

Connected Heegaard Floer homology of sums of Seifert fibrations

IRVING DAI

We compute the connected Heegaard Floer homology (defined by Hendricks, Hom, and Lidman) for a large class of 3–manifolds, including all linear combinations of Seifert fibered homology spheres. We show that for such manifolds, the connected Floer homology completely determines the local equivalence class of the associated ι –complex. Some identities relating the rank of the connected Floer homology to the Rokhlin invariant and the Neumann–Siebenmann invariant are also derived. Our computations are based on combinatorial models inspired by the work of Némethi on lattice homology.

57M27, 57R58

1 Introduction and results

Understanding the structure of the homology cobordism group $\Theta_{\mathbb{Z}}^3$ has been a motivating problem in the development of 3– and 4–manifold topology over the last several decades; see Fintushel and Stern [4], Furuta [6] and Frøyshov [5]. Most recently, Manolescu [13] introduced a $\text{Pin}(2)$ –equivariant version of Seiberg–Witten Floer homology and used this to refute the triangulation conjecture by ruling out 2–torsion in $\Theta_{\mathbb{Z}}^3$ with $\mu(Y) = 1$. Further variants and applications of Manolescu’s construction have been studied extensively by several authors; see eg Lin [12; 11; 10], Stoffregen [18; 19] and Dai [1].

Hendricks and Manolescu [8] defined an analogue of $\text{Pin}(2)$ –equivariant Seiberg–Witten Floer homology in the Heegaard Floer setting, called involutive Heegaard Floer homology. This associates to a rational homology sphere Y (with selfconjugate spin^c –structure \mathfrak{s}) an algebraic object called an ι –complex, whose mapping cone is a well-defined 3–manifold invariant $\text{HFI}^-(Y, \mathfrak{s})$.¹ Hendricks, Manolescu and Zemke [9] constructed a group $\mathfrak{I}_{\mathbb{Q}}$ consisting of all possible ι –complexes modulo an algebraic

¹Involutive Floer homology is defined for all three-manifolds, but in this paper we will only consider the case of rational homology spheres.

equivalence relation called local equivalence. This notion is modeled on the relation of homology cobordism, in the sense that if two (integer) homology spheres are homology cobordant, then their ι -complexes are locally equivalent. Hendricks–Manolescu–Zemke obtain a homomorphism

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

One can thus obstruct the existence of a homology cobordism between two homology spheres by computing their local equivalence classes in $\mathfrak{J}_{\mathbb{Q}}$ and showing that they are not equal.

Manolescu and the author [2] exhibited a \mathbb{Z}^{∞} -subgroup of $\Theta_{\mathbb{Z}}^3$ by establishing the linear independence of a certain infinite family of ι -complexes $\{X_i\}_{i=1}^{\infty}$.² (This follows a similar proof in the $\text{Pin}(2)$ -case due to Stoffregen [19]; see also Fintushel and Stern [4] and Furuta [6].) Let \mathfrak{X} be the subgroup of $\mathfrak{J}_{\mathbb{Q}}$ spanned by the X_i . Stoffregen and the author [3] showed that (up to grading shift) $h(Y) \in \mathfrak{X}$ for a large class of 3-manifolds Y , including any sum of Seifert fibered homology spheres. For such Y , understanding $h(Y)$ thus reduces to studying linear combinations of the basis elements X_i . The author is currently unaware of any example with $h(Y) \notin \mathfrak{X}$; the existence of such a manifold would prove that Seifert fibered spaces do not generate the homology cobordism group.³

Following a construction of Stoffregen [18] in the Seiberg–Witten case, Hendricks, Hom and Lidman [7] used involutive Heegaard Floer homology to define a new invariant of homology cobordism, called connected Heegaard Floer homology. This is an invariant of the local equivalence class of (Y, \mathfrak{s}) which associates to (Y, \mathfrak{s}) a particular (grading-shifted) summand $\text{HF}_{\text{conn}}(Y, \mathfrak{s})$ of the usual Heegaard Floer homology $\text{HF}^-(Y, \mathfrak{s})$. While a priori the connected Floer homology of Y is a strictly weaker invariant than its local equivalence class, it is considerably easier to state and understand, as the obstruction coming from HF_{conn} takes the form of a homology group in the usual sense.

The goal of this paper is to use the results of [3] to compute the connected Floer homology of any 3-manifold with local equivalence class in \mathfrak{X} . Our main result (Theorem 1.1) establishes the possible forms that HF_{conn} can take, and shows that (for such manifolds) HF_{conn} is sufficient to recover the entire local equivalence class. We

²These differ from the basis elements presented in [2, Theorem 1.7] by an overall grading shift. See the discussion following [3, Lemma 4.3].

³There are many algebraic examples of ι -complexes which are not locally equivalent to anything in \mathfrak{X} , but it is unclear whether these occur as the ι -complexes of actual 3-manifolds.

then rephrase several computations from [3] to obtain identities relating the rank of HF_{conn} to other homology cobordism invariants, such as the Rokhlin invariant and the Neumann–Siebenmann invariant (see Neumann [15] and Siebenmann [17]).

Finally, a brief word on the methods used in this paper. The bulk of our approach consists of developing combinatorial models for understanding the local equivalence classes of different ι -complexes. We think that these models, which we call “geometric complexes”, are interesting in their own right, as they allow various algebraic constructions which are otherwise opaque to be translated into topological operations on finite-dimensional cell complexes. (See also [2, Section 7] and [3, Section 2.4].) Our construction is heavily inspired by work of Némethi [14] on lattice homology (see Ozsváth and Szabó [16]). Although here we only use geometric complexes as a convenient combinatorial tool, we hope that this may be indicative of deeper connections to Floer theory as a whole.

1.1 Statement of results

We first establish some notation. Let Y be a rational homology sphere, and let \mathfrak{s} be a selfconjugate spin^c -structure on Y . We denote by h the map sending (Y, \mathfrak{s}) to the local equivalence class of its (grading-shifted) ι -complex:

$$(Y, \mathfrak{s}) \mapsto h(Y, \mathfrak{s}) = (\text{CF}^-(Y, \mathfrak{s}), \iota)[-2].$$

See Section 2.1 for details. In this paper, we will be concerned with 3-manifolds for which $h(Y, \mathfrak{s})$ decomposes (up to grading shift) as a linear combination of the basis complexes X_i . Since each X_i has d -invariant equal to zero, it is easily checked that in this situation the grading shift is given by $d(Y, \mathfrak{s})$, ie

$$h(Y, \mathfrak{s}) = \left(\sum_i c_i X_i \right) [-d(Y, \mathfrak{s})].$$

See Section 2.1 for a definition of the X_i and a discussion of which 3-manifolds are known to satisfy this decomposition. For now, the reader may take Y to be a sum of Seifert fibered rational homology spheres. We begin by defining an algorithm that takes an element of \mathbb{Z}^∞ (written as a linear combination of the X_i) and outputs a \mathbb{Z} -graded $\mathbb{F}[U]$ -module.

Let $\mathcal{T}_d(l)$ be the U -torsion tower $\mathbb{F}[U]/U^l$, shifted in grading so that its uppermost element has grading d . Note that $\mathcal{T}_d(l)$ has l elements; we refer to l as the *length* of $\mathcal{T}_d(l)$. We will additionally identify certain towers as *downwards pointing* and other

towers as *upwards pointing*. In the first case, we define the *head* of $\mathcal{T}_d(l)$ to be the element of maximal grading d , and the *tail* of $\mathcal{T}_d(l)$ to be the element of minimal grading $d - 2(l - 1)$. 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 $\mathcal{T}_d(l)$.

Now 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 canceling pairs $X_i - X_i$. We use the indices appearing in X to construct a \mathbb{Z} -graded $\mathbb{F}[U]$ -module, as follows. 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.

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:

- (1) 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.
- (2) 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.
- (3) 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 1.

