

**From the holdings of the Department of Special Collections , Princeton University Library,
Princeton, NJ**

These documents are either in the public domain, or can only be used for educational and research purposes (Fair use) as per U.S. Copyright law (text below). By accessing this file, all users agree that their use falls in the public domain or within fair use as defined by the copyright law. We do not charge any permission or usage fees for the publication of images of material in our collections, including those provided by Princeton University Library, via our website, catalog, or directly from staff. For more information regarding copyright, credit, and citation please review the information on the Copyright, Credit and Citation Guidelines page of our website.

<https://library.princeton.edu/special-collections/policies/copyright-credit-and-citation-guidelines>. If you have any questions, including how to order higher-quality images from the collection, contact Special Collections Public Services via our Ask Us! Form. <https://library.princeton.edu/special-collections/ask-us>

Doctoral Dissertations, 2011-2026

Mathematics

2019

Involutive Heegaard Floer Homology and Homology Cobordism

Dai, Irving

Princeton, NJ : Princeton University

<http://arks.princeton.edu/ark:/88435/dsp0102870z743>

U.S. Copyright law: The copyright law of the United States (Title 17, §108, United States Code) governs the making of photocopies or other reproductions of copyrighted material. Under certain conditions specified in the law, libraries and archives are authorized to furnish a photocopy or other reproduction. One of these specified conditions is that the photocopy or reproduction is not to be used for any purpose other than private study, scholarship, or research. If a user makes a request for, or later uses, a photocopy or reproduction for purposes in excess of fair use that user may be liable for copyright infringement. This institution reserves the right to refuse to accept a copying order if, in its judgment, fulfillment of the order would involve a violation of copyright law. <https://www.copyright.gov/title17/title17.pdf>

INVOLUTIVE HEEGAARD FLOER HOMOLOGY AND HOMOLOGY COBORDISM

IRVING DAI

A DISSERTATION
PRESENTED TO THE FACULTY
OF PRINCETON UNIVERSITY
IN CANDIDACY FOR THE DEGREE
OF DOCTOR OF PHILOSOPHY

RECOMMENDED FOR ACCEPTANCE
BY THE DEPARTMENT OF
MATHEMATICS
ADVISER: ZOLTÁN SZABÓ

JUNE 2019

© Copyright by Irving Dai, 2019.

All Rights Reserved

Abstract

In this thesis, we present several results regarding the application of Heegaard Floer theory to the homology cobordism group. The majority of our work is concerned with giving a structural understanding of the involutive Heegaard Floer homology for linear combinations of Seifert fibered spaces. As an application, we show that if Y is a linear combination of Seifert fibered homology spheres with $\mu(Y) = 1$, then Y is not torsion in the homology cobordism group. We also discuss what can be said about the $\text{Pin}(2)$ -equivariant monopole Floer homology of Seifert fibered spaces using our techniques. These results give a possible approach towards showing that Seifert fibered spaces do not generate the homology cobordism group.

Acknowledgements

Firstly, I would like to thank my adviser, Zoltán Szabó, for his continued support and encouragement throughout my time as a graduate student at Princeton. Without his extended guidance and mentorship, this thesis would not have been possible. I would also like to take this opportunity to thank all of my professors over the course of my mathematical career, both at Princeton and otherwise. Special thanks goes to Peter Kronheimer for mentoring me during my time as an undergraduate. I would also like to thank Ciprian Manolescu, Peter Ozsváth, Cesar Silva, and Clifford Taubes for their support during various critical junctures of my academic development.

I would also like to thank Jennifer Hom, Ciprian Manolescu, Matthew Stoffregen, and Linh Truong for being my collaborators and co-authors. Their input has been instrumental in helping me learn and produce mathematics. In addition, I would like to thank the students, faculty, and staff (especially Jill LeClair) of the Princeton mathematics department for helping foster an environment in which I could freely discuss and conduct research. I am immeasurably grateful for the opportunities that they have created.

Last but not least, I would like to thank my family and friends, especially my father, mother, and sister. Without their continued support, I would not be where I am today. It is thus my great pleasure to close this chapter of my life with them at my side. Finally, I would like to extend a special thanks to Maggie Miller (for making my time at Princeton significantly happier, as well as being a great source of mathematical and moral encouragement) and Sifan Zhang (for a long history of friendship).

To my family, especially Helen.

Contents

Abstract	iii
Acknowledgements	iv
1 Introduction	1
1.1 Historical overview	2
1.2 Recent developments	4
1.3 Statement of results	6
2 Background Material	9
2.1 Involutive Heegaard Floer homology	9
2.2 Local equivalence	11
2.3 Connected Heegaard Floer homology	15
2.4 Lattice homology	16
3 Preliminary Notions	20
3.1 Graded roots and standard complexes	20
3.2 Geometric complexes	24
4 Involutive Floer Homology of Seifert Fibered Spaces	29
4.1 Identification of J_0	29
4.2 Computation of the involutive Floer homology	31
4.3 Correction terms	34
4.4 The monotone subroot algorithm	36
5 The Structure of $h(\Theta_{\mathbf{SF}})$	42
5.1 Good families	42
5.2 The swapping lemma	47

5.3	Proof of the main theorem	58
5.4	The basis decomposition algorithm	61
5.5	Examples and applications	62
6	Connected Heegaard Floer Homology	67
6.1	The connected homology algorithm	67
6.2	Preliminary examples	71
6.3	Geometric operations	74
6.4	Algebraic details	78
6.5	Sketch of proof	80
6.6	Examples and applications	84
7	Pin(2)-Equivariant Monopole Floer Homology	87
7.1	Monopole Floer homology	87
7.2	The Gysin sequence	89
8	Generalizations and Further Work	96
8.1	Lattice homology for other plumbings	96
8.2	Generalization of results	98
8.3	Other work	99

Chapter 1

Introduction

The aim of this thesis is to present several (joint) results of the author regarding the application of Heegaard Floer theory to the homology cobordism group $\Theta_{\mathbb{Z}}^3$. Historically, the homology cobordism group has occupied a central place in the development of low-dimensional topology, and has correspondingly enjoyed a long and fruitful relationship with Floer theory. The most recent example of this is Manolescu's celebrated disproof of the triangulation conjecture, which proceeded by establishing constraints on the existence of 2-torsion in $\Theta_{\mathbb{Z}}^3$ [Man16]. Manolescu's argument utilized a new suite of homology cobordism invariants coming from $\text{Pin}(2)$ -equivariant Seiberg-Witten Floer homology. In this paper, we will likewise use a recent involutive refinement of Heegaard Floer homology (introduced by Hendricks and Manolescu in [HM17]) to derive structural results concerning $\Theta_{\mathbb{Z}}^3$.

There are two main theoretical and computational tools which will be used in this thesis. The first, which we have already mentioned, is involutive Heegaard Floer homology. The second is a combinatorial construction due to Ozsváth-Szabó [OS03b] and Némethi [Ném08] called lattice homology. This is a homology theory for plumbed 3-manifolds which is known to be isomorphic to Heegaard Floer homology in a wide range of cases. One of the fundamental results we discuss in this paper is the fact that in many situations, the lattice homology of Y also be used to compute its involutive Heegaard Floer homology. Through an extended series of calculations, this will allow us to constrain various homology cobordism invariants for a large family of plumbed 3-manifolds.

The results described in this thesis are drawn from [DM17], [DS17], [Dai18a], and [Dai18b]. Where appropriate, the author is indebted to Ciprian Manolescu and Matthew Stoffregen (his co-authors). Throughout, we have opted to emphasize concrete examples and algorithms, referring

occasionally to the original papers for details of proofs. For the sake of exposition, we have also refrained from stating our results in their fullest generality. For example, many of our theorems hold (in appropriate form) for rational homology spheres, rather than integer homology spheres. A discussion of these and other generalizations is included at the end of the paper.

1.1 Historical overview

We begin with a brief history of the homology cobordism group. Although this section is mostly orthogonal to the rest of the thesis, we hope that our discussion here will help contextualize our results for readers less familiar with the study of $\Theta_{\mathbb{Z}}^3$.

Recall that a (closed, oriented) 3-manifold Y is said to be an integer homology sphere if $H_*(Y) \cong H_*(S^3)$. Two homology spheres Y_1 and Y_2 are said to be homology cobordant if there exists a smooth 4-manifold W such that (a) $\partial W = -Y_1 \sqcup Y_2$, and (b) $H_*(W) \cong H_*(S^3 \times I)$.¹ We then define the homology cobordism group by setting

$$\Theta_{\mathbb{Z}}^3 = \{(\text{oriented}) \text{ integer homology spheres}\} / \text{homology cobordism}.$$

This has a group structure, with group operation given by connected sums, inverses given by orientation reversal, and identity given by the class of S^3 . Note that the class of a homology sphere is zero in $\Theta_{\mathbb{Z}}^3$ if and only if it bounds a (smooth) integer homology ball.

The simplest example of a non-trivial integer homology sphere is given by the Poincaré sphere, which is usually defined as the quotient of S^3 by the binary icosahedral group. This fits into a larger family of examples called Brieskorn spheres, which may be defined as follows. Given coprime natural numbers (p, q, r) , let $\Sigma(p, q, r)$ be the intersection

$$\{x^p + y^q + z^r = 0\} \cap S_{\epsilon}^5 \subseteq \mathbb{C}^3,$$

where S_{ϵ}^5 is a small sphere centered around the origin. It turns out that these 3-manifolds always have the same homology as S^3 , and can be distinguished from each other (and S^3) by a fundamental group computation. Brieskorn spheres are themselves a special case of a larger class of integer homology spheres, called Seifert fibered homology spheres. These will form a fundamental family of examples in this thesis.

Although Brieskorn spheres for different triples (p, q, r) are distinct up to diffeomorphism, it turns out that understanding their behavior with respect to homology cobordism is significantly more complicated. For example, let W be the smooth 4-manifold constructed by attaching a single 1-handle to B^4 and then a 2-handle which winds algebraically (but not geometrically) once around

¹Here, all homologies are taken with integer coefficients.

this 1-handle, as in Figure 1.1. Then W is an integer homology ball, and using elementary Kirby calculus one can prove that the boundary of W is diffeomorphic to $\Sigma(2, 5, 7)$.² Thus, even in the simplest cases, determining whether a given homology sphere is zero in $\Theta_{\mathbb{Z}}^3$ is a subtle and difficult question. For further examples of such constructions, see [Maz61, AK79, CH81].




Figure 1.1: A non-trivial homology 4-ball. Figure taken from [AD05].

The converse problem of obtaining invariants for obstructing homology cobordism likewise has a long history in low-dimensional topology. The first instance of such an invariant was given by the Rokhlin homomorphism, which may be defined as follows. Given a homology sphere Y , it is a standard fact that Y bounds a smooth 4-manifold W with even (unimodular) intersection form and $b_1(W) = 0$. By the classification of lattices, this necessarily has signature divisible by 8. The Rokhlin invariant of Y is then given by

$$\mu(Y) = \sigma(W)/8 \pmod{2}.$$

This is independent of the choice of W due to Rokhlin's theorem, which states that the signature of any (smooth) closed spin 4-manifold is divisible by 16. A similar argument proves that μ descends to a well-defined homomorphism

$$\mu : \Theta_{\mathbb{Z}}^3 \rightarrow \mathbb{Z}/2\mathbb{Z},$$

called the Rokhlin homomorphism. As a simple example, note that $\Sigma(2, 3, 5)$ bounds a 4-manifold W_{E_8} whose intersection form is given by the E_8 lattice. Since this has signature -8 , we see that $\mu(\Sigma(2, 3, 5)) = 1$. Hence the class of $\Sigma(2, 3, 5)$ is not zero in $\Theta_{\mathbb{Z}}^3$. Note, however, that the Rokhlin invariant does not prevent $\Sigma(2, 3, 5)$ from being 2-torsion in $\Theta_{\mathbb{Z}}^3$.

A more refined argument can be obtained by using Donaldson's diagonalization theorem. Consider again $\Sigma(2, 3, 5)$. If this were the boundary of a smooth homology ball W , then we could use W to cap off the previously mentioned W_{E_8} to obtain a (closed) smooth 4-manifold with intersection form E_8 . But E_8 is not diagonalizable over $\mathbb{Z}$, so this is impossible. The same argument shows that

²Choosing different framings gives, for example, $\Sigma(2, 3, 13)$ and $\Sigma(3, 4, 5)$. These homology balls are due to Mazur [Maz61], and are an example of what are called Mazur manifolds.

in fact any $\#_n \Sigma(2, 3, 5)$ is non-zero in $\Theta_{\mathbb{Z}}^3$, so $\Sigma(2, 3, 5)$ is not torsion in $\Theta_{\mathbb{Z}}^3$. Similar lattice-theoretic obstructions have now been adapted in a wide range of other contexts, including rational homology cobordism and the study of the slice-ribbon conjecture [Lis07].

There are thus several deep connections between $\Theta_{\mathbb{Z}}^3$ and the theory of smooth 4-manifolds. To further reinforce this correspondence, note that according to a celebrated result of Freedman, every homology sphere bounds a contractible *topological* 4-manifold. Relaxing the smoothness condition on the cobordism W in the definition of $\Theta_{\mathbb{Z}}^3$ thus leads to the trivial group. Hence, the structure of $\Theta_{\mathbb{Z}}^3$ (in some sense) reflects the dichotomy between smooth and continuous topology in dimension four. Given this, it is not surprising that many of the current methods for understanding $\Theta_{\mathbb{Z}}^3$ are drawn from gauge theory and Floer homology. Historically, one of the first successes of this approach was due to Furuta [Fur90] and Fintushel-Stern [FS90], who used Yang-Mills theory on orbifolds to show that $\Theta_{\mathbb{Z}}^3$ admits a $\mathbb{Z}^\infty$ -subgroup.

Despite several advances, however, many problems involving the nature of the homology cobordism group remain open. We close our historical discussion with some broad questions about $\Theta_{\mathbb{Z}}^3$ which will motivate the results of this thesis:

1. What can we say about the algebraic structure of $\Theta_{\mathbb{Z}}^3$? For example, does $\Theta_{\mathbb{Z}}^3$ have any torsion?
2. Can we restrict the geometric properties of representatives in $\Theta_{\mathbb{Z}}^3$? For example, do Seifert fibered spaces generate $\Theta_{\mathbb{Z}}^3$?

Regarding (1), note that according to Manolescu's theorem [Man16], there is no 2-torsion in $\Theta_{\mathbb{Z}}^3$ with Rokhlin invariant one. In general, it is not known whether there is torsion in $\Theta_{\mathbb{Z}}^3$ of any kind (Rokhlin invariant one or otherwise). Regarding (2), we note that by a theorem of Myers [Mye83], every homology sphere is homology cobordant to a hyperbolic manifold; hence hyperbolic manifolds span $\Theta_{\mathbb{Z}}^3$. It is known that not every homology sphere is homology cobordant to a single Seifert fibered space (see for example [Sto17]), but the question of whether Seifert fibered spaces span $\Theta_{\mathbb{Z}}^3$ is still open. We use Θ_{SF} to denote the subgroup of $\Theta_{\mathbb{Z}}^3$ generated by Seifert fibered spaces, so that (2) asks whether the inclusion $\Theta_{\text{SF}} \subseteq \Theta_{\mathbb{Z}}^3$ is strict.

1.2 Recent developments

We now give a brief description of some of the more recent developments in Floer theory in order to properly state our results. Precise definitions and further details can be found in Chapter 2.

In [Man16], Manolescu introduced an equivariant refinement of Seiberg-Witten Floer/monopole Floer homology, called Pin(2)-equivariant Seiberg-Witten Floer homology. This is a homology theory

for rational homology spheres which utilizes an extra symmetry in the Chern-Simons-Dirac functional. Using the $\text{Pin}(2)$ -homology of Y , Manolescu defined a set of three integer-valued homology cobordism invariants, α, β , and γ , and used these to prove:

Theorem 1.2.1. [Man16, Corollary 1.2] Let Y be a homology sphere with $\mu(Y) = 1$. Then Y is not 2-torsion in $\Theta_{\mathbb{Z}}^3$.

See also work of Lin [Lin18a]. In light of Manolescu's theorem, it is natural to ask whether there is any torsion in $\Theta_{\mathbb{Z}}^3$ with Rokhlin invariant one. In this thesis, one of our principal applications will be to answer this question negatively for linear combinations of Seifert fibered spaces.

In [HM17], Hendricks and Manolescu introduced a similar refinement on the Heegaard Floer side, called involutive Heegaard Floer homology.³ This takes into account the conjugation symmetry on a Heegaard diagram given by interchanging the α - and β -curves. The involutive Heegaard Floer homology of Y likewise comes equipped with two numerical homology cobordism invariants $\bar{d}$ and $\underline{d}$. In this paper, however, we will usually work with the homotopy type of the underlying chain complex, which consists of the usual Heegaard Floer complex $CF^-(Y)$, together with the data of a homotopy involution ι on $CF^-(Y)$.⁴ Such a pair $(CF^-(Y), \iota)$ is called an ι -complex.

In [HMZ18], Hendricks, Manolescu, and Zemke define a certain equivalence relation on the set of ι -complexes, called local equivalence. This should be thought of as the algebraic analogue of homology cobordism, in the sense that if two homology spheres are homology cobordant, then their associated ι -complexes are locally equivalent. Letting

$$\mathfrak{J} = \{\text{all possible } \iota\text{-complexes}\} / \text{local equivalence},$$

we thus obtain a map

$$h : \Theta_{\mathbb{Z}}^3 \rightarrow \mathfrak{J}$$

given by sending Y to the local equivalence class of its (grading-shifted) ι -complex. In [HMZ18], it is shown that h is a homomorphism, where the group structure on the right-hand side is given by tensor product. (See Section 2.2 for a precise definition of $\mathfrak{J}$.) Our approach will thus be to study $\Theta_{\mathbb{Z}}^3$ by understanding the structure of $\mathfrak{J}$. One of the main results of this thesis will be an explicit computation of $h(\Theta_{\text{SF}})$, which will allow us to constrain homology cobordisms between linear combinations of Seifert fibered spaces.

We also discuss a related construction due to Hendricks, Hom, and Lidman called connected Heegaard Floer homology. This is an invariant of the local equivalence class of Y which takes the

³This is actually the Heegaard Floer analogue of $\mathbb{Z}/4\mathbb{Z}$ -equivariant Seiberg-Witten Floer homology, rather than $\text{Pin}(2)$ -equivariant Seiberg-Witten Floer homology. See [HM17, Section 3].

⁴The involutive Heegaard Floer homology is then the homology of the mapping cone of $1 + \iota$. See Section 2.1.

form of a summand of the (usual) Heegaard Floer homology $HF^-(Y)$. (Since the connected Floer homology is simply an $\mathbb{F}[U]$ -module, this is somewhat easier to understand than the ι -complex of Y .) Roughly speaking, the connected Heegaard Floer homology is composed of certain U -torsion towers in $HF^-(Y)$ that are preserved under the maps induced by a homology cobordism. In this paper, we will use our computation of $h(\Theta_{\text{SF}})$ to constrain the form of the connected Floer homology for any linear combination of Seifert fibered spaces, thus opening up a possible approach towards showing that $\Theta_{\text{SF}} \neq \Theta_{\mathbb{Z}}^3$.

Finally, a brief word on lattice homology. Given a (negative-definite) plumbed 3-manifold Y , Ozsváth-Szabó [OS03b] and Némethi [Ném08] associate to Y a combinatorial homology theory called lattice homology, computed from the plumbing diagram for Y . This takes the form of a graded $\mathbb{F}[U]$ -module. In many cases (in particular, for Seifert fibered spaces), the lattice homology of Y is known to be isomorphic to its Heegaard Floer homology. This has been used by several authors to prove results in the Heegaard Floer setting; see for example [Kar14, KL15]. However, it turns out that the lattice homology of Y contains some additional information on top of being isomorphic to the Heegaard Floer homology. In this paper, we will leverage this fact to compute the involutive Heegaard Floer homology whenever Y is a Seifert fibered homology sphere.

1.3 Statement of results

We now give a brief summary of the results discussed in this thesis. In Chapter 4, we present joint work with Ciprian Manolescu from [DM17] regarding the involutive Heegaard Floer homology of Seifert fibered spaces. The main theorem we prove here is the following:

Theorem 1.3.1. [DM17, Theorem 1.1] Let Y be a (positively oriented) Seifert fibered homology sphere. Then both the involutive Heegaard Floer homology $HFI^-(Y)$ and its associated local equivalence class can be computed from the lattice homology of Y .

In the statement of Theorem 1.3.1, we have been intentionally vague about what it means to compute the local equivalence class of Y . In Section 4.4, we will show that the local equivalence class of any Seifert fibered space can be parameterized by a finite tuple of integers, and we provide a straightforward algorithm for determining these from the lattice homology. Theorem 1.3.1 will moreover establish an explicit model for the ι -complex of such manifolds, which will be convenient for future computations. As an application of Theorem 1.3.1, we prove the following result regarding the involutive Floer correction terms:

Theorem 1.3.2. [DM17, Theorem 1.2] Let Y be a (positively oriented) Seifert fibered homology sphere. Then

$$\bar{d}(Y) = d(Y) \quad \text{and} \quad \underline{d}(Y) = -2\bar{\mu}(Y).$$

Here, $\bar{\mu}(Y)$ is the Neumann-Siebenmann invariant of Y (see Definition 4.3.1).

In Chapter 5, we present joint work with Matthew Stoffregen from [DS17] generalizing the results of [DM17] to linear combinations of Seifert fibered spaces. The main theorem we have here is an explicit computation of $h(\Theta_{\text{SF}})$:

Theorem 1.3.3. [DS17, Theorem 1.1] Let Θ_{SF} denote the subgroup of $\Theta_{\mathbb{Z}}^3$ generated by Seifert fibered spaces. Then

$$h(\Theta_{\text{SF}}) \cong \mathbb{Z}^{\infty}.$$

An explicit basis for $h(\Theta_{\text{SF}})$ may be given by the images of the Brieskorn spheres

$$Y_i = h(\Sigma(2i+1, 4i+1, 4i+3))$$

for $i \geq 1$, together with the image of the Poincaré homology sphere $h(\Sigma(2, 3, 5))$.

The proof of Theorem 1.3.3 will include an algorithm for decomposing the image of any linear combination of Seifert fibered spaces in terms of the above basis. Clearly, this immediately gives a way of obstructing the existence of a homology cobordism between two such 3-manifolds. Several examples and applications are given in Section 5.5. We also prove that the Neumann-Siebenmann invariant can be computed from $h(Y)$ whenever Y is a linear combination of Seifert fibered spaces. This shows:

Theorem 1.3.4. [DS17, Corollary 1.3] The Neumann-Siebenmann invariant descends to a well-defined homomorphism

$$\bar{\mu} : \Theta_{\text{SF}} \rightarrow \mathbb{Z}.$$

As an application, we have the following non-torsion result:

Theorem 1.3.5. [DS17, Corollary 1.4] Let $Y \in \Theta_{\text{SF}}$. If $\mu(Y) = 1$, then Y is not torsion in $\Theta_{\mathbb{Z}}^3$.

In Chapter 6, we summarize the results of [Dai18a]. Our main theorem links the connected Heegaard Floer homology of any $Y \in \Theta_{\text{SF}}$ with the basis decomposition of Theorem 1.3.3:

Theorem 1.3.6. [Dai18a, Theorem 1.1] Let $Y \in \Theta_{\text{SF}}$. Then the connected Heegaard Floer homology $HF_{\text{conn}}(Y)$ can be explicitly computed from the basis decomposition $h(Y) \in \mathbb{Z}^{\infty}$. Conversely, the basis decomposition of $h(Y) \in \mathbb{Z}^{\infty}$ can be determined from $HF_{\text{conn}}(Y)$.

We use this to describe a possible avenue for finding a homology sphere that does not lie in Θ_{SF} . Throughout, our computations will be loosely informed by the combinatorial model for lattice homology developed by Némethi in [Ném08]. This will provide a pleasing geometric interpretation of the algebraic calculations carried out in this paper.

Finally, in Chapter 7, we briefly discuss what can be said in the setting of $\text{Pin}(2)$ -equivariant monopole Floer homology. The results of this section are taken from [Dai18b].⁵ Our main result is the following analogue of Theorem 1.3.1:

Theorem 1.3.7. [Dai18b, Theorem 1.1] Let Y be a (positively oriented) Seifert fibered homology sphere. Then the $\text{Pin}(2)$ -equivariant monopole Floer homology $\widetilde{HS}(-Y)$ can be computed from the lattice homology of Y .

We do not attempt to deal with the underlying $\text{Pin}(2)$ -chain complexes, since the algebra in the $\text{Pin}(2)$ -case is rather more complicated than in the involutive setting. However, we expect that a similar result as in Theorem 1.3.3 holds in the $\text{Pin}(2)$ -setting also. We likewise have the following analogue of Theorem 1.3.2:

Theorem 1.3.8. [Dai18b, Theorem 1.3] Let Y be a (positively oriented) Seifert fibered homology sphere. Then

$$\alpha(Y) = \begin{cases} d(Y)/2 & \text{if } d(Y)/2 = \bar{\mu}(Y) \pmod{2} \\ d(Y)/2 + 1 & \text{if } d(Y)/2 + 1 = \bar{\mu}(Y) \pmod{2} \end{cases}$$

and $\beta(Y) = \gamma(Y) = -\bar{\mu}(Y)$.

See also [Sto15, Corollary 1.2]. Theorem 1.3.8 was conjectured by Manolescu in [Man16], and is originally due to Stoffregen [Sto15], although with a rather different proof. Our proof of Theorem 1.3.8 is more general and extends to Seifert fibered rational homology spheres (as well as a larger class of plumbed 3-manifolds).

Conventions. For convenience of the reader, we list some of the most easily-confused conventions adopted in this paper. Unless otherwise specified, we always work over the field $\mathbb{F} = \mathbb{Z}/2\mathbb{Z}$. We say that a Seifert fibered space is *positively oriented* if it is oriented as the boundary of its canonical negative-definite plumbing. If M is a graded module, then we use $M[i]$ to denote the shifted module in which an element $x \in M$ with grading g now has grading $g - i$ (as an element of $M[i]$).

⁵The results of [Dai18b] are actually chronologically prior to the other works discussed in this thesis.

Chapter 2

Background Material

In this chapter, we give the necessary background required for the rest of the paper. In the first three sections, we give a brief overview of involutive Heegaard Floer homology and the related notions of local equivalence and connected Heegaard Floer homology. We then discuss lattice homology and indicate its relevance to the current project.

2.1 Involutive Heegaard Floer homology

We begin with a review of involutive Heegaard Floer homology, as defined by Hendricks and Manolescu in [HM17]. We assume that the reader is familiar with the basic set-up of Heegaard Floer homology; see for example [OS04a, OS04b, OS03a, OS06]. The exposition here is taken from [DM17, DS17]. For simplicity, we restrict ourselves to the case where our 3-manifold Y is an integer homology sphere. Note that this means Y has a single spin^c -structure $\mathfrak{s}$, which is necessarily self-conjugate.

Let $\mathcal{H} = (H, J)$ be a Heegaard pair for Y , consisting of a pointed Heegaard splitting $H = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, z)$ of Y , together with a family of almost-complex structures J on $\text{Sym}^g(\Sigma)$. Associated to $\mathcal{H}$, we have the Heegaard Floer complex $CF^-(\mathcal{H})$, which is a $\mathbb{Z}$ -graded, free $\mathbb{F}[U]$ -module generated by the intersection points $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$ in $\text{Sym}^g(\Sigma)$. There are also three other variants of this complex, related by

$$\widehat{CF} = CF^-(U=0), CF^\infty = U^{-1}CF^-, \text{ and } CF^+ = CF^\infty / CF^-.$$

We denote these by CF° for $\circ = \widehat{}, \infty, +, -$. If $\mathcal{H}$ and $\mathcal{H}'$ are two Heegaard pairs for Y with the same basepoint z , then by work of Juhász and Thurston [JT12], any sequence of Heegaard moves

relating $\mathcal{H}$ and $\mathcal{H}'$ defines a homotopy equivalence

$$\Phi(\mathcal{H}, \mathcal{H}') : CF^\circ(\mathcal{H}) \rightarrow CF^\circ(\mathcal{H}').$$

This assignment is itself unique up to chain homotopy, in the sense that any two sequences of Heegaard moves define chain-homotopic maps Φ ; see also [HM17, Proposition 2.3]. This justifies the use of the notation $CF^\circ(Y)$, rather than $CF^\circ(\mathcal{H})$. Taking the homology of $CF^\circ(Y)$ yields the Heegaard Floer homology $HF^\circ(Y)$ constructed by Ozsváth and Szabó in [OS04a], [OS04b].

Now consider the conjugate Heegaard pair $\overline{\mathcal{H}} = (\overline{H}, \overline{J})$. This is defined by reversing the orientation on Σ and interchanging the α - and β -curves to give the Heegaard splitting

$$\overline{H} = (-\Sigma, \beta, \alpha, z),$$

and taking the conjugate family $\overline{J}$ of almost-complex structures on $\text{Sym}^g(-\Sigma)$. The points of $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$ are in obvious correspondence with the points of $\mathbb{T}_\beta \cap \mathbb{T}_\alpha$, and J -holomorphic disks with boundary on $(\mathbb{T}_\alpha, \mathbb{T}_\beta)$ are in bijection with $\overline{J}$ -holomorphic disks with boundary on $(\mathbb{T}_\beta, \mathbb{T}_\alpha)$. This yields a canonical isomorphism

$$\eta : CF^\circ(\mathcal{H}) \rightarrow CF^\circ(\overline{\mathcal{H}}),$$

where here we have used the fact that $\mathfrak{s} = \overline{\mathfrak{s}}$ is necessarily self-conjugate. Note that this is *not* the map $\Phi(\mathcal{H}, \overline{\mathcal{H}})$ defined in the previous paragraph. Instead, defining

$$\iota = \Phi(\overline{\mathcal{H}}, \mathcal{H}) \circ \eta,$$

we obtain a chain map from $CF^\circ(\mathcal{H})$ to itself. In [HM17, Section 2.2] it is shown that ι is a homotopy involution and is independent (up to the notion of homotopy equivalence defined later in this section) of the choice of $\mathcal{H}$. The involutive Heegaard Floer complex is then defined to be the mapping cone

$$CFI^\circ(Y) = \left(CF^\circ(Y) \xrightarrow{Q(1+\iota)} Q \cdot CF^\circ(Y)[-1] \right). \quad (2.1)$$

Here, Q is a formal variable marking the right-hand copy of $CF^\circ(Y)$.¹ Taking the homology of $CFI^\circ(Y)$ yields the involutive Heegaard Floer homology $HFI^\circ(Y)$. For $\circ = \infty, +, -$, this may be thought of as a module over the ring $\mathbb{F}[U, Q]/(Q^2)$, where U and Q have degrees -2 and -1 , respectively. In [HM17, Proposition 2.8] it is shown that the quasi-isomorphism type of $CFI^\circ(Y)$ is an invariant of Y .

