \mathbb{S}^1$ with knot G in it.

The same holds if we reverse the orientation on the manifold. We perform a (+1)-contact surgery along K' to obtain (Y, ξ) . Denote by W the induced cobordism. By Proposition 3.4.1 this induces a map

$$\Gamma_{-W}: \widehat{\text{HFK}}(-Y \# (\mathbb{S}^2 \times \mathbb{S}^1), K') \longrightarrow \widehat{\text{HF}}(-Y)$$

with $c(Y, \xi) = \Gamma_{-W}(\widehat{\mathcal{L}}(K'))$. So, the contact element vanishes, since $\widehat{\mathcal{L}}(K') = 0$. \square

Chapter 4

Holomorphic Discs and Surgery Exact Triangles

In this chapter we will refer to the sequences given in Corollaries 3.2.2 and 3.2.6 as the **Dehn Twist sequences**. The second part of this thesis, i.e. this chapter, mainly focuses on the relationship between the Dehn Twist sequences and the surgery exact triangle in knot Floer homology (cf. §2.6).

In this chapter we will begin proving that in the situation given in §3.2 we can set up an exact sequence by using maps defined by counting holomorphic triangles, i.e. with the cobordism maps:

$$\dots \xrightarrow{\partial_*} \widehat{\text{HFK}}(Y, K) \xrightarrow{\widehat{F}_{W_1}^w} \widehat{\text{HF}}(Y_{-1}(K)) \xrightarrow{\widehat{F}_{W_2}^w} \widehat{\text{HFK}}(Y_0(K), \mu) \xrightarrow{\partial_*} \dots \quad (4.0.1)$$

Of course, this strongly resembles the surgery exact sequence in knot Floer homology. However, the maps $\widehat{F}_{W_1}^w$ and $\widehat{F}_{W_2}^w$ are defined slightly different than in the situation of the knot Floer homology surgery exact sequence: the point w – encoding the knot – is used differently in the definition of these maps. Moreover, with this slight modification we see, that the Sequence (4.0.1) stays in a strong relationship with the Dehn Twist sequence from §3.2: we get the following diagram where all triangles and boxes

commute (cf. Theorem 4.1.6).

$$\begin{array}{ccccccc}
 & & & & \widehat{\text{HFK}}(Y_0(K), \mu) & \xrightarrow{f_*} & \dots \\
 & & & \nearrow \Gamma_2 & \downarrow \widehat{F}_{\alpha\delta\bar{\delta}}^w & & \uparrow \\
 \dots & \xrightarrow{\partial_*} & \widehat{\text{HFK}}(Y, K) & \xrightarrow{\widehat{F}_{W_1}^w} & \widehat{\text{HF}}(Y_{-1}(K)) & \xrightarrow{\widehat{F}_{W_2}^w} & \widehat{\text{HFK}}(Y_0(K), \mu) & \xrightarrow{\partial_*} & \dots \\
 & & \downarrow \widehat{F}_{\alpha\beta\bar{\beta}}^w & & & & & & \\
 \dots & \xrightarrow{f_*} & \widehat{\text{HFK}}(Y, K) & \xrightarrow{\Gamma_1} & \widehat{\text{HF}}(Y_{-1}(K)) & & & &
 \end{array} \quad (4.0.2)$$

As a consequence, the Dehn Twist sequences can be defined with coherent orientations and refined with respect to Spin^c -structures. Moreover, the connecting morphism f_* of the Dehn Twist sequence and the connecting morphism ∂_* fit into the following commutative square.

$$\begin{array}{ccc}
 \widehat{\text{HFK}}(Y, K) & \xrightarrow{f_*} & \widehat{\text{HFK}}(Y_0(K), \mu) \\
 \downarrow \widehat{F}_{Y \times I}^w & & \uparrow \widehat{F}_{Y_0(K) \times I}^w \\
 \widehat{\text{HFK}}(Y, K) & \xrightarrow{\partial_*} & \widehat{\text{HFK}}(Y_0(K), \mu)
 \end{array} \quad (4.0.3)$$

By looking at the mapping cone proof of the surgery exact sequence of Ozsváth and Szabó we see that (4.0.1) can be modified to give a surgery exact sequence where ∂_* is replaced by $\widehat{F}_{W_3}^w$. In consequence, the image and kernel of $\widehat{F}_{W_3}^w$ and ∂_* coincide. The composition law will show that this fact implies that the image and kernel of $\widehat{F}_{W_3}^w$ and f_* coincide. It follows immediately that the rank of the image and kernel of $\widehat{F}_{W_3}^w$ can be computed combinatorially. Of course, a more general result is already known by work of Lipshitz, Manolescu and Wang (see [25]). However, the relation we derived provides a new proof of this fact – at least in the knot Floer homology case – and gives rise to an alternative algorithm for the combinatorial computation. As a matter of fact the map f_* is defined by counting holomorphic discs in a suitable Heegaard diagram and this map carries information of the map $\widehat{F}_{W_3}^w$ which is defined by counting holomorphic triangles in a Heegaard triple diagram. To us, it seems that this fact makes it interesting to study properties of f_* . In §4.2.1 we will discuss in what situations the map f_* can be defined and study properties of them. These maps fulfill properties very similar to the properties of the cobordism maps: they fit into a surgery exact triangle and preserve contact geometric information when induced by (+1)-contact surgeries.

4.1 Surgery Exact Triangle and Dehn Twist Sequence

The shape of the Dehn Twist sequence strongly resembles the known surgery exact triangle in knot Floer homology (cf. §2.6). We will try to investigate and derive their relationship.

Given an abstract open book (P, ϕ) , let $\delta \subset P$ be a homologically essential, simple closed curve. Let $(\Sigma, \alpha, \beta, z)$ be an induced Heegaard diagram such that δ intersects β_1 once, transversely and is disjoint from the other β -circles. We define the following sets of attaching circles

$$\begin{aligned}\beta' &= \{\beta'_1, \dots, \beta'_g\} \\ \tilde{\delta} &= \{\tilde{\delta}, \beta''_2, \dots, \beta''_g\},\end{aligned}$$

where $\beta'_1 = D_\delta^+(\beta_1)$ and D_δ^+ denotes a positive Dehn Twist along δ . The β'_i , $i \geq 2$, are isotopic push-offs of the β_i such that β_i and β'_i intersect in a cancelling pair of intersection points. Furthermore, let β''_i , $i \geq 2$, be push-offs of the β'_i . As above, the push-offs are chosen such that the β''_i and β'_i intersect in a cancelling pair of intersection points. The curve $\tilde{\delta}$ is given as a perturbation (cf. Figure 4.1) of the curve δ , like indicated in Figure 4.1.




Figure 4.1: The relevant attaching circles.

Using the defined attaching circles we may form a sequence

$$\widehat{\text{CFK}}(\Sigma, \alpha, \beta, z, w) \xrightarrow{\widehat{F}_{\alpha\beta\beta'}^w} \widehat{\text{CF}}(\Sigma, \alpha, \beta', z) \xrightarrow{\widehat{F}_{\alpha\beta'\tilde{\delta}}^w} \widehat{\text{CFK}}(\Sigma, \alpha, \tilde{\delta}, z, w). \quad (4.1.1)$$

In the following we will use the notation \widehat{F} for both the map induced by a Heegaard triple on the homological level and the map induced on the chain level. Which one we are referring to will always be clear from the context. The superscript- w indicates that the map counts holomorphic triangles with $n_w = 0$. Using the mapping cone proof of Ozsváth and Szabó (cf. [42] or cf. §2.6), we can show that these fit into a surgery exact triangle

$$\begin{array}{ccc}
 & \widehat{\mathrm{HF}}(Y_{-1}(K)) & \\
 \widehat{F}_{W_1}^w \nearrow & & \searrow \widehat{F}_{W_2}^w \\
 \widehat{\mathrm{HFK}}(Y, K) & \xleftarrow{\widehat{F}_{W_3}^w} & \widehat{\mathrm{HFK}}(Y_0(K), \mu)
 \end{array} \tag{4.1.2}$$

where $\widehat{F}_{W_1}^w$ and $\widehat{F}_{W_2}^w$ correspond to the maps $\widehat{F}_{\alpha\beta\beta'}^w$ and $\widehat{F}_{\alpha\beta'\tilde{\delta}}^w$ in sequence (4.1.1). The map $\widehat{F}_{W_3}^w$ is the map on homology induced by the doubly-pointed triple diagram $(\Sigma, \alpha, \tilde{\delta}, \beta, w, z)$ (cf. §2.6). We will focus our attention on the sequence (4.1.1) and discuss the behavior of the maps \widehat{F} therein with methods similar to those used in §3.2.