We are now ready to state our main theorem:

Theorem 1.1 *Let Y be a rational homology sphere, and let \mathfrak{s} be a selfconjugate spin^c -structure on Y . Suppose that*

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




Figure 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. Maslov gradings are indicated on the left of the diagram.

where we assume that the above linear combination is maximally simplified with $i_1 \geq i_2 \geq \dots \geq i_n$. Then the connected Heegaard Floer homology of (Y, \mathfrak{s}) is given by

$$\text{HF}_{\text{conn}}(Y, \mathfrak{s}) = \mathcal{H}(X)[-d(Y, \mathfrak{s}) + 1].$$

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

We will prove Theorem 1.1 by establishing an explicit model for the connected complex (see Definition 2.7) of (Y, \mathfrak{s}) . Conversely, suppose that we have some $h(Y, \mathfrak{s}) \in \mathfrak{X}$ for which $\text{HF}_{\text{conn}}(Y, \mathfrak{s})$ is known. Then it is not hard to check that one can recover the decomposition of $h(Y, \mathfrak{s})$ into the X_i . Hence:

Theorem 1.2 Suppose that (up to grading shift) $h(Y, \mathfrak{s}) \in \mathfrak{X}$. Then the local equivalence class of $h(Y, \mathfrak{s})$ in $\mathfrak{J}_{\mathbb{Q}}$ is completely determined by the isomorphism class of $\text{HF}_{\text{conn}}(Y, \mathfrak{s})$, together with $d(Y, \mathfrak{s})$.

If Y_1 and Y_2 are two manifolds as above, whose connected Floer homologies and d -invariants are known, then Theorems 1.1 and 1.2 allow us to immediately compute

the connected Floer homology of $Y_1 \# Y_2$. To do this, we first apply Theorem 1.2 to write $h(Y_1)$ and $h(Y_2)$ as a linear combination of the X_i ; then we sum these together, remove all canceling pairs, and apply Theorem 1.1 to the result.

We also derive some identities involving the rank of $\text{HF}_{\text{conn}}(Y)$. For these applications, we specialize to the class of almost-rational (AR) plumbed 3-manifolds. This family (which includes all Seifert fibered rational homology spheres) is defined by placing certain combinatorial restrictions on the usual three-dimensional plumbing construction. See [14] (also [2, Definition 2.7]) for details.

Corollary 1.3 *Let Y be a linear combination of AR plumbed 3-manifolds, and let \mathfrak{s} be a selfconjugate spin^c -structure on Y . Then*

$$\bar{\mu}(Y, \mathfrak{s}) + \frac{1}{2}d(Y, \mathfrak{s}) = \text{signed rank HF}_{\text{conn}}(Y, \mathfrak{s}).$$

Here, $\bar{\mu}(Y, \mathfrak{s})$ is the Neumann–Siebenmann invariant of (Y, \mathfrak{s}) . 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.

Note that the orientations of the towers in $\text{HF}_{\text{conn}}(Y, \mathfrak{s})$ are not intrinsically defined, but must instead be deduced from the overall structure of $\text{HF}_{\text{conn}}(Y, \mathfrak{s})$. (See the proof of Theorem 1.2.) This can only be done when $\text{HF}_{\text{conn}}(Y, \mathfrak{s})$ takes the form prescribed by Theorem 1.1, which is consistent with the fact that $\bar{\mu}$ is not defined for all 3-manifolds.

Corollary 1.4 *Let Y be a linear combination of AR plumbed homology spheres. Then*

$$\mu(Y) + \frac{1}{2}d(Y) = \text{rank HF}_{\text{conn}}(Y) \pmod{2},$$

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

Finally, we end with the following question:

Question 1.5 *Do there exist (integer) homology spheres Y for which $\text{HF}_{\text{conn}}(Y)$ is not isomorphic to an $\mathbb{F}[U]$ -module of the form described in Theorem 1.1?*

An affirmative answer would produce an element of $\Theta_{\mathbb{Z}}^3$ not homology cobordant to any sum of Seifert fibered spaces. There are certainly algebraic examples of ι -complexes whose connected homology does not take the form of Theorem 1.1 (see for instance Example 6.2), but it is unclear whether these occur as the ι -complexes of actual 3-manifolds.

Organization of the paper In Section 2, we review the algebraic formalism underlying involutive Heegaard Floer homology and connected Floer homology. In Section 3, we discuss geometric complexes and outline several related constructions. Sections 4 and 5 are devoted to proving Theorem 1.1. Finally, in Section 6, we prove Theorem 1.2, as well as Corollaries 1.3 and 1.4.

Acknowledgements The author would like to thank his advisor, Zoltán Szabó, for his continued support and guidance. He would also like to thank Kristen Hendricks, Jen Hom, Tye Lidman, Ciprian Manolescu and Matt Stoffregen for helpful conversations. In addition, the author would like to thank the referee for several helpful comments.

The author was partially supported by NSF grant DGE-1148900.

2 Background

2.1 Involutive Floer homology

In this section, we review the involutive Heegaard Floer package defined by Hendricks and Manolescu [8]. For the present application, we will only need to understand the algebraic formalism of the output; we refer the interested reader to [8] instead for the topological details of the construction.

Definition 2.1 [9, Definition 8.1] An ι -complex is a pair (C, ι) consisting of:

- A \mathbb{Q} -graded, finitely generated, free chain complex C over the ring $\mathbb{F}[U]$, where $\deg(U) = -2$. Moreover, we ask that there is some $\tau \in \mathbb{Q}$ such that the complex C is supported in degrees differing from τ by integers. We also require that there is a relatively graded isomorphism

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

and that $U^{-1}H_*(C)$ is supported in degrees differing from τ by even integers.

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

Two ι -complexes are said to be *homotopy equivalent* if there are (grading-preserving) homotopy equivalences between them which are ι -equivariant up to homotopy.

Given a rational homology sphere Y with selfconjugate spin^c -structure \mathfrak{s} , Hendricks and Manolescu define a homotopy involution ι on the Heegaard Floer complex $\text{CF}^-(Y, \mathfrak{s})$ by using the involution on the Heegaard diagram interchanging the α - and β -curves. They then show that the map sending

$$(Y, \mathfrak{s}) \mapsto (\text{CF}^-(Y, \mathfrak{s}), \iota)$$

is well-defined up to homotopy equivalence of ι -complexes. The involutive Heegaard Floer homology of (Y, \mathfrak{s}) is constructed by taking the mapping cone of $\text{CF}^-(Y, \mathfrak{s})$ with respect to the map $1 + \iota$. In this paper, we will always work with the ι -complexes themselves, rather than their involutive homology.

Definition 2.2 For any positive integer i , let X_i be the chain complex generated over $\mathbb{F}[U]$ by α , $\iota\alpha$ and β , with

$$\partial\beta = U^i(\alpha + \iota\alpha)$$

and $\partial(\alpha) = \partial(\iota\alpha) = 0$. The Maslov gradings are given by $M(\alpha) = M(\iota\alpha) = 0$ and $M(\beta) = -2i + 1$. The action of ι interchanges α and $\iota\alpha$ and sends β to itself. Note that here ι is an actual involution (rather than a homotopy involution); this will be true for all complexes considered in this paper.

In order to study homology cobordism, we will need the following weaker notion of equivalence between ι -complexes:

Definition 2.3 Two ι -complexes (C, ι) and (C', ι') are called *locally equivalent* if there exist (grading-preserving) chain maps

$$f: C \rightarrow C', \quad g: C' \rightarrow C$$

between them such that

$$f \circ \iota \simeq \iota' \circ f, \quad g \circ \iota' \simeq \iota \circ g,$$

and f and g induce isomorphisms on homology after inverting the action of U . This is strictly weaker than the notion of homotopy equivalence between ι -complexes.

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. In [9], it is shown that if Y_1 and Y_2 are homology cobordant, then their respective ι -complexes are locally equivalent.

In [9, Section 8], it is shown that the set of ι -complexes up to local equivalence forms a group, with the group operation being given by tensor product. We call this group the *involutive Floer group* and denote it by $\mathfrak{I}_{\mathbb{Q}}$.

Definition 2.4 [9, Proposition 8.8] Let $\mathfrak{J}_{\mathbb{Q}}$ 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 $\mathfrak{J}_{\mathbb{Q}}$ is given by the trivial complex \mathcal{I}_0 consisting of a single $\mathbb{F}[U]$ -tower starting in grading zero, together with the identity map on this complex. Inverses in $\mathfrak{J}_{\mathbb{Q}}$ are given by dualizing. There is an obvious subgroup of $\mathfrak{J}_{\mathbb{Q}}$ generated by the set of ι -complexes which are \mathbb{Z} -graded; we denote this by \mathfrak{J} . Clearly, $\mathfrak{J}_{\mathbb{Q}}$ consists of an infinite number of copies of \mathfrak{J} , one for each $[\tau] \in \mathbb{Q}/2\mathbb{Z}$.