Using $HFI^-(Y)$, we may define a set of numerical homology cobordism invariants, called the *involutive Heegaard Floer correction terms*. Let $\mathcal{U}_r^-$ be the infinite U -tower $\mathbb{F}[U]$, shifted so that the element 1 has grading r . Recall that for Y an integer homology sphere, the usual Heegaard Floer

¹Note that Q has degree -1 . This means that the right-hand copy of $CF^\circ(Y)$ has the same grading as the usual $CF^\circ(Y)$ complex, while the left-hand copy of $CF^\circ(Y)$ is shifted up in grading by one (as is standard in the mapping cone construction).

homology decomposes as the direct sum of a single infinite tower and a U -torsion part:

$$HF^-(Y) = \mathcal{U}_r^- \oplus U\text{-torsion}.$$

The *Ozsváth-Szabó d -invariant* is then defined to be $r + 2$. Note that r is well-defined, even though the U -nontorsion summand is not canonically determined. Similarly, in the involutive setting, we have that

$$HFI^-(Y) = \mathcal{U}_l^- \oplus \mathcal{U}_u^- \oplus U\text{-torsion}$$

for some $\mathcal{U}_l^-$ lying in odd gradings and $\mathcal{U}_u^-$ lying in even gradings. Again, l and u are well-defined, even though the U -nontorsion summand is not canonically determined. Note that in sufficiently low gradings, multiplication by Q maps $\mathcal{U}_l^-$ isomorphically onto $\mathcal{U}_u^-$. The involutive Heegaard Floer correction terms are then defined by

$$\underline{d}(Y) = l + 1 \text{ and } \bar{d}(Y) = u + 2.$$

The shifts by one and two in these definitions are chosen so that $d = \underline{d} = \bar{d} = 0$ for $Y = S^3$. Like the Ozsváth-Szabó d -invariant, the involutive Floer correction terms are invariant under homology cobordism, although it should be noted that they are not homomorphisms; see [HM17, Theorem 1.3]. Moreover, we have

$$\underline{d}(Y) \leq d(Y) \leq \bar{d}(Y).$$

See Sections 4.2 and 4.3 for an example of some computations of HFI^- , as well as $\bar{d}$ and $\underline{d}$.

2.2 Local equivalence

In this paper, we will usually deal with the underlying chain complex $CF^-(Y)$ and the homotopy involution ι , rather than the mapping cone $CFI^-(Y)$. We formalize this in the following definition:

Definition 2.2.1. [HMZ18, Definition 8.1] An ι -complex is a pair (C, ι) , consisting of

- a $\mathbb{Z}$ -graded, finitely generated, free chain complex C over the ring $\mathbb{F}[U]$, where $\deg(U) = -2$.

We also require that there is a (grading-preserving) isomorphism

$$U^{-1}H_*(C) \cong \mathbb{F}[U, U^{-1}],$$

where $U^{-1}H_*(C)$ is supported in even gradings; and,

- a grading-preserving chain homomorphism $\iota : C \rightarrow C$, such that ι^2 is chain homotopic to the identity.

We say that two ι -complexes are *homotopy equivalent* if there exist grading-preserving, U -equivariant chain homotopy equivalences

$$F : C \rightarrow C', \quad G : C' \rightarrow C$$

that are homotopy inverses to each other, such that

$$F \circ \iota \simeq \iota' \circ F, \quad G \circ \iota' \simeq \iota \circ G,$$

where $\simeq$ denotes $\mathbb{F}[U]$ -equivariant chain homotopy.

Given a homology sphere Y , the pair $(CF^-(Y), \iota)$ is evidently an ι -complex. Note that if (C_1, ι_1) and (C_2, ι_2) are ι -complexes, then their tensor product $C_1 \otimes C_2$ is also an ι -complex, with homotopy involution given by $\iota_1 \otimes \iota_2$. The following important structural result about involutive Floer homology was proven by Hendricks, Manolescu, and Zemke in [HMZ18]:

Theorem 2.2.2. [HMZ18, Theorem 1.1] Suppose Y_1 and Y_2 are homology spheres, and let ι_1, ι_2 and ι denote the conjugation involutions on the Floer complexes $CF^-(Y_1)$, $CF^-(Y_2)$, and $CF^-(Y_1 \# Y_2)$, respectively. Then, the equivalence class of the ι -complex $(CF^-(Y_1 \# Y_2), \iota)$ is the same as that of

$$(CF^-(Y_1) \otimes_{\mathbb{F}[U]} CF^-(Y_2)[-2], \iota_1 \otimes \iota_2),$$

where $[-2]$ denotes a grading shift.

In this paper, we will be concerned with the following fundamental equivalence relation, discussed in [HMZ18, Section 8]:

Definition 2.2.3. [HMZ18, Definition 8.5] Two ι -complexes (C, ι) and (C', ι') are called *locally equivalent* if there exist (grading-preserving) chain maps

$$F : C \rightarrow C', \quad G : C' \rightarrow C$$

such that

$$F \circ \iota \simeq \iota' \circ F, \quad G \circ \iota' \simeq \iota \circ G,$$

and F and G induce isomorphisms between $U^{-1}H_*(C_1)$ and $U^{-1}H_*(C_2)$. We call a map F as above a *local map* from (C, ι) to (C', ι') , and similarly we refer to G as a local map in the other direction.

Definition 2.2.3 should be thought of as the algebraic analogue of homology cobordism. In [HMZ18], it is shown that if two homology spheres are homology cobordant, then their ι -complexes are locally equivalent. Moreover, it is shown that the set of ι -complexes modulo local equivalence forms a group:

Definition 2.2.4. [HMZ18, Proposition 8.8] Let $\mathfrak{J}$ be the set of ι -complexes up to local equivalence. This has a multiplication given by tensor product, which sends (the local equivalence classes of) two ι -complexes (C_1, ι_1) and (C_2, ι_2) to (the local equivalence class of) their tensor product complex $(C_1 \otimes C_2, \iota_1 \otimes \iota_2)$. The identity element of $\mathfrak{J}$ is given by the trivial complex (consisting of a single $\mathbb{F}[U]$ -tower starting in grading zero, together with the identity map on this complex). Inverses in $\mathfrak{J}$ are given by dualizing. See [HMZ18, Section 8] for further discussion.

In light of Theorem 2.2.2, we thus obtain a homomorphism

$$h : \Theta_{\mathbb{Z}}^3 \rightarrow \mathfrak{I}$$

given by sending Y to the local equivalence class of its (grading-shifted) ι -complex:

$$h(Y) = (CF^-(Y), \iota)[-2].$$

We can further restrict h to the subgroup of $\Theta_{\mathbb{Z}}^3$ generated by Seifert fibered spaces, which we denote by Θ_{SF} . As stated previously, one of the central results of this paper will be an explicit computation of $h(\Theta_{\text{SF}})$.

In order to gain some intuition for Definition 2.2.3, observe that if Y_1 and Y_2 are homology cobordant, then the cobordism maps coming from (non-involutive) Heegaard Floer theory induce grading-preserving isomorphisms after passing to $U^{-1}HF^-$. Thus, the main import of Definition 2.2.3 is the additional requirement that F and G homotopy commute with ι . This is illustrated by the following example:

Example 2.2.5. In Figure 2.1, we have displayed two ι -complexes, corresponding to S^3 and $\Sigma(2, 3, 7)$, respectively. (See Chapter 4 for justification.) The reader should verify that these are *not* locally equivalent. Note that there are obvious maps F and G between them which induce isomorphisms on $U^{-1}H_*$: for example, we may take $F(x) = v$, together with $G(v) = G(\iota v) = x$ and $G(\alpha) = 0$. However, F does not commute (even up to homotopy) with the action of ι . This shows that the involutive Heegaard Floer homology is powerful enough to distinguish S^3 and $\Sigma(2, 3, 7)$, even though their Ozsváth-Szabó d -invariants are equal.




Figure 2.1: Two ι -complexes corresponding to S^3 (left) and $\Sigma(2, 3, 7)$ (right). Note that x and v have the same grading. In both complexes, ι is an actual involution and is given by reflection across the obvious vertical axis. Vertical lines represent multiplication by U and arrows represent the action of ∂ ; both ι and ∂ are U -equivariant.

Example 2.2.6. In Figure 2.2, we have displayed two ι -complexes, corresponding to $\Sigma(2, 3, 7)$ and

$\Sigma(2, 5, 11)$, respectively. (See Chapter 4 for justification.) The reader should verify that these two ι -complexes are locally equivalent. The map F in one direction given by the obvious identification of the left-hand side as a subcomplex of the right-hand side; this sends $F(v) = w_2$ and $F(\alpha) = \beta_2$, extending ι - and U -equivariantly. The map G is given by sending

$$G(w_1) = Uv \quad \text{and} \quad G(w_2) = v$$

$$G(\beta_1) = 0 \quad \text{and} \quad G(\beta_2) = \alpha,$$

again extending ι - and U -equivariantly. The reader is encouraged to draw the homology of these two complexes and check that the induced maps send U -nontorsion elements to U -nontorsion elements. (Thus we obtain isomorphisms after inverting the action of U .)




Figure 2.2: Two ι -complexes corresponding to $\Sigma(2, 3, 7)$ (left) and $\Sigma(2, 5, 11)$ (right). Note that v and w_2 have the same grading. In both complexes, ι is an actual involution and is given by reflection across the obvious vertical axis. Vertical lines represent multiplication by U and arrows represent the action of ∂ ; both ι and ∂ are U -equivariant.

We will also occasionally talk of the involutive Floer correction terms of an (abstract) ι -complex (C, ι) . By this, we of course mean that one should form the mapping cone complex as in (2.1) and apply the definitions of $\bar{d}$ and $\underline{d}$ given in the previous section. In this paper, however, we will actually subtract two from these quantities when discussing the correction terms of an (abstract) ι -complex, in order to account for the grading shift in the definition of h . This convention means that $\bar{d}(Y) = \bar{d}(h(Y))$ and $\underline{d}(Y) = \underline{d}(h(Y))$. Similarly, we define the d -invariant of an abstract $\mathbb{F}[U]$ -complex by subtracting two from the usual algebraic definition, so that d is additive with respect to tensor products.

2.3 Connected Heegaard Floer homology

We now discuss the construction of connected Heegaard Floer homology, due to Hendricks, Hom, and Lidman in [HHL18]. The exposition here is taken from [Dai18a].

Definition 2.3.1. [HHL18, Definition 3.1] Let (C, ι) be an ι -complex. A *self-local equivalence* is a local map from (C, ι) to itself; that is, a (grading-preserving) chain map

$$f : C \rightarrow C,$$

such that

$$f \circ \iota \simeq \iota \circ f,$$

and f induces an isomorphism on homology after inverting the action of U .

Hendricks, Hom, and Lidman define a preorder on the set of self-local equivalences of (C, ι) by declaring $f \lesssim g$ whenever $\ker f \subseteq \ker g$. A self-local equivalence f is then said to be *maximal* if $g \gtrsim f$ implies $g \lesssim f$ for any self-local equivalence g . Maximal self-local equivalences should heuristically be thought of as producing a local equivalence between C and some very small subcomplex of C given by $\operatorname{im} f$. (This subcomplex is small because $\ker f$ is large.) Note, however, that since ι is a homotopy involution and f commutes with ι only up to homotopy, the action of ι need not preserve $\operatorname{im} f$, and similarly for the relevant homotopy maps. However, according to [HHL18, Lemma 3.7], we can modify ι to produce an actual homotopy involution ι_f on $\operatorname{im} f$ such that (C, ι) and $(\operatorname{im} f, \iota_f)$ are locally equivalent.

Hendricks, Hom, and Lidman further show that a maximal self-local equivalence always exists [HHL18, Lemma 3.3], and that any two maximal self-local equivalences f and g give homotopy equivalent ι -complexes $(\operatorname{im} f, \iota_f)$ and $(\operatorname{im} g, \iota_g)$ [HHL18, Lemma 3.8]. We can thus make the unambiguous definition:

Definition 2.3.2. [HHL18, Definition 3.9] Let (C, ι) be an ι -complex. The *connected complex* of (C, ι) , which we denote by $(C_{\text{conn}}, \iota_{\text{conn}})$, is defined to be $(\operatorname{im} f, \iota_f)$, where f is any maximal self-local equivalence. This is well-defined up to homotopy equivalence of ι -complexes.

Note that since (C, ι) and $(C_{\text{conn}}, \iota_{\text{conn}})$ are locally equivalent, the connected complex is in fact an invariant of the local equivalence class of (C, ι) [HHL18, Proposition 3.10]. Indeed, in this paper, we will think of the connected complex as the simplest possible local representative of (C, ι) . The connected homology is then defined to be (more or less) the usual homology of this complex:

Definition 2.3.3. [HHL18, Definition 3.13] Let (C, ι) be an ι -complex. The *connected homology* of (C, ι) is defined to be the U -torsion submodule of $H_*(C_{\text{conn}})$, shifted up in grading by one.

Here, $H_*(C_{\text{conn}})$ is the usual homology of C_{conn} as a $\mathbb{F}[U]$ -complex. The *connected Heegaard Floer homology* of Y is of course defined to be the connected homology of $(CF^-(Y), \iota)$. The grading shift by one is enacted so that HF_{conn} is a summand of HF_{red} , viewed as a quotient of HF^+ .

Example 2.3.4. The reader should check that the only self-local equivalences of the ι -complex of $\Sigma(2, 3, 7)$ are isomorphisms. Thus $HF_{\text{conn}}(\Sigma(2, 3, 7))$ is easily seen to consist of a single generator. Since the ι -complex of $\Sigma(2, 5, 11)$ is locally equivalent to that of $\Sigma(2, 3, 7)$ (as per Example 2.2.6), the connected homology of $\Sigma(2, 5, 11)$ is likewise one-dimensional.

In [HHL18, Proposition 4.1], it is shown that the connected homology behaves well under dualization. Note, however, that the connected homology does *not* behave well under tensor products. Roughly speaking, this is due to the fact that even if (C_1, ι_1) and (C_2, ι_2) are the simplest representatives of their respective local equivalence classes, the tensor product $(C_1 \otimes C_2, \iota_1 \otimes \iota_2)$ need not be. (For example, the tensor product of any ι -complex with its dual is locally equivalent to the trivial complex, and thus has zero connected homology.) In Chapter 6, we will compute the connected homology of linear combinations of Seifert fibered spaces. Note that *a priori*, the isomorphism class of the connected Floer homology (as an $\mathbb{F}[U]$ -module) is strictly weaker than the full local equivalence class of its underlying ι -complex. However, we will see in Chapter 6 that for elements of Θ_{SF} , the two contain the same information.

2.4 Lattice homology

We now turn to a discussion of lattice homology. This was first defined by Ozsváth and Szabó in [OS03b], but in this paper we will present a combinatorial re-formulation due to Némethi [Ném08]. There are several possible conventions depending on which flavor of lattice homology (or cohomology) we are considering; here, we present the theory corresponding to $HF^-(Y)$. See Sections 3.2 and 4.1 for a discussion of other variants.

Definition 2.4.1. Let Γ be a decorated tree, and denote the decoration of any vertex v in Γ by $m(v) \in \mathbb{Z}$. Let L_Γ be the integer lattice spanned by the vertices of Γ . We define a bilinear pairing on L_Γ by setting

$$(v, w) = \begin{cases} m(v) & \text{if } v = w, \\ 1 & \text{if } v \text{ and } w \text{ are connected by an edge, and} \\ 0 & \text{otherwise,} \end{cases}$$

and extending linearly to L_Γ . We call this pairing the *intersection form* associated to Γ . We denote by $W = W(\Gamma)$ the four-manifold obtained by attaching two-handles to B^4 along unknots which are linked (and framed) according to Γ . In this situation, we call $Y = \partial W$ a *plumbed 3-manifold*. Note that Y is oriented as the boundary of W .

Throughout, we will further suppose that the intersection form on L_Γ is negative definite; this is equivalent to the claim that Y is a rational homology sphere realized as the link of a normal surface singularity (see [Ném08] for further discussion). Since we will be dealing with integer homology spheres, we will also assume that the intersection form on L_Γ is unimodular. It is a standard fact that any Seifert fibered homology sphere (of the appropriate orientation) can be expressed as such a negative-definite plumbing; see for example [OS03b].

Definition 2.4.2. Let $Y = \partial W(\Gamma)$ be a unimodular, negative-definite plumbed manifold, as above. Fix any characteristic vector k on L_Γ . This means that $(k, x) \equiv (x, x) \pmod{2}$ for all $x \in L_\Gamma$. We define a *weight function* w on the set of vectors in L_Γ by

$$w(x) = (k, x) + (x, x).$$

This is a negative-definite quadratic form on L_Γ .² We extend w to all side-length-one lattice cubes (of any dimension) in L_Γ by setting

$$w(\square) = \min\{w(x) \mid x \text{ a vertex of } \square\}.$$

For each integer n , let S_n be the cubical complex formed by the union over all lattice cubes with weight greater than or equal to n . This should be thought of as a discretization of the (super)level set $w(x) \geq n$ in $L_\Gamma \otimes \mathbb{R}$, which is just the interior of an ellipsoid. (See Figure 2.3 for a schematic example.) Note that S_n is empty for n sufficiently large, while in general S_n is a subset of S_{n-1} .




Figure 2.3: A schematic example of successive cubical complexes S_n with n decreasing from left to right. The level sets $w(x) = n$ in $L_\Gamma \otimes \mathbb{R}$ are represented by dotted ellipses.

Roughly speaking, lattice homology captures the changing topology of successive cubical complexes S_n . We assemble this information into a $\mathbb{Z}$ -graded $\mathbb{F}[U]$ -module denoted by $\mathbb{H}^-(\Gamma, k)$, as follows. In its simplest form, lattice homology simply counts the number of connected components

²Our definition of w here differs from that in e.g. [DM17] by a factor of two. We have chosen this convention because it is more convenient for the rest of the paper. Note that since k is characteristic, w takes values in $2\mathbb{Z}$.

in each S_n :

Definition 2.4.3. Let Γ be a plumbing diagram, and let k be a characteristic vector for L_Γ . We define $\mathbb{H}^-(\Gamma, k)$ in grading n to be:

$$\mathbb{H}_n^-(\Gamma, k) = \begin{cases} H_0(S_n; \mathbb{F}) & \text{if } n \text{ is even} \\ 0 & \text{if } n \text{ is odd.} \end{cases}$$

For n very negative (and even), $\mathbb{H}_n^-(\Gamma, k)$ is one-dimensional, but if Γ has a sufficiently skew intersection form, then there may be values of n for which S_n has multiple connected components.³ (See Figure 2.3.) Multiplication by U is defined to be equal to the induced inclusion map

$$i_*: H_0(S_n; \mathbb{F}) \rightarrow H_0(S_{n-2}; \mathbb{F}).$$

Note that U has degree -2 . For a fixed Y , it is shown in [OS03b] that $\mathbb{H}^-(\Gamma, k)$ is independent of the choice of Γ and k (up to overall grading shift). One can also define $\mathbb{H}^-(\Gamma, k)$ as the homology of a chain complex in the usual sense; see [Ném08]. This will be discussed further in Section 3.2.

The fundamental theorem which will concern us here is the following:

Theorem 2.4.4 (Theorem 1.2 [OS03b]; Theorem 5.2.2 [Ném08]). Let Y be a Seifert fibered homology sphere, and let Γ be any negative-definite plumbing graph for Y . Then for any characteristic vector k , we have an isomorphism of $\mathbb{F}[U]$ -modules:

$$HF^-(Y) \cong \mathbb{H}^-(\Gamma, k)[-(|\Gamma| + k^2)/4 + 2].$$

In light of the fact that S_n has multiple connected components whenever the pairing on Γ is sufficiently skew, this provides a simple but pleasing link between the Heegaard Floer homology of Y and the “exoticness” of the intersection form on L_Γ . Theorem 2.4.4 actually holds in rather greater generality than for Seifert fibered homology spheres; this is discussed in Section 8.1.

The isomorphism of Theorem 2.4.4 has been used by several authors to prove results on the Heegaard Floer side; see for example [Kar14, KL15]. As a practical matter, it turns out that lattice homology can be used to compute the Heegaard Floer homology of Brieskorn spheres extremely efficiently in terms of the parameters (p, q, r) [CK14]. Lattice homology has also been utilized extensively by Némethi in his work on singularity theory. See [Ném07, Ném05] for an extended discussion from this point of view.

There is one more very important property of lattice homology which will concern us. If Γ is a plumbing graph equipped with a characteristic vector k , then we have the following symmetry:

³In general, more complicated versions of lattice homology take into account the higher homology groups of each S_n . In these cases, we do not define $\mathbb{H}_n^-(\Gamma, k)$ to be zero for n odd. However, it turns out that for all manifolds in the statement of Theorem 2.4.4, the relevant cubical complexes S_n consist of contractible connected components. See Section 3.2.

Definition 2.4.5. We define the *lattice involution* J_0 on L_Γ by

$$J_0x = -x - k.$$

Geometrically, this corresponds to reflection through the point $-k/2$. Clearly, J_0 extends to an involution on the set of side-length one lattice cubes (of any dimension) in L_Γ . (Note that our lattice cubes are unoriented.) It is straightforward to check that the weight function w (defined using k) is invariant under J_0 , and thus that J_0 extends to an involution from $\mathbb{H}^-(\Gamma, k)$ to itself.

The application of lattice homology in this paper will thus be twofold. Firstly, in addition to providing a convenient computational tool, Theorem 2.4.4 actually constrains the form of $HF^-(Y)$ whenever Y is a Seifert fibered space. More precisely, in the next section we will define the concept of a *graded root*; it will be evident that the lattice homology is always of this form. Secondly, in Section 4.1 we show that under the isomorphism of Theorem 2.4.4, the homological action of ι_* corresponds to the lattice involution J_0 . As we will see, these two facts taken together will allow us to compute the involutive Heegaard Floer homology of Seifert fibered spaces.

Chapter 3

Preliminary Notions

We now present two fundamental constructions that will be used repeatedly throughout the remainder of this paper. The exposition here is a compilation of background material taken from [DM17, DS17, Dai18a].

3.1 Graded roots and standard complexes

This material is taken from [DM17]. We begin with the following definition:

Definition 3.1.1. A *graded root* is an infinite tree R , together with a height function gr from the vertices of R to $2\mathbb{Z}$ such that:

- $\text{gr}(u) - \text{gr}(v) = \pm 2$ for any edge (u, v) ,
- $\text{gr}(u) < \max\{\text{gr}(v), \text{gr}(w)\}$ for all edges (u, v) and (u, w) with $v \neq w$,
- gr is bounded above, and
- $\text{gr}^{-1}(k)$ is finite for any $k \in 2\mathbb{Z}$, and in fact $|\text{gr}^{-1}(k)| = 1$ for all even $k \ll 0$.

Geometrically, a graded root is simply an upwards-opening tree with infinite stem below, as shown in Figure 3.1. We use the function gr to record the heights of vertices in R , with the convention that any pair of vertices connected by an edge differ by a height of two. We normalize gr so that all vertices have even height.

Definition 3.1.2. Given a graded root R , we define a graded $\mathbb{F}[U]$ -module $\mathbb{H}^-(R)$ by giving $\mathbb{H}^-(R)$ one generator for each vertex v of R , located in the same grading as $\text{gr}(v)$. (These are generators of $\mathbb{H}^-(R)$ over $\mathbb{F}$.) We define the U -action by specifying

$$U \cdot v = w \text{ if } (v, w) \text{ is an edge and } \text{gr}(v) - \text{gr}(w) = 2.$$

Clearly, U has degree -2 .

A graded root can thus be thought of as a graphical representation of an $\mathbb{F}[U]$ -module. Since R and $\mathbb{H}^-(R)$ evidently contain precisely the same information, we will sometimes abuse notation and refer to them interchangeably. In our case, we will be considering graded roots which have an additional symmetry corresponding to the action of the lattice involution J_0 . To this end, we define:

Definition 3.1.3. A *symmetric graded root* is a graded root R with an involution J_0 satisfying the following properties:

- $\text{gr}(v) = \text{gr}(J_0 v)$ for any vertex v ,
- (v, w) is an edge in R if and only if $(J_0 v, J_0 w)$ is an edge in R , and
- for every $k \in 2\mathbb{Z}$, there is at most one J_0 -invariant vertex v with $\text{gr}(v) = k$.

Geometrically, a symmetric graded root is simply a graded root which is symmetric under reflection about the obvious central vertical axis. It is clear that the action of J_0 on R may be viewed as an involution (which we also denote by J_0) on $\mathbb{H}^-(R)$ commuting with the action of U .

It is clear from the construction of lattice homology that $\mathbb{H}^-(\Gamma, k)$ takes the form of a symmetric graded root. (Note that the lattice homology comes with a preferred basis given by the connected components of each level set S_n .) Thus Theorem 2.4.4 implies:

Theorem 3.1.4 (Theorem 5.2.2 [Ném08]). Let Y be a Seifert fibered homology sphere. Then $HF^-(Y)$ takes the form of a symmetric graded root.

For the moment, we simply *define* the involution on the graded root to be the lattice involution J_0 (Definiton 2.4.5). However, in Section 4.1, we will show that this coincides with the homological action of ι_* on $HF^-(Y)$ under the isomorphism of Theorem 2.4.4.




Figure 3.1: Graded root (with leaf and angle labels) representing the Floer homology of $\Sigma(2, 7, 15)$. Vertices represent generators (over $\mathbb{F}$) and edges represent multiplication by U .

A graded root R has an obvious (free, finitely-generated) model chain complex which replicates its homology. Following [DM17, Section 4], we construct a free $\mathbb{F}[U]$ -complex whose homology is $\mathbb{H}^-(R)$, as follows. Let $v_1, v_2, \dots, v_n$ be the leaves of R , enumerated in left-to-right lexicographic order. We also enumerate by $\alpha_1, \alpha_2, \dots, \alpha_{n-1}$ the $n - 1$ upward-opening angles in left-to-right lexicographic order. (See Figure 3.1.) We denote by $\text{gr}(v_i)$ the grading of vertex v_i and by $\text{gr}(\alpha_i)$ the grading of the vertex supporting angle α_i .

The generators (over $\mathbb{F}[U]$) of our complex are given as follows. For each leaf v_i , we place a single generator in grading $\text{gr}(v_i)$, which by abuse of notation we also denote by v_i . Note that v_i is a generator of our complex over $\mathbb{F}[U]$, so that as an abelian group we are introducing an entire tower of generators $\mathbb{F}[U]v_i$. For each angle α_i , we similarly place a single generator in grading $\text{gr}(\alpha_i) + 1$, denoting this by α_i . (The gradings of these generators are motivated by the notion of a *geometric complex*, to appear in Section 3.2.) We define our differential to be identically zero on the v_i , and set

$$\partial\alpha_i = U^{(\text{gr}(v_i) - \text{gr}(\alpha_i))/2}v_i + U^{(\text{gr}(v_{i+1}) - \text{gr}(\alpha_i))/2}v_{i+1}$$

on the α_i , extending to the entire complex linearly and U -equivariantly. We call this complex the *standard complex of R* and denote it by $C_*(R)$. It is not difficult to see that the homology of $C_*(R)$ is precisely $\mathbb{H}^-(R)$: the angle generators α_i serve to “glue together” (on the level of homology) the U -towers generated by each of the v_i . See Figure 3.2.

If R is moreover symmetric, then there is an obvious involution J_0 on $C_*(R)$ whose induced action on $\mathbb{H}^-(R)$ is the desired symmetry. (Here, we have abused notation slightly and denoted our chain-level involution by J_0 also.) This sends $J_0v_i = v_{n+1-i}$ and $J_0\alpha_i = \alpha_{n-i}$, extending linearly and U -equivariantly. For convenience of notation, we define the *absolute value* of the index of a generator to be $\min(i, n + 1 - i)$ for leaf generators v_i and $\min(i, n - i)$ for angle generators α_i . Where no confusion is possible, we denote this by $|i|$, so that J_0 identifies generators of the same absolute value. It is clear that $(C_*(R), J_0)$ is an ι -complex.

The fundamental property which will allow our analysis to go through is the fact that up to homotopy, $(C_*(R), J_0)$ is the *unique* ι -complex model for $\mathbb{H}^-(R)$. More precisely, we have:

Lemma 3.1.5 (Lemma 4.2 [DM17]). Let R be a symmetric graded root, and suppose that (C', J') is any ι -complex for which:

1. $H_*(C') \cong \mathbb{H}^-(R)$ as $\mathbb{F}[U]$ -modules, and
2. The induced action of J' on $H_*(C')$ coincides with the action of J_0 on $\mathbb{H}^-(R)$ under the above isomorphism.




Figure 3.2: Standard complex corresponding to the symmetric graded root of Figure 3.1.

Then (C', J') is homotopy equivalent to the standard complex $(C_*(R), J_0)$.

Proof. We first use the routine fact that any two (freely generated) $\mathbb{F}[U]$ -complexes which have the same homology are homotopy equivalent to each other. This means that (forgetting about the involutions) C' and $C_*(R)$ are homotopy equivalent $\mathbb{F}[U]$ -complexes. Let $F : C' \rightarrow C_*(R)$ and $G : C_*(R) \rightarrow C'$ be homotopy inverses effecting this equivalence, and let $J : C_*(R) \rightarrow C_*(R)$ be given by

$$J = F \circ J' \circ G.$$

Then it is easily checked that J is a homotopy involution on $C_*(R)$ and that (C', J') and $(C_*(R), J)$ are homotopy equivalent ι -complexes. Note that since F and G can be chosen to effect the given isomorphism between $H_*(C')$ and $\mathbb{H}^-(R)$, we can assume that J and J_0 coincide on the level of homology. It thus suffices to show that J and J_0 are homotopy equivalent maps from $C_*(R)$ to itself, as this easily implies $(C_*(R), J)$ and $(C_*(R), J_0)$ are homotopy equivalent ι -complexes.

We construct a chain homotopy H between J and J_0 as follows. Let v be any homogenous chain of even grading in $C_*(R)$. Since v is a cycle and J and J_0 induce the same map on homology, we know that Jv and J_0v differ by a boundary. Because ∂ has zero kernel on the odd part of $C_*(R)$, this boundary has a unique odd-chain preimage under ∂ , which we denote by Hv . Thus we have $Jv + J_0v = \partial Hv$. Note that since J and J_0 are linear and U -equivariant, the uniqueness of Hv shows that H is a linear and U -equivariant map from the even chains to the odd chains of $C_*(R)$. We define our chain homotopy to be equal to H on the even part of $C_*(R)$ and identically zero on the odd part of $C_*(R)$.