By abuse of notation, we will denote by δ the set of attaching circles $\{\delta, \beta'_2, \dots, \beta'_g\}$, too. The work done in §3.2 shows that we have a short exact sequence of chain complexes

$$0 \longrightarrow \widehat{\mathrm{CFK}}(\Sigma, \alpha, \tilde{\beta}, z, w) \xrightarrow{\Gamma_1} \widehat{\mathrm{CF}}(\Sigma, \alpha, \beta', z) \xrightarrow{\Gamma_2} \widehat{\mathrm{CFK}}(\Sigma, \alpha, \delta, z, w) \longrightarrow 0. \tag{4.1.3}$$

The sequences (4.1.1) and (4.1.3) are designed to coincide at the middle term, namely at $\widehat{\mathrm{CF}}(\Sigma, \alpha, \beta', z)$.




Figure 4.2: Heegaard triple diagram defining $\widehat{F}_{\alpha, \beta, \beta'}^w$.

Lemma 4.1.1. *The maps $\widehat{F}_{\alpha\beta\beta'}^w$ and $\widehat{F}_{\alpha\beta'\delta}^w$ respect the splitting of $\widehat{\text{CF}}(\Sigma, \alpha, \beta', z)$, given in Proposition 3.2.1, i.e. given by the sequence (4.1.3).*

Proof. We show that the claim is true for the map $\widehat{F}_{\alpha\beta\beta'}^w$. We look at Figure 4.2 and try to show that there is no holomorphic triangle from an $\alpha\beta$ -intersection to an $\alpha\delta$ -intersection (cf. §3.2.1) that contributes to $\widehat{F}_{\alpha\beta\beta'}^w$:




Figure 4.3: Here we can see that $\widehat{F}_{\alpha\beta\beta'}^w$ respects the splitting.

Let ϕ be a triangle that connects a point $x \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ with a point $y \in \mathbb{T}_\alpha \cap \mathbb{T}_\delta \subset \mathbb{T}_\alpha \cap \mathbb{T}_{\beta'}$. The triangle ϕ connects y with $\widehat{\Theta}$ along its β' -boundary. In Figure 4.3 we illustrate the two possible ways to do that. In both cases the β' -boundary of ϕ follows the black arrow pictured there. We either cause a non-negative intersection number n_w (cf. left of Figure 4.3) or a non-negative intersection number n_z (cf. right part of Figure 4.3). Thus, $n_w(\phi) \neq 0$ or $n_z(\phi) \neq 0$, which shows that ϕ does not contribute to $\widehat{F}_{\alpha\beta\beta'}^w$. A similar line of arguments can be used to prove the claim for $\widehat{F}_{\alpha\beta'\delta}^w$. \square

It is a consequence of the last lemma that

$$\widehat{F}_{\alpha\beta'\tilde{\delta}}^w \circ \widehat{F}_{\alpha\beta\beta'}^w = 0.$$

Using the given attaching circles $\alpha, \beta, \tilde{\beta}, \delta$ and $\tilde{\delta}$ we may introduce the maps $\widehat{F}_{\alpha\beta\tilde{\beta}}^w$ and $\widehat{F}_{\alpha\tilde{\delta}\delta}^w$.

Lemma 4.1.2. *The diagram*

$$\begin{array}{ccc} \widehat{\text{CFK}}(\Sigma, \alpha, \beta, z, w) & \xrightarrow{\widehat{F}_{\alpha\beta\beta'}^w} & \widehat{\text{CF}}(\Sigma, \alpha, \beta', z) \\ \widehat{F}_{\alpha\beta\tilde{\beta}}^w \downarrow & \nearrow \iota & \\ \widehat{\text{CFK}}(\Sigma, \alpha, \tilde{\beta}, z, w) & & \end{array}$$

commutes where ι denotes the inclusion induced by a natural identification of generators.

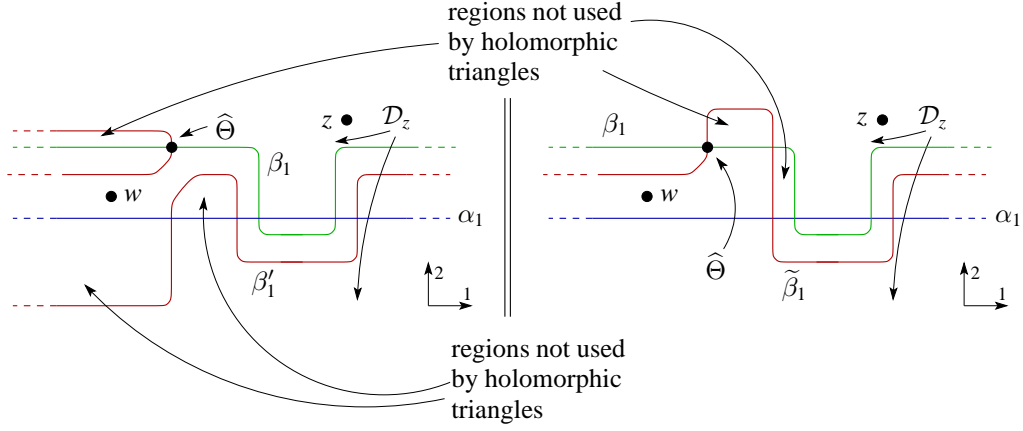


Figure 4.4: Comparing the boundary conditions of $\widehat{F}_{\alpha\beta\beta'}^w$ and $\widehat{F}_{\alpha\beta\tilde{\beta}}^w$.

Let us denote by h the map $\widehat{F}_{\alpha\beta\beta'}^w$ and by g the map $\widehat{F}_{\alpha\beta\tilde{\beta}}^w$. There is a canonical inclusion

$$\iota: \widehat{\text{CFK}}(\Sigma, \alpha, \tilde{\beta}, z, w) \longrightarrow \widehat{\text{CF}}(\Sigma, \alpha, \beta', z, w)$$

induced by an identification of intersection points. Namely, observe that

$$\begin{aligned} \mathbb{T}_\alpha \cap \mathbb{T}_{\beta'} &= \mathbb{T}_\alpha \cap \mathbb{T}_\beta \sqcup \mathbb{T}_\alpha \cap \mathbb{T}_\delta \\ &= \mathbb{T}_\alpha \cap \mathbb{T}_{\tilde{\beta}} \sqcup \mathbb{T}_\alpha \cap \mathbb{T}_\delta \end{aligned}$$

in case $\tilde{\beta}$ is a suitable perturbation of β we will define in a moment. We define $\tilde{\beta}_i = \beta_i$, for all $i \geq 2$, and $\tilde{\beta}_1$ as indicated in Figure 4.4 (see also Figure 4.1). We would like to show that $h = \iota \circ g$.

Definition 4.1.3. Let $(\Sigma, \alpha, \beta, z)$ be a Heegaard diagram and denote by $\mathcal{D}_1, \dots, \mathcal{D}_k$ the components of $\Sigma \setminus \{\alpha \cup \beta\}$. We say that a Whitney disc ϕ **does not use** a domain \mathcal{D}_i , $i \in \{1, \dots, k\}$, if the domain \mathcal{D}_i does not appear in $\mathcal{D}(\phi)$, i.e. writing $\mathcal{D}(\phi)$ as

$$\mathcal{D}(\phi) = \sum_{j=1}^k d_j \cdot \mathcal{D}_j,$$

the coefficient d_i vanishes. We also say that the domain $\mathcal{D}(\phi)$ **does not use** \mathcal{D}_i .

The main idea is to first prove that given intersections $x, y \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, all positive domains \mathcal{D} , i.e. all coefficients in \mathcal{D} are greater than or equal to 0, connecting x and y , with $n_w(\mathcal{D}) = n_z(\mathcal{D}) = 0$, do not use certain components of $\Sigma \setminus \{\alpha \cup \beta\}$ or $\Sigma \setminus \{\alpha \cup \tilde{\beta}\}$.