Let h be the map sending a pair (Y, \mathfrak{s}) to the local equivalence class of its (grading-shifted) ι -complex,

$$h(Y, \mathfrak{s}) = (\text{CF}^-(Y, \mathfrak{s}), \iota)[-2].$$

The main result of [3] is then:⁴

Theorem 2.5 [3, Theorem 1.2] *Let Y be a linear combination of almost-rational plumbed 3-manifolds, and let \mathfrak{s} be a selfconjugate spin^c -structure on Y . Then $h(Y, \mathfrak{s})$ decomposes uniquely as*

$$h(Y, \mathfrak{s}) = \left(\sum_i c_i X_i \right) [-d(Y, \mathfrak{s})].$$

Here, we define the family of almost-rational (AR) plumbed 3-manifolds by placing certain combinatorial restrictions on the usual three-dimensional plumbing construction. See [14] (also [2, Definition 2.7]) for a precise discussion. Note that the class of AR plumbed 3-manifolds includes all Seifert fibered rational homology spheres.

There are other classes of manifolds for which the conclusion of Theorem 2.5 holds, most notably linear combinations of manifolds whose local equivalence classes are already given by the X_i . This is true (in particular) for negative surgeries on L -space knots; see the proof of [7, Theorem 1.8].

2.2 Connected Floer homology

We now turn to the definition of connected Floer homology [7]. We begin with:

⁴This is not exactly the statement of [3, Theorem 1.2]; the basis elements presented in Theorem 2.5 differ from the ones in [2] and [3] by an overall grading shift.

Definition 2.6 [7, Definition 3.1] Let (C, ι) be an ι -complex. A *selflocal equivalence* is a local map from (C, ι) to itself, that is, a (grading-preserving) chain map

$$f: C \rightarrow C \quad \text{such that} \quad f \circ \iota \simeq \iota \circ f$$

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

Hendricks, Hom and Lidman define a preorder on the set of selflocal equivalences of (C, ι) by declaring $f \lesssim g$ whenever $\ker f \subseteq \ker g$. A selflocal equivalence f is then said to be *maximal* if $g \gtrsim f$ implies $g \lesssim f$ for any selflocal equivalence g .

Maximal selflocal equivalences should heuristically be thought of as producing a local equivalence between C and some very small subcomplex of C given by $\text{im } f$. (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 $\text{im } f$, and similarly for the relevant homotopy maps. However, according to [7, Lemma 3.7], we can modify ι to produce an actual homotopy involution ι_f on $\text{im } f$ such that (C, ι) and $(\text{im } f, \iota_f)$ are locally equivalent.

Hendricks–Hom–Lidman further show that a maximal selflocal equivalence always exists [7, Lemma 3.3], and that any two maximal selflocal equivalences give homotopy equivalent ι -complexes $(\text{im } f, \iota_f)$ and $(\text{im } g, \iota_g)$ [7, Lemma 3.8]. We can thus make the unambiguous definition:

Definition 2.7 [7, 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 $(\text{im } f, \iota_f)$, where f is any maximal selflocal 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) [7, Proposition 3.10]. Indeed, the connected complex should be thought of 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 2.8 [7, 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 an $\mathbb{F}[U]$ -complex.

The *connected Heegaard Floer homology of* (Y, \mathfrak{s}) is then of course defined to be the connected homology of $(CF^-(Y, \mathfrak{s}), \iota)$. The grading shift by one is enacted so that HF_{conn} is a summand of HF_{red} , viewed as a quotient of HF^+ .

Example 2.9 The homology of X_i is given by the direct sum of a single nontorsion tower (starting in grading zero) and a U -torsion tower $\mathcal{T}_0(i)$ of length i . It is not difficult to show that the connected complex of X_i is all of X_i , and hence that the connected homology of X_i is precisely $\mathcal{T}_0(i)$ (shifted up in grading by one). See [7, Lemma 7.2].

In [7, Proposition 4.1], it is shown that the connected Floer homology behaves well under dualization. Note, however, that the connected homology does *not* behave well under tensor product. 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.

In this paper, we will not actually use any properties of the connected Floer homology beyond the fact that it is a (grading-shifted) summand of the usual Floer homology and that it dualizes well. Instead, in light of Theorem 2.5, we will focus on producing “simple” representatives of linear combinations of the X_i . While each of the X_i is individually very straightforward, we will see that it is actually quite subtle to construct good local representatives of their tensor products.

3 Geometric operations

In this section, we review the definition of a geometric complex and outline several related constructions which we will use to describe the ι -complexes of sums of Seifert fibered spaces. See [2, Section 7] and [3, Section 2.4] for background on geometric complexes and associated applications.

3.1 Geometric complexes

Let C_{cell} be a finite cell complex. Suppose that we are given a function gr from the cells of C_{cell} to a coset of $2\mathbb{Z}$ in \mathbb{Q} , 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 (relatively) \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

$$\partial e_i = \sum_{e_{i-1} \in \text{cellular boundary of } e_i} U^{\frac{1}{2}(\text{gr}(e_{i-1}) - \text{gr}(e_i))} e_{i-1}.$$

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 nonnegative 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*.

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 geometric complexes will come equipped with a cellular chain map J on C_{cell} with $J^2 = 1$. Geometrically, this corresponds 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). Moreover, in all of the cases we will be discussing, the action of J will have 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*.

Example 3.1 The prototypical example of a (split) geometric complex is the complex X_i of Definition 2.2. The skeleton of X_i is given by a (rather trivial) one-dimensional cell complex, consisting of two 0-cells corresponding to α and $J\alpha$, and a single 1-cell corresponding to β . The action of J interchanges α and $J\alpha$ and sends β to itself. This complex, together with its associated grading function gr , is displayed on the right in Figure 2. See [3, Section 2.4] for further discussion.




Figure 2: The complex X_i (left) and its associated skeleton (right). On the left, solid lines represent multiplication by U and dashed lines represent the action of ∂ .

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. See [3, Lemma 3.3].

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 . If C_1 and C_2 are split, then $C_1 \otimes C_2$ is also split, with the involution on $C_1 \otimes C_2$ acting coordinatewise. See [3, Section 2.4].

We now give two important families of geometric complexes.

Example 3.2 Consider an alternating sum of X_i , ie

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

with $i_1 > i_2 > \dots > i_n$. By results of [3], this is locally equivalent to a one-dimensional geometric complex. (According to [3, Theorem 4.2], the above sum is locally equivalent




Figure 3: Geometric complexes of alternating sums with n even (top) and n odd (bottom).

to a monotone root, which in turn can be described by a one-dimensional geometric complex as constructed in [2, Section 7].) There are two possible forms for this complex, both of which are displayed in Figure 3. 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). The grading function gr is given by

$$\begin{aligned} \text{gr}(v_k) &= \text{gr}(Jv_k) = 2(-i_1 + i_2 - \dots + i_k), \\ \text{gr}(w_k) &= \text{gr}(Jw_k) = 2(-i_1 + i_2 - \dots - i_k), \end{aligned}$$

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 1.1. (To see this, it is easier to think of the complex as a monotone root, as in [2, Section 6].) The result is schematically depicted in Figure 4. By [7, Proposition 7.5], this is indeed equal (after grading shift) to the connected homology of $X_{i_1} - X_{i_2} + \dots \pm X_{i_n}$.

Example 3.3 Consider a sum of X_i with the same orientation, ie

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

with $i_1 \geq i_2 \geq \dots \geq i_n$. By [3, Lemma 5.2], this is locally equivalent to a geometric complex whose skeleton is a J -symmetric decomposition of the n -ball. This is displayed in Figure 5. For each $0 \leq k < n$, there are two cells of dimension k , which we have labeled e_k and Je_k , while there is a single J -symmetric cell e_n of dimension n . The grading function gr is given by

$$\text{gr}(e_k) = \text{gr}(Je_k) = 2(-i_1 - i_2 - \dots - i_k) \quad \text{with} \quad \text{gr}(e_0) = \text{gr}(Je_0) = 0.$$

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 1.1. (Generators of the




Figure 4: 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 $X_5 - X_4 + X_2$.




Figure 5: Geometric complex of same-sign sums.

torsion towers are given by the obvious cycles $e_k + Je_k$.) The result is schematically depicted in Figure 6. By [7, Proposition 7.1], this is indeed equal (after grading shift) to the connected homology of $X_{i_1} + X_{i_2} + \dots + X_{i_n}$.

Examples 3.2 and 3.3 give geometric representatives for the local equivalence classes of alternating sums and same-sign sums of 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 for the remainder of this paper will be to produce similar representatives for arbitrary linear combinations of the X_i .

Before we continue, we will need to make one more definition involving geometric complexes. Let C be a geometric complex with skeleton C_{cell} and grading function gr .




Figure 6: Homology of same-sign sums. Dotted lines are separated by Maslov grading two, with the uppermost line having Maslov grading zero. This example is $X_4 + X_3 + X_2$.