Because ∂ is identically zero on the even part of $C_*(R)$, the above equality establishes the desired chain homotopy identity on even chains. Thus, let α be an odd chain. Putting $v = \partial\alpha$ yields the equality $J\partial\alpha + J_0\partial\alpha = \partial H\partial\alpha$. Since J and J_0 are chain maps, the left-hand side is equal to

$\partial(J\alpha + J_0\alpha)$, and since ∂ has zero kernel on the odd part of $C_*(R)$, this shows that $J\alpha + J_0\alpha = H\partial\alpha$. Having defined H to be zero on odd chains, this establishes the chain homotopy identity for odd chains, completing the proof. $\square$

Remark 3.1.6. More generally, let f and g be any two grading-preserving chain maps between standard complexes of graded roots. If f and g coincide on the level of homology, then the proof of Lemma 3.1.5 shows that f and g are chain homotopic. The reader may thus be (rightly) suspicious of Lemma 3.1.5, since two maps which coincide on homology generally are *not* chain homotopic. The crucial point here is that standard complexes come with a canonical splitting into even and odd chains, which is automatically preserved due to grading reasons. More precisely, suppose C_1 and C_2 are two arbitrary $\mathbb{F}[U]$ -complexes, and let f and g be maps between them that coincide on homology. Attempting to replicate the proof of Lemma 3.1.5, choose splittings $C_1 = Z_1 \oplus B_1$ and $C_2 = Z_2 \oplus B_2$, where $Z_i = \ker \partial$ and ∂ is injective on B_i . For any $z_1 \in Z_1$, we have $f(z_1) + g(z_1) = \partial b_2$ for some unique $b_2 \in B_2$. We thus set $H z_1 = b_2$; as above, this defines H on Z_1 . The difficulty comes from defining H on B_1 . Indeed, for any $b_1 \in B_1$, we have $f(\partial b_1) + g(\partial b_1) = \partial H(\partial b_1)$. Thus

$$\partial(f(b_1) + g(b_1)) = \partial H(\partial b_1).$$

However, even though ∂ has zero kernel on B_2 , we *cannot* conclude from this that $f(b_1) + g(b_1) = H\partial b_1$. This is because f and g may not preserve the chosen splitting; i.e., $f(b_1) + g(b_1)$ may not lie in B_2 .

The import of Lemma 3.1.5 is that we may talk interchangeably about the information contained in $C_*(R)$ and the information contained in $\mathbb{H}^-(R)$. In particular, it is now clear that in order to compute the ι -complex of a Seifert fibered space, it suffices to verify that the homological action of ι_* coincides with the reflection involution J_0 on $\mathbb{H}^-(R)$ under the isomorphism of Theorem 2.4.4. Lemma 3.1.5 will then imply that the model described above is homotopy equivalent to the one coming from Heegaard Floer theory.

3.2 Geometric complexes

We now take a brief interlude to describe a combinatorial trick which will be especially useful for organizing $\mathbb{F}[U]$ -complex data. The material here is taken from [DS17, Dai18a]. For the original motivation behind these ideas, the reader should consult the work of Némethi [Ném08].

Let C_{cell} be a cell complex. Suppose that we are given a function gr from the cells of C_{cell} to $2\mathbb{Z}$, with the property that whenever a cell e_{i-1} appears in the cellular boundary of e_i , we have

$\text{gr}(e_{i-1}) \geq \text{gr}(e_i)$. Let $C = C_{\text{cell}} \otimes \mathbb{F}[U]$. We turn C into a $\mathbb{Z}$ -graded chain complex over $\mathbb{F}[U]$ as follows. For any cell e_i of dimension i , we define the Maslov grading $M(e_i)$ of e_i to be

$$M(e_i) = \text{gr}(e_i) + i.$$

We extend this to all of C by declaring U to be of degree -2 . The differential on C is defined to be

$$\partial e_i = \sum_{e_{i-1} \in \text{cellular bdry of } e_i} U^{(\text{gr}(e_{i-1}) - \text{gr}(e_i))/2} e_{i-1}. \quad (3.1)$$

Here, we have simply taken the usual cellular differential and multiplied by the appropriate powers of U everywhere so as to be of degree -1 with respect to the Maslov grading. Note that each exponent of U is non-negative due to the condition on gr . We call a chain complex which is defined in this manner a *geometric complex*, and we refer to the underlying cell complex C_{cell} as the *skeleton* of C . Note that C comes equipped with an additional grading coming from the usual $\mathbb{Z}$ -grading on C_{cell} , which may be extended to all of C by declaring U to have degree zero (with respect to this grading). We refer to this as the *dimensional grading*.

Remark 3.2.1. More conceptually, all we are doing is presenting a convenient bookkeeping device for thinking about C . We view the data of C as separated into two distinct parts – first, the skeleton C_{cell} , which is a chain complex in the usual sense over $\mathbb{F}$; and second, the data consisting of the Maslov gradings of the generators of C , which is encoded in the function gr (together with the dimensional grading). We think of the skeleton as giving the “shape” of differential on C , with the correct powers of U in the differential being forced by the knowledge of the Maslov gradings.

In this paper, all of our complexes will come equipped with a cellular chain map J on C_{cell} with $J^2 = 1$, corresponding to reflection through some point in C_{cell} . If the grading function gr is invariant with respect to J , then by imposing U -equivariance we can extend J to a grading-preserving chain involution on all of C (which we also denote by J). If C_{cell} has trivial homology, then it is easy to check that the pair $(C = C_{\text{cell}} \otimes \mathbb{F}[U], J)$ forms an ι -complex. When we say that an ι -complex is locally equivalent to some geometric complex with involution J , we of course mean that our ι -complex is locally equivalent to (C, J) in the usual sense.

If C_1 and C_2 are two geometric complexes, then their tensor product is also a geometric complex, constructed as follows. The skeleton of $C_1 \otimes C_2$ is given by the usual cellular product of the skeleta of C_1 and C_2 , and the grading function gr is defined by setting $\text{gr}(x \times y) = \text{gr}(x) + \text{gr}(y)$ for any pair of cells x from C_1 and y from C_2 . It is evident that if J_1 and J_2 are involutions on C_1 and C_2 , then $J_1 \times J_2$ is an involution on $C_1 \otimes C_2$, which geometrically corresponds to reflection in each coordinate.

Example 3.2.2. If $C_*(R)$ is the standard complex associated to a graded root, then $C_*(R)$ is a geometric complex whose skeleton is (rather trivially) given by a line segment. Indeed, the twisted differential (3.1) is exactly the same as the differential given in Section 3.1, with leaf generators corresponding to 0-cells and angle generators corresponding to 1-cells. See Figure 3.3.

The main use of the geometric complex point of view will come from the following principle. Let X and Y be two geometric complexes. Suppose that we have a cellular map f from the skeleton of X to the skeleton of Y . Thus f takes each e -cell in the skeleton of X to a sum of e -cells in the skeleton of Y , and preserves the usual cellular differential.¹ We wish to know if f can be promoted to a true grading-preserving chain map between X and Y . Suppose that

$$f(x) = \sum_i y_i,$$

where x is a e -cell in the skeleton of X and the y_i are distinct e -cells in the skeleton of Y . Then there is an obvious choice of grading-preserving lift given by

$$\tilde{f}(x) = \sum_i U^{(\text{gr}(y_i) - \text{gr}(x))/2} y_i, \quad (3.2)$$

where now we think of x and the y_i as actual generators in X and Y and the same set of y_i appears in the sum. In order for this to be well-defined, we must check that each e -cell y_i appearing in $f(x)$ satisfies the inequality $\text{gr}(y_i) \geq \text{gr}(x)$. If this condition holds for every x in the skeleton of X , then f evidently lifts to a grading-preserving map from X to Y , and it is easily verified that this preserves the twisted differential (3.1). We formalize this in the following lemma:

Lemma 3.2.3 (Lemma 3.3 [DS17]). Let X and Y be two geometric complexes. Suppose that X_{cell} and Y_{cell} have trivial homology and are equipped with gr-preserving involutions J , so that X and Y are ι -complexes. Let f be a cellular map between X_{cell} and Y_{cell} satisfying the following conditions:

- (a) (Lifting condition.) For any cell $x \in X_{\text{cell}}$, we have $\text{gr}(y_i) \geq \text{gr}(x)$ whenever y_i appears in $f(x)$,
- (b) The map f is J -equivariant, and
- (c) The map f takes each 0-cell in X_{cell} to an odd number of 0-cells in Y_{cell} .

Then the lift (3.2) of f is a local map from X to Y .

Proof. All we need to verify is that the lifted map $\tilde{f}$ induces an isomorphism on homology after inverting the action of U . This easily follows from the fact that f takes each 0-cell in X_{cell} to an

¹Here, we abuse language slightly and refer to any chain map between the cellular chain complexes of skeleta as a “cellular map”, in order to distinguish these from chain maps between $\mathbb{F}[U]$ -complexes. Thus, e -cells are allowed to be taken to sums of e -cells and/or zero. In all of the examples considered in this paper, these maps will have a clear topological interpretation (which somewhat justifies the terminology).

odd number of 0-cells in Y_{cell} . □

We will thus often construct local maps between geometric complexes by describing cellular maps between their skeleta and verifying the lifting condition separately.

Example 3.2.4. The map F of Example 2.2.6 may be thought of as the inclusion of one interval into another. Similarly, the map G of Example 2.2.6 may be thought of as a “crushing map” which (roughly speaking) retracts one interval onto another. Although this example is rather trivial, we hope that in later sections (especially Sections 5 and 6) we can convince the reader of the usefulness of viewing computations in terms of geometric complexes.

Finally, we explain the connection between the notion of a geometric complex and the construction of lattice homology given in Section 2.4. Let S_n denote the cellular subcomplex of C_{cell} consisting of all cells with $\text{gr} \geq n$. For each (super)level set S_n , we form an absolutely graded vector space

$$H_{tot}(S_n; \mathbb{F})[-n] = \bigoplus_{i \geq 0} H_i(S_n; \mathbb{F})[-n].$$

This is just the usual homology of S_n (summed over all homological gradings), but with Maslov grading given by the homological grading plus n . That is, a cellular e -cycle in S_n has Maslov grading $e + n$. Note that we also have a dimensional grading given by the usual homological grading e . Then it is not difficult to check that:

Lemma 3.2.5. [Ném08, Theorem 3.1.12] There a U -equivariant isomorphism

$$H_*(C) \cong \bigoplus_{n \text{ even}} H_{tot}(S_n; \mathbb{F})[-n],$$

which preserves both the Maslov grading and the dimensional grading. Note that we only sum over even n , since gr takes values in $2\mathbb{Z}$. The action of U on the right is given by the induced inclusion map

$$i_* : H_{tot}(S_n; \mathbb{F}) \rightarrow H_{tot}(S_{n-2}; \mathbb{F}).$$

It may be useful for the reader to consider the case of (the standard complex of) a graded root, schematically displayed in Figure 3.3. Note that in this case, the homology of each (super)level set is supported in dimensional grading zero. We can thus think about a geometric complex either in terms of its skeleton and grading function, or in terms of sequences of level sets.

If we consider the cell complex corresponding to the standard cubical subdivision of $\mathbb{R}^n$ and define gr to be the weight function of Section 2.4, then we obtain a geometric complex which we denote by $\mathbb{CF}^-(\Gamma, k)$. The lattice homology as defined in Section 2.4 is then the part of the homology of




Figure 3.3: Graded root (left) and successive (super)level sets S_n (right), with n decreasing from top to bottom. The skeleton of the standard complex is given by the lowermost displayed S_n . The reader should check that applying the prescription of Lemma 3.2.5 to the right-hand side yields the graded root given on the left. Compare with Figure 3.2.

$\mathbb{CF}^-(\Gamma, k)$ lying in dimensional grading zero. According to [Ném08, Theorem 5.2.2], however, this is all of the homology of $\mathbb{CF}^-(\Gamma, k)$ in the case that $Y = \partial W(\Gamma)$ is a Seifert fibered space. In general, we thus define the lattice homology to be the full homology of $\mathbb{CF}^-(\Gamma, k)$.

Remark 3.2.6. We can form a plus version of the above construction by tensoring C_{cell} with $\mathbb{F}[U]/U\mathbb{F}[U]$ and still keeping the same differential (3.1), in obvious analogy with the construction of HF^+ . In the case of lattice homology, we denote this complex by $\mathbb{CF}^+(\Gamma, k)$. This will briefly be useful in Section 4.1, but in this paper we will usually work with the minus version.

Chapter 4

Involutive Floer Homology of Seifert Fibered Spaces

In this chapter, we present the results of [DM17]. In the first two sections, we compute the involutive Heegaard Floer homology $HFI^-(Y)$ for any Seifert fibered space Y . We then prove Theorem 1.3.2 regarding the involutive Floer correction terms $\bar{d}$ and $\underline{d}$. Finally, we introduce the notion of a *monotone root*, and use this to parameterize the set of local equivalence classes of Seifert fibered spaces. This is joint work with Ciprian Manolescu.

4.1 Identification of J_0

We begin by explaining why the homological action of ι_* coincides with the action of J_0 under the isomorphism of Theorem 2.4.4. To see this, we first sketch the original definition of lattice (co)homology due to Ozsváth and Szabó in [OS03b]. This is simply a combinatorial re-formulation of the theory described in Section 2.4, but we replicate it here for convenience of the reader.

As in Section 2.4, let Y be given by a negative-definite plumbing $W = W(\Gamma)$. We view W as a cobordism between $-Y$ and S^3 . Let $\text{Char}(\Gamma)$ denote the set of characteristic vectors on $H_2(W) \cong L_\Gamma$. Note that since any two characteristic vectors differ by an element of $2L_\Gamma$, $\text{Char}(\Gamma)$ is isomorphic to L_Γ as an affine space. Let $\mathcal{U}_0^+$ denote the graded $\mathbb{F}[U]$ -module $HF^+(S^3) = \mathbb{F}[U, U^{-1}]/U\mathbb{F}[U]$, and fix any element $y \in HF^+(-Y)$. We define a map $\phi_y : \text{Char}(\Gamma) \rightarrow \mathcal{U}_0^+$ by evaluation on y :

$$\phi_y(k) = F_{W, PD[k]}^+(y).$$

Here, $F_{W, PD[k]}^+$ is the Heegaard Floer cobordism map from $HF^+(-Y)$ to $HF^+(S^3) \cong \mathcal{U}_0^+$ associated

to the spin^c structure on W with $c_1 = PD[k]$. It is shown in [OS03b] that the map ϕ_y satisfies a set of “adjunction relations” relating the values of ϕ_y on different elements of $\text{Char}(\Gamma)$. More precisely, let b be a basis vector in L_Γ corresponding to a vertex of Γ , and let

$$n(k, b) = ((k, b) + (b, b))/2.$$

Then

$$U^{n(k, b)} \cdot \phi_y(k + 2b) = \phi_y(k), \quad \text{for } n(k, b) \geq 0$$

and

$$\phi_y(k + 2b) = U^{-n(k, b)} \cdot \phi_y(k), \quad \text{for } n(k, b) < 0.$$

The set of all maps from $\text{Char}(\Gamma)$ to $\mathcal{U}_0^+$ satisfying these relations forms a graded $\mathbb{F}[U]$ -module, which we denote temporarily by $\mathbb{H}^+(\Gamma)$. (For the definition of the absolute grading, see [OS03b, Section 1].) Theorem 1.2 of [OS03b] says that for Y a Seifert fibered space, the correspondence $y \mapsto \phi_y$ effects an $\mathbb{F}[U]$ -module isomorphism between $HF^+(-Y)$ and $\mathbb{H}^+(\Gamma)$.

We now sketch how $\mathbb{H}^+(\Gamma)$ is related to the definition of lattice homology given in Section 2.4. Fix a reference characteristic vector k_0 , so that $\text{Char}(\Gamma)$ is identified with L_Γ through the map sending $k_0 + 2x \in \text{Char}(\Gamma)$ to $x \in L_\Gamma$. Recall that our weight function on L_Γ is given by

$$w(x) = (k_0, x) + (x, x).$$

Here, we are defining w with respect to the chosen characteristic vector k_0 . Let $x \in L_\Gamma$, and let b be a basis vector in L_Γ corresponding to a vertex of Γ . Consider the lattice edge e_1 in L_Γ with endpoints $e_0 = x$ and $e'_0 = x + b$. A straightforward computation shows that

$$w(e'_0) - w(e_0) = (k_0, b) + (x, b) + (b, x) + (b, b) = 2n(k_0 + 2x, b).$$

Note that $k_0 + 2x$ is the element of $\text{Char}(\Gamma)$ identified with x . We can thus write the twisted differential (3.1) as

$$\partial e_1 = \begin{cases} e_0 + U^{n(k_0 + 2x, b)} e'_0 & \text{if } n(k_0 + 2x, b) \geq 0 \\ U^{-n(k_0 + 2x, b)} e_0 + e'_0 & \text{if } n(k_0 + 2x, b) < 0. \end{cases}$$

Comparing this with the definition of $\mathbb{H}^+(\Gamma)$, we see that under our identification of $\text{Char}(\Gamma)$ with L_Γ , the $\mathbb{F}[U]$ -module $\mathbb{H}^+(\Gamma)$ is (up to overall grading shift) precisely the zero-dimensional *cohomology* of the complex $\mathbb{CF}^+(\Gamma, k)$ defined in Remark 3.2.6.¹ Applying the usual duality statements in Heegaard Floer theory to the isomorphism $HF^+(-Y) \cong \mathbb{H}^+(\Gamma)$, we obtain (after appropriate grading shift) the isomorphism between $HF^-(Y)$ and $\mathbb{H}^-(\Gamma, k)$, as stated in Theorem 2.4.4.

The desired identification of ι_* now follows easily from the equivariance properties of the involution on Heegaard Floer homology with respect to cobordism maps:

¹It is thus perhaps slightly inconsistent to denote the cohomology (rather than the homology) of $\mathbb{CF}^+$ by $\mathbb{H}^+$. We have done so to follow the notation of [OS03b].

Theorem 4.1.1. [DM17, Theorem 3.1] Let Y be a Seifert fibered homology sphere. Then, under the isomorphism of Theorem 2.4.4, the induced action ι_* on $HF^-(Y)$ coincides with the reflection involution J_0 on $\mathbb{H}^-(\Gamma, k)$.

Proof. It is easy to check that the (dual) involution J_0 on $\mathbb{H}^+(\Gamma)$ is given by precomposing an element $\phi_y : \text{Char}(G) \rightarrow \mathcal{U}_0^+$ with the reflection symmetry $k \mapsto -k$. To check that ι_* on $HF^+(-Y)$ corresponds to J_0 , it thus suffices to show that

$$\phi_y(k) = \phi_{\iota_*(y)}(-k)$$

for any $y \in HF^+(-Y)$ and $k \in \text{Char}(G)$. In other words, we must establish the equality

$$F_{W,PD}^+[k] = F_{W,-PD}^+[k] \circ \iota_*.$$

This follows from [OS06, Theorem 3.6], which says that the cobordism maps (on the level of homology) commute with the conjugation involutions. We are also using here the fact that the involution on $HF^+(S^3)$ is the identity, which is clear because $HF_g^+(S^3)$ has dimension at most one in every grading g .

We have thus identified the involution ι_* on $HF^+(-Y)$ with the (dual) involution J_0 on $\mathbb{H}^+(\Gamma)$. It is shown in [HM17, Section 4.2] that the involutions ι_* on $HF^+(-Y)$ and $HF^-(Y)$ are taken to each other under the duality isomorphism. The desired characterization of the involution for the minus version follows. $\square$

Remark 4.1.2. It should be noted that the isomorphism between lattice homology and Heegaard Floer homology is established on the homological level. More precisely, [OS03b, Theorem 1.2] is proven by using the surgery exact sequence in Heegaard Floer homology and induction on the size of Γ , rather than by establishing a homotopy equivalence between the Heegaard Floer complex and the lattice complex. It is thus difficult to understand the chain-level action of ι directly, which is why we have instead used Lemma 3.1.5.

4.2 Computation of the involutive Floer homology

We have now shown that the ι -complex of any Seifert fibered space can be computed by finding the relevant graded root R , together with the involution on R coming from the lattice homology. This can be done extremely efficiently for any Brieskorn sphere $\Sigma(p, q, r)$ via the algorithm described by Can and Karakurt in [CK14]; their algorithm uses only the numerology of the semigroup generated by p , q , and r . If desired, it is also easy to directly compute the involutive Heegaard Floer homology using the mapping cone formula (2.1), as follows.

Let $\ker(1 + J_0)$ be the kernel of the map

$$1 + J_0 : \mathbb{H}^-(R) \rightarrow \mathbb{H}^-(R),$$

and let

$$\operatorname{coker}(1 + J_0) = \mathbb{H}^-(R) / \operatorname{im}(1 + J_0).$$

Note that both of these are graded $\mathbb{F}[U]$ -modules. Observing that $\operatorname{im}(1 + J_0) \subseteq \ker(1 + J_0)$, we have a map $\pi_* : \ker(1 + J_0) \rightarrow \operatorname{coker}(1 + J_0)$ given by composing the inclusion $\ker(1 + J_0) \rightarrow \mathbb{H}^-(R)$ with the quotient by $\operatorname{im}(1 + J_0)$. Then:

Theorem 4.2.1. [DM17, Theorem 1.1] Let Y be a Seifert fibered space with symmetric graded root R and involution J_0 . Then

$$HFI^-(Y) \cong \ker(1 + J_0)[-1] \oplus \operatorname{coker}(1 + J_0).$$

Multiplication by Q is identified with the map from left-hand summand onto the right-hand summand given by π_* .

Proof. Write the mapping cone complex of $C_*(R)$ as the complex $C_*(R)[-1] \otimes \mathbb{F}[Q]/(Q^2)$, together with the differential $\partial' = \partial + Q(1 + J_0)$. Recall that $C_*(R)$ splits into even and odd chains, which we denote by C_{even} and C_{odd} , respectively. Expanding the mapping cone complex with respect to this splitting, we obtain the complex

$$\begin{array}{ccc} C_{\text{odd}}[-1] & \xrightarrow{Q(1 + J_0)} & QC_{\text{odd}}[-1] \\ \downarrow \partial & & \downarrow \partial \\ C_{\text{even}}[-1] & \xrightarrow{Q(1 + J_0)} & QC_{\text{even}}[-1] \end{array}$$

with the action of ∂' on each term given by the sum of all outgoing arrows. Since ∂ is injective on C_{odd} , we see from this that $\ker \partial'$ is the direct sum of $QC_{\text{even}}[-1]$ and the submodule K of $C_{\text{even}}[-1] \oplus QC_{\text{odd}}[-1]$ defined by

$$K = \{v + Q\alpha : v \in C_{\text{even}} \text{ and } \alpha \in C_{\text{odd}} \text{ and } (1 + J_0)v = \partial\alpha\}.$$

After quotienting out by $\operatorname{im} \partial'$, the term $QC_{\text{even}}[-1]$ gives us a summand isomorphic to $\operatorname{coker}(1 + J_0)$. To understand what happens to K , observe that for a fixed even chain v , there exists an odd chain α such that $v + Q\alpha \in K$ if and only if $(1 + J_0)v$ is nullhomologous in $C_*(R)$. If such an α exists, it is moreover unique since ∂ is injective on odd chains. Identifying $v + Q\alpha \in K$ with v , it is not hard to check that after quotienting out by $\operatorname{im} \partial'$, the term K gives rise to a summand isomorphic to $\ker(1 + J_0)$. Moreover, it is clear that under this identification, the action of Q takes $\ker(1 + J_0)$ to $\ker(1 + J_0) / \operatorname{im}(1 + J_0) \subseteq \operatorname{coker}(1 + J_0)$, since a representative of the form $v + Q\alpha$ is mapped simply to Qv . $\square$

We graphically represent the involutive Heegaard Floer homology as follows. We begin by drawing a graded root for the $\mathbb{F}[U]$ -module $\text{coker}(1 + J_0)$. Passing from $\mathbb{H}^-(R)$ to $\text{coker}(1 + J_0)$ identifies pairs of vertices that are symmetric across the infinite stem, with the caveat that vertices lying on the stem itself remain in singleton equivalence classes. Since this identification is U -equivariant, we obtain a graded root whose corresponding homology is $\text{coker}(1 + J_0)$ by simply folding R in half across its stem, as in Figure 4.1.

Now consider $\ker(1 + J_0)$. As a vector space, this is isomorphic to $\text{coker}(1 + J_0)$, as it has the same rank in each grading. We thus depict $\ker(1 + J_0)$ similarly by folding R in half. However, the interpretation of the vertices in this half-root is now slightly different. As before, vertices lying on the infinite stem in the half-root are identified with their counterparts in R , but vertices not on the infinite stem correspond to sums of symmetric pairs of vertices in R . To see the effect that this has on the U -structure, suppose that v is a vertex of the latter kind such that in our original graded root, the product $w = Uv$ lies on the infinite stem. Then $J_0 w = w$, so $U(v + J_0 v) = 2w = 0$. Hence in our half-root representation of $\ker(1 + J_0)$, the vertices v and w are *not* connected by an edge. A graphical representation of $\ker(1 + J_0)$ is thus obtained by folding R in half and deleting each edge coming off of the infinite stem. Again, see Figure 4.1.

Finally, to compute the action of Q from $\ker(1 + J_0)$ to $\text{coker}(1 + J_0)$, observe that an element of $\ker(1 + J_0)$ consisting of the sum of two symmetric vertices maps to zero in $\text{coker}(1 + J_0)$, while a vertex lying on the infinite stem has non-zero image. Thus the action of Q is an isomorphism from the infinite stem of $\ker(1 + J_0)$ onto the infinite stem of $\text{coker}(1 + J_0)$, and is zero everywhere else.




Figure 4.1: Original graded root (left) and involutive Heegaard Floer homology (right) of $\Sigma(2, 7, 15)$. Deleted edges in $\ker(1 + J_0)$ are represented by dotted lines. Multiplication by Q is represented by the dashed lines.

4.3 Correction terms

We now turn to a computation of the involutive Heegaard Floer correction terms $\bar{d}$ and $\underline{d}$. Our goal will be to relate these to the Neumann-Siebenmann invariant, defined in [Neu80, Sie80]:

Definition 4.3.1. Let Y be a plumbed three-manifold with spin structure $\mathfrak{s}$ and plumbing graph Γ . Let L_Γ be the integer lattice spanned by the vertices of Γ . Amongst the characteristic vectors on L_Γ corresponding to spin^c structures limiting to $\mathfrak{s}$ on Y , there is a unique vector w_0 (called the *Wu vector*) whose coordinates in the natural basis of L_Γ are all zero or one. Define the *Neumann-Siebenmann invariant* of $(Y, \mathfrak{s})$ by

$$\bar{\mu}(Y, \mathfrak{s}) = \frac{1}{8} (\sigma(\Gamma) - w_0^2),$$

where $\sigma(\Gamma)$ is the signature of Γ and w_0^2 is the self-pairing of w_0 .

In the case that Y is an integer homology sphere, there is only one spin structure $\mathfrak{s}$ on Y , and we accordingly denote the Neumann-Siebenmann invariant by $\bar{\mu}(Y)$. In [Neu80] it is shown that $\bar{\mu}(Y)$ is an integer lift of the Rokhlin invariant (defined for plumbed three-manifolds) and is independent of the choice of plumbing. However, it is *not* known in general whether the Neumann-Siebenmann invariant is a homology cobordism invariant.

Before proceeding, it will be helpful for us to introduce the following combinatorial lemma regarding the Wu vector. The set of basis vectors having non-zero coefficient in w_0 is referred to as the *Wu set*. It is a standard fact that the Wu set always consists of pairwise nonadjacent vertices in Γ . For completeness, we give a proof of this below:

Lemma 4.3.2. Let Γ be any plumbing tree, and let w_0 be a characteristic vector for Γ such that all of the coordinates of w_0 in the natural basis of L_Γ are either zero or one. After appropriate permutation, we may assume that $w_0 = b_1 + b_2 + \cdots + b_\ell$ for some $0 \leq \ell \leq |\Gamma|$, where the b_i are basis vectors. Then $(b_i, b_j) = 0$ for all $1 \leq i \neq j \leq \ell$; that is, the Wu set for w_0 consists of pairwise nonadjacent vertices in Γ .

Proof. Let Γ' be the induced subgraph of Γ spanned by the vertices corresponding to $b_1, b_2, \dots, b_\ell$. Because w_0 is characteristic, we have that

$$(b_1 + b_2 + \cdots + b_\ell, b_i) \equiv (b_i, b_i) \pmod{2}$$

for all $1 \leq i \leq \ell$. Now, two distinct vertices have pairing one precisely when connected by an edge in Γ , and have pairing zero otherwise. Hence the above equality shows that in our induced subgraph Γ' , every vertex has an even number of adjacencies. But Γ (and thus Γ') has no cycles, so the only

way for this to be possible is for every vertex to be isolated. This proves the lemma. $\square$

We now turn to the proof of Theorem 1.3.2:

Proof of Theorem 1.3.2. Let Y be a Seifert fibered space equipped with a negative-definite plumbing graph Γ , and fix a characteristic vector k for L_Γ . Let R be the symmetric graded root associated to the (shifted) lattice homology $\mathbb{H}^-(\Gamma, k)[-(|\Gamma| + k^2)/4 + 2]$, so that the involutive Heegaard Floer homology of Y is given by the prescription of Section 4.2 applied to R .

It is straightforward to identify $\bar{d}$ and $\underline{d}$ in terms of the graphical presentation given in Figure 4.1. Indeed, it is clear from Figure 4.1 that the relevant generator for computing $\bar{d}$ is given by the uppermost vertex in the $\text{coker}(1 + J_0)$ -summand of $HFI^-(Y)$. This shows that $\bar{d}(Y) = d(Y)$. Likewise, the relevant generator for computing $\underline{d}$ is given by the uppermost vertex on the infinite stem of the $\ker(1 + J_0)$ -summand. Taking into account the grading shift in the definition of $\underline{d}$, this similarly yields

$$\underline{d}(Y) = g + 2,$$

where g is the grading of the uppermost J_0 -invariant vertex in the original root R .

Recall that when computing the lattice homology of Y , we are free to choose any characteristic vector k for L_Γ . We thus consider what happens when we use the Wu vector w_0 for k . The grading shift in the statement of Theorem 2.4.4 is then given by

$$-(|\Gamma| + w_0^2)/4 + 2 = -(-\sigma(\Gamma) + w_0^2)/4 + 2 = 2\bar{\mu}(Y) + 2.$$

It thus suffices to prove that for the graded root representing the *unshifted* lattice homology $\mathbb{H}^-(\Gamma, w_0)$, the J_0 -invariant vertex has grading zero. Indeed, after applying the above grading shift, this vertex will have grading $-2\bar{\mu}(Y) - 2$; adding two (as per the previous paragraph) then gives the desired result.

Recall that the vertices in $\mathbb{H}^-(\Gamma, w_0)$ of grading n (for n even) correspond to connected components of the (super)level set S_n in the language of Section 2.4. According to [Ném08, Theorem 4.3.3], if Y is a Seifert fibered space, then each connected component of S_n has zero higher homology. It follows from the Lefschetz fixed point theorem that a connected component is sent to itself under the action of J_0 if and only if it contains the point of reflection $-w_0/2$. Hence it suffices to show that $-w_0/2 \in S_n$ precisely when $n \leq 0$.

