Geometry & Topology

Volume 27 (2023)

A calculus for bordered Floer homology

JONATHAN HANSELMAN
LIAM WATSON

msp

msp

Geometry & Topology **27**:3 (2023) 823–924
DOI: 10.2140/gt.2023.27.823
Published: 8 June 2023

A calculus for bordered Floer homology

JONATHAN HANSELMAN
LIAM WATSON

We consider a class of manifolds with torus boundary admitting bordered Heegaard Floer homology of a particularly simple form; namely, the type D structure may be described graphically by a disjoint union of loops. We develop a calculus for studying bordered invariants of this form and, in particular, provide a complete description of slopes giving rise to L–space Dehn fillings as well as necessary and sufficient conditions for L–spaces resulting from identifying two such manifolds along their boundaries. As an application, we show that Seifert-fibred spaces with torus boundary fall into this class, leading to a proof that, among graph manifolds containing a single JSJ torus, the property of being an L–space is equivalent to non-left-orderability of the fundamental group and to the nonexistence of a coorientable taut foliation.

57M27

1 Introduction

This paper is concerned with developing Heegaard Floer theory with a view to better understanding the relationship between coorientable taut foliations, left-orderable fundamental groups and manifolds that do not have *simple* Heegaard Floer homology. Recall that manifolds with simplest possible Heegaard Floer homology, called L–spaces, are rational homology spheres Y for which $\dim \widehat{HF}(Y) = |H_1(Y;\mathbb{Z})|$ (all Heegaard Floer–theoretic objects in this work will take coefficients in $\mathbb{Z}/2\mathbb{Z}$). On the other hand, a group G is left-orderable if there exists a nonempty set $P \subset G$, called a positive cone, that is a closed subsemigroup of G and gives a partition of the group in the sense that $G = P \amalg \{1\} \amalg P^{-1}$. This auxiliary structure is equivalent to G admitting an effective action on $\mathbb{R}$ by order-preserving homeomorphisms. The work of Boyer, Rolfsen and Wiest is a good introduction to left-orderable groups in the context of three-manifold topology, including the interaction with taut foliations [4]; for closed, orientable, irreducible, three-manifolds it is conjectured that being an L–space is equivalent to having a non-left-orderable fundamental group; see Boyer, Gordon and Watson [3]. This conjecture holds for Seifert-fibred spaces and, as a natural extension of this case, graph manifolds are a key family of interest; see Boileau and Boyer [1], Boyer and Clay [2], Clay, Lidman and Watson [5], Hanselman [8; 9] and Mauricio [21]. Towards establishing the conjecture for graph manifolds, we prove:

Theorem 1.1 *Suppose that Y is a graph manifold with a single JSJ torus; that is, Y is constructed by identifying two Seifert-fibred manifolds with torus boundary along their boundaries. Then the following are equivalent:*

(i) *Y is an L–space.*

(ii) *$\pi_1(Y)$ is not left-orderable.*

(iii) *Y does not admit a coorientable taut foliation.*

The equivalence between (ii) and (iii) is due to Boyer and Clay [2]; the focus of this paper is understanding the behaviour of Heegaard Floer homology in this setting. To do this, we make use of bordered Heegaard Floer homology, a variant of Heegaard Floer homology adapted to cut-and-paste arguments. Briefly, this theory assigns a differential graded module over a particular algebra to each manifold with torus boundary. A chain complex for the Heegaard Floer homology of the associated closed manifold is obtained from a pairing theorem due to Lipshitz, Ozsváth and Thurston [20].

Our approach to this problem is to work in a more general setting. We consider a particular class of differential graded modules which we call loop-type (Definition 3.2), and introduce a calculus for studying loops; the bulk of this paper is devoted to developing this calculus in detail. Given a three-manifold M with torus boundary, M is called loop-type if its associated bordered invariants are loop-type up to homotopy (Definition 3.13). Recall that a slope in ∂M is the isotopy class of an essential simple closed curve in ∂M, and denote by $M(\gamma)$ the closed three manifold resulting from Dehn filling along a slope γ. The set of slopes may be (noncanonically) identified with the extended rationals $\mathbb{Q} \cup \{\frac{1}{0}\}$. Of central interest is the subset $\mathcal{L}_M$ consisting of those slopes giving rise to L–spaces after Dehn filling. We prove:

Theorem 1.2 (detection) *Suppose that M is a loop-type rational homology torus. Then there is a complete, combinatorial description of the set $\mathcal{L}_M$ in terms of loop calculus. In particular, $\mathcal{L}_M$ may be identified with the restriction to $\mathbb{Q} \cup \{\frac{1}{0}\}$ of a connected interval in $\mathbb{R} \cup \{\infty\}$. Moreover, if M is simple loop-type, this interval has rational endpoints.*

We remark that this is the expected behaviour from the foliations/orderability vantage point, at least for graph manifolds. It is interesting that the analogous behaviour on the Heegaard Floer side appears to be intrinsic to the algebraic structures that arise. Namely, this is a statement that makes sense for loop-type bordered invariants, without reference to any three-manifold. The introduction of the technical notion of a *simple loop* (Definition 4.20) also allows us to state and prove a gluing theorem. Let $\mathcal{L}_M^\circ$ denote the interior of $\mathcal{L}_M$.

Theorem 1.3 (gluing) *Suppose that M_1 and M_2 are simple loop-type rational homology tori and neither is solid torus-like. Then, given a homeomorphism $h\colon \partial M_1 \to \partial M_2$, the closed manifold $M_1 \cup_h M_2$ is an L–space if and only if, for every slope γ in ∂M_1, either $\gamma \in \mathcal{L}_{M_1}^\circ$ or $h(\gamma) \in \mathcal{L}_{M_2}^\circ$.*

Note that Theorem 1.3 does not hold if either M_1 or M_2 is a solid torus. If M_1 is a solid torus with meridian $m = \partial D^2 \times \{\mathrm{pt}\}$, then, according to the definition of an L–space slope, $M_1 \cup_h M_2$ is an L–space if and only if $h(m) \in \mathcal{L}_{M_2}$; equivalently, the statement of Theorem 1.3 holds if we use $\mathcal{L}_{M_i}$ in place of $\mathcal{L}_{M_i}^\circ$. When both M_1 and M_2 are solid tori, this simply amounts to the construction of lens spaces interpreted in our notation. More generally, however, there are bordered invariants that arise in the loop setting that behave just like solid tori with respect to gluing. We will need to deal with these explicitly; this amounts to defining a class of manifolds and loops which are referred to as solid torus-like (Definition 3.20).

Theorems 1.2 and 1.3 follow from working with loops in the abstract. Towards the proof of Theorem 1.1, and in the interest of establishing an existence result, a key class of loop-type manifolds is provided by Seifert-fibred spaces.

Theorem 1.4 *Suppose M is a rational homology solid torus admitting a Seifert-fibred structure. Then M has simple loop-type bordered Heegaard Floer homology.*

Let N denote the twisted I–bundle over the Klein bottle, and recall that the rational longitude λ for this manifold with torus boundary may be identified with a fibre in a Seifert structure over the Möbius band. As a Seifert-fibred rational homology solid torus, N is a simple loop-type manifold; compare [3]. The twisted I–bundle over the Klein bottle allows for an alternative detection statement for L–space slopes.

Theorem 1.5 (detection via the twisted I–bundle over the Klein bottle) *Suppose that M is a loop-type rational homology torus. Then $\gamma \in \mathcal{L}^\circ_M$ if and only if $N \cup_h M$ is an L–space, where $h(\lambda) = \gamma$.*

This answers a question of Boyer and Clay in the case of connected boundary; compare [2, Question 1.8]. In particular, our notion of detection aligns precisely with the characterization given by Boyer and Clay [2]. Indeed, we prove that, more generally, the twisted I–bundle over the Klein bottle used in Theorem 1.5 may be replaced with any simple loop-type manifold for which every nonlongitudinal filling is an L–space; see Theorem 7.3. There are many examples of these provided by *Heegaard Floer homology solid tori*; for more on this class of manifolds, see Hanselman, Rasmussen and Watson [12, Section 1.5].

Note that the interior of $\mathcal{L}_M$, denoted by $\mathcal{L}^\circ_M$, is the set of strict L–space slopes. The complement of $\mathcal{L}^\circ_M$, according to Theorem 1.5, corresponds to the set of *non-L–space* (*NLS*) *detected slopes* in the sense of Boyer and Clay [2, Definition 7.16].

According to Theorem 1.4, the exterior of any torus knot in the three-sphere gives an example of a simple loop-type manifold (indeed, this follows from work of Lipshitz, Ozsváth and Thurston [20]). More generally, if K is a knot in the three-sphere admitting an L–space surgery (an L–space knot), then $S^3 \smallsetminus \nu(K)$ is a simple loop-type manifold. In this setting it is well known that $\mathcal{L}_M$ is (the restriction to the rationals of) $[2g-1,\infty]$, where g is the Seifert genus of K, whenever K admits a positive L–space surgery. Specializing Theorem 1.3 to this setting, we have:

Theorem 1.6 *Let K_i be an L–space knot and write $M_i = S^3 \smallsetminus \nu(K)$ for $i = 1, 2$. Given a homeomorphism $h\colon \partial M_1 \to \partial M_2$, the closed manifold $M_1 \cup_h M_2$ is a L–space if and only if $h(\mathcal{L}^\circ_{M_1}) \cup \mathcal{L}^\circ_{M_2} \cong \mathbb{Q} \cup \{\frac{1}{0}\}$.*

Note that this solves [2, Problem 1.11] by resolving (much more generally and in the affirmative) one direction of [9, Conjecture 1] (a special case of [2, Conjecture 1.10] or [5, Conjecture 4.3]). It should also be noted that, owing to the existence of hyperbolic L–space knots, Theorem 1.6 provides additional food for thought regarding

[2, Question 1.12]. Namely, one would like to know if the order-detected/foliation-detected slopes in the boundary of the exterior of an L–space knot K coincide with the complement of $\mathcal{L}_M^\circ$, where $M = S^3 \smallsetminus \nu(K)$. For integer homology three-spheres resulting from surgery on a knot in S^3, there has been considerable progress in this vein; see Li and Roberts for foliations [18], Boileau and Boyer for left-orders [1], and also Hedden and Levine for splicing results [14].

Related work

Nonsimplicity is somewhat sporadic; however, it can be shown explicitly that nonsimple graph manifolds exist. On the other hand, it is clear from examples that certain graph manifolds with torus boundary do not give rise to loop-type bordered invariants, at least not in any obvious way. These subtleties seem particularly interesting when contrasted with the behaviour of foliations and orders: for more complicated graph manifolds, more subtle notions of foliations and orders (relative to the boundary tori) need to be considered [2]. On the other hand, it is now clear from work of Rasmussen and Rasmussen [27] that simplicity (for loops) is more than a convenience: the class of simple loop-type manifolds is equivalent to the class of (exteriors of) Floer-simple knots; see Hanselman, Rasmussen, Rasmussen and Watson [10, Proposition 6] in particular. As a result, Theorem 1.3 may be recast in terms of Floer-simple manifolds [10, Theorem 7]. This observation gives rise to an extension of Theorem 1.1 to the case of general graph manifolds; this is the main result of a joint paper of the authors with Rasmussen and Rasmussen [10], illustrating a fruitful overlap between these two independent projects. We note that the methods used by Rasmussen and Rasmussen are different from those used in our work. They appeal to knot Floer homology instead of bordered Heegaard Floer homology. This leads to somewhat divergent results and emphasis: while Rasmussen and Rasmussen give a clear picture of the interval of L–space slopes in terms of classical invariants, our machinery is better suited to gluing results.

Structure of the paper

Section 2 collects the essentials of bordered Heegaard Floer homology and puts in place our conventions. In particular, the definitions of L–space and strict L–space slope are found here.

Section 3 is devoted to defining loops and loop calculus. This represents the key tool; loop calculus provides a combinatorial framework for studying bordered Heegaard Floer homology. Note that we define and work in an a priori broader setting of abstract loops. It seems likely that many of the loops considered do not represent the type D structure of any bordered three-manifold. This calculus is applied towards two distinct ends: detection and gluing.

Section 4 gives characterizations of L–space slopes and strict L–space slopes in the loop setting. In particular, we prove Theorem 1.2. The proof of the second part of this result — identifying the set of L–space slopes as the restriction of an interval with rational endpoints — requires a detailed study of non-L–space slopes. While interesting in its own right, this fact is essential to the gluing results that follow.

Section 5 proves Theorem 1.3 by first establishing the appropriate gluing statements for abstract loops. The section concludes with a complete characterization of L–spaces resulting from generalized splices of L–space knots in the three-sphere, proving Theorem 1.6.

Section 6 turns to the study of loop-type manifolds. We describe an algorithm for constructing rational homology sphere graph manifolds by way of three moves, and determine the effect of these moves on bordered invariants. Using this, we establish a class of graph manifold rational homology tori, subsuming Seifert-fibred rational homology tori, for which we now have a complete understanding of gluing in Heegaard Floer homology according to the material in the preceding two sections.

Section 7 collects all of the forgoing material in order to prove our main results. This section includes the proofs of Theorems 1.4 and 1.5, and from these results we prove Theorem 1.1.

Addendum: additional context and further developments

This paper was first posted to arXiv on August 22, 2015. In the intervening years there have been many closely related developments, depending on or growing directly out of this project. In hopes that this aids the present-day reader while maintaining the paper in its original form, we will expand on these here.

As described at the start of the paper, the genesis for this work was an interest in understanding the so-called L–space conjecture in the presence of an essential torus [2; 3]. While a comprehensive list of references detailing the support for this conjecture now seems too long to compile, we point to Dunfield [6] for more recent computational evidence. Theorem 1.1 provided a first step towards establishing the conjecture for graph manifolds. Indeed, the synchronous work of Rasmussen and Rasmussen [27] provided a key result that allowed the “simple loop-type” hypothesis to be replaced by

"Floer-simple"; then the proof of Theorem 1.1, which is the focus of Section 7, provides the scaffold used in our four-author project to establish the L–space conjecture for graph manifolds in full [10]. (Note that S Rasmussen gave a different, independent proof [28].)

One idea that is central to this work, but perhaps somewhat hidden from view, is the following: given a property of a bordered invariant that is unchanged under Dehn twists, one obtains a property associated with a three-manifold. This is used in Section 3, where we establish our class of interest, namely loop-type three-manifolds; see Definition 3.13.

Our strategy, at the beginning of this project, was to ignore the pathology of possible non-loop-type manifolds — based on empirical evidence that such manifolds never arose. The result of this, on the one hand, was the presence of an annoying technical assumption, but, on the other, provided the correct lens through which to view the bordered invariants for a manifold with torus boundary. Indeed, on completing this project and [10], it became clear that bordered invariants, subject to the loop-type condition, could be interpreted as immersed curves in a once-punctured torus. Moreover, gluing could now be stated in terms of Lagrangian intersection Floer homology in a once-punctured torus, giving rise to a gluing theorem that was then best-possible: Theorem 1.3 is restated for Floer-simple manifolds in [10, Theorem 7], and then further upgraded to loop-type manifolds with no simplicity hypothesis at all in the initial arXiv version of [11].

When viewed in these terms, loop calculus can be extended to general bordered invariants, giving rise to a graphical interpretation of the box tensor product in terms of intersections between train tracks — these latter are immersions of the graphs considered in this paper into the torus. Combined with new work of Haiden, Katzarkov and Kontsevich [7], this indicated that in the general setting our curves should carry local systems, which provides a means of handling graphs that cannot be converted to curves. So, once interpreted in the bordered setting, the loop-type hypothesis is simply the restriction to the trivial local system case. This is explained in detail in [12]. Note that it remains unclear what the role local systems play in general is — we are still not aware of an explicit example. (It is worth noting that local systems do play a role in Khovanov homology, which admits an analogous immersed curves description for decompositions along Conway spheres; see Kotelskiy, Watson and Zibrowius [16].) Nevertheless, these are easy to manage in practice, and up to minor modifications proofs go through. In particular, an alternative graphical approach to our work in Sections 4 (compare [11, Section 5]) and 5 (compare [11, Section 6]) now exists.

Despite these advances, however, it remains the case that, given a bordered invariant, one needs an algebrocombinatorial tool in order to work with and/or feed it into a computer. And, indeed, some operations are still only manageable using the tools described here. This is the case for annular gluing, for example, as this uses the "merge" described in Section 6, which can be adapted to the non-loop-type setting (that is, to handle nontrivial local systems). This is explained in more recent work, and, as an application, we establish a complete description of the effect of cabling on bordered invariants for knots in S^3 in terms of immersed curves in [13]. The proofs depend on the merge operation and loop calculus in an essential way.

As described in Section 3, then, "loop calculus" is meant to describe the tools needed to study bordered invariants for manifolds with torus boundary. In particular, we view this calculus as a toolkit for studying type D structures over the torus algebra in general. With hindsight, this is a viable algebrocombinatorial machinery for working with objects in the Fukaya category of the once-punctured torus. This was our attempt to standardize the seemingly ad hoc techniques that appear in earlier work of the authors; see [3; 9], for example.

Acknowledgements

We are very grateful to Jake and Sarah Rasmussen for their interest in this work and for their willingness to share their own [27]. This strengthened the present work considerably and facilitated the joint paper at the intersection of the two projects [10].

Hanselman was partially supported by NSF RTG grant DMS-1148490. Watson was partially supported by a Marie Curie Career Integration Grant (HFFUNDGRP).

2 Background and conventions

We begin by briefly recalling the essentials of bordered Heegaard Floer homology; for details see [20]. We restrict attention to compact, orientable three-manifolds M with torus boundary. Let $\mathbb{F}$ denote the two-element field.

2.1 Bordered structures

A bordered three-manifold, in this setting, is a pair (M, Φ) where $\Phi\colon S^1 \times S^1 \to \partial M$ is a fixed homeomorphism satisfying $\Phi(S^1 \times \{\mathrm{pt}\}) = \alpha$ and $\Phi(\{\mathrm{pt}\} \times S^1) = \beta$ for slopes α and β in ∂M satisfying $\Delta(\alpha, \beta) = 1$. Recall that a slope in ∂M is the isotopy class of an essential simple closed curve in ∂M or, equivalently, a primitive class in $H_1(\partial M; \mathbb{Z})/\{\pm 1\}$. The distance $\Delta(\cdot\,, \cdot)$ is measured by considering the minimal geometric intersection between slopes; thus, the requirement that $\Delta(\alpha,\beta)=1$ ensures that the pair $\{\alpha,\beta\}$, having chosen orientations, forms a basis for $H_1(\partial M;\mathbb{Z})$.

As a result, any bordered manifold (M,Φ) may be represented by the ordered triple (M,α,β), with the understanding that (M,α,β) and (M,β,α) differ as bordered manifolds (that is, these represent different bordered structures on the same underlying manifold M). We will adhere to this convention for describing bordered manifolds as it makes clear that bordered manifolds come with a pair of preferred slopes.

Definition 2.1 For a given bordered manifold (M,α,β), the slope α is referred to as the standard slope and the slope β is referred to as the dual slope.

2.2 The torus algebra

The torus algebra $\mathcal{A}$ is generated (as a vector space over $\mathbb{F}$) by elements

$$\iota_0,\quad \iota_1,\quad \rho_1,\quad \rho_2,\quad \rho_3,\quad \rho_{12},\quad \rho_{23},\quad \rho_{123}$$

with multiplication defined by

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

(all other products $\rho_I\rho_J$ vanish) and

$$\iota_0\rho_1=\rho_1=\rho_1\iota_1,\quad \iota_1\rho_2=\rho_2=\rho_2\iota_0,\quad \iota_0\rho_3=\rho_3=\rho_3\iota_1,$$

so that $\iota_0+\iota_1$ is a unit. Denote by $\mathcal{I}$ the subring of idempotents in $\mathcal{A}$ generated by ι_0 and ι_1. This algebra has various geometric interpretations; see [20]. The bordered Heegaard Floer invariants of (M,α,β) are modules of various types over $\mathcal{A}$, as we now describe.

2.3 Type *D* structures

A (left) type D structure over $\mathcal{A}$ is an $\mathbb{F}$–vector space N equipped with a left action of the idempotent subring $\mathcal{I}$ so that $N=\iota_0N\oplus\iota_1N$, together with a map

$$\delta^1\colon N\to\mathcal{A}\otimes_{\mathcal{I}}N$$

satisfying a compatibility condition with the multiplication on $\mathcal{A}$ [20]. The compatibility ensures that the map

$$\partial\colon \mathcal{A}\otimes_{\mathcal{I}}N\to\mathcal{A}\otimes_{\mathcal{I}}N,\quad a\otimes x\mapsto a\cdot\delta^1(x),$$

promotes $\mathcal{A} \otimes N$ to a left differential module over $\mathcal{A}$ (in particular, $\partial^2 = 0$), where $a \cdot (b \otimes y) = ab \otimes y$. While we will generally confuse type D structures and their associated differential modules, the advantage of the type D structure is in an iterative definition

$$\delta^k : N \to \mathcal{A}^{\otimes k} \otimes_{\mathcal{I}} N,$$

where $\delta^{k+1} = (\mathrm{id}_{\mathcal{A}^{\otimes k}} \otimes \delta^1) \circ \delta^k$ for $k > 1$. The type D structure is bounded if all δ^k are identically zero for sufficiently large k.

Given a basis for N, this structure may be described graphically. An $\mathcal{A}$–decorated graph is a directed graph with vertex set labelled by $\{\bullet, \circ\}$ and edge set labelled by elements of $\mathcal{A}$ consistent with the edge orientations. The labelling of the vertices specifies the splitting of the generating set according to the idempotents, while the edge set encodes the differential. For example, the $\mathcal{A}$–decorated graph

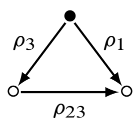

encodes the fact that there is a single generator x in the ι_0–idempotent with $\delta^1(x) = \rho_1 \otimes u + \rho_3 \otimes v$ (or $\partial(x) = \rho_1 u + \rho_3 v$). The higher maps in the type D structure can be extracted from following directed paths in the graph; for example, in the graph above we have $\delta^2(x) = \rho_3 \otimes \rho_{23} \otimes u$. By convention we drop the label on edges labelled by the identity element of $\mathcal{A}$.

An $\mathcal{A}$–decorated graph determines a type D structure with respect to a particular basis. In general we do not have a preferred choice of basis and we care about type D structures only up to homotopy equivalence, so there are many $\mathcal{A}$–decorated graphs which we should deem to be equivalent. Choosing a different basis for N leaves the vertex set unchanged and changes the arrows in a predictable way; for example, the basis change replacing x with $x + y$ adds an arrow out of the vertex labelled by x for each arrow out of the vertex labelled by y, and an arrow into the vertex labelled by y for each arrow into the vertex labelled by x. We can also replace an $\mathcal{A}$–decorated graph with one for a homotopy equivalent type D structure by using edge reduction (or its inverse) as described for example by Levine [17, Section 2.6]. Briefly, any segment of a graph of the form

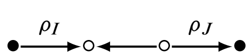

may be replaced by a single edge

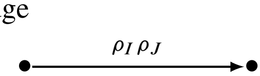

where the edge is simply deleted if the product $\rho_I \rho_J$ vanishes (we have chosen specific vertex labelling for illustration only). Again, we follow the convention that an unlabelled edge represents an edge labelled by the identity element of $\mathcal{A}$. Note that, by repeatedly cancelling such edges, we may always find a graph with no unlabelled edges.

Definition 2.2 An $\mathcal{A}$–decorated graph is reduced if no edge is labelled by the identity element of $\mathcal{A}$.

It turns out that any two $\mathcal{A}$–decorated graphs for homotopy equivalent type D structures can be related by a sequence of basis changes and edge reductions or insertions.

Given a bordered three-manifold (M, α, β), Lipshitz, Ozsváth and Thurston define a type D structure $\widehat{CFD}(M, \alpha, \beta)$ over $\mathcal{A}$ that is an invariant of the bordered manifold up to quasi-isomorphism [20]. As explained above, we will sometimes regard this object as a differential module over $\mathcal{A}$. This invariant splits over spin^c structures of M; that is,

$$\widehat{CFD}(M, \alpha, \beta) = \bigoplus_{\mathfrak{s} \in \mathrm{Spin}^c(M)} \widehat{CFD}(M, \alpha, \beta; \mathfrak{s}).$$

Note that $\mathrm{Spin}^c(M)$ can be identified with $H^2(M) \simeq H_1(M, \partial M)$.

2.4 Type A structures

A (right) type A structure over $\mathcal{A}$ is an $\mathbb{F}$–vector space M equipped with a right action of the idempotent subring $\mathcal{I}$ so that $M = M\iota_0 \oplus M\iota_1$, together with maps

$$m_{k-1} \colon M \otimes_{\mathcal{I}} \mathcal{A}^{\otimes k} \to M$$

satisfying the $\mathcal{A}_\infty$ relations; see [20]. That is, a type A structure is a right $\mathcal{A}_\infty$–module over $\mathcal{A}$. A type A structure is bounded if the m_k vanish for all sufficiently large k. Given a bordered three-manifold (M, α, β), Lipshitz, Ozsváth and Thurston define a type A structure $\widehat{CFA}(M, \alpha, \beta)$ over $\mathcal{A}$ that is an invariant of the bordered manifold up to quasi-isomorphism.

There is a similar graphical representation for type A structures. Indeed, owing to a duality between type D and type A structures for three-manifolds, $\widehat{CFA}(M, \alpha, \beta)$ may be deduced from the graph describing $\widehat{CFD}(M, \alpha, \beta)$ by appealing to an algorithm described by Hedden and Levine [14]. This algorithm takes subscripts $1 \mapsto 3$ and $3 \mapsto 1$ while fixing 2, with the convention that a conversion of the form $23 \mapsto 21$ must be parsed as $2, 1$ (this example is shown below, as it occurs in the conversion of the sample graph shown previously). In this type A context, sequences of directed edges must be concatenated in order to obtain all of the multiplication maps. For example, labelling the generators as before, the graph

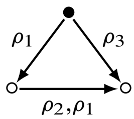

encodes operations $m_2(x, \rho_3) = v$, $m_2(x, \rho_1) = u$ and $m_3(u, \rho_2, \rho_1) = v$, as well as $m_3(x, \rho_{12}, \rho_2) = v$.

2.5 Pairing

Consider a closed, orientable three-manifold Y decomposed along a (possibly essential) torus so that $Y = M_1 \cup_h M_2$ for some homeomorphism $h\colon \partial M_1 \to \partial M_2$. If h has the property that $h(\alpha_1) = \beta_2$ and $h(\beta_1) = \alpha_2$, then we will write this decomposition as $Y = (M_1, \alpha_1, \beta_1) \cup (M_2, \alpha_2, \beta_2)$. The reason for this convention is to ensure compatibility with the paring theorem established by Lipshitz, Ozsváth and Thurston [20]. In particular, they prove that

$$\widehat{CF}(Y) \cong \widehat{CFA}(M_1, \alpha_1, \beta_1) \boxtimes \widehat{CFD}(M_2, \alpha_2, \beta_2),$$

where $\widehat{CF}(Y)$ is a chain complex with homology $\widehat{HF}(Y)$. As a vector space over $\mathbb{F}$, this chain complex is generated by tensors (over $\mathcal{I} \in \mathcal{A}$) of the form $x \otimes_{\mathcal{I}} y$, where $x \in \widehat{CFA}(M_1, \alpha_1, \beta_1)$ and $y \in \widehat{CFD}(M_2, \alpha_2, \beta_2)$, with differential

$$\partial^{\boxtimes}(x \otimes_{\mathcal{I}} y) = \sum_{k=0}^{\infty} (m_{k+1} \otimes \mathrm{id})(x \otimes_{\mathcal{I}} \delta^k(y)),$$

which is a finite sum provided at least one of the modules in the pairing is bounded.

As a particular special case, consider the pairing theorem in the context of Dehn filling. Given a three-manifold M with torus boundary, write $M(\alpha)$ for the result of Dehn filling M along the slope α; that is, $M(\alpha) = (D^2 \times S^1) \cup_h M$, where the homeomorphism h is determined by $h(\partial D^2 \times \{\mathrm{pt}\}) = \alpha$. In particular,

$$\widehat{CF}(M(\alpha)) \cong \widehat{CFA}(D^2 \times S^1, l, m) \boxtimes \widehat{CFD}(M, \alpha, \beta),$$

where $m = \partial D^2 \times \{\mathrm{pt}\}$. In this context, any choice of slopes l dual to m and β dual to α will do, since the family $(D^2 \times S^1, l + nm, m)$ are all homeomorphic as bordered manifolds. This is due to the Alexander trick; the Dehn twist along m in $\partial(D^2 \times S^1)$ extends to a homeomorphism of $D^2 \times S^1$.

Since every bordered manifold is equipped with a preferred choice of slopes, it will be important to distinguish between the two Dehn fillings along these slopes.

Definition 2.3 Let (M,α,β) be a bordered three-manifold. The standard filling is the Dehn filling
$$M(\alpha)=(D^2\times S^1,l,m)\cup(M,\alpha,\beta)$$
(that is, $m\mapsto\alpha$) and the dual filling is the Dehn filling
$$M(\beta)=(D^2\times S^1,m,l)\cup(M,\alpha,\beta)$$
(that is, $m\mapsto\beta$).

It is easy to compute $\widehat{HF}$ of these two fillings from $\widehat{CFD}(M,\alpha,\beta)$. There is a representative of $\widehat{CFA}(D^2\times S^1,\ell,m)$ which has a single generator x with idempotent ι_0 and operations of the form $m_{3+i}(x,\rho_3,\rho_{23},\dots,\rho_{23},\rho_2)=x$, where there are $i\ge 0$ copies of ρ_{23}. It follows that $\widehat{CF}(M(\alpha))$ is generated by the ι_0–generators of $\widehat{CFD}(M,\alpha,\beta)$ with a differential coming from any chain of the form

$$\bullet\xrightarrow{\rho_3}\circ\xrightarrow{\rho_{23}}\circ\ \cdots\ \circ\xrightarrow{\rho_{23}}\circ\xrightarrow{\rho_2}\bullet$$

Similarly, $\widehat{CFA}(D^2\times S^1,m,\ell)$ has a representative with a single generator y with idempotent ι_1 and operations of the form $m_{3+i}(y,\rho_2,\rho_{12},\dots,\rho_{12},\rho_1)=y$, where there are $i\ge 0$ copies of ρ_{12}. It follows that $\widehat{CF}(M(\alpha))$ is generated by the ι_1–generators of $\widehat{CFD}(M,\alpha,\beta)$ with a differential coming from any chain of the form

$$\circ\xrightarrow{\rho_2}\bullet\xrightarrow{\rho_{12}}\bullet\ \cdots\ \bullet\xrightarrow{\rho_{12}}\bullet\xrightarrow{\rho_1}\circ$$

More generally, given (M,α,β), we would like to compute $\widehat{HF}(M(\gamma))$ for any slope γ expressed in terms of α and β. In particular, we will always make a choice of orientations so that $\alpha\cdot\beta=+1$, resulting in slopes of the form $\gamma=\pm(p\alpha+q\beta)\in H_1(\partial M;\mathbb{Z})/\{\pm1\}$. As is familiar, the fixed choice $\{\alpha,\beta\}$ gives rise to an identification of the set of slopes and the extended rational numbers $\hat{\mathbb{Q}}:=\mathbb{Q}\cup\{\frac{1}{0}\}$. Our convention is that the slope $p\alpha+q\beta$ is identified with $p/q\in\hat{\mathbb{Q}}$. We will return to a detailed description of the pairing theorem for an arbitrary Dehn filling in the next section, since the following definition will be of central importance:

Definition 2.4 An L–space is a rational homology sphere Y for which $\dim\widehat{HF}(Y)=|H_1(Y;\mathbb{Z})|$. An L–space slope is a slope γ in ∂M for which the result of Dehn filling $M(\gamma)$ is an L–space. For any M with torus boundary, let $\mathcal{L}_M$ denote the set of L–space slopes in ∂M.

We will need to distinguish certain L–space slopes. To do this, consider the natural inclusion

$$\widehat{\mathbb{Q}} \hookrightarrow \widehat{\mathbb{R}} = \mathbb{R} \cup \left\{\tfrac{1}{0}\right\}$$

arising from orienting the basis slopes so that $\alpha \cdot \beta = +1$, and endow $\widehat{\mathbb{Q}}$ with the subspace topology. With this identification, $\mathcal{L}_M \subset \widehat{\mathbb{Q}}$.

Definition 2.5 The set of strict L–space slopes, denoted by $\mathcal{L}_M^\circ$, is the interior of the subset $\mathcal{L}_M$.