For each cell e_i of C_{cell} , we define

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

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, ie

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

Note that this means any term appearing in the boundary map ∂ on C must appear with a U -power of at least $\frac{1}{2} \text{width}(C)$.

3.2 Doubling 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 4.1 for the precise details.

Consider the cell complexes displayed in Figure 7. On the left, we have displayed the skeleta of several split complexes, while on the right, we have displayed a corresponding




Figure 7: Doublings of split complexes.

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.

Let X be any split complex, and let Δ be any nonnegative 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
- (3) $\text{gr}(\theta) = \text{gr}(\eta) - 2\Delta$.

In Section 4.1 we will show that $X^d(\Delta)$ is a well-defined split geometric complex (ie 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 3.2, 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 3.3, 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 7.

The main technical result of this paper will be to show that the doubled complex $X^d(\Delta)$ is locally equivalent to $X \otimes X_\Delta$ in general. This will allow 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 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.

We should stress here that for very simple complexes, it is straightforward to describe the cellular boundary map of the doubled complex in terms of the original. For example, the double of a spherical complex is a spherical complex of one higher dimension. However, sometimes the relation between the two cellular boundary maps is slightly more subtle. For instance, in the fourth example of Figure 7, we must modify the boundary map on the 2-cells incident to our original J -invariant 0-cell 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 8 shows a split complex with two putative doublings. In Section 4.1 we will provide a more precise definition of the doubling operation by explicitly defining the doubled complex and its boundary map; this will turn out to depend on an additional set of choices. However, we will in fact show that *any* choice of doubling is locally equivalent to $X \otimes X_\Delta$, so for our purposes this ambiguity is unimportant.

3.3 Halving operations

Let us briefly review the procedure for dualizing ι -complexes given in [9, Section 8.3]. If $\{x_i\}$ is a homogenous $\mathbb{F}[U]$ -basis for X , then the dual complex X^* has a corresponding $\mathbb{F}[U]$ -basis given by $\{x_i^*\}$, where the Maslov grading of x_i^* is $M(x_i^*) = -M(x_i)$. The differential on X^* is defined so that $U^k x_b^*$ appears in $\partial^* x_a^*$ if and only if $U^k x_a$ appears in ∂x_b . Similarly, an element $U^k x_b^*$ appears in $\iota^* x_a^*$ if and only if $U^k x_a$




Figure 8: A split complex which appears to have two possible doublings.

appears in ix_b . Note that the action of U on X^* still has degree -2 , and the differential still drops the Maslov grading by one.

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

$$\partial^*(\alpha^*) = \partial^*(J\alpha^*) = U^i \beta^* \quad \text{and} \quad \partial^* \beta^* = 0.$$

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 9. 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 (finite) \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 4.1, there will be no trouble in defining the doubling operation for these complexes also. If X has skeleton C_{cell} with generators e_i , 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 . This is equipped with the grading function $\text{gr}(e_i^*) = -\text{gr}(e_i)$, and the dimensional grading is defined so that e_i^* has dimensional grading $-i$. The usual construction then exhibits X^* as a geometric complex with skeleton C_{cell}^* .




Figure 9: The skeleton of X_i^* (left), duals of Example 3.2 (middle) and dual of Example 3.3 (right). Here, crosses or dashed lines indicate the “relative” part of the cell complex. For example, on the right-hand side, we have $\partial^* e_0^* = e_1^* + J e_1^*$.

To reconcile this with the geometric picture, observe that if we shift the dimensional grading up by n while shifting gr down by n , the Maslov grading is preserved. Hence enacting complementary grading shifts does not change the complex X . We could thus have equally well defined $gr(e_i^*) = -gr(e_i) - n$ and set the dimensional grading of e_i^* to be $n - i$, where n is the dimension of the original skeleton C_{cell} . If C_{cell} is an actual cell complex, then the complex C_{cell}^* with this (now nonnegative) dimensional grading may be represented by the dual cell complex of C_{cell} in the geometric sense, as in Figure 9. Dualizing Examples 3.2 and 3.3 yields geometric complexes whose skeleta are given by the relative cell complexes of Figure 9. See [3, Remark 5.3].

The dual of a geometric complex is thus also a geometric complex, with the same width and the same number of generators. If the original complex is split, then its dual is split also.

While not absolutely necessary for the results stated in Section 1, it is perhaps useful for us to describe the operation dual to doubling. This corresponds to tensoring with X_Δ^* (rather than X_Δ). Consider the cell complexes displayed in Figure 10. 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.

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 11. 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




Figure 10: Halvings of split complexes.

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 “halve” the J -invariant cell, which is now one-dimensional. (We simply shift the dimensional grading up by one and decrease gr by one, as described previously.) See Figure 11.




Figure 11: 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 9.

Now let X be any split complex, and Δ any nonnegative 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) $\text{gr}(\omega^*) = \text{gr}(J\omega^*) = \text{gr}(\eta^*)$.
- (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 3.2 and (the dual of) Example 3.3.

4 Split complexes and doubling

In this section, we give a precise construction of the doubling operation. We then prove that for any split ι -complex X , the doubled complex $X^d(\Delta)$ is locally equivalent to $X \otimes X_\Delta$.

4.1 Doubling operations, revisited

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 C . Then C_{cell} (as a module over \mathbb{F}) decomposes as a direct sum

$$C_{\text{cell}} = C \oplus JC \oplus \eta,$$

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 C .

It will be useful to also consider the related splitting $C \oplus (1 + J)C \oplus \eta$, in which case we can write the differential ∂x of any $x \in C$ uniquely as

$$a + (1 + J)b \quad \text{or} \quad a + (1 + J)b + \eta$$

for $a, b \in C$. We will sometimes use parentheses to denote the possibility of adding η , ie $\partial x = a + (1 + J)b (+\eta)$. Similarly, the fact that J commutes with ∂ and that $J\eta = \eta$ implies that

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

for some $\zeta \in C$. Note that a, b and ζ are allowed to be sums of cells and/or zero. We will continue to use the notation $\partial x = a + (1 + J)b (+\eta)$ throughout the rest of this section.

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

Definition 4.1 Let X be a split complex with a particular choice of splitting C , and let Δ be a nonnegative 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

$$(C \oplus \omega) \oplus (JC \oplus J\omega) \oplus \theta.$$

Here, C and JC 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 $(C \oplus \omega)$.

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

Definition 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 $x \in C$. 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) $\partial^d \omega = \partial \eta = (1 + J)\zeta$.

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

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

The three cases for $\partial^d x$ above correspond to different situations as in Figure 7. 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. (We will sometimes write this as $\partial^d x = \partial x$.) 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 7.) 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 7.)

We now verify that ∂^d is a differential (squares to zero). The reader who already believes that the geometric doubling operation is well-defined (and given by the above algebraic construction) may skip this verification, but should be cautioned that the casework below will turn out to be useful in later sections.

Lemma 4.3 *The differential ∂^d turns $X^d(\Delta)$ into a chain complex, ie it squares to zero.*

Proof We proceed using copious amounts of casework:

(1) Suppose $x \in C$. We first divide into two cases based on whether or not ∂x contains η :

(1a) Suppose ∂x does not contain η , so that $\partial x = a + (1 + J)b$. In order to understand $(\partial^d)^2x$, let us consider the possible forms of ∂a and ∂b . Write $\partial a = p + (1 + J)q (+\eta)$ and $\partial b = r + (1 + J)s (+\eta)$, with $p, q, r, s \in C$. Here, we have written $(+\eta)$ to keep open the possibility of ∂a and/or ∂b containing η . Since we know $\partial^2x = 0$, however, we immediately see that ∂a cannot contain η and that $p = 0$ and $q = r$. Thus we have

$$\partial a = (1 + J)r \quad \text{and} \quad \partial b = r + (1 + J)s (+\eta).$$

Now, the fact that $\partial^2b = 0$ implies that ∂r cannot contain η . Hence (referring to the definition of ∂^d) we have $\partial^d a = \partial a$. We further subdivide into three remaining subcases:

(i) Suppose ∂b does not contain η , and ∂s does not contain η . Then $\partial^d x = \partial x$ and $\partial^d b = \partial b$. In this case we obviously have $(\partial^d)^2x = \partial^2x = 0$.

(ii) Suppose ∂b does not contain η , but ∂s does contain η . Then $\partial^d b = \partial b + \theta$, and

$$(\partial^d)^2x = \partial a + (1 + J)(\partial b + \theta) = \partial^2x + (1 + J)\theta,$$

which is zero, since $J\theta = \theta$.