In order to illustrate the intuition behind the proof, suppose for a moment that w_0 is in fact zero (which is equivalent to the pairing of Γ being even). Then the weight w of $-w_0/2$ (as defined in Section 2.4) is zero, and this immediately implies the claim. If w_0 is not zero, however, then $-w_0/2$ does not lie in the lattice L_Γ , and in order to determine whether $-w_0/2$ lies in a sublevel set, we

must instead compute the weight of the smallest lattice cube of L_Γ containing $-w_0/2$. After an appropriate permutation of the natural basis of L_Γ , we may assume that

$$w_0 = b_1 + b_2 + \cdots + b_\ell$$

for some $0 \leq \ell \leq |\Gamma|$. Let C_ℓ be the vertices of the ℓ -dimensional lattice cube containing the sides $b_1, b_2, \dots, b_\ell$; that is, define

$$C_\ell = \{c_1 b_1 + c_2 b_2 + \cdots + c_\ell b_\ell : \text{each } c_i = 0 \text{ or } 1\}.$$

Then the smallest lattice cube containing $-w_0/2$ is ℓ -dimensional and is given by the translate $-w_0 + C_\ell$. Since the weight of a lattice cube is defined to be the minimum over the weights of its vertices, it suffices to show that $w(-w_0 + x) = 0$ for all $x \in C_\ell$. Now, since

$$w(-w_0 + x) = (-w_0 + x, -w_0 + x) + (-w_0 + x, w_0) = -(w_0, x) + (x, x),$$

this is equivalent to showing that $(w_0, x) = (x, x)$ for all $x \in C_\ell$. (Note that since w_0 is characteristic, this equality is always true modulo two, but it is strict equality that we must establish.) By Lemma 4.3.2, however, we have that $(b_i, b_j) = 0$ whenever $1 \leq i \neq j \leq \ell$. Expanding the pairings (w_0, x) and (x, x) immediately establishes the equality and completes the proof. $\square$

Remark 4.3.3. Since $\underline{d}$ is a homology cobordism invariant, Theorem 1.3.2 shows that $\bar{\mu}$ is a homology cobordism invariant for all Seifert fibered spaces. For other proofs of this result, see [FFU01, Sav02].

4.4 The monotone subroot algorithm

As we have seen in Example 2.2.6, there are certainly non-trivial local equivalences between different graded roots. In this section, we construct a special family of graded roots which will turn out to parameterize the possible local equivalence classes of Seifert fibered spaces.

Definition 4.4.1. Fix a positive integer n , and let $h_1, \dots, h_n$ and $r_1, \dots, r_n$ be sequences of even integers such that:

1. $h_1 > h_2 > \cdots > h_n$,
2. $r_1 < r_2 < \cdots < r_n$, and
3. $h_n \geq r_n$.

We construct a graded root $M = M(h_1, r_1; \dots; h_n, r_n)$ associated to this data as follows. First, we form the stem of our graded root by drawing a single infinite U -tower with uppermost vertex in grading r_n . For each $1 \leq i < n$, we then introduce a symmetric pair of leaves v_i and $J_0 v_i$ in grading h_i and connect these to the stem by using a pair of paths meeting the stem in degree r_i . If $h_n > r_n$,

we similarly introduce a pair of vertices v_n and $J_0 v_n$ in grading h_n and connect these to the stem at r_n . In the “degenerate” case that $h_n = r_n$, we instead declare v_n to be the J_0 -invariant vertex of grading r_n lying on the stem and take no further action. See Figure 4.2. We call a symmetric graded root constructed in this way a *monotone root*.




Figure 4.2: $M(4, -8; 2, -4; 0, -2)$ (left) and $M(4, -4; 0, 0)$ (right). We have suppressed marking the vertices.

Note that if M is a monotone root, then the only vertices of M which have valency greater than two lie on the stem and have grading equal to some r_i . Geometrically, a monotone root is a symmetric graded root such that:

1. Each branch coming off of the stem has exactly one leaf; and,
2. The gradings of the leaves of M form a strictly decreasing sequence as we move inwards towards the stem.

We refer to a monotone root $M = M(h_1, r_1; \dots; h_n, r_n)$ with $2n$ parameters as a *type- n root*. Note that due to the ordering of the r_i , our indexing of the leaves v_i is consistent with the lexicographic left-to-right ordering described in Section 3.1. Hence if $C_*(M)$ is the standard complex associated to M , then the gradings gr of the generators of $C_*(M)$ are precisely the parameters h_i and r_i .

We now claim that every symmetric graded root R has a monotone subroot M . In general there may be multiple such subroots, but here we construct one which in some sense captures as much of R as possible. We begin with some terminology. Let v be any vertex in R , and denote by γ_v the half-infinite path in R starting at v and running down the length of the stem. We define the *base* $b(v)$ of v to be the grading at which γ_v first meets the stem; i.e., the maximal grading of any vertex in the intersection of the stem and γ_v . We denote the collection of all vertices in R with a specified base b by C_b , and call this set the *cluster* of R based at b . Geometrically, C_b is just the union over all the branches of R meeting the stem at degree b . In general, we will use the term “cluster” to

refer to a non-trivial cluster (i.e., one with more than one vertex). If C_b is a non-trivial cluster, then it contains a pair of distinct leaves v and $J_0 v$ whose gradings are maximal amongst the vertices of C_b . We call these leaves the *tip(s)* of C_b . If C_b has more than one such pair, we arbitrarily select one to be the tip(s). Which one we select is not important, but we will use the fact that we have a specially marked pair of such vertices in each cluster.

Algorithm 4.4.2 (Monotone subroot algorithm). We construct a subset S of the leaves of R using the following “greedy” algorithm:

1. Let r be the degree of the uppermost J_0 -invariant vertex v in R . If the cluster C_r is trivial, then we add v to S ; otherwise, we add the two tips of C_r to S .
2. Next, let b be the greatest even integer strictly less than r for which C_b is non-trivial. If the tips of C_b have grading strictly greater than all of the vertices currently in S , then we add them to S ; otherwise, we do not.
3. Continue on down the stem in this manner, adding tips of clusters precisely when they have height strictly greater than any of the vertices currently in S .

We now define M to be the smallest subroot of R that contains the leaves in S after this process terminates. That is, identifying R with $\mathbb{H}^-(R)$, M is defined to be the $\mathbb{F}[U]$ -submodule spanned by the elements of S . It is clear that the isomorphism class of M (disregarding its embedding in R) is independent of our choice of tips. See Figure 4.3.




Figure 4.3: Graded root (left) and monotone subroot (right). We have suppressed marking the vertices.

We now show that every graded root is locally equivalent to its monotone subroot as given by Algorithm 4.4.2. It may be helpful for the reader to think of this as a systematic generalization of Example 2.2.6.

Theorem 4.4.3. [DM17, Theorem 6.1] Let R be a symmetric graded root. Then $(C_*(R), J_0)$ is locally equivalent to $(C_*(M), J_0)$, where M is the monotone subroot of R constructed above.

Proof. Let R be a symmetric graded root, and select a pair of tips from each non-trivial cluster as described above. For the purpose of exposition, it will be convenient to assume that R has the following technical property with regard to the ordering of its leaves. Let C_b be any non-trivial cluster in R . We require that out of all of the leaves of R contained in C_b , the tips of C_b are closest to the vertical line of symmetry in the lexicographic ordering (i.e., have indices closest to n). Such a property can easily be achieved by permuting the order of the leaves within each cluster in such a way that does not change the isomorphism class of R . If this property is satisfied, we say that R is *nicely ordered*.

In order to provide some intuition for the proof, we first construct a pair of maps between $\mathbb{H}^-(M)$ and $\mathbb{H}^-(R)$ on the homological level. There is an obvious $\mathbb{F}[U]$ -equivariant map $F : \mathbb{H}^-(M) \rightarrow \mathbb{H}^-(R)$ given by the inclusion of M into R . To define a map $G : \mathbb{H}^-(R) \rightarrow \mathbb{H}^-(M)$, let v be any vertex of R lying to the left of the vertical line of symmetry. If v already lies in M , then we define $G(v) = v$. Otherwise, we define $G(v)$ to be the closest vertex in M with grading $\text{gr}(v)$ that lies to the right of v . See Figure 4.4.




Figure 4.4: Homological action of G . Vertices not marked with an arrow are fixed by G . The monotone subroot in question is spanned by v_3 and v_5 .

This defines G for vertices to the left of the stem; we extend G to all the vertices of R in the obvious J_0 -equivariant way. The fact that M is monotone (as well as the assumption that R is nicely ordered) ensures that G is equivariant with respect to the action of U . We think of G as “folding” R into M by collapsing all of the vertices of R outside of M towards the center line. It is clear that F and G take U -nontorsion elements in $\mathbb{H}^-(M)$ to U -nontorsion elements in $\mathbb{H}^-(R)$, and thus induce isomorphisms after inverting the action of U .

We now show that F and G are induced by actual J_0 -equivariant chain maps between $C_*(M)$ and $C_*(R)$. Let $v_1, \dots, v_n$ be the generators of $C_*(R)$ lying to the left of the center line (inclusive), and let $i_1 < i_2 < \dots < i_k = n$ be the indices of the leaves in the left half of R which lie in S . In

order to remember the fact that M is a subroot of R , we enumerate the even generators of $C_*(M)$ by $v_{i_1}, \dots, v_{i_k}$ (and their reflections), rather than by using the indices 1 through k . We denote the odd generators of $C_*(M)$ likewise by $\beta_{i_1}, \dots, \beta_{i_{k-1}}$ (and their reflections), together with possibly a J_0 -invariant generator β_n .

We now construct a chain map $f : C_*(M) \rightarrow C_*(R)$. On the even part of $C_*(M)$, we define this to be the inclusion sending v_{i_j} as a generator of $C_*(M)$ to v_{i_j} as a generator of $C_*(R)$. For β_{i_j} an odd generator of $C_*(M)$, we define

$$f(\beta_{i_j}) = \sum_{s=i_j}^{i_{j+1}-1} U^{(\text{gr}(\alpha_s) - \text{gr}(\beta_{i_j}))/2} \alpha_s$$

for $1 \leq j < k$. One can check that $\text{gr}(\beta_{i_j}) = \text{gr}(\alpha_{i_j}) \leq \text{gr}(\alpha_s)$ for all $i_j \leq s < i_{j+1}$ due to the monotonicity of M and the fact that R is nicely ordered. If necessary, we also define $f(\beta_n) = \alpha_n$. This defines f on generators of $C_*(R)$ to the left of the center line; we extend f to generators on the right side of the center line in the obvious J_0 -equivariant way. Then f gives a J_0 -equivariant chain map from $C_*(M)$ to $C_*(R)$ which induces the map F on homology. The reader may find it helpful to consult Figure 4.5.

Finally, we construct a chain map $g : C_*(R) \rightarrow C_*(M)$. Let v_s be an even generator of $C_*(R)$ lying to the left of the stem, and let i_j be the least element of $\{i_1, \dots, i_k\}$ greater than or equal to s . Again due to the structure of M and R , it is not hard to see that $\text{gr}(v_s) \leq \text{gr}(v_{i_j})$. We thus define

$$g(v_s) = U^{(\text{gr}(v_{i_j}) - \text{gr}(v_s))/2} v_{i_j},$$

extending J_0 -equivariantly for generators to the right of the center line. Now let α_s be an odd generator of $C_*(R)$. If s is an element of $\{i_1, \dots, i_{k-1}\}$ (or if α_n is J_0 -invariant and $s = n$), we define $g(\alpha_s) = \beta_s$; otherwise, we define $g(\alpha_s) = 0$ (again extending to all of $C_*(R)$ J_0 -equivariantly). We leave it to the reader to check that g is a chain map which induces G . See Figure 4.5. $\square$

Example 4.4.4. In Figure 4.5, we have displayed a schematic example of the maps f and g in the proof of Theorem 4.4.3. Here, we have

$$f(v_1) = v_1 \quad \text{and} \quad f(v_3) = v_3$$

$$f(\beta_1) = \alpha_1 + U^2 \alpha_2,$$

and

$$g(v_1) = v_1, \quad g(v_2) = v_3, \quad \text{and} \quad g(v_3) = v_3$$

$$g(\alpha_1) = \beta_1 \quad \text{and} \quad g(\alpha_2) = 0.$$

Remark 4.4.5. It may be instructive for the reader to try to re-interpret the maps f and g in the

proof of Theorem 4.4.3 as maps between intervals, as in Example 3.2.4.




Figure 4.5: Half of a graded root R (left), together with (half of) the complexes $C_*(M)$ and $C_*(R)$ (right). In this example, $i_1 = 1$ and $i_2 = 3$, so M is spanned by v_1 (and $J_0 v_1$) and v_3 .

Finally, it turns out that any two monotone roots which are locally equivalent must be identical:

Theorem 4.4.6. [DM17, Theorem 6.2] Let M and M' be two monotone roots with involutions J_0 and J'_0 . If $(C_*(M), J_0)$ and $(C_*(M'), J'_0)$ are locally equivalent, then $M = M'$.

This will actually be implied by other results established later in this paper, although it can also be proven directly. We omit the direct proof here; for details, see [DM17, Section 6].

We thus see that the local equivalence classes of Seifert fibered spaces are parameterized precisely by the set of monotone roots, which are described by finite tuples of integers as in Definition 4.4.1. To determine the local equivalence class of a Seifert fibered space, we thus: (a) compute the symmetric graded root using lattice homology, and then (b) extract the monotone subroot specified by Algorithm 4.4.2. This justifies our claim in Theorem 1.3.1 that the lattice homology of Y computes its local equivalence class.

Remark 4.4.7. Note that if $M = M(h_1, r_1; \dots; h_n, r_n)$ is a monotone root, then the proof of Theorem 1.3.2 shows $\bar{d}(M) = h_1$ and $\underline{d}(M) = r_n$. The involutive Floer correction terms of a Seifert fibered space can thus immediately be read off from the parameterization of its local equivalence class.

Chapter 5

The Structure of $h(\Theta_{\text{SF}})$

In this chapter, we present the results of [DS17]. In the first four sections, we prove that $h(\Theta_{\text{SF}}) \cong \mathbb{Z}^\infty$ and give an explicit algorithm for decomposing elements of $h(\Theta_{\text{SF}})$ in terms of the basis of Theorem 1.3.3. Our approach relies on an in-depth analysis of tensor products of type-one roots. In the last section, we give several examples and list some further results and applications from [DS17]. This is joint work with Matthew Stoffregen.

5.1 Good families

We now introduce a special family of monotone roots that will play a fundamental role over the course of the next several sections.

Definition 5.1.1. Let $M = M(h_1, r_1)$ be a type-one root. Following [Sto15, Fact 5.6], it will be convenient for us to define¹

$$\tilde{\delta}(M) = h_1 - r_1.$$

Now let $\mathcal{M} = \{M_i\}_{i \geq 1}$ be a family of type-one roots such that $\tilde{\delta}(M_i) = 2i$ for all $i \geq 1$. If this holds, then we call $\mathcal{M}$ a *good family of roots*. Note that there are two obvious choices for good families of roots. These are given by $\mathcal{X} = \{X_i\}_{i \geq 1}$ and $\mathcal{Y} = \{Y_i\}_{i \geq 1}$, where

$$X_i = M(0, -2i)$$

and

$$Y_i = M(2i, 0).$$

¹This convention differs from $\tilde{\delta}$ in [Sto17] in that the present $\tilde{\delta}$ is twice that of [Sto17].

Example 5.1.2. The family $\mathcal{Y}$ has the benefit that it is explicitly realized by the Brieskorn spheres

$$\Sigma(2i+1, 4i+1, 4i+3)$$

for $i \geq 1$. More precisely, it is a consequence of [Sto17, Theorem 1.8] that we have the local equivalence

$$h(\Sigma(2i+1, 4i+1, 4i+3)) = M(2i, 0),$$

as displayed in Figure 5.1. Hence as i varies, we obtain all of the elements of $\mathcal{Y}$. Note that we can similarly realize $\mathcal{X}$ by taking connected sums with $\Sigma(2, 3, 5)$.




Figure 5.1: Monotone roots $Y_i = h(\Sigma(2i+1, 4i+1, 4i+3))$.

We now show that any good family of roots is linearly independent in $\mathfrak{J}$. In particular, this proves that the Brieskorn spheres $\Sigma(2i+1, 4i+1, 4i+3)$ are linearly independent in $\Theta_{\mathbb{Z}}^3$. This was first established by Stoffregen in [Sto17] using $\text{Pin}(2)$ -equivariant Seiberg-Witten Floer homology; a similar argument was given in [DM17] using involutive Heegaard Floer homology. Here, we present the latter proof. This establishes linear independence as a consequence of the computation of $\bar{d}$ and $\underline{d}$ for connected sums of type-one roots. We thus begin by giving a simplified model complex for the local equivalence classes of such sums.

Let B^n be the standard n -dimensional unit ball, and let J_0 be the usual involution on B^n given by reflection through the origin. There is an obvious J_0 -equivariant cellular decomposition of B^n , consisting of exactly one symmetric pair of cells in each dimension 0 through $n-1$, and a single J_0 -invariant cell in dimension n . We denote the cells lying in dimension i by e_i and $J_0 e_i$, with the understanding that $e_n = J_0 e_n$. The cellular boundary operator is given by

$$\partial(e_i) = \partial(J_0 e_i) = e_{i-1} + J_0 e_{i-1},$$

with the convention that $e_i = J_0 e_i = 0$ for $i < 0$. See Figure 5.2.

Now let $d \in 2\mathbb{Z}$, and let $(\Delta_1, \Delta_2, \dots, \Delta_n)$ be a sequence of non-negative even integers. We associate a geometric complex to these parameters as follows. The skeleton of our complex is defined to be the above cellular decomposition of B^n . The function gr is defined by setting

$$\text{gr}(e_0) = \text{gr}(J_0 e_0) = d$$

and

$$\text{gr}(e_i) = \text{gr}(J_0 e_i) = d - \left(\sum_{j=1}^i \Delta_j \right)$$

for $1 \leq i \leq n$. Note that this means the differential on our complex is given by

$$\partial(e_i) = \partial(J_0 e_i) = U^{\Delta_i/2}(e_{i-1} + J_0 e_{i-1}).$$

We call such a complex a *spherical complex* and denote it by

$$S = S(d, n; \Delta_1, \dots, \Delta_n).$$

See Figure 5.2.




Figure 5.2: Spherical complex $S(d, n; \Delta_1, \dots, \Delta_n)$ with $n = 2$.

Lemma 5.1.3. [DS17, Lemma 5.2] Let $M_1, \dots, M_n$ be a sequence of type-one roots parameterized by $M_i = (h^i, r^i)$ for $1 \leq i \leq n$. Without loss of generality, suppose that

$$\tilde{\delta}(M_1) \geq \tilde{\delta}(M_2) \geq \dots \geq \tilde{\delta}(M_n).$$

Then we have the local equivalence

$$M_1 + \dots + M_n = S(d, n; \tilde{\delta}(M_1), \dots, \tilde{\delta}(M_n)),$$

where $d = h^1 + \dots + h^n$.

Proof. We proceed by induction on n . If $n = 1$, then the spherical complex in question is the same as the standard complex defined in Section 3.1. To establish the inductive step, it suffices to prove the local equivalence

$$M(h, r) + S(d, n; \Delta_1, \dots, \Delta_n) = S(d + h, n + 1; h - r, \Delta_1, \dots, \Delta_n),$$

where without loss of generality we have assumed that $h - r \geq \Delta_i$ for all Δ_i . Denote the cells of $M(h, r)$ by v , $J_0 v$, and $\alpha = J_0 \alpha$. For convenience, we denote the cells of $S(d, n; \Delta_1, \dots, \Delta_n)$ by x_i for $0 \leq i \leq n$, and the cells of $S(d + h, n + 1; h - r, \Delta_1, \dots, \Delta_n)$ by y_i for $0 \leq i \leq n + 1$. Note that $x_n = J_0 x_n$ and $y_{n+1} = J_0 y_{n+1}$.

We now apply Lemma 3.2.3. The skeleta of the product complex on the left and the spherical complex on the right are displayed in Figure 5.3. To define a J_0 -equivariant cellular map from the

left to the right, let

$$f(v \times x_0) = f(v \times J_0 x_0) = y_0,$$

$$f(v \times x_i) = f(v \times J_0 x_i) = 0 \text{ for } 1 \leq i \leq n, \text{ and}$$

$$f(\alpha \times x_i) = y_{i+1} \text{ for } 0 \leq i \leq n,$$

extending J_0 -equivariantly. To interpret this geometrically, note that the complex on the left in Figure 5.3 is a cylinder; f corresponds (roughly) to crushing each end of this cylinder to a single point. It is easily checked that this preserves the cellular differential. To see that f satisfies the lifting condition, we check that f preserves gr wherever it is non-zero:

$$\text{gr}(v \times x_0) = \text{gr}(v \times J_0 x_0) = h + d = \text{gr}(y_0)$$

and

$$\text{gr}(\alpha \times x_i) = r + \left(d - \sum_{j=1}^i \Delta_j \right) = (d + h) - \left((h - r) + \sum_{j=1}^i \Delta_j \right) = \text{gr}(y_{i+1}).$$

This gives the local map in one direction. To define the map in the other direction, set

$$g(y_i) = v \times x_i + \alpha \times J_0 x_{i-1} \text{ for } 0 \leq i \leq n, \text{ and}$$

$$g(y_{n+1}) = \alpha \times x_n,$$

again extending J_0 -equivariantly. Geometrically, this effects a J_0 -equivariant homeomorphism from the solid ball onto the solid cylinder. The reader may explicitly verify that this preserves the cellular differential (over $\mathbb{F}$) by computing

$$\begin{aligned} \partial g(y_i) &= \partial(v \times x_i + \alpha \times J_0 x_{i-1}) \\ &= v \times (x_{i-1} + J_0 x_{i-1}) + (v + J_0 v) \times J_0 x_{i-1} + \alpha \times (x_{i-2} + J_0 x_{i-2}) \\ &= v \times x_{i-1} + J_0 v \times J_0 x_{i-1} + \alpha \times x_{i-2} + \alpha \times J_0 x_{i-2} \end{aligned}$$

and

$$\begin{aligned} g\partial(y_i) &= g(y_{i-1} + J_0 y_{i-1}) \\ &= v \times x_{i-1} + \alpha \times J_0 x_{i-2} + J_0 v \times J_0 x_{i-1} + \alpha \times x_{i-2} \end{aligned}$$

for $1 \leq i \leq n$. A similar computation holds for y_{n+1} , utilizing the fact that $J_0 x_n = x_n$. To see that g satisfies the lifting condition, consider the two summands of $g(y_i)$ given by $v \times x_i$ and $\alpha \times J_0 x_{i-1}$. The grading of the latter is equal to $\text{gr}(y_i)$ (as computed previously). The grading of the former is given by

$$\text{gr}(v \times x_i) = h + d - \left(\sum_{j=1}^i \Delta_j \right).$$

This is greater than or equal to $\text{gr}(y_i)$, using the fact that $h - r \geq \Delta_i$ for all Δ_i . Applying Lemma 3.2.3

establishes the local equivalence and completes the proof. $\square$




Figure 5.3: Skeleta of the complexes considered in the proof of Lemma 5.1.3.

We stress that Lemma 5.1.3 only applies to connected sums of type-one roots with positive coefficients. In Chapter 6, we will describe a model complex for the local equivalence class of a general linear combination of type-one roots. However, Lemma 5.1.3 suffices for the application currently at hand:

Lemma 5.1.4. Let $S = S(d, n; \Delta_1, \dots, \Delta_n)$ be a spherical complex. Then the involutive Floer correction terms are given by $\bar{d}(S) = d(S) = d$ and $\underline{d}(S) = d - \Delta_1$.

Proof. This is a straightforward computation from the definitions. It is useful to note that e_0 and $J_0 e_0$ are U -nontorsion in $H_*(S)$, and that any U -nontorsion element in $H_*(S)$ must in fact be homologous to U -powers of one of these. The relevant cycle for computing $\bar{d}$ is then easily seen to be $Q \cdot e_0$. The relevant cycle for computing $\underline{d}$ is given by $U^{\Delta_1/2} e_0 + Q \cdot e_1$. $\square$

Lemma 5.1.5. Let $M_1, \dots, M_n$ be any sequence of type-one roots. Then

$$d(M_1 + \dots + M_n) - \underline{d}(M_1 + \dots + M_n) = \max \bar{\delta}(M_i).$$

Proof. This follows immediately from Lemmas 5.1.3 and 5.1.4. $\square$

Following the argument of [Sto15, Theorem 1.6], we thus finally obtain:

Theorem 5.1.6. [DM17, Theorem 1.7] Let $\mathcal{M} = \{M_i\}_{i \geq 1}$ be a good family of roots. Then the elements of $\mathcal{M}$ are linearly independent.

Proof. Suppose that we had some linear relation between the M_i . Moving terms around as necessary, we obtain a local equivalence

$$c_{i_1} M_{i_1} + \dots + c_{i_p} M_{i_p} = c_{j_1} M_{j_1} + \dots + c_{j_q} M_{j_q},$$

where all the coefficients are positive and no index appears on both the left and the right. Now, since d and $\underline{d}$ are both local equivalence invariants, their difference $d - \underline{d}$ must be the same on the left-hand side as it is on the right. Applying Lemma 5.1.5 then gives a contradiction, since in our case the $\tilde{\delta}(M_i)$ are different for distinct M_i . $\square$

5.2 The swapping lemma

We now begin our study of linear combinations of graded roots in earnest. As a warm-up exercise, we give an example of a non-trivial local equivalence between a linear combination of monotone roots:

Example 5.2.1. Let x and y be two even integers for which $x < y < 0$. We establish the local equivalence

$$M(0, x; y, y) + M(0, x - y) = M(0, x).$$

One can of course instead write $M(0, x; y, y) = M(0, x) - M(0, x - y)$; this expresses the type-two root on the left as the tensor product of a type-one root and the dual of a type-one root. The monotone roots in question are displayed in Figure 5.4, together with their standard complexes. We have also displayed the geometric realization of $M(0, x; y, y) + M(0, x - y)$ on the left-hand side in Figure 5.5, together with the gradings of its cells. For convenience, we refer to the skeleton of $M(0, x; y, y) + M(0, x - y)$ as A and the skeleton of $M(0, x)$ as B .




Figure 5.4: Monotone roots of Example 5.2.1 (above) and the geometric realizations of their standard complexes (below). Gradings of leaf and angle generators are marked in both parts of the diagram.

We now describe a local map from $M(0, x; y, y) + M(0, x - y)$ to $M(0, x)$ by defining a cellular




Figure 5.5: Geometric realizations of $M(0, x; y, y) \otimes M(0, x - y)$ (left) and $M(0, x)$ (right).

map between their skeleta. We send the bottom-left, top-left, and top-middle 0-cells of A to the left-hand 0-cell in B . We similarly map the bottom-right, top-right, and bottom-middle 0-cells of A to the right-hand 0-cell in B . Note that since $y < 0$, this correspondence satisfies the lifting condition of Lemma 3.2.3. We now send the three red 1-cells in A to the single 1-cell in B , and send all the other 1- and 2-cells to zero. It is easily checked that this is a J_0 -equivariant cellular map (over $\mathbb{F}$). Applying Lemma 3.2.3, we obtain the local map in question.

Defining a local map from $M(0, x)$ to $M(0, x; y, y) + M(0, x - y)$ is even simpler: we send the left-hand 0-cell in B to the bottom-left 0-cell in A , and the right-hand 0-cell in B to the top-right 0-cell in A . We then send the single 1-cell in B to the sum of the three red 1-cells in A . Applying Lemma 3.2.3, we obtain the desired local equivalence.

We now proceed to the main theorem of the section. In order to formulate this, we will need to set up an extra bit of notation. First, we define a slight modification of the notion of a monotone root:

Definition 5.2.2. Fix a positive integer n , and let $h_1, \dots, h_n$ and $r_1, \dots, r_n$ be sequences of even integers such that:

1. $h_1 \geq h_2 \geq \dots \geq h_n$,
2. $r_1 < r_2 < \dots < r_n$, and
3. $h_n \geq r_n$.

We construct a graded root associated to this data following the same prescription as in Definition 4.4.1. We refer to such a root as a *weakly monotone root* and likewise denote it by $M(h_1, r_1; \dots; h_n, r_n)$. Geometrically, a weakly monotone root is exactly the same as a monotone root, except that the gradings of the leaves are nonincreasing (rather than strictly decreasing) as we move towards the stem.

Now let X and Y be two weakly monotone roots. We say that two roots X' and Y' are obtained

from X and Y via a *swap operation* if they are related as in Figure 5.6. More precisely, suppose that X is a type- m root, and fix any $1 \leq a \leq m$. Consider the set of leaves in X whose indices are strictly greater than a in absolute value. Let S_X be the subgraph of X formed by taking the union over the paths connecting these leaves to the vertex supporting the angle α_a in X . This subgraph is marked in blue in Figure 5.6. Letting Y be a type- n root and $1 \leq b \leq n$, we define S_Y similarly. We then form X' and Y' by cutting out S_X and S_Y from X and Y , respectively, and interchanging them. We say that X' and Y' are a *valid swapped pair* if X' and Y' are still weakly monotone.




Figure 5.6: Swap operation on two monotone roots X and Y . In this example, $m = 3$, $n = 2$, and $a = b = 1$.

To see what this does to the parameters of X and Y , let

$$X = M(p_1, s_1; \dots; p_m, s_m) \text{ and}$$

$$Y = M(q_1, t_1; \dots; q_n, t_n).$$

Fix any pair of integers $1 \leq a \leq m$ and $1 \leq b \leq n$, and let $\Delta = s_a - t_b$. Then the swapped roots X' and Y' are defined by the parameters

$$X' = M(p_1, s_1; \dots; p_a, s_a; q_{b+1} + \Delta, t_{b+1} + \Delta; \dots; q_n + \Delta, t_n + \Delta) \text{ and}$$

$$Y' = M(q_1, t_1; \dots; q_b, t_b; p_{a+1} - \Delta, s_{a+1} - \Delta; \dots; p_m - \Delta, s_m - \Delta).$$

Geometrically, this reflects the fact that the parameters corresponding to S'_Y are the same as the parameters corresponding to S_X , except uniformly shifted in grading by Δ . (Similarly for S'_X and S_Y .) The condition that X' and Y' are a valid swapped pair is equivalent to the pair of inequalities

$$\begin{aligned} p_a &\geq q_{b+1} + \Delta \text{ and} \\ q_b &\geq p_{a+1} - \Delta, \end{aligned} \tag{5.1}$$

since the other monotone inequalities are automatically satisfied. Note that we allow the swapping operation even if one of the subgraphs is empty, which occurs if $a = m$ or $b = n$. In this case, the monotonicity condition for one of the swapped roots is then trivial. The reader should familiarize themselves with the relations between the parameters above, since these will be of significant

importance shortly. Our main claim is now:

Theorem 5.2.3. [DS17, Theorem 3.5] Let X and Y be two weakly monotone roots, and suppose X' and Y' are a valid swapped pair obtained from X and Y . Then $X \otimes Y$ is locally equivalent to $X' \otimes Y'$.