Recall that, if $a, b \in \widehat{\mathbb{Q}}$, then the subsets $(a, b) \cap \widehat{\mathbb{Q}}$ and $[a, b] \cap \widehat{\mathbb{Q}}$ are open and closed, respectively, in $\widehat{\mathbb{Q}}$. By abuse, we will write simply (a, b) and $[a, b]$ with the understanding that these describe subsets of $\widehat{\mathbb{Q}}$.

A key example to consider is that of the exterior of a nontrivial knot K in S^3, with $M = S^3 \setminus \nu K$. In this case it is well known that, if M admits a nontrivial L–space filling, then $\mathcal{L}_M$ is either $[2g-1, \infty]$ or $[\infty, 1-2g]$ relative to the preferred basis consisting of the knot meridian μ (corresponding to $1/0$) and the Seifert longitude λ (corresponding to 0), where g denotes the Seifert genus of K. Notice that μ and $(2g-1)\mu + \lambda$ are nonstrict L–space slopes by definition. On the other hand, if K is the trivial knot then $M \cong D^2 \times S^1$ and $\mathcal{L}_M = \mathcal{L}_M^\circ = \widehat{\mathbb{Q}} \setminus \{0\}$ since these are precisely the fillings that give lens spaces.

Every bordered manifold comes with a preferred identification of the set of slopes with $\widehat{\mathbb{Q}}$; in particular, the notation $p/q \in \mathcal{L}(M, \alpha, \beta)$ should be understood to mean the slope $\pm(p\alpha + q\beta) \in \mathcal{L}_M$. We will adhere to this convention, and use the two interchangeably where there is no potential for confusion.

2.6 Gluing via change of framing

In the interest of determining the set $\mathcal{L}_M$ we will need a means of describing any slope γ in ∂M in terms of a fixed basis of slopes $\{\alpha, \beta\}$. Suppose (M, α, β) is given; we would like to calculate $\widehat{HF}(M(p\alpha + q\beta))$. Then, according to the pairing theorem,

$$\widehat{CF}(M(p\alpha + q\beta)) \cong \widehat{CFA}(D^2 \times S^1, m, l) \boxtimes \widehat{CFD}(M, r\alpha + s\beta, p\alpha + q\beta),$$

where $\left(\begin{smallmatrix} q & p \\ s & r \end{smallmatrix}\right) \in \mathrm{SL}_2(\mathbb{Z})$. Notice that $p = 0$ recovers a chain complex for the dual filling.

Fixing a basis so that $\left(\begin{smallmatrix} 1 \\ 0 \end{smallmatrix}\right)$ represents the standard slope and $\left(\begin{smallmatrix} 0 \\ 1 \end{smallmatrix}\right)$ represents the dual slope, notice that the Dehn twist along α carrying $\beta \mapsto \alpha + \beta$ is encoded by the matrix $\left(\begin{smallmatrix} 1 & 1 \\ 0 & 1 \end{smallmatrix}\right)$. Call this the standard Dehn twist and denote it by T_{st}. Similarly, the Dehn twist along β carrying $\alpha \mapsto \alpha + \beta$ is encoded by the matrix $\left(\begin{smallmatrix} 1 & 0 \\ 1 & 1 \end{smallmatrix}\right)$. Call this the (negative) dual Dehn twist and denote it by T_{du}^{-1}. Associated with each Dehn twist is a mapping cylinder and to this Lipshitz, Ozsváth and Thurston assign a bimodule. Bimodules are defined similarly to type D and type A structures, except that there are two separate actions which may each be either type D or type A, so that bimodules can have type DD, DA, AD or AA; see [19]. For the Dehn twists $T_{\mathrm{st}}^{\pm 1}$ and $T_{\mathrm{du}}^{\pm 1}$, we use $\widehat{T}_{\mathrm{st}}^{\pm 1}$ and $\widehat{T}_{\mathrm{du}}^{\pm 1}$ to denote the corresponding type DA bimodules.[1]

We can take box tensor products of bimodules or of a module and a bimodule by pairing one type D action with one type A action. The convention is that type D actions are always left actions while type A actions are always right actions, and since the sidedness of the actions is unambiguous we will use the ordering of modules in a product to indicate which actions are involved. That is, for (bi)modules N and M we will understand $N \boxtimes M$ to mean the result of pairing the last action on N with the first action on M, even if the former is a left action and the latter is a right action. In practice this will be clear from the context.

Given an odd-length continued fraction expansion $p/q = [a_1, a_2, \dots, a_n]$, we obtain a decomposition according to Dehn twists,

$$\begin{pmatrix} q & p \\ s & r \end{pmatrix} = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}^{a_n} \cdots \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}^{a_2} \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}^{a_1}.$$

The bordered manifold $(M, r\alpha + s\beta, p\alpha + q\beta)$ is obtained from (M, α, β) by attaching the mapping cylinder of a homeomorphism h with representative $h_* = \left(\begin{smallmatrix} q & p \\ s & r \end{smallmatrix}\right)$, and so

$$\widehat{CFD}(M, r\alpha + s\beta, p\alpha + q\beta) \cong \widehat{T}_{\mathrm{st}}^{a_n} \boxtimes \cdots \boxtimes \widehat{T}_{\mathrm{du}}^{-a_2} \boxtimes \widehat{T}_{\mathrm{st}}^{a_1} \boxtimes \widehat{CFD}(M, \alpha, \beta),$$

where $\widehat{T}_{\mathrm{st}}^{n} = \underbrace{\widehat{T}_{\mathrm{st}} \boxtimes \cdots \boxtimes \widehat{T}_{\mathrm{st}}}_{n}$.

More generally, given bordered manifolds (M_1, α_1, β_1) and (M_2, α_2, β_2), we can calculate

$$\widehat{CFA}(M, \alpha_1, \beta_1) \boxtimes \widehat{CFD}(M, r\alpha_2 + s\beta_2, p\alpha_2 + q\beta_2)$$

by considering a homeomorphism h as above. This gives a complex computing $\widehat{HF}(Y)$, where

$$Y \cong M_1 \cup_h M_2 \cong (M_1, \alpha_1, \beta_1) \cup (M_2, r\alpha_2 + s\beta_2, p\alpha_2 + q\beta_2)$$

and $h \colon \partial M_1 \to \partial M_2$ is specified by

$$\alpha_1 \mapsto r\alpha_2 + s\beta_2, \quad \beta_1 \mapsto p\alpha_2 + q\beta_2.$$

[1] In the notation of [19, Section 10], we have $\widehat{T}_{\mathrm{st}} = \widehat{CFDA}(\tau_m)$ and $\widehat{T}_{\mathrm{du}} = \widehat{CFDA}(\tau_\ell^{-1})$.

With these gluing conventions in place, we make an observation that will be of use in the sequel:

Proposition 2.6 *Given type D modules N_1 and N_2, let $N_1' = \widehat{T}_{\mathrm{st}}^n \boxtimes N_1$ and $N_2' = \widehat{T}_{\mathrm{du}}^{-n} \boxtimes N_2$ for some $n \in \mathbb{Z}$. There is a homotopy equivalence*

$$N_1 \boxtimes \widehat{CFAA}(\mathbb{I}) \boxtimes N_2 \cong N_1' \boxtimes \widehat{CFAA}(\mathbb{I}) \boxtimes N_2',$$

where $\widehat{CFAA}(\mathbb{I})$ is the type AA identity bimodule.

Remark 2.7 The type AA identity bimodule is defined in [19] and, in particular, gives rise to $\widehat{CFA}(M,\alpha,\beta) \cong \widehat{CFAA}(\mathbb{I}) \boxtimes \widehat{CFD}(M,\alpha,\beta)$. This observation leads to the algorithm described in Section 2.4.

Proof of Proposition 2.6 The right-hand side of the claimed equivalence can be written as

$$(\widehat{T}_{\mathrm{st}}^n \boxtimes N_1) \boxtimes \widehat{CFAA}(\mathbb{I}) \boxtimes \widehat{T}_{\mathrm{du}}^{-n} \boxtimes N_2.$$

The DA–bimodule $\widehat{T}_{\mathrm{st}}$ is homotopy equivalent to the AD–bimodule

$$\widehat{CFAA}(\mathbb{I}) \boxtimes \widehat{T}_{\mathrm{du}} \boxtimes \widehat{CFDD}(\mathbb{I}).$$

To see this, note that the Heegaard diagram for $T_{\mathrm{st}} = \tau_m$ in [19, Figure 25] can be obtained from the Heegaard diagram for $T_{\mathrm{du}} = \tau_l$ by rotating 180 degrees. Thus the right side above simplifies to

$$\begin{aligned} N_1 \boxtimes (\widehat{CFAA}(\mathbb{I}) \boxtimes \widehat{T}_{\mathrm{du}} \boxtimes \widehat{CFDD}(\mathbb{I}))^n \boxtimes \widehat{CFAA}(\mathbb{I}) \boxtimes \widehat{T}_{\mathrm{du}}^{-n} \boxtimes N_2 & \\ \cong N_1 \boxtimes \widehat{CFAA}(\mathbb{I}) \boxtimes \widehat{T}_{\mathrm{du}}^{n} \boxtimes \widehat{T}_{\mathrm{du}}^{-n} \boxtimes N_2 & \\ \cong N_1 \boxtimes \widehat{CFAA}(\mathbb{I}) \boxtimes N_2. & \end{aligned}$$

□

Remark 2.8 If $N_1 = \widehat{CFD}(M_1,\alpha_1,\beta_1)$ and $N_2 = \widehat{CFD}(M_2,\alpha_2,\beta_2)$, the homotopy equivalence above corresponds to the fact that, with the gluing conventions described above,

$$(M_1,\alpha_1,\beta_1) \cup (M_2,\alpha_2,\beta_2) \cong (M_1,\alpha_1,\beta_1 + n\alpha_1) \cup (M_2,\alpha_2 + n\beta_2,\beta_2)$$

for any $n \in \mathbb{Z}$.

2.7 Grading

We conclude this discussion with a description of the relative $(\mathbb{Z}/2\mathbb{Z})$–grading on the bordered invariants, summarizing the discussion in [9, Section 2.2]. For more details and developments, see [15; 26]

The torus algebra $\mathcal{A}$ may be promoted to a graded algebra by defining $\mathrm{gr}(\rho_1) = \mathrm{gr}(\rho_3) = 0$ and $\mathrm{gr}(\rho_2) = 1$, and extending according to $\mathrm{gr}(\rho_I\rho_J) \equiv \mathrm{gr}(\rho_I) + \mathrm{gr}(\rho_J)$ reduced modulo 2 (we will always drop this reduction from the notation). In particular, the grading is 1 on all the remaining nonzero products ρ_{12}, ρ_{23} and ρ_{123}.

The relative $(\mathbb{Z}/2\mathbb{Z})$–grading on elements x of a type D structure is determined by

$$\mathrm{gr}(\rho_I \cdot x) \equiv \mathrm{gr}(\rho_I) + \mathrm{gr}(x) \quad \text{and} \quad \mathrm{gr}(\partial x) \equiv \mathrm{gr}(x) + 1.$$

In particular, given a connected $\mathcal{A}$–decorated graph this relative grading is determined by choosing the grading on a given vertex, and then noting that only edges labelled by ρ_1, ρ_3 or the identity element of $\mathcal{A}$ alter the grading.

The relative $(\mathbb{Z}/2\mathbb{Z})$–grading on elements x of a type A structure is determined by

$$\mathrm{gr}(m_{k+1}(x, \rho_{I_1}, \dots, \rho_{I_k})) - k - 1 \equiv \mathrm{gr}(x) + \sum_{j=1}^{k} \mathrm{gr}(\rho_{I_j}).$$

Notice that, given a type D structure with a choice of relative grading, a relative grading on the associated type A structure is obtained by switching the grading of each generator with idempotent ι_0.

If $Y \cong (M_1, \alpha_1, \beta_2) \cup (M_2, \alpha_2, \beta_2)$, choices of relative $(\mathbb{Z}/2\mathbb{Z})$–gradings on each of the objects $\widehat{CFA}(M_1, \alpha_1, \beta_2)$ and $\widehat{CFD}(M_2, \alpha_2, \beta_2)$ give rise to a relative grading on $\widehat{CF}(Y)$ via the pairing theorem and the rule $\mathrm{gr}(x \otimes y) = \mathrm{gr}(x) + \mathrm{gr}(y)$. This agrees with the usual relative $(\mathbb{Z}/2\mathbb{Z})$–grading on $\widehat{CF}(Y)$, so that, in particular, $|\chi(\widehat{CF}(Y))| \geq |H_1(Y;\mathbb{Z})|$ for any rational homology sphere Y. At the level of homology, it follows immediately that Y is an L–space if and only if every generator of $\widehat{HF}(Y)$ has the same grading. More generally, the rule $\mathrm{gr}(x \otimes y) = \mathrm{gr}(x) + \mathrm{gr}(y)$ determines the relative $(\mathbb{Z}/2\mathbb{Z})$–grading on tensor products involving bimodules.

Where required, we will make use of an additional marking on the vertices of an $\mathcal{A}$–decorated graph by $\{\bullet^+, \bullet^-, \circ^+, \circ^-\}$ to indicate a choice of relative $(\mathbb{Z}/2\mathbb{Z})$–grading on the underlying differential module. In addition, given a type D structure N with a choice of relative $(\mathbb{Z}/2\mathbb{Z})$–grading, we define the idempotent Euler characteristics

$$\chi_\bullet(N) = \chi(\iota_0 N) \quad \text{and} \quad \chi_\circ(N) = \chi(\iota_1 N).$$

Note that $\chi_\bullet$ and $\chi_\circ$ are well defined on homotopy classes of type D structures, since changing the basis of N leaves the vertex set of the corresponding graph unchanged and edge reduction removes a pair of $\bullet$–vertices or a pair of $\circ$–generators with opposite gradings.

The following lemma records how $\chi_\bullet$ and $\chi_\circ$ change under reparametrization of the boundary:

Lemma 2.9 *Let* $\left(\begin{smallmatrix} q & p \\ s & r \end{smallmatrix}\right) \in \mathrm{SL}_2(\mathbb{Z})$. *Then*

$$\pm\chi_\bullet\widehat{CFD}(M, r\alpha+s\beta, p\alpha+q\beta) = r\chi_\bullet\widehat{CFD}(M,\alpha,\beta) + s\chi_\circ\widehat{CFD}(M,\alpha,\beta),$$
$$\pm\chi_\circ\widehat{CFD}(M, r\alpha+s\beta, p\alpha+q\beta) = p\chi_\bullet\widehat{CFD}(M,\alpha,\beta) + q\chi_\circ\widehat{CFD}(M,\alpha,\beta).$$

Here the $\pm$ *choice depends on the choice of absolute* $(\mathbb{Z}/2\mathbb{Z})$*–grading on* $\widehat{CFD}(M,\alpha,\beta)$ *and* $\widehat{CFD}(M, r\alpha+s\beta, p\alpha+q\beta)$.

Proof This is true when $r=q=1$ and $s=p=0$. Assuming the claim is true for given r, s, p and q, we check that it is true when p and q are replaced by $p+r$ and $q+s$ by examining the effect of the standard Dehn twist. The relevant bimodule $\widehat{T}_{\mathrm{st}}$ is pictured in Figure 2. It has three generators: one pairs with ι_0–generators to produce ι_0–generators, one pairs with ι_1–generators to produce ι_1–generators, and one pairs with ι_0–generators to produce ι_1–generators. All three generators have the same relative $(\mathbb{Z}/2\mathbb{Z})$–grading; we will assume they have grading 0. It follows that

$$\begin{aligned}\chi_\bullet\widehat{CFD}(M, r\alpha+s\beta, (p+r)\alpha+(q+s)\beta) &= \chi_\bullet(\widehat{T}_{\mathrm{st}} \boxtimes \widehat{CFD}(M, r\alpha+s\beta, p\alpha+q\beta)) \\ &= \chi_\bullet\widehat{CFD}(M, r\alpha+s\beta, p\alpha+q\beta) \\ &= r\chi_\bullet\widehat{CFD}(M,\alpha,\beta) + s\chi_\circ\widehat{CFD}(M,\alpha,\beta),\end{aligned}$$

$$\begin{aligned}&\chi_\circ\widehat{CFD}(M, r\alpha+s\beta, (p+r)\alpha+(q+s)\beta) \\ &\quad= \chi_\circ(\widehat{T}_{\mathrm{st}} \boxtimes \widehat{CFD}(M, r\alpha+s\beta, p\alpha+q\beta)) \\ &\quad= \chi_\circ\widehat{CFD}(M, r\alpha+s\beta, p\alpha+q\beta) + \chi_\bullet\widehat{CFD}(M, r\alpha+s\beta, p\alpha+q\beta) \\ &\quad= (p+r)\chi_\bullet\widehat{CFD}(M,\alpha,\beta) + (q+s)\chi_\circ\widehat{CFD}(M,\alpha,\beta).\end{aligned}$$

Similarly, we check the claim for when r and s are replaced with $r-p$ and $s-q$ by considering the dual Dehn twist bimodule $\widehat{T}_{\mathrm{du}}$ (pictured in Figure 4). This bimodule has three generators with relative $(\mathbb{Z}/2\mathbb{Z})$–gradings 0, 0 and 1: the first pairs with ι_0–generators to produce ι_0–generators, the second pairs with ι_1–generators to produce ι_1–generators, and the third pairs with ι_1–generators to produce ι_0–generators. It follows that

$$\begin{aligned}&\chi_\bullet\widehat{CFD}(M, (r-p)\alpha+(s-q)\beta, p\alpha+q\beta) \\ &\quad= \chi_\bullet(\widehat{T}_{\mathrm{du}} \boxtimes \widehat{CFD}(M, r\alpha+s\beta, p\alpha+q\beta)) \\ &\quad= \chi_\bullet\widehat{CFD}(M, r\alpha+s\beta, p\alpha+q\beta) - \chi_\circ\widehat{CFD}(M, r\alpha+s\beta, p\alpha+q\beta) \\ &\quad= (r-p)\chi_\bullet\widehat{CFD}(M,\alpha,\beta) + (s-q)\chi_\circ\widehat{CFD}(M,\alpha,\beta),\end{aligned}$$

$$\begin{aligned}\chi_\circ\widehat{CFD}(M,(r-p)\alpha+(s-q)\beta,p\alpha+q\beta) &= \chi_\circ(\widehat{T}_{\mathrm{du}}\boxtimes\widehat{CFD}(M,r\alpha+s\beta,p\alpha+q\beta))\\ &= \chi_\circ\widehat{CFD}(M,r\alpha+s\beta,p\alpha+q\beta)\\ &= p\chi_\circ\widehat{CFD}(M,\alpha,\beta)+q\chi_\circ\widehat{CFD}(M,\alpha,\beta).\end{aligned}$$

The argument extends easily to multiples and inverses of Dehn twists. The claim follows by decomposing an arbitrary element of $\mathrm{SL}_2(\mathbb{Z})$ into Dehn twists. □

For rational homology spheres, we will be interested in a distinguished slope, referred to as the *rational longitude*. Given a bordered structure (α,β) for M, this slope $\gamma=p\alpha+q\beta$ is characterized by the property that some collection of like-oriented copies of γ bounds a surface in M. This slope is recorded by the fraction p/q, and the $(\mathbb{Z}/2\mathbb{Z})$–grading allows us to recover this data from $\widehat{CFD}(M,\alpha,\beta)$ as follows:

Proposition 2.10 *Fix a spin^c–structure $\mathfrak{s}$ on a bordered manifold (M,α,β), and let $\chi_\bullet$ and $\chi_\circ$ denote the Euler characteristics of $\widehat{CFD}(M,\alpha,\beta;\mathfrak{s})$ in the appropriate idempotent. If $\chi_\bullet=\chi_\circ=0$, then M is not a rational homology solid torus. Otherwise, M is a rational homology solid torus, the nullhomologous curves in ∂M are the multiples of $\chi_\circ\alpha-\chi_\bullet\beta$, and, in particular, the rational longitude is $-\chi_\circ/\chi_\bullet$.*

Proof Let $\mathfrak{s}'$ be a spin^c–structure on $M(p\alpha+q\beta)$ that restricts to $\mathfrak{s}$ on M. The Euler characteristic $\chi\widehat{HF}(M(p\alpha+q\beta);\mathfrak{s}')$ is nonzero if and only if $M(p\alpha+q\beta)$ is a rational homology sphere [24, Proposition 5.1]. The closed manifold $M(p\alpha+q\beta)$ is the dual filling of $(M,r\alpha+s\beta,p\alpha+q\beta)$ for any r and s with $rq-ps=1$. Thus, up to sign, $\chi\widehat{HF}(M(p\alpha+q\beta);\mathfrak{s}')$ is

$$\chi_\circ\widehat{CFD}(M,r\alpha+s\beta,p\alpha+q\beta;\mathfrak{s}')=p\chi_\bullet+q\chi_\circ$$

by applying Lemma 2.9. Note that if $\chi_\bullet$ and $\chi_\circ$ are both zero, then $M(p\alpha+q\beta)$ is not a rational homology sphere for any p/q and hence M is not a rational homology solid torus. Otherwise, $M(p\alpha+q\beta)$ is a rational homology sphere unless $p/q=-\chi_\circ/\chi_\bullet\in\widehat{\mathbb{Q}}$. It follows that when M is a rational homology solid torus the slope of the rational longitude is $-\chi_\circ/\chi_\bullet$.

For the remainder of the proof, let $p/q=-\chi_\circ/\chi_\bullet$, where p and q are relatively prime. The multicurve $np\alpha+nq\beta$ is nullhomologous in M for some $n>0$. To determine the smallest such n consider the standard filling of $(M,r\alpha+s\beta,p\alpha+q\beta)$ for some r and s with $rq-ps=1$. This manifold is obtained from M by adding a solid torus identifying the longitude l with $p\alpha+q\beta$; the resulting first homology group is a direct sum of $H_1(M,\partial M)$ and $H_1(D^2\times S^1)$ (which is generated by the longitude l) with the relation $nl = np\alpha + nq\beta = 0$. Since $H_1(M, \partial M)$ indexes the spin^c–structures on M, for each spin^c structure $\mathfrak{s}$ on M, there are n elements of $H_1(M(r\alpha + s\beta))$. It follows that $\chi\widehat{HF}(M(r\alpha + s\beta); \mathfrak{s}') = n$ for each $\mathfrak{s}'$. By Lemma 2.9, we have

$$n = \chi_\bullet \widehat{CFD}(M, r\alpha + s\beta, p\alpha + q\beta; \mathfrak{s}') = r\chi_\bullet + s\chi_\circ,$$

which may be combined with the fact that $n = nqr - nps$. So, up to sign, $(\chi_\bullet, \chi_\circ) = (nq, -np)$ and hence the minimal nullhomologous multicurve in ∂M is $\chi_\circ\alpha - \chi_\bullet\beta$, as claimed. $\square$

Remark 2.11 The minimal integer n appearing in this proof, multiplied by the number of spin^c structures, corresponds to the constant c_M associated with M described in [29, Section 3.1].

3 Loop calculus

A focus of this work is the development of a calculus for studying bordered invariants. This will be achieved by restricting to a class of manifolds whose bordered invariants can be represented by certain valence two $\mathcal{A}$–decorated graphs, which we will call loops.

3.1 Loops and loop-type manifolds

Towards defining loops, consider the following arrows which may appear in an $\mathcal{A}$–decorated graph:

$$\begin{array}{cccc}
\bullet\xrightarrow{\rho_1} & \bullet\xrightarrow{\rho_3} & \circ\xrightarrow{\rho_2} & \circ\xleftarrow{\rho_3} \\
\bullet\xrightarrow{\rho_{12}} & \bullet\xleftarrow{\rho_2} & \circ\xrightarrow{\rho_{23}} & \circ\xleftarrow{\rho_{23}} \\
\bullet\xrightarrow{\rho_{123}} & \bullet\xleftarrow{\rho_{12}} & \circ\xleftarrow{\rho_1} & \circ\xleftarrow{\rho_{123}} \\
\mathbf{I}_\bullet & \mathbf{II}_\bullet & \mathbf{I}_\circ & \mathbf{II}_\circ
\end{array}$$

We will be interested in valence two $\mathcal{A}$–decorated graphs subject to the following restriction:

($\star$) Each ι_0–vertex is adjacent to an edge of type $\mathbf{I}_\bullet$ and an edge of type $\mathbf{II}_\bullet$, and each ι_1–vertex is adjacent to an edge of type $\mathbf{I}_\circ$ and an edge of type $\mathbf{II}_\circ$.

Definition 3.1 A loop is a connected valence two $\mathcal{A}$–decorated graph satisfying ($\star$).

Since any loop describes a differential module over $\mathcal{A}$, a loop may be promoted to a graded loop via the $(\mathbb{Z}/2\mathbb{Z})$–grading described in Section 2.7. In particular, where needed, the vertex set will be extended to $\{\bullet^+, \bullet^-, \circ^+, \circ^-\}$.

As abstract combinatorial objects loops provide a tractable structure to work with; this section develops a calculus for doing so. The results derived from this calculus apply to the following class of bordered invariants:

Definition 3.2 The bordered invariant $\widehat{CFD}(M,\alpha,\beta)$ is said to be of loop-type if, up to homotopy, it may be represented by a collection of loops, that is, by a (possibly disconnected) $\mathcal{A}$–decorated valence two graph satisfying ($\star$). For simplicity, in this paper we will make the additional assumption that the number of connected components of the valence two graph describing $\widehat{CFD}(M,\alpha,\beta)$ coincides with $|\mathrm{Spin}^c(M)|$.

We will refer to the bordered manifold (M,α,β) being of loop-type when the associated bordered invariant has this property. Some motivation for this definition is provided by the following:

Proposition 3.3 *Let $\mathcal{H}$ be a bordered Heegaard diagram describing (M,α,β) and suppose $\widehat{CFD}(M,\alpha,\beta)=\widehat{CFD}(\mathcal{H})$ is reduced and represented by a valence two $\mathcal{A}$–decorated graph having a single connected component per* spin^c*–structure. Then $\widehat{CFD}(M,\alpha,\beta)$ is loop-type.*

Proof The hypothesis $\widehat{CFD}(M,\alpha,\beta)=\widehat{CFD}(\mathcal{H})$ allows us to use the notion of generalized coefficient maps developed in [20, Section 11.6], which force restrictions on the type D modules that can occur as invariants of manifolds with torus boundary. Briefly, generalized coefficient maps are extra differentials obtained by counting holomorphic curves that run over the basepoint. The torus algebra is extended with additional Reeb chords: ρ_0, ρ_{01}, ρ_{30}, ρ_{012}, ρ_{301}, ρ_{230}, ρ_{0123}, ρ_{3012}, ρ_{2301} and ρ_{1230}. With these additional differentials, $\partial^2=0$ is no longer satisfied. However, we have instead that

$$\partial^2(x)=\rho_{1230}x+\rho_{3012}x \quad \text{for any } x \text{ in } \iota_0\widehat{CFD}(M,\alpha,\beta),$$
$$\partial^2(x)=\rho_{2301}x+\rho_{1230}x \quad \text{for any } x \text{ in } \iota_1\widehat{CFD}(M,\alpha,\beta).$$

Let x be a generator with idempotent ι_0. Since $\partial^2(x)$ contains the term $\rho_{1230}x$, $\partial(x)$ contains the term $\rho_{I_1}y$ and $\partial(y)$ contains the term $\rho_{I_2}x$ for some $y\in\widehat{CFD}(M,\alpha,\beta)$ and some (generalized) Reeb chords ρ_{I_1} and ρ_{I_2} such that $\rho_{I_1}\rho_{I_2}=\rho_{1230}$. Since we assumed that $\widehat{CFD}(M,\alpha,\beta)$ is reduced, I_1 and I_2 are not $\varnothing$. It follows that $I_1\in\{1,12,123\}$ and, in the graphical representation, the vertex corresponding to x has an incident edge of type $\mathbf{I}_\bullet$.

Since $\partial^2(x)$ contains the term $\rho_{3012}x$, $\partial(x)$ contains the term $\rho_{I_1}y$ and $\partial(y)$ contains the term $\rho_{I_2}x$ for some $y\in\widehat{CFD}(M,\alpha,\beta)$ and some (generalized) Reeb chords ρ_{I_1} and ρ_{I_2} such that $\rho_{I_1}\rho_{I_2}=\rho_{3012}$. It follows that either $I_1=3$ or $I_2\in\{2,12\}$ and, in the graphical representation, the vertex corresponding to x has an incident edge of type $\mathbf{II}_\bullet$. Since any vertex has valence two by assumption, the vertex corresponding to x must have exactly one edge from each of $\mathbf{I}_\bullet$ and $\mathbf{II}_\bullet$.

The argument when x has idempotent ι_1 is similar. Since $\partial^2(x)$ contains the term $\rho_{2301}x$, it follows that the vertex corresponding to x is adjacent an edge of type $\mathbf{I}_\circ$; and since $\partial^2(x)$ contains the term $\rho_{0123}x$, it follows that the corresponding vertex is adjacent to an edge of type $\mathbf{II}_\circ$. □

Remark 3.4 While this argument required a particular choice of Heegaard diagram, it seems likely that this hypothesis is not a necessary. However, for the purposes of this work the more general statement will not be required — we simply restrict to the graphs satisfying ($\star$) by definition — and leave developing the necessary algebra to future work.

It is important to note that loops in the abstract need not be related to three-manifolds: it is not necessarily true that every loop (or disjoint union of loops) arises as the bordered invariant of some three-manifold with torus boundary. This section concludes with an explicit example for illustration.

3.2 Standard and dual notation

It is natural to decompose a loop into pieces by breaking along vertices corresponding to one of the two idempotents; the constraint ($\star$) suggests that the pieces resulting from such a decomposition are quite limited. Indeed, breaking along generators with idempotent ι_0, five essentially different chains are possible and these possibilities are listed in Figure 1, left. Since ($\star$) also puts restrictions on how these segments can be concatenated, each piece is depicted with puzzle-piece ends. For instance, a type a piece can be followed by a type b piece, but not by a type c piece. Any given piece may also appear backwards; we denote this with a *bar*.

Segments of type a and b behave differently from the other segment types in several ways. We will call type a and b segments *stable chains* and type c, d and e segments *unstable chains*. This terminology comes from [20, Theorem 11.26], where the notion of unstable chains was introduced in order to describe a procedure for extracting the bordered invariants from knot Floer homology. Three types of unstable chains were described, which correspond precisely to our type c, d and e segments. Type a and b segments also appeared as the segments coming from horizontal arrows and vertical arrows, respectively, in knot Floer homology. The motivation for this terminology is

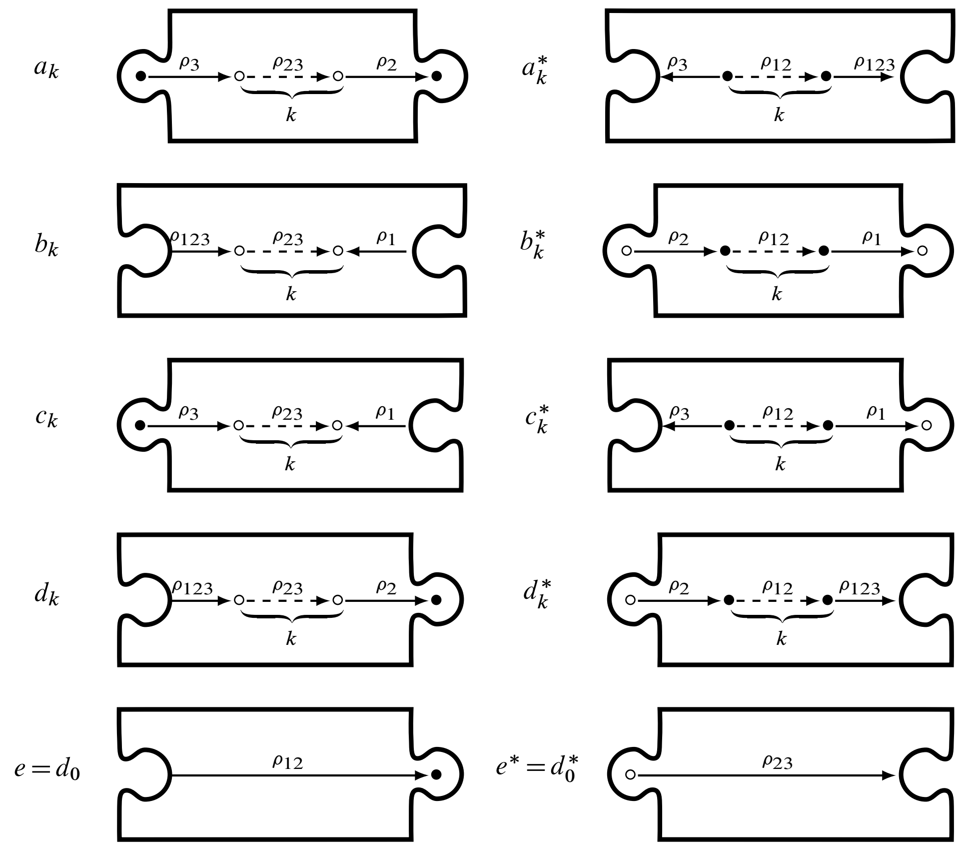

Figure 1: Possible segments: Standard notation (left) is obtained by breaking a loop along ι_0 idempotents ($\bullet$–vertices) and dual notation (right) is obtained by breaking a loop along ι_1 idempotents ($\circ$–vertices). Note that the integer k records the number of interior vertices, so that the ρ_{23} (for standard notation) or ρ_{12} (for dual notation) appears $k-1$ times.

that, as we will see in Section 3.4, stable chains are preserved by certain Dehn twists while unstable chains are not.