(iii) Suppose ∂b does contain η . Then we have

$$\partial^d x = a + (1 + J)b + \theta \quad \text{and} \quad \partial^d b = r + (1 + J)s + \omega.$$

Hence

$$\begin{aligned} (\partial^d)^2x &= \partial^d(a + (1 + J)b + \theta) \\ &= (1 + J)r + (1 + J)(r + (1 + J)s + \omega) + \partial^d\theta = (1 + J)\omega + \partial^d\theta, \end{aligned}$$

which is zero, by definition of $\partial^d\theta$.

(1b) Now suppose ∂x does contain η , so that $\partial x = a + (1 + J)b + \eta$. Then

$$\partial^d x = a + (1 + J)b + \omega.$$

Due to the dimensional grading, neither a nor b can contain η in their differential, and similarly for any cell appearing in ∂a or ∂b . Hence $\partial^d a = \partial a$ and $\partial^d b = \partial b$. Since $\partial^d \omega = \partial \eta$, it easily follows that $(\partial^d)^2 x = \partial^2 x = 0$.

(2) Due to the dimensional grading, ζ cannot contain η in its differential, and similarly for any cell appearing in $\partial \zeta$. Hence $\partial^d \zeta = \partial \zeta$. It follows that $(\partial^d)^2 \omega = \partial^2 \eta = 0$.

(3) We evidently have that $(\partial^d)^2 \theta = (1 + J)^2 \zeta = 0$.

By J -equivariance, we have that ∂^d squares to zero on JC and $J\omega$ also. This completes the verification that ∂^d is indeed a differential. \square

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

Definition 4.4 We define gr as follows:

- (1) If $x \in C$ (or JC), then $\text{gr}(x)$ is the same in $X^d(\Delta)$ as it is in X .
- (2) $\text{gr}(\omega) = \text{gr}(J\omega) = \text{gr}(\eta)$.
- (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 C and JC 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. (Compare with Section 3.2.)

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 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.

This completes the verification that $X^d(\Delta)$ is a well-defined geometric complex. 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.

4.2 Local equivalence

Let X be a split complex with a choice of splitting given by C , and let $X^d(\Delta)$ be its double as defined in Section 4.1. The main result of this section will be to show that $X^d(\Delta)$ is locally equivalent to $X \otimes X_\Delta$.

We employ the following principle from [2, Lemma 3.3]. Let C_1 and C_2 be two geometric complexes graded by the same coset of $2\mathbb{Z}$ in \mathbb{Q} , and let f be a cellular chain map from the skeleton of C_1 to the skeleton of C_2 . Suppose that for any pair of cells x from C_1 and y from C_2 with y appearing in $f(x)$, we have $\text{gr}(y) \geq \text{gr}(x)$. Then we can lift f to a grading-preserving chain map from C_1 to C_2 by multiplying through by the appropriate powers of U , ie

$$\tilde{f}(x) = \sum_{y_i \text{ appearing in } f(x)} U^{\frac{1}{2}(\text{gr}(y_i) - \text{gr}(x))} y_i.$$

To construct maps between $X^d(\Delta)$ and $X \otimes X_\Delta$, we will produce chain maps between their skeleta and check the monotonicity condition on their gradings separately.

Recall that X_Δ has three generators over $\mathbb{F}[U]$, which we denote by α , $J\alpha$ and β . These have gradings given by $\text{gr}(\alpha) = \text{gr}(J\alpha) = 0$ and $\text{gr}(\beta) = -2\Delta$; the skeleton of X_Δ has differential $\partial\beta = \alpha + J\alpha$.

We begin by constructing a chain map from $X^d(\Delta)$ to $X \otimes X_\Delta$. Here, for convenience we use \cdot to denote products of cells, eg $x \cdot \alpha$ for the product of x in X and α in X_Δ . It may be helpful for the reader to consult Example 4.8 and Figure 12.

Lemma 4.5 *We can construct a chain map f from $X^d(\Delta)$ to $X \otimes X_\Delta$.*

Proof We define our map f on generators $x \in C$, ω and θ of $X^d(\Delta)$, and extend by requiring J -equivariance.

- (1) Let $x \in C$, and suppose $\partial x = a + (1 + J)b (+\eta)$. Define $f(x) = x \cdot \alpha + (Jb) \cdot \beta$.
- (2) Define $f(\omega) = \eta \cdot \alpha + (J\zeta) \cdot \beta$.
- (3) Define $f(\theta) = \eta \cdot \beta$.

Let us first verify the monotonicity condition on the gradings. If $x \in C$, we see that the first term appearing in $f(x)$ has the same grading as x , since $\text{gr}(x \cdot \alpha) = \text{gr}(x)$. All the cells appearing in b have grading greater than or equal to that of x (since b appears in ∂x), and the gradings of the cells appearing in $(Jb) \cdot \beta$ are precisely these gradings shifted down by 2Δ . Since $2\Delta \leq \text{width}(X)$, it easily follows that the gradings of the cells appearing in $(Jb) \cdot \beta$ are still greater than or equal to that of x . The other cases are checked similarly.

We now show that f is a chain map. As in the proof that ∂^d is a differential, we have several cases to consider:

(1) Let $x \in C$. We divide into two cases based on whether or not ∂x contains η :

(1a) Let $\partial x = a + (1 + J)b$. As before, we have

$$\partial a = (1 + J)r \quad \text{and} \quad \partial b = r + (1 + J)s (+\eta).$$

If ∂b does not contain η , then we compute

$$\begin{aligned} f(\partial^d x) &= f(a + (1 + J)b) = f(a) + f(b) + Jf(b) \\ &= a \cdot \alpha + (Jr) \cdot \beta + b \cdot \alpha + (Js) \cdot \beta + (Jb) \cdot (J\alpha) + s \cdot \beta \\ &= (a + b) \cdot \alpha + (Jb) \cdot (J\alpha) + (Jr + Js + s) \cdot \beta, \\ \partial f(x) &= \partial(x \cdot \alpha + (Jb) \cdot \beta) = (\partial x) \cdot \alpha + J(\partial b) \cdot \beta + (Jb) \cdot (\partial \beta) \\ &= (a + (1 + J)b) \cdot \alpha + (Jr + (1 + J)s) \cdot \beta + (Jb) \cdot (\alpha + J\alpha) \\ &= (a + b) \cdot \alpha + (Jb) \cdot (J\alpha) + (Jr + Js + s) \cdot \beta, \end{aligned}$$

as desired. If ∂b does contain η , then the above computation must be slightly modified, and we have an extra term:

$$\begin{aligned} f(\partial^d x) &= f(a + (1 + J)b + \theta) \\ &= (a + b) \cdot \alpha + (Jb) \cdot (J\alpha) + (Jr + Js + s) \cdot \beta + f(\theta) \\ &= (a + b) \cdot \alpha + (Jb) \cdot (J\alpha) + (Jr + Js + s) \cdot \beta + \eta \cdot \beta. \end{aligned}$$

Similarly,

$$\begin{aligned} \partial f(x) &= \partial(x \cdot \alpha + (Jb) \cdot \beta) \\ &= (a + (1 + J)b) \cdot \alpha + (Jr + (1 + J)s + \eta) \cdot \beta + (Jb) \cdot (\alpha + J\alpha) \\ &= (a + b) \cdot \alpha + (Jb) \cdot (J\alpha) + (Jr + Js + s) \cdot \beta + \eta \cdot \beta, \end{aligned}$$

as desired.

(1b) Let $\partial x = a + (1 + J)b + \eta$. We again write $\partial a = p + (1 + J)q$ and $\partial b = r + (1 + J)s$, noting that ∂a and ∂b cannot contain η due to the dimensional grading. Imposing the condition that $\partial^2 x = 0$ yields $p = 0$ and $q + r + \zeta = 0$. We thus compute

$$\begin{aligned} f(\partial^d x) &= f(a + (1 + J)b + \omega) = f(a) + f(b) + Jf(b) + f(\omega) \\ &= a \cdot \alpha + (Jq) \cdot \beta + b \cdot \alpha + (Js) \cdot \beta + (Jb) \cdot (J\alpha) + s \cdot \beta + \eta \cdot \alpha + (J\zeta) \cdot \beta \\ &= (a + b + \eta) \cdot \alpha + (Jb) \cdot (J\alpha) + (Jq + J\zeta + Js + s) \cdot \beta, \\ \partial f(x) &= \partial(x \cdot \alpha + (Jb) \cdot \beta) \\ &= (a + (1 + J)b + \eta) \cdot \alpha + (Jr + (1 + J)s) \cdot \beta + (Jb) \cdot (\alpha + J\alpha) \\ &= (a + b + \eta) \cdot \alpha + (Jb) \cdot (J\alpha) + (Jr + Js + s) \cdot \beta, \end{aligned}$$

with equality following from the fact that $Jq + J\zeta = Jr$, since $q + r + \zeta = 0$.