Proof. Let X and Y be two weakly monotone roots as above, and assume that X' and Y' are a valid swapped pair. In order to communicate the idea of the proof, we will assume that $m = n$ and $a = b = 1$. For the purposes of exposition, we will also assume that X and Y both have an even number of leaves, although this is easily seen to be unimportant.

For ease of notation, let the leaf and angle generators of X be denoted by v_i and α_i , respectively, and let the leaf and angle generators of Y be denoted by w_i and β_i . Likewise, let the generators of X' be v'_i and α'_i , and let the generators of Y' be w'_i and β'_i . In case we need to refer to a generator of X or X' without specifying whether it is a leaf or an angle generator, we will use the notation x_i or x'_i (and similarly y_i or y'_i for Y or Y'). Note that there is an obvious gr-preserving identification between the unprimed and primed generators of index $|i| = 1$. For generators of index $|i| > 1$, we think of x'_i as corresponding to y_i , except with shifted grading

$$\text{gr}(x'_i) = \text{gr}(y_i) + \Delta.$$

Likewise, y'_i may be identified with x_i , except that²

$$\text{gr}(y'_i) = \text{gr}(x_i) - \Delta.$$

As in Example 5.2.1, we proceed by analyzing the skeleta of $X \otimes Y$ and $X' \otimes Y'$. Because $m = n$, these two skeleta are identical to each other, and are in fact just $2m \times 2m$ squares. Hence the subtlety will come from understanding the grading function gr defined on each. We have displayed the skeleta of $X \otimes Y$ and $X' \otimes Y'$ in Figure 5.7, together with some notational colorings, which we now explain.

Consider the skeleton of $X \otimes Y$, displayed on the left in Figure 5.7. We divide this into several parts, as follows. First, instead of explicitly drawing the 0-, 1-, and 2-cells appearing in the center of the complex, we have drawn a $2(m-1) \times 2(m-1)$ square box enveloping them. We refer to the collection of cells appearing inside this square as the *inner box*. These are the cells coming from products of leaf generators and angle generators internal to S_X and S_Y ; that is, cells of the form

$$x_i \times y_j, \text{ with } |i|, |j| > 1.$$

²Here, we use the term “identification” rather loosely, as we have not (yet) made any claim about constructing maps between the unprimed and primed complexes. As the leaf and angle generators of various complexes are just 0- and 1-cells, we obviously have a tautological “identification” between any two generators of the same dimension, although this may not preserve gr . However, given the geometric nature of the swapping operation, it is clear that certain generators in the unprimed and primed complexes should be naturally thought of as related to each other.




Figure 5.7: Skeleta of $X \otimes Y$ and $X' \otimes Y'$. On the left, the inner box represents cells of the form $x_i \times y_j$ with $|i|, |j| > 1$; similarly for cells $x'_i \times y'_j$ on the right.

Also appearing in the picture are the four corners of the overall square, together with the eight 1-cells and four 2-cells incident to them. These are cells of the form

$$x_i \times y_j, \text{ with } |i| = |j| = 1.$$

We refer to the collection of these cells as the *corner cells* and have colored/shaded them black in Figure 5.7. Next, there are the 0- and 1-cells which occur along the faces of the overall square, but are not corner cells. These are of the form

$$v_i \times y_j, \text{ with } |i| = 1 \text{ and } |j| > 1, \text{ or}$$

$$x_i \times w_j, \text{ with } |i| > 1 \text{ and } |j| = 1.$$

We call these cells the *side panels* and color them purple. Finally, there are also the 1- and 2-cells running between the inner box and the side panels. These are either of the form

$$\alpha_i \times y_j, \text{ with } |i| = 1 \text{ and } |j| > 1, \text{ or}$$

$$x_i \times \beta_j, \text{ with } |i| > 1 \text{ and } |j| = 1.$$

We refer to these as the *bridge cells*. In Figure 5.7, we have colored/shaded the bridge cells of the first kind blue and the second kind red. We define a similar decomposition of $X' \otimes Y'$, displayed on the right in Figure 5.7, except that we switch the coloring of the red and blue cells. (The reason for this change will become clear presently.)

Now let us compare the tensor products $X \otimes Y$ and $X' \otimes Y'$. First, observe that because of the natural identification of unprimed and primed generators when $|i| = 1$, the obvious bijection between the corner cells of $X \otimes Y$ and the corner cells of $X' \otimes Y'$ preserves the grading function gr . Explicitly, one can check that in both cases, the following gradings hold: all four 0-cells have $\text{gr} = p_1 + q_1$, the four horizontal 1-cells have $\text{gr} = s_1 + q_1$, the four vertical 1-cells have $\text{gr} = p_1 + t_1$, and the four 2-cells have $\text{gr} = s_1 + t_1$.

Next, consider the inner boxes of the two complexes. For $|i|, |j| > 1$, we have

$$\text{gr}(x'_i \times y'_j) = (\text{gr}(y_i) + \Delta) + (\text{gr}(x_j) - \Delta) = \text{gr}(y_i \times x_j).$$

Hence the correspondence between the two inner boxes defined by sending $x'_i \times y'_j$ to $y_i \times x_j$ for $|i|, |j| > 1$ preserves gr . We may think of this geometrically as identifying the inner boxes via reflection across the diagonal. Note that since each complex is symmetric about the vertical and horizontal axes, we can also effect a gr -preserving correspondence by using a ninety-degree rotation in either the clockwise or counter-clockwise direction.

Finally, consider the bridge cells. For $X \otimes Y$, we have already described the possible forms that these can take. In $X' \otimes Y'$, the bridge cells are of the form

$$\alpha'_i \times y'_j, \text{ with } |i| = 1 \text{ and } |j| > 1, \text{ or}$$

$$x'_i \times \beta'_j, \text{ with } |i| > 1 \text{ and } |j| = 1.$$

Gradings of cells of the first kind are given by

$$s_1 + \text{gr}(y'_j) = s_1 + \text{gr}(x_j) - \Delta = \text{gr}(x_j) + t_1$$

for $|j| > 1$. These are precisely the same as the gradings of bridge cells of the second kind in $X \otimes Y$.

Similarly, gradings of cells of the second kind in $X' \otimes Y'$ are given by

$$\text{gr}(x'_i) + t_1 = \text{gr}(y_i) + \Delta + t_1 = s_1 + \text{gr}(y_i),$$

for $|i| > 1$, which are the same as the gradings of bridge cells of the first kind in $X \otimes Y$. Hence we see that, as in the inner-box case, we again have a gr -preserving correspondence between the two sets of bridge cells, given by a ninety-degree rotation in either direction. To represent this identification, we have interchanged the red and blue colorings in $X' \otimes Y'$, so that the obvious bijection between the red (respectively, blue) cells in $X \otimes Y$ and the red (respectively, blue) cells in $X' \otimes Y'$ given by either rotation preserves gr .

These observations give some intuition for how to define a map between the skeleton of $X \otimes Y$ and the skeleton of $X' \otimes Y'$. Roughly speaking, we map the corner cells to each other via the obvious identity bijection, while using a ninety-degree rotation on the inner box and bridge cells. There are several evident difficulties with this suggestion. First, it is unclear how to extend this to an actual cellular map from one skeleton to the other while still preserving the cellular differential. Second, we have not specified how to define our map on the side panels. Indeed, one can check that the gradings of the cells in the side panels do *not* work out to produce any such similar correspondence, ninety-degree rotation or otherwise. We thus proceed instead by first constructing two auxiliary complexes which are locally equivalent to $X \otimes Y$ and $X' \otimes Y'$. As we will see, these auxiliary complexes will turn out to be more amenable to the intuition we have built up over the last several paragraphs.

We begin by defining a geometric complex A which is locally equivalent to $X \otimes Y$. The skeleton of A is displayed on the right in Figure 5.8. Like $X \otimes Y$, the skeleton of A has a $2(m-1) \times 2(m-1)$ inner box, which we define to be exactly same as that of $X \otimes Y$. That is, we let the cells in the inner box of A be duplicates of those in $X \otimes Y$, and we set gr on the inner box of A to be the same as it is on $X \otimes Y$. Next, the skeleton of A contains a collection of red and blue 1- and 2-cells. These are in obvious bijection with the red and blue bridge cells in $X \otimes Y$, as indicated by the correspondence in Figure 5.9. We define gr on these cells as being equal to gr of their counterparts in $X \otimes Y$. Finally, A contains a total of twelve black 0-, 1-, and 2-cells. We define gr on these cells as follows: the four 0-cells have $\text{gr} = p_1 + q_1$, the two horizontal 1-cells have $\text{gr} = s_1 + q_1$, the two vertical 1-cells have $\text{gr} = p_1 + t_1$, and the four 2-cells have $\text{gr} = s_1 + t_1$. These gradings should be compared with the gradings of the corner cells in $X \otimes Y$. Note that A has an obvious involution given by reflection in each coordinate.




Figure 5.8: Auxiliary complex A (right), locally equivalent to $X \otimes Y$ (left).

This defines a geometric complex A which we claim is locally equivalent to $X \otimes Y$. We begin by constructing a local map from $X \otimes Y$ to A by defining a map f between their skeleta. We let f be given by the identity on the inner box, and define f on the red and blue cells via the bijection indicated in Figure 5.9. Next, we send the four corners of $X \otimes Y$ to the four corners of A , and the four black 2-cells of $X \otimes Y$ to the four black 2-cells of A via the correspondence shown in Figure 5.9. It remains to define f on the black 1-cells and the side panels.

Consider (for example) the uppermost face of $X \otimes Y$. We send all of the purple 0-cells contained in this face to the nearest counter-clockwise corner 0-cell, which in this case is the top-left corner of A . We define f to be zero on all of the 1-cells contained in this face, except for the right-hand black 1-cell, which we map to the uppermost face in A . One should think of f restricted to the top face of $X \otimes Y$ as a map from the interval to itself, fixing the endpoints, where everything to the left




Figure 5.9: Schematic depiction of a map from $X \otimes Y$ to A . Solid arrows denote maps of 1-cells and dashed arrows denote maps of 2-cells. We have only indicated the map on a single red 1-cell and 2-cell; the map on the other red cells is defined similarly. Extension to the rest of $X \otimes Y$ (not displayed in the figure) is achieved by repeatedly rotating the picture by ninety degrees.

of the right-hand black 1-cell is crushed to the left endpoint and the remaining 1-cell is “stretched out” over the full interval. We define f similarly in a “counter-clockwise” manner on the other three sides of $X \otimes Y$. It is straightforward to check that this map preserves the cellular differential.

Note that f actually maps each cell of $X \otimes Y$ either to a single cell of A or to zero. Because of the way gr is defined on A , in the former case f actually preserves gr everywhere except for possibly on the purple 0-cells. Now, the purple 0-cells contained in the top face of $X \otimes Y$ have gradings

$$\text{gr} = p_i + q_1, \text{ with } 2 \leq i \leq m.$$

Due to the (weak) monotonicity of X , these are all less than or equal to $p_1 + q_1$, which is the grading of the top-left corner of A . A similar argument for the other purple 0-cells using the monotonicity of Y shows that f satisfies the lifting condition of Lemma 3.2.3. Since f is evidently J_0 -equivariant and maps each 0-cell of $X \otimes Y$ to exactly one 0-cell of A , applying Lemma 3.2.3 yields the desired local map.

We now construct a local map from A to $X \otimes Y$ by defining a map g from the skeleton of A to the skeleton of $X \otimes Y$. We set g to be the inverse of f on the inner box, as well as all black, red, and blue 2-cells. We define g on the corner 0-cells via the obvious identity correspondence, and define g on the black 1-cells by sending each face of A to the sum of 1-cells constituting the corresponding face in $X \otimes Y$. Finally, we define g on the red and blue 1-cells of A according to the prescription of Figure 5.10. More precisely, we send a red or blue 1-cell in A to the corresponding red or blue 1-cell in $X \otimes Y$, summed together with the chain of purple and black 1-cells running from that 1-cell to the nearest counter-clockwise corner 0-cell. With some care, one can check that g preserves the cellular differential (over $\mathbb{F}$).

We now check that g satisfies the lifting condition of Lemma 3.2.3. Note that, except in the case




Figure 5.10: Slanted red 1-cell in A mapped to the sum of 1-cells in $X \otimes Y$ displayed on the left.

of the black 1-cells and bridge 1-cells of A , g maps each cell of A to a single cell in $X \otimes Y$ of the same grading. We thus consider the black 1-cells first. As described, g takes the uppermost face of A to the sum of 1-cells constituting the top face of $X \otimes Y$. The gradings of these latter 1-cells are given by

$$\text{gr} = s_i + q_1, \text{ with } 1 \leq i \leq m.$$

Due to the monotonicity of X , these are all greater than or equal to $s_1 + q_1$, which is the grading of the top face of A . A similar computation using the monotonicity of Y holds for the other three sides, verifying the lifting condition for the black 1-cells.

Now consider a red 1-cell in A , as in Figure 5.10. This has grading $\text{gr} = p_i + t_1$ for some fixed $|i| > 1$. Note that this is less than or equal to $p_2 + t_1$, by the monotonicity of X . Now, the red 1-cell in the image of g has exactly the same grading as the original, while the remaining 1-cells in the image of g all lie in the uppermost face of $X \otimes Y$, and have gradings bounded below by $s_1 + q_1$ (by the previous paragraph). Thus, in this instance, to verify the monotonicity condition it suffices to establish the inequality

$$p_2 + t_1 \leq s_1 + q_1.$$

However, rearranging this yields the inequality $q_1 \geq p_2 - \Delta$, which is precisely the condition that Y' is weakly monotone, as in (5.1). A similar argument for the blue cells using the fact that X' is weakly monotone proves that g satisfies the lifting condition for all cells. Applying Lemma 3.2.3 then yields the desired local equivalence.

We have thus constructed a somewhat simpler geometric complex A which is locally equivalent to $X \otimes Y$. Note that (roughly speaking) A consists entirely of the corner cells, inner box, and bridge cells of $X \otimes Y$, with the side panels having “disappeared”. We now construct a similar geometric complex B which is locally equivalent to $X' \otimes Y'$. This is displayed on the right in Figure 5.11. Note that B is defined in exactly the same way as A , except with one important modification. Whereas before we mapped the purple 0-cells in $X \otimes Y$ to the nearest *counter-clockwise* corner 0-cell, we now map the purple 0-cells in $X' \otimes Y'$ to the nearest *clockwise* corner 0-cell. (We likewise reverse the chirality of the various other maps needed to define the local equivalence.) This results in the

complex B shown in Figure 5.11. Again, the fact that B is locally equivalent to $X' \otimes Y'$ follows from the fact that X , Y , X' , and Y' are all weakly monotone.




Figure 5.11: Geometric complexes A and B , locally equivalent to $X \otimes Y$ and $X' \otimes Y'$. Compare with Figure 5.8.

We now observe, however, that there is actually a gr-preserving isomorphism between A and B . This is most easily described geometrically, as follows. We view the black 0- and 1-cells as constituting a fixed outer frame, and we think of the black 2-cells together with the red and blue cells as being malleable strings/rubber sheets connecting the frame to the inner box. To map A onto B , we rotate the inner box by ninety degrees counter-clockwise, while keeping the outer frame fixed. This “drags” the remaining cells of A into bijection with the cells of B . Moreover, because of our discussion at the beginning of the proof, this correspondence preserves gr. Similarly, to map B onto A , we rotate the inner box clockwise by ninety degrees. This shows that A and B are actually isomorphic complexes, completing the proof.

We now sketch the generalization of the proof to the case when m, n, a , and b are allowed to vary. For the moment, we will still assume that $m - a = n - b$. As before, we again divide $X \otimes Y$ into four types of cells. These have been color-coded and displayed on the left in Figure 5.12. Some cell types are exactly analogous to their counterparts in the previous case; for example, the inner box consists of cells of the form

$$x_i \times y_j, \text{ with } |i| > a \text{ and } |j| > b.$$

However, the corner cells are now replaced with a larger collection of black cells of the form

$$x_i \times y_j, \text{ with } |i| \leq a \text{ and } |j| \leq b,$$

In Figure 5.12, we have avoided drawing most of the corner cells, and have instead covered the majority of them with four black corner wedges. Similarly, the side panels generalize to a larger collection of purple cells. As before, it is easily checked that the corner cells of $X \otimes Y$ and $X' \otimes Y'$ are identical, while the inner boxes and bridge cells are related by a ninety-degree rotation.




Figure 5.12: The case where $m - a = n - b$. Auxiliary complex A (right), locally equivalent to $X \otimes Y$ (left). We have avoided shading the 2-cells in order to prevent clutter in the diagram.

We will not re-formulate the entire argument here in the more general case. The idea is, of course, to construct an auxiliary complex A which (roughly speaking) retains the corner cells, inner box, and bridge cells of $X \otimes Y$, while throwing out the side panels. Such a complex is displayed on the right in Figure 5.12. The map from $X \otimes Y$ to A should again be thought of as taking (for example) the uppermost group of side panels and crushing them directly left towards the top-left corner wedge. In contrast to the previous case, there are now some purple 1-cells (the vertical ones) on which our map is non-zero, as well as some new black and purple 2-cells that must be considered. However, the desired maps on these cells are defined using the same intuition as before. We leave it to the reader to complete the generalization of the proof.

Finally, we consider the case when $m - a \neq n - b$. This initially seems more difficult, since if $m - a \neq n - b$, then the obvious inner box is no longer square, and hence cannot exhibit the desired rotational symmetry. However, if (without loss of generality) $m - a < n - b$, then we may pad the complex of X with redundant parameters, in order to create a homotopy equivalent complex with a larger number of generators. Roughly speaking, we take the standard complex associated to X , which has the parameters

$$(p_1, s_1; \dots; p_m, s_m)$$

and replace it with a similar complex that has the parameters

$$(p_1, s_1; \dots; p_m, s_m; s_m, s_m; \dots; s_m, s_m).$$

This is obtained by inserting redundant 0- and 1-cells (all of grading s_m) into the center of the skeleton of X . Replacing our original complex with this “fattened” one does not alter the homotopy class or change the parameters a and b of the swapping operation. Hence we can always reduce to the case when $m - a = n - b$. This completes the proof. $\square$

5.3 Proof of the main theorem

We now show that every monotone root is locally equivalent to a linear combination of type-one roots (up to overall grading shift). Combined with the fact that any good family is linearly independent, this will imply Theorem 1.3.3.

Theorem 5.3.1. [DS17, Theorem 4.2] Let $M = M(h_1, r_1; \dots; h_n, r_n)$ be a monotone root. Then we have the local equivalence

$$M = \left(\sum_{i=1}^n M(h_i, r_i) \right) - \left(\sum_{i=1}^{n-1} M(h_{i+1}, r_i) \right).$$

Proof. We proceed by induction on n . If $n = 1$, then M is already a type-one root, and there is nothing to prove. Otherwise, we have the local equivalence

$$M = M + M(h_2, r_1) - M(h_2, r_1)$$

trivially obtained by taking the tensor product of M with $M(h_2, r_1)$ and its dual. We now apply Theorem 5.2.3 to $X = M$ and $Y = M(h_2, r_1)$, with the swapping parameters $a = b = 1$. This yields the valid swapped pair

$$X' = M(h_1, r_1)$$

and

$$Y' = M(h_2, r_1; h_2, r_2; \dots; h_n, r_n),$$

as displayed in Figure 5.13. Note that in this case the subgraph S_Y is empty. We thus have the chain of local equivalences

$$\begin{aligned} M &= M + M(h_2, r_1) - M(h_2, r_1) \\ &= M(h_1, r_1) + M(h_2, r_1; h_2, r_2; \dots; h_n, r_n) - M(h_2, r_1) \\ &= M(h_1, r_1) + M(h_2, r_2; \dots; h_n, r_n) - M(h_2, r_1), \end{aligned}$$

where in the last line we have applied Algorithm 4.4.2 to obtain a monotone root of type $n - 1$. This establishes the inductive step and completes the proof. See Figure 5.13. $\square$

Note that this immediately implies that the ι -complex of any connected sum Y of Seifert fibered spaces is locally equivalent to a linear combination of type-one roots. More precisely, let $\mathcal{M} = \{M_i\}_{i \geq 1}$ be any good family of roots. Since every (non-trivial) type-one root is equal to an element of $\mathcal{M}$ up to grading shift, it follows that $h(Y)$ is equal to a linear combination of type-one roots drawn from $\mathcal{M}$, together with an overall grading shift Δ :

$$h(Y) = \left(\sum_i c_i M_i \right) [\Delta]. \tag{5.2}$$

We thus obtain a proof of Theorem 1.3.3:




Figure 5.13: Proof of Theorem 5.3.1.

Proof of Theorem 1.3.3. We can effect any even grading shift by tensoring with an appropriate multiple of $h(\Sigma(2, 3, 5))$. It follows from Theorem 5.3.1 that if Y is any individual Seifert fibered space, then $h(Y)$ lies in the span of $h(\Sigma(2, 3, 5))$ and the $h(\Sigma(2i + 1, 4i + 1, 4i + 3))$. Hence these certainly generate $h(\Theta_{\text{SF}})$. Linear independence follows from Theorem 5.1.6. $\square$

We now make a more careful study of the grading shift in the decomposition (5.2). Fix any good family of roots $\mathcal{M} = \{M_i\}$, and let Y be a linear combination of Seifert fibered spaces. We define $\Delta_{\mathcal{M}}(Y) \in 2\mathbb{Z}$ to be the grading shift associated to $h(Y)$ when using the family $\mathcal{M}$ in (5.2), so that

$$h(Y) = \left(\sum_i c_i M_i \right) [\Delta_{\mathcal{M}}(Y)].$$

Note that due to the uniqueness of the above decomposition (Theorem 5.1.6), $\Delta_{\mathcal{M}}(Y)$ is well-defined and descends to a homomorphism from the subgroup of $\mathfrak{J}$ spanned by (standard complexes of) graded roots to $2\mathbb{Z}$. More precisely, the grading shift Δ of (5.2) clearly changes sign under

orientation reversal and is additive under connected sum. Moreover, $\Delta_{\mathcal{M}}(Y)$ implicitly depends only on the local equivalence class $h(Y)$.

Now let us consider the specific families $\mathcal{X}$ and $\mathcal{Y}$. If M is a single monotone root, then we may rearrange the decomposition of Theorem 5.3.1 to read

$$\begin{aligned} M &= \left(\sum_{i=1}^n M(0, r_i - h_i)[-h_i] \right) - \left(\sum_{i=1}^{n-1} M(0, r_i - h_{i+1})[-h_{i+1}] \right) \\ &= \left(\sum_{i=1}^n M(0, r_i - h_i) - \sum_{i=1}^{n-1} M(0, r_i - h_{i+1}) \right) [-h_1]. \end{aligned}$$

Hence in the case of a single monotone root, the overall grading shift for the family $\mathcal{X}$ is given by $-h_1$. Meanwhile,

$$\begin{aligned} M &= \left(\sum_{i=1}^n M(h_i - r_i, 0)[-r_i] \right) - \left(\sum_{i=1}^{n-1} M(h_{i+1} - r_i, 0)[-r_i] \right) \\ &= \left(\sum_{i=1}^n M(h_i - r_i, 0) - \sum_{i=1}^{n-1} M(h_{i+1} - r_i, 0) \right) [-r_n]. \end{aligned}$$

Thus, the overall grading shift if we use $\mathcal{Y}$ is given by $-r_n$. Now suppose that Y is a single Seifert fibered space, and let M be the monotone root corresponding to $h(Y)$. Recall from Remark 4.4.7 that $h_1 = d(Y)$ and $r_n = -2\bar{\mu}(Y)$. It follows that if Y is a single Seifert fibered space, then

$$\Delta_{\mathcal{X}}(Y) = -d(Y)$$

and

$$\Delta_{\mathcal{Y}}(Y) = 2\bar{\mu}(Y).$$

Consider the former. Since $\Delta_{\mathcal{X}}$ and $-d$ are both homomorphisms, the fact that they agree on each individual Seifert fibered space implies that they are equal on any linear combination of Seifert fibered spaces. A similar argument for $\Delta_{\mathcal{Y}}$ yields a proof of Theorem 1.3.4:

Proof of Theorem 1.3.4. It is a standard fact that the Neumann-Siebenmann invariant changes sign under orientation reversal and is additive under connected sum; that is,

$$\bar{\mu}(-Y) = -\bar{\mu}(Y)$$

and

$$\bar{\mu}(Y_1 \# Y_2) = \bar{\mu}(Y_1) + \bar{\mu}(Y_2).$$

Even though we do not *a priori* know that $\bar{\mu}$ is a homology cobordism invariant, we do know that $\Delta_{\mathcal{Y}}$ is similarly additive under connected sum and changes sign under orientation reversal. Hence the fact that $\Delta_{\mathcal{Y}}(Y)$ coincides with $2\bar{\mu}(Y)$ for each individual Seifert fibered space establishes the equality in general for connected sums of Seifert fibered spaces. The fact that $\Delta_{\mathcal{Y}}$ depends only on the local equivalence class $h(Y)$ then implies the same for $\bar{\mu}(Y)$. $\square$

Proof of Theorem 1.3.5. The Neumann-Siebenmann reduces to the Rokhlin invariant mod 2. $\square$

5.4 The basis decomposition algorithm

We now take a brief step back to collect together the results needed to explicitly compute the basis decomposition of Theorem 1.3.3:

Algorithm 5.4.1 (Basis decomposition algorithm). Let $Y = \Sigma_1 \# \Sigma_2 \# \cdots \# \Sigma_m$ be a linear combination of Seifert fibered spaces. Here, the Σ_j are not required to be distinct, and some of the Σ_j may have reversed orientation. We compute $h(Y)$ as follows:

1. For each Σ_j , compute the symmetric graded root for Σ_j (or $-\Sigma_j$, in the case of reversed orientation).
2. Extract the monotone subroot according to Algorithm 4.4.2.
3. Write down the (grading-shifted) basis decomposition afforded by Theorem 5.3.1. In terms of the basis $\{Y_k\}_{k \in \mathbb{N}}$, this can be obtained by following the schematic of Figure 5.14:
 - (a) The indices of the positive terms are given by the parameters $(h_i - r_i)/2$. These are represented by the lengths of the blue line segments (divided by two). Note that if $h_n = r_n$, then the final parameter is zero and does not contribute to the basis decomposition.
 - (b) The indices of the negative terms are given by the parameters $(h_{i+1} - r_i)/2$. These are represented by the lengths of the red line segments (divided by two).
 - (c) As per the discussion at the end of Section 5.3, we apply an overall grading shift equal to $-r_n$.
4. Sum together the decompositions (and grading shifts) for the Σ_j , multiplying a decomposition by -1 if Σ_j has reversed orientation. Note that in this case, the grading shift is also multiplied by -1 .

This yields the desired basis decomposition

$$h(Y) = \left(\sum_i c_i Y_i \right) [2\bar{\mu}(Y)].$$

Note that the overall grading shift is equal to $2\bar{\mu}(Y)$. If desired, we can remove the grading shift by taking connected sums with $\Sigma(2, 3, 5)$.

Remark 5.4.2. The reader may (rightly) object that Algorithm 5.4.1 can only be applied if the decomposition of Y into individual Seifert fibered spaces is given explicitly. In Chapter 6, we will show that $h(Y)$ can (in some sense) be computed even when this is not the case, although this




Figure 5.14: Blue line segments have lengths $h_i - r_i$. Red (dashed) line segments have lengths $h_{i+1} - r_i$.

requires knowledge of the connected Floer homology.

5.5 Examples and applications

For concreteness, we now give a few explicit examples:

Example 5.5.1. Let $Y = \Sigma(5, 8, 13)$. We have drawn part of the graded root associated to $h(Y)$ on the left in Figure 5.15; the monotone subroot is displayed on the right. (The full graded root corresponding to $h(Y)$ has several other length-one branches occurring in lower gradings which we have suppressed for brevity.) Applying Algorithm 5.4.1, we see that

$$h(\Sigma(5, 8, 13)) = M(4, 0; 2, 2) = (Y_2 - Y_1)[-2]$$




Figure 5.15: Graded root and monotone subroot for $h(\Sigma(5, 8, 13))$. See also [Sto15].

Example 5.5.2. Let $Y = \Sigma(13, 21, 34)$. We have drawn part of the graded root associated to $h(Y)$ on the left in Figure 5.16; the monotone subroot is displayed on the right. (The full graded

root is significantly more complicated, but we have displayed only the portion which is relevant for computing the local equivalence class.) Applying Algorithm 5.4.1, we see that

$$h(\Sigma(13, 21, 34)) = M(12, 0; 10, 2) = (Y_6 + Y_4 - Y_5)[-2]$$




Figure 5.16: Graded root and monotone subroot for $h(\Sigma(13, 21, 34))$.

As evidenced by the above examples, most “small” Brieskorn spheres have a fairly simple decomposition into the Y_i . Indeed, several familiar classes of Brieskorn spheres are actually of type-one type, as discussed in [Sto15]. However, it appears that arbitrarily high complexity can be achieved by considering Brieskorn spheres with sufficiently large parameters; for example, using parameters on the order of $100 \sim 500$, one can construct examples that decompose into a linear combination of upwards of twenty basis elements.³

As a rough statistical observation, the author has found that if $\mathcal{F}$ is a family of Brieskorn spheres in which only one parameter changes,⁴ then the elements of $\mathcal{F}$ generally appear to fall into at most one or two local equivalence classes. (See for example work of Tweedy [Twe13].) This occurs even though the Heegaard Floer homology becomes successively more and more complicated. Indeed, as evidenced by Example 5.5.2, passing from the Heegaard Floer homology to its monotone subroot generally results in a vast reduction in complexity. Unfortunately, this means that Theorem 1.3.3 is not very effective at distinguishing “small” families of Seifert fibered spaces.

Note that if Y is a single Seifert fibered space, then $h(Y)$ takes a particular form with respect to the basis decomposition of Theorem 1.3.3:

³These, as well as Example 5.5.2 above, are originally due to Duncan McCoy.

⁴For example, $\mathcal{F} = \{\Sigma(p, q, pqn + k)\}_{n \in \mathbb{N}}$.

Theorem 5.5.3. [DS17, Theorem 1.2] Let Y be a Seifert fibered homology sphere (of either orientation), and write the terms in the basis decomposition in order of decreasing index. Then each non-zero coefficient c_i is either $+1$ or -1 , and the sum is alternating. If Y has the usual orientation, then the first coefficient is $+1$; if Y has reversed orientation, then the first coefficient is -1 .

Theorem 5.5.3 is evident from Theorem 5.3.1. It is thus easily checked that a linear combination of non-trivial monotone roots is generally *not* locally equivalent to a single monotone root; this easily produces many examples of linear combinations of Seifert fibered spaces which are not homology cobordant to any single Seifert fibered space.