Denote the standard alphabet by

$$\mathfrak{A} = \{a_i, b_i, c_j, d_j\},$$

where $i \in \mathbb{Z} \smallsetminus \{0\}$ and $j \in \mathbb{Z}$. The letters in the standard alphabet correspond directly to the segments depicted in Figure 1, with the relationships

$$a_{-i} = \bar{a}_i \quad \text{and} \quad b_{-i} = \bar{b}_i$$

for $i > 0$ as well as

$$c_{-j} = \bar{d}_j \quad \text{and} \quad d_{-j} = \bar{c}_j$$

for $j \geq 0$ with $d_0 = e = \bar{c}_0$ and $\bar{d}_0 = \bar{e} = c_0$. Throughout this paper we will always assume these shorthand relationships and, by abuse, we will often not distinguish between segments and letters. For instance, the symbols d_0 and e will be used interchangeably and may refer either to a letter in the standard alphabet or the corresponding $\mathcal{A}$–decorated graph segment. As a result of this equivalence, the notation $\bar{s}$ makes sense for any standard letter s, with the understanding that $\bar{\bar{s}} = s$.

We will be interested in the set of cyclic words $W_{\mathfrak{A}}$ on the alphabet $\mathfrak{A}$ that are consistent with the puzzle-piece notation of Figure 1, which encodes the restriction ($\star$). Thus, for example, the cyclic word c_1 is an element of $W_{\mathfrak{A}}$ while the cyclic word a_1 is not. Another immediate restriction is that an element of $W_{\mathfrak{A}}$ must contain an equal number of a–type letters and b–type letters.

Proposition 3.5 *A loop $\boldsymbol{\ell}$ with at least one ι_0–generator may be represented as a cyclic word $\boldsymbol{w_\ell}$ in $W_{\mathfrak{A}}$. Moreover, this representation is unique up to overall reversal of the word $\boldsymbol{w_\ell}$, that is, writing $\boldsymbol{w_\ell}$ with the opposite cyclic ordering and replacing each letter s with $\bar{s}$.*

Proof This is immediate from the definitions. □

As a result, we will be interested in the equivalence class that identifies $\boldsymbol{w_\ell}$ and $\bar{\boldsymbol{w}}_{\boldsymbol{\ell}}$, where $\bar{\boldsymbol{w}}_{\boldsymbol{\ell}}$ denotes the reversal of $\boldsymbol{w_\ell}$. Denote this equivalence class of cyclic words by $(\boldsymbol{w_\ell})$; by abuse we will continue referring to $(\boldsymbol{w_\ell})$ as a cyclic word. By Proposition 3.5, this sets up a one-to-one correspondence $\boldsymbol{\ell} \leftrightarrow (\boldsymbol{w_\ell})$ between loops with at least one ι_0–decorated vertex and (equivalence classes of) cyclic words in $W_{\mathfrak{A}}$.

Definition 3.6 When $\boldsymbol{\ell}$ is represented by a cyclic word $(\boldsymbol{w_\ell})$ using the standard alphabet, we say $\boldsymbol{\ell}$ is written in standard notation.

A loop cannot be written in standard notation if it does not contain an ι_0–decorated vertex. This suggests it will sometimes be useful to break a loop along ι_1–decorated vertices. There are again five types of chains possible, as listed in Figure 1, right. As before, the first two types will be referred to as stable chains and the rest as unstable chains. Denote the dual alphabet by

$$\mathfrak{A}^* = \{a_i^*, b_i^*, c_j^*, d_j^*\}$$

($i \in \mathbb{Z} \smallsetminus \{0\}$ and $j \in \mathbb{Z}$), where, as before, $a^*_{-i} = \bar{a}^*_i$, $b^*_{-i} = \bar{b}^*_i$ (for $i > 0$), $c^*_{-j} = \bar{d}^*_j$, $d^*_{-j} = \bar{c}^*_j$ (for $j \geq 0$) and $d^*_0 = e^* = \bar{c}^*_0$. Proceeding as before, let $W_{\mathfrak{A}^*}$ denote the set of cyclic words in the alphabet $\mathfrak{A}^*$ which are consistent with the puzzle-piece notation.

Proposition 3.7 *A loop $\boldsymbol{\ell}$ with at least one ι_1–generator may be represented as a cyclic word $\boldsymbol{w}^*_{\boldsymbol{\ell}}$ in $W_{\mathfrak{A}^*}$. Moreover, this representation is unique up to overall reversal of the word $\boldsymbol{w}^*_{\boldsymbol{\ell}}$.*

Proof This is immediate from the definitions. □

Again, we denote by $(\boldsymbol{w}^*_{\boldsymbol{\ell}})$ the equivalence class identifying $\boldsymbol{w}^*_{\boldsymbol{\ell}}$ and $\bar{\boldsymbol{w}}^*_{\boldsymbol{\ell}}$, and remark that this gives rise to a second one-to-one correspondence $\boldsymbol{\ell} \leftrightarrow (\boldsymbol{w}^*_{\boldsymbol{\ell}})$ for loops $\boldsymbol{\ell}$ with at least one ι_1–decorated vertex.

Definition 3.8 When $\boldsymbol{\ell}$ is represented by a cyclic word $(\boldsymbol{w}^*_{\boldsymbol{\ell}})$ using the dual alphabet, we say $\boldsymbol{\ell}$ is written in dual notation.

Collecting these observations, notice that, whenever $\boldsymbol{\ell}$ contains instances of both idempotents, the pair of correspondences $(\boldsymbol{w}^*_{\boldsymbol{\ell}}) \leftrightarrow \boldsymbol{\ell} \leftrightarrow (\boldsymbol{w}_{\boldsymbol{\ell}})$ sets up a natural map between the standard representation and the dual representation. In particular, this allows us to set $(\boldsymbol{w}_{\boldsymbol{\ell}})^* = (\boldsymbol{w}^*_{\boldsymbol{\ell}})$ in a well-defined way; we refer to $(\boldsymbol{w}_{\boldsymbol{\ell}})^*$ as the dual of $(\boldsymbol{w}_{\boldsymbol{\ell}})$ and remark that $(\boldsymbol{w}_{\boldsymbol{\ell}})^{**} = (\boldsymbol{w}_{\boldsymbol{\ell}})$.

This description of loops as cyclic words is the essential starting point for loop calculus. A typical loop will have vertices in both idempotents, so it is expressible in both standard and dual notation; switching between the two will be a key part of the loop calculus. We now make this process explicit.

First notice that, given $(\boldsymbol{w}_{\boldsymbol{\ell}})$ representing a loop $\boldsymbol{\ell}$ in standard notation, there is a natural normal form $(\boldsymbol{w}_{\boldsymbol{\ell}}) = (\boldsymbol{u}_1 \boldsymbol{v}_1 \boldsymbol{u}_2 \boldsymbol{v}_2 \dots \boldsymbol{u}_n \boldsymbol{v}_n)$, where

(N1) the subword $\boldsymbol{u}_i$ is a standard letter with subscript $k_i \neq 0$; and

(N2) the subword $\boldsymbol{v}_i$ is $n_i \geq 0$ consecutive copies of either d_0 or c_0 (here n_i may be zero).

This normal form makes sense for any $\boldsymbol{w}_{\boldsymbol{\ell}}$ that does not consist only of $d_0 = e$ or $c_0 = \bar{e}$ letters; we may safely ignore these sporadic examples since, in these cases, $\boldsymbol{\ell}$ has no ι_1–decorated vertices and cannot be expressed in dual notation.

Now the dual word $\boldsymbol{w}^*_{\boldsymbol{\ell}}$ is obtained as follows:

(D1) Replace each $\boldsymbol{v}_i$ by a dual letter with subscript $n_i + 1$ and type[2] determined by (the types of) the ordered pair $\{\boldsymbol{u}_i, \boldsymbol{u}_{i+1}\}$ (written $\boldsymbol{u}_i\boldsymbol{u}_{i+1}$ for brevity) from the subword $\boldsymbol{u}_i\boldsymbol{v}_i\boldsymbol{u}_{i+1}$ according to the rules

$$\begin{aligned} \bar{a}b, \bar{a}d, \bar{c}b, \bar{c}d &\to a^*, & \bar{b}a, \bar{b}c, \bar{d}a, \bar{d}c &\to \bar{a}^*, \\ a\bar{b}, a\bar{c}, d\bar{b}, d\bar{c} &\to b^*, & b\bar{a}, b\bar{d}, c\bar{a}, c\bar{d} &\to \bar{b}^*, \\ \bar{a}\bar{b}, \bar{a}\bar{c}, \bar{c}\bar{b}, \bar{c}\bar{c} &\to c^*, & ba, bc, ca, cc &\to \bar{c}^*, \\ ab, ad, db, dd &\to d^*, & \bar{b}\bar{a}, \bar{b}\bar{d}, \bar{d}\bar{a}, \bar{d}\bar{d} &\to \bar{d}^*. \end{aligned}$$

Note that indices are taken mod n, so $\boldsymbol{u}_{n+1}$ is identified with $\boldsymbol{u}_1$.

(D2) Replace each $\boldsymbol{u}_i$ with the subword consisting of $k_i - 1$ consecutive d_0^* letters if $k_i > 0$ or $1 - k_i$ consecutive c_0^* letters if $k_i < 0$.

To convert from dual notation to standard notation, we use exactly the same procedure (interchanging the words *standard* and *dual* and adding/removing stars where appropriate in the discussion above). Note that letters in dual (resp. standard) notation correspond to consecutive pairs of letters in standard (resp. dual) notation after we ignore letters with subscript 0. We make the following observation about this correspondence:

Observation 3.9 *Stable chains in dual* (*resp. standard*) *notation correspond precisely to consecutive pairs of letters in standard* (*resp. dual*) *notation*, *ignoring letters with subscript* 0, *whose subscripts have opposite signs.*

With the forgoing in place we will not distinguish between loops in standard or in dual notation; for a given loop $\boldsymbol{\ell}$ containing (vertices decorated by) both idempotents, it is always possible to choose a representative for $\boldsymbol{\ell}$ in either of the alphabets $\mathfrak{A}$ or $\mathfrak{A}^*$. In summary: We will regard a loop as both a graph-theoretic object and as an equivalence class of words $(\boldsymbol{w}_{\boldsymbol{\ell}})$ and $(\boldsymbol{w}_{\boldsymbol{\ell}}^*)$ modulo dualizing. In particular, we will adopt the abuse of notation that $\boldsymbol{\ell}$ is such an equivalence class of words, where convenient, and think of loops as graph-theoretic objects and word-theoretic objects interchangeably. In particular, we will let the notation $\boldsymbol{\ell}$ stand in for a choice of cyclic word representative (in either alphabet).

We will often refer to a subword $\boldsymbol{w}$ of a given loop $\boldsymbol{\ell}$, so the notation $(\boldsymbol{w})$ indicating the cyclic closure of a word (when it exists) will be used to distinguish subwords from

[2] Recall that a letter of type a, b, c or d with negative subscript can also be written as a letter of type $\bar{a}$, $\bar{b}$, $\bar{d}$ or $\bar{c}$ with positive subscript. Here the *type* of a standard letter (other than d_0 and c_0) refers to the element of $\{a, \bar{a}, b, \bar{b}, c, \bar{c}, d, \bar{d}\}$ corresponding to the representation with positive subscript. Similarly, a dual letter other than d_0^* and c_0^* has a well-defined type in $\{a^*, \bar{a}^*, b^*, \bar{b}^*, c^*, \bar{c}^*, d^*, \bar{d}^*\}$.

(equivalence classes of) cyclic words. The length of a subword is the number of pieces (or letters) in the word.

To conclude with a particular example, let M denote the complement of the left-hand trefoil and consider the bordered manifold (M, μ, λ), where μ is the knot meridian and λ is the Seifert longitude. Following [20, Chapter 11], $\widehat{CFD}(M, \mu, \lambda)$ is described by a loop, as shown in

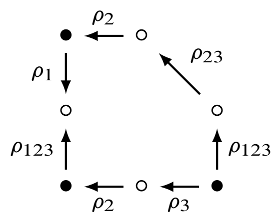

We may express this in standard notation as $(a_1 b_1 \bar{d}_2)$ or in dual notation as $(a_1^* e^* b_1^* \bar{d}_1^*)$. In particular, according to the discussion above, this pair of words is a representative of the same equivalence specifying the loop $\boldsymbol{\ell}$ shown. That is, we regard $\boldsymbol{\ell} \sim (a_1 b_1 \bar{d}_2) \sim (a_1^* e^* b_1^* \bar{d}_1^*)$. More generally, the homotopy type of the invariant $\widehat{CFD}(M, \mu, \lambda - (2+n)\mu)$ is represented by the loop $(a_1 b_1 c_n)$ for $n \in \mathbb{Z}$, following the convention that $c_0 = \bar{e}$ and $c_n = \bar{d}_{-n}$ when $n < 0$.

3.3 Operations on loops

We now define two abstract operations on loops: T and H. These operations are easy to describe for loops in terms of standard and dual notation, respectively, and we will see in the next subsection that they correspond to important bordered Floer operations.

If a loop $\boldsymbol{\ell}$ cannot be written in standard notation (that is, it is a collection of only e^* segments), then $\mathrm{T}(\boldsymbol{\ell}) = \boldsymbol{\ell}$. Otherwise, express $\boldsymbol{\ell}$ in standard notation and consider the operation T determined on individual letters via

$$\mathrm{T}(a_i) = a_i, \quad \mathrm{T}(b_i) = b_i, \quad \mathrm{T}(c_j) = c_{j-1}, \quad \mathrm{T}(d_j) = d_{j+1}$$

for any $i \in \mathbb{Z} \smallsetminus \{0\}$ and $j \in \mathbb{Z}$. For collections of loops, we also define $\mathrm{T}(\{\boldsymbol{\ell}_i\}_{i=1}^n) = \{\mathrm{T}(\boldsymbol{\ell}_i)\}_{i=1}^n$. The operation T is invertible; denote the inverse by T^{-1}.

Note that the letters a_i and b_i are fixed by T. As a result, these are sometimes referred to as stable chains (or standard stable chains). The (standard) unstable chains are the letters c_j and d_j. This is consistent with the notion of stable and unstable chains from [20].

The operator H is defined similarly, but with respect to dual notation. If a loop ℓ cannot be written in dual notation, then $\mathrm{H}(\ell) = \ell$. Otherwise, express ℓ in dual notation and consider the operation H on ℓ defined on individual dual letters via

$$\mathrm{H}(a_i^*) = a_i^*, \quad \mathrm{H}(b_i^*) = b_i^*, \quad \mathrm{H}(c_j^*) = c_{j-1}^*, \quad \mathrm{H}(d_j^*) = d_{j+1}^*$$

for any $i \in \mathbb{Z} \setminus \{0\}$ and $j \in \mathbb{Z}$. For collections of loops, we also define $\mathrm{H}(\{\ell_i\}_{i=1}^n) = \{\mathrm{H}(\ell_i)\}_{i=1}^n$. As with T, H is an invertible operation with inverse H^{-1}.

Note that a_i^* and b_i^* are fixed by H. When we have need for it, a^*– and b^*–type letters will be referred to as dual stable chains, while c^*– and d^*–type letters will be referred to as dual unstable chains.

The operation T is easy to define for loops in standard notation, but it would be difficult to describe purely in terms of dual letters. Similarly, the operation H is simple to define in dual notation, but would be complicated in terms of standard letters. This suggests the need to comfortably switch between the two; in particular, given a loop ℓ expressed in standard notation, finding $\mathrm{H}(\ell)$ in standard notation is a three-step process: dualize, apply H and dualize again.

We will often need to apply combinations of the operations T and H to a loop. The composition $\mathrm{T} \circ \mathrm{H}^{-1} \circ \mathrm{T}$, in particular, appears often; it will be convenient to regard this composition as another loop operation, which we call E. The following lemma describes the effect of E on a cyclic word in standard notation; this provides a convenient shortcut compared with computing the operations T, H^{-1} and T individually. We state the lemma for general loops, but we will only prove it in a special case.

Lemma 3.10 *If ℓ is written in standard notation, then $\mathrm{E}(\ell)$ is determined by the following action of* E *on standard letters*:

$$\mathrm{E}(a_k) = a_{-k}^*, \quad \mathrm{E}(b_k) = b_{-k}^*, \quad \mathrm{E}(c_k) = c_{-k}^*, \quad \mathrm{E}(d_k) = d_{-k}^*.$$

Proof We give the proof in the special case that ℓ consists only of d_k segments with $k \geq 0$. The general proof is left to the reader.

Let $\ell = (d_{k_1} d_{k_2} \dots d_{k_n})$, with $k_i \geq 0$. Then $\mathrm{T}(\ell) = (d_{k_1+1} \dots d_{k_n+1})$. Writing this in dual notation and applying H^{-1}, we have

$$\mathrm{T}(\ell) = (d_1^* \underbrace{d_0^* \dots d_0^*}_{k_1} d_1^* \underbrace{d_0^* \dots d_0^*}_{k_2} \dots d_1^* \underbrace{d_0^* \dots d_0^*}_{k_n}),$$

$$\mathrm{H}^{-1} \circ \mathrm{T}(\ell) = (d_0^* \underbrace{d_{-1}^* \dots d_{-1}^*}_{k_1} d_0^* \underbrace{d_{-1}^* \dots d_{-1}^*}_{k_2} \dots d_0^* \underbrace{d_{-1}^* \dots d_{-1}^*}_{k_n}).$$

Dualizing and twisting again gives

$$\begin{aligned}
\mathrm{H}^{-1} \circ \mathrm{T}(\boldsymbol{\ell}) &= (c_2 \underbrace{c_1 \dots c_1}_{k_1-1} c_2 \underbrace{c_1 \dots c_1}_{k_2-1} \dots c_2 \underbrace{c_1 \dots c_1}_{k_n-1}),\\
\mathrm{T} \circ \mathrm{H}^{-1} \circ \mathrm{T}(\boldsymbol{\ell}) &= (c_1 \underbrace{c_0 \dots c_0}_{k_1-1} c_1 \underbrace{c_0 \dots c_0}_{k_2-1} \dots c_1 \underbrace{c_0 \dots c_0}_{k_n-1}),\\
\mathrm{T} \circ \mathrm{H}^{-1} \circ \mathrm{T}(\boldsymbol{\ell}) &= (d^*_{-k_1} d^*_{-k_2} \dots d^*_{-k_n}). \qquad \square
\end{aligned}$$

As suggested by the example of the left-hand trefoil exterior given at the end of the last subsection, expressing loops as cyclic words gives rise to clean way of describing a type D structure. This extends to reparametrization: consulting [20, Chapter 11], when M is the exterior of the left-hand trefoil, $\widehat{CFD}(M, \mu, \lambda - (2+n)\mu)$ is represented by the loop $(a_1 b_1 c_n)$ for $n \in \mathbb{Z}$. Motivating the next subsection, this loop may be expressed as $\mathrm{T}^{-n}(\boldsymbol{\ell})$, where $\boldsymbol{\ell} = (a_1 b_1 c_0)$.

3.4 Dehn twists

The operators T and H naturally encode the effect of a Dehn twist on a loop representing the type D structure of a bordered manifold. Recall that

$$\widehat{T}_{\mathrm{st}}^{\pm 1} \boxtimes \widehat{CFD}(M, \alpha, \beta) \cong \widehat{CFD}(M, \alpha, \beta \pm \alpha)$$

and, more generally, for any type D structure N over $\mathcal{A}$ the pair of type D structures $\widehat{T}_{\mathrm{st}}^{\pm 1} \boxtimes N$ are well defined.

Proposition 3.11 *If $\boldsymbol{\ell}$ is a loop with corresponding type D structure $N_{\boldsymbol{\ell}}$, then $N_{\boldsymbol{\ell}}^{\pm} = \widehat{T}_{\mathrm{st}}^{\pm 1} \boxtimes N_{\boldsymbol{\ell}}$ is a loop-type module represented by the loop $\mathrm{T}^{\pm}(\boldsymbol{\ell})$.*

Proof We compute the box tensor product $\widehat{T}_{\mathrm{st}} \boxtimes N_{\boldsymbol{\ell}}$ by considering the effect on one segment at a time in a standard representative for $\boldsymbol{\ell}$.

Recall from [19, Section 10] (see Figure 2) that the type DA bimodule $\widehat{T}_{\mathrm{st}}$ has three generators p, q and r (denoted by $\bullet$, $\circ$ and $*$, respectively, in Figure 2) with idempotents determined by $p = \iota_0 p \iota_0$, $q = \iota_1 q \iota_1$ and $r = \iota_1 r \iota_0$. Thus each generator $x \in \iota_1 N_{\boldsymbol{\ell}}$ gives rise to a single generator $q \otimes x \in \widehat{T}_{\mathrm{st}} \boxtimes N_{\boldsymbol{\ell}}$, and each generator $x \in \iota_0 N_{\boldsymbol{\ell}}$ gives rise to two generators $p \otimes x$ and $r \otimes x$ in the box tensor product. There is always a ρ_2 arrow from $r \otimes x$ to $p \otimes x$; that is, $\partial(r \otimes x)$ has a summand $\rho_2 \cdot (p \otimes x)$.

If generators $x, y \in \iota_0 N_{\boldsymbol{\ell}}$ are connected by a single segment s (in the loop $\boldsymbol{\ell}$), we will consider the portion of a loop representing $N_{\boldsymbol{\ell}}^{+}$ between $p \otimes x$ and $p \otimes y$ and show

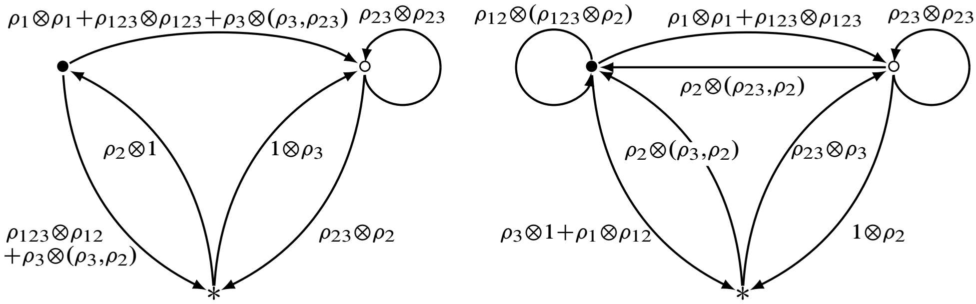

Figure 2: Graphical representations of the Dehn twist bimodules $\widehat{T}_{\text{st}}$ (left) and $\widehat{T}_{\text{st}}^{-1}$ (right), following [19, Section 10].

that (up to homotopy equivalence) it is the segment $\mathrm{T}(s)$. To talk about the portion *between*, we need a (cyclic) ordering on the elements of $N_{\boldsymbol{\ell}}^+$. This is inherited from a choice of cyclic ordering on the elements of $\boldsymbol{\ell}$, together with a specified order of $p\otimes x$ and $r\otimes x$ for each $x\in\iota_0 N_{\boldsymbol{\ell}}$. If there is a segment s from x to y, we say that $r\otimes x$ is between $p\otimes x$ and $p\otimes y$ if the puzzle piece shape of s on the x end agrees with the shape of a_k, that is, if s is a_k, $\bar{a}_k$, c_k, $\bar{d}_k$ or $\bar{e}$.

Consider first a segment a_k from x to y, where x and y are generators of $\iota_0 N_{\boldsymbol{\ell}}$. The cases of $k=1$ and $k\geq 2$ are slightly different; both are pictured in Figure 3. In either case, the effect of tensoring the segment with $\widehat{T}_{\text{st}}$ is pictured. There is a differential starting at $r\otimes x$; after cancelling this differential, the result is a segment of type $a_k=\mathrm{T}(a_k)$ from $p\otimes x$ to $p\otimes y$.

Consider next a segment b_k from x to y. Tensoring the segment with $\widehat{T}_{\text{st}}$, we see that the portion between $p\otimes x$ and $p\otimes y$ is simply a segment of type $b_k=\mathrm{T}(b_k)$. Note that the generators $r\otimes x$ and $r\otimes y$ are not included in this new segment; they must be included in the segments on either side.

If x and y are connected by a segment c_k, then there is a differential starting at $r\otimes x$. If $k>1$, then cancelling this differential leaves a segment from $p\otimes x$ to $p\otimes y$ of type $c_{k-1}=\mathrm{T}(c_k)$. If $k=1$, then cancelling the differential produces a new ρ_{12} arrow, and thus there is a segment $\bar{e}=\mathrm{T}(c_1)$ from $p\otimes x$ to $p\otimes y$. If x and y are connected by segments d_k or e, we see in Figure 3 that the portion of $N_{\boldsymbol{\ell}}^+$ from $x\otimes p$ to $y\otimes p$ is a segment $d_{k+1}=\mathrm{T}(d_k)$ or $d_1=\mathrm{T}(e)$, respectively.

Segments with the opposite orientation behave the same way. A segment of $\bar{s}$ from x to y is the same as a segment s from y to x. In the tensor product, this produces a segment $\mathrm{T}(s)$ from $p\otimes y$ to $p\otimes x$, or a segment $\overline{\mathrm{T}(s)}=\mathrm{T}(\bar{s})$ from $p\otimes x$ to $p\otimes y$.

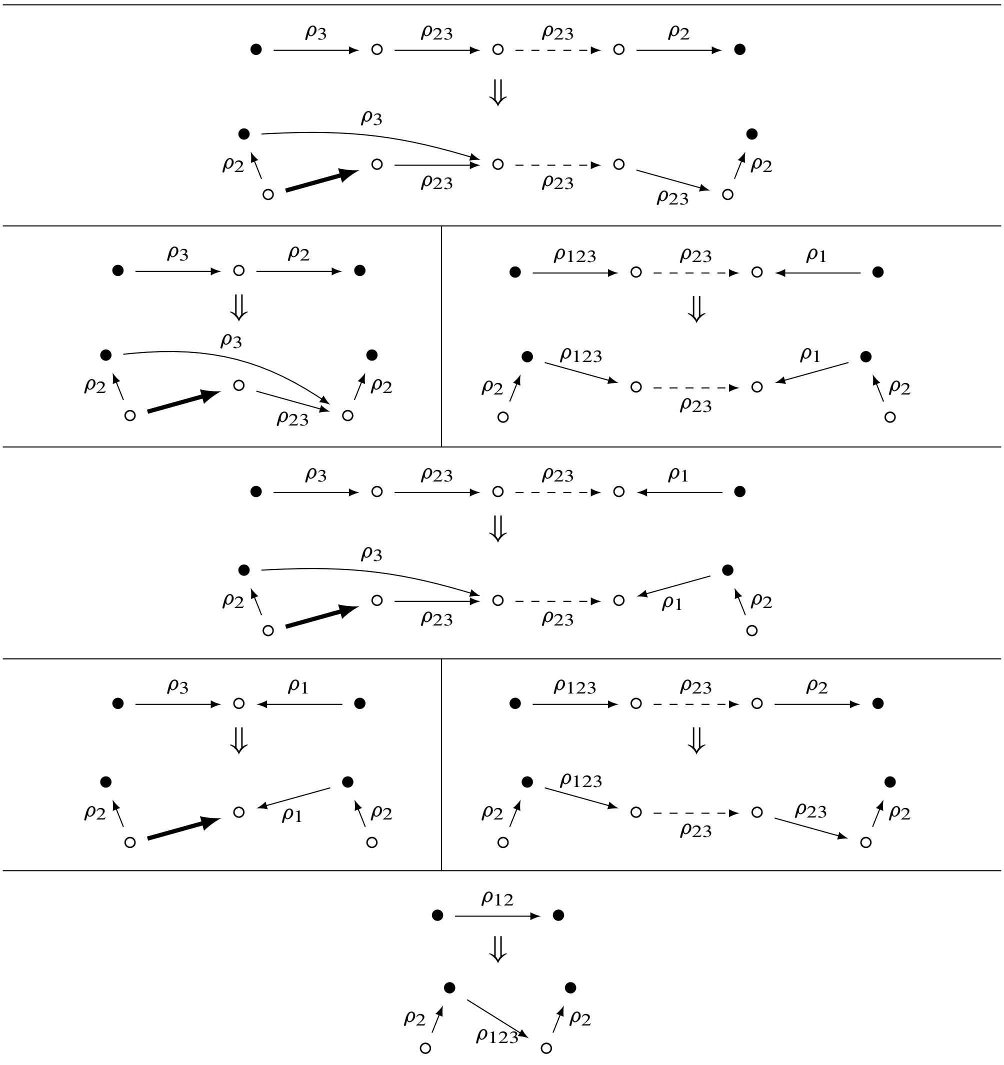

Figure 3: Illustrating the proof of Proposition 3.11: The effect of box tensoring with $\widehat{T}_{\mathrm{st}}$ on each of the possible segments occurring in a loop expressed in standard notation. Unmarked edges, which are eliminated using edge reduction described in Section 2.3, are highlighted.

It remains to check that a loop consisting only of e^* segments represents a type D structure N that is fixed by the action of the standard Dehn twist. This is easy to see, since in this case the only relevant operation in $\widehat{T}_{\mathrm{st}}$ is

$$m_2(q, \rho_{23}) = \rho_{23} \cdot q.$$

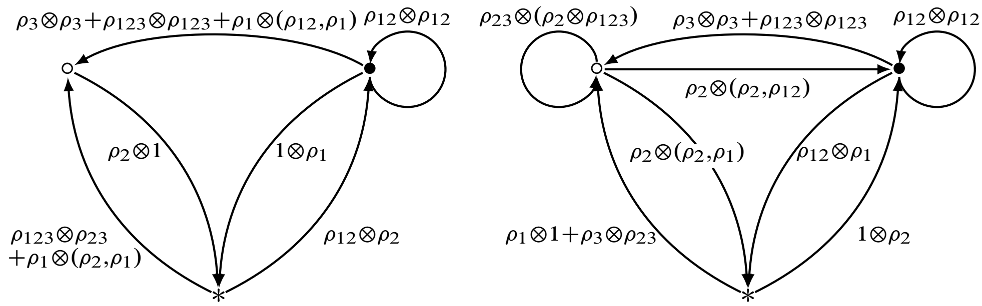

Figure 4: Graphical representations of the Dehn twist bimodules $\widehat{T}_{\mathrm{du}}^{-1}$ (left) and $\widehat{T}_{\mathrm{du}}$ (right), following [19, Section 10].

Each generator of $\iota_1 N = N$ gives rise to one generator of N^+ and each ρ_{23} arrow in N gives a ρ_{23} arrow in N^+.

The case of $\widehat{T}_{\mathrm{st}}^{-1}$ can be deduced from the case of $\widehat{T}_{\mathrm{st}}$. Let N' be a type D module represented by the loop $\mathrm{T}^{-1}(\boldsymbol{\ell})$. We have just shown

$$\widehat{T}_{\mathrm{st}} \boxtimes N' \cong N$$

and it follows that

$$\widehat{T}_{\mathrm{st}}^{-1} \boxtimes N \cong \widehat{T}_{\mathrm{st}}^{-1} \boxtimes \widehat{T}_{\mathrm{st}} \boxtimes N' \cong N'$$

since $\widehat{T}_{\mathrm{st}}^{-1} \boxtimes \widehat{T}_{\mathrm{st}}$ is homotopy equivalent to the identity bimodule. □

Proposition 3.12 *If $\boldsymbol{\ell}$ is a loop with corresponding type D structure $N_{\boldsymbol{\ell}}$, then $N_{\boldsymbol{\ell}}^{\mp} = \widehat{T}_{\mathrm{du}}^{\mp 1} \boxtimes N_{\boldsymbol{\ell}}$ is a loop-type module represented by the loop $\mathrm{H}^{\pm}(\boldsymbol{\ell})$.*

Proof The proof is similar, with the relevant bimodules $\widehat{T}_{\mathrm{du}}$ and $\widehat{T}_{\mathrm{du}}^{-1}$ shown in Figure 4. The result of tensoring $\widehat{T}_{\mathrm{du}}$ with each type of dual segment is shown in Figure 5. We see that a^* and b^* segments are fixed, c_k^* segments become c_{k+1}^* segments, and d_k^* segments become d_{k-1}^* segments. In other words, tensoring $\widehat{T}_{\mathrm{du}}$ with a dual segment $\boldsymbol{s}^*$ gives $\mathrm{H}^{-1}(\boldsymbol{s}^*)$. Thus, for a loop $\boldsymbol{\ell}$, $\widehat{T}_{\mathrm{du}} \boxtimes \boldsymbol{\ell}$ is the loop $\mathrm{H}^{-1}(\boldsymbol{\ell})$. Since $\widehat{T}_{\mathrm{du}}^{-1}$ is the inverse of $\widehat{T}_{\mathrm{du}}$, we conclude that $\widehat{T}_{\mathrm{du}}^{-1} \boxtimes \boldsymbol{\ell}$ is the loop $\mathrm{H}(\boldsymbol{\ell})$. □

We conclude this discussion by observing that the notion of a manifold M (with torus boundary) being of loop-type is now well defined. In particular, since any reparametrization of a loop gives rise to a loop, it follows that the property of loop-type (or having loop-type bordered invariants) is independent of the bordered structure and hence a property of the underlying (unbordered) manifold. Indeed:

Definition 3.13 A compact, orientable, connected three-manifold M with torus boundary is loop-type if $\widehat{CFD}(M, \alpha, \beta)$ is of loop-type for any choice of basis slopes α and β.

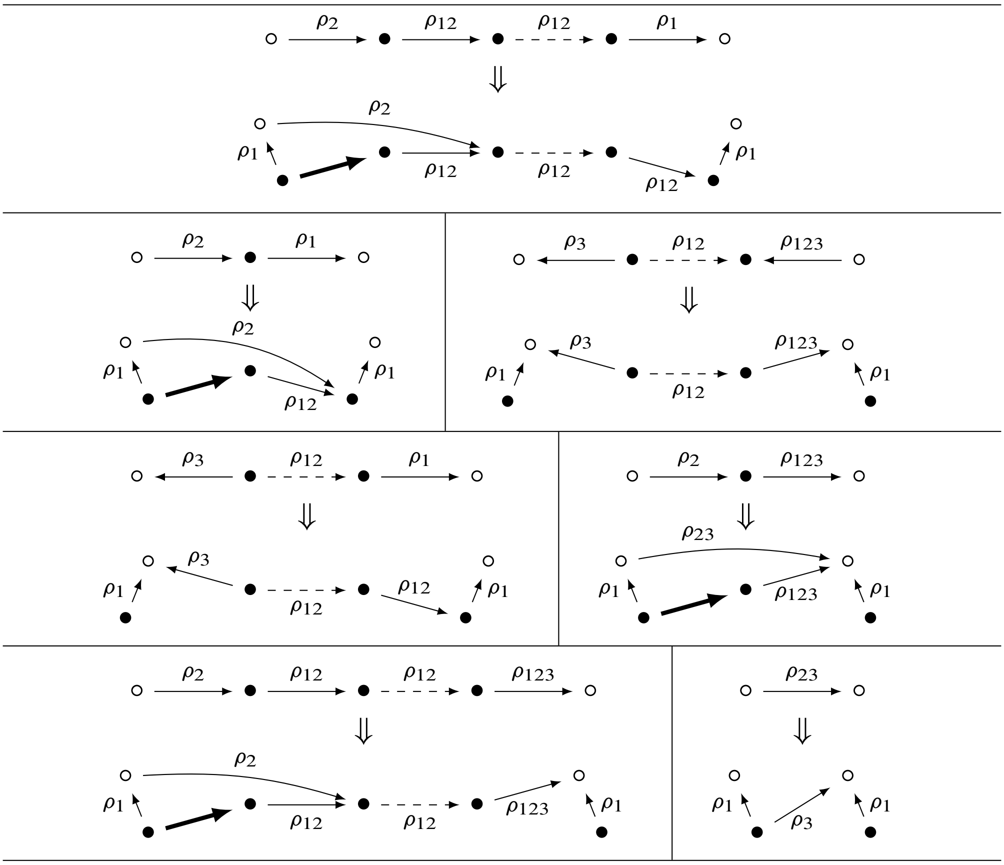

Figure 5: Illustrating the proof of Proposition 3.12: The effect of the Dehn twist $\widehat{T}_{\mathrm{du}}$ on each of the possible segments occurring in a loop expressed in dual notation. Unmarked edges, which are eliminated using edge reduction described in Section 2.3, are highlighted.

3.5 Solid tori

As a simple example of the loop operations described above, we now describe the computation of $\widehat{CFD}$ of a solid torus with arbitrary framing. By a p/q–framed solid torus, we will mean a bordered solid torus $(D^2 \times S^1, \alpha, \beta)$ such that the meridian is $p\alpha + q\beta$. Recall that for the standard (0–framed) and dual (∞–framed) solid tori we have

$$\widehat{CFD}(D^2 \times S^1, l, m) = (e), \quad \widehat{CFD}(D^2 \times S^1, m, l) = (e^*).$$

The bordered invariants for solid tori with other framings can be computed by applying Dehn twists as described in Section 2.6.

Example 3.14 We compute $\widehat{CFD}$ of the $\frac{7}{2}$–framed solid torus using the continued fraction expansion $-\frac{2}{7} = [-1, 2, -2, 3]$:

$$\begin{array}{rccc}
\widehat{CFD}(D^2\times S^1, l, m): & (e) & & \\
& \Downarrow_{\mathrm{T}^3} & & \\
\widehat{CFD}(D^2\times S^1, l, m+3l): & (d_3) & \sim & (d_1^* d_0^* d_0^*) \\
& & & \Downarrow_{\mathrm{H}^{-2}} \\
\widehat{CFD}(D^2\times S^1, -2m-5l, m+3l): & (c_1c_1c_0c_1c_0) & \sim & (d_{-1}^* d_{-2}^* d_{-2}^*) \\
& \Downarrow_{\mathrm{T}^2} & & \\
\widehat{CFD}(D^2\times S^1, -2m-5l, -3m-7l): & (c_{-1}c_{-1}c_{-2}c_{-1}c_{-2}) & \sim & (c_{-1}^* c_{-1}^* c_{-1}^* c_0^* c_{-1}^* c_{-1}^* c_0^*) \\
& & & \Downarrow_{\mathrm{H}^{-1}} \\
\widehat{CFD}(D^2\times S^1, m+2l, -3m-7l): & (d_{-4}d_{-3}) & \sim & (c_0^* c_0^* c_0^* c_1^* c_0^* c_0^* c_1^*)
\end{array}$$

It is easy to check $m = 7(m+2l) + 2(-3m-7l)$, and so $(D^2 \times S^1, m+2l, -3m-7l)$ is a $\frac{7}{2}$–framed solid torus. Note that given $\widehat{CFD}$ of a solid torus we can also check the framing by using Proposition 2.10. If we choose the $(\mathbb{Z}/2\mathbb{Z})$–grading so that endpoints of d_k chains in standard notation have grading 0 and endpoints of c_k chains have grading 1, it is not difficult to see that the Euler characteristic $\chi_\bullet$ of a loop in standard notation is the number of d_k segments minus the number of c_k segments and $\chi_\circ$ is given by the sum of the subscripts. Thus $(D^2 \times S^1, m+2l, -3m-7l)$ has rational longitude $-\chi_\circ/\chi_\bullet = -\frac{-7}{2}$.

$\widehat{CFD}$ for arbitrarily framed solid tori can be computed by a similar procedure. The result is always a loop of a particularly simple form.

Lemma 3.15 *If $q \neq 0$, then $\widehat{CFD}(D^2 \times S^1, pm+ql, rm+sl)$ can be represented by a single loop $\ell = (d_{k_1} d_{k_2} \dots d_{k_m})$. Moreover, the difference between $\max\{k_i\}$ and $\min\{k_i\}$ is at most 1.*

Proof If $p = 0$ (this implies that $q = r = 1$) then

$$\widehat{CFD}(D^2 \times S^1, l, m+sl) = \widehat{T}_{\mathrm{st}}^s \boxtimes \widehat{CFD}(D^2 \times S^1, l, m) \sim \mathrm{T}^s((e)) = (d_s).$$

Otherwise, let $[a_1, \dots, a_{2n}]$ be an even length continued fraction for q/p and choose a_{2n+1} so that $[a_1, \dots, a_{2n+1}]$ is a continued fraction for s/r. Then

$$\begin{aligned}\widehat{CFD}(D^2 \times S^1, pm+ql, rm+sl) &\cong \widehat{T}_{\mathrm{st}}^{a_{2n+1}} \boxtimes \cdots \boxtimes \widehat{T}_{\mathrm{du}}^{-a_2} \boxtimes \widehat{T}_{\mathrm{st}}^{a_1} \boxtimes \widehat{CFD}(M, l, m) \\ &\sim \mathrm{T}^{a_{2n+1}} \circ \mathrm{H}^{a_{2n}} \circ \cdots \circ \mathrm{H}^{a_2} \circ \mathrm{T}^{a_1}((d_0)).\end{aligned}$$

If p/q is positive, we may assume that $a_i > 0$ for all $1 \le i \le 2n$. We will show by induction on n that $\mathrm{H}^{a_{2n}} \circ \cdots \circ \mathrm{H}^{a_2} \circ \mathrm{T}^{a_1}((d_0))$ is a loop in standard notation consisting only of d_0 and d_1 chains. The base case of $n = 0$ is immediate. Assuming that $\mathrm{H}^{a_{2n-2}} \circ \cdots \circ \mathrm{H}^{a_2} \circ \mathrm{T}^{a_1}((d_0))$ is a loop in standard notation consisting only of d_0 and d_1 chains, applying the twist $\mathrm{T}^{a_{2n-1}}$ produces a loop consisting of d_k chains with only positive subscripts. In dual notation, such a loop involves only d_0^* and d_1^* segments. Applying the twist $\mathrm{H}^{a_{2n}}$ produces a loop consisting of d_k^* segments with positive subscripts. Switching to standard notation, this loop contains only d_0 and d_1 segments. Similarly, if p/q is negative, we assume that $a_i < 0$ for $1 \le i \le n$ and observe by induction that $\mathrm{H}^{a_{2n}} \circ \cdots \circ \mathrm{H}^{a_2} \circ \mathrm{T}^{a_1}((d_0))$ is a loop in standard notation consisting only of d_0 and d_{-1} chains.

Finally, applying the twist $\mathrm{T}^{a_{2n+1}}$ preserves the fact that the loop consists of type d_k unstable chains. It also preserves the relative differences of the subscripts, so the difference between the maximum and minimum subscript remains at most one. $\square$

3.6 Abstract fillings and abstract slopes

Recall that a loop $\boldsymbol{\ell}$ represents a type D structure, which by slight abuse of notation we denote by $\boldsymbol{\ell}$. Since $\boldsymbol{\ell}$ is reduced, there is an associated type A structure as described in Section 2.4, which we denote by $\boldsymbol{\ell}^A$. As a result, given loops $\boldsymbol{\ell}_1$ and $\boldsymbol{\ell}_2$ we can represent the chain complex produced by the box tensor product of the associated modules by $\boldsymbol{\ell}_1^A \boxtimes \boldsymbol{\ell}_2$. Since loops are connected $\mathcal{A}$–decorated graphs, the type D and type A structures associated to a loop have a well-defined relative $(\mathbb{Z}/2\mathbb{Z})$–grading, as described in Section 2.7. The gradings on $\boldsymbol{\ell}_1^A$ and $\boldsymbol{\ell}_2$ induce a relative $(\mathbb{Z}/2\mathbb{Z})$–grading on $\boldsymbol{\ell}_1^A \boxtimes \boldsymbol{\ell}_2$.

Consider the loops $\boldsymbol{\ell}_\bullet = (e)$ and $\boldsymbol{\ell}_\circ = (e^*)$. Given any loop $\boldsymbol{\ell}$, we have a pair of chain complexes

$$C_{\boldsymbol{\ell}}\left(\tfrac{1}{0}\right) = \boldsymbol{\ell}_\bullet^A \boxtimes \boldsymbol{\ell} \quad \text{and} \quad C_{\boldsymbol{\ell}}(0) = \boldsymbol{\ell}_\circ^A \boxtimes \boldsymbol{\ell},$$

which, noting that $\boldsymbol{\ell}_\bullet$ and $\boldsymbol{\ell}_\circ$ represent type D structures of the standard and dual (bordered) solid torus, respectively, might be regarded as abstract standard and dual Dehn filings of $\boldsymbol{\ell}$, respectively.

Remark 3.16 We do not need to assume $\boldsymbol{\ell}$ is bounded, since, if it is not, it is homotopy equivalent to a modified loop which is bounded and has the same box tensor product with $\boldsymbol{\ell}_\bullet^A$. Such a modified loop can be obtained by replacing either a ρ_{12} arrow with the homotopy equivalent (but not reduced) sequence

$$\bullet \xrightarrow{\rho_1} \circ \longleftarrow \circ \xrightarrow{\rho_2} \bullet$$

or a ρ_{123} arrow with

$$\bullet \xrightarrow{\rho_1} \circ \longleftarrow \circ \xrightarrow{\rho_{23}} \bullet$$

If $\boldsymbol{\ell}$ cannot be written in standard notation — that is, it is a collection of e^* segments — then $H_*(\boldsymbol{\ell}_\bullet^A \boxtimes \boldsymbol{\ell})$ must have two generators of opposite grading. To see this, replace one e^* segment with

$$\circ \xrightarrow{\rho_2} \bullet \longleftarrow \bullet \xrightarrow{\rho_3} \circ$$

to produce a bounded modified loop.

The standard filling picks out the $\bullet$–idempotent, in practice, and adds a differential for each type a_k chain. The dual filling picks out the $\circ$–idempotent and adds a differential for each b_k^* chain. For instance, when M is the exterior of the left-hand trefoil and $\widehat{CFD}(M,\mu,\lambda)$ is represented by $\boldsymbol{\ell} \sim (a_1 b_1 \bar{d}_2) \sim (a_1^* e^* b_1^* \bar{d}_1^*)$, the resulting complexes $\widehat{CF}(S^3) \cong \widehat{CF}(M(\mu)) \cong \boldsymbol{\ell}_\bullet^A \boxtimes \boldsymbol{\ell}$ and $\widehat{CF}(M(\lambda)) \cong \boldsymbol{\ell}_\circ^A \boxtimes \boldsymbol{\ell}$ are

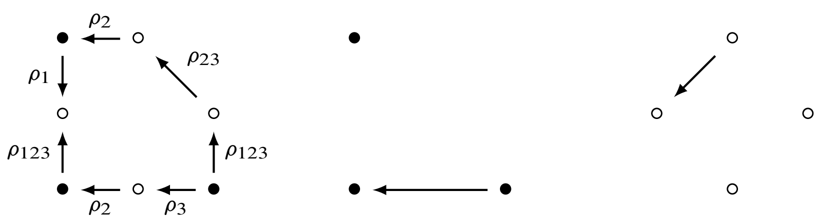

We regard the chain complexes $C_{\boldsymbol{\ell}}\left(\frac{1}{0}\right)$ and $C_{\boldsymbol{\ell}}(0)$ as the result of abstract Dehn fillings along a pair of abstract slopes in $\boldsymbol{\ell}$, which we identify with $\infty = \frac{1}{0}$ (corresponding to the standard filling) and 0 (corresponding to the dual filling). In fact, a given loop $\boldsymbol{\ell}$ gives rise to a natural $\widehat{\mathbb{Q}}$–family of chain complexes: choosing an even-length continued fraction $p/q = [a_1, a_2, \dots, a_n]$, let $\boldsymbol{\ell}^{p/q} = \mathrm{H}^{a_n} \circ \cdots \circ \mathrm{T}^{a_3} \circ \mathrm{H}^{a_2} \circ \mathrm{T}^{a_1}(\boldsymbol{\ell})$ and define

$$C_{\boldsymbol{\ell}}\left(\frac{p}{q}\right) = \boldsymbol{\ell}_\bullet^A \boxtimes \boldsymbol{\ell}^{p/q}.$$

We will regard the complex $C_{\boldsymbol{\ell}}(p/q)$ as an abstract Dehn filling of the loop $\boldsymbol{\ell}$ along the abstract slope $p/q \in \widehat{\mathbb{Q}}$.

The reason for this definition is illustrated as follows:

Proposition 3.17 *If a given loop $\boldsymbol{\ell}$ represents $\widehat{CFD}(M,\alpha,\beta)$ for some bordered manifold (M,α,β), the chain complex $C_{\boldsymbol{\ell}}(p/q)$ is (homotopy equivalent to) the chain complex $\widehat{CF}(M(p\alpha+q\beta))$. That is, abstract filling along abstract slopes corresponds to Dehn filling along slopes whenever $\boldsymbol{\ell}$ describes the type D structure of a bordered three-manifold.*

Proof This follows immediately from the definitions; note that $\boldsymbol{\ell}_{p/q}$ represents $\widehat{CFD}(M, p\alpha+q\beta, r\alpha+s\beta)$, where $p/q = [a_1, a_2, \dots, a_n]$ is an even length continued fraction and $r/s = [a_1, a_2, \dots, a_{n-1}]$. $\square$

Fixing a relative $(\mathbb{Z}/2\mathbb{Z})$–grading for a loop $\boldsymbol{\ell}$, consider the idempotent Euler characteristics $\chi_\bullet(\boldsymbol{\ell})$ and $\chi_\circ(\boldsymbol{\ell})$. If $\chi_\bullet(\boldsymbol{\ell})$ and $\chi_\circ(\boldsymbol{\ell})$ are not both 0, then $\boldsymbol{\ell}$ has a preferred slope $-\chi_\circ(\boldsymbol{\ell})/\chi_\bullet(\boldsymbol{\ell})$; we call this slope the *abstract rational longitude*. By Proposition 2.10, if $\boldsymbol{\ell}$ represents $\widehat{CFD}(M, \alpha, \beta)$ for a rational homology solid torus (M, α, β), then the abstract rational longitude for $\boldsymbol{\ell}$ is the rational longitude for (M, α, β).

Recall that a slope p/q is an L–space slope for a bordered manifold (M, α, β) if Dehn filling along the curve $p\alpha + q\beta$ yields an L–space. We end this section by defining a similar notion of L–space slopes for loops.

Definition 3.18 Given a loop $\boldsymbol{\ell}$, we say an abstract slope p/q in $\widehat{\mathbb{Q}}$ is an L–space slope for $\boldsymbol{\ell}$ if the relatively $(\mathbb{Z}/2\mathbb{Z})$–graded chain complex $C_{\boldsymbol{\ell}}(p/q)$ is an L–space chain complex in the sense that $\dim H_*(C_{\boldsymbol{\ell}}(p/q)) = |\chi(C_{\boldsymbol{\ell}}(p/q))| \neq 0$.

Remark 3.19 With these notions in place, we will now drop the modifier *abstract* when treating loops despite the fact that a given loop may or may not describe the type D structure of a three-manifold. In particular, we will not make a distinction between slopes and abstract slopes in the sequel.

By considering loops in the abstract, a particular class of loops is singled out.

Definition 3.20 A loop $\boldsymbol{\ell}$ is solid torus-like if it may be obtained from the loop $(ee\cdots e)$ via applications of $\mathrm{T}^{\pm 1}$ and $\mathrm{H}^{\pm 1}$.

Note that $\chi_\bullet(ee\cdots e)$ counts the number of e segments appearing and $\chi_\circ(ee\cdots e) = 0$ identifies the rational longitude of a solid torus-like loop. In particular, $\boldsymbol{\ell}_\bullet$ (representing $\widehat{CFD}(D^2 \times S^1, l, m)$) is solid torus-like. Justifying the chosen terminology, we have the following behaviour:

Proposition 3.21 *If $\boldsymbol{\ell}$ is solid torus-like with $\chi_\circ(\boldsymbol{\ell}) = 0$, then $\boldsymbol{\ell}^A \boxtimes \boldsymbol{\ell}' \cong \bigoplus_{\chi_\bullet(\boldsymbol{\ell})} \boldsymbol{\ell}_\bullet^A \boxtimes \boldsymbol{\ell}'$ for any loop $\boldsymbol{\ell}'$.*

Proof Recall that $\boldsymbol{\ell}_\bullet^A$ has a single generator x and operations

$$m_{3+i}(x, \rho_3, \underbrace{\rho_{23}, \dots, \rho_{23}}_{i \text{ times}}, \rho_2) = x,$$

so that, for generators $x \otimes u, x \otimes v \in \boldsymbol{\ell}_{\bullet}^{A} \boxtimes \boldsymbol{\ell}'$, with $x \otimes v$ a summand of $\partial(x \otimes u)$, there must be

$$\delta^{i+2}(u) = \rho_3 \otimes \underbrace{\rho_{23} \otimes \cdots \otimes \rho_{23}}_{i \text{ times}} \otimes \rho_2 \otimes v$$

in the type D structure for $\boldsymbol{\ell}'$. Now consider the type A structure described by $\boldsymbol{\ell}^{A}$. This has $n = \chi_{\bullet}(\boldsymbol{\ell})$–generators $x_0, \ldots, x_{n-1}$ and operations

$$m_{3+i}(x_j, \rho_3, \underbrace{\rho_{23}, \ldots, \rho_{23}}_{i \text{ times}}, \rho_2) = x_{i+j+1},$$

where the subscripts are understood to be reduced modulo n. (Note that the cyclic ordering on the generators is determined by $m_3(x_j, \rho_3, \rho_2) = x_{j+1}$.) So, given generators u and v in the type D structure for $\boldsymbol{\ell}'$ as above, $x_{i+j+1} \otimes v$ is in the image of $\partial(x_j \otimes u)$ for each $j \in \{0, \ldots, n-1\}$. This achieves the desired splitting. $\square$

Corollary 3.22 *Suppose $\boldsymbol{\ell}$ is solid torus-like with $\chi_{\circ}(\boldsymbol{\ell}) = 0$. Then $\boldsymbol{\ell}^{A} \boxtimes \boldsymbol{\ell}'$ is an L–space complex if and only if ∞ is an L–space slope for the loop $\boldsymbol{\ell}'$.*

As a result, with respect to gluing, solid torus-like loops will need to be treated like solid tori. Consequently, manifolds with solid torus-like invariants (should they exist) will need to be singled out.

Definition 3.23 A loop-type manifold M is solid torus-like if it is a rational homology solid torus and every loop in the representation of $\widehat{CFD}(M, \alpha, \beta)$ is solid torus-like.

We conclude this section with an explicit example, demonstrating that not every abstract loops arises as the type D structure of a bordered three-manifold. Consider the loop $\boldsymbol{\ell}$, described by the cyclic word $(a_1 b_1 \bar{a}_1 \bar{b}_1)$,

$$\begin{tikzcd}
\bullet \arrow[d, "\rho_1"'] & \circ \arrow[l, "\rho_2"'] & \bullet \arrow[l, "\rho_3"'] \arrow[d, "\rho_1"] \\
\circ & & \circ \\
\bullet \arrow[u, "\rho_{123}"] & \circ \arrow[l, "\rho_2"] & \bullet \arrow[l, "\rho_3"] \arrow[u, "\rho_{123}"']
\end{tikzcd}$$

Suppose that $\boldsymbol{\ell}$ describes the invariant $\widehat{CFD}(M, \alpha, \beta)$ for some orientable three-manifold with torus boundary M. Then the abstract Dehn filling $\boldsymbol{\ell}_{\bullet}^{A} \boxtimes \boldsymbol{\ell}$ yields the chain complex $\widehat{CF}(M(\alpha))$. However, observe that $H_*(\boldsymbol{\ell}_{\bullet}^{A} \boxtimes \boldsymbol{\ell}) = 0$ (in particular, this slope is not an L–space slope according to Definition 3.18). This shows that no such M exists: in general, for a closed orientable three-manifold Y, $\widehat{HF}(Y)$ does not vanish [24].

However, we remark that this particular loop arises as a component of the graph describing $\widehat{CFD}(M,\mu,\lambda)$, where $M = S^3 \smallsetminus \nu(K)$ and K is any thin knot in S^3 that does not admit L–space surgeries; see [25].

4 Characterizing slopes

We now turn to an application of the loop calculus developed above. For a given loop $\boldsymbol{\ell}$ we give an explicit description of the L–space and non-L–space slopes. As a consequence, we will prove the following:

Theorem 4.1 *Given a loop $\boldsymbol{\ell}$, the set of non-L–space slopes is an interval in $\widehat{\mathbb{Q}}$, that is, the restriction of a connected subset in $\widehat{\mathbb{R}}$. As a result, if M is a loop-type manifold then there is a decomposition $\widehat{\mathbb{R}} = U \cup V$ into disjoint, connected subsets $U, V \in \widehat{\mathbb{R}}$ such that $\mathcal{L}_M = U \cap \widehat{\mathbb{Q}}$ and $\mathcal{L}^c_M = V \cap \widehat{\mathbb{Q}}$.*

Remark 4.2 The subset U determining $\mathcal{L}_M$ in this statement may be empty.

4.1 L–space slopes

A loop $\boldsymbol{\ell}$ has two preferred slopes, 0 and ∞; we begin by giving explicit conditions on the loop $\boldsymbol{\ell}$ (in terms of its representatives in standard or dual notation) under which these are L–space slopes.

Proposition 4.3 *Given a loop $\boldsymbol{\ell}$, ∞ is an L–space slope if and only if $\boldsymbol{\ell}$ can be written in standard notation with at least one d_k letter and no c_k letters (where k can be any integer). Similarly, 0 is an L–space slope if and only if $\boldsymbol{\ell}$ can be written in dual notation with at least one d^*_k letter and no c^*_k letters. Moreover, by reversing the orientation of $\boldsymbol{\ell}$, the statements hold with the roles of d_k and c_k (or the roles of d^*_k and c^*_k) reversed.*

Proof The slope ∞ is an L–space slope if $C\left(\frac{1}{0}\right) = \boldsymbol{\ell}^A_\bullet \boxtimes \boldsymbol{\ell}$ is an L–space complex, that is, if $H_*\big(C\left(\frac{1}{0}\right)\big)$ is nontrivial and each generator of $H_*\big(C\left(\frac{1}{0}\right)\big)$ has the same $(\mathbb{Z}/2\mathbb{Z})$–grading.

Recall that $\boldsymbol{\ell}^A_\bullet$ has a single generator x with idempotent ι_0 and operations

$$m_{3+i}(x,\rho_3,\underbrace{\rho_{23},\dots,\rho_{23}}_{i\text{ times}},\rho_2) = x$$

for each $i \geq 0$. The box tensor product of this module with $\boldsymbol{\ell}$ is easy to describe if $\boldsymbol{\ell}$ is written in standard notation. There is one generator $x \otimes y$ for each $\bullet$–vertex y in $\boldsymbol{\ell}$ (by abuse y is both the vertex in the loop $\boldsymbol{\ell}$ and the corresponding generator in the type D structure corresponding to ℓ). Given a cyclic word representing ℓ in standard notation, each letter represents a chain between adjacent $\bullet$–vertices; for each type a_k chain from y_1 to y_2 there is a differential from $x \otimes y_1$ to $x \otimes y_2$. Since sequential occurrences of type a chains are impossible in any loop, the differentials on $C\left(\frac{1}{0}\right)$ are isomorphisms mapping single generators to single generators. Therefore, the contribution to $H_*(\ell_\bullet^A \boxtimes \ell)$ (with relative $(\mathbb{Z}/2\mathbb{Z})$–grading) is simply given by the $\bullet$–vertices of ℓ which are not an endpoint of a type a segment.

The $(\mathbb{Z}/2\mathbb{Z})$–grading on $C\left(\frac{1}{0}\right)$ can be recovered as follows: the generators corresponding to two adjacent $\bullet$–vertices in ℓ have the same grading if the vertices are connected by an unstable chain and opposite gradings if the vertices are connected by a stable chain. It follows that endpoints of chains of type d_k all have the same grading and endpoints of chains of type c_k all have the opposite grading. This is because d_k segments must be separated from each other by an even number of stable chains, and from c_k segments by an odd number of stable chains.

Suppose ℓ can be written in standard notation with at least one d_k and no c_k. Every generator of $C\left(\frac{1}{0}\right)$ either comes from an endpoint of a d_k or from the common endpoint of two stable chains. The latter generators vanish in homology, since one of the two stable chains must be type a, and the former generators all have the same $(\mathbb{Z}/2\mathbb{Z})$–grading. Thus, in this case, $C\left(\frac{1}{0}\right)$ is an L–space complex.

Suppose now that ℓ contains both d_k and c_k segments. For any unstable chain, at least one of the two endpoints must correspond to a generator of $C\left(\frac{1}{0}\right)$ that survives in homology, since an unstable chain can be adjacent to a type a chain on at most one side. Since endpoints of type d_k and type c_k unstable chains produce generators of opposite grading, it follows that $C\left(\frac{1}{0}\right)$ has generators of both gradings that survive in homology, and thus $C\left(\frac{1}{0}\right)$ is not an L–space complex.

If ℓ has no unstable chains when written in standard notation, then it has only stable chains, which alternate between types a and b. In this case every $\bullet$–vertex is the endpoint of a type a chain. Thus $H_*\big(C\left(\frac{1}{0}\right)\big)$ is trivial, and $C\left(\frac{1}{0}\right)$ is not an L–space complex. Finally, if ℓ cannot be written in standard notation, then it consists only of e^* segments. $H_*(\ell_\bullet^A \boxtimes \ell)$ has two generators with opposite gradings; it follows that ∞ is not an L–space slope.

The proof for 0–filling is almost identical: $\ell_\circ^A$ has a single generator x with idempotent ι_1 and operations

$$m_{3+i}(x, \rho_2, \underbrace{\rho_{12}, \dots, \rho_{12}}_{i \text{ times}}, \rho_1) = x$$

for each $i \geq 0$. The box tensor product of this module with $\boldsymbol{\ell}$ has one generator for each generator y of $\boldsymbol{\ell}$ with idempotent ι_1 (that is, each $\circ$–vertex) and a differential for each type a^* segment. Expressing $\boldsymbol{\ell}$ in dual notation, the contribution to $H_*(\boldsymbol{\ell}_\circ^A \boxtimes \boldsymbol{\ell})$ (with relative $(\mathbb{Z}/2\mathbb{Z})$–grading) is given by the ι_0–generators of $\boldsymbol{\ell}$ which do not lie at the end of a type a^* segment. If $\widehat{CFD}(Y,\mathfrak{s})$ is a loop that cannot be written in dual notation — that is, it is a collection of e segments — then the contribution to $H_*(\boldsymbol{\ell}_\circ^A \boxtimes \boldsymbol{\ell}^D)$ is two generators of opposite grading.

Since a^* and b^* segments change the $(\mathbb{Z}/2\mathbb{Z})$–grading while c^* and d^* segments do not, the rest of the proof is completely analogous to the proof for the ∞–filling. $\square$

Combining the two conditions in Proposition 4.3 gives a stronger condition on the loop.

Proposition 4.4 *Given a loop $\boldsymbol{\ell}$, both ∞ and 0 are L–space slopes if and only if the following equivalent conditions hold*:

(i) *$\boldsymbol{\ell}$ can be written in standard notation with at least one d_k letter and no c_k letters (with $k \in \mathbb{Z}$), and $\boldsymbol{\ell}$ contains a subword from exactly one of the sets*
$$A_+ = \{b_i a_j, a_i e^n b_j, a_i e^n d_j, d_i e^n b_j, d_i e^n d_j, d_\ell, a_\ell, b_\ell \mid i,j \geq 1, n \geq 0, \ell \geq 2\},$$
$$A_- = \{b_i a_j, a_i e^n b_j, a_i e^n d_j, d_i e^n b_j, d_i e^n d_j, d_\ell, a_\ell, b_\ell \mid i,j \leq 1, n \geq 0, \ell \leq 2\}.$$
We will say that $\boldsymbol{\ell}$ satisfies condition (i)$_\pm$ *if it satisfies condition* (i) *with a subword in $A_\pm$.*

(ii) *$\boldsymbol{\ell}$ can be written in dual notation with at least one d_k^* letter and no c_k^* letters (with $k \in \mathbb{Z}$), and $\boldsymbol{\ell}$ contains a subword from exactly one of the sets*
$$A_+^* = \bigl\{b_i^* a_j^*, a_i^*(e^*)^n b_j^*, a_i^*(e^*)^n d_j^*, d_i^*(e^*)^n b_j^*, d_i^*(e^*)^n d_j^*, d_\ell^*, a_\ell^*, b_\ell^* \mid i,j \geq 1, n \geq 0, \ell \geq 2\bigr\},$$
$$A_-^* = \bigl\{b_i^* a_j^*, a_i^*(e^*)^n b_j^*, a_i^*(e^*)^n d_j^*, d_i^*(e^*)^n b_j^*, d_i^*(e^*)^n d_j^*, d_\ell^*, a_\ell^*, b_\ell^* \mid i,j \leq 1, n \geq 0, \ell \leq 2\bigr\}.$$
We will say that $\boldsymbol{\ell}$ satisfies condition (ii)$_\pm$ *if it satisfies condition* (ii) *with a subword in $A_\pm^*$.*

Moreover, condition (i)$_+$ *is equivalent to condition* (ii)$_+$ *and condition* (i)$_-$ *is equivalent to condition* (ii)$_-$.

Proof By Proposition 4.3, 0 and ∞ are both L–space slopes for $\boldsymbol{\ell}$ if and only if $\boldsymbol{\ell}$ contains either d_k segments or c_k segments in standard notation, but not both, and $\boldsymbol{\ell}$ contains either d_k^* segments or c_k^* segments in dual notation, but not both. Reading a loop with the opposite orientation takes type d segments to type c segments, so, possibly after switching the orientation on $\boldsymbol{\ell}$, we may assume such a loop contains d_k segments and not c_k segments in standard notation. Similarly, up to switching the orientation on $\boldsymbol{\ell}$, we could arrange that there are d_k^* segments and no c_k^* segments in dual notation. Note however that we cannot ensure both of these simultaneously, as it will be convenient to choose standard notation and dual notation representatives for $\boldsymbol{\ell}$ with respect to the same orientation on the loop.

First assume that $\boldsymbol{\ell}$ contains standard unstable chains of type d_k but not c_k. Under this assumption, the set A_+ is precisely the subwords in standard notation that dualize to give a d_k^* chain. This follows from the discussion on dualizing in Section 3.2. A d_{n+1}^* segment with $n \geq 0$ arises from a subword $a_i e^n b_j$, $a_i e^n d_j$, $d_i e^n b_j$ or $d_i e^n d_j$ with $i, j \geq 1$. A segment d_{-n-1}^* arises from a subword $b_i \bar{e}^n a_j$, $b_i \bar{e}^n c_j$, $c_i \bar{e}^n a_j$ or $c_i \bar{e}^n c_j$ with $i, j \geq 1$; since we assume that $\boldsymbol{\ell}$ contains no c_k segments, including $\bar{e} = c_0$, the only relevant case is $b_i a_j$. Finally, a $d_0^* = e^*$ segment arises from a standard letter of type a_ℓ, b_ℓ, c_ℓ, or d_ℓ with subscript $\ell > 1$; we can ignore the case of c_ℓ by assumption. By similar reasoning, we can check that, under the assumption that $\boldsymbol{\ell}$ contains no c_k letters, the set A_- is precisely the subwords in standard notation that dualize to give c_k^* chains.

If we instead assume that $\boldsymbol{\ell}$ contains dual unstable chains of type d_k^* but not c_k^*, the argument is similar. Note that the sets $A_\pm^*$ are the same as the sets $A_\pm$ with stars added to each letter, and the process for switching from dual to standard notation is the same as the process for switching from standard to dual (up to adding/removing stars from letters). We have immediately that A_+^* is the set of subwords in that dualize to give a d_k letter, and A_-^* is the set of subwords that dualize to give a c_k letter.

We have shown that conditions (i) and (ii) are both equivalent to $\boldsymbol{\ell}$ containing an unstable chain of exactly one of the two types in both standard notation and dual notation, and this is equivalent to both 0 and ∞ being L–space slopes. Finally, if $\boldsymbol{\ell}$ satisfies condition (i)$_+$ then $\boldsymbol{\ell}$ contains both a d_k segment and a d_k^* segment (with respect to the same orientation of the loop). It follows that $\boldsymbol{\ell}$ satisfies condition (ii)$_+$. If $\boldsymbol{\ell}$ satisfies condition (i)$_-$ then it contains both a d_k segment and a c_k^* segment. After reversing the orientation of $\boldsymbol{\ell}$, it contains a c_k segment and a d_k^* segment; thus $\boldsymbol{\ell}$ satisfies condition (ii)$_-$. □

Corollary 4.5 *If* 0 *and* ∞ *are L–space slopes for a loop $\boldsymbol{\ell}$, then the set of L–space slopes for $\boldsymbol{\ell}$ contains either all positive slopes* (*that is*, *the interval* $[0, \infty]$) *or all negative slopes* (*that is*, *the interval* $[-\infty, 0]$).

Proof Let L_+ be the set of loops satisfying condition (i)$_+$ and (ii)$_+$ in Proposition 4.4, and let L_- be the set of loops satisfying conditions (i)$_-$ and (ii)$_-$. It is not difficult to see that condition (i)$_+$ is preserved by the operation T since, for any $\boldsymbol{w} \in A_+$, $\mathrm{T}(\boldsymbol{w})$ is in A_+ or contains an element of A_+ as a subword. Similarly, condition (ii)$_+$ is preserved by the operation H. Therefore the set L_+ is preserved by T and H. In the same way, the set L_- is preserved by T^{-1} and H^{-1}.

If $p/q > 0$, let $[a_1, \dots, a_n]$ be a continued fraction for p/q of even length with all positive terms. To see if p/q is an L–space slope, we reparametrize by

$$\mathrm{H}^{a_n} \circ \cdots \circ \mathrm{T}^{a_3} \circ \mathrm{H}^{a_2} \circ \mathrm{T}^{a_1}$$

taking the slope p/q to the slope ∞. It follows that, if a loop is in L_+, then any $p/q > 0$ is an L–space slope, since the reparametrized loop is also in L_+.

If $p/q < 0$, let $[-a_1, \dots, -a_n]$ be a continued fraction for p/q of even length with all negative terms. To see if p/q is an L–space slope, we reparametrize by

$$\mathrm{H}^{-a_n} \circ \cdots \circ \mathrm{T}^{-a_3} \circ \mathrm{H}^{-a_2} \circ \mathrm{T}^{-a_1}$$

taking the slope p/q to the slope ∞. It follows that if a loop is in L_-, then any $p/q < 0$ is an L–space slope. □

4.2 Non-L–space slopes

The goal of this section is to establish easily checked conditions certifying that the standard and dual slope of a given loop are non-L–space slopes. Our focus will be on the following result:

Proposition 4.6 *Suppose that $\boldsymbol{\ell}$ is a loop for which both standard and dual fillings give rise to non-L–spaces. Then the standard and dual slopes bound an interval of non-L–space slopes in $\mathcal{L}_M^c$. That is, one of $[0, \infty]$ or $[-\infty, 0]$ consists entirely of non-L–space slopes.*

The proof of this result is similar to the proof of Corollary 4.5 but is more technical and will therefore be built up in a series of lemmas.

To begin, note that Proposition 4.3 can be restated in terms of non-L–space slopes:

Proposition 4.7 *The slope ∞ is a non-L–space slope for a loop $\boldsymbol{\ell}$ if and only if either*

(i) *$\boldsymbol{\ell}$ contains both c_k and d_k unstable chains in standard notation,*

(ii) *$\boldsymbol{\ell}$ contains no unstable chains in standard notation, or*

(iii) *$\boldsymbol{\ell}$ cannot be written in standard notation.*

The slope 0 *is a non-*L*–space slope if and only if either*

(i) ℓ *contains both* c_k^* *and* d_k^* *unstable chains in dual notation*,

(ii) ℓ *contains no unstable chains in dual notation*, *or*

(iii) ℓ *cannot be written in dual notation.*

It will be helpful to give conditions in standard notation under which 0 is a non-L–space slope. Consider the sets

$$A_1 = \{a_k, b_k, c_k, c_i c_0^n c_j, c_i c_0^n a_j, b_i c_0^n c_j, b_i c_0^n a_j\},$$
$$A_2 = \{a_{-k}, b_{-k}, d_{-k}, d_{-i} d_0^n d_{-j}, d_{-i} d_0^n b_{-j}, a_{-i} d_0^n d_{-j}, a_{-i} d_0^n b_{-j}\},$$
$$A_3 = \{a_k, b_k, d_k, d_i d_0^n d_j, d_i d_0^n b_j, a_i d_0^n d_j, a_i d_0^n b_j\},$$
$$A_4 = \{a_{-k}, b_{-k}, c_{-k}, c_{-i} c_0^n c_{-j}, c_{-i} c_0^n a_{-j}, b_{-i} c_0^n c_{-j}, b_{-i} c_0^n a_{-j}\},$$

where the indices and exponents run over all integers satisfying $k > 1$, $i, j > 0$ and $n \geq 0$. These four sets may be interpreted as follows:

Lemma 4.8 *A loop* ℓ *written in standard notation contains a word in* $A_1 \cup A_3$ *if and only if it contains* d_n^* *in dual notation for some integer* n. *It contains a word in* $A_2 \cup A_4$ *if and only if it contains* c_n^* *in dual notation for some integer* n.

Proof We prove the first statement and leave the second to the reader.

First notice that if $k > 1$ then any letter a_k, b_k, c_k or d_k contains at least one instance of $\circ \xrightarrow{\rho_{23}} \circ$, which, in dual notation, gives an $e^* = d_0^*$. If $i, j > 0$ then each word of the form $c_i c_0^n c_j$, $c_i c_0^n a_j$, $b_i c_0^n c_j$ or $b_i c_0^n a_j$ gives an instance of

$$\circ \xleftarrow{\rho_1} \bullet \overset{\rho_{12}}{\dashrightarrow} \bullet \xrightarrow{\rho_3} \circ$$

which, in dual notation, gives $\bar{c}_{n+1}^* = d_{-n-1}^*$. Similarly, each word of the form $d_i d_0^n d_j$, $d_i d_0^n b_j$, $a_i d_0^n d_j$ or $a_i d_0^n b_j$ gives an instance of

$$\circ \xrightarrow{\rho_2} \bullet \overset{\rho_{12}}{\dashrightarrow} \bullet \xrightarrow{\rho_{123}} \circ$$

which, in dual notation, gives d_{n+1}^*. For the converse, observe that the segments d_0^*, d_{n+1}^* and d_{-n-1}^* with $n \geq 0$ in dual notation can only arise from the words mentioned above in standard notation. $\square$

Corollary 4.9 *The slope* 0 *is not an* L*–space slope for* ℓ *if and only if either*

(i) ℓ *contains a subword from* $A_1 \cup A_3$ *and a subword from* $A_2 \cup A_4$,

(ii) ℓ *does not contain a subword from any* A_i, *or*

(iii) ℓ *cannot be written in dual notation.*

Proof This follows immediately from Proposition 4.7 and Lemma 4.8. □

Based on this observation, the proof of Proposition 4.6 will reduce to several cases, which are treated in the following lemmas.

Lemma 4.10 *If $\boldsymbol{\ell}$ is a loop containing a subword from the set A_1 then $\mathrm{T}^{-1}(\boldsymbol{\ell})$ and $\mathrm{H}^{-1}(\boldsymbol{\ell})$ also contain a subword from A_1. If $\boldsymbol{\ell}$ contains a subword from the set A_2 then $\mathrm{T}^{-1}(\boldsymbol{\ell})$ and $\mathrm{H}^{-1}(\boldsymbol{\ell})$ also contain a subword from A_2.*

Proof First consider $\mathrm{T}^{-1}(\boldsymbol{w})$ for any subword $\boldsymbol{w} \in A_1$. We have $\mathrm{T}^{-1}(a_k) = a_k$, $\mathrm{T}^{-1}(b_k)=b_k$, $\mathrm{T}^{-1}(c_k)=c_{k+1}$, $\mathrm{T}^{-1}(c_i\bar{e}^n c_j)=c_{i+1}c_1^n c_{j+1}$, $\mathrm{T}^{-1}(c_i\bar{e}^n a_j)=c_{i+1}c_1^n a_j$, $\mathrm{T}^{-1}(b_i\bar{e}^n c_j) = b_i c_1^n c_{j+1}$ and $\mathrm{T}^{-1}(b_i\bar{e}^n a_j) = b_i c_1^n a_j$. In each case, $\mathrm{T}^{-1}(\boldsymbol{w})$ is in A_1 or contains a subword in A_1. Thus the operation T^{-1} preserves the property that $\boldsymbol{\ell}$ contains a word from A_1. On the other hand, revisiting the proof of Lemma 4.8, we see that each of a_k, b_k and c_k produces at least one e^* when dualized, while each of $c_i\bar{e}^n c_j$, $c_i\bar{e}^n a_j$, $b_i\bar{e}^n c_j$ and $b_i\bar{e}^n a_j$ produces a $\bar{c}_{n+1}^*$. Since $\mathrm{H}^{-1}(e^*) = \bar{c}_1^*$ and $\mathrm{H}^{-1}(\bar{c}_{n+1}^*) = \bar{c}_{n+2}^*$, it is enough to observe that

$$\{c_i\bar{e}^n c_j, c_i\bar{e}^n a_j, b_i\bar{e}^n c_j, b_i\bar{e}^n a_j\} \subset A_1$$

is the set of subwords that give rise to a $\bar{c}_{n+1}^*$ under dualizing (where $n \geq 0$ and $i, j > 0$). It follows that the operation H^{-1} preserves the property that $\boldsymbol{\ell}$ contains a word from A_1.

Next consider $\mathrm{T}^{-1}(\boldsymbol{w})$ for any $\boldsymbol{w} \in A_2$. Proceeding as above, observe that $\mathrm{T}^{-1}(\boldsymbol{w})$ is in the set

$$\{\bar{a}_k, \bar{b}_k, \bar{c}_{k+1}, \bar{c}_{i+1}\bar{c}_1^n\bar{c}_{j+1}, \bar{c}_{i+1}\bar{c}_1^n\bar{b}_j, \bar{a}_i\bar{c}_1^n\bar{c}_{j+1}, \bar{a}_i\bar{c}_1^n\bar{b}_j\}.$$

Each of these words is in A_2 or contains a subword in A_2, so the operation T^{-1} preserves the property that $\boldsymbol{\ell}$ contains a word from A_2. Each of $\bar{a}_k$, $\bar{b}_k$ and $\bar{c}_k$ produces at least one $\bar{e}^*$ when dualized, while each of $\bar{c}_i e^n \bar{c}_j$, $\bar{c}_i e^n \bar{b}_j$, $\bar{a}_i e^n \bar{c}_j$ and $\bar{a}_i e^n \bar{b}_j$ produces a c_{n+1}^*. Since $\mathrm{T}^{-1}(\bar{e}^*) = c_1^*$ and $\mathrm{T}^{-1}(c_{n+1}^*) = c_{n+2}^*$, it is enough to observe that

$$\{\bar{c}_i e^n \bar{c}_j, \bar{c}_i e^n \bar{b}_j, \bar{a}_i e^n \bar{c}_j, \bar{a}_i e^n \bar{b}_j\} \subset A_2$$

is the set of subwords that give rise to a c_{n+1}^* under dualizing (where $n \geq 0$ and $i, j > 0$). It follows that the operation H^{-1} preserves the property that $\boldsymbol{\ell}$ contains a word from A_2. □

Lemma 4.11 *If $\boldsymbol{\ell}$ is a loop containing a subword from the set A_3 then $\mathrm{T}(\boldsymbol{\ell})$ and $\mathrm{H}(\boldsymbol{\ell})$ also contain a subword from A_3. If $\boldsymbol{\ell}$ contains a subword from the set A_4, then $\mathrm{T}(\boldsymbol{\ell})$ and $\mathrm{H}(\boldsymbol{\ell})$ also contain a subword from A_4.*

Proof This is analogous to the proof of Lemma 4.10 and left to the reader. □

We next consider the case that ℓ contains a subword from each of A_1 and A_4 but does not contain a subword from A_2 or A_3. While this third case is ultimately similar to the two cases we have already treated, it is more technical owing in part to three subcases.

Lemma 4.12 *Suppose ℓ is a loop containing a subword from each of the sets A_1 and A_4 but that ℓ does not contain a subword from either of the sets A_2 or A_3. If ℓ contains a $\bar{c}_1$ then the loop $\mathrm{T}^{-1}(\ell)$ contains a subword from each of A_1 and A_2, while the loop $\mathrm{H}^{-1}(\ell)$ either*

(i) *contains a subword from each of A_1 and A_2, or*

(ii) *contains a $\bar{c}_1$ and a subword from each of the sets A_1 and A_4 but does not contain a subword from A_2 or from A_3.*

Proof First notice that $\mathrm{T}^{-1}(\bar{c}_1) = \bar{c}_2 \in A_2$. Since the set A_1 is closed under T^{-1} (compare Lemma 4.10), it follows that $\mathrm{T}^{-1}(\ell)$ contains a subword from both A_1 and A_2.

Turning now to the behaviour under H^{-1}: Since ℓ does not contain a word from A_2, any occurrence of $\bar{c}_1$ can be preceded only by a_1e^i or d_1e^i (for $i \geq 0$), and can be followed only by e^jb_1 or e^jd_1 (for $j \geq 0$). (Note that instances of a_k or b_k for $k > 1$ are ruled out since ℓ does not contain a word from A_3.) Regardless of which arises, this ensures that a $\bar{c}_1$ is part of a segment in ℓ of the form

$$\circ \xrightarrow{\rho_2} \bullet \overset{\rho_{12}}{\dashrightarrow} \bullet \xrightarrow{\rho_1} \circ \xleftarrow{\rho_3} \bullet \overset{\rho_{12}}{\dashrightarrow} \bullet \xrightarrow{\rho_{123}} \circ$$

and hence a subword $b^*_{i+1}a^*_{j+1}$ when ℓ is expressed in dual notation. It follows that the property that ℓ contains a $\bar{c}_1$ is closed under H^{-1}, under the assumption that ℓ does not contain a subword from A_2 or A_3.

Now observe that the set A_1 is closed under H^{-1} by Lemma 4.10, so it must be the case that $\mathrm{H}^{-1}(\ell)$ contains subwords from A_1 and A_2 or, if this is not the case, from A_1 and A_4. In the latter case, suppose that $\mathrm{H}^{-1}(\ell)$ contains a word from A_3. Then $\mathrm{H}(\mathrm{H}^{-1}(\ell)) = \ell$ contains a word from A_3 also, which is a contradiction. □

Lemma 4.13 *Suppose ℓ is a loop containing a subword from each of the sets A_1 and A_4 but that ℓ does not contain a subword from either of the sets A_2 or A_3. If ℓ contains a d_1 then the loop $\mathrm{T}(\ell)$ contains a subword from each of A_3 and A_4, while the loop $\mathrm{H}(\ell)$ either*

(i) *contains a subword from each of A_3 and A_4, or*

(ii) *contains a d_1 and a subword from each of the sets A_1 and A_4 but does not contain a subword from A_2 or from A_3.*

Proof This is similar to the proof of Lemma 4.12. The key observation is that the instance of d_1 is part of a segment in $\boldsymbol{\ell}$ of the form

$$\circ \xrightarrow{\rho_3} \bullet \overset{\rho_{12}}{\dashrightarrow} \bullet \xrightarrow{\rho_{123}} \circ \xrightarrow{\rho_2} \bullet \overset{\rho_{12}}{\dashrightarrow} \bullet \xleftarrow{\rho_1} \circ$$

which, in dual form, gives $a^*_{j+1} b^*_{i+1}$. Thus the property that $\boldsymbol{\ell}$ contains a d_1 implies that $\mathrm{H}(\boldsymbol{\ell})$ contains a d_1. □

Lemma 4.14 *Suppose $\boldsymbol{\ell}$ is a loop containing a subword from each of the sets A_1 and A_4 but that $\boldsymbol{\ell}$ does not contain a subword from either of the sets A_2 or A_3, and suppose that ∞ is a non-L–space filling for $\boldsymbol{\ell}$. If $\boldsymbol{\ell}$ contains neither a $\bar{c}_1 = d_{-1}$ nor a d_1 then two cases arise: either $\boldsymbol{\ell}$ contains the subword $\boldsymbol{w}_1 = a_1 d_0^n \bar{b}_1$ or $\boldsymbol{\ell}$ contains the subword $\boldsymbol{w}_2 = \bar{a}_1 d_0^n b_1$ for some $n \geq 1$. Regardless of which occurs,*

(i) *$\mathrm{T}(\boldsymbol{\ell})$ contains subwords from both A_3 and A_4,*

(ii) *$\mathrm{T}^{-1}(\boldsymbol{\ell})$ contains subwords from both A_1 and A_2,*

(iii) *$\mathrm{H}(\boldsymbol{\ell})$ either contains subwords A_3 and A_4 or it contains subwords from A_1 and A_4 and contains $\boldsymbol{w}_1$ or $\boldsymbol{w}_2$, and*

(iv) *$\mathrm{H}^{-1}(\boldsymbol{\ell})$ either contains subwords from A_1 and A_2 or it contains subwords from A_1 and A_4 and contains $\boldsymbol{w}_1$ or $\boldsymbol{w}_2$.*

Proof Suppose $\boldsymbol{\ell}$ contains subwords from A_1 and A_4 but not from A_2 or A_3, and $\boldsymbol{\ell}$ contains neither d_1 nor d_{-1}. By Proposition 4.7, if $\boldsymbol{\ell}$ has a non-L–space standard filling, then either $\boldsymbol{\ell}$ contains no standard unstable chains or $\boldsymbol{\ell}$ contains both c_k and d_k unstable chains. In the former case $\boldsymbol{\ell}$ consists of alternating a_k and b_k segments; in the latter case, $\boldsymbol{\ell}$ must contain a sequence of d_0 segments preceded by an $a_{\pm 1}$ segment and followed by a $b_{\pm 1}$ segment (since a_k, b_k and d_k are in A_2 or A_3 for $|k| > 1$, and by assumption $\boldsymbol{\ell}$ also contains no $d_{\pm 1}$ segments). In either case, we see that the subscripts on the type a and b stable chains must have opposite signs (since $a_i d_0^n b_j$ is in A_3 and $a_{-i} d_0^n b_{-j}$ is in A_2), and thus that $\boldsymbol{\ell}$ contains $\boldsymbol{w}_1 = a_1 d_0^n \bar{b}_1$ or $\boldsymbol{w}_2 = \bar{a}_1 d_0^n b_1$.

Note then that $\mathrm{T}(\boldsymbol{w}_1) = a_1 d_1^n \bar{b}_1$, which gives a word in A_3; $\mathrm{T}^{-1}(\boldsymbol{w}_1) = a_1 \bar{c}_1^n \bar{b}_1$, which gives a word in A_2; $\mathrm{T}(\boldsymbol{w}_2) = \bar{a}_1 d_1^n b_1$, which gives a word in A_3; and $\mathrm{T}^{-1}(\boldsymbol{w}_2) = \bar{a}_1 \bar{c}_1^n b_1$, which gives a word in A_2. Since T preserves the property that there is a subword from A_4 (Lemma 4.11) and T^{-1} preserves the property that there is a subword from A_1 (Lemma 4.10), we have shown that conclusions (i) and (ii) hold.

Consider the subword $\boldsymbol{w}_1 = a_1 d_0^n \bar{b}_1$. If a loop $\boldsymbol{\ell}$ contains $\boldsymbol{w}_1$, then the graph for $\boldsymbol{\ell}$ contains the segment

$$\bullet \xrightarrow{\rho_3} \circ \xrightarrow{\rho_2} \bullet \overset{\rho_{12}}{\dashrightarrow} \bullet \xrightarrow{\rho_1} \circ \xleftarrow{\rho_{123}} \bullet$$

It follows that in dual notation $\boldsymbol{\ell}$ contains a b^*_{n+1}. Moreover, this b^*_{n+1} is preceded by a segment ending with a ρ_3 arrow (that is, a segment of type $\bar{a}^*$ or $\bar{c}^*$) and followed by a segment beginning with a backwards ρ_{123} arrow (that is, a segment of type $\bar{a}^*$ or $\bar{d}^*$). Thus $\boldsymbol{\ell}$ contains $\boldsymbol{w}_1$ in standard notation if and only if it contains one of the following in dual notation: $\bar{a}^*_i b^*_{n+1} \bar{a}^*_j$, $\bar{a}^*_i b^*_{n+1} \bar{d}^*_j$, $\bar{c}^*_i b^*_{n+1} \bar{a}^*_j$ or $\bar{c}^*_i b^*_{n+1} \bar{d}^*_j$, where $i, j > 0$. Note that the set of such words is closed under the action of H, except for the last two words if the index i is 1 since then H replaces the $\bar{c}_1$ with $\bar{c}_0 = e = d_0$. Thus, if $\boldsymbol{\ell}$ contains one of these subwords, then $\mathrm{H}(\boldsymbol{\ell})$ also contains one of these subwords (and thus the subword $\boldsymbol{w}_1$ in standard notation) or it contains $d^*_0 b^*_{n+1}$. In the latter case, $\boldsymbol{\ell}$ contains a ρ_{23} arrow followed by a ρ_2 arrow. This means that $\boldsymbol{\ell}$ contains an a_k or d_k segment with $k > 1$, each of which are in A_3. Similarly, $\mathrm{H}^{-1}(\boldsymbol{\ell})$ either contains $\boldsymbol{w}_1$ in standard notation or it contains $b^*_{n+1} \bar{d}^*_0$, which implies that $\boldsymbol{\ell}$ contains a subword in A_2.

The details of a similar argument for the subword $\boldsymbol{w}_2 = \bar{a}_1 d_0^n \bar{b}_1$ are left to the reader. We similarly find that, if $\boldsymbol{\ell}$ contains $\boldsymbol{w}_2$, then $\mathrm{H}(\boldsymbol{\ell})$ contains either $\boldsymbol{w}_2$ or a subword in A_3 and $\mathrm{H}^{-1}(\boldsymbol{\ell})$ contains either $\boldsymbol{w}_2$ or a subword in A_2.

Since by assumption $\boldsymbol{\ell}$ contains subwords from A_1 and A_4, $\mathrm{H}(\boldsymbol{\ell})$ contains a subword from A_4 and from either A_1 or A_3 (Lemma 4.13). Since $\boldsymbol{\ell}$ also contains a subword $\boldsymbol{w}_1$ or $\boldsymbol{w}_2$, the discussion above implies that $\mathrm{H}(\boldsymbol{\ell})$ contains $\boldsymbol{w}_1$, $\boldsymbol{w}_2$ or a subword in A_3; that is, (iii) is satisfied. Similarly, (iv) is satisfied since $\mathrm{H}^{-1}(\boldsymbol{\ell})$ contains a subword in A_2 and a subword in either A_1 or A_3 (Lemma 4.12), and a subword $\boldsymbol{w}_1$ or $\boldsymbol{w}_2$ if there is no subword from A_1. □

Lemma 4.15 *Suppose $\boldsymbol{\ell}$ contains no subwords in any A_i and that $\boldsymbol{\ell}$ contains both a d_k segment and a c_k segment. Then $\mathrm{T}^{-1}(\boldsymbol{\ell})$ contains a subword from each of A_1 and A_2 and $\mathrm{T}(\boldsymbol{\ell})$ contains a subword from each of A_3 and A_4.*

Proof The assumption that $\boldsymbol{\ell}$ contains no words from any of the sets A_i implies that in standard notation $\boldsymbol{\ell}$ consists of the segments $a_{\pm 1}$, $b_{\pm 1}$, $c_{\pm 1}$, $d_{\pm 1}$, c_0 and d_0, and the nonzero subscripts alternate between $+1$ and -1. By assumption, $\boldsymbol{\ell}$ must contain at least one d_1, d_0 or d_{-1} segment. If $\boldsymbol{\ell}$ contains d_1 then $\mathrm{T}(\boldsymbol{\ell})$ contains $d_2 \in A_3$. Moreover, d_1 must be preceded by d_0, d_{-1} or a_{-1} and followed by d_0, d_{-1} or b_{-1}. It follows that $\mathrm{T}^{-1}(\boldsymbol{\ell})$ contains a d_0 preceded by a d_{-1}, d_{-2} or a_{-1} and followed by a d_{-1}, d_{-2} or b_{-1}; thus $\mathrm{T}^{-1}(\boldsymbol{\ell})$ contains a subword in A_2. Similarly if $\boldsymbol{\ell}$ contains d_{-1} then $\mathrm{T}^{-1}(\boldsymbol{\ell})$ contains $d_{-2} \in A_2$, and, since d_{-1} must be preceded by d_0, d_1 or a_1 and followed by d_0, d_1 or b_1, $\mathrm{T}(\boldsymbol{\ell})$ contains a subword in A_3. Finally, suppose $\boldsymbol{\ell}$ contains a d_0 segment but no d_1 or d_{-1}. If the d_0 segment is adjacent to another d_0 segment, then $\mathrm{T}(\boldsymbol{\ell})$ contains $d_1 d_1 \in A_3$ and $\mathrm{T}^{-1}(\boldsymbol{\ell})$ contains $d_{-1} d_{-1} \in A_2$. Otherwise, $\boldsymbol{\ell}$ must contain either $a_{-1} d_0 b_1$ or $a_1 d_0 b_{-1}$. It is easy to check that $\mathrm{T}(\boldsymbol{\ell})$ contains either $d_1 b_1$ or $a_1 d_1$, both elements of A_3, and that $\mathrm{T}^{-1}(\boldsymbol{\ell})$ contains either $a_{-1} d_{-1}$ or $d_{-1} b_{-1}$ in A_2.

A similar argument shows that, if $\boldsymbol{\ell}$ contains c_1, c_0 or c_{-1}, then $\mathrm{T}(\boldsymbol{\ell})$ contains a subword in A_4 and $\mathrm{T}^{-1}(\boldsymbol{\ell})$ contains a subword in A_1. This can also be deduced by reversing the orientation on $\boldsymbol{\ell}$, which takes d_k segments to c_{-k} segments and takes the sets A_2 and A_3 to A_1 and A_4. □

Proof of Proposition 4.6 Let $\boldsymbol{\ell}$ be a loop for which both standard and dual filling gives a non-L–space. According to Corollary 4.9, the fact that dual filling is a non-L–space implies that either $\boldsymbol{\ell}$ contains a word from A_1 or A_3 and from A_2 or A_4, $\boldsymbol{\ell}$ contains no words from any A_i, or $\boldsymbol{\ell}$ cannot be written in dual notation. In the last case, $\boldsymbol{\ell}$ consists only of type e segments and it is clear that standard filling is an L–space. Thus we have the following cases to consider:

(1) $\boldsymbol{\ell}$ contains a subword from A_1 and from A_2.

(2) $\boldsymbol{\ell}$ contains a subword from A_3 and from A_4.

(3) $\boldsymbol{\ell}$ contains a subword from A_1 and from A_4, but not A_2 or A_3.

(4) $\boldsymbol{\ell}$ does not contain a subword from any A_i.

Note that, a priori, there is another case similar to (3) in which $\boldsymbol{\ell}$ contains subwords in A_2 and A_3; however, this is equivalent to (3) after reversing orientation on $\boldsymbol{\ell}$.

In case (1), Lemma 4.10 implies that applying any combination of T^{-1} and H^{-1} to $\boldsymbol{\ell}$ gives rise to a loop $\boldsymbol{\ell}'$ which contains a subword in each of A_1 and A_2, and thus has non-L–space dual filling. The p/q filling on $\boldsymbol{\ell}$ is given by dual filling $\mathrm{T}^{a_n} \circ \cdots \circ \mathrm{H}^{a_2} \circ \mathrm{T}^{a_1}(\boldsymbol{\ell})$, where $[a_1, \dots, a_n]$ is an odd length continued fraction for p/q. For any $-\infty < p/q < 0$, we can choose a continued fraction with each $a_i \le 0$, and so p/q is not an L–space slope for $\boldsymbol{\ell}$. Similarly, in case (2) we choose a continued fraction with positive terms and use Lemma 4.11 to see that any $0 < p/q < \infty$ is a non-L–space.

Case (3) has three subcases, depending on whether $\boldsymbol{\ell}$ contains a $\bar{c}_1$, a d_1 or neither. If $\boldsymbol{\ell}$ contains a $\bar{c}_1$, then, by Lemmas 4.12 and 4.10, applying any combination of T^{-1} and H^{-1} to $\boldsymbol{\ell}$ with at least one application of T^{-1} produces a loop $\boldsymbol{\ell}'$ with a subword in A_1 and A_2. Given any $-\infty < p/q < 0$ we can choose an odd length continued fraction $[a_1, \dots, a_n]$ for which each a_i is strictly negative except that a_1 may be 0, and if $a_1 = 0$ then $n \ge 3$. It follows that p/q is a non-L–space slope for $\boldsymbol{\ell}$. Similarly, if $\boldsymbol{\ell}$ contains a d_1, then Lemmas 4.13 and 4.11 imply that any $0 < p/q < \infty$ is a non-L–space slope. If ℓ contains neither a d_1 nor a $\bar{c}_1$, then any p/q is a non-L–space slope by Lemma 4.14.

In case (4), ℓ has no dual unstable chains. Since dual stable chains are fixed by $\mathrm{H}^{\pm 1}$, we have immediately that $\mathrm{H}(\ell) = \mathrm{H}^{-1}(\ell) = \ell$. Since ∞ is a non-L–space slope, ℓ must have either no standard unstable chains or unstable chains of both type c_k and type d_k. If ℓ has no standard unstable chains then $\mathrm{T}(\ell) = \mathrm{T}^{-1}(\ell) = \ell$; it follows that any combination of Dehn twists preserves ℓ, and so any slope p/q is a non-L–space slope (in fact in this case any filling of ℓ has trivial homology). If ℓ has standard unstable chains of both types, then Lemma 4.15 implies that any slope p/q is a non-L–space slope. □

4.3 Proof of Theorem 4.1

We are now ready to prove that the set of L–space slopes for a loop ℓ is a (possibly empty) interval in $\widehat{\mathbb{Q}}$. We will use the fact that applying T and H to a loop ℓ changes the set of L–space slopes in a controlled way: the slope p/q for ℓ is equivalent to the slope $(p-nq)/q$ for $\mathrm{T}^n(\ell)$ and the slope $p/(q-np)$ for $\mathrm{H}^n(\ell)$. These transformations preserve the cyclic ordering on $\widehat{\mathbb{Q}}$, and thus preserve the connectedness of the set of L–space slopes.

Remark 4.16 In particular, while Corollary 4.5 and Proposition 4.6 are stated in terms of the 0 and ∞ slopes, analogous statements hold for any two slopes of distance 1 since such pair of slopes can be taken to 0 and ∞ by a combination of $\mathrm{T}^{\pm 1}$ and $\mathrm{H}^{\pm 1}$. More precisely, for any two slopes of distance 1, if both are L–space slopes then one of the two intervals between them consists entirely of L–space slopes, and if both are non-L–space slopes then one of the two intervals between them consists entirely of non-L–space slopes.

In addition to Corollary 4.5 and Proposition 4.6, we will need the following lemma:

Lemma 4.17 *For any loop ℓ, there do not exist integers $n_1 < n_2 < n_3 < n_4$ such that n_1 and n_3 are L–space slopes and n_2 and n_4 are not L–space slopes, or vice versa.*

Proof An integer n is an L–space slope for ℓ if and only if 0 is an L–space slope for $\mathrm{T}^n(\ell)$. By Lemma 4.8, this happens when $\mathrm{T}^n(\ell)$ contains a word from the set $A_1 \cup A_3$ or from the set $A_2 \cup A_4$, but not both. Recall from the proof of Lemmas 4.10 and 4.11 that A_1 and A_2 are closed under T^{-1} and A_3 and A_4 are closed under T.

Suppose that $n_1 < n_2 < n_3$, where n_2 is an L–space slope for ℓ while n_1 and n_3 are non-L–space slopes; we will show that all $n < n_1$ and all $n > n_3$ are non-L–space slopes. Up to reversing the loop, we can assume $\mathrm{T}^{n_2}(\boldsymbol{\ell})$ contains a word in $A_1 \cup A_3$ and not in $A_2 \cup A_4$. We consider the case that $\mathrm{T}^{n_2}(\boldsymbol{\ell})$ contains a word in A_1; the case that it contains a word in A_3 is similar and left to the reader.

By the closure property mentioned above (Lemma 4.10), $\mathrm{T}^{n_1}(\boldsymbol{\ell})$ contains a word in A_1. Since n_1 is a non-L–space slope, $\mathrm{T}^{n_1}(\boldsymbol{\ell})$ must contain a word from $A_2 \cup A_4$. It does not contain a word in A_4, since $\mathrm{T}^{n_2}(\boldsymbol{\ell})$ does not and the property of containing a word in A_4 is preserved by T (Lemma 4.11), and so it must contain a word in A_2. It follows that, for any $n < n_1$, $\mathrm{T}^{n}(\boldsymbol{\ell})$ contains words in A_1 and A_2 and thus n is a non-L–space slope for $\boldsymbol{\ell}$.

Since the property of containing a word in A_2 is preserved by T^{-1} (Lemma 4.10), we know that $\mathrm{T}^{n_3}(\boldsymbol{\ell})$ does not contain a word in A_2. We would like to show that $\mathrm{T}^{n_3}(\boldsymbol{\ell})$ contains a word in A_3, since this would imply it also contains a word in A_4; it would follow that for all $n > n_3$ the loop $\mathrm{T}^{n}(\boldsymbol{\ell})$ contains words from both A_3 and A_4 and thus n is a non-L–space slope for $\boldsymbol{\ell}$. Suppose, to the contrary, that $\mathrm{T}^{n_3}(\boldsymbol{\ell})$ does not contain a word in A_3. Then we see immediately that T^{n_2} does not contain a word in A_3. It also does not contain a word in A_2, so it must consist only of the segments

$$\{a_1, a_{-1}, b_1, b_{-1}, d_1, d_0, d_{-1}, c_k\},$$

where k can be any integer. Moreover, since $n_2 + 1 \le n_3$, $\mathrm{T}^{n_2+1}(\boldsymbol{\ell})$ does not contain a word in A_3. In particular this implies that $\mathrm{T}^{n_2}(\boldsymbol{\ell})$ does not contain d_1. Any d_{-1} cannot be preceded by d_{-1} or a_{-1} or followed by a d_{-1} or b_{-1}, since this would give a word in A_2. The alternative, that d_{-1} is preceded by a d_0 or a_1 and followed by a d_0 or b_1, is also ruled out since this would produce a word in A_3 in $\mathrm{T}^{n_2+1}(\boldsymbol{\ell})$; thus $\mathrm{T}^{n_2}(\boldsymbol{\ell})$ does not contain d_{-1}. If $\mathrm{T}^{n_2}(\boldsymbol{\ell})$ contains d_0, then it contains either $a_1 d_0^m b_1$, $a_{-1} d_0^m b_1$, $a_1 d_0^m b_{-1}$ or $a_{-1} d_0^m b_1$. None of these are possible; the first two are words in A_3 and A_2, respectively, and the last two give rise to words in A_3 in $\mathrm{T}^{n_2+1}(\boldsymbol{\ell})$. We have now shown that $\mathrm{T}^{n_2}(\boldsymbol{\ell})$ does not contain any d_k segments, but, since it also does not contain $a_{-1}b_{-1} \in A_2$ this contradicts the fact that $\mathrm{T}^{n_1}(\boldsymbol{\ell})$ has a word in A_2. Therefore, $\mathrm{T}^{n_3}(\boldsymbol{\ell})$ contains a word in A_3 and a word in A_4, and n is a non-L–space slope for all $n > n_3$. □

Proof of Theorem 4.1 Any two distinct slopes r/s and p/q determine two intervals in $\widehat{\mathbb{Q}}$; we need to show that if r/s and p/q are L–space slopes then one of these intervals lies entirely in the set of L–space slopes for $\boldsymbol{\ell}$. By reparametrizing with T and H as discussed above, we may assume that $r/s = 0$. We will assume $p/q > 0$ below;

similar arguments apply when $p/q < 0$. If $p/q = \infty$, the result holds by Corollary 4.5, so we will assume $p/q < \infty$.

Suppose that 0 and p/q are L–space slopes and there exist non-L–space slopes $u \in (0, p/q)$ and $v \in [0, p/q]^c$. To reach a contradiction, we choose a continued fraction $[a_1, \dots, a_n]$ for p/q with $a_1 \geq 0$ and $a_i > 0$ for $i > 1$ and proceed by induction on the length n. In the base case of $n = 1$, $p/q = a_1$ is an integer. We claim that there is an integer $u' \in (0, a_1)$ that is a non-L–space slope. To see this, if u is not an integer, consider the slopes $\lfloor u \rfloor$ and $\lceil u \rceil$. These slopes are distance 1 and both intervals between them contain non-L–space slopes ($u \in (\lfloor u \rfloor, \lceil u \rceil)$ and $v \in [\lfloor u \rfloor, \lceil u \rceil]^c$). Thus, by Corollary 4.5 (taking into account Remark 4.16), at least one of them is a non-L–space slope. Similarly, if $v \neq \infty$ then either $\lfloor v \rfloor$ or $\lceil v \rceil$ is a non-L–space slope. If $v = \infty$ is a non-L–space slope, note that by Proposition 4.3 either $\boldsymbol{\ell}$ contains no unstable chains or both types of unstable chains in standard notation or $\boldsymbol{\ell}$ cannot be written in standard notation. If $\boldsymbol{\ell}$ cannot be written in standard notation (in which case it is a collection of e^* segments) or if $\boldsymbol{\ell}$ contains no unstable chains in standard notation, then $\boldsymbol{\ell}$ is fixed by $\mathrm{T}^{\pm 1}$. It follows that all integer slopes behave the same, but this contradicts the fact that 0 is an L–space slope and u' is not. Thus $\boldsymbol{\ell}$ must contain both types of unstable chains in standard notation. Applying T^m for m sufficiently large produces a loop where all unstable chains are of type d or $\bar{d}$ (with large index), and there is at least one of each type. It follows from Corollary 4.9 that m is a non-L–space slope, since $d_k \in A_3$ and $\bar{d}_k = c_{-k} \in A_4$. Thus we have integers $0 < u' < a_1$ and $v' < 0$ or $v' > a_1$ such that 0 and a_1 are L–space slopes and u' and v' are not; this contradicts Lemma 4.17.

Suppose the continued fraction $[a_1, \dots, a_n]$ for p/q has length $n > 1$ and $a_1 = 0$. We consider the loop $\boldsymbol{\ell}' = \mathrm{H}^{a_2}(\boldsymbol{\ell})$. Note that the slope p/q for $\boldsymbol{\ell}$ corresponds to the slope $p'/q' = p/(q - a_2 p)$ for $\boldsymbol{\ell}'$, which has continued fraction $[a_3, \dots, a_n]$. The slope 0 for $\boldsymbol{\ell}$ corresponds to 0 for $\boldsymbol{\ell}'$, and the slopes u and v for $\boldsymbol{\ell}$ correspond to slopes u' and v' for $\boldsymbol{\ell}'$. Thus $\boldsymbol{\ell}'$ has L–space slopes 0 and p'/q' and non-L–space slopes $u' \in (0, p'/q')$ and $v' \in (-\infty, 0) \cup (p'/q', \infty]$. Since p'/q' has a continued fraction of length $n-2$, this produces a contradiction by induction.

Suppose p/q has continued fraction $[a_1, \dots, a_n]$ of length $n > 1$ with $a_1 > 0$. Consider the distance 1 slopes $a_1 = \lfloor p/q \rfloor$ and $a_1 + 1 = \lceil p/q \rceil$; both intervals between these slopes contain L–space slopes (either 0 or p/q), so by Proposition 4.6 one of them must be an L–space slope. First suppose that a_1 is an L–space slope. By the base case of induction, we cannot have both $u \in (0, a_1)$ and $v \in [0, p/q]^c \subset [0, a_1]^c$, and so we must have $u \in (a_1, p/q)$. Consider the loop $\boldsymbol{\ell}' = \mathrm{T}^{a_1}(\boldsymbol{\ell})$. The slopes 0 and

$p'/q' = p/q - a_1$ are L–space slopes for ℓ' while the slopes $u' = u - a_2 \in (0, p'/q')$ and $v' = v - a_2 \in [0, p'/q']^c$ are non-L–space slopes. A continued fraction for p'/q' is $[0, a_2, \dots, a_n]$; as shown above, this produces a contradiction.

Finally, suppose instead that a_1 is a non-L–space slope for ℓ and $a_1 + 1$ is an L–space slope. The base case of induction rules out the possibility that $v \in [0, a_1 + 1]^c$, so we must have $v \in (p/q, a_1 + 1)$. Consider the loop $\ell' = \mathrm{H}^{a_2 - 1} \circ \mathrm{T}^{a_1+1}(\ell)$. The slope a_1 for ℓ corresponds to the slope 0 for ℓ', and the slope p/q corresponds to a slope p'/q' with continued fraction $[-1, 1, a_3, \dots, a_n] \sim [0, -a_3, \dots, -a_n]$. There are non-$L$–space slopes for ℓ' in both intervals between 0 and p'/q'. By induction (using the analogue of the above cases when $p/q < 0$), this is a contradiction.

This proves that the set of L–space slopes for a loop ℓ is an interval. To prove the statement for bordered manifolds (M, α, β) of loop-type, we simply observe that if M is a loop-type manifold, $M(p\alpha + q\beta)$ is an L–space if and only if p/q is an L–space slope for each loop ℓ in $\widehat{CFD}(M, \alpha, \beta)$. The set of L–space slopes for M is the intersection of the intervals of L–space slopes for each loop, and hence an interval. □

4.4 Strict L–space slopes

The notion of L–space slope is fairly natural, however, in the context of the theorems in this paper, it is not quite the right condition. We will be interested primarily in slopes that are not only L–space slopes but are also surrounded by a neighbourhood of L–space slopes.

Definition 4.18 A slope in the boundary of a three-manifold M with torus boundary is a strict L–space slope if it is an L–space slope in the interior of $\mathcal{L}_M$. Denote the set of strict L–space slopes by $\mathcal{L}_M^\circ$.

We will give a geometric interpretation of strict L–space slopes in Section 7. For now, we will use loop notation and the results in this section about L–space slopes to determine when an L–space slope is strict.

Proposition 4.19 *Given a loop ℓ:*

(1) *The slope ∞ is a strict L–space slope for ℓ if and only if ℓ can be written in standard notation using only the subwords*

$$d_k, \quad b_1 a_{-1} \quad \text{and} \quad b_{-1} a_1,$$

where k can be any integer, with at least one d_k and such that $b_1 a_{-1}$ is never adjacent to $b_{-1} a_1$.

(2) *The slope* 0 *is a strict* L*–space slope for* ℓ *if and only if* ℓ *can be written in dual notation using only the subwords*

$$d_k^*, \quad b_1^* a_{-1}^* \quad \text{and} \quad b_{-1}^* a_1^*,$$

where k *can be any integer, with at least one* d_k^* *and such that* $b_1^* a_{-1}^*$ *is never adjacent to* $b_{-1}^* a_1^*$*.*

Proof Suppose ∞ is a strict L–space slope. In particular, it is an L–space slope, so, by Proposition 4.3, ℓ can be written in standard notation with no c_k and at least one d_k. Since it does not contain c_k, ℓ can be broken into pieces of the form d_k or $b_i a_j$ for any integer k and nonzero integers i and j, with at least one d_k. We also have that n and $-n$ are L–space slopes for sufficiently large n. Equivalently, 0 is an L–space slope after we apply T^n or T^{-n} for sufficiently large n. The twist T^n has the effect of replacing all unstable chains with chains of type d_k for $k \gg 0$, while the twist T^{-n} replaces all unstable chains with d_k for $k \ll 0$.

Consider a loop consisting of the pieces d_k with $k \gg 0$ and $b_i a_j$ with any nonzero integers i and j, with at least one d_k. Referring to the discussion of dualizing in Section 3.2, note that the loop certainly contains $d_0^* = e^*$, since it contains d_k with $k > 1$. Thus, by Proposition 4.3, 0 is an L–space slope if and only if the loop contains no c_k^* segments. The loop contains c_k^* with $k > 0$ if and only if it contains $a_i b_j$ with $i, j < 0$. It contains c_k^* with $k < 0$ if and only if it contains $b_i a_j$ with $i, j < 0$. It contains $c_0^* = \bar{e}^*$ segments if and only if it contains a_ℓ or b_ℓ with $\ell < -1$.

Similarly, consider a loop consisting of the pieces d_k with $k \ll 0$ and $b_i a_j$ with any nonzero integers i and j, with at least one d_k. The loop contains $c_0^* = \bar{e}^*$, so 0 is an L–space slope if and only if the loop contains no d_k^* segments. Moreover, the loop contains d_k^* if and only if it contains $b_i a_j$ or $a_i b_j$ with $i, j > 0$ or a_ℓ or b_ℓ with $\ell > 1$.

Therefore, given a loop for which ∞ is an L–space slope, n and $-n$ are also L–space slopes for all sufficiently large n if and only if the loop contains no chains of type a_k or b_k with $k > 1$, a_1 segments are never adjacent to b_1 segments, and a_{-1} segments are never adjacent to b_{-1} segments.

The proof of part (2) is completely analogous. Given that 0 is an L–space slope, we must check that $1/n$ and $-1/n$ are L–space slopes for sufficiently large n, which is equivalent to checking that ∞ is an L–space slope after applying H^n or H^{-n}. These twists have the effect of replacing unstable chains in dual notation with d_k^* segments with either $k \gg 0$ or $k \ll 0$. The rest of the proof is identical to the proof above after adding/removing stars on each segment. □

4.5 Simple loops

We will often restrict to a special class of loops, which we call simple.

Definition 4.20 A loop $\boldsymbol{\ell}$ is *simple* if there is a loop $\boldsymbol{\ell}'$ consisting only of unstable chains such that $\boldsymbol{\ell}$ can be obtained from $\boldsymbol{\ell}'$ using the operations $\mathrm{T}^{\pm 1}$ and $\mathrm{H}^{\pm 1}$. A collection of loops $\{\boldsymbol{\ell}_i\}_{i=1}^n$ is said to be simple if, possibly after applying a sequence of the operations $\mathrm{T}^{\pm 1}$ and $\mathrm{H}^{\pm 1}$, every loop consists only of unstable chains.

Remark 4.21 In the above definition, we do not specify whether the unstable chains are in standard or dual notation. In fact, it does not matter: if $\boldsymbol{\ell}$ contains only unstable chains in standard notation then $\mathrm{T}^n(\boldsymbol{\ell})$ contains only unstable chains in dual notation for sufficiently large n, and if $\boldsymbol{\ell}$ contains only dual unstable chains then $\mathrm{H}^n(\boldsymbol{\ell})$ contains only standard unstable chains. Definition 4.20 can be stated in terms of standard notation only (compare [10, Definition 4]), but it is sometimes convenient to check the condition in dual notation.

The notion of simple loops gives rise to a refinement of loop-type manifolds: M is said to be of *simple loop-type* if it is loop-type and, for some choice of α and β, the loops representing $\widehat{CFD}(M,\alpha,\beta)$ consist only of unstable chains. Equivalently, M is of simple loop-type if, for any choice of α and β, $\widehat{CFD}(M,\alpha,\beta)$ is represented by a simple collection of loops.

Note that solid torus-like loops (Definition 3.20) are examples of simple loops. Although we do not give an explicit description of all simple loops, we prove the following useful property:

Proposition 4.22 *If $\boldsymbol{\ell}$ is simple then, up to reorienting the loop, $\boldsymbol{\ell}$ has no a_k or b_k segments with $k < 0$.*

The proof makes use of the following two observations.

Lemma 4.23 *If $\boldsymbol{\ell}$ contains no a_k, b_k or d_k segments with $k < 0$, then the same is true for $\mathrm{T}(\boldsymbol{\ell})$ and $\mathrm{H}(\boldsymbol{\ell})$.*

Proof Recall that an a_k segment with $k < 0$ corresponds to a segment of type $\bar{a}$, a b_k with $k < 0$ corresponds to type $\bar{b}$, and a d_k with $k < 0$ corresponds to type $\bar{c}$. The statement for $\mathrm{T}(\boldsymbol{\ell})$ is obvious, since T fixes a_k and b_k segments and increases the subscript on d_k segments by one. For the second, we observe that $\boldsymbol{\ell}$ contains no a_i^*, b_i^* or c_i^* segments in dual notation with $i > 0$. Indeed, an a_i^* or a c_i^* contains a backwards

ρ_3 arrow, which in standard notation implies the presence of a segment of type $\bar{a}$ or $\bar{c}$ (equivalently, a segment of type a_k or d_k with $k < 0$), and a b^* segment contains a forward ρ_1 segment, which implies the presence of type $\bar{b}$ or $\bar{c}$ segment (equivalently, of a b_k or d_k with $k < 0$). It is clear that this property is preserved by H, which fixes a^* and b^* segments and decreases the index on c^* segments. Since $\mathrm{H}(\boldsymbol{\ell})$ has no a_i^*, b_i^* or c_i^* segments with $i > 0$, it is straightforward to check that $\mathrm{H}(\boldsymbol{\ell})$ has no forward ρ_1 arrows or backward ρ_3 arrows, and thus it does not contain any a_k, b_k or d_k segments with $k < 0$. $\square$

Lemma 4.24 *If $\boldsymbol{\ell}$ contains no a_k, b_k or c_k segments with $k < 0$, then the same is true for $\mathrm{T}^{-1}(\boldsymbol{\ell})$ and $\mathrm{H}^{-1}(\boldsymbol{\ell})$.*

Proof The proof is completely analogous to the previous lemma. $\square$

Proof of Proposition 4.22 Since $\boldsymbol{\ell}$ is simple, there is some $\boldsymbol{\ell}'$ consisting of only unstable chains such that $\boldsymbol{\ell}$ is obtained from $\boldsymbol{\ell}'$ by a sequence of Dehn twists. That sequence of Dehn twists determines an element of the mapping class group, which can be represented by the matrix $\left(\begin{smallmatrix} p & q \\ r & s \end{smallmatrix}\right)$. Let $[k_1, \dots, k_{2n}]$ be a continued fraction for p/q of even length; the element of the mapping class group above can be decomposed as the sequence of Dehn twists

$$\mathrm{T}^m \circ \mathrm{H}^{k_{2n}} \circ \mathrm{T}^{k_{2n-1}} \circ \cdots \circ \mathrm{H}^{k_2} \circ \mathrm{T}^{k_1},$$

where m is an integer. Let

$$\boldsymbol{\ell}'' = \mathrm{H}^{k_{2n}} \circ \mathrm{T}^{k_{2n-1}} \circ \cdots \circ \mathrm{H}^{k_2} \circ \mathrm{T}^{k_1}(\boldsymbol{\ell}).$$

First suppose that p/q is positive. We may assume that each k_i is positive for $1 \le i \le 2n$. Then $\boldsymbol{\ell}'$ can be written with no a_k, b_k or d_k segments with $k < 0$, and, by Lemma 4.23, this property is closed under all positive twists, so $\boldsymbol{\ell}''$ can also be written with no a_k, b_k or d_k segments with $k < 0$. It follows that $\boldsymbol{\ell} = \mathrm{T}^m(\boldsymbol{\ell}'')$ can be written with no a_k or b_k segments with $k < 0$.

If instead p/q is negative, we may chose a continued fraction with each k_i negative. Then $\boldsymbol{\ell}'$ can be written with no a_k, b_k or c_k segments with $k < 0$, and, by Lemma 4.24, the same is true for $\boldsymbol{\ell}''$. It follows that $\boldsymbol{\ell} = \mathrm{T}^m(\boldsymbol{\ell}'')$ can be written with no a_k or b_k segments with $k < 0$. $\square$

The argument above can be repeated in dual notation, taking a continued fraction for r/s instead of p/q, to prove the following:

Proposition 4.25 *If $\boldsymbol{\ell}$ is simple then, up to reorienting the loop, $\boldsymbol{\ell}$ has no a_k^* or b_k^* segments with $k < 0$.*

The main advantage of restricting to simple loops is that it greatly simplifies the conditions under which a given slope is a strict L–space slope. Recall that, by Proposition 4.19, ∞ is a strict L–space slope for $\boldsymbol{\ell}$ if and only if $\boldsymbol{\ell}$ can be decomposed into words of the form d_k, $b_1 a_{-1}$ or $b_{-1} a_1$. If $\boldsymbol{\ell}$ is a simple loop, then the last two words cannot appear by Proposition 4.22, and ∞ is a strict L–space slope if and only if $\boldsymbol{\ell}$ consists only of unstable chains in standard notation. Equivalently (using Observation 3.9), ∞ is a strict L–space slope if and only if $\boldsymbol{\ell}$ does not contain both positive and negative subscripts in dual notation. Similarly, 0 is a strict L–space slope if and only if $\boldsymbol{\ell}$ consists only of unstable chains in dual notation, which is equivalent to $\boldsymbol{\ell}$ having only nonnegative or only nonpositive subscripts in standard notation.

We conclude this section by refining Theorem 4.1 in the case of simple loops. Theorem 4.1 states that set of L–space slopes for a loop $\boldsymbol{\ell}$ is a (possibly empty) interval in $\widehat{\mathbb{Q}}$; the endpoints of this interval could be irrational, a priori, but we find that only rational endpoints are possible.

Proposition 4.26 *Let $\boldsymbol{\ell}$ be a simple loop. The set of L–space slopes for $\boldsymbol{\ell}$ is either*

(i) *identified with $\mathbb{Q}$ (in practice, every slope other than the rational longitude); or*

(ii) *the restriction to $\widehat{\mathbb{Q}}$ of a closed interval U in $\widehat{\mathbb{R}}$ with rational endpoints.*

Proof Note that applying the twist operations T and H to a loop preserves the cyclic ordering on abstract slopes, and so properties (i) and (ii) are preserved under these operations. Given an arbitrary simple loop $\boldsymbol{\ell}$, we will describe an algorithm for applying twists to produce a loop $\boldsymbol{\ell}'$ with one of the following properties:

(1) $\boldsymbol{\ell}'$ contains only d_0 segments;

(2) $\boldsymbol{\ell}'$ contains only d_1, d_0 and d_{-1} segments, with d_1 and d_{-1} segments alternating (ignoring d_0 segments); or

(3) $\boldsymbol{\ell}'$ contains only d_k segments with $k \geq -1$, including at least one d_{-1} segment, such that each d_{-1} is followed by $(d_0)^j d_k$ for some $j \geq 0$ and $k > 0$, and $\boldsymbol{\ell}'$ does not satisfy (2).

By definition, some sequence of twists produces a loop $\boldsymbol{\ell}_1$ which can be written in standard notation using only type d_k segments. We may assume that the minimum subscript for the d_k segments in $\boldsymbol{\ell}_1$ is 0. Given $\boldsymbol{\ell}_i$ (starting with $i = 1$), the algorithm proceeds as follows: If the subscripts in $\boldsymbol{\ell}_i$ are all 0, then $\boldsymbol{\ell}_i$ satisfies (1) and the algorithm stops. Otherwise, consider the loop $\mathrm{T}^{-1}(\boldsymbol{\ell}_i)$; this loop consists of d_k segments with $k \geq -1$ including at least one d_{-1}. If $\mathrm{T}^{-1}(\boldsymbol{\ell}_1)$ satisfies (2) or (3), the algorithm stops. Otherwise, consider the loop $\mathrm{E}(\boldsymbol{\ell}_i)$, which by Lemma 3.10 contains only d_k^* segments with $k \leq 0$ in dual notation. It follows that $\mathrm{E}(\boldsymbol{\ell}_i)$ contains only c_k segments with $k \geq 0$ in standard notation; reversing the loop, we have that $\mathrm{E}(\boldsymbol{\ell}_i)$ can be written with only d_k segments with $k \leq 0$ in standard notation. Let m_i denote the minimum subscript for a d_k segment in $\mathrm{E}(\boldsymbol{\ell}_i)$, and define $\boldsymbol{\ell}_{i+1}$ to be $\mathrm{T}^{-m_i} \circ \mathrm{E}(\boldsymbol{\ell}_i)$. Note that $\boldsymbol{\ell}_{i+1}$ consists of d_k segments with $k \geq 0$, including at least one d_0. We now repeat the algorithm using $\boldsymbol{\ell}_{i+1}$.

To see that this algorithm terminates, let κ_i denote the number of d_0 segments in the loop $\boldsymbol{\ell}_i$. Observe that d_0 segments in $\boldsymbol{\ell}_{i+1}$ come from minimal subscripts in $\mathrm{E}(\boldsymbol{\ell}_i)$, which come from maximal sequences of d_0's in $\boldsymbol{\ell}_i$; in particular, there is at least one d_0 in $\boldsymbol{\ell}_i$ for each d_0 in $\boldsymbol{\ell}_{i+1}$, and so $\kappa_{i+1} \leq \kappa_i$. Moreover, the inequality is strict unless $\boldsymbol{\ell}_i$ has no consecutive d_0 segments. If $\mathrm{T}^{-1}(\boldsymbol{\ell}_i)$ satisfies (2) or (3) then the algorithm terminates; otherwise, $\boldsymbol{\ell}_i$ must contain $d_0(d_1)^j d_0$ for some $j \geq 0$. Let λ_i denote the minimal such j. As observed above, if $\lambda_i = 0$ then $\kappa_{i+1} < \kappa_i$. If $\lambda_i > 0$, then the subword $d_0(d_1)^{\lambda_i} d_0$ in $\boldsymbol{\ell}_i$ gives rise to the subword $d_{-2}(d_{-1})^{\lambda_i - 1} d_{-2}$ in $\mathrm{E}(\boldsymbol{\ell}_i)$ and the subword $d_0(d_1)^{\lambda_i - 1} d_0$ in $\boldsymbol{\ell}_{i+1}$. Thus, at each step in the algorithm, either $\kappa_{i+1} < \kappa_i$ or $\kappa_{i+1} = \kappa_i$ and $\lambda_{i+1} < \lambda_i$. Since κ_i and λ_i are nonnegative integers, the algorithm must terminate after finitely many steps.

Let $\boldsymbol{\ell}'$ be the result of the algorithm above. In cases (1) and (2), we can easily check that condition (i) is satisfied. First note that ∞ is an L–space slope for $\boldsymbol{\ell}'$ since it contains only unstable chains in standard notation. Moreover, for any $n \in \mathbb{Z}$, $1/n$ is an L–space slope for $\boldsymbol{\ell}'$ if and only if ∞ is an L–space slope for $\mathrm{H}^n(\boldsymbol{\ell}')$. But $\mathrm{H}^n(\boldsymbol{\ell}') = \boldsymbol{\ell}'$ since $\boldsymbol{\ell}'$ either cannot be written in dual notation (case (1)) or contains only stable chains in dual notation (case (2)). Since the set of L–space slopes is an interval and it contains slopes arbitrarily close to 0 on both sides, it must be all of $\widehat{\mathbb{Q}} \setminus \{0\}$. In case (3), note that writing $\boldsymbol{\ell}'$ in dual notation produces stable chains, but each b^* segment is immediately followed by an a^* segment. Moreover, since $\boldsymbol{\ell}'$ does not satisfy (2), it has at least one unstable chain in dual notation. It follows from Propositions 4.3 and 4.19 that 0 is an L–space slope for $\boldsymbol{\ell}'$ but not a strict L–space slope. Thus 0 must be a boundary of the interval of L–space slopes. Since $\boldsymbol{\ell}'$ was obtained from $\boldsymbol{\ell}$ by applying a finite number of twists, the slope 0 for $\boldsymbol{\ell}'$ can be expressed as a rational slope for $\boldsymbol{\ell}$.

In case (3), we have found one rational boundary of the interval of L–space slopes. In fact, we can check that it is the left boundary. In this case $\chi_\circ(\boldsymbol{\ell}')$ and $\chi_\bullet(\boldsymbol{\ell}')$ have the same sign, and so the slope of the rational longitude $-\chi_\circ/\chi_\bullet$ is negative. By Corollary 4.5 it follows that the set of L–space slopes for ℓ' contains $[0,\infty]$, and so 0 must be the left boundary of the interval of L–space slopes for ℓ'. A similar algorithm, with opposite signs for subscripts in property (3), shows that the right endpoint is also rational. □

The proof of Theorem 1.2 is now complete: it follows from (and is made more precise by) Theorem 4.1 in combination with Proposition 4.26

5 Gluing results

This section is devoted to proving Theorem 1.3. We first prove the analogous result on the level of abstract loops, and then deduce the gluing theorem for simple loop-type manifolds. We end the section with an application to generalized splicing of L–space knot complements and give the proof of Theorem 1.6.

5.1 A gluing result for abstract loops

We will say that two loops ℓ_1 and ℓ_2 are *L–space aligned* if, for every slope $r/s \in \widehat{\mathbb{Q}}$, either r/s is a strict L–space slope for ℓ_1 or s/r is a strict L–space slope for ℓ_2. This section is devoted to proving the following proposition:

Proposition 5.1 *If ℓ_1 and ℓ_2 are simple loops which are not solid-torus-like, then $\ell_1^A \boxtimes \ell_2$ is an L–space chain complex if and only if ℓ_1 and ℓ_2 are L–space aligned.*

An essential observation is that ℓ_1 and ℓ_2 are L–space aligned if and only if $\textsc{t}(\ell_1)$ and $\textsc{h}(\ell_2)$ are L–space aligned, since $\textsc{t}$ takes the slope r/s to $(r+s)/s$ while $\textsc{h}$ takes the slope s/r to $s/(r+s)$. More generally, we can apply a sequence of twists $\{\textsc{t}^{k_1}, \textsc{h}^{k_2}, \dots, \textsc{t}^{k_{2n-1}}, \textsc{h}^{k_{2n}}\}$ to ℓ_1 and a corresponding sequence of twists $\{\textsc{h}^{k_1}, \textsc{t}^{k_2}, \dots, \textsc{h}^{k_{2n-1}}, \textsc{t}^{k_{2n}}\}$ to ℓ_2 without changing whether or not the pair of loops is L–space aligned. By Proposition 2.6, the quasi-isomorphism type of $\ell_1^A \boxtimes \ell_2$ is also unchanged. Thus, in proving Proposition 5.1, we may first reparametrize the pair of loops to get a more convenient form.

Proof of Proposition 5.1, "only if" direction Suppose ℓ_1 and ℓ_2 are not L–space aligned; that is, there exists a slope p/q that is not a strict L–space slope for ℓ_1 and q/p is not a strict L–space slope for ℓ_2. In fact, we may assume that $p/q = \infty$; if not, we reparametrize as described above, replacing ℓ_1 with

$$\ell_1' = \textsc{h}^{k_{2n}} \circ \textsc{t}^{k_{2n-1}} \circ \cdots \circ \textsc{h}^{k_2} \circ \textsc{t}^{k_1}(\ell_1)$$

so that the slope p/q for $\boldsymbol{\ell}_1$ becomes the slope ∞ for $\boldsymbol{\ell}_1'$, and replacing $\boldsymbol{\ell}_2$ with

$$\boldsymbol{\ell}_2' = \mathrm{T}^{k_{2n}} \circ \mathrm{H}^{k_{2n-1}} \circ \cdots \circ \mathrm{T}^{k_2} \circ \mathrm{H}^{k_1}(\boldsymbol{\ell}_2).$$

Furthermore, we may assume that $\boldsymbol{\ell}_1$ contains no segments of type d_k, $\bar{d}_k$, e or $\bar{e}$, and that $\boldsymbol{\ell}_2$ contains no segments of type d_k^*, $\bar{d}_k^*$, e^* or $\bar{e}^*$; if necessary, we replace $\boldsymbol{\ell}_1$ with $\mathrm{T}^{-n}(\boldsymbol{\ell}_1)$ and $\boldsymbol{\ell}_2$ with $\mathrm{H}^{-n}(\boldsymbol{\ell}_2)$ for sufficiently large n.

Since ∞ is not a strict L–space slope for $\boldsymbol{\ell}_1$, the loop $\boldsymbol{\ell}_1$ must contain stable chains in standard notation (note that, since $\boldsymbol{\ell}_1$ is not solid torus-like, it can be written in standard notation). In particular, after possibly reversing the loop, $\boldsymbol{\ell}_1$ contains a b_k segment. The corresponding segment in $\boldsymbol{\ell}_1^A$ is

$$\begin{tikzcd} \underset{x_1}{\bullet} \arrow[r, "\rho_3{,}\rho_2{,}\rho_1"] & \underset{y_1}{\circ} \arrow[r, dashed, "\rho_2{,}\rho_1"] & \underset{y_k}{\circ} & \underset{x_2}{\bullet} \arrow[l, "\rho_3"'] \end{tikzcd}$$

Since 0 is not a strict L–space slope for $\boldsymbol{\ell}_2$, this loop must contain stable chains in dual notation. In particular it contains an a_ℓ^* segment; we label the corresponding generators

$$\begin{tikzcd} \underset{w_1}{\circ} & \underset{z_1}{\bullet} \arrow[l, "\rho_3"'] \arrow[r, dashed, "\rho_{12}"] & \underset{z_\ell}{\bullet} \arrow[r, "\rho_{123}"] & \underset{w_2}{\circ} \end{tikzcd}$$

Consider the generator y_k in $\boldsymbol{\ell}_1^A$, which has no outgoing $\mathcal{A}_\infty$ operations. To determine the possible incoming operations, note that the segment b_k must be followed by either a type c segment or a type a segment. This is because we assumed that $\boldsymbol{\ell}_1$ contains no $\bar{e}$ or $\bar{d}_j$ segments, and $\boldsymbol{\ell}_1$ cannot contain both b_k and $\bar{a}_j$ segments since it is simple. In either case, x_2 has an outgoing ρ_3 labelled arrow in $\boldsymbol{\ell}_1^A$ and no incoming arrows. It follows that y_k has only the incoming $\mathcal{A}_\infty$ operations

$$\begin{aligned} m_2(x_2, \rho_3) &= y_k, \\ m_{2+i}(y_{k-i}, \rho_2, \rho_{12}, \dots, \rho_{12}, \rho_1) &= y_k \quad \text{for } 1 \le i < k, \\ m_{3+k}(x_1, \rho_3, \rho_2, \rho_{12}, \dots, \rho_{12}, \rho_1) &= y_k, \end{aligned}$$

and possibly more operations whose inputs end with $\rho_2, \rho_{12}, \dots, \rho_{12}, \rho_1$.

Consider the generator w_2 in $\boldsymbol{\ell}_2$. Note that a_ℓ^* must be followed by either a b^* segment or a $\bar{c}^*$ segment. It follows that the only incoming sequences of arrows consist of a ρ_{123} or ρ_1 arrow preceded by some number of ρ_{12} arrows. Comparing this with the $\mathcal{A}_\infty$ operations terminating at y_k described above, it is clear that the generator $y_k \otimes w_2$ in $\boldsymbol{\ell}_1^A \boxtimes \boldsymbol{\ell}_2$ has no incoming differentials. It also has no outgoing differentials, since y_k has no outgoing $\mathcal{A}_\infty$ operations. Thus $y_k \otimes w_2$ survives in homology.

Similarly, consider the generator x_1 in ℓ_1^A and z_1 in ℓ_2. The generator z_1 has no incoming sequences of arrows, and the only outgoing sequences consist of a single ρ_3 arrow or begin with some number of ρ_{12} arrows followed by a ρ_{123}. Here we use the fact that the segment a_ℓ^* in the simple loop ℓ_2 can only be preceded by a b^* segment or a c^* segment, so the outgoing ρ_3 arrow cannot be followed by another outgoing arrow. In ℓ_1^A, the segment b_k must be preceded by an a segment or a $\bar{c}$ segment. It follows that, for any nontrivial operation $m_{n+1}(x_1, \rho_{I_1}, \dots, \rho_{I_n})$, we have

- $\rho_{I_1} \neq \rho_{123}$;
- if $\rho_{I_1} = \rho_{12}$, then $\rho_{I_i} = \rho_{12}$ for $1 \leq i \leq n-1$ and $\rho_{I_n} = \rho_1$;
- if $\rho_{I_1} = \rho_3$, then $n > 1$.

We see that no $\mathcal{A}_\infty$ operations starting at x_1 match with the δ^n maps starting at z_1. Thus the generator $x_1 \otimes z_1$ in $\ell_1^A \boxtimes \ell_2$ has no incoming or outgoing differentials and survives in homology.

Finally, we observe that $\mathrm{gr}(z_1) = \mathrm{gr}(w_2)$ since z_1 and w_2 are connected by only ρ_{12} and ρ_{123} arrows, but $\mathrm{gr}(y_k) = \mathrm{gr}(y_1) = -\mathrm{gr}(x_1)$, since arrows labelled (ρ_2, ρ_1) preserve grading but arrows labelled (ρ_3, ρ_2, ρ_1) flip grading. It follows that $y_k \otimes w_2$ and $x_1 \otimes z_1$ have opposite $(\mathbb{Z}/2\mathbb{Z})$–grading. Since both survive in homology, $\ell_1^A \boxtimes \ell_2$ is not an L–space complex. $\square$

To prove the converse we will use the fact that ℓ_1 and ℓ_2 are L–space aligned to put strong restrictions on the segments that may appear in the loops ℓ_1 and ℓ_2. Once again, we can apply twists to ℓ_1 and ℓ_2 to obtain ℓ_1' and ℓ_2' with convenient parametrizations, such that ℓ_1' and ℓ_2' are still L–space aligned and $\ell_1^A \boxtimes \ell_2$ is homotopy equivalent to $(\ell_1')^A \boxtimes \ell_2'$. The set of strict L–space slopes for ℓ_2 is some nonempty open interval in $\widehat{\mathbb{Q}}$. This interval is either all of $\widehat{\mathbb{Q}}$ except the rational longitude or it has distinct rational endpoints; see Proposition 4.26. In the latter case, we can reparametrize so that for ℓ_2' these boundaries have slope 0 and p/q for some $1 < p/q \leq \infty$. To see this, take $p > 0$ and choose n so that $0 \leq q + np < p$ and $1 < p/(q+np) \leq \infty$; we can replace p/q with $p/(q+np)$ by applying dual twists, in particular, leaving the slope 0 fixed. Now the set of strict L–space slopes for ℓ_2' is exactly $(0, p/q)$. The fact that ℓ_1' and ℓ_2' are L–space aligned then implies that the set of strict L–space slopes for ℓ_1' contains $[-\infty, q/p]$. If the set of strict L–space slopes for ℓ_2 is all of $\widehat{\mathbb{Q}}$ except the rational longitude, then we can choose a parametrization such that the rational longitude of ℓ_2' is 0 and such that the set of strict L–space slopes for ℓ_1' contains $[-\infty, 0]$.

Lemma 5.2 *If $q/p \in [0,1)$ and ℓ is a simple loop for which the interval of strict L–space slopes contains $[-\infty, q/p]$, then ℓ consists only of segments c_k^* with $0 \leq k \leq \lceil p/q \rceil$.*

Proof Since 0 is a strict L–space slope, ℓ can be written with only dual unstable chains. Up to reading the loop in reverse order, we can assume the unstable chains are c_k^* segments. Moreover, the fact that ∞ is a strict L–space slope implies that ℓ cannot contain c_k^* segments with both positive and negative subscripts. Since the rational longitude is given by $-\chi_\circ(\ell)/\chi_\bullet(\ell)$ and falls in the interval $(q/p, \infty)$, we must have that $\chi_\circ(\ell)$ and $\chi_\bullet(\ell)$ have opposite signs. This only happens if ℓ is composed of c_k^* segments with $k \geq 0$. Let $n = \lceil p/q \rceil$. Observe that ℓ must contain at least one c_k^* with $0 \leq k < n$ (recall that $c_0 = \bar{e}^*$), since otherwise the rational longitude is less than $1/n$. Finally, the fact that $1/n$ is a strict L–space slope implies that ∞ is a strict L–space slope for the loop $\mathrm{H}^n(\ell)$. Since ℓ contains $\bar{e}^*$ or c_k^* with $k < n$, $\mathrm{H}^n(\ell)$ contains at least one c_k^* with $k < 0$ and therefore does not contain any c_k^* with $k > 0$. Therefore ℓ does not contain c_k^* with $k > n$. □

Lemma 5.3 *If $p/q \in (1, \infty]$ and ℓ is a simple loop that is not solid torus-like for which the interval of strict L–space slopes contains $(0, p/q)$, then ℓ consists only of a_k, b_k, c_k and d_k segments (for $k > 0$) and e and $\bar{c}_1$ segments. Moreover,*

- *ℓ contains no two $\bar{c}_1 = d_{-1}$ segments separated only by $e = d_0$ segments,*
- *ℓ contains no c_k segments with $k < \lceil p/q \rceil - 1$, and*
- *if 0 is not a strict L–space slope for ℓ then there is at least one $\bar{c}_1$ segment.*

Proof Since 1 is a strict L–space slope, $\mathrm{T}(\ell)$ can be written with no $\bar{a}$, $\bar{b}$, $\bar{c}$ or $\bar{d}$ segments. It follows that ℓ can be written with no $\bar{a}$, $\bar{b}$, $\bar{d}$ or $\bar{e}$ segments and no $\bar{c}_k$ segments with $k > 1$.

Let $n = \lceil p/q \rceil$. Since $n - 1 < p/q$ is a strict L–space slope, $\mathrm{T}^{n-1}(\ell)$ does not contain both barred and unbarred segments (ignoring e's). Suppose ℓ contains c_k with $k < n-1$. Then $\mathrm{T}^{n-1}(\ell)$ contains at least one $\bar{d}$ segment and cannot contain any unbarred segments. Any a, b or c_k segments with $k > n-1$ in ℓ produces an unbarred segment in $\mathrm{T}^{n-1}(\ell)$, so we must have that ℓ consists only of c_k segments with $k \leq n-1$. However, in this case it is easy to see that $\chi_\bullet(\ell)$ and $\chi_\circ(\ell)$ have opposite signs and $|\chi_\circ(\ell)| \leq (n-1)|\chi_\bullet(\ell)|$, which contradicts the fact that the rational longitude $-\chi_\circ/\chi_\bullet$ does not fall in the interval $(0, p/q)$. Thus ℓ does not contain c_k with $k < n-1$.

Now $\boldsymbol{\ell}$ must contain an a_k, b_k, c_k or d_k segment with $k > 1$ or two segments of type a_1, b_1, c_1 or d_1 separated only by e's. Otherwise, $\boldsymbol{\ell}$ would consist of only $\bar{c}_1$, e and d_1 segments with at least one $\bar{c}_1$ for each d_1; in this case, $\chi_\bullet(\boldsymbol{\ell})$ and $\chi_\circ(\boldsymbol{\ell})$ have opposite signs and $|\chi_\circ(\boldsymbol{\ell})| \le |\chi_\bullet(\boldsymbol{\ell})|$, so the rational longitude falls in $(0, 1]$. It follows that in dual notation $\boldsymbol{\ell}$ contains $\bar{c}_k^*$, e^* or d_k^*, and thus $\mathrm{H}^m(\boldsymbol{\ell})$ contains a d^* segment for sufficiently large m. Since $1/m$ is a strict L–space slope for sufficiently large m, we must have that $\mathrm{H}^m(\boldsymbol{\ell})$ does not contain a $\bar{d}^*$ segment. Therefore $\boldsymbol{\ell}$ does not contain any c_k^* segments, and thus in standard notation it does not have two $\bar{c}_1$ segments separated only by e's.

Finally, if 0 is not a strict L–space slope for $\boldsymbol{\ell}$, then $\boldsymbol{\ell}$ must contain both barred and unbarred segments in standard notation; it follows that $\boldsymbol{\ell}$ must contain at least one $\bar{c}_1$. □

The two previous lemmas only depend on $\lceil p/q \rceil$. If p/q is not an integer, it is possible to give further restrictions on subwords that can appear in the loop. We will prove one such restriction using two properties for a pair of loops. For an integer $r \ge 0$, we will say that two loops $\boldsymbol{\ell}_1$ and $\boldsymbol{\ell}_2$ satisfy property λ (or property λ for $r \ge 0$) if:

(λ1) $\boldsymbol{\ell}_1$ consists only of c_k^* with $0 \le k \le n$, for some n, with at least one c_n^*.

(λ2) $\boldsymbol{\ell}_2$ consists only of a_k, b_k $(k > 0)$, d_k $(k \ge -1)$ and c_l $(l \ge m)$ segments for some $m > 0$, with at least one c_m and at least one $\bar{c}_1 = d_{-1}$.

(λ3) There is an integer $N > 0$ and subscripts $k_i \in \{N, N-1\}$ for $1 \le i \le r$ such that $\boldsymbol{\ell}_1$ contains the subword $c_N^* c_{k_1}^* \dots c_{k_r}^* c_N^*$ and $\boldsymbol{\ell}_2$ contains the subword $c_{N-1} c_{k_1} \dots c_{k_r} c_{N-1}$.

The integer $r \ge 0$ appearing in condition (λ3) is the complexity of the pair $(\boldsymbol{\ell}_1, \boldsymbol{\ell}_2)$ satisfying property λ; when we appeal to pairs satisfying property λ the aim will be to decrease this complexity. Similarly, two loops $\boldsymbol{\ell}_1$ and $\boldsymbol{\ell}_2$ satisfy property property λ^* (or property λ^* for $r \ge 0$) if:

(λ^*1) $\boldsymbol{\ell}_1$ consists only of $\bar{d}_k$ with $0 \le k \le n$, for some n, with at least one $\bar{d}_n$.

(λ^*2) $\boldsymbol{\ell}_2$ consists only of $\bar{a}_k^*$, $\bar{b}_k^*$ (for $k > 0$), $\bar{c}_k^*$ (for $k > -1$) and $\bar{d}_l^*$ (for $l \ge m$) segments for some m, with at least one $\bar{d}_m^*$ and at least one $d_1^* = \bar{c}_{-1}^*$.

(λ^*3) There is an integer $N > 0$ and subscripts $k_i \in \{N, N-1\}$ for $1 \le k \le r$ such that $\boldsymbol{\ell}_1$ contains the subword $\bar{d}_N \bar{d}_{k_1} \dots \bar{d}_{k_r} \bar{d}_N$ and $\boldsymbol{\ell}_2$ contains the subword $\bar{d}_{N-1}^* \bar{d}_{k_1}^* \dots \bar{d}_{k_r}^* \bar{d}_{N-1}^*$.

Lemma 5.4 *If two simple loops $\boldsymbol{\ell}_1$ and $\boldsymbol{\ell}_2$ satisfy either property λ or property λ^*, then $\boldsymbol{\ell}_1$ and $\boldsymbol{\ell}_2$ are not L–space aligned.*

$$\begin{array}{cccccc}
\boldsymbol{\ell}_1 & \bar{e}^* c_3^* c_2^* c_2^* c_1^* c_2^* & \xrightarrow{\mathrm{H}^2} & \bar{d}_2^* c_1^* \bar{e}^* \bar{e}^* \bar{d}_1^* \bar{e}^* & \boldsymbol{\ell}_1' \\
\boldsymbol{\ell}_2 & a_1 \bar{c}_1 b_1 c_1 c_2 c_1 c_1 & \xrightarrow{\mathrm{T}^2} & a_1 d_1 b_1 \bar{d}_1 \bar{e} \bar{d}_1 \bar{d}_1 & \boldsymbol{\ell}_2'
\end{array}$$

Figure 6: Loops $\boldsymbol{\ell}_1$ and $\boldsymbol{\ell}_2$ satisfying property λ for $n = 3$ and $m = 1$. The relevant subwords, with $N = r = 2$, have been highlighted. Note that this illustrates the key step in the proof of Lemma 5.4 when $n > m+1$: one checks that 0 is not a strict L–space slope for $\boldsymbol{\ell}_1'$ (this loop contains a $\bar{d}^*$ segment) while the set of strict L–space slopes for $\boldsymbol{\ell}_2'$ does not intersect $[0, \infty]$ thus $\boldsymbol{\ell}_1'$ and $\boldsymbol{\ell}_2'$ (and hence $\boldsymbol{\ell}_1$ and $\boldsymbol{\ell}_2$) are not L–space aligned.

Proof Suppose first that property λ is satisfied, where n is the maximum subscript for c_k^* segments in $\boldsymbol{\ell}_1$ and m is the minimum subscript of c_k segments in $\boldsymbol{\ell}_2$. Since $\boldsymbol{\ell}_1$ contains c_N^* and $\boldsymbol{\ell}_2$ contains c_{N-1} for some N, we have that $n \geq m+1$. Consider the reparametrized loops $\boldsymbol{\ell}_1' = \mathrm{H}^{m+1}(\boldsymbol{\ell}_1)$ and $\boldsymbol{\ell}_2' = \mathrm{T}^{m+1}(\boldsymbol{\ell}_2)$; $\boldsymbol{\ell}_1'$ and $\boldsymbol{\ell}_2'$ are L–space aligned if and only if $\boldsymbol{\ell}_1$ and $\boldsymbol{\ell}_2$ are. We now have:

- $\boldsymbol{\ell}_1'$ consists only of c_k^* with $-m \leq k \leq n-m-1$ with at least one c_{n-m-1}^* (where, as usual, $c_0^* = \bar{e}^*$).
- $\boldsymbol{\ell}_2'$ consists only of a_k, b_k (for $k > 0$), c_k (for $k \geq -1$) and d_l segments with $l \geq m$, with at least one $c_{-1} = \bar{d}_1$ and at least one d_m.

Note that, for $\boldsymbol{\ell}_2'$, 0 is not a strict L–space slope because the loop contains at least one barred segment, $\bar{d}_1$, and at least one unbarred segment, d_m. Note also that ∞ is not a strict L–space slope, since $\boldsymbol{\ell}_2'$ contains unstable chains with both orientations. The slope -1 is a strict L–space slope, since $\mathrm{T}^{-1}(\boldsymbol{\ell}_2')$ has no barred segments (ignoring e's). It follows that $\boldsymbol{\ell}_2'$ has no strict L–space slopes in $[0, \infty]$.

First consider the case that $n > m+1$. In this case, $\boldsymbol{\ell}_1'$ contains at least one c_k^* segment with $k > 0$. If $\boldsymbol{\ell}_1'$ also contains a c_k^* segment with $k < 0$, then 0 is not a strict L–space slope for $\boldsymbol{\ell}_1'$. Since ∞ is not a strict L–space slope for $\boldsymbol{\ell}_2'$, $\boldsymbol{\ell}_1'$ and $\boldsymbol{\ell}_2'$ are not L–space aligned. If $\boldsymbol{\ell}_1'$ does not contain a c_k^* segment with $k < 0$, then it consists only of c_k^* segments with $k \geq 0$. It follows that the rational longitude $-\chi_\circ(\boldsymbol{\ell}_1')/\chi_\bullet(\boldsymbol{\ell}_1')$ is positive. Since the rational longitude is not a strict L–space slope, and all positive slopes for $\boldsymbol{\ell}_2'$ are not strict L–space slopes, $\boldsymbol{\ell}_1'$ and $\boldsymbol{\ell}_2'$ are not L–space aligned.

In the case that $n = m+1$, we have $N = n$ in the statement of property λ (with complexity r) for $\boldsymbol{\ell}_1$ and $\boldsymbol{\ell}_2$. The shifted loops $\boldsymbol{\ell}_1'$ and $\boldsymbol{\ell}_2'$ satisfy an additional condition:

- There are subscripts $k_i' \in \{0,1\}$ for $1 \le i \le r$ such that ℓ_1' contains the subword $\bar{e}^*\bar{d}^*_{k_1'} \dots \bar{d}^*_{k_r'}\bar{e}^*$ and ℓ_2' contains the subword $\bar{d}_1\bar{d}_{k_1'} \dots \bar{d}_{k_r'}\bar{d}_1$, following (as usual) the convention that $\bar{d}_0 = \bar{e}$ and $\bar{d}_0^* = \bar{e}^*$.

The next step is to write ℓ_1' in standard notation and ℓ_2' in dual notation. First, ℓ_1' consists only of $\bar{e}^*$ and $\bar{d}^*$ segments, so in standard notation it consists only of $\bar{e}$ and $\bar{d}$ segments. There is some maximum subscript on the $\bar{d}$ segments; call it n'. Note that this says that ℓ_1' and ℓ_2' satisfy condition (λ^*1) for property λ^*. Dualizing ℓ_2' is slightly harder. Types of segments in dual notation are determined by adjacent pairs of standard segments (ignoring e's and $\bar{e}$'s); see Section 3.2. The possible pairs of segments in ℓ_2' are ab, ad, dd, db, ba, bc, $b\bar{d}$, ca, cc, $c\bar{d}$, $\bar{d}a$, $\bar{d}c$ and $\bar{d}\bar{d}$. These correspond to dual segments d^*, d^*, d^*, d^*, $\bar{c}^*$, $\bar{c}^*$, $\bar{b}^*$, $\bar{c}^*$, $\bar{c}^*$, $\bar{b}^*$, $\bar{a}^*$, $\bar{a}^*$ and $\bar{d}^*$, respectively. Since ℓ_2' has no e segments, the subscripts on the d^* segments are at most 1. Since the only bar segments in ℓ_2' in standard notation have subscript 1, there are no $\bar{e}^*$ segments in ℓ_2'. Thus ℓ_2' consists only of $\bar{a}_k^*$, $\bar{b}_k^*$, $\bar{c}_k^*$, $\bar{d}_k^*$ (for $k > 0$), e^* and d_1^* segments. There is some minimum subscript on the $\bar{d}$ segments; call it m'. Moreover, since ℓ_2' contains at least one d_k segment, it also contains at least one d_1^* segment. Note that this says that ℓ_1' and ℓ_2' satisfy condition (λ^*2) of property λ^*.

If $n' > m'+1$, then we proceed in a similar fashion to the $n > m+1$ case for property λ. We replace ℓ_1' with $\ell_1'' = \mathrm{T}^{-m'-1}(\ell_1')$ and ℓ_2' with $\ell_2'' = \mathrm{H}^{-m'-1}(\ell_2')$. We can observe that ℓ_2'' contains at least one c_1^* segment and at least one $\bar{c}_{m'}$ segment; thus 0 and ∞ are not strict L–space slopes for ℓ_2'', and in fact no slope in $[-\infty, 0]$ is a strict L–space slope. We can also observe that ℓ_1'' consists of c, $\bar{e}$ and $\bar{d}$ segments with at least one $\bar{d}$. Thus either 0 is not a strict L–space slope or the rational longitude is negative. In either case, ℓ_1'' and ℓ_2'' are not L–space aligned.

Now assume that $n' = m' + 1$. Consider the sequence $k_1', \dots, k_r'$. This sequence consists of some number of (possibly empty) strings of 0's each separated by a single 1; let $l_1, \dots, l_s$ be the sequence of the lengths of these strings of 0's. Note that s is at most $r+1$. Since ℓ_2' contains the word $\bar{d}_1\bar{d}_{k_1'} \dots \bar{d}_{k_r'}\bar{d}_1$, it follows that it contains the dual word $\bar{d}^*_{l_1}\bar{d}^*_{l_2} \dots \bar{d}^*_{l_{s-1}}\bar{d}^*_{l_s}$. Similarly, ℓ_1' contains the dual word $\bar{e}^*\bar{d}^*_{k_1'} \dots \bar{d}^*_{k_r'}\bar{e}^*$, which must be followed and proceeded by more $\bar{d}^*$ segments, possibly with additional $\bar{e}^*$ segments in between. It follows that ℓ_1' contains the word $\bar{d}_{l_1'}\bar{d}_{l_2} \dots \bar{d}_{l_{s-1}}\bar{d}_{l_s'}$, where $l_1' > l_1$ and $l_s' > l_s$. Since we have assumed that $n' = m' + 1$, we must have that $l_1' = l_1 + 1$ and $l_s' = l_s + 1$, and moreover that $l_1 = l_s$ and $l_i \in \{l_1, l_1 + 1\}$ for every other i. Let $N' = l_1 + 1$, so that ℓ_1' contains the subword $\bar{d}_{N'}\bar{d}_{l_2} \dots \bar{d}_{l_{s-1}}\bar{d}_{N'}$ and ℓ_2'

contains the subword $\bar{d}^*_{N'-1}\bar{d}^*_{l_1}\dots\bar{d}^*_{l_{s-1}}\bar{d}^*_{N'-1}$. In other words, the property (λ^*3) of property λ^* is satisfied with complexity $r' = s-2$. Note that $r' < r$, since $s \le r+1$. We also have that $r' \ge 0$, since if $s = 1$ then each k_i' is 0 for each $1 \le i \le r$; it would follow that ℓ_2' contains $\bar{d}^*_{r+1}$ and ℓ_1' contains a $\bar{d}$ segment with subscript at least $r+3$, and so $n' > m'+1$.

We have shown that if ℓ_1 and ℓ_2 satisfy property λ (for some integer $r \ge 0$), then either they are not L–space aligned or they can be modified to ℓ_1' and ℓ_2' which satisfy property λ^* (for some integer $r' \ge 0$), where $0 \le r' < r$. A similar proof shows that if ℓ_1 and ℓ_2 satisfy property λ^* then either they are not L–space aligned or they can be modified to ℓ_1' and ℓ_2' which satisfy property λ with $0 \le r' < r$. In both cases, the complexity is reduced, so, by induction on r, we have that ℓ_1 and ℓ_2 are not L–space aligned if they satisfy either property. $\square$

Now that we have placed restrictions on loops ℓ_1 and ℓ_2 that are L–space aligned, we can complete the proof of Proposition 5.1 by analyzing the box tensor product segment by segment. In the proof, we will determine when certain generators in $\ell_1^A \boxtimes \ell_2$ cancel in homology. To aid in this, we introduce the following terminology: we refer to differentials starting or ending at $x \otimes y$ in $\ell_1^A \boxtimes \ell_2$ as being "on the right" or "on the left" depending on whether the type D operations in ℓ_2 that give rise to the differential in the tensor product are to the right or left of y, relative to the cyclic ordering on ℓ_2. In

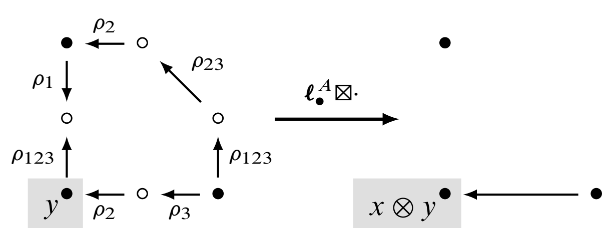

having fixed the loop $(d_2\bar{b}_1\bar{a}_1)$ (that is, reading the loop counterclockwise) the generator $x \otimes y$ cancels on the right when paired with the standard solid torus. (Recall that this example may be identified with the trivial surgery on the right-hand trefoil.) This terminology is motivated by picturing the tensor product on a grid with rows indexed by generators of ℓ_1^A and columns indexed by generators of ℓ_2; compare Figure 7. With this terminology in place, we observe:

Lemma 5.5 *For any loops ℓ_1 and ℓ_2 and any generator $x \otimes y \in \ell_1^A \boxtimes \ell_2$, there is at most one differential into or out of $x \otimes y$ on the right, and at most one on the left.*

Proof This follows from examining the arrows of type $\mathbf{I}_\bullet$, $\mathbf{II}_\bullet$, $\mathbf{I}_\circ$ and $\mathbf{II}_\circ$ introduced in Section 3.1 and the corresponding type A arrows. For instance, suppose y has a type $\mathbf{I}_\bullet$ arrow on the right. A differential on the right of $x \otimes y$ must arise from an $\mathcal{A}_\infty$ operation starting at x with inputs starting with ρ_1, ρ_{12} or ρ_{123}. Such $\mathcal{A}_\infty$ operations correspond to sequences of arrows in $\boldsymbol{\ell}_1^A$ starting at x in which the first arrow is labelled by ρ_1. There is at most one such arrow from x in $\boldsymbol{\ell}_1^A$, since the corresponding vertex in $\boldsymbol{\ell}_1$ is adjacent to only one arrow of type $\mathbf{II}_\bullet$. A directed sequence of arrows in $\boldsymbol{\ell}_1^A$ of length k beginning with this ρ_1 labelled arrow gives rise to a collection of k $\mathcal{A}_\infty$ operations with first input ρ_1, ρ_{12} or ρ_{123}, each corresponding to the first i arrows in the sequence for some $i \leq k$. We observe that none of these operations have inputs that are the first n inputs of another operation in the collection. The reason for this is, when reading off the inputs for an $\mathcal{A}_\infty$ operation from a sequence of arrows following the conventions in Section 2.4, the effect of adding one more arrow is to leave all but the last input in the list unchanged, multiply the last input by something nontrivial and possibly add more inputs at the end of the list. For example, for a sequence of arrows, in $\boldsymbol{\ell}_1^A$,

$$\underset{x}{\bullet} \xrightarrow{\rho_1} \circ \xrightarrow{\rho_2,\rho_1} \circ \xrightarrow{\rho_2} \bullet$$

the first arrow yields a contribution to $m_2(x, \rho_1)$, the first two arrows yield a contribution to $m_3(x, \rho_{12}, \rho_1)$ and all three arrows yields a contribution to $m_3(x, \rho_{12}, \rho_{12})$. Thus at most one operation in the collection can pair with a sequence of arrows on the right of y in $\boldsymbol{\ell}_2$. It follows that $x \otimes y$ has at most one differential on the right.

Similar arguments show that there is at most one differential on the right if the arrow on the right of y is type $\mathbf{II}_\bullet$, $\mathbf{I}_\circ$ or $\mathbf{II}_\circ$, and the same is true on the left. □

Thus, all differentials in $\boldsymbol{\ell}_1^A \boxtimes \boldsymbol{\ell}_2$ appear in linear chains and $x \otimes y$ will cancel in homology if it has an odd length chain of differentials on either side; we say that $x \otimes y$ cancels on the right (resp. left) if there is an odd length chain of differentials on the right (resp. left). Note that if there is an odd length chain of differentials on one side of $x \otimes y$, successively applying edge reduction to cancel the outermost differential on that side eventually results in cancelling $x \otimes y$ without using the differentials on the other side. If there is an even length chain, this chain can be removed using edge reduction leaving no differentials on that side of $x \otimes y$. If $x \otimes y$ does not cancel on the right or the left, we say $x \otimes y$ does not cancel in homology, since it, or potentially a linear combination of it with other generators of the same $(\mathbb{Z}/2\mathbb{Z})$–grading, survives in homology.

Proof of Proposition 5.1, "if" direction Suppose $\boldsymbol{\ell}_1$ and $\boldsymbol{\ell}_2$ are L–space aligned. Up to changing parametrization, we can assume that the interval of strict L–space slopes for $\boldsymbol{\ell}_2$ contains $(0, p/q)$ for some $1 < p/q \leq \infty$ and does not contain 0 and the interval of strict L–space slopes for $\boldsymbol{\ell}_1$ contains $[-\infty, q/p]$. By Lemmas 5.2 and 5.3, $\boldsymbol{\ell}_1$ can be written with only e^* and c_k^* segments and $\boldsymbol{\ell}_2$ consists of a_k, b_k, c_k, d_k, e and $\bar{c}_1$ segments. We will fix the $(\mathbb{Z}/2\mathbb{Z})$–grading on each loop so that every generator of $\boldsymbol{\ell}_1^A$ has grading 0 and all ι_1–generators in $\boldsymbol{\ell}_2$ have grading 0 except those coming from $\bar{c}_1$ segments.

Consider a generator $x \otimes y$ in $\boldsymbol{\ell}_1^A \boxtimes \boldsymbol{\ell}_2$. Then x belongs to a segment s_1 in $\boldsymbol{\ell}_1^A$ and y belongs to a segment s_2 in $\boldsymbol{\ell}_2$. We will consider cases depending on the type of the segments s_1 and s_2 and in each case show that either $x \otimes y$ has grading 0 or it cancels in homology. Therefore $\boldsymbol{\ell}_1^A \boxtimes \boldsymbol{\ell}_2$ is an L–space complex.

First note that $\mathrm{gr}(x)$ is always 0 by assumption. If s_2 is an e, d_k or b_k segment, then y also has grading 0 and the grading of $x \otimes y$ is 0. If s_1 is $\bar{e}^*$ then x is in idempotent ι_1 and so y must also have idempotent ι_1. All ι_1–generators of $\boldsymbol{\ell}_2$ have grading 0 except those in $\bar{c}_1$ segments, so $x \otimes y$ has grading 0 if s_1 is $\bar{e}^*$ and s_2 is not $\bar{c}_1$.

Suppose that s_1 is c_k^* and s_2 is a_ℓ with generators labelled by

$$s_1 = \quad \xleftarrow{\rho_1} \underset{x_1}{\bullet^{+}} \xrightarrow{\rho_3,\rho_2} \underset{x_2}{\bullet^{+}} \overset{\rho_3,\rho_2}{\dashrightarrow} \underset{x_k}{\bullet^{+}} \xrightarrow{\rho_3} \underset{x_{k+1}}{\circ^{+}}$$

$$s_2 = \quad \underset{y_0}{\bullet^{-}} \xrightarrow{\rho_3} \underset{y_1}{\circ^{+}} \xrightarrow{\rho_{23}} \underset{y_2}{\circ^{+}} \overset{\rho_{23}}{\dashrightarrow} \underset{y_\ell}{\circ^{+}} \xrightarrow{\rho_2} \underset{y_{\ell+1}}{\bullet^{+}}$$

Every generator in each segment has grading 0 except for y_0, so $x \otimes y$ has grading 0 unless $y = y_0$ and $x = x_i$ for $i \in \{1, \dots, k\}$. For each i, $\boldsymbol{\ell}_1^A$ has the operation

$$m_{2+k-i}(x_i, \rho_3, \rho_{23}, \dots, \rho_{23}) = x_{k+1}.$$

If $k - i < \ell$, it follows that in the tensor product there is a differential from $x_i \otimes y_0$ to $x_{k+1} \otimes y_{k-i+1}$. It is not difficult to check that $x_{k+1} \otimes y_{k-i+1}$ does not cancel from the right; the arrow in $\boldsymbol{\ell}_2$ to the right of y_{k-i+1} is outgoing and s_1 is followed by a c^* or $\bar{e}^*$ segment, so x_{k+1} has only incoming $\mathcal{A}_\infty$ operations. Therefore, in this case $x_i \otimes y_0$ cancels in homology. If instead $k - i \geq \ell$ then $\boldsymbol{\ell}_1^A$ has the operation

$$m_{2+\ell}(x_i, \rho_3, \rho_{23}, \dots, \rho_{23}, \rho_2) = x_{i+\ell}$$

and there is a differential in the box tensor product from $x_i \otimes y_0$ to $x_{i+\ell} \otimes y_{\ell+1}$. Again it is not difficult to see that $x_{i+\ell} \otimes y_{\ell+1}$ does not cancel from the right, since the arrow in $\boldsymbol{\ell}_2$ to the right of $y_{\ell+1}$ must be an outgoing ρ_1, ρ_{12} or ρ_{123} arrow and $x_{i+\ell}$ has no outgoing $\mathcal{A}_\infty$ operations with first input ρ_1, ρ_{12} or ρ_{123}. Thus $x_i \otimes y_0$ cancels in homology.

Suppose that s_1 is c_k^*, with generators labelled as above, and s_2 is $\bar{c}_1$ with generators labelled as follows (note that the generator y_0 is not actually part of the segment s_2):

$$s_2 = \quad \underset{y_0}{\bullet^{+}} \xrightarrow{\quad\rho_1\quad} \underset{y_1}{\circ^{-}} \xleftarrow{\quad\rho_3\quad} \underset{y_2}{\bullet^{+}}$$

The generator y_2 has grading 0 and the generator y_1 has grading 1, so $x \otimes y$ only has grading 1 if $x = x_{k+1}$ and $y = y_1$. In this setting, $x_{k+1} \otimes y_1$ has an incoming differential on the right which starts from $x_k \otimes y_2$. To ensure that $x_{k+1} \otimes y_1$ cancels in homology we need to check that $x_k \otimes y_2$ does not cancel on the right. To the right of y_2 in $\boldsymbol{\ell}_2$ is an outgoing sequence of arrows that starts with some number of ρ_{12} arrows followed by a ρ_{123} arrow (here we use that a $\bar{c}_1$ segment in $\boldsymbol{\ell}_2$ is not followed by another $\bar{c}_1$ segment with only e's in between). If $k > 1$ then x_k has no outgoing $\mathcal{A}_\infty$ operations except $m_2(x_k, \rho_3) = x_{k+1}$. If $k = 1$, then x_k has additional operations, but, since s_1 is preceded by some number of $\bar{e}^*$ segments and then a c^* segment, the inputs for these operations can only be some number of ρ_{12}'s followed by a ρ_1. In either case, it is clear that $x_k \otimes y_2$ does not cancel on the right.

Suppose that s_1 is $\bar{e}^*$ and s_2 is $\bar{c}_1$. In this case, x is the only generator of $\bar{e}^*$, y is the only generator of $\bar{c}_1$ with idempotent ι_1, and $x \otimes y$ has grading 1. In $\boldsymbol{\ell}_2$, s_2 is preceded by some number i of e segments, which are preceded by a b_k or d_k segment. Thus to the left of y in $\boldsymbol{\ell}_2$ there is an incoming sequence of arrows that ends with ρ_2, i ρ_{12}'s and ρ_1. In $\boldsymbol{\ell}_1^A$, s_1 is followed by some number j of $\bar{e}^*$ segments followed by a c_k^*. If $j > i$ then x has an incoming operation with inputs $(\rho_2, \rho_{12}, \dots, \rho_{12}, \rho_1)$ with i ρ_{12}'s, and if $j \le i$ then x has an incoming operation with inputs $(\rho_{12}, \dots, \rho_{12}, \rho_1)$ with j ρ_{12}'s. In either case, these operations give rise to a differential in the box tensor product ending at $x \otimes y$. In both cases it is also easy to check that the initial generator of this differential has no other differentials, and so $x \otimes y$ cancels in homology.

The only case remaining is the case that s_1 is c_k^* and s_2 is c_ℓ. This case is depicted in Figure 7 for $k = 3$ and $\ell = 2$. We label the generators of s_1 sequentially as $x_1, \dots, x_{k+1}$, and we label the generators of s_2 as $y_0, \dots, y_\ell$. For the remainder of this proof, assume that s_1 is c_k^*, s_2 is c_ℓ, $x \otimes y$ has grading 1 and $x \otimes y$ does not cancel in homology; we will produce a contradiction, proving the proposition.

Since $x \otimes y$ has grading 1, y is y_0 and x is x_i with $i \in \{1, \dots, k\}$. For each $i > k - \ell$, $x_i \otimes y_0$ has an outgoing differential which ends at $x_{k+1} \otimes y_{k-i+1}$. If $i > k - \ell + 1$, then

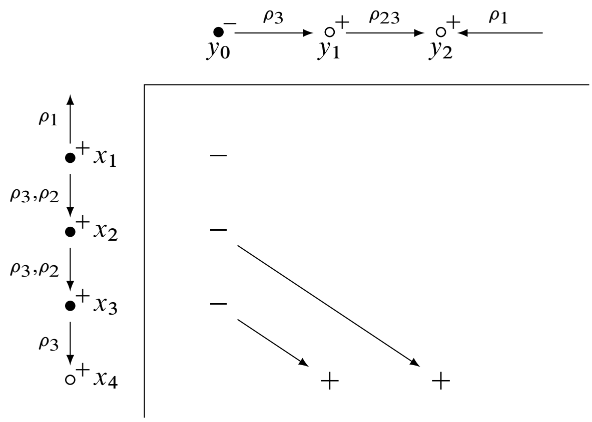

Figure 7: A portion of the chain complex $\boldsymbol{\ell}_1^A \boxtimes \boldsymbol{\ell}_2$ coming from a segment c_3^* in $\boldsymbol{\ell}_1^A$ (left edge) and a segment c_2 in $\boldsymbol{\ell}_2$ (top edge). Signs $+$ and $-$ indicate generators in the box tensor product with grading 0 and 1, respectively.

$x_{k+1} \otimes y_{k-i+1}$ does not cancel on the right, since the arrow to the right of y_{k-i+1} is an outgoing ρ_{23} arrow but x_{k+1} has only incoming $\mathcal{A}_\infty$ operations. Thus, if $i > k-\ell+1$, $x_i \otimes y_0$ cancels on the right. By Lemmas 5.2 and 5.3, $k \le n$ and $\ell \ge n-1$; it follows that $k-\ell$ is at most 1 and $x_i \otimes y_0$ cancels on the right for any $i > 2$. Thus x must be either x_1 or x_2.

If $k \le \ell$ then $x_2 \otimes y_0$ cancels on the right. If $k = n$ and $\ell = n-1$, $x_2 \otimes y_0$ potentially cancels from the right. It has an outgoing differential on the right ending at $x_{k+1} \otimes y_\ell$, so it cancels from the right if and only if $x_{k+1} \otimes y_\ell$ does not cancel on the right. Now $x_{k+1} \otimes y_\ell$ has a differential on the right only if $s_1 = c_k^*$ is followed by a segment $c_{k'}^*$, in which case the differential starts with the generator $x_1' \otimes y_0'$, where x_1' is the first generator of the $c_{k'}^*$ segment following s_1 and y_0' is the first generator of the segment following s_2. If $x_1' \otimes y_0'$ cancels on the right then so does $x_2 \otimes y_0$, and $x_1' \otimes y_0'$ automatically cancels on the right if s_2 is followed by a type a segment. Thus, if $x = x_2$ we must have that s_1 is followed by $c_{k'}^*$, s_2 is followed by c_ℓ', and $x_1' \otimes y_0'$ does not cancel from the right.

If $k < \ell$ then $x_1 \otimes y_0$ cancels on the right, and if $k > \ell$ then $x_1 \otimes y_0$ does not cancel on the right. If $k = \ell = n$ or $k = \ell = n-1$, then $x_1 \otimes y_0$ potentially cancels on the right. By the same reasoning as above, it does not cancel on the right if and only if s_1 is followed by $c_{k'}^*$, s_2 is followed by c_ℓ', and $x_1' \otimes y_0'$ does not cancel from the right. Now $x_1 \otimes y_0$ may also cancel from the left. It has an outgoing differential on the left ending at $x_0 \otimes y_{\ell'}'$, where x_0 is the last generator in the segment immediately preceding s_1

and $y'_{\ell'}$ is the last generator in the segment preceding s_2. If s_2 is preceded by $b_{\ell'}$ then it is easy to see that $x_0 \otimes y'_{\ell'}$ has no differentials on the left, so $x_1 \otimes y_0$ cancels from the left. If instead s_2 is preceded by $c_{\ell'}$, $x_0 \otimes y'_{\ell'}$ cancels from the left unless s_1 is preceded by $c^*_{k'}$ and $k' < \ell'$ or $k' = \ell'$ and $x'_1 \otimes y'_0$ cancels from the left.

Suppose that $x \otimes y = x_2 \otimes y_0$ does not cancel in homology (in particular it does not cancel from the right). Then we have shown that $k = n$, $\ell = n-1$, $s_1 = c^*_k$ is followed by $c^*_{k'}$, $s_2 = c_\ell$ is followed by $c_{\ell'}$, and the generator $x'_1 \otimes y'_0$ does not cancel from the right. Furthermore, this last fact implies that either $k' = n$ and $\ell' = n-1$ or that $k' = \ell'$, $c^*_{k'}$ is followed by $c^*_{k''}$, $c_{\ell'}$ is followed by $c_{\ell''}$, and $x''_1 \otimes y''_0$ does not cancel from the right, where x''_1 and y''_0 are the appropriate generators of $c^*_{k''}$ and $c_{\ell'}$. Repeating this argument, we see that $s_1 = c^*_n$ is followed by a sequence of c^*_n and c^*_{n-1} segments ending with a c^*_n and $s_2 = c_{n-1}$ is followed by a sequence of c_n and c_{n-1} segments ending with c_{n-1} but with indices otherwise the same as the indices in the sequence of c^* segments following s_1.

Suppose that $x \otimes y = x_1 \otimes y_0$ does not cancel in homology. The fact that it does not cancel from the right implies that $k = \ell = n$ or $k = \ell = n-1$, $s_1 = c^*_k$ is followed by $c^*_{k'}$, $s_2 = c_\ell$ is followed by $c_{\ell'}$, and the generator $x'_1 \otimes y'_0$ does not cancel from the right. As in the preceding paragraph, this implies that $s_1 = c^*_n$ is followed by a sequence of c^*_n and c^*_{n-1} segments ending with a c^*_n and $s_2 = c_{n-1}$ is followed by a corresponding sequence of c_n and c_{n-1} segments ending with c_{n-1}. The fact that $x \otimes y = x_1 \otimes y_0$ does not cancel from the left implies that $s_1 = c^*_k$ is preceded by $c^*_{k'}$, $s_2 = c_\ell$ is preceded by $c_{\ell'}$, and either $k' = n$ and $\ell' = n-1$ or $k' = \ell'$ and the generator $x'_1 \otimes y'_0$ does not cancel from the left. Repeating the argument, we see that s_1 is preceded by a sequence of c^*_n and c^*_{n-1} segments starting with c^*n and s_2 is preceded by a sequence of c_n and c_{n-1} segments with matching sequence of indices except that the initial segment is c_{n-1}.

Regardless of whether $x = x_1$ or $x = x_2$, we find that there is a sequence of indices $k_1, k_2, \dots, k_r$ with $k_i \in \{n, n-1\}$ such that $\boldsymbol{\ell}_1^A$ contains $c^*_n c^*_{k_1} c^*_{k_2} \dots c^*_{k_r} c^*_n$ and $\boldsymbol{\ell}_2$ contains $c_{n-1} c_{k_1} c_{k_2} \dots c_{k_r} c_{n-1}$. It follows that $\boldsymbol{\ell}_1$ and $\boldsymbol{\ell}_2$ satisfy property λ. By Lemma 5.4, this implies that $\boldsymbol{\ell}_1$ and $\boldsymbol{\ell}_2$ are not L–space aligned, a contradiction. □

As noted previously, the gluing statement in Proposition 5.1 needs to be modified if either loop is solid torus-like. Note that for a solid torus-like loop, all slopes are L–space slopes except the rational longitude. If $\boldsymbol{\ell}_1$ is solid torus-like with rational longitude r/s and $\boldsymbol{\ell}_2$ is simple, then $\boldsymbol{\ell}_1$ and $\boldsymbol{\ell}_2$ are L–space aligned if and only if s/r

is a strict L–space slope for $\boldsymbol{\ell}_2$. This is a sufficient, but not a necessary, condition for $\boldsymbol{\ell}_1^A \boxtimes \boldsymbol{\ell}_2$ to be an L–space complex.

Proposition 5.6 *If $\boldsymbol{\ell}_1$ and $\boldsymbol{\ell}_2$ are simple loops and $\boldsymbol{\ell}_1$ is solid torus-like with rational longitude r/s, then $\boldsymbol{\ell}_1^A \boxtimes \boldsymbol{\ell}_2$ is an L–space chain complex if and only if s/r is an L–space slope for $\boldsymbol{\ell}_2$.*

Proof We may choose a framing so that $\boldsymbol{\ell}_1$ is a collection of e segments, that is, so that the rational longitude is represented by 0. Correspondingly, the slope s/r for $\boldsymbol{\ell}_2$ is represented by ∞. The result now follows from Corollary 3.22. □

5.2 A gluing result for loop-type manifolds

Returning to loop-type manifolds, we are now in a position to collect the material proved in this section and, in particular, apply Proposition 5.1 to establish a gluing theorem.

Theorem 5.7 *Let (M_1, α_1, β_1) and (M_2, α_2, β_2) be simple loop-type bordered manifolds with torus boundary which are not solid torus-like, and let Y be the closed manifold $(M_1, \alpha_1, \beta_1) \cup (M_2, \alpha_2, \beta_2)$ (with the gluing map $\alpha_1 \mapsto \beta_2$, $\beta_1 \mapsto \alpha_2$, as in Section 2.5). Then Y is an L–space if and only if every essential simple closed curve on $\partial M_1 = \partial M_2 \subset Y$ determines a strict L–space slope for either M_1 or M_2.*

Remark 5.8 There is an alternative statement of the conclusion on Theorem 5.7 using the notation laid out in this paper: the closed manifold $M_1 \cup M_2$ is an L–space if and only if, for each rational p/q, either $p/q \in \mathcal{L}^\circ(M_1, \alpha_1, \beta_1)$ or $q/p \in \mathcal{L}^\circ(M_2, \alpha_2, \beta_2)$. Recall that, following the conventions in Section 2.5, $p/q \in \mathcal{L}^\circ(M, \alpha, \beta)$ if and only if $\pm(p\alpha + q\beta) \in \mathcal{L}_M^\circ$.

Note that Theorem 5.7 implies Theorem 1.3.

Remark 5.9 If either bordered manifold in the statement of Theorem 5.7 is solid torus-like, then Y is an L–space if and only if every essential simple closed curve on $\partial M_1 = \partial M_2 \subset Y$ determines an L–space slope for either M_1 or M_2. This is an immediate consequence of Proposition 5.6, and amounts to replacing $\mathcal{L}^\circ(M_i, \alpha, \beta)$ by $\mathcal{L}(M_i, \alpha, \beta)$. Notice that Dehn surgery — by definition of an L–space slope — is a special case of this version of the gluing result, since the solid torus is solid torus-like, in the sense of Definition 3.23.

Proof of Theorem 5.7 Let $\widehat{CFD}(M_1,\alpha_1,\beta_1)$ be represented by simple loops $\ell_1^1,\dots,\ell_1^n$ and let $\widehat{CFD}(M_2,\alpha_2,\beta_2)$ be represented by simple loops $\ell_2^1,\dots,\ell_2^m$. A given slope is a (strict) L–space slope for (M_i,α_i,β_i) if and only if it is a (strict) L–space slope (abstractly) for each loop ℓ_i^k. Y is an L–space if and only if $\ell_1^k \boxtimes \ell_2^j$ is and L–space complex (again, abstractly) for each $1 \le k \le n$ and $1 \le j \le m$.

Suppose there is a slope p/q such that $p/q \notin \mathcal{L}^\circ(M_1,\alpha_1,\beta_1)$ and $q/p \notin \mathcal{L}^\circ(M_2,\alpha_2,\beta_2)$. Then p/q is not a strict L–space slope for ℓ_1^k for some k and q/p is not a strict L–space slope for ℓ_2^j for some j. We may assume that ℓ_1^k is not solid torus-like; if it is solid torus-like, then p/q must be the rational longitude for (M_1,α_1,β_1) and thus not a strict L–space slope for any of the loops $\ell_1^1,\dots,\ell_1^n$. Similarly, we may assume that ℓ_2^j is solid torus-like. Since ℓ_1^k and ℓ_2^j are not L–space aligned, $\ell_1^k \boxtimes \ell_2^j$ is not an L–space complex, and Y is not an L–space.

If every slope is a strict L–space slope for either M_1 or M_2, then every slope is a strict L–space slope either for each ℓ_1^k or for each ℓ_2^j. It follows that, for every pair (k,j), ℓ_1^k and ℓ_2^j are L–space aligned, and thus Y is an L–space. □

5.3 L–space knot complements

In light of Theorem 5.7, it is natural to ask which 3–manifolds M with torus boundary have simple loop-type bordered invariants. We first observe that complements of L–space knots in S^3, that is, those knots admitting an L–space surgery, have this property.

Proposition 5.10 *If K is an L–space knot in S^3 and $M = S^3 \setminus \nu(K)$, then for any choice of parametrizing curves α and β, $\widehat{CFD}(M,\alpha,\beta)$ can be represented by a single simple loop.*

Proof We need only show that, for some choice of parametrizing curves, $\widehat{CFD}(M,\alpha,\beta)$ is a loop which consists of only unstable chains in either standard notation or dual notation. Consider then $\widehat{CFD}(M,\mu,n\mu+\lambda)$, where μ is the meridian of the knot K and λ is the Seifert longitude of K. Suppose that $|n|$ is chosen sufficiently large that the result of Dehn surgery $S^3_n(K) = M(n\mu+\lambda)$ is an L–space.

The knot Floer homology $CFK^-(K)$ has a staircase shape of alternating horizontal and vertical arrows (see [25], for example). By the construction in [20, Section 11.5], $\widehat{CFD}(M,\mu,n\mu+\lambda)$ consists of alternating horizontal and vertical chains, with the ends connected by a single unstable chain. Thus $\widehat{CFD}(M,\mu,n\mu+\lambda)$ is a loop. In standard loop notation, the horizontal chains are type a segments, the vertical chains are type b segments, and the unstable chain is a type d segment if $n>0$ and a type c

segment if $n < 0$. More precisely,

$$\widehat{CFD}(M,\mu,n\mu+\lambda) = \begin{cases} b_{k_1}a_{k_2}b_{k_3}a_{k_4}\cdots b_{k_{2r-1}}a_{k_{2r}}d_j & \text{if } n > 0, \\ a_{k_1}b_{k_2}a_{k_3}b_{k_4}\cdots a_{k_{2r-1}}b_{k_{2r}}c_j & \text{if } n < 0. \end{cases}$$

In either case, the loop has no bar segments in standard notation, so when we switch to dual notation it has no stable chains. Thus, $\widehat{CFD}(M,\mu,n\mu+\lambda)$ is a simple loop. □

Note that the behaviour established in this application of Theorem 5.7 is expected in general and, in particular, should not require the hypothesis that the knot complement be a loop-type manifold; see [9, Conjecture 1; 5, Conjecture 4.3].

Conjecture 5.11 *For knots K_1 and K_2 in S^3, with $M_i = S^3 \setminus \nu(K)$ for $i = 1,2$, the generalized splicing $M_1 \cup_h M_2$ is an L–space if and only if either $\gamma \in \mathcal{L}^\circ_{M_1}$ or $h(\gamma) \in \mathcal{L}^\circ_{M_2}$ for every slope γ.*

On the other hand, it is clear that the case where two L–space knot complements are identified in the generalized splice is the interesting case, and therefore the loop restriction in this setting is quite natural. In this case, we remark that the question is now settled.

Corollary 5.12 *Conjecture 5.11 holds when the K_i are L–space knots.*

Proof This is immediate on combining Theorem 5.7 and Proposition 5.10. □

Note that Corollary 5.12 implies Theorem 1.6.

Some authors define L–space knots to be the class of knots for which some *positive* surgery yields an L–space. We will not follow this convention because, while this is a natural definition for knots in S^3 in that certain statements become simpler, our interest is in the more general setting of manifolds with torus boundary admitting L–space fillings, wherein the distinction seems to be less meaningful. In particular, the next section is concerned with this more general setting.

6 Graph manifolds

Note that knots in the three-sphere admitting L–space surgeries contain torus knots — those knots admitting a Seifert structure on their complement — as a strict subset. Our goal now is to establish another class of loop-type manifolds: Seifert-fibred rational homology solid tori. This should be viewed as the natural geometric enlargement of the class of torus knot exteriors. To do so, we study these as a subset of graph manifolds, and establish sufficient conditions for a rational homology solid torus admitting a graph manifold structure to be of loop-type.

6.1 Preliminaries on graph manifolds

We will represent a graph manifold rational homology sphere Y by a plumbing tree, following the notation developed in [22]. A plumbing tree is an acyclic graph with integer weights associated to each vertex. Such a graph specifies a graph manifold as follows: To each vertex v_i of weight e_i and valence d_i we assign the Euler number e_i circle bundle over the sphere minus d_i disks, and for each edge connecting vertices v_i and v_j glue the corresponding bundles along a torus boundary component by a gluing map that takes the fibre of one bundle to a curve in the base surface of the other bundle, and vice versa.

To allow for graph manifolds with boundary, we associate an additional integer $b_i \geq 0$ to each vertex. To construct the corresponding manifold, we associate to each vertex a bundle over S^2 minus $(d_i + b_i)$; d_i boundary components of this bundle will glue to bundles corresponding to other vertices, but the remaining b_i boundary components remain unglued. The resulting graph manifold has $\sum_i b_i$ toroidal boundary components. In diagrams of plumbing trees, we will indicate the presence of boundary tori by drawing b_i half-edges (dotted lines which do not connect to another vertex) at each vertex v_i.

Given a plumbing tree Γ with a single boundary half-edge, let M_Γ denote the corresponding graph manifold. The torus ∂M_Γ has a natural choice of parametrizing curves: one corresponds to a fibre in the S^1 bundle containing the boundary, and one corresponds to a curve in the base surface of that bundle (note that these are precisely the curves used above to specify the gluing maps in the construction based on a given graph). Call these two slopes α and β, respectively. These slopes do not have a preferred orientation, but reversing the orientation on the bundle associated to every vertex in the construction of M_Γ gives a diffeomorphism between $(M_\Gamma, \alpha, \beta)$ and $(M_\Gamma, -\alpha, -\beta)$ as bordered manifolds.

Thus $(M_\Gamma, \alpha, \beta)$ is a canonical bordered manifold associated to Γ. Since we will be interested in the bordered invariants of such graph manifolds, for ease of notation we will refer to $\widehat{CFD}(M_\Gamma, \alpha, \beta)$ simply as $\widehat{CFD}(\Gamma)$.

We will make use of three important operations on single boundary plumbing trees. These are summarized in Figure 8, and described as follows:

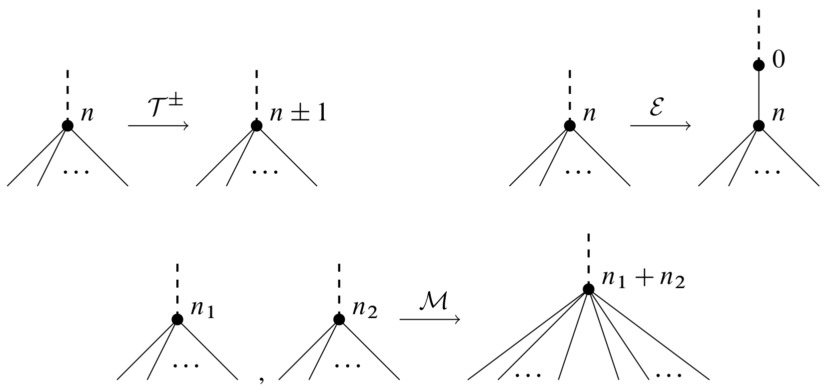

Figure 8: Three operations on graphs for constructing and graph manifold with torus boundary.

Twist The operation $\mathcal{T}^{\pm}$ adds ± 1 to the weight of the vertex containing the boundary edge.

Extend The operation $\mathcal{E}$ inserts a new 0–framed vertex between the boundary vertex and the boundary edge.

Merge The operation $\mathcal{M}$ takes two single boundary plumbing trees and identifies their boundary vertices, removing one of the two boundary edges to produce a new single boundary plumbing tree; the weight of the new boundary vertex is the sum of the weights of the original boundary vertices.

With these operations we can, in particular, construct any single-boundary plumbing tree.

6.2 Operations on bordered manifolds

The graph operations described above may be thought of as operations on the corresponding (bordered) graph manifolds. In particular, we will abuse notation and write $M_{\mathcal{T}(\Gamma)} = \mathcal{T}(M_\Gamma)$. $\mathcal{T}(M_\Gamma)$ is obtained from M_Γ by attaching the mapping cylinder of a positive (standard) Dehn twist to the boundary of M_Γ. More precisely, by considering the gluing conventions prescribed by the plumbing tree, the effect on the bordered manifold is $\mathcal{T}(M_\Gamma, \alpha, \beta) = (M_\Gamma, \alpha, \beta + \alpha)$.

Similarly, $\mathcal{E}(M_\Gamma) = M_{\mathcal{E}(\Gamma)}$ is obtained from M_Γ by attaching the bimodule corresponding to the two-boundary plumbing tree

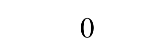

The graph manifold assigned to this plumbing tree is the trivial S^1–bundle over the annulus, or $T^2 \times [0,1]$. Recall that the bordered structure on each boundary component of this manifold is given by the convention that α is an S^1 fibre and β lies in the base surface. One checks that the resulting bordered manifold can be realised as the mapping cylinder for the diffeomorphism represented by the matrix $\left(\begin{smallmatrix} 0 & 1 \\ -1 & 0 \end{smallmatrix}\right) \in \mathrm{SL}_2(\mathbb{Z})$. The effect on the level of bordered manifolds is $\mathcal{E}(M_\Gamma, \alpha, \beta) = (M_\Gamma, -\beta, \alpha)$.

Finally, $\mathcal{M}(M_{\Gamma_1}, M_{\Gamma_2})$ is the result of attaching the bundle $S^1 \times \mathcal{P}$, where $\mathcal{P}$ is a pair of pants. The bordered structure on $S^1 \times \mathcal{P}$ (which we suppress from the notation) is determined as follows: the two "input" boundary components which glue to M_{Γ_1} and M_{Γ_2} are parametrized by pairs (α, β) with β a fibre and α a curve in the base surface $\mathcal{P}$, while the third "output" boundary component is parametrized by (α, β) with α a fibre and β a curve in $\mathcal{P}$. As usual, the bordered structure on $S^1 \times \mathcal{P}$ determines the gluing as well as the bordered structure on the resulting three-manifold.

To see that $\mathcal{M}(M_{\Gamma_1}, M_{\Gamma_2})$ is indeed the manifold corresponding to $\mathcal{M}(\Gamma_1, \Gamma_2)$, note that $S^1 \times \mathcal{P}$ with the specified bordered structure is the manifold associated with the three-boundary plumbing tree $\Gamma_{\mathcal{M}}$,

$\Gamma_{\mathcal{M}} =$ 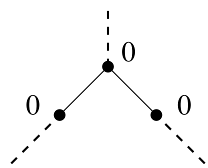

On the other hand, attaching Γ_1 and Γ_2 to the two lower boundary edges of $\Gamma_{\mathcal{M}}$ produces a single-boundary plumbing tree that is equivalent to $\mathcal{M}(\Gamma_1, \Gamma_2)$ (by equivalent, we mean that the corresponding manifolds are diffeomorphic; to see this equivalence, use rule R3 of [22] to contract the 0–framed valence two vertices).

From this discussion, it should be clear that the operations $\mathcal{T}$, $\mathcal{E}$ and $\mathcal{M}$ may be extended to natural operations on arbitrary bordered manifolds. We are interested in the effect of these operations on bordered invariants. Following the conventions laid out, we have that $\mathcal{T}^{\pm}(M, \alpha, \beta) = (M, \alpha, \beta \pm \alpha)$, so that, at the level of bordered invariants,

$$\widehat{CFD}(\mathcal{T}^{\pm}(M, \alpha, \beta)) \cong \widehat{T}_{\mathrm{st}}^{\pm 1} \boxtimes \widehat{CFD}(M, \alpha, \beta). \tag{1}$$

Similarly, the action of $\left(\begin{smallmatrix} 0 & 1 \\ -1 & 0 \end{smallmatrix}\right) \in \mathrm{SL}_2(\mathbb{Z})$, giving $\mathcal{E}(M, \alpha, \beta) = (M, -\beta, \alpha)$, is realized by

$$\widehat{CFD}(\mathcal{E}(M, \alpha, \beta)) \cong \widehat{T}_{\mathrm{st}} \boxtimes \widehat{T}_{\mathrm{du}} \boxtimes \widehat{T}_{\mathrm{st}} \boxtimes \widehat{CFD}(M, \alpha, \beta) \tag{2}$$

on type D structures.