(2) The proof that f is a chain map on ω is handled similarly. Let $\partial \zeta = r + (1 + J)s$. Imposing the condition that $\partial^2 \eta = 0$, we see that $r = 0$. Hence we compute

$$\begin{aligned} f(\partial^d \omega) &= f((1 + J)\zeta) = \zeta \cdot \alpha + (Js) \cdot \beta + (J\zeta) \cdot (J\alpha) + s \cdot \beta, \\ \partial f(\omega) &= \partial(\eta \cdot \alpha + (J\zeta) \cdot \beta) = (\zeta + J\zeta) \cdot \alpha + (1 + J)s \cdot \beta + (J\zeta) \cdot (\alpha + J\alpha) \\ &= \zeta \cdot \alpha + s \cdot \beta + (Js) \cdot \beta + (J\zeta) \cdot (J\alpha), \end{aligned}$$

as desired.

(3) Finally we verify that f is a chain map on θ . We have

$$f(\partial^d \theta) = f((1 + J)\omega) = \eta \cdot \alpha + (J\zeta) \cdot \beta + \eta \cdot (J\alpha) + \zeta \cdot \beta$$

and

$$\partial f(\theta) = \partial(\eta \cdot \beta) = (\zeta + J\zeta) \cdot \beta + \eta \cdot (\alpha + J\alpha),$$

which are evidently equal.

This completes the verification that f is a chain map. □

We now construct a chain map g from $X \otimes X_\Delta$ to $X^d(\Delta)$. It may be useful to consult Example 4.8 and Figure 12.

Lemma 4.6 *We can construct a chain map g from $X \otimes X_\Delta$ to $X^d(\Delta)$.*

Proof As usual, we define g piecewise on different generators of $X \otimes X_\Delta$.

(1) We first define g on $C \otimes \alpha$. Let $x \in C$. We define

$$g(x \cdot \alpha) = x,$$

where on the right-hand side we mean the corresponding element $x \in (C \oplus \omega)$ in $X^d(\Delta)$. By imposing J -equivariance, this defines g on $(JC) \otimes (J\alpha)$. Defining g on $C \otimes (J\alpha)$ is similar, but with an extra special case:

$$g(x \cdot (J\alpha)) = \begin{cases} x & \text{if } \partial x \text{ does not contain } \eta, \\ x + \theta & \text{if } \partial x \text{ contains } \eta. \end{cases}$$

By imposing J -equivariance, this also defines g on $(JC) \otimes \alpha$.

(2) We define

$$g(\eta \cdot \alpha) = \omega.$$

Note that imposing J -equivariance forces $g(\eta \cdot (J\alpha)) = J\omega$.

(3) We define g on $C \otimes \beta$ to be zero. By J -equivariance, this defines g on $(JC) \otimes \beta$ to also be zero. We set $g(\eta \cdot \beta) = \theta$.

As usual, we must verify that g satisfies the usual property involving the gradings. In all cases except when we send $x \cdot (J\alpha)$ to $x + \theta$, it is evident that g as written preserves gr , whereas in the exceptional case the claim follows from the fact that $\text{gr}(\theta) = \text{gr}(\eta) - 2\Delta$ and $2\Delta \leq \text{width}(X)$.

We now verify that g is a chain map. As usual, we proceed using copious amounts of casework:

(1) We first check that g is a chain map on $C \otimes \alpha$ and $C \otimes (J\alpha)$.

(1a) First suppose that ∂x does not contain η , so that $\partial x = a + (1 + J)b$. If ∂b does not contain η , we have

$$g(\partial(x \cdot \alpha)) = g((a + (1 + J)b) \cdot \alpha) = g(a \cdot \alpha) + g(b \cdot \alpha) + Jg(b \cdot (J\alpha)) = a + (1 + J)b,$$

while

$$\partial^d g(x \cdot \alpha) = \partial^d x = a + (1 + J)b,$$

as desired. A similar computation holds for $x \cdot (J\alpha)$. (Note that here we have used the fact that ∂a cannot contain η , since $\partial^2 x = 0$.) Now suppose that ∂b does contain η . Then $g(b \cdot (J\alpha)) = b + \theta$, rather than just Jb . Hence, in this case,

$$g(\partial(x \cdot \alpha)) = g((a + (1 + J)b) \cdot \alpha) = a + (1 + J)b + \theta,$$

while

$$\partial^d g(x \cdot \alpha) = \partial^d x = a + (1 + J)b + \theta,$$

as desired. Again, a similar computation holds for $x \cdot (J\alpha)$.

(1b) Now suppose $\partial x = a + (1 + J)b + \eta$. Then we compute

$$g(\partial(x \cdot \alpha)) = g((a + (1 + J)b + \eta) \cdot \alpha) = a + (1 + J)b + \omega,$$

while

$$\partial^d g(x \cdot \alpha) = \partial^d x = a + (1 + J)b + \omega,$$

as desired. Here, we have used the fact that a and b cannot have η in their differential, due to the dimensional grading. Similarly,

$$g(\partial(x \cdot (J\alpha))) = g((a + (1 + J)b + \eta) \cdot (J\alpha)) = a + (1 + J)b + J\omega,$$

while

$$\partial^d g(x \cdot (J\alpha)) = \partial^d (x + \theta) = a + (1 + J)b + \omega + (\omega + J\omega) = a + (1 + J)b + J\omega,$$

as desired.

(2) We check that

$$g(\partial(\eta \cdot \alpha)) = g((1 + J)\zeta \cdot \alpha) = (1 + J)\zeta,$$

while

$$\partial^d g(\eta \cdot \alpha) = \partial^d \omega = (1 + J)\zeta,$$

as desired.

(3) For any $x \in C$, we have that $\partial^d g(x \cdot \beta) = 0$. On the other hand,

$$g(\partial(x \cdot \beta)) = g((\partial x) \cdot \beta + x \cdot (\alpha + J\alpha)).$$

If ∂x does not contain η , then $g((\partial x) \cdot \beta) = 0$ and $g(x \cdot \alpha) = g(x \cdot (J\alpha)) = x$, so this is zero. If ∂x does contain η , then $g((\partial x) \cdot \beta) = \theta$, while $g(x \cdot \alpha) = x$ and $g(x \cdot (J\alpha)) = x + \theta$, so this is again zero. Finally, we easily verify that

$$g(\partial(\eta \cdot \beta)) = g((\zeta + J\zeta) \cdot \beta + \eta \cdot (\alpha + J\alpha)) = \omega + J\omega,$$

while

$$\partial^d g(\eta \cdot \beta) = \partial^d (\theta) = (1 + J)\omega,$$

as desired.

This completes the verification that g is a chain map. □

Now suppose that X is a split ι -complex, so that $U^{-1}H_*(X) = \mathbb{F}[U, U^{-1}]$. Technically, we have not yet shown that the doubled complex $X^d(\Delta)$ is also an ι -complex, in the sense that we have not formally verified $U^{-1}H_*(X^d(\Delta)) = \mathbb{F}[U, U^{-1}]$. This

property is clear in the geometric picture, since doubling a complex does not change the underlying homotopy type of the skeleton. (Hence doubling a cell complex with trivial homology should give a cell complex which still has trivial homology.) We will provide a more concrete proof in the next section when we explicitly describe the homology of $X^d(\Delta)$. For now, let us state:

Theorem 4.7 *Let X be a split ι -complex, and let Δ be a nonnegative even integer such that $2\Delta \leq \text{width}(X)$. Then $X^d(\Delta)$ is locally equivalent to $X \otimes X_\Delta$.*

Proof Observe that $g(f(\omega)) = \omega$, $g(f(\theta)) = \theta$ and $g(f(x)) = x$ for any $x \in C$. Hence the composition

$$g \circ f: X^d(\Delta) \rightarrow X \otimes X_\Delta$$

is the identity. Since X and X_Δ are both ι -complexes, their tensor product is an ι -complex, and we have $U^{-1}H_*(X \otimes X_\Delta) = \mathbb{F}[U, U^{-1}]$. Modulo the fact that $U^{-1}H_*(X^d(\Delta)) = \mathbb{F}[U, U^{-1}]$, it easily follows that f and g must be isomorphisms between $U^{-1}H_*(X^d(\Delta))$ and $U^{-1}H_*(X \otimes X_\Delta)$. \square

Note that although, a priori, the definition of $X^d(\Delta)$ depends on a choice of splitting C for X , Theorem 4.7 does not depend on this auxiliary data.