We close this section by summarizing some further results from [DS17] regarding the involutive Floer correction terms $\bar{d}$ and $\underline{d}$. In Section 4.3, we computed $\bar{d}$ and $\underline{d}$ for connected sums of type-one roots of the same sign. In light of Theorem 1.3.3, it is clear that to compute $\bar{d}$ and $\underline{d}$ for elements of Θ_{SF} in general, it suffices to understand linear combinations of type-one roots. Thus, consider a linear combination in $\mathfrak{J}$ of the form

$$(Y_{s_1} + Y_{s_2} + \dots + Y_{s_m}) - (Y_{t_1} + Y_{t_2} + \dots + Y_{t_n}),$$

where the Y_i are the basis elements of Theorem 1.3.3. Without loss of generality, we assume that the s_i are distinct from the t_i , and that $s_1 \geq s_2 \geq \dots \geq s_m$ and $t_1 \geq t_2 \geq \dots \geq t_n$. Then we have:

Theorem 5.5.4. [DS17, Theorem 1.5] Let Y be a linear combination of Seifert fibered homology spheres. Suppose that $h(Y)$ decomposes as a linear combination

$$(Y_{s_1} + Y_{s_2} + \dots + Y_{s_m}) - (Y_{t_1} + Y_{t_2} + \dots + Y_{t_n})$$

as above, shifted by some grading Δ . Define the quantities

$$P_i = 2 \left(\sum_{j=1}^i t_j - \sum_{j=1}^i s_j \right) \text{ for } 0 \leq i \leq \min(m, n), \text{ and}$$

$$Q_i = 2 \left(\sum_{j=1}^i t_j - \sum_{j=1}^{i+1} s_j \right) \text{ for } 0 \leq i \leq \min(m-1, n).$$

Note that $P_0 = 0$ and $Q_0 = -2s_1$. Then

$$\begin{aligned} \underline{d}(Y) = & d(Y) + \max\{ \min(P_0, Q_0), \\ & \min(P_0, P_1, Q_1), \\ & \vdots \\ & \min(P_0, P_1, P_2, \dots, P_{\min(m, n)}, Q_{\min(m, n)}) \}, \end{aligned}$$

with the understanding that if $\min(m, n) = m$, the $Q_{\min(m, n)}$ in the last line should be deleted.

Since orientation reversal corresponds to negation in $\mathfrak{J}$, changing Y to $-Y$ interchanges the index

sets $\{s_i\}$ and $\{t_i\}$. Due to the fact that $\bar{d}(Y) = -\underline{d}(-Y)$, Theorem 5.5.4 can also be used to compute $\bar{d}(Y)$.

We will not give the proof of Theorem 5.5.4 here. The original proof given in [DS17] proceeds by analyzing the spherical complex corresponding to the sum of the Y_{s_i} , tensored with the dual of the spherical complex corresponding to the sum of the Y_{t_i} . We record the statement of Theorem 5.5.4 mainly to demonstrate the fact that $\bar{d}$ and $\underline{d}$ can be determined explicitly from the basis decomposition of $h(Y)$. We also briefly describe (without proof) two applications of this computation. The first is an asymptotic characterization of the involutive Floer correction terms:

Corollary 5.5.5. [DS17, Corollary 1.6] Let Y be a linear combination of Seifert fibered homology spheres. Suppose that $h(Y)$ is given by a linear combination

$$(Y_{s_1} + Y_{s_2} + \dots + Y_{s_m}) - (Y_{t_1} + Y_{t_2} + \dots + Y_{t_n})$$

as in the statement of Theorem 5.5.4. Then the involutive Floer correction terms of $\#_k Y$ for k sufficiently large are as follows. If $t_1 > s_1$, then for k sufficiently large, we have

$$\bar{d}(\#_k Y) = k \cdot d(Y) + 2t_1 \text{ and}$$

$$\underline{d}(\#_k Y) = k \cdot d(Y).$$

If $s_1 > t_1$, then for k sufficiently large, we have

$$\bar{d}(\#_k Y) = k \cdot d(Y) \text{ and}$$

$$\underline{d}(\#_k Y) = k \cdot d(Y) - 2s_1.$$

Thus, even though the general form for $\bar{d}$ and $\underline{d}$ is quite complicated, the involutive Floer correction terms of $\#_k Y$ for k sufficiently large only depend on the original d -invariant of Y , together with the basis element of greatest index appearing in $h(Y)$. We currently do not know whether a similar “stabilization” result holds more generally for all homology spheres; if not, then Corollary 5.5.5 provides an interesting obstruction to being a connected sum of Seifert fibered spaces. Corollary 5.5.5 should be compared with [Sto17, Theorem 1.3], which shows that (in general) $\alpha(\#_k Y)/k$, $\beta(\#_k Y)/k$, and $\gamma(\#_k Y)/k$ all limit to $d(Y)$ as $k \rightarrow \infty$. See also [HHL18, Theorem 1.2], which establishes the analogous result for $\bar{d}$ and $\underline{d}$.

We also have the following realization result:

Corollary 5.5.6. [DS17, Corollary 1.7] Let $d, \bar{d}$, and $\underline{d}$ be any triple of even integers such that $\underline{d} \leq d \leq \bar{d}$ and $d, \bar{d}$, and $\underline{d}$ are not all equal to each other. In addition, let $\bar{\mu}$ be any integer. Then there are infinitely many distinct classes in Θ_{SF} with the invariants $d, \bar{d}, \underline{d}$, and $\bar{\mu}$.

Given Theorem 1.3.3, Corollary 5.5.6 is not particularly surprising, although it does establish that

$\bar{\mu}$ is almost entirely independent from $\bar{d}$ and $\underline{d}$. Note that if $\underline{d} = d = \bar{d}$, then it is not hard to show the corresponding local equivalence class must be trivial (up to grading shift); in our case, this also implies $d = -2\bar{\mu}$. For the analogous result on the $\text{Pin}(2)$ -side, see work of Lin [Lin18b].

Chapter 6

Connected Heegaard Floer Homology

In this chapter, we present the results of [Dai18a]. Our object of study here will be the connected Heegaard Floer homology of elements in Θ_{SF} . In the first section, we give a straightforward algorithm for computing the connected Floer homology from the basis decomposition of Theorem 1.3.3. This will depend on an explicit model for the connected complex of linear combinations of type-one roots. We thus spend several sections describing this construction using geometric complexes. As an application, we briefly indicate a possible approach towards showing that $\Theta_{\text{SF}} \neq \Theta_{\mathbb{Z}}^3$.

6.1 The connected homology algorithm

We begin with the connection between the basis decomposition of Theorem 1.3.3 and the connected Heegaard Floer package developed by Hendricks, Hom, and Lidman. Recall that the connected homology of an ι -complex is defined to be the usual $\mathbb{F}[U]$ -homology of its connected complex, which should be thought of as the simplest possible representative of its local equivalence class. Roughly speaking, the connected Heegaard Floer homology of Y records which U -torsion towers in $HF^-(Y)$ are invariant under homology cobordism. These are preseved by virtue of being “connected” to the U -nontorsion tower via the action of ι , although this is of course an imprecise notion.

If we restrict our attention to linear combinations of Seifert fibered spaces, then Theorem 1.3.3 already affords a complete understanding of Θ_{SF} from the point of view of involutive Heegaard Floer homology. For other families of 3-manifolds, however, we may not have an explicit understanding of

the set of local equivalence classes. In these situations, the connected Floer homology provides a convenient way of obstructing local equivalence between two ι -complexes. In particular, the connected homology allows us to compare elements of Θ_{SF} with other 3-manifolds whose ι -complexes may not *a priori* lie in $h(\Theta_{\text{SF}})$. We can also use the connected homology to study elements of Θ_{SF} whose decomposition into individual Seifert fibered spaces is not known explicitly; see Remark 5.4.2.

In this section, it will be slightly more convenient for us to work with the basis $\mathcal{X} = \{X_i\}_{i \geq 1}$ for $h(\Theta_{\text{SF}})$, rather than our usual basis $\mathcal{Y} = \{Y_i\}_{i \geq 1}$. This means that if $Y \in \Theta_{\text{SF}}$, we have

$$h(Y) = \left(\sum_i c_i X_i \right) [-d(Y)].$$

We begin by defining an algorithm that takes an element of $\mathbb{Z}^\infty$ (written as a linear combination of the X_i) and outputs a $\mathbb{Z}$ -graded $\mathbb{F}[U]$ -module.

Definition 6.1.1. Let $\mathcal{T}_r(l)$ be the U -torsion tower $\mathbb{F}[U]/U^l$, shifted in grading so that the element 1 has grading r . Note that $\mathcal{T}_r(l)$ has l elements; we refer to l as the *length* of $\mathcal{T}_r(l)$. We will additionally identify certain towers as *downwards pointing* and other towers as *upwards pointing*. In the first case, we define the *head* of $\mathcal{T}_r(l)$ to be the element of maximal grading r , and the *tail* of $\mathcal{T}_r(l)$ to be the element of minimal grading $r - 2(l - 1)$. In the second, we interchange these conditions, so that the head is the element of minimal grading and the tail is the element of maximal grading. Note that in both cases, the action of U still decreases grading by two; these definitions are a notational aid and are not intrinsic to $\mathcal{T}_r(l)$.

Algorithm 6.1.2 (Connected homology algorithm). Let X be a linear combination of the form

$$X = \pm X_{i_1} \pm X_{i_2} \pm X_{i_3} \pm \cdots \pm X_{i_n},$$

with $i_1 \geq i_2 \geq \cdots \geq i_n$. Without loss of generality, we may assume that X is *maximally simplified*, meaning that there are no cancelling pairs $X_i - X_i$. We use the indices appearing in X to construct a $\mathbb{Z}$ -graded $\mathbb{F}[U]$ -module, as follows:

1. We begin by drawing a U -torsion tower of length i_1 . If the sign of X_{i_1} is positive, then we identify this tower as downwards pointing, and place it so that its head has grading zero. If the sign of X_{i_1} is negative, then we identify the tower as upwards pointing, and place it so that its head has grading one.
2. We continue drawing U -torsion towers inductively, where at each step k we draw a tower of length i_k . If the sign of X_{i_k} is positive, then we identify this tower as downwards pointing; otherwise, we identify the tower as upwards pointing. We place the tower according to the following criteria:

- (a) If the sign of X_{i_k} is opposite the sign of $X_{i_{k-1}}$, then the head of the i_k -tower should have the same grading as the tail of the i_{k-1} -tower.
- (b) If the signs of X_{i_k} and $X_{i_{k-1}}$ are both positive, then the head of the i_k -tower should have grading one less than the tail of the i_{k-1} -tower.
- (c) If the signs of X_{i_k} and $X_{i_{k-1}}$ are both negative, then the head of the i_k -tower should have grading one greater than the tail of the i_{k-1} -tower.

We denote the $\mathbb{F}[U]$ -module constructed in this way by $\mathcal{H}(X)$. An example of this algorithm is presented in Figure 6.1.




Figure 6.1: Example of $\mathcal{H}(X)$. Here, arrows represent orientations of towers, where we think of an arrow as pointing from the head to the tail of a tower. Note that the action of U has degree -2 throughout; the arrows are simply a notational device to aid with the placement of the towers. Maslov gradings are indicated on the left of the diagram.

We are now ready to state our main theorem:

Theorem 6.1.3. [Dai18a, Theorem 1.1] Let Y be an integer homology sphere. Suppose that we have the local equivalence

$$h(Y) = (\pm X_{i_1} \pm X_{i_2} \pm X_{i_3} \pm \cdots \pm X_{i_n})[-d(Y)],$$

where we assume that the above linear combination is maximally simplified with $i_1 \geq i_2 \geq \cdots \geq i_n$.

Then the connected Heegaard Floer homology of Y is given by

$$HF_{\text{conn}}(Y) = \mathcal{H}(X)[-d(Y) + 1].$$

Here, $X = \pm X_{i_1} \pm \cdots \pm X_{i_n}$ is the linear combination obtained in the decomposition of $h(Y)$, and $\mathcal{H}(X)$ is the $\mathbb{F}[U]$ -module constructed using Algorithm 6.1.2.

We will not give the full proof of Theorem 6.1.3 in this paper. Instead, we will spend most of our time describing an explicit model for the connected complex of $h(Y)$ in the above situation (although we will not give a complete proof for this either). This will have a nice interpretation in terms of geometric complexes, which we hope will provide sufficient motivation in case the reader wishes to read [Dai18a].

Note that if we know beforehand that $h(Y) \in h(\Theta_{\text{SF}})$, then it is easy to recover the basis decomposition of Theorem 1.3.3 using the connected homology. More precisely:

Theorem 6.1.4. [Dai18a, Theorem 1.2] Suppose that $h(Y) \in h(\Theta_{\text{SF}})$. Then the the basis decomposition of $h(Y)$ can be determined from the isomorphism class of $HF_{\text{conn}}(Y)$, together with the knowledge of $d(Y)$.

Proof. We “invert” the procedure of Algorithm 6.1.2. Explicitly, let i_1 be the maximal length of any U -torsion tower in $HF_{\text{conn}}(Y)$. Since the form of $HF_{\text{conn}}(Y)$ is constrained to lie in the output of Algorithm 6.1.2, the collection of towers of length i_1 in $HF_{\text{conn}}(Y)$ necessarily forms a sequence

$$\bigoplus_{j=1}^{c_{i_1}} \mathcal{T}_{r-(2i_1-1)(j-1)}(i_1)$$

for some uppermost grading r and natural number c_{i_1} . The coefficient of X_{i_1} in $h(Y)$ is then evidently given by $\pm c_{i_1}$. The sign of c_{i_1} can be determined by the placement of the above sequence of towers relative to $d(Y)$. Proceeding inductively allows us to similarly compute the other coefficients of $h(Y)$. $\square$

In particular, this shows that for $Y \in \Theta_{\text{SF}}$, the connected homology contains precisely the same information as the full local equivalence class of Y . Moreover, if Y_1 and Y_2 are two manifolds as above whose connected Floer homologies and d -invariants are known, then Theorems 6.1.3 and 6.1.4 allow us to immediately compute the connected Floer homology of $Y_1 \# Y_2$. To do this, we first apply Theorem 6.1.4 to write $h(Y_1)$ and $h(Y_2)$ as a linear combination of the X_i ; then we sum these together, remove all cancelling pairs, and apply Theorem 6.1.3 to the result.

Finally, we also derive some identities involving the rank of $HF_{\text{conn}}(Y)$.

Corollary 6.1.5. [Dai18a, Corollary 1.3] Let Y be a linear combination of Seifert fibered spaces. Then

$$\bar{\mu}(Y) + \frac{d(Y)}{2} = \text{signed rank } HF_{\text{conn}}(Y),$$

where $\bar{\mu}(Y)$ is the Neumann-Siebenmann invariant of Y . The signed rank of HF_{conn} is computed by assigning each element in a downwards pointing tower a positive multiplicity, and each element

in an upwards pointing tower a negative multiplicity.

Proof. As discussed in Section 5.3, we have both

$$h(Y) = (\pm Y_{i_1} \pm Y_{i_2} \pm Y_{i_3} \pm \cdots \pm Y_{i_n})[2\bar{\mu}(Y)]$$

and

$$h(Y) = (\pm X_{i_1} \pm X_{i_2} \pm X_{i_3} \pm \cdots \pm X_{i_n})[-d(Y)].$$

Since $Y_i = X_i[-2i]$, substituting this into the first equality and comparing with the second yields

$$2\bar{\mu}(Y) - 2(\pm i_1 \pm i_2 \pm \cdots \pm i_n) = -d(Y),$$

where the signs in front of the i_k are the same as the signs appearing in the decomposition of $h(Y)$. Each index i_k contributes a tower of length i_k to $HF_{\text{conn}}(Y)$, which is downwards pointing if the sign in front of X_{i_k} is positive and upwards pointing if the sign in front of X_{i_k} is negative. The corollary easily follows. $\square$

Note that the orientations of the towers in $HF_{\text{conn}}(Y)$ are not intrinsically defined, but must instead be deduced from the overall structure of $HF_{\text{conn}}(Y)$. (See the proof of Theorem 6.1.4.) This can only be done when $HF_{\text{conn}}(Y)$ takes the form prescribed by Theorem 6.1.3, which is consistent with the fact that $\bar{\mu}$ is not defined for all 3-manifolds.

Corollary 6.1.6. [Dai18a, Corollary 1.4] Let Y be a linear combination of Seifert fibered spaces. Then

$$\mu(Y) + \frac{d(Y)}{2} = \text{rank } HF_{\text{conn}}(Y) \pmod{2},$$

where $\mu(Y)$ is the Rokhlin invariant of Y .

Proof. This follows immediately from reducing both sides of the previous corollary by two. $\square$

6.2 Preliminary examples

As a warm-up, we begin by recalling some of the examples investigated in previous sections:

Example 6.2.1. Consider an alternating sum of X_i ; i.e.,

$$X_{i_1} - X_{i_2} + X_{i_3} - \cdots \pm X_{i_n}$$

with $i_1 > i_2 > \cdots > i_n$. This is locally equivalent to the one-dimensional geometric complex corresponding to the monotone root of Theorem 5.3.1. There are two possible forms for this complex, both of which are displayed in Figure 6.2. If n is even (so that the sign in front of X_{i_n} is a minus), then the skeleton of the relevant complex has a J -invariant 0-cell; otherwise, the skeleton has a

J -invariant 1-cell. We have labeled the 0-cells of our complex as v_k and Jv_k (with k even) and the 1-cells of our complex as w_k and Jw_k (with k odd).¹ The grading function gr is given by

$$\text{gr}(v_k) = \text{gr}(Jv_k) = 2(-i_1 + i_2 - \cdots + i_k), \text{ and}$$

$$\text{gr}(w_k) = \text{gr}(Jw_k) = 2(-i_1 + i_2 - \cdots - i_k),$$

with $\text{gr}(v_0) = \text{gr}(Jv_0) = 0$. It is straightforward to compute the (usual) homology of this complex and show that its torsion part is given by the procedure described in Theorem 6.1.3. The result is schematically depicted in Figure 6.3. The reader can check that all self-local equivalences of our complex are isomorphisms; hence (after grading shift) this is the connected homology of $X_{i_1} - X_{i_2} + \cdots \pm X_{i_n}$.




Figure 6.2: Geometric complexes of alternating sums with n even (top) and n odd (bottom).




Figure 6.3: Homology of an alternating sum (n odd). Dotted lines are separated by Maslov grading two, with the uppermost line having Maslov grading zero. This example is $X_5 - X_4 + X_2$.

Example 6.2.2. Consider a sum of X_i with the same orientation; i.e.,

$$X_{i_1} + X_{i_2} + X_{i_3} + \cdots + X_{i_n}$$

with $i_1 \geq i_2 \geq \cdots \geq i_n$. By Lemma 5.1.3 this is locally equivalent to a geometric complex whose skeleton is the J -symmetric decomposition of an n -ball. This is displayed in Figure 6.4. For each $0 \leq k < n$, there are two cells of dimension k , which we have labelled e_k and Je_k , while there is a

¹Note the slight change in indexing from previous examples. For brevity, we also drop the subscript on the involution J for the rest of this section.

single J -symmetric cell e_n of dimension n . The grading function gr is given by

$$\text{gr}(e_k) = \text{gr}(Je_k) = 2(-i_1 - i_2 - \cdots - i_k),$$

with $\text{gr}(e_0) = \text{gr}(Je_0) = 0$. Again, it is straightforward to compute the (usual) homology of this complex and show that its torsion part is given the procedure described in Theorem 6.1.3. (Generators of the torsion towers are given by the obvious cycles $e_k + Je_k$.) The result is schematically depicted in Figure 6.5. The reader can check that all self-local equivalences of our complex are isomorphisms; hence (after grading shift) this is the connected homology of $X_{i_1} + X_{i_2} + \cdots + X_{i_n}$.




Figure 6.4: Geometric complex of same-sign sums.




Figure 6.5: Homology of same-sign sums. Dotted lines are separated by Maslov grading two, with the uppermost line having Maslov grading zero. This example is $X_4 + X_3 + X_2$.

Examples 6.2.1 and 6.2.2 give geometric representatives for the local equivalence classes of alternating sums and same-sign sums of the X_i , respectively. In both cases, the given representatives have the property that the connected homology is given by (the torsion part of) the usual homology. This is not very surprising, as one might guess that the complexes in question are indeed the simplest possible representatives of their local equivalence classes. (Taking the tensor product of n of the X_i without any simplification yields a complex with 3^n generators; our model complex has $2n + 1$ generators.) The goal of future sections will be to produce similar representatives for arbitrary linear

combinations of the X_i .

Note that in the above examples, the action of J has exactly one fixed cell η with $J\eta = \eta$; all other cells occur in J -symmetric pairs. We refer to a geometric complex with this property as a *split complex*. If C_1 and C_2 are two split complexes, then their product is also obviously a split complex. We will further need the following definition involving geometric complexes:

Definition 6.2.3. Let C be a geometric complex with skeleton C_{cell} and grading function gr . For each cell e_i of C_{cell} , define

$$\text{width}(e_i) = \min_{e_{i-1} \in \text{cellular bdry of } e_i} \{\text{gr}(e_{i-1}) - \text{gr}(e_i)\},$$

with the understanding that if $\partial e_i = 0$ then $\text{width}(e_i) = \infty$. We define the *width* of the complex C to be the minimum width over all its cells; i.e.,

$$\text{width}(C) = \min\{\text{width}(e_i)\}.$$

Note that this means any term appearing in the boundary map ∂ on C must appear with a U -power of at least $\text{width}(C)/2$.

6.3 Geometric operations

In this section, we sketch some constructions which will allow us to inductively produce local representatives for linear combinations of the X_i . Our exposition here is meant to convey intuition and form a geometric overview; see Section 6.4 for further details.

Consider the cell complexes displayed in Figure 6.6. On the left, we have displayed the skeleta of several split complexes, while on the right, we have displayed a corresponding series of cell complexes obtained in each case by “doubling” the single J -invariant cell on the left. More precisely, to pass from the left- to the right-hand side, we replace the single J -invariant cell η on the left with a new pair of symmetric cells (denoted by ω and $J\omega$), each of which have the same boundary (if any) as η . We then fill in the resulting “gap” with a single J -invariant cell θ of one higher dimension, whose boundary is the sum of our two new cells.

Now let X be any split complex, and let Δ be any non-negative integer with $2\Delta \leq \text{width}(X)$. We define a new split complex $X^d(\Delta)$, called the *double* of X (with parameter Δ), as follows. The skeleton of $X^d(\Delta)$ is given by applying the above doubling operation to the skeleton of X . The function gr on $X^d(\Delta)$ is defined as follows:

1. On every cell except ω , $J\omega$, and θ , the function gr is equal to its value on the corresponding cell from X ;
2. $\text{gr}(\omega) = \text{gr}(J\omega) = \text{gr}(\eta)$; and,




Figure 6.6: Doublings of split complexes.

$$3. \text{gr}(\theta) = \text{gr}(\eta) - 2\Delta.$$

In Section 6.4 we will show (or at least, claim) that $X^d(\Delta)$ is a well-defined split geometric complex (i.e., that the function gr satisfies the usual monotonicity condition with respect to the cellular differential) and explain the condition $2\Delta \leq \text{width}(X)$. For now, the reader should take a moment to verify that for the representative complexes of Example 6.2.1, the alternating sum $X_{i_1} - X_{i_2} + \cdots + X_{i_n}$ (for n odd) is the double (with parameter i_n) of $X_{i_1} - X_{i_2} + \cdots - X_{i_{n-1}}$. Similarly, for the complexes of Example 6.2.2, we have that $X_{i_1} + X_{i_2} + \cdots + X_{i_n}$ is the double (with parameter i_n) of $X_{i_1} + X_{i_2} + \cdots + X_{i_{n-1}}$. Compare with the first two examples in Figure 6.6.

The main technical claim, of course, is that the doubled complex $X^d(\Delta)$ is locally equivalent to $X \otimes X_\Delta$ in general. This allows us to construct local representatives of linear combinations of the X_i via induction. Note that passing from $X \otimes X_\Delta$ to $X^d(\Delta)$ represents a significant reduction in complexity, since if X has g generators, then $X \otimes X_\Delta$ has $3g$ generators, while $X^d(\Delta)$ has $g + 2$ generators.

Remark 6.3.1. The relation between the cellular boundary maps for X and $X^d(\Delta)$ can be somewhat subtle. For instance, in the fourth example of Figure 6.6, we must modify the boundary map on the 2-cells incident to our original J -invariant 0-cell η to include the new J -invariant 1-cell θ . The reader may thus (rightly) object that we have not really defined how to construct the skeleton of

$X^d(\Delta)$. While in the examples above it is more or less obvious how to “double” the J -invariant cell of X , one might wonder whether this is always well-defined, or indeed whether there may be multiple ways to carry out this operation. For example, Figure 6.7 shows a split complex with two putative doublings. In Section 6.4, we will explicitly write down the doubled complex and its boundary map; this will turn out to depend on an additional set of choices. However, it turns out that *any* choice of doubling is locally equivalent to $X \otimes X_\Delta$, so for our purposes this ambiguity is unimportant.




Figure 6.7: A split complex which appears to have two possible doublings.

We now briefly consider the process of dualizing a geometric complex. If X is a geometric complex, then X^* is not quite a geometric complex in the sense defined previously. For example, the dual of X_i has three generators α^* , $J\alpha^*$, and β^* , with the differential

$$\partial^*(\alpha^*) = \partial^*(J\alpha^*) = U^i \beta^*$$

and $\partial^* \beta^* = 0$. Hence the putative “skeleton” of X_i^* has two 1-cells (corresponding to α^* and $J\alpha^*$), each having boundary consisting of only one 0-cell (corresponding to β^*). While this evidently does not correspond to a cell complex in the usual sense, it is not difficult to see that we may view this as a *relative* cell complex, as in Figure 6.8. Here, the homology is no longer isomorphic to the usual singular homology of our cell complex, but rather the relative homology of our cell complex with respect to some subcomplex.

We thus relax the definition of a geometric complex slightly by not requiring the skeleton to be an actual cell complex, but simply a $\mathbb{Z}$ -graded chain complex over $\mathbb{F}$. (We will still refer to generators of the skeleton as cells, and the $\mathbb{Z}$ -grading as the dimensional grading.) All previous notions such as split complexes, width, and so on, apply equally well to this more general class of complexes; as we will see in Section 6.4, there will be no trouble in defining the doubling operation for these complexes also. If X has skeleton C_{cell} with generators e_i , then we may consider the dual complex C_{cell}^* spanned by the obvious generators e_i^* , where e_b^* appears in $\partial^* e_a^*$ if and only if e_a appears in ∂e_b . It is easy to check that for an appropriate choice of gr , X^* is then a (generalized) geometric complex

over C_{cell}^* , which we can often think of as a relative cell complex. The duals of Example 6.2.1 and Example 6.2.2 are shown in Figure 6.8. See [Dai18a, Section 3.3] for details.




Figure 6.8: The skeleton of X_i^* (left); duals of Example 6.2.1 (middle); dual of Example 6.2.2 (right). Here, crosses or dashed lines indicate the “relative” part of the cell complex. For example, on the right-hand side, we have $\partial^* e_0^* = e_1^* + J e_1^*$.

We now geometrically describe the operation dual to doubling. This corresponds to tensoring with X_Δ^* (rather than X_Δ). Consider the cell complexes displayed in Figure 6.9. On the left, we have displayed the skeleta of several split complexes, while on the right we have displayed a corresponding sequence of complexes obtained in each case by “halving” the J -invariant cell on the left. This consists of replacing the old J -invariant cell η^* with the sum of a new pair of symmetric cells (denoted by ω^* and $J\omega^*$), which meet along a new J -invariant cell θ^* of one lower dimension.




Figure 6.9: Halvings of split complexes.

There is a slight additional subtlety in the case that η^* is a 0-cell. In this situation, we must first take the product of our complex with the one-dimensional complex displayed on the left in Figure 6.10. This has one generator, which we think of as a 1-cell instead of a 0-cell. The value of gr on this generator is defined to be -1 , so that its Maslov grading is zero. Taking the product thus algebraically does not change our chain complex, but boosting the dimension by one in the geometric picture allows us to “half” the J -invariant cell, which is now one-dimensional. (We are simply shifting the dimensional grading up by one and decreasing gr by one, so that the Maslov

grading remains the same.) See Figure 6.10.




Figure 6.10: Halving a complex X with a J -invariant 0-cell. First, we cross X with the complex on the left to obtain an algebraically equivalent complex X' , but which we view as having increased dimension. Note that in this example, the result of halving gives a complex equivalent to the one displayed on the right in Figure 6.8.

Now let X be any split complex, and let Δ be any non-negative integer with $2\Delta \leq \text{width}(X)$. The *half* of X (with parameter Δ), denoted by $X^h(\Delta)$, is constructed by applying the above halving operation to the skeleton of X and defining gr as follows:

1. On every cell except ω^* , $J\omega^*$, and θ^* , the function gr is equal to its value on the corresponding cell from X ;
2. $\text{gr}(\omega^*) = \text{gr}(J\omega^*) = \text{gr}(\eta^*)$; and,
3. $\text{gr}(\theta^*) = \text{gr}(\eta^*) + 2\Delta$.

Again, the reader may legitimately worry that we have not rigorously defined the halving operation on skeleta. For now, it is perhaps better to take this as a motivational picture and verify that the halving operation is well-defined and produces the desired complexes in the case of Example 6.2.1 and (the dual of) Example 6.2.2.

6.4 Algebraic details

For completeness, we now give a precise construction of the doubling operation. Let X be a split geometric complex with skeleton C_{cell} . Recall this means that C_{cell} has exactly one J -invariant cell η , and all the other cells of C_{cell} occur in J -symmetric pairs. We partition the cells of C_{cell} by choosing one cell from each such pair and denoting the span of these cells by S . Then C_{cell} (as a module over $\mathbb{F}$) decomposes as a direct sum

$$C_{\text{cell}} = S \oplus JS \oplus \eta,$$

where we have written η to mean the (single-element) submodule spanned by η . Note that this splitting is not preserved by ∂ . Usually when we discuss split complexes we will have in mind a particular splitting obtained by choosing some S .

It will be useful to also consider the related splitting $S \oplus (1 + J)S \oplus \eta$, in which case we can write

the differential of any $x \in S$ uniquely as

$$\partial x = \begin{cases} a + (1 + J)b, & \text{or} \\ a + (1 + J)b + \eta \end{cases}$$

for $a, b \in S$. Similarly, the fact that J commutes with ∂ and that $J\eta = \eta$ implies that

$$\partial\eta = (1 + J)\zeta$$

for some $\zeta \in S$. Note that a , b , and ζ are allowed to be sums of cells and/or zero.