Which domains are expected not to be used is indicated in Figure 4.4, the left part illustrating the situation for h , the right part illustrating the situation for g . With this information, we compare the boundary conditions of holomorphic triangles for h and g . The conclusion will be that, with its β' -boundary, the holomorphic triangles counted by h always stay inside $\mathbb{T}_{\beta'} \cap \mathbb{T}_{\tilde{\beta}}$. And, with its $\tilde{\beta}$ -boundary, holomorphic triangles counted by g stay inside $\mathbb{T}_{\tilde{\beta}} \cap \mathbb{T}_{\beta'}$. Thus, we are able to identify the moduli spaces of holomorphic triangles contributing to h and g with arguments similar to those used in the proof of Proposition 3.2.1.

Proof. Figure 4.4 shows the part of the Heegaard triple diagrams where the boundary conditions for the holomorphic triangles involved in the definition of h and g differ. The picture illustrates which regions are not used by holomorphic triangles that contribute to h and g . This has to be shown in the following: We start our discussion with the map h and look at Figure 4.5. Each part of Figure 4.5 covers one of the cases which we will discuss in the following. The different parts of Figure 4.5 show parts of the Heegaard diagram pictured in the left of Figure 4.4. We focused on those parts important to our arguments. Denote by ϕ a holomorphic triangle that contributes to h . The domains, which we want to show not to be used by ϕ , will be denoted by \mathcal{D}_{x_i} , $i = 1, 2, 3$. In each of these regions we fix a point x_i , $i = 1, 2, 3$. If ϕ uses one of the domains \mathcal{D}_{x_i} , the associated intersection number n_{x_i} is non-zero.




Figure 4.5: Here we see why n_{x_i} , $i = 1, 2, 3$ have to be trivial.

Suppose the domain $\mathcal{D}(\phi)$ has non-trivial intersection number n_{x_1} (cf. left part of Figure 4.5). This means we generate a β' -boundary pointing inside \mathcal{D}_w , as indicated by the black arrow in the left part of Figure 4.5. Consequently, n_w has to be non-zero.

Suppose the domain $\mathcal{D}(\phi)$ has non-trivial intersection number n_{x_2} (cf. middle part of Figure 4.5). As we can see from the middle part of Figure 4.5 (by following the black

arrow), this forces n_z to be non-zero, since we generate a β' -boundary that has to run to $\hat{\Theta}$.

Suppose the domain $\mathcal{D}(\phi)$ has non-trivial intersection number n_{x_3} (cf. right part of Figure 4.5). This generates a β' -boundary emanating from $\hat{\Theta}$. Since n_z vanishes, the boundary has to run once along β'_1 . But then n_w is non-zero, as indicated by the black arrow.

This shows that every holomorphic triangle that contributes to h has trivial intersection number n_{x_i} , $i = 1, 2, 3$.




Figure 4.6: Here we see why n_{x_i} , $i = 1, 2$ have to be trivial.

We continue arguing that holomorphic triangles contributing to g , cannot use the domains indicated in the right part of Figure 4.4. Let ϕ be a holomorphic triangle contributing to g . Analogous to the discussion done for h , we denote the regions not expected to be used by ϕ with \mathcal{D}_{x_i} , $i = 1, 2$. In each of the domains we fix a point x_i . We want to show that n_{x_i} to be non-zero implies $n_w \neq 0$ or $n_z \neq 0$. The different parts of Figure 4.6 show parts of the Heegaard diagram pictured at the right of Figure 4.4.

Suppose the domain $\mathcal{D}(\phi)$ has non-trivial intersection number n_{x_1} (cf. left part of Figure 4.6). Since $n_w = 0$, we generate a β -boundary pointing inside \mathcal{D}_z , as it is indicated in the left part of Figure 4.6 (the boundary follows the black arrow). We see that $n_z \neq 0$.

Suppose the domain $\mathcal{D}(\phi)$ has non-trivial intersection number n_{x_2} (cf. right part of Figure 4.6). Since $n_z = 0$, we generate a β -boundary pointing inside \mathcal{D}_w (cf. right part of Figure 4.6) forcing n_w to be non-zero.

Thus, using arguments that are similar to those applied in the proof of Proposition 3.2.1, we can identify the moduli spaces of holomorphic triangles that contribute to h and g . \square

Lemma 4.1.4. *The diagram*

$$\begin{array}{ccc}
 & & \widehat{\text{CFK}}(\Sigma, \alpha, \delta, z, w) \\
 & \nearrow \pi & \downarrow \widehat{F}_{\alpha\delta\tilde{\delta}}^w \\
 \widehat{\text{CF}}(\Sigma, \alpha, \beta', z) & \xrightarrow{\widehat{F}_{\alpha\beta'\tilde{\delta}}^w} & \widehat{\text{CFK}}(\Sigma, \alpha, \tilde{\delta}, z, w)
 \end{array}$$

commutes where π is the projection induced by a natural identification of generators.

Proof. The proof is analogous to the proof of Lemma 4.1.2. Analogous to ι we can define the projection π by identifying

$$\begin{aligned}
 \mathbb{T}_\alpha \cap \mathbb{T}_{\beta'} &= \mathbb{T}_\alpha \cap \mathbb{T}_\beta \sqcup \mathbb{T}_\alpha \cap \mathbb{T}_\delta \\
 &= \mathbb{T}_\alpha \cap \mathbb{T}_\beta \sqcup \mathbb{T}_\alpha \cap \mathbb{T}_{\tilde{\delta}},
 \end{aligned}$$

i.e. by identifying $\mathbb{T}_\alpha \cap \mathbb{T}_\delta$ with $\mathbb{T}_\alpha \cap \mathbb{T}_{\tilde{\delta}}$. This induces a projection π between the respective chain modules.

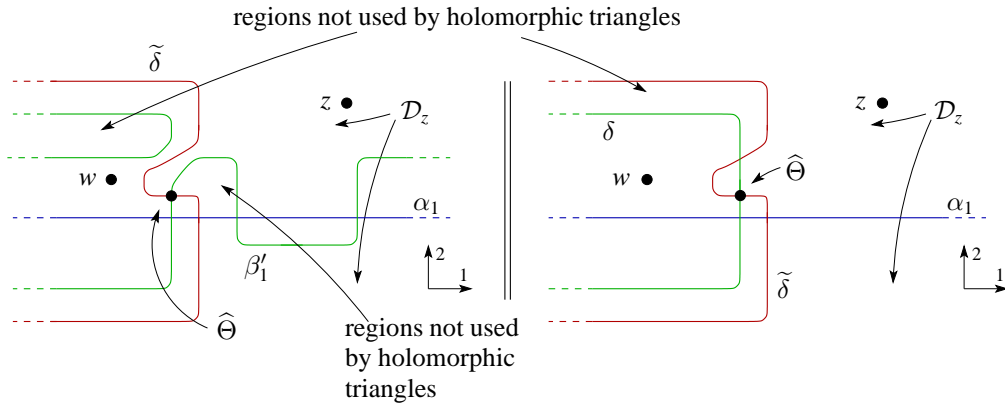


Figure 4.7: Comparing the boundary conditions of $\widehat{F}_{\alpha\beta'\tilde{\delta}}^w$ and $\widehat{F}_{\alpha\delta\tilde{\delta}}^w$.

In the following we will denote by h the map $\widehat{F}_{\alpha\beta'\tilde{\delta}}^w$ and by g the map $\widehat{F}_{\alpha\delta\tilde{\delta}}^w$. This time, we would like to show that $h = g \circ \pi$. Figure 4.7 indicates which domains are not used by holomorphic triangles (in the sense of Definition 4.1.3) that contribute to g and h . This has to be shown in the following discussion. Observe that each part of Figure 4.8 shows a part of the Heegaard diagrams pictured in Figure 4.7. Each of these portions will be relevant in one of the cases we will have to investigate. There are two domains not to be used by holomorphic triangles contributing to g (cf. left part of Figure 4.7).