Example 4.8 It is helpful to work through a concrete example of the maps f and g . In Figure 12, we have displayed a sample split complex X , together with its double $X^d(\Delta)$ and the skeleton of the tensor product $X \otimes X_\Delta$. We have marked the J -invariant cell η in purple, and chosen a splitting C given by the span of the red cells. The corresponding cells of C in $X^d(\Delta)$ are also marked in red, while we have marked ω in purple and θ in blue. The reader who wishes to gain intuition for the construction of the maps f and g should stare at Figure 12 and attempt to construct J -equivariant chain maps between $X^d(\Delta)$ and $X \otimes X_\Delta$ by hand.

In the second row, we have displayed the action of f . 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—whose boundary contains a nonzero component lying in $(1+J)C$ —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 to each other in the obvious way. The reader should verify that extending via J -equivariance yields a cellular chain map.

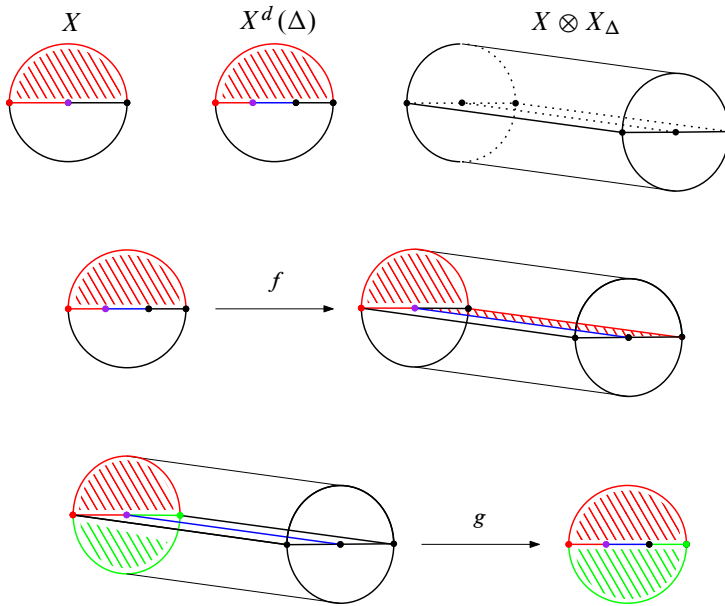


Figure 12: Local maps between $X^d(\Delta)$ and $X \otimes X_\Delta$.

In the third row, we have displayed the action of g . 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 that extending via J -equivariance yields a cellular chain map.

5 Connected homology

At this point, it is clear that we can produce a representative complex for any linear combination of the X_i by inductively using Theorem 4.7 and dualizing whenever needed. More precisely, consider a linear combination

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

with $i_1 \geq i_2 \geq \dots \geq i_n$. We start with $\pm X_{i_1}$, which is a split complex with width i_1 . At each step k , we either add or subtract X_{i_k} . In the former case, we use Theorem 4.7 (as usual) to obtain a new split complex with width i_k . In the latter, we first dualize our original complex to produce a split complex with the same width, and then add X_{i_k} to the dualized complex and dualize again. This similarly yields a new split complex with width i_k , and we proceed by induction. (Alternatively, we could explicitly write

down the halving operation and show that it is locally equivalent to subtracting X_{i_k} .) Note that the representative complex we obtain at the end of this process has $2n + 1$ generators, rather than the 3^n generators obtained by naively taking the full tensor product.

We now investigate the homology of the doubled complex.

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 .

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 and 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

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

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 nonzero 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 nonzero 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, ie

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

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

Note that Lemma 5.1 shows that if $U^{-1}H_*(X) = \mathbb{F}[U, U^{-1}]$, then $U^{-1}H_*(X^d(\Delta)) = \mathbb{F}[U, U^{-1}]$, completing the proof of Theorem 4.7.

Lemma 5.1 easily implies the following precursor to Theorem 1.1. Below, we still allow canceling pairs of basis elements (ie $X_i - X_i$) to appear in our linear combinations.

Lemma 5.2 *Consider a linear combination*

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

with $i_1 \geq i_2 \geq \dots \geq i_n$. Let Σ be the representative complex for X obtained by inductively applying Theorem 4.7 (dualizing when necessary). Then the homology of Σ has torsion part given by the algorithm of Theorem 1.1 (allowing for canceling pairs $X_i - X_i$).

Proof We proceed by induction. Let Σ_k denote the representative complex for the first k terms in the above linear combination. We simultaneously show that (a) $H_*(\Sigma_k)$ is given by the desired construction, and (b):

- (1) If the sign of X_{i_k} is positive, then the J -invariant cell of Σ_k has Maslov grading g_k one lower than the tail of the i_k -tower.
- (2) If the sign of X_{i_k} is negative, then the J -invariant cell of Σ_k has Maslov grading g_k equal to the tail of the i_k -tower.

The statements are easily established in the case of only one summand $\pm X_{i_1}$. Note that dualizing reflects torsion homology towers over the horizontal line of Maslov grading $\frac{1}{2}$, so it easily follows that (a) and (b) hold for Σ_k if and only if they hold for $-\Sigma_k$. Proceeding inductively, we assume that (a) and (b) hold for Σ_k (and also $-\Sigma_k$). Consider Σ_{k+1} . If the sign of $X_{i_{k+1}}$ is positive, then (a) follows from Lemma 5.1, together with (a) and (b) for Σ_k . We easily compute $g_{k+1} = g_k - 2i_{k+1} + 1$, since the dimensional grading of the J -invariant cell increases by one while the value of gr decreases by $2i_{k+1}$. This establishes (b). If the sign of $X_{i_{k+1}}$ is negative, then we apply the preceding argument to establish (a) and (b) for $-\Sigma_{k+1} = -\Sigma_k + X_{i_{k+1}}$. \square

Lemma 5.2 does not quite give the connected homology, since (for example) any canceling pair of basis elements $X_i - X_i$ introduces two towers in the homology of Σ , while in local equivalence $X_i - X_i$ is zero. The content of Theorem 1.1 is that this is the only further simplification in the homology which can happen. The author would like to thank Jennifer Hom for helpful comments regarding the following proof.

Proof of Theorem 1.1 We proceed by induction on n . For $X = c_1 X_{i_1}$, the claim is a result of Example 3.3. Now suppose that we have a maximally simplified linear combination

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

with $i_1 > i_2 > \cdots > i_n$, and 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, $H_{\text{conn}}(A)$ has 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, $H_{\text{conn}}(A)$ does not contain any towers of

length 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 Lemma 5.2, we have a representative complex Σ for X whose torsion homology is already given by the construction claimed in the statement of Theorem 1.1. Note that $H_*(\Sigma)$ has exactly $|c_k|$ towers of length i_k for each $1 \leq k \leq n$. The previous paragraphs then imply that the connected homology of X must be all of this homology, as desired. The grading shift in the statement of the theorem follows from taking into account the shift by -2 in the map h , together with the shift by $-d$ in Theorem 2.5 and the shift by -1 in the definition of connected Floer homology. \square

6 Examples and applications

In this section, we prove Theorem 1.2, as well as Corollaries 1.3 and 1.4. We begin with a proof of Theorem 1.2.

Proof of Theorem 1.2 Let l and k be positive integers, and let d be any integer. Consider the $\mathbb{F}[U]$ -module

$$\bigoplus_{i=0}^{k-1} \mathcal{T}_{d-i(2l-1)}(l)$$

formed by concatenating k different U -torsion towers of length l , where the towers are placed so that the uppermost element of each tower has grading one less than the lowermost element of the previous tower. We refer to this as a *chain of towers (of length l)*. As in Section 1.1, we may identify chains of towers as downwards pointing or upwards pointing. In the former case, we define the head of a chain to be the element with maximal grading and the tail to be the element with minimal grading; in the latter, we reverse these conditions. It is clear from the algorithm of Theorem 1.1 that if we group together all the U -torsion towers in HF_{conn} of a fixed length, they must form a chain.

Now sort these chains in descending order according to their tower lengths l . According to the algorithm of Theorem 1.1, the first chain in this list will either point downwards and have its head in grading $d(Y) - 1$, or point upwards and have its head in grading $d(Y)$. The orientation of each subsequent chain can then be inductively determined by its position relative to the tail of the previous chain. The theorem easily follows. \square

We now rephrase the computations of [3] in terms of the connected Floer homology.