We now make a proper definition of $X^d(\Delta)$. We begin with the skeleton of $X^d(\Delta)$ as a set:

Definition 6.4.1. Let X be a split complex with a particular choice of splitting S , and let Δ be a non-negative integer with $2\Delta \leq \text{width}(X)$. As a module over $\mathbb{F}$, the skeleton of $X^d(\Delta)$ is defined to be the direct sum

$$(S \oplus \omega) \oplus (JS \oplus J\omega) \oplus \theta.$$

Here, S and JS are copies (as modules over $\mathbb{F}$) of their corresponding summands in C_{cell} , while ω , $J\omega$, and θ are new generators in the skeleton of $X^d(\Delta)$. (We have again abused notation and written individual generators to mean their span over $\mathbb{F}$.) The action of J interchanges C and JC as before, interchanges ω and $J\omega$, and sends θ to itself. Note that this new skeleton is also split, with a choice of splitting given by $(S \oplus \omega)$.

The differential on $X^d(\Delta)$ is defined as follows:

Definition 6.4.2. The differential ∂^d on the skeleton of $X^d(\Delta)$ is defined by modifying the original differential ∂ on the skeleton of X to take into account the new cells. We have the following casework:

1. Let $x \in S$. Then there are three cases:

$$\partial^d x = \begin{cases} a + (1 + J)b, & \text{if } \partial x = a + (1 + J)b \text{ and } \partial b \text{ does not contain } \eta; \\ a + (1 + J)b + \theta, & \text{if } \partial x = a + (1 + J)b \text{ and } \partial b \text{ does contain } \eta; \\ a + (1 + J)b + \omega, & \text{if } \partial x = a + (1 + J)b + \eta. \end{cases}$$

2. $\partial^d \omega = \partial\eta = (1 + J)\zeta$, and

3. $\partial^d \theta = (1 + J)\omega$.

The differential on JS and $J\omega$ is defined by requiring J -equivariance.

We will not prove that ∂^d is an actual differential here; the proof is an exercise in straightforward but tedious casework. See [Dai18a, Lemma 4.3].

Remark 6.4.3. The three cases for $\partial^d x$ above correspond to different situations as in Figure 6.6. The first case is the most straightforward, where x is either sufficiently “far away” from η or is of

very different dimension, and the boundary map on x does not need to be modified. The second case corresponds to a situation in which the boundary of x must be modified to include the new J -invariant cell θ . (Consider the 2-cell in the fourth example of Figure 6.6.) Finally, the third case corresponds to the situation where the original boundary of x contains η , and η must be replaced by the corresponding cell ω . (Consider the 2-cell in the third example of Figure 6.6.)

To complete our definition of $X^d(\Delta)$, we finally define the grading function gr on $X^d(\Delta)$:

Definition 6.4.4. We set:

1. If $x \in S$ (or JS), then $\text{gr}(x)$ is the same in $X^d(\Delta)$ as it is in X ;
2. $\text{gr}(\omega) = \text{gr}(J\omega) = \text{gr}(\eta)$; and
3. $\text{gr}(\theta) = \text{gr}(\eta) - 2\Delta$.

Note that the grading of θ is the only place where we use the parameter Δ . We remind the reader that the Maslov grading is equal to the sum of the dimensional grading and gr . The dimensions of the cells in S and JS are the same as before, while the dimension of ω is equal to that of η , and the dimension of θ is equal to the dimension of η plus one.

In order to verify that this data forms a well-defined geometric complex, we must check the usual monotonicity condition on gr with respect to ∂^d . In all situations except the second case of (1) in Definition 6.4.2, this is either immediate or it follows from the fact that the original differential ∂ has this property. In the remaining case, it follows from the fact that $\text{gr}(\theta) = \text{gr}(\eta) - 2\Delta$ and $2\Delta \leq \text{width}(X)$. More precisely, in this situation we know that some summand b_0 of b has η appearing in its differential. Hence $\text{gr}(x) \leq \text{gr}(b_0) \leq \text{gr}(\eta)$. The condition on Δ ensures that $\text{gr}(\theta)$ lies between $\text{gr}(b_0)$ and $\text{gr}(\eta)$, so in particular $\text{gr}(x) \leq \text{gr}(\theta)$, as desired. Note that $X^d(\Delta)$ is obviously split, and has width 2Δ . We think of the skeleton of $X^d(\Delta)$ as being obtained from the skeleton of X via the doubling operation described above, although the reader who is uncomfortable with this point of view can of course rely solely on the algebraic definition.

6.5 Sketch of proof

We now briefly explain how to deduce Theorem 6.1.3 from the machinery that we have established so far. Our first claim is of course:

Theorem 6.5.1. [Dai18a, Theorem 4.7] Let X be a split geometric complex, and let $2\Delta \leq \text{width}(X)$. Then $X^d(\Delta)$ is locally equivalent to $X \otimes X_\Delta$.

The proof of Theorem 6.5.1 can be found in see [Dai18a]. The argument proceeds as in the proof of Lemma 5.1.3, except that explicitly writing down the maps between skeleta is rather more complicated. Establishing the lifting condition is then similar to the verification of the monotonicity condition on ∂^d . The reader may examine Example 6.5.2, which gives some flavor of the argument.

Example 6.5.2. In Figure 6.11, we have displayed a sample split complex X , together with its double $X^d(\Delta)$ and the skeleton of the tensor product $X \otimes X_\Delta$. (The latter is just a cylinder over the skeleton of X .) In X , we have colored the J -invariant cell η purple, and chosen a splitting S given by the span of the red cells. The corresponding cells of S in $X^d(\Delta)$ are similarly colored red; we have colored ω purple and θ blue.

In the second row, we have displayed a cellular map f from the skeleton of $X^d(\Delta)$ to the skeleton of $X \otimes X_\Delta$. Here, the red 0-cell and red horizontal 1-cell are sent to their obvious counterparts in $X \otimes X_\Delta$. However, the semicircular red 1-cell is sent to the sum of its corresponding 1-cell in $X \otimes X_\Delta$, plus an additional red 1-cell running down the length of the cylinder. Similarly, the red 2-cell is sent to the sum of the two red 2-cells indicated on the right. The purple and blue cells are mapped bijectively to each other. The reader should verify (carefully) that extending via J -equivariance yields a cellular chain map. Roughly speaking, one can think of f as smearing $X^d(\Delta)$ over the cylinder in a J -equivariant fashion.

In the third row, we have displayed a map g in the other direction. Here, the red, green, purple, and blue cells are all mapped to each other in the obvious way, with the exception that the green 1-cell is mapped to the sum of the green and blue 1-cells on the right. All cells (other than the blue 1-cell) which run along the length of the cylinder are sent to zero. Again, the reader should verify (carefully) that extending via J -equivariance yields a cellular chain map.

If one accepts that Theorem 6.5.1 holds, then we immediately obtain a simplified representative for any linear combination of the X_i , simply by repeatedly applying the doubling operation and dualizing whenever necessary. It is not difficult to prove that the homology of the resulting complex is given by the procedure of Algorithm 6.1.2. Indeed, a straightforward exact sequence argument shows that doubling adds a single U -torsion tower of length Δ to the homology:

Lemma 6.5.3. [Dai18a, Lemma 5.1] Let X be a split complex whose single J -invariant cell has Maslov grading g . Let $2\Delta \leq \text{width}(X)$, and let $X^d(\Delta)$ be the double of X with parameter Δ . Then

$$H_*(X^d(\Delta)) = H_*(X) \oplus \mathcal{T}_g(\Delta).$$

That is, the homology of $X^d(\Delta)$ is isomorphic to the direct sum of the homology of X , plus a single U -torsion tower of length Δ which starts in grading g .




Figure 6.11: Local maps between $X^d(\Delta)$ and $X \otimes X_\Delta$.

Proof. Let X' be the subcomplex of $X^d(\Delta)$ spanned over $\mathbb{F}[U]$ by the two generators $\omega + J\omega$ and θ . Since $\omega + J\omega$ has Maslov grading g and $\partial^d \theta = U^\Delta(\omega + J\omega)$, it is clear that the homology of X' is the U -torsion tower $\mathcal{T}_g(\Delta)$. Now observe that the quotient complex $X^d(\Delta)/X'$ is isomorphic to X via the map which sends ω and $J\omega$ (which are identified in the quotient complex) to η, θ (which is zero in the quotient complex) to zero, and is the obvious correspondence on all the other generators. We thus have a long exact sequence of U -equivariant maps:

$$\cdots \rightarrow H_*(X') \rightarrow H_*(X^d(\Delta)) \rightarrow H_*(X^d(\Delta)/X') = H_*(X) \rightarrow \cdots$$

We begin by showing that the boundary maps in this sequence are zero. To see this, it suffices to prove that the induced inclusion map $i_* : H_*(X') \rightarrow H_*(X^d(\Delta))$ on the level of homology is injective. The non-zero elements of $H_*(X')$ are given by the classes $[U^i(\omega + J\omega)]$ with $0 \leq i < \Delta$. Viewed as elements of $X^d(\Delta)$, the cycles $U^i(\omega + J\omega)$ are still non-zero in homology, since any element in the image of the boundary map ∂^d on $X^d(\Delta)$ must appear with a U -power of at least Δ . This shows that we have a short exact sequence of U -equivariant maps.

Proving that the sequence splits is also straightforward. We know that the inclusion of X' into $X^d(\Delta)$ induces a $\mathbb{F}[U]$ -equivariant isomorphism of $H_*(X')$ onto its image in $H_*(X^d(\Delta))$. To show that the sequence splits, it suffices to show that this image is in fact a direct summand of $H_*(X^d(\Delta))$. This will be true if the class $[\omega + J\omega]$ in $H_*(X^d(\Delta))$ is not in the image of multiplication by U . (In

such a situation, it is easy to see that we can extend $[\omega + J\omega]$ to an $\mathbb{F}[U]$ -basis for all of $H_*(X^d(\Delta))$.) But the only way for this to fail would be for $\omega + J\omega$ to be homologous to a U -power of some other cycle; i.e.,

$$\partial x = (\omega + J\omega) + Uy$$

for some cycle y and bordism element x . Because the width of $X^d(\Delta)$ is Δ , this is impossible unless $\Delta = 0$, in which case we have $H_*(X^d(\Delta)) = H_*(X)$ and the theorem is trivially verified. $\square$

According to [HHL18, Proposition 4.1], the connected homology behaves well under dualizing. Keeping track of gradings, Lemma 6.5.3 then easily implies that the homology of our model complex is given by Algorithm 6.1.2. For details, see [Dai18a, Section 5].

It thus remains to show that our model complex actually computes the connected homology. Note that since the connected homology is obviously a submodule of the usual $\mathbb{F}[U]$ -homology, the above result immediately shows the connected homology of X must be a submodule of the procedure given in Algorithm 6.1.2. Hence to prove Theorem 6.1.3, it suffices to show that the connected homology must be at least as large as the output of Algorithm 6.1.2.

Proof of Theorem 6.1.3. Let X be a maximally simplified linear combination of the X_i . Write

$$X = c_1 X_{i_1} + c_2 X_{i_2} + \cdots + c_n X_{i_n}$$

with $i_1 > i_2 > \cdots > i_n$. We proceed by induction on n . For $X = c_1 X_{i_1}$, the claim is the result of Example 6.2.2. Thus, assume that we have established the claim for

$$A = c_1 X_{i_1} + c_2 X_{i_2} + \cdots + c_{n-1} X_{i_{n-1}}.$$

Write $B = c_n X_{i_n}$, so that $X = A \otimes B$.

We begin by showing that $H_{\text{conn}}(X)$ must have at least $|c_k|$ towers of length i_k , for each $1 \leq k \leq n-1$. To see this, first consider the local equivalence

$$A = X \otimes (-B).$$

This shows that $H_{\text{conn}}(A)$ is a summand of $H_*(X \otimes (-B))$, where (without loss of generality) we may replace X with a locally equivalent complex so that $H_{\text{conn}}(X)$ is given by the torsion part of $H_*(X)$. Now, by the inductive hypothesis, $H_{\text{conn}}(A)$ consists of exactly $|c_k|$ towers of length i_k , for each $1 \leq k \leq n-1$. The torsion homology of $H_*(-B)$ consists of exactly $|c_n|$ towers, each of length i_n . Since i_n is not equal to any other i_k , it follows from an application of the Künneth theorem and an examination of the Tor functor that the only way for $H_{\text{conn}}(A)$ to be a summand of $H_*(X \otimes (-B))$ is for $H_{\text{conn}}(X)$ to contain at least $|c_k|$ towers of length i_k , for each $1 \leq k \leq n-1$.

We now show that $H_{\text{conn}}(X)$ must have at least $|c_n|$ towers of length i_n . This follows similarly

by observing that we have the local equivalence

$$B = (-A) \otimes X.$$

Indeed, this implies that $H_{\text{conn}}(B)$ is a summand of $H_*((-A) \otimes X)$, where we may replace each factor with a locally equivalent complex whose connected homology is given by its usual homology. Now, $H_{\text{conn}}(B)$ consists of exactly $|c_n|$ towers, each having length i_n . By the inductive hypothesis, all torsion towers appearing in $H_{\text{conn}}(A)$ have lengths distinct from i_n . It follows from the Künneth theorem and an examination of the Tor functor that $H_{\text{conn}}(X)$ must contain at least $|c_n|$ towers of length i_n , as desired.

By our previous work, we have a representative complex Σ for X whose torsion homology is already given by the construction claimed in the statement of Theorem 6.1.3. Note that $H_*(\Sigma)$ has exactly $|c_k|$ towers of length i_k , for each $1 \leq k \leq n$. The above paragraphs then imply that the connected homology of X must be all of this homology, as desired. $\square$

6.6 Examples and applications

We now give an illustration of Algorithm 6.1.2, in which we give the actual complexes constructed via the doubling and halving operations of Section 6.3:

Example 6.6.1. Consider a maximally simplified linear combination of the form

$$X = X_{i_1} + X_{i_2} - X_{i_3} - X_{i_4} + X_{i_5}.$$

This is locally equivalent to the complex constructed via a sequence of iterated doubling and halving operations, displayed on the left in Figure 6.12. We begin with the geometric complex of X_{i_1} , double to obtain $X_{i_1} + X_{i_2}$, half to obtain $X_{i_1} + X_{i_2} - X_{i_3}$, and so on. The values of gr are similarly defined inductively; the final result is displayed in the upper-right in Figure 6.12.

One can verify by hand that the homology of this complex has the form prescribed by Algorithm 6.1.2, as follows. The first torsion tower is generated by the sum of the two zero-graded 0-cells, while the second tower is generated by the sum of the two semicircular 1-cells. The other generators are more subtle: the third tower is generated by a closed loop running around either the upper or the lower hemisphere, consisting of the sum of (the appropriate U -powers of) four 1-cells. The fourth tower is generated by either of the interior 0-cells, plus the appropriate U -power of either of the zero-graded 0-cells. Finally, the fifth tower is generated by the sum of the two interior 0-cells. In general, it is rather tricky to identify the generators of the torsion towers directly in the representative complex, which is why we have instead used Lemma 6.5.3 for the proof of Theorem 6.1.3.




Figure 6.12: Doubling and halving operations (left); final values of gr (upper-right); homology (lower-right).

Finally, we close with an example of an ι -complex whose connected homology is not of the form prescribed by Theorem 6.1.3:

Example 6.6.2. Consider the geometric complex whose skeleton is given by the usual J -equivariant cellular decomposition of D^2 , as in Figure 6.13. While this is the same skeleton as for $X_{i_1} + X_{i_2}$, we put a “misordered” grading function on D^2 , as follows. Fix integers $0 < x < y$. Let the 0-cells have grading zero, the 1-cells have grading $-2x$, and the 2-cell have grading $-2(x + y)$. Explicitly, this means that our ι -complex is generated over $\mathbb{F}[U]$ by the set $\{e_0, Je_0, e_1, Je_1, e_2 = Je_2\}$, with

$$\partial e_1 = \partial Je_1 = U^x(e_0 + Je_0) \text{ and}$$

$$\partial e_2 = U^y(e_1 + Je_1),$$

where $x < y$. The resulting connected homology is displayed (modulo the grading shift by one) on the right in Figure 6.13. This has two U -torsion towers of lengths x and y , but the tower of shorter length occurs in higher grading, in contrast to the connected homology of $X_{i_1} + X_{i_2}$. It is easily checked that the connected homology of Figure 6.13 (keeping in mind its placement in relation to the d -invariant) cannot be achieved through the procedure of Algorithm 6.1.2.




Figure 6.13: A connected homology not arising via the procedure of Algorithm 6.1.2.

Again, we stress that it is not known whether this is realized as the ι -complex of any actual 3-manifold. Indeed, such an example would show that Seifert fibered spaces do not generate the homology cobordism group. One possible source of examples includes plumbed homology spheres with two bad vertices, in the sense of [OS03b]. For these manifolds, the lattice homology (together with the lattice involution) is frequently locally equivalent to a geometric complex with skeleton D^2 . Moreover, the usual Heegaard Floer homology (forgetting the involution) is known to be isomorphic to the (generalized) lattice homology, as defined in Section 3.2.

It thus seems reasonable to conjecture that the ι -complexes of these manifolds are locally equivalent to the models coming from their lattice homology. It is then a matter of searching for a plumbing as in Example 6.6.2, with misordered U -torsion towers. Unfortunately, the author is not currently aware of any such example. See [Ném08] for several sample calculations.

Chapter 7

Pin(2)-Equivariant Monopole Floer Homology

In this chapter, we present the results of [Dai18b] on the Pin(2)-equivariant monopole Floer homology of Seifert fibered spaces. Although the results given in this chapter are significantly less comprehensive than in the involutive Floer case, we will still see that a surprising amount of information can be obtained from knowledge of the lattice homology.

7.1 Monopole Floer homology

We begin by reviewing the essential tenets of Pin(2)-equivariant monopole Floer homology, as defined by Lin in [Lin18a]. As usual, for simplicity we assume that Y is an integer homology sphere. Recall that the monopole Floer homology groups are defined by studying the chain complex generated by critical points of the (perturbed) Chern-Simons-Dirac functional on a certain “configuration space” associated to Y . (See e.g. [KM07] for details.) This is a gauge-theoretic invariant which assigns to Y three groups fitting into the long exact sequence

$$\cdots \rightarrow \overline{HM}(Y) \rightarrow \widetilde{HM}(Y) \rightarrow \widehat{HM}(Y) \rightarrow \cdots .$$

Each of these groups are $\mathbb{Z}$ -graded modules over the ring $\mathbb{F}[U]$, which may be thought of as the S^1 -equivariant cohomology of a point. The action of U has degree -2 .

It furthermore turns out that the (unperturbed) Chern-Simons-Dirac functional has an extra $\mathbb{Z}/2\mathbb{Z}$ -symmetry, which allows us to consider the subcomplex consisting of $\mathbb{Z}/2\mathbb{Z}$ -invariant chains of critical points and flows between such chains. The main analytical difficulty in doing this is

that for perturbations of the Chern-Simons-Dirac functional which preserve the $\mathbb{Z}/2\mathbb{Z}$ -symmetry, critical points necessarily occur in entire submanifolds, rather than isolated points. In [Lin18a], Lin constructs an appropriate theory using a family of such “Morse-Bott” perturbations. This yields a similar triple of gauge-theoretic invariants fitting into the analogous long exact sequence

$$\cdots \rightarrow \overline{HS}(Y) \rightarrow \widetilde{HS}(Y) \rightarrow \widehat{HS}(Y) \rightarrow \cdots .$$

Each of these groups are modules over the ring

$$\mathcal{R} = \mathbb{F}[V, Q]/(Q^3),$$

which may be thought of as the $\text{Pin}(2)$ -equivariant cohomology of a point. If Y is a homology sphere, then (as in the non-equivariant case) the $\text{Pin}(2)$ -equivariant monopole Floer homology inherits an absolute $\mathbb{Z}$ -grading. The actions of V and Q have degrees -4 and -1 , respectively.

As in the involutive Heegaard Floer case, we can use this extra algebraic structure to define an additional suite of homology cobordism invariants, called the *Manolescu correction terms*. Let $\mathcal{V}_r^+$ be the $\mathbb{F}[V]$ -module $\mathbb{F}[V^{-1}, V]/V\mathbb{F}[V]$, shifted so that the element 1 has grading r . Then the $\text{Pin}(2)$ -equivariant monopole Floer homology decomposes into the direct sum of a finite part and the sum of three infinite towers

$$\widetilde{HS}(Y) = \mathcal{V}_a^+ \oplus \mathcal{V}_b^+ \oplus \mathcal{V}_c^+ \oplus \text{finite part},$$

where the action of Q sends the c -tower onto the b -tower and the b -tower onto the a -tower. Note that a , b , and c are well-defined, even though the finite part of the splitting is non-canonical. We define $\alpha(Y) \geq \beta(Y) \geq \gamma(Y)$ to be the integers such that

$$a = 2\alpha(Y), \quad b = 2\beta(Y) + 1, \quad \text{and} \quad c = 2\gamma(Y) + 2.$$

These correction terms can also be defined in terms of $\widehat{HS}(Y)$, in which case the definitions more closely parallel those for $\bar{d}$ and $\underline{d}$. See [Man16] for Manolescu’s disproof of the triangulation conjecture using these invariants.

In this paper, we will actually only need to understand some simple formal properties of the $\text{Pin}(2)$ -equivariant monopole Floer homology. The most important of these is the relation between the usual and $\text{Pin}(2)$ -equivariant monopole Floer homologies. This is expressed by the Gysin exact sequence (see Section 4.3 of [Lin18a]):

$$\cdots \xrightarrow{\cdot Q} \widetilde{HS}(Y) \xrightarrow{\iota_*} \widetilde{HM}(Y) \xrightarrow{\pi_*} \widetilde{HS}(Y) \xrightarrow{\cdot Q} \widetilde{HS}(Y) \xrightarrow{\iota_*} \cdots, \quad (7.1)$$

which should be thought of as analogous to the usual Gysin sequence in algebraic topology for S^0 -bundles. (There are of course similar sequences for the other two flavors of monopole Floer homology.) This is a map of graded $\mathcal{R}$ -modules, where V acts on the usual monopole Floer homology as U^2 and Q acts as zero. The maps ι_* and π_* in the Gysin sequence preserve the grading, while multiplication by Q has grading -1 . In very simple cases, one can use the Gysin sequence to determine

the $\text{Pin}(2)$ -equivariant monopole Floer homology from the usual monopole Floer homology. This is done for example by Lin in [Lin17, Section 4]. Here, we will show that with a little bit of additional information coming from the lattice homology, the Gysin sequence determines the $\text{Pin}(2)$ -homology for all Seifert fibered spaces.

As in the Heegaard Floer case, if Y is a Seifert fibered homology sphere, then the monopole Floer homology is known to be isomorphic to the lattice homology. Following the exposition of [Dai18b], we will actually work with the lattice cohomology as discussed in Section 4.1, although it is not difficult to carry out the argument analogously for $\mathbb{H}^-(\Gamma, k)$. Recall that we have:

Theorem 7.1.1 (Theorem 1.2 [OS03b]; Theorem 5.2.2 [Ném08]). Let Y be a Seifert fibered homology sphere, and let Γ be any negative-definite plumbing graph for Y . Then we have an isomorphism of $\mathbb{F}[U]$ -modules:

$$\widetilde{HM}(-Y) \cong \mathbb{H}^+(\Gamma).$$

Note that $\mathbb{H}^+(\Gamma)$ can be obtained by taking the usual graded root coming from $\mathbb{H}^-(\Gamma, k)[-(|\Gamma| + k^2)/4 + 2]$ and reflecting it over the horizontal line of absolute grading -1 .

It should be noted that the proof of Theorem 7.1.1 follows exactly the same argument as given for the Heegaard Floer homology in [OS03b]. More precisely, the isomorphism of Theorem 7.1.1 is effected by sending an element y in the monopole Floer homology of $-Y$ to the evaluation map $\phi_y : \text{Char}(\Gamma) \rightarrow \mathcal{U}_0^+$ defined by

$$\phi_y(k) = \widetilde{HM}(W(\Gamma), PD[k])(y)$$

for any characteristic vector k . Here,

$$\widetilde{HM}(W(\Gamma), PD[k]) : \widetilde{HM}(-Y) \rightarrow \widetilde{HM}(S^3)$$

is the monopole Floer cobordism map associated to the spin^c structure on W with $c_1 = PD[k]$.

Note that we have identified the monopole Floer homology of S^3 with $\mathcal{U}_0^+$ (see e.g. [KM07]).

7.2 The Gysin sequence

As in the involutive Heegaard Floer case, our computation of the $\text{Pin}(2)$ -equivariant monopole Floer homology will rely on an understanding the lattice involution J_0 . To this end, we have:

Theorem 7.2.1. [Dai18a, Lemma 2.5] Let Y be a Seifert fibered homology sphere. Then, under the isomorphism of Theorem 7.1.1, the composition of maps $\iota_* \circ \pi_*$ in the Gysin sequence for $-Y$ coincides with the action of $1 + J_0$ on $\mathbb{H}^+(\Gamma)$.

Proof. Denote by j the involution on the chain groups of the usual monopole Floer homology of $-Y$ coming from the $\mathbb{Z}/2\mathbb{Z}$ -symmetry of the Chern-Simons-Dirac functional.¹ We claim that the induced action of j_* on the monopole Floer homology coincides with the action of J_0 under the isomorphism of Theorem 7.1.1. This follows from the same argument as in Theorem 4.1.1. Indeed, the monopole Floer cobordism map is defined by studying the moduli space of solutions to the four-dimensional Seiberg-Witten equations over the pair $(W(\Gamma), PD[k])$ with some specified limiting behavior at the ends of the cobordism (see e.g. [KM07]). As in the three-dimensional case, there is a j -symmetry of the four-dimensional Seiberg-Witten equations, which (roughly speaking) restricts to the three-dimensional j -symmetry at the ends of the cobordism. However, since k need not be self-conjugate, in general this symmetry merely identifies the two moduli spaces of the four-dimensional equations corresponding to $(W(\Gamma), PD[k])$ and $(W(\Gamma), -PD[k])$. Points in the former limit to y at one end, while points in the latter limit to jy . It follows that we have the equality

$$\phi_y(k) = \phi_{jy}(-k)$$

for all $y \in \widetilde{HM}(-Y)$ and $k \in \text{Char}(\Gamma)$. Thus the correspondence $y \mapsto \phi_y$ takes the involution j on the monopole Floer homology to the map on the lattice cohomology given by precomposing with the reflection involution $k \mapsto -k$.

We now claim that the map on $\widetilde{HM}(-Y)$ induced by $1 + j$ is equal to the composition of maps $\iota_* \circ \pi_*$ in the Gysin sequence (7.1). To see this, we simply observe that the Gysin sequence is the homology exact sequence associated to the short exact sequence of chain complexes

$$0 \rightarrow C_*^{inv} \xrightarrow{i} C_* \xrightarrow{1+j} (1+j)C_* \rightarrow 0.$$

See [Lin18a, Proposition 3.10]. Here, C_* is the chain complex for the usual monopole Floer homology of $-Y$, and C_*^{inv} is the subcomplex of j -invariant chains. The map i is given by inclusion, and the subcomplex $(1+j)C_*$ is easily shown to be quasi-isomorphic to C_*^{inv} via the inclusion of $(1+j)C_*$ in C_*^{inv} . Thus $\iota_* \circ \pi_*$ is evidently induced by the map $1 + j$ on C_* . $\square$

We now make a precise re-formulation of Theorem 1.3.7. As in Theorem 4.2.1, this will make use of a few auxiliary objects derived from the lattice cohomology. Let $\ker(1 + J_0)$ and $\text{coker}(1 + J_0)$ denote the kernel and cokernel of the map $1 + J_0$ on the lattice cohomology $\mathbb{H}^+(\Gamma)$. Note that these are $\mathbb{F}[U]$ -modules. For convenience, we also define

$$\mathbb{H}' = \ker(1 + J_0) / \text{im}(1 + J_0).$$

Because $\mathbb{H}^+(\Gamma)$ takes the form of an (inverted) graded root, it is immediate that $\mathbb{H}'$ is at most one-

¹In the usual set-up of [KM07], the perturbation of the Chern-Simons-Dirac functional is not taken to be j -equivariant. Hence it is not clear that after perturbation, we have a symmetry map j on the monopole Floer chain complex. By Corollary 3.7 of [Lin18a], however, the monopole Floer homology defined using a j -equivariant Morse-Bott perturbation is canonically isomorphic to the usual monopole Floer homology.

dimensional in each grading and is in fact isomorphic to $\mathcal{U}_\rho^+$, where ρ is the grading of the lowest J_0 -invariant vertex. It follows from the proof of Theorem 1.3.2 that

$$\rho = 2\bar{\mu}(Y).$$

In what follows, it will also be helpful for us to fix a particular basis for $\mathbb{H}^+(\Gamma)$. First, let F be the span (over $\mathbb{F}$) of the vertices lying on the infinite stem. As an abelian group, this may be identified with $\mathbb{H}'$. Secondly, there are vertices which occur in pairs E_i and $J_0 E_i$; we denote the span of the $\{E_i\}$ and the $\{J_0 E_i\}$ by E . We thus obtain a decomposition

$$\mathbb{H}^+(\Gamma) = E \oplus F = \text{span}\{E_i\} \oplus \text{span}\{(1 + J_0)E_i\} \oplus F.$$

Note that the above splitting does not respect the U -action. While it is easily seen that U maps the span of $\{(1 + J_0)E_i\}$ into itself, the image of F under U in general lies in the subspace $\text{span}\{(1 + J_0)E_i\} \oplus F$. Observing that $\ker(1 + J_0) = \text{span}\{(1 + J_0)E_i\} \oplus F$ and $\text{im}(1 + J) = \text{span}\{(1 + J_0)E_i\}$, we evidently have an isomorphism of graded $\mathbb{F}[U]$ -modules

$$\mathbb{H}' \cong (\text{span}\{(1 + J_0)E_i\} \oplus F) / \text{span}\{(1 + J_0)E_i\},$$

where we have written the right-hand side as a quotient to emphasize the U -module structure.

In order to describe the $\text{Pin}(2)$ -homology as an $\mathcal{R}$ -module, it will be convenient for us to break it up into four $\mathbb{F}[V]$ -submodules. To this end, we establish the following notation. Let A and B be two $\mathbb{Z}$ -graded $\mathbb{F}[V]$ -modules with graded parts A_q and B_q , and let $n \in \mathbb{Z}/4\mathbb{Z}$. We write $A =_{[n]} B$ if the submodules

$$A_{[n]} = \bigoplus_{q \equiv n \pmod{4}} A_q \quad \text{and} \quad B_{[n]} = \bigoplus_{q \equiv n \pmod{4}} B_q$$

consisting of the sums of the groups in gradings congruent to n modulo 4 are isomorphic as graded $\mathbb{F}[V]$ -modules. We refer to the above $\mathbb{F}[V]$ -submodules as the $[n]$ -submodules of A and B . Note that if A is in fact a $\mathcal{R}$ -module, then in addition we have an action of Q which maps $A_{[n]}$ into $A_{[n-1]}$ for each n .