In each of these domains we fix a point x_i and denote the associated domain by \mathcal{D}_{x_i} , $i = 1, 2$ (cf. left and middle part of Figure 4.8). There is one domain not to be used by triangles contributing to h (cf. right part of Figure 4.7). We fix a point x_3 in this domain and denote the associated domain by \mathcal{D}_{x_3} (cf. right of Figure 4.8). Let ϕ be a holomorphic triangle that contributes to g .

Suppose the domain $\mathcal{D}(\phi)$ has non-trivial intersection n_{x_1} (cf. left part of Figure 4.8). This generates a β' -boundary like indicated by the black arrow in the left portion of Figure 4.8. This boundary cannot be killed, i.e. cannot be interpreted as sitting in the interior of $\mathcal{D}(\phi)$, since $n_w = 0$. This β' -boundary, thus, has to emanate from $\widehat{\Theta}$ forcing it to follow the black arrow like indicated. Thus, n_z is non-zero.

Suppose the domain $\mathcal{D}(\phi)$ has non-trivial intersection n_{x_2} (cf. middle part of Figure 4.8). We create a β' -boundary like indicated by the black arrow in the middle portion of Figure 4.8. This boundary points towards $\widehat{\Theta}$. But recall that the β' -boundary of ϕ has to emanate from $\widehat{\Theta}$, as can be seen by looking at the triangle pictured at the top of the left and middle part of Figure 4.8. Thus, we have to generate a β' -boundary going along β' once, completely. But this implies n_w to be non-zero.

Now suppose that ϕ is a holomorphic triangle that contributes to g . Assume the domain $\mathcal{D}(\phi)$ has non-trivial intersection n_{x_3} (cf. right part of Figure 4.8). This time we generate $\widetilde{\delta}$ -boundary like indicated by the black arrow in the right portion of Figure 4.8. This boundary cannot be killed, since $n_z = 0$. This boundary has to emanate from $\widehat{\Theta}$ as can be seen by looking at the triangle pictured at the top of the right part of Figure 4.8. But this is impossible, since $n_w = 0$.

We have seen that holomorphic triangles, that contribute to h or g , do not use the domains indicated in Figure 4.7. Again, using arguments that are similar to those applied in the proof of Proposition 3.2.1, we can identify the moduli spaces of holomorphic triangles that contribute to h and g . This shows that $h = g \circ \pi$. \square

From Lemma 4.1.2 and Lemma 4.1.4 we see that (4.1.1) is a short exact sequence of chain complexes (since (4.1.3) is) and, thus, it induces a long exact sequence

$$\begin{array}{ccc}
 & \widehat{\text{HF}}(Y_{-1}(K)) & \\
 \nearrow \widehat{F}_{W_1}^w & & \searrow \widehat{F}_{W_2}^w \\
 \widehat{\text{HFK}}(Y, K) & \xleftarrow{\partial_*} & \widehat{\text{HFK}}(Y_0(K), \mu)
 \end{array} \tag{4.1.4}$$




Figure 4.8: Here we see why n_{x_i} , $i = 1, 2$ have to be trivial for f and why n_{x_3} has to be trivial for g .

between the homologies. When comparing with the sequence (4.1.2), we immediately see that

$$\begin{aligned} \text{im}(\widehat{F}_{W_3}^w) &= \text{im}(\partial_*) \\ \ker(\widehat{F}_{W_3}^w) &= \ker(\partial_*) \end{aligned} \quad (4.1.5)$$

Moreover, putting together both Lemma 4.1.2 and Lemma 4.1.4, we derive a strong relationship between the sequences (4.1.1) and (4.1.3).

Theorem 4.1.5. *All triangles and boxes in the following diagram commute.*

$$\begin{array}{ccccccc} & & & & \widehat{\text{HFK}}(Y_0(K), \mu) & \xrightarrow{f_*} & \dots \\ & & & \nearrow \Gamma_2 & \downarrow \widehat{F}_{\alpha\delta\tilde{\delta}}^w & & \uparrow \\ \dots & \xrightarrow{\partial_*} & \widehat{\text{HFK}}(Y, K) & \xrightarrow{\widehat{F}_{W_1}^w} & \widehat{\text{HF}}(Y_{-1}(K)) & \xrightarrow{\widehat{F}_{W_2}^w} & \widehat{\text{HFK}}(Y_0(K), \mu) & \xrightarrow{\partial_*} & \dots \\ & & \downarrow \widehat{F}_{\alpha\beta\tilde{\beta}}^w & \nearrow \Gamma_1 & & & & & \\ \dots & \xrightarrow{f_*} & \widehat{\text{HFK}}(Y, K) & & & & & & \end{array} \quad (4.1.6)$$

Proof. We put together Lemma 4.1.2 and Lemma 4.1.4 to get two short exact sequences of chain complexes that are related like claimed, i.e. we have

$$\begin{array}{ccccccc} & & & & \widehat{\text{CFK}}(\Sigma, \alpha, \delta, z, w) & \longrightarrow & 0 \\ & & & \nearrow \pi & \downarrow \widehat{F}_{\alpha\delta\tilde{\delta}}^w & & \\ 0 & \longrightarrow & \widehat{\text{CFK}}(\Sigma, \alpha, \beta, z, w) & \xrightarrow{\widehat{F}_{\alpha\beta\beta'}^w} & \widehat{\text{CF}}(\Sigma, \alpha, \beta', z) & \xrightarrow{\widehat{F}_{\alpha\beta'\delta}^w} & \widehat{\text{CFK}}(\Sigma, \alpha, \tilde{\delta}, z, w) & \longrightarrow & 0 \\ & & \downarrow \widehat{F}_{\alpha\beta\tilde{\beta}}^w & \nearrow \nu & & & & & \\ 0 & \longrightarrow & \widehat{\text{CFK}}(\Sigma, \alpha, \tilde{\beta}, z, w) & & & & & & \end{array}$$

To identify the diagonal sequence, i.e. the sequence

$$0 \longrightarrow \widehat{\text{CFK}}(\Sigma, \alpha, \tilde{\beta}, z, w) \xrightarrow{\iota} \widehat{\text{CF}}(\Sigma, \alpha, \beta', z) \xrightarrow{\pi} \widehat{\text{CFK}}(\Sigma, \alpha, \delta, z, w) \longrightarrow 0$$

with the Dehn Twist sequence given in Corollary 3.2.2, we have to isotope $\tilde{\beta}_1$ a bit. Observe that $\tilde{\beta}_1$ does not match with the situation given in Corollary 3.2.2 or with the situation given in Proposition 3.2.1 (cf. Figure 4.1 and the proof of Proposition 3.2.1). The isotopy, however, is supported within $\mathcal{D}_z \cup \mathcal{D}_w$. Furthermore, recall that an isotopy not generating/cancelling intersection points, acts on the Heegaard Floer homology as a perturbation $\mathcal{J}_{s,t}$ of the path of almost complex structures $\mathcal{J}_{s,0}$ (cf. §2.3.3) used in the definition of the Heegaard Floer homologies. We have to see that the induced map $\widehat{\Phi}_{\mathcal{J}_{s,t}}$ (cf. §2.3.2 and §2.3.3) is the identity on the chain level: In the definition $\widehat{\Phi}_{\mathcal{J}_{s,t}}$ we count 0-dimensional components of holomorphic discs with $n_w = n_z = 0$. The family $\mathcal{J}_{s,t}$ coincides with $\mathcal{J}_{s,0}$ outside of a set, which is contained in $(\mathcal{D}_z \cup \mathcal{D}_w) \times \text{Sym}^{g-1}(\Sigma)$, since the isotopy perturbing $\tilde{\beta}_1$ is supported in $\mathcal{D}_z \cup \mathcal{D}_w$. Thus, for $x, y \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, we have an identification

$$\left(\mathcal{M}_{\mathcal{J}_{s,t}}(x, y) \right)_{n_z=n_w=0}^{\mu=0} = \left(\mathcal{M}_{\mathcal{J}_{s,0}}(x, y) \right)_{n_z=n_w=0}^{\mu=0}, \quad (4.1.7)$$

where the notation should indicate that we are interested in moduli spaces with Maslov index 0 and whose elements satisfy $n_z = n_w = 0$. The moduli space on the right of Equation (4.1.7), in the following denoted by \mathcal{M} , is empty unless $x = y$: Suppose there is a holomorphic Whitney disc ϕ connecting x with y . Assuming x and y are not equal, the disc ϕ is non-constant. So, because of the translation action (cf. §2.1.2) the disc ϕ comes in a 1-dimensional family. Thus, ϕ cannot be an element of \mathcal{M} . If x and y are the same point, the moduli space \mathcal{M} contains the constant holomorphic disc. But it does not contain non-constant holomorphic discs by the same reasoning done for $x \neq y$.

Consequently, the map $\widehat{\Phi}_{\mathcal{J}_{s,t}}$ is the identity on the chain level. We know from §2.3.2 that the map $\widehat{\Phi}_{\mathcal{J}_{s,t}}$ is a chain map, i.e. we have

$$0 = \widehat{\partial}_{\mathcal{J}_{s,1}} \circ \widehat{\Phi}_{\mathcal{J}_{s,t}} - \widehat{\Phi}_{\mathcal{J}_{s,t}} \circ \widehat{\partial}_{\mathcal{J}_{s,0}} = \widehat{\partial}_{\mathcal{J}_{s,1}} - \widehat{\partial}_{\mathcal{J}_{s,0}}.$$

Thus, the signed count of holomorphic discs with Maslov index 1 in both

$$\widehat{\text{CFK}}(\Sigma, \alpha, \beta, z, w) \quad \text{and} \quad \widehat{\text{CFK}}(\Sigma, \alpha, \tilde{\beta}, z, w)$$

equals for each homotopy class admitting holomorphic representatives. Thus, we may replace the map ι with Γ_1 . The map π already equals Γ_2 . \square

Corollary 4.1.6. *The following equalities hold*

$$\begin{aligned}\mathrm{im}(\widehat{F}_{W_3}^w) &= \mathrm{im}(f_*) \\ \mathrm{ker}(\widehat{F}_{W_3}^w) &= \mathrm{ker}(f_*),\end{aligned}$$

where f is the map defined in Corollary 3.2.3.

Proof. Consider the commutative diagram

$$\begin{array}{ccc}\widehat{\mathrm{HFK}}(Y, K) & \xrightarrow{f_*} & \widehat{\mathrm{HFK}}(Y_0(K), \mu) \\ \downarrow \widehat{F}_{Y \times I}^w & & \uparrow \widehat{F}_{Y_0(K) \times I}^w \\ \widehat{\mathrm{HFK}}(Y, K) & \xrightarrow{\partial_*} & \widehat{\mathrm{HFK}}(Y_0(K), \mu)\end{array} \quad (4.1.8)$$

which is the square from sequence (4.1.6) and which commutes according to Theorem 4.1.5. Note that the vertical maps are induced by the triples $\alpha\beta\tilde{\beta}$ and $\alpha\delta\tilde{\delta}$, which can be associated to the trivial cobordisms $Y \times I$ and $Y_0(K) \times I$. As we have observed in (4.1.5), the kernel and the image of $\widehat{F}_{W_3}^w$ coincide with the kernel and image of ∂_* . Thus, we may write $\widehat{F}_{W_3}^w$ instead of ∂_* at the lower arrow. Doing so, the box does not commute anymore but the composition of $\widehat{F}_{W_3}^w$ with the vertical maps yields a map whose kernel and image coincides with the kernel and image of f_* . By the composition law of the maps induced by cobordisms the composition is again a map associated to a cobordism. Denote this cobordism by W . The following square indicates the situation.

$$\begin{array}{ccc}\widehat{\mathrm{HFK}}(Y, K) & \xrightarrow{f_*} & \widehat{\mathrm{HFK}}(Y_0(K), \mu) \\ \downarrow \widehat{F}_{Y \times I}^w & \widehat{F}_W^w & \uparrow \widehat{F}_{Y_0(K) \times I}^w \\ \widehat{\mathrm{HFK}}(Y, K) & \xrightarrow{\partial_*} & \widehat{\mathrm{HFK}}(Y_0(K), \mu) \\ & \widehat{F}_{W_3}^w & \end{array} \quad (4.1.9)$$

Using the composition law we get

$$\widehat{F}_W^w = \widehat{F}_{Y \times I}^w \circ \widehat{F}_{W_3}^w \circ \widehat{F}_{Y_0(K) \times I}^w = \widehat{F}_{Y \times I \cup W \cup Y_0(K) \times I}^w = \widehat{F}_{W_3}^w$$

giving the desired result. \square

So, basically, instead of counting holomorphic triangles, we can count holomorphic discs to gain information about the map $\widehat{F}_{W_3}^w$. Especially, given that $(\Sigma, \alpha, \beta', z)$ is a **nice** Heegaard diagram (in the sense of Sarkar and Wang, see [47] or cf. Definition 2.1.28). In this case the map f can be computed combinatorially. In this way we get

an algorithm to combinatorially compute the rank of the kernel and the image of $\widehat{F}_{W_3}^w$. Note that Lipshitz, Manolescu and Wang in [25] determine an algorithm to do that for cobordism maps in the hat-theory in case of \mathbb{Z}_2 -coefficients. Their result is more general than ours, however, in case of knot Floer homologies we are able to present a different algorithm.

Corollary 4.1.7 ([25]). *The rank of the kernel and image of $\widehat{F}_{W_3}^w$ can be computed combinatorially.*

Proof. To compute the ranks it suffices to compute the ranks of the kernel and image of f_* combinatorially. Recall that f_* is part of the boundary $\widehat{\partial}_{\alpha,\beta'}$. It remains to show that there is a nice Heegaard diagram (cf. Definition 2.1.28) induced by an open book decomposition and which is adapted to the setup used to define the sequences given in Corollaries 3.2.2 and 3.2.6. In [45] Plamenevskaya shows that the Sarkar-Wang algorithm (see [47]) can be modified to apply for open books by just using isotopies of the monodromy. This means that a given open book (P, ϕ) can be modified to an isotopic open book (P, ϕ') such that the associated Heegaard diagram is nice. To give some more details: Start with an open book (P, ϕ) and choose a cut system to define an associated Heegaard diagram (Σ, α, β) . Use finger moves (see [47]) of the β -curves inside the page $P \times \{1\}$ to obtain a nice Heegaard diagram (cf. Definition 2.1.28). These finger moves add up to give an isotopy φ_t of the page P . This isotopy, by construction, is the identity near the boundary P . The resulting diagram is adapted to the curve $\phi_1(\delta)$. \square

The following corollary is an immediate consequence of Theorem 4.1.5. We will not outline the proof, since the construction is lengthy but straightforward. The horizontal part of the sequence given in Theorem 4.1.5 can be defined with coherent orientations, and it refines with respect to Spin^c -structures (in the sense of [40]). The diagonal part, i.e. the Dehn Twist sequence, commutes with the horizontal part, so, we can use the refinements and the coherent orientations on the horizontal part to generate refinements and coherent orientations on the Dehn Twist sequence.

Corollary 4.1.8. *The Dehn Twist sequences, i.e. the sequences given in Corollaries 3.2.2 and 3.2.6, can be defined with coherent orientations. Furthermore, these sequences refine with respect to Spin^c -structures.* \square

4.2 Chain Maps and Holomorphic Discs

The last paragraph enlightened a connection between counting holomorphic triangles in doubly-pointed Heegaard triple diagrams and counting holomorphic discs in doubly-pointed Heegaard diagrams. This connection gave rise to an alternative algorithm to compute ranks of cobordism maps combinatorially. We will focus our attention on the maps f , as defined in Corollary 3.2.3, and try to answer the following questions: Is it possible to give a definition of f in general situations? What properties do these maps have?

4.2.1 General Definition

To give a general definition of the map f , suppose we are given a pair (Y, K) where Y is a 3-manifold and $K \subset Y$ a knot. Let $(\Sigma, \alpha, \beta, z)$ be a subordinate Heegaard diagram, i.e. we write

$$T^2 \# \Sigma' = \Sigma$$

such that K is the core of the first torus component, i.e. of T^2 . We apply the notation from Proposition 3.2.1. Let μ be a meridian of T^2 and define $\tilde{\beta}_1$ as $\lambda + n \cdot \mu$ where $\lambda + n \cdot \mu$ represents the surgery framing of K . The left part of Figure 4.9 illustrates the situation:




Figure 4.9: Heegaard diagrams suitable for defining f .

The diagram $(\Sigma, \alpha, \delta, w, z)$ represents the pair (Y, K) and $(\Sigma, \alpha, \tilde{\beta}, w, z)$ represents the surgered manifold Y_K , and in it, a knot μ_K^Y . These two diagrams fit into an exact triangle (cf. Corollary 3.2.2)

$$\begin{array}{ccc}
 \widehat{\text{HF}}(\Sigma, \alpha, \delta, z, w) & \xrightarrow{f_*} & \widehat{\text{HF}}(\Sigma, \alpha, \tilde{\beta}, z, w) \\
 \searrow \Gamma_2 & & \swarrow \Gamma_1 \\
 \widehat{\text{HF}}(\Sigma, \alpha, \beta', z) & &
 \end{array} \tag{4.2.1}$$

where β' is defined by applying to $\tilde{\beta}_1$ a positive Dehn Twist along δ . With Proposition 3.2.1 the sequence is defined (analogous to the sequence given in Corollary 3.2.2) and by Corollary 3.2.3 we get a definition of f within (4.2.1). Namely, for $x \in \mathbb{T}_\alpha \cap \mathbb{T}_\delta$ we define

$$f_{\alpha\delta,w}(x) = \sum_{y \in \mathbb{T}_\alpha \cap \mathbb{T}_{\tilde{\beta}}} \sum_{\phi \in H^*(x,y,1)} \#\widehat{\mathcal{M}}_\phi \cdot y,$$

where $H(x, y, 1) \subset \pi_2^{\alpha\beta'}(x, y)$ are the homotopy classes with Maslov index 1 and such that the pair $(n_*(\phi), n_{**}(\phi))$ does not equal $(0, 0)$. As defined in Corollary 3.2.3, we denote by $\pi_2^{\alpha\beta'}$ the Whitney discs associated to the diagram $(\Sigma, \alpha, \beta', z)$. We define $f_* = (f_{\alpha\delta,w})_*$.

Proposition 4.2.1. *Let K' be a push-off of K (with respect to its framing) in Y . The knot μ_K^Y is the knot K' interpreted as sitting in Y_K .*




Figure 4.10: Determining the knot μ_K^Y

Proof. The manifold Y_K is given by

$$Y_K = Y \setminus (\mathbb{S}^1 \times D^2) \cup_{\varphi} \mathbb{S}^1 \times D^2$$

$$\begin{array}{ccc} n \cdot \mu + \lambda & \longleftarrow & \mu_0 \\ -\mu & \longleftarrow & \lambda_0 \end{array}$$

where λ is the longitude determining the framing given by the tubular neighborhood of K , μ_0 is a meridian and λ_0 the standard longitude of $\mathbb{S}^1 \times D^2$. The knot μ_Y^K is determined via the pair (w, z) in the diagram $(\Sigma, \alpha, \tilde{\beta})$. By definition of the pair (w, z) , the induced knot, i.e. μ_K^Y , intersects the co-core of the 2-handle determined by K once, transversely and is disjoint from all other 2-handles. Hence, in the decomposition above, the curve μ_K^Y equals the longitude λ_0 . The gluing map sends the curve λ_0 to a

meridian of K . Hence, the situation described in the left part of Figure 4.10 applies. Sliding μ_K^Y over the 2-handle determined by K (cf. Figure 4.10), we see that in Y_K the knot μ_K^Y is isotopic to a push-off of K that determines the surgery framing on K . \square

Thus, we obtain a map

$$f_* : \widehat{\text{HFK}}(Y, K) \longrightarrow \widehat{\text{HFK}}(Y_K, \mu_K^Y).$$

Theorem 4.2.2. *The map f_* does not depend on the choices made in its definition. It just depends on the cobordism induced by the surgery along K .*

Proof. This immediately follows from the invariance properties derived in §3.3). \square

4.2.2 Properties

Given a 3-manifold and a knot $K \subset Y$ with framing n , let us do surgery along K with its specified framing and denote by W_1 the induced cobordism. With the discussion done in paragraph §4.2.1, we can associate to the cobordism W_1 a map

$$f_{W_1} : \widehat{\text{HFK}}(Y, K) \longrightarrow \widehat{\text{HFK}}(Y_K^n, K'),$$

where $K' = \mu_{K'}^Y$ is a meridian of K in Y interpreted as sitting in Y_K^n . We continue to form a surgery exact triangle (cf. §2.6), i.e. we do (-1) -surgery along K' , and denote its induced cobordism by W_2 . We obtain a map

$$f_{W_2} : \widehat{\text{HFK}}(Y_K^n, K') \longrightarrow \widehat{\text{HFK}}(Y_K^{n+1}, K''),$$

where $K'' = \mu_{K''}^{Y_K^n}$. Interpreted as sitting in Y , the knot K'' is a meridian of K' , and it is not linked with K . Surgery along K'' with framing (-1) yields the manifold Y , again. Denote the associated cobordism by W_3 .

Theorem 4.2.3. *The maps f_{W_i} , $i = 1, \dots, 3$, fit into the following surgery exact sequence*

$$\begin{array}{ccc} \widehat{\text{HFK}}(Y, K) & \xrightarrow{f_{W_1}} & \widehat{\text{HFK}}(Y_K^n, K') \\ & \searrow f_{W_3} & \swarrow f_{W_2} \\ & \widehat{\text{HFK}}(Y_K^{n+1}, K'') & \end{array} \quad (4.2.2)$$




Figure 4.11: Determining the type of K''' .

Proof. First of all we have to see that $K''' = \mu_{K''}^{Y_K^{n+1}}$ is isotopic to K in Y . The left part of Figure 4.11 pictures the knot K''' in the surgery diagram of Y induced by the surgery triangle. With a Rolfsen twist, a handle slide of K''' over K and a slam dunk we show that K''' is a copy of K . Thus, f_{W_3} is, indeed, a map as indicated in (4.2.2). It remains to show exactness of the sequence, given in the theorem: In the present situation we intend to show that the cobordisms W_i , $i = 1, \dots, 3$, induce an exact triangle

$$\begin{array}{ccc}
 \widehat{\text{HFK}}(Y, K) & \xrightarrow{\widehat{F}_{W_1}} & \widehat{\text{HFK}}(Y_K^n, K') \\
 & \searrow \widehat{F}_{W_3} & \swarrow \widehat{F}_{W_2} \\
 & \widehat{\text{HFK}}(Y_K^{n+1}, K'') &
 \end{array} \tag{4.2.3}$$

To do that, we have to see that the cobordisms fit topologically into a surgery exact triangle (cf. §2.6). This is done in Figure 4.12. The left portion shows the moves done to produce sequence (4.2.2). We start with a pair (Y, K) and topologically do a surgery along K with framing n . Comparing this move with the corresponding move pictured in the right part of Figure 4.12, we see that both are equivalent after a handle slide, as indicated in the picture. Following the second and the third arrow in the left portion of Figure 4.12, we perform the same recipe, i.e. we compare with the right portion of Figure 4.12 and detect equality after a suitable handle slide. Since we are in a suitable topological situation, with a straightforward adaption of the proof of the surgery exact triangle in knot Floer homology given by Ozsváth and Szabó, we see that (4.2.3) is, indeed, an exact sequence. Using Corollary 4.1.6 (especially Diagram (4.1.9)) at each arrow of the sequence (4.2.3), we can replace the maps \widehat{F}_{W_i} with f_{W_i} , $i = 1, \dots, 3$, without affecting exactness. Thus, we get (4.2.2). \square




Figure 4.12: The left portion determines the topological moves done to produce the triangle, given in (4.2.2). The right portion those moves done to produce (4.2.3).

Theorem 4.2.4. *Let (Y, ξ) be a contact manifold, $L \subset Y$ a Legendrian knot and let W be the cobordism induced by a $(+1)$ -contact surgery along L . Then the map*

$$f_{-w}: \widehat{\text{HFK}}(-Y, L) \longrightarrow \widehat{\text{HFK}}(-Y_L^+, L')$$

preserves the contact geometric information, i.e. $\widehat{f}_{-w}(\widehat{\mathcal{L}}(L)) = \widehat{\mathcal{L}}(L')$. Here L' is a push-off of L in Y interpreted in Y_L^+ .

Proof. The top row and the bottom row of Figure 4.13 illustrate the situation for both possible orientations of L . Choose an open book decomposition (P, ϕ) adapted to the contact structure ξ such that L sits on a page of the open book with the contact framing coinciding with the page framing. We may choose a cut system in such a way that L intersects the first β -circle once and is disjoint from the other β -circles. The β -circle, having a non-trivial intersection with L , should be denoted by δ . We obtain a set of attaching circles

$$\delta = \{\delta, \beta_2, \dots, \beta_g\}.$$

This set of attaching circles can be used to define the Heegaard diagram $(\Sigma, \alpha, \delta, z)$. We include an additional point w such that the pair (z, w) determines L as an oriented knot. The left column of Figure 4.13 illustrates both possibilities, i.e. the positioning for both potential orientations on L . We define β_1 to be the curve, obtained after

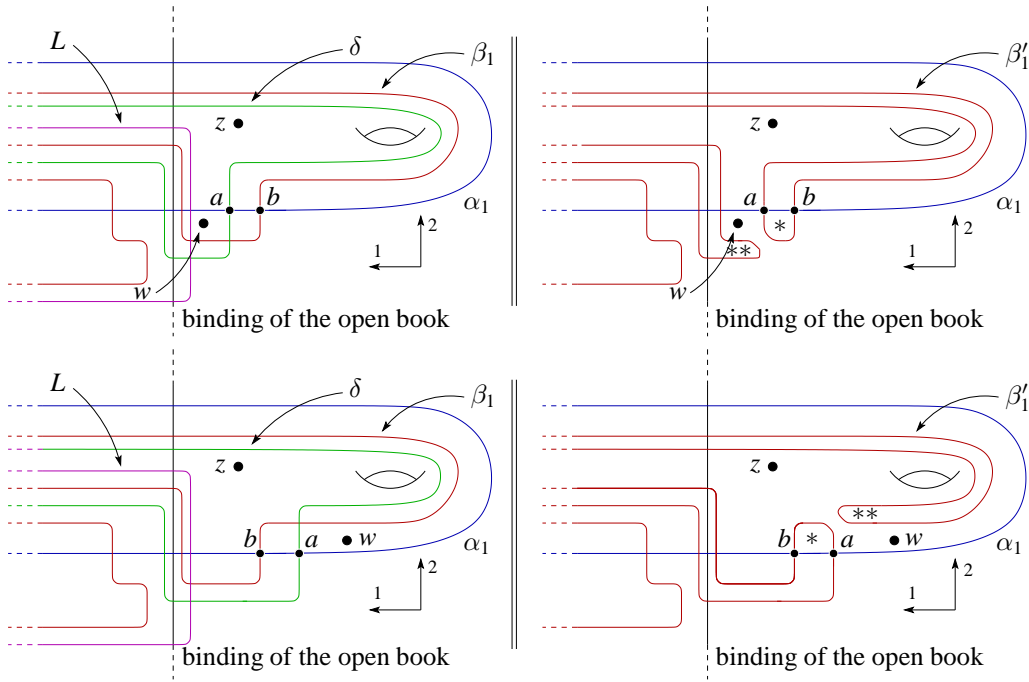


Figure 4.13: The top row and bottom row illustrate the situations for both orientations on the Legendrian knot.

applying to δ a negative Dehn Twist along L . We obtain a third set of attaching circles

$$\beta = \{\beta_1, \beta_2, \dots, \beta_g\}.$$

Observe that the cobordism W given by the triple $(\Sigma, \alpha, \delta, \beta)$ is induced by a $(+1)$ -contact surgery along L . Furthermore, observe that the data $(\Sigma, \alpha, \beta, w, z)$ and the curve $\delta \subset \Sigma$ are suitable for applying Proposition 3.2.1. We get an exact sequence as given in Corollaries 3.2.2 and 3.2.6. The induced connecting morphism is denoted by

$$f_{-W}: \widehat{\text{HFK}}(-Y, L) \longrightarrow \widehat{\text{HFK}}(-Y_L^+, L'),$$

where L' is a push-off of L interpreted as sitting in $-Y_L^+$. The map f_{-W} is induced by the map $f_{\alpha\beta, w}$ (see Corollary 3.2.3) which is defined for $x \in \mathbb{T}_\alpha \cap \mathbb{T}_\delta$ by

$$f_{\alpha\beta, w}(x) = \sum_{y \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} \sum_{\phi \in H(x, y, 1)} \#\widehat{\mathcal{M}}_\phi \cdot y,$$

where $H(x, y, 1) \subset \pi_2^{\alpha\beta'}(x, y)$ are the homotopy classes of Whitney discs in $(\Sigma, \alpha, \beta', z)$ with $\mu = 1$ and $(n_*(\phi), n_{**}(\phi)) \neq (0, 0)$. Hence, the right column of Figure 4.13 ap-

plies. There are specific intersections $x_i \in \alpha_i \cap \beta_i$, $i \geq 2$, such that the point $\{a, x_2, \dots, x_g\}$ represents the Legendrian knot invariant $\widehat{\mathcal{L}}(L)$ in $\widehat{\text{HFK}}(-Y, L)$ and such that $\{b, x_2, \dots, x_g\}$ represents the knot invariant $\widehat{\mathcal{L}}(L')$ in $\widehat{\text{HFK}}(-Y_L^+, L')$ (cf. Figure 4.13). There is only one holomorphic disc ϕ connecting $\{a, x_2, \dots, x_g\}$ and $\{b, x_2, \dots, x_g\}$. This disc satisfies $n_*(\phi) = 1$, and, hence, it appears in the definition of $f_{\alpha\beta, w}$. The positions of the points w and z circumvent the existence of any other holomorphic disc emanating from $\{a, x_2, \dots, x_g\}$. Thus, we see that

$$f_{\alpha\beta, w}(\{a, x_2, \dots, x_g\}) = \{b, x_2, \dots, x_g\}$$

completing the proof. □

Bibliography

- [1] G. E. Bredon, *Geometry and Topology*, Graduate Texts in Mathematics, Vol. **139**, Springer-Verlag, 1993.
- [2] F. Ding and H. Geiges, *Symplectic fillability of tight contact structures on torus bundles*, *Algebr. Geom. Topol.* **1** (2001), 153–172.
- [3] ———, *A Legendrian surgery presentation of contact 3-manifolds*, *Math. Proc. Cambridge Philos. Soc.* **136** (2004), 583–598.
- [4] ———, *Handle moves in contact surgery diagrams*, *J. Topol.* **2** (2009), 105–122.
- [5] F. Ding, H. Geiges, and A. I. Stipsicz, *Surgery diagrams for contact 3-manifolds*, *Turkish J. Math.* **28** (2004), 41–74, (Proceedings of the 10th Gökova Geometry-Topology Conference, 2003).
- [6] ———, *Lutz twist and contact surgery*, *Asian J. of Math.* **9** (2005), 57–64.
- [7] Ya. Eliashberg, *Classification of overtwisted contact structures on 3-manifolds*, *Invent. Math.* **98** (1989), 623–637.
- [8] ———, *Topological characterization of Stein manifolds of dimension > 2* , *Internat. J. Math.* **1** (1990), 29–46.
- [9] J. B. Etnyre, *Lectures on open-book decompositions and contact structures*, *Amer. Math. Soc.* **5** (2006), 103–142, (Proceedings of the Clay Mathematics Summer School).
- [10] A. Floer, *A relative index for the symplectic action*, *Comm. Pure Appl. Math.* **41** (1988), 393–407.
- [11] ———, *Morse theory for Lagrangian intersections*, *J. Differ. Geom.* **28** (1988), 513–547.

- [12] ———, *The unregularised gradient flow of the symplectic action*, *Comm. Pure Appl. Math.* **41** (1988), 775–813.
- [13] ———, *Symplectic fixed points and holomorphic spheres*, *Commun. Math. Phys.* **120** (1989), 576–611.
- [14] ———, *Witten’s complex and infinite dimensional Morse theory*, *J. Differ. Geom.* **30** (1989), 207–221.
- [15] H. Geiges, *A contact geometric proof of the Whitney-Graustein theorem*, *Enseign. Math.* **55**(2) (2009), 93–102.
- [16] ———, *An Introduction to Contact Topology*, *Cambridge Studies in Advanced Mathematics*, Vol. **109**, Cambridge University Press, 2008.
- [17] P. Ghiggini, P. Lisca, and A. I. Stipsicz, *Tight contact structures on some small Seifert fibered 3-manifolds*, *Amer. J. Math.* **129** (2007), 1403–1447.
- [18] R. E. Gompf and A. I. Stipsicz, *4-Manifolds and Kirby Calculus*, *Graduate Studies in Mathematics*, Vol. **20**, American Mathematical Society, 1999.
- [19] M. Gromov, *Pseudo-holomorphic curves in symplectic manifolds*, *Inv. Math.* **82** (1985), 307–347.
- [20] K. Honda, W. H. Kazez, and G. Matić, *On the contact class in Heegaard Floer homology*, *J. Diff. Geom.*, to appear.
- [21] K. Honda, W. H. Kazez, and G. Matić, *The contact invariant in Sutured Floer homology*, arXiv:0705.2828.
- [22] A. Juhász, *Holomorphic discs and sutured manifolds*, *Alg. Geom. Topol.* **6** (2006), 1429–1457.
- [23] P. Kronheimer and T. Mrowka, *Monopoles and Three-Manifolds*, *New Mathematical Monographs*, Vol. **10**, Cambridge University Press, 2008.
- [24] R. Lipshitz, *A cylindrical reformulation of Heegaard-Floer homology*, *Geom. Topol.* **10** (2006), 955–1097.
- [25] R. Lipshitz, C. Manolescu, and J. Wang, *Combinatorial cobordism maps in the hat Heegaard Floer theory*, *Duke Math. J.* **145**(2) (2008), 207–247.
- [26] R. Lipshitz, P. Ozsváth, and D. Thurston, *Bordered Heegaard Floer homology: Invariance and pairing*, arXiv:0810.0687v2.

- [27] P. Lisca, P. Ozsváth, A. I. Stipsicz, and Z. Szabó, *Heegaard Floer invariants of Legendrian knots in contact three-manifolds*, JEMS, to appear.
- [28] P. Lisca and A. I. Stipsicz, *Ozsváth-Szabó invariants and tight contact manifolds I*, *Geom. Topol.* **8** (2004), 925–945.
- [29] ———, *Notes on the contact Ozsváth-Szabó invariants*, *Pacific J. Math.* **228**(2) (2006), 277–295.
- [30] ———, *Ozsváth-Szabó invariants and tight contact manifolds II*, *J. Diff. Geom.* **75** (2007), 109–141.
- [31] ———, *Ozsváth-Szabó invariants and tight contact manifolds III*, *J. Symplectic Topology* **5** (2007), 357–384.
- [32] I. G. MacDonald, *Symmetric products of an algebraic curve*, *Topology* **1** (1962), 319–343.
- [33] D. McDuff and D. Salamon, *j -Holomorphic Curves and Symplectic Topology*, *Colloquium Publications*, Vol. **52**, American Mathematical Society, 2004.
- [34] B. Ozbagci and A. I. Stipsicz, *Surgery on Contact 3-Manifolds and Stein Surfaces*, *Bolyai Society Mathematical Studies*, Vol. **13**, Springer-Verlag, 2004.
- [35] P. Ozsváth and A. I. Stipsicz, *Contact surgeries and the transverse invariant in knot Floer homology*, arXiv:0803.1252v1 (2008).
- [36] P. Ozsváth and Z. Szabó, *Heegaard diagrams and holomorphic disks*, *Diff. faces of Geom., Int. Math. Series*, 301–348.
- [37] ———, *Heegaard Floer homology and alternating knots*, *Geom. Topol.* **7** (2003), 225–254.
- [38] ———, *Holomorphic disks and knot invariants*, *Adv. Math.* **186** (2004), 58–116.
- [39] ———, *Holomorphic disks and three-manifold invariants: Properties and applications*, *Ann. of Math.* **159**(3) (2004), 1159–1245.
- [40] ———, *Holomorphic disks and topological invariants for closed three-manifolds*, *Ann. of Math.* **159**(3) (2004), 1027–1158.
- [41] ———, *Heegaard Floer homologies and contact structures*, *Duke Math. J.* **129**(1) (2005), 39–61.

- [42] ———, *On the Heegaard Floer homology of branched double-covers*, Adv. in Math. **194** (2005), 1–33.
- [43] ———, *An introduction to Heegaard Floer homology*, Amer. Math. Soc. **5** (2006), 3–28, (Proceedings of the Clay Mathematics Summer School).
- [44] ———, *Holomorphic triangles and invariants of smooth four-manifolds*, Adv. Math. **202**(2) (2006), 326–400.
- [45] O. Plamenevskaya, *A combinatorial description of the Heegaard Floer contact invariant*, arXiv:math.GT/0612322v1 (2006).
- [46] V. V. Prasolov and A. B. Sossinskiy, *Knots, Links, Braids and 3-Manifolds*, Translations of Mathematical Monographs, Vol. **154**, American Mathematical Society, 1997.
- [47] S. Sarkar and J. Wang, *An algorithm for computing some Heegaard Floer homologies*, (2008), to appear in Ann. of Math.
- [48] A. I. Stipsicz and V. Vertesi, *On invariants for Legendrian knots*, arXiv:0806.1436.
- [49] A. I. Stipsicz, *Surgery diagrams and open books decompositions of contact 3-manifolds*, Acta Math. Hungar. **108**(1-2) (2005), 71–86.
- [50] C. H. Taubes, *Embedded contact homology and Seiberg-Witten Floer cohomology I*, arXiv.org:0811.3985v2.

Ich versichere, dass ich die von mir vorgelegte Dissertation selbständig angefertigt, die benutzten Quellen und Hilfsmittel vollständig angegeben und die Stellen der Arbeit —einschliesslich Tabellen, Karten und Abbildungen—, die anderen Werken im Wortlaut oder dem Sinn nach entnommen sind, in jedem Einzelfall als Entlehnung kenntlich gemacht habe; dass diese Dissertation noch keiner anderen Fakultät oder Universität zur Prüfung vorgelegen hat; dass sie —abgesehen von unten angegebenen Teilpublikationen— noch nicht veröffentlicht worden ist sowie dass ich eine Teilveröffentlichung vor Abschluss des Promotionsverfahrens nicht vornehmen werde.

Die Bestimmungen der Promotionsordnung sind mir bekannt. Die von mir vorgelegte Dissertation ist von Herrn Prof. Hansjörg Geiges, Ph.D. (Cantab) betreut worden.

Teilpublikationen: Die wesentlichen Inhalte aus Kapitel 3 wurden in Form eines Artikels mit dem Titel 'Dehn Twists in Heegaard Floer Homology' auf dem Preprint-Server www.arXiv.org veröffentlicht.

Lebenslauf

Name	Bijan Sahamie
Geburtsdatum	31. Oktober 1979
Geburtsort	Oberhausen
Staatsangehörigkeit	deutsch
Familienstand	verheiratet
Schulbildung	
1986-1990	Grundschule Nesselrodestr. in Köln
1990-1999	Erich-Kästner-Gymnasium in Köln
1999-2000	Grundwehrdienst
Studium	
10/2000	Einschreibung zu den Studiengängen Mathematik und Physik auf Diplom an der Universität zu Köln
02/2003	Vordiplome in Physik und Mathematik
09/2005	Diplom in Mathematik an der Universität zu Köln Thema der Diplomarbeit: Classification of Plane- and Morse-Smale Fields on 3-Manifolds
11/2005 - 09/2009	Promotionsstudium der Mathematik an der Universität zu Köln
09/2008-02/2009	Visiting Scholar an der Columbia University in New York (DAAD gefördert) bei P. S. Ozsváth
Berufstätigkeit	
10/2002 - 07/2005	Studentische Hilfskraft am Mathematischen Institut der Universität zu Köln
11/2005 - 09/2009	Wissenschaftlicher Mitarbeiter am Mathematischen Institut der Universität zu Köln
seit 10/2009	Wissenschaftlicher Assistent am Mathematischen Institut der LMU München bei D. Kotschick.