Proof of Corollary 1.3 According to [3, Theorem 1.2] as written, we have

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

where $Y_i = X_i[-2i]$; see the discussion following [3, Lemma 4.3]. Our previous decomposition

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

thus shows that

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

where the signs in front of the i_k are the same as the signs appearing in the decomposition of $h(Y, \mathfrak{s})$. Each index i_k contributes a tower of length i_k to $\text{HF}_{\text{conn}}(Y, \mathfrak{s})$, 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

Proof of Corollary 1.4 The Rokhlin invariant reduces to the Neumann–Siebenmann invariant modulo 2. \square

We now give an illustration of Theorem 1.1 in which we also give the actual complexes constructed via the doubling and halving operations of Section 3.

Example 6.1 Consider a maximally simplified linear combination of the form

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

The representative complex of Lemma 5.2 may be constructed via a sequence of iterated doubling and halving operations, as displayed on the left in Figure 13. We begin with the geometric complex of X_{i_1} , double to obtain $X_{i_1} + X_{i_2}$, halve 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 13.

One can verify by hand that the homology of this complex has the form prescribed by Theorem 1.1, 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

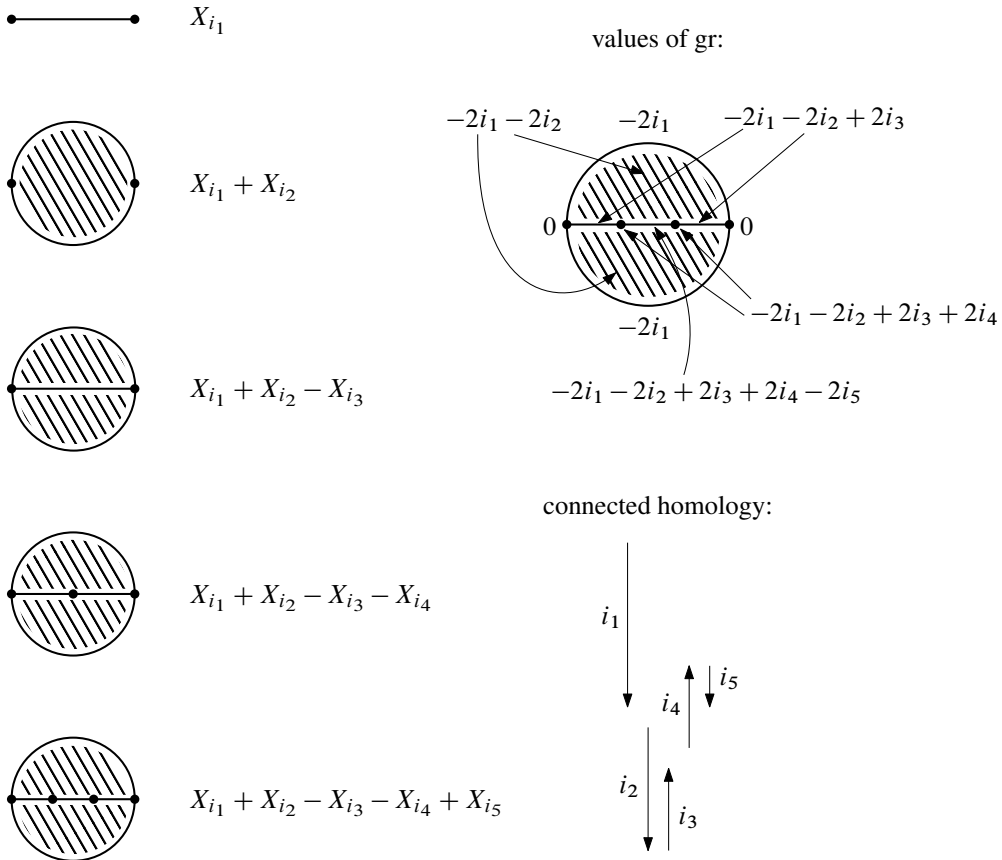


Figure 13: Doubling and halving operations (left), final values of gr (upper-right), and homology (lower-right).

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 5.1 for the proof of Theorem 1.1.

Finally, we close with an example of an ι -complex whose connected homology is not of the form prescribed by Theorem 1.1. Again, we stress that it is not known whether this is realized as the ι -complex of any actual 3-manifold; such an example would show that Seifert fibered spaces do not generate the homology cobordism group.

Example 6.2 Consider the geometric complex whose skeleton is given by the usual J -equivariant cellular decomposition of D^2 , as in Figure 14. While this is the same




Figure 14: A connected homology not arising via the procedure of Theorem 1.1.

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

$$\partial e_1 = \partial Je_1 = U^x(e_0 + Je_0) \quad \text{and} \quad \partial e_2 = U^y(e_1 + Je_1),$$

where $x < y$. The resulting connected homology is displayed (modulo the grading shift by one) on the right in Figure 14. 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 14 (keeping in mind its placement in relation to the d –invariant) cannot be achieved through the procedure of Theorem 1.1.

References

- [1] **I Dai**, *On the Pin(2)–equivariant monopole Floer homology of plumbed 3–manifolds*, Michigan Math. J. 67 (2018) 423–447 MR
- [2] **I Dai, C Manolescu**, *Involutive Heegaard Floer homology and plumbed three–manifolds*, J. Inst. Math. Jussieu 18 (2019) 1115–1155
- [3] **I Dai, M Stoffregen**, *On homology cobordism and local equivalence between plumbed manifolds*, Geom. Topol. 23 (2019) 865–924 MR
- [4] **R Fintushel, R J Stern**, *Instanton homology of Seifert fibred homology three spheres*, Proc. London Math. Soc. 61 (1990) 109–137 MR

- [5] **K A Frøyshov**, *Equivariant aspects of Yang–Mills Floer theory*, *Topology* 41 (2002) 525–552 MR
- [6] **M Furuta**, *Homology cobordism group of homology 3–spheres*, *Invent. Math.* 100 (1990) 339–355 MR
- [7] **K Hendricks, J Hom, T Lidman**, *Applications of involutive Heegaard Floer homology*, *J. Inst. Math. Jussieu* (online publication April 2019)
- [8] **K Hendricks, C Manolescu**, *Involutive Heegaard Floer homology*, *Duke Math. J.* 166 (2017) 1211–1299 MR
- [9] **K Hendricks, C Manolescu, I Zemke**, *A connected sum formula for involutive Heegaard Floer homology*, *Selecta Math.* 24 (2018) 1183–1245 MR
- [10] **F Lin**, *Pin(2)–monopole Floer homology, higher compositions and connected sums*, *J. Topol.* 10 (2017) 921–969 MR
- [11] **F Lin**, *The surgery exact triangle in Pin(2)–monopole Floer homology*, *Algebr. Geom. Topol.* 17 (2017) 2915–2960 MR
- [12] **F Lin**, *A Morse–Bott approach to monopole Floer homology and the triangulation conjecture*, *Mem. Amer. Math. Soc.* 1221, Amer. Math. Soc., Providence, RI (2018) MR
- [13] **C Manolescu**, *Pin(2)–equivariant Seiberg–Witten Floer homology and the triangulation conjecture*, *J. Amer. Math. Soc.* 29 (2016) 147–176 MR
- [14] **A Némethi**, *Lattice cohomology of normal surface singularities*, *Publ. Res. Inst. Math. Sci.* 44 (2008) 507–543 MR
- [15] **W D Neumann**, *An invariant of plumbed homology spheres*, from “Topology Symposium Siegen 1979” (U Koschorke, W D Neumann, editors), *Lecture Notes in Math.* 788, Springer (1980) 125–144 MR
- [16] **P Ozsváth, Z Szabó**, *On the Floer homology of plumbed three-manifolds*, *Geom. Topol.* 7 (2003) 185–224 MR
- [17] **L Siebenmann**, *On vanishing of the Rohlin invariant and nonfinitely amphicheiral homology 3–spheres*, from “Topology Symposium Siegen 1979” (U Koschorke, W D Neumann, editors), *Lecture Notes in Math.* 788, Springer (1980) 172–222 MR
- [18] **M Stoffregen**, *Pin(2)–equivariant Seiberg–Witten Floer homology of Seifert fibrations*, preprint (2015) arXiv
- [19] **M Stoffregen**, *Manolescu invariants of connected sums*, *Proc. Lond. Math. Soc.* 115 (2017) 1072–1117 MR

Department of Mathematics, Princeton University
 Princeton, NJ, United States
 idai@math.princeton.edu

Received: 4 May 2018 Revised: 8 October 2018