Theorem 7.2.2. [Dai18b, Theorem 2.3] Let Y be a Seifert fibered homology sphere. Then we have the following set of $\mathbb{F}[V]$ -module isomorphisms:

$$\begin{aligned} \widetilde{HS}(-Y)_{[\rho+3]} &= 0, \\ \widetilde{HS}(-Y)_{[\rho+2]} &= \ker(1 + J_0), \\ \widetilde{HS}(-Y)_{[\rho+1]} &= \mathbb{H}'[-1], \text{ and} \\ \widetilde{HS}(-Y)_{[\rho]} &= \text{coker}(1 + J_0), \end{aligned}$$

where V acts on all groups on the right-hand side by U^2 . The action of Q from the $[\rho + 2]$ - to the

$[\rho + 1]$ -submodule may be identified with the quotient map

$$\ker(1 + J_0) \rightarrow \mathbb{H}' = \ker(1 + J_0) / \text{im}(1 + J_0),$$

followed with multiplication by U in $\mathbb{H}'$. The action of Q from the $[\rho + 1]$ - to the $[\rho]$ -submodule may be identified with the inclusion map

$$\mathbb{H}' = \ker(1 + J_0) / \text{im}(1 + J_0) \rightarrow \text{coker}(1 + J_0).$$

Multiplication by Q is zero in all other cases.

Proof. Let q be a fixed grading of the monopole Floer homology $\widetilde{HS}(-Y)$. We begin by determining the $\text{Pin}(2)$ -homology in grading q when $q = \rho$ modulo 2. Consider the splitting of the lattice cohomology in grading q given by

$$\text{span}\{E_i\}_{i=1}^n \oplus \text{span}\{(1 + J_0)E_i\}_{i=1}^n \oplus F,$$

where F is at most one-dimensional. Substituting the lattice homology into the Gysin sequence using Theorem 7.1.1, it follows from Theorem 7.2.1 that the composition $\iota_* \circ \pi_*$ maps the first summand above isomorphically onto the second. Let $E' = \pi_*(\text{span}\{E_i\})$. Then E and E' fit into the subcomplex

$$0 \rightarrow E' \xrightarrow{\iota_*} E = \text{span}\{(1 + J_0)E_i\} \oplus \text{span}\{E_i\} \xrightarrow{\pi_*} E' \rightarrow 0. \quad (7.2)$$

Here, note that $\text{span}\{E_i\}$ maps isomorphically onto E' via π_* , and E' maps isomorphically onto the subspace $\text{span}\{(1 + J_0)E_i\}$ via ι_* . This subcomplex is of course simply n copies of the exact sequence

$$0 \rightarrow \mathbb{F} \rightarrow \mathbb{F} \oplus \mathbb{F} \rightarrow \mathbb{F} \rightarrow 0. \quad (7.3)$$

Note that although the choice of $\{E_i\}$ is not canonical, the subspace E' is canonically defined, since $\pi_*(E) = \pi_*(\text{span}\{E_i\})$.

Now let us consider F , which we assume for the moment to be non-zero in grading q . The fact that $(\iota_* \circ \pi_*)F = 0$ does not immediately allow us to determine how F fits into the Gysin sequence. Instead, there are two cases. First, suppose that $\pi_*F \neq 0$. Since $(\iota_* \circ \pi_*)F = 0$, this implies that π_*F is in the image of Q . Denote π_*F by F' . Since the monopole Floer homology of $-Y$ is zero in grading $q + 1$, a simple diagram chase shows that F and F' lie in some subcomplex

$$\begin{array}{ccccc} \mathbb{F} & & * & & \mathbb{F} \\ & \searrow \cdot Q & & \nearrow & \\ \mathbb{F} & & 0 & & \mathbb{F} \\ & \searrow \cdot Q & & \nearrow & \\ F' & & F & \xrightarrow{\pi_*} & F' \end{array} \quad (7.4)$$

On the other hand, if π_*F is zero, then ι_* surjects onto F . Choose any preimage F' of F under ι_* .

Since F does not lie in the image of $1 + J_0$, we see that F' cannot be in the image of π_* . Thus the action of Q on F' is non-zero, and we have the subcomplex

$$\begin{array}{ccccc}
 F' & \xrightarrow{\iota_*} & F & & F' \\
 & & \cdot Q & \swarrow & \\
 \mathbb{F} & & 0 & & \mathbb{F} \\
 & & \cdot Q & \swarrow & \\
 \mathbb{F} & & * & & \mathbb{F}
 \end{array} \tag{7.5}$$

Note that in this case F' is not canonically determined, but must be chosen as a preimage of F .

We now claim that the $\text{Pin}(2)$ -homology in grading q is precisely equal to $E' \oplus F'$ in either of the two cases described above. Indeed, suppose we had an element x of the $\text{Pin}(2)$ -homology in grading q lying outside of $E' \oplus F'$. We claim that without loss of generality we may assume $\iota_* x = 0$. Indeed, suppose not. Since $(\iota_* \circ \pi_*)(\iota_* x) = 0$, we see that $\iota_* x$ must lie in $\ker(1 + J_0) = \text{span}\{(1 + J_0)E_i\} \oplus F$. By subtracting off elements of E' from x , we may thus assume that $\iota_* x$ lies in F . If $\iota_* x$ is still non-zero, then $\pi_* F$ must be zero, and we are in the second case above where $F = \iota_* F'$. Subtracting off the non-zero element of F' from x , we obtain an element lying outside of $E' \oplus F'$ such that $\iota_* x = 0$.

Since the image of π_* certainly lies in $E' \oplus F'$, we know that x is not in the image of π_* . Thus $Qx \neq 0$. Putting everything together, we thus have that x lies in the subcomplex

$$\begin{array}{ccccc}
 \mathbb{F} & & 0 & & \mathbb{F} \\
 & & \cdot Q & \swarrow & \\
 x & & * & & x \\
 & & \cdot Q & \swarrow & \\
 \mathbb{F} & & 0 & & \mathbb{F}
 \end{array}$$

But this contradicts the fact that $Q^3 = 0$. Hence the Gysin sequence in grading q is the direct sum of copies of (7.2) and either (7.4) or (7.5). A similar diagram chase, together with the fact that the monopole Floer homology of $-Y$ is supported only in even dimensions, shows that there can be no other elements of the $\text{Pin}(2)$ -homology in gradings $q + 1$ or $q - 1$. Since F (as an abelian group) is isomorphic to a single U -tower, it follows that the entire Gysin sequence must be isomorphic to the

direct sum of copies of (7.3) (all in even gradings) and a tower of repeated copies of

$$\begin{array}{ccccc}
 \mathbb{F} & \xrightarrow{\iota_*} & \mathbb{F} & & \mathbb{F} \\
 & & \cdot Q & \nearrow & \\
 \mathbb{F} & \longleftarrow & 0 & & \mathbb{F} \\
 & & \cdot Q & \nearrow & \\
 \mathbb{F} & \longleftarrow & \mathbb{F} & \xrightarrow{\pi_*} & \mathbb{F}
 \end{array}$$

stacked on top of each other with a grading shift of four. The lowest line of this tower has grading ρ .

This determines the $\text{Pin}(2)$ -homology as an abelian group, and also specifies the Q -action. In order to determine the V -action, we must be a bit more circumspect as to the precise nature of the maps ι_* and π_* . Consider the four $\mathbb{F}[V]$ -submodules in the statement of the theorem. In gradings $q = \rho + 3 \bmod 4$, we see that the $\text{Pin}(2)$ -homology is indeed identically zero. In gradings $q = \rho + 2 \bmod 4$, the map ι_* is an $\mathbb{F}[V]$ -module isomorphism from the $\text{Pin}(2)$ -homology onto the subspace $\ker(1 + J_0) = \text{span}\{(1 + J_0)E_i\} \oplus F$, which establishes the second equality. In gradings $q = \rho \bmod 4$, the map π_* surjects onto the $\text{Pin}(2)$ -homology with kernel $\text{im}(1 + J_0) = \text{span}\{(1 + J_0)E_i\}$, proving the fourth equality.

It remains to establish the third equality and express the Q -action in terms of the claimed isomorphisms. We begin by identifying the a -tower inside $\text{coker}(1 + J_0)$. This may be done simply by selecting an element in $\text{coker}(1 + J_0)$ of appropriately high grading and understanding the multiplication by U . It is then not difficult to see that the a -tower may be identified with (the appropriate $\mathbb{F}[V]$ -submodule of) $\mathbb{H}' \subseteq \text{coker}(1 + J_0)$. The above decomposition of the Gysin sequence shows that multiplication by Q is an isomorphism from the $[\rho + 1]$ -submodule onto this tower; hence the $[\rho + 1]$ -submodule is a single V -tower which we may identify with $\mathbb{H}'[-1]$. This proves the third equality and gives the Q -action from the $[\rho + 1]$ - to the $[\rho]$ -submodule. Finally, the Q -action from the $[\rho + 2]$ - to the $[\rho + 1]$ -submodule is injective on $F' \cong F \subseteq \ker(1 + J_0)$ and zero otherwise, which is the description of the Q -action given in the statement of the theorem. $\square$

Example 7.2.3. On the left in Figure 7.1, we have displayed the (inverted) graded root corresponding to the lattice cohomology of $-\Sigma(2, 7, 15)$. In this example, the grading of the lowermost vertex is $d(Y) = 0$, and the grading of the lowermost J_0 -invariant vertex is $\rho = 2\bar{\mu}(Y) = 4$. In the center, we have displayed (inverted) graded roots corresponding to $\text{coker}(1 + J_0)$ and $\ker(1 + J_0)$. These may be constructed as in Section 4.3, although the reader should take care to observe that in the lattice

cohomology case, we delete branches coming off the stem of $\text{coker}(1 + J_0)$, rather than $\ker(1 + J_0)$.

On the right, the corresponding $\text{Pin}(2)$ -equivariant monopole Floer homology is displayed. The leftmost summand coincides with the $[\rho]$ submodule of $\text{coker}(1 + J_0)$. The center summand is simply a copy of the infinite stem (or rather, the appropriate $\mathbb{F}[V]$ -submodule of the infinite stem). Finally, the rightmost summand coincides with the $[\rho + 2]$ submodule of $\ker(1 + J_0)$. Here, $a = 4$, $b = 5$, and $c = 2$. Note that the a -tower does not extend “all the way down” – it stops at grading $a = 4$, even though there is an element below it in the same mod 4 grading. Explicitly,

$$\widetilde{HS}(-\Sigma(2, 7, 15)) = (\mathcal{V}_4^+ \oplus \mathcal{V}_5^+ \oplus \mathcal{V}_2^+) \oplus \mathbb{F}_{(0)} \oplus \mathbb{F}_{(2)} \oplus \mathbb{F}_{(6)}.$$




Figure 7.1: Lattice cohomology (left) and $\text{Pin}(2)$ -equivariant monopole Floer homology (right) of $-\Sigma(2, 7, 15)$. Multiplication by Q is represented by dashed lines. The rightmost summand has been drawn in such way so as to emphasize the fact that it is a mod 4 submodule of $\ker(1 + J_0)$. Note the deletion of vertices on the right-hand side.

Theorem 7.2.2 easily allows us to establish Theorem 1.3.8:

Proof. Theorem 7.2.2 identifies the b -tower, and shows that multiplication by Q is an isomorphism from the b -tower onto the a -tower. It follows that $a = 2\alpha(-Y) = \rho$ and $b = 2\beta(-Y) + 1 = \rho + 1$. Moreover, an analysis of multiplication by U on $\ker(1 + J_0)$ shows that the c -tower extends as far down as possible. It easily follows that

$$c = 2\gamma(-Y) + 2 = \begin{cases} d(-Y) + 2 & \text{if } d(-Y) = \rho \pmod{4} \\ d(-Y) & \text{if } d(-Y) = \rho + 2 \pmod{4}. \end{cases}$$

The claim then follows from the fact that $\alpha(-Y) = -\gamma(Y)$ and $\beta(-Y) = -\beta(Y)$, together with our previous assertion that $\rho = 2\bar{\mu}(Y)$. $\square$

Chapter 8

Generalizations and Further Work

In this chapter, we briefly describe several generalizations of the results presented in this thesis. We also discuss the connection between the ideas we have seen so far and recent joint work with Jennifer Hom, Linh Truong, and Matthew Stoffregen regarding an infinite-rank summand of $\Theta_{\mathbb{Z}}^3$ [DHST19].

8.1 Lattice homology for other plumbings

Although our discussion in this paper has been in terms of Seifert fibered homology spheres, it turns out that we can use the same arguments to study the involutive Heegaard Floer homology of a larger class of plumbed 3-manifolds. To this end, we state a slightly more general version of Theorem 2.4.4:

Theorem 8.1.1 (Theorem 1.2 [OS03b]; Theorem 5.2.2 [Ném08]). Let Y be an almost-rational plumbed 3-manifold, and let $\mathfrak{s}$ be a spin^c -structure on Y . Let k be a characteristic vector limiting to $\mathfrak{s}$ on Y . Then we have an isomorphism of $\mathbb{F}[U]$ -modules:

$$HF^-(Y, \mathfrak{s}) \cong \mathbb{H}^-(\Gamma, k)[-(|\Gamma| + k^2)/4 + 2].$$

Moreover, $\mathbb{H}^-(\Gamma, k)$ is supported only in dimensional grading zero.

This generalizes Theorem 2.4.4 in two directions. Firstly, Y is now only required to be a rational homology sphere. Secondly, we have introduced the notion of an almost-rational plumbed 3-manifold. We discuss both of these in turn.

Recall that in Section 2.1, we defined the involutive Heegaard Floer homology for Y any integer homology sphere. If Y is now a rational homology sphere, then we need the additional data of a self-conjugate spin^c -structure $\mathfrak{s}$ on Y . The involutive Floer package then assigns to the pair $(Y, \mathfrak{s})$ an ι -complex as before, and the definition of $CFI^-(Y, \mathfrak{s})$ and $HFI^-(Y, \mathfrak{s})$ proceeds verbatim. The

only difference is that the complexes in question are now $\mathbb{Q}$ -graded, rather than $\mathbb{Z}$ -graded, and the involutive Floer correction terms $\bar{d}(Y, \mathfrak{s})$ and $\underline{d}(Y, \mathfrak{s})$ likewise take values in $\mathbb{Q}$. A similar construction defines the $\text{Pin}(2)$ -equivariant monopole Floer homology for rational homology spheres.

On the lattice homology side, if Y is no longer an integer homology sphere, then the intersection form on L_Γ is no longer unimodular. In this case, we may identify the dual lattice $L'_\Gamma = \text{Hom}(L_\Gamma, \mathbb{Z})$ with $H_2(W, Y)$, and the exact sequence

$$0 \rightarrow H_2(W) \rightarrow H_2(W, Y) \rightarrow H_1(Y) \rightarrow 0$$

shows that we can view L_Γ as a sublattice of L'_Γ . The intersection form on L_Γ extends to a $\mathbb{Q}$ -valued intersection form on L'_Γ . The characteristic vector $k \in L_\Gamma$ from Section 2.4 should thus be replaced more generally by a characteristic vector in the dual lattice L'_Γ ; that is, an element $k \in L'_\Gamma$ such that $(k, x) = (x, x) \bmod 2$ for all $x \in L_\Gamma$. Note that L_Γ still acts on L'_Γ by sending $k \mapsto k + 2x$ for $x \in L_\Gamma$. As before, characteristic vectors k are identified with spin^c -structures on W , but now equivalence classes $[k]$ under the action of L_Γ are identified with spin^c -structures $\mathfrak{s}$ on Y . The lattice homology is then defined exactly as before (by using cubical subcomplexes of L_Γ), except that we obtain different answers depending on which equivalence class $[k]$ we use to construct the weight function. Under the isomorphism of Theorem 8.1.1, these correspond to the splitting of HF^- along different spin^c -structures on Y .

We now turn to the definition of an almost-rational plumbed 3-manifold. As usual, we assume that Γ is a negative-definite plumbing graph for Y . We define the *canonical characteristic element* to be the element $K \in \text{Hom}(L_\Gamma, \mathbb{Z})$ given by

$$K(v) = m(v) - 2$$

for all vertices v of Γ . For $x \in L_\Gamma$, we write $x > 0$ if all of the coefficients of x are non-negative and $x \neq 0$. We say that Γ is a *rational plumbing graph* if

$$-(K(x) + (x, x))/2 \geq 1$$

for all $x > 0$. This notion comes from the study of singularity theory; see for example [Ném08] for further discussion. We then have:

Definition 8.1.2. [Ném05, Definition 8.1] We say that Γ is an *almost-rational plumbing graph* if there exists a single vertex v_0 such that replacing the value of the decoration $m(v_0)$ with some $m'_0 \leq m(v_0)$ yields a rational plumbing graph. In this situation, we say that Y is an *almost-rational plumbed 3-manifold*, or an *AR plumbed manifold* (for short).

We will not discuss the topological significance of Definition 8.1.2 here; see [Ném08] for a description from the point of view of algebraic geometry. For our purposes, it is best to think of this family simply

as a class of rational homology spheres whose Heegaard Floer homology is given by the prescription of Theorem 8.1.1. (Note that the set of AR plumbed manifolds includes the set of Seifert fibered rational homology spheres.) See [Ném05, Section 8] for several examples of AR plumbed manifolds.

8.2 Generalization of results

Let Y be an AR plumbed manifold, and let $\mathfrak{s}$ be a self-conjugate spin^c -structure on Y . Consider the lattice homology defined using a characteristic vector k limiting to $\mathfrak{s}$ on Y . Note that since $\mathfrak{s}$ is self-conjugate, we have $[k] = [-k]$. It follows that k actually lies in the sublattice $L_\Gamma \subseteq L'_\Gamma$, and hence that reflection through the point $-k/2$ maps L_Γ onto itself. We thus again obtain an involution J_0 on the lattice homology $\mathbb{H}^-(\Gamma, k)$ which makes it into a symmetric graded root.¹ The same argument as in Theorem 4.1.1 then shows that J_0 coincides with ι_* under the isomorphism of Theorem 8.1.1.

It easily follows that our computation of the involutive Heegaard Floer homology (and associated ι -complex) holds equally well for this more general class of plumbed manifolds. Hence we may generalize Theorems 1.3.1 and 1.3.2 (and similarly Theorems 1.3.7 and 1.3.8) simply by replacing Y throughout with $(Y, \mathfrak{s})$, where Y is an AR plumbed manifold and $\mathfrak{s}$ is a self-conjugate spin^c -structure on Y . Similarly, if we let Θ_{AR} be the subgroup of $\Theta_{\mathbb{Z}}^3$ generated by AR plumbed (integer) homology spheres, then Theorems 1.3.3, 1.3.4, 1.3.5 and 1.3.6 generalize immediately by replacing Θ_{SF} with Θ_{AR} .

Remark 8.2.1. Note that our results do *not* translate directly into the setting of rational homology cobordism. This is because two rational homology spheres which are rationally homology cobordant do not necessarily have to be locally equivalent. As an example of this, note that $\Sigma(2, 3, 7)$ bounds a rational homology ball [FS84], but has nontrivial local equivalence class. (Indeed, the correction terms $\bar{d}$ and $\underline{d}$ are in general not rational homology cobordism invariants, unlike the d -invariant.) Instead, the cobordisms in question must be equipped with self-conjugate spin^c -structures. One can also instead study $\mathbb{Z}/2\mathbb{Z}$ -homology spheres with the relation of $\mathbb{Z}/2\mathbb{Z}$ -homology cobordism; this is briefly discussed in [HM17, HMZ18].

Finally, we note that our results evidently generalize to any situation in which the Heegaard Floer homology (together with the induced action of ι_*) is known to take the form of a symmetric graded root. For example, this is the case in the setting of [HHL18], where the authors study negative

¹This is now graded by a coset of $2\mathbb{Z}$ in $\mathbb{Q}$.

surgeries on L-space knots. In particular, Hendricks, Hom, and Lidman show that (-1) -surgeries on L-space knots are locally equivalent to type-one roots. Thus, one can of course consider an appropriate set of such manifolds in place of the Brieskorn spheres of Theorem 1.3.3. We leave other generalizations of this type to the reader.

8.3 Other work

We close this chapter with a brief discussion linking the results of this thesis to more recent joint work with Jennifer Hom, Linh Truong, and Matthew Stoffregen regarding an infinite-rank summand of $\Theta_{\mathbb{Z}}^3$ [DHST19]. The results of [DHST19] are based on several ideas coming from the work of Hom on knot Floer homology [Hom14], and a full treatment of [DHST19] is beyond the scope of this paper. However, we give a brief summary of [DHST19] from the point of view of our discussion so far.

Given Theorem 1.3.3, note that we obtain an infinite family of linearly independent homomorphisms from Θ_{SF} to $\mathbb{Z}$ simply by taking the coefficients of the basis decomposition into the Y_i . A slightly more intrinsic description of this can be given by using Theorem 1.3.6 and passing to the connected Floer homology. Indeed, from the procedure of Algorithm 6.1.2, we see that these homomorphisms simply count U -torsion towers of a given length appearing in $HF_{\text{conn}}(Y)$. Crucial in this count, however, is the notion that each U -torsion tower comes with a particular sign. In Theorem 1.3.6, this is determined from the overall structure of $HF_{\text{conn}}(Y)$, but in general there is no reason to assume that the connected Floer homology is of this form. Hence it is not entirely clear if these homomorphisms can be extended from Θ_{SF} to all of $\Theta_{\mathbb{Z}}^3$.

The approach of [DHST19] may thus be viewed as an attempt to define a signed count of torsion towers in general, without any foreknowledge of the shape of the connected complex. To do this, one considers a slightly weaker notion than local equivalence, called *(almost) local equivalence*, in which two actions of ι are considered to be equivalent if they are congruent modulo U .² This gives a modified local equivalence group $\widehat{\mathfrak{J}}$, together with a homomorphism from $\Theta_{\mathbb{Z}}^3$ factoring through $\mathfrak{J}$:

$$\widehat{h} : \Theta_{\mathbb{Z}}^3 \xrightarrow{h} \mathfrak{J} \rightarrow \widehat{\mathfrak{J}}.$$

It turns out that under this weaker notion of almost local equivalence, the set of possible equivalence classes in $\widehat{\mathfrak{J}}$ can be parameterized by a family of explicitly described model complexes. Using the structure of these, it is possible to define a signed count of U -torsion towers of any fixed length. The bulk of [DHST19] is then devoted to proving that these counts are well-defined, and are additive

²This is a rough discussion; see [DHST19] for precise details.

under tensor products. Once this is done, we obtain:

Theorem 8.3.1. [DHST19, Theorem 1.1] The homology cobordism group contains a direct summand isomorphic to $\mathbb{Z}^\infty$.

This can be viewed as a strengthening of Theorem 1.3.3, in the sense that the homomorphisms used in the proof of Theorem 8.3.1 cover the basis coefficient homomorphisms of Theorem 1.3.3 when restricted to Θ_{SF} .³

³Again, it is not known whether $\Theta_{\text{SF}} = \Theta_{\mathbb{Z}}^3$, or even if $h(\Theta_{\text{SF}}) = h(\Theta_{\mathbb{Z}}^3)$. If this were true, then of course Theorem 1.3.3 would provide an infinite family of linearly independent homomorphisms from $\Theta_{\mathbb{Z}}^3$ to $\mathbb{Z}$. The approach of [DHST19] sidesteps this question by developing some delicate algebraic machinery to understand abstract ι -complexes in general.

Bibliography

- [AD05] Selman Akbulut and Selahi Durusoy, *An involution acting nontrivially on Heegaard-Floer homology*, Geometry and topology of manifolds, Fields Inst. Commun., vol. 47, Amer. Math. Soc., Providence, RI, 2005, pp. 1–9.
- [AK79] Selman Akbulut and Robion Kirby, *Mazur manifolds*, Michigan Math. J. **26** (1979), no. 3, 259–284.
- [CH81] Andrew J. Casson and John L. Harer, *Some homology lens spaces which bound rational homology balls*, Pacific J. Math. **96** (1981), no. 1, 23–36.
- [CK14] Mahir B. Can and Çağrı Karakurt, *Calculating Heegaard-Floer homology by counting lattice points in tetrahedra*, Acta Math. Hungar. **144** (2014), no. 1, 43–75.
- [Dai18a] Irving Dai, *Connected Heegaard Floer homology of sums of Seifert fibrations*, 2018, preprint, arXiv:1804.06058.
- [Dai18b] ———, *On the $\text{Pin}(2)$ -equivariant monopole Floer homology of plumbed 3-manifolds*, Michigan Math. J. **67** (2018), no. 2, 423–447.
- [DHST19] Irving Dai, Jennifer Hom, Matthew Stoffregen, and Linh Truong, *An infinite-rank summand of the homology cobordism group*, 2019, preprint, arXiv:1810.06145.
- [DM17] Irving Dai and Ciprian Manolescu, *Involutive Heegaard Floer homology and plumbed three-manifolds*, 2017, preprint, arXiv:1704.02020.
- [DS17] Irving Dai and Matthew Stoffregen, *On homology cobordism and local equivalence between plumbed manifolds*, 2017, preprint, arXiv:1710.08055.
- [FFU01] Yoshihiro Fukumoto, Mikio Furuta, and Masaaki Ue, *W -invariants and Neumann-Siebenmann invariants for Seifert homology 3-spheres*, Topology Appl. **116** (2001), no. 3, 333–369.

- [FS84] Ronald Fintushel and Ronald J. Stern, *A μ -invariant one homology 3-sphere that bounds an orientable rational ball*, Four-manifold theory (Durham, N.H., 1982), Contemp. Math., vol. 35, Amer. Math. Soc., Providence, RI, 1984, pp. 265–268.
- [FS90] ———, *Instanton homology of Seifert fibred homology three spheres*, Proc. London Math. Soc. (3) **61** (1990), no. 1, 109–137.
- [Fur90] Mikio Furuta, *Homology cobordism group of homology 3-spheres*, Invent. Math. **100** (1990), no. 2, 339–355.
- [HHL18] Kristen Hendricks, Jennifer Hom, and Tye Lidman, *Applications of involutive Heegaard Floer homology*, 2018, preprint, arXiv:1802.02008.
- [HM17] Kristen Hendricks and Ciprian Manolescu, *Involutive Heegaard Floer homology*, Duke Math. J. **166** (2017), no. 7, 1211–1299.
- [HMZ18] Kristen Hendricks, Ciprian Manolescu, and Ian Zemke, *A connected sum formula for involutive Heegaard Floer homology*, Selecta Math. (N.S.) **24** (2018), no. 2, 1183–1245.
- [Hom14] Jennifer Hom, *The knot Floer complex and the smooth concordance group*, Comment. Math. Helv. **89** (2014), no. 3, 537–570.
- [JT12] András Juhász and Dylan Thurston, *Naturality and mapping class groups in Heegaard Floer homology*, 2012, preprint, arXiv:1210.4996.
- [Kar14] Çağrı Karakurt, *Contact structures on plumbed 3-manifolds*, Kyoto J. Math. **54** (2014), no. 2, 271–294.
- [KL15] Çağrı Karakurt and Tye Lidman, *Rank inequalities for the Heegaard Floer homology of Seifert homology spheres*, Trans. Amer. Math. Soc. **367** (2015), no. 10, 7291–7322.
- [KM07] Peter Kronheimer and Tomasz Mrowka, *Monopoles and three-manifolds*, New Mathematical Monographs, vol. 10, Cambridge University Press, Cambridge, 2007.
- [Lin17] Francesco Lin, *The surgery exact triangle in $\text{Pin}(2)$ -monopole Floer homology*, Algebr. Geom. Topol. **17** (2017), no. 5, 2915–2960.
- [Lin18a] ———, *A Morse-Bott approach to monopole Floer homology and the triangulation conjecture*, Mem. Amer. Math. Soc. **255** (2018), no. 1221, v+162.

- [Lin18b] ———, *Non-formality in $\text{Pin}(2)$ -monopole Floer homology*, 2018, preprint, arXiv:1809.01637.
- [Lis07] Paolo Lisca, *Lens spaces, rational balls and the ribbon conjecture*, *Geom. Topol.* **11** (2007), 429–472.
- [Man16] Ciprian Manolescu, *$\text{Pin}(2)$ -equivariant Seiberg-Witten Floer homology and the triangulation conjecture*, *J. Amer. Math. Soc.* **29** (2016), no. 1, 147–176.
- [Maz61] Barry Mazur, *A note on some contractible 4-manifolds*, *Ann. of Math. (2)* **73** (1961), 221–228.
- [Mye83] Robert Myers, *Homology cobordisms, link concordances, and hyperbolic 3-manifolds*, *Trans. Amer. Math. Soc.* **278** (1983), no. 1, 271–288.
- [Ném05] András Némethi, *On the Ozsváth-Szabó invariant of negative definite plumbed 3-manifolds*, *Geom. Topol.* **9** (2005), 991–1042.
- [Ném07] ———, *Graded roots and singularities*, *Singularities in geometry and topology*, World Sci. Publ., Hackensack, NJ, 2007, pp. 394–463.
- [Ném08] ———, *Lattice cohomology of normal surface singularities*, *Publ. Res. Inst. Math. Sci.* **44** (2008), no. 2, 507–543.
- [Neu80] Walter D. Neumann, *An invariant of plumbed homology spheres*, *Topology Symposium, Siegen 1979 (Proc. Sympos., Univ. Siegen, Siegen, 1979)*, *Lecture Notes in Math.*, vol. 788, Springer, Berlin, 1980, pp. 125–144.
- [OS03a] Peter Ozsváth and Zoltán Szabó, *Absolutely graded Floer homologies and intersection forms for four-manifolds with boundary*, *Adv. Math.* **173** (2003), no. 2, 179–261.
- [OS03b] ———, *On the Floer homology of plumbed three-manifolds*, *Geom. Topol.* **7** (2003), 185–224.
- [OS04a] ———, *Holomorphic disks and topological invariants for closed three-manifolds*, *Ann. of Math. (2)* **159** (2004), no. 3, 1027–1158.
- [OS04b] ———, *Holomorphic disks and three-manifold invariants: properties and applications*, *Ann. of Math. (2)* **159** (2004), no. 3, 1159–1245.

- [OS06] ———, *Holomorphic triangles and invariants for smooth four-manifolds*, Adv. Math. **202** (2006), no. 2, 326–400.
- [Sav02] Nikolai Saveliev, *Fukumoto-Furuta invariants of plumbed homology 3-spheres*, Pacific J. Math. **205** (2002), no. 2, 465–490.
- [Sie80] Laurence Siebenmann, *On vanishing of the Rohlin invariant and nonfinitely amphicheiral homology 3-spheres*, Topology Symposium, Siegen 1979 (Proc. Sympos., Univ. Siegen, Siegen, 1979), Lecture Notes in Math., Springer, Berlin, 1980, pp. 172–222.
- [Sto15] Matthew Stoffregen, *Pin(2)-equivariant Seiberg-Witten Floer homology of Seifert fibrations*, 2015, preprint, arXiv:1505.03234.
- [Sto17] ———, *Manolescu invariants of connected sums*, Proc. Lond. Math. Soc. (3) **115** (2017), no. 5, 1072–1117.
- [Twe13] Eamonn Tweedy, *Heegaard Floer homology and several families of Brieskorn spheres*, Topology Appl. **160** (2013), no. 4, 620–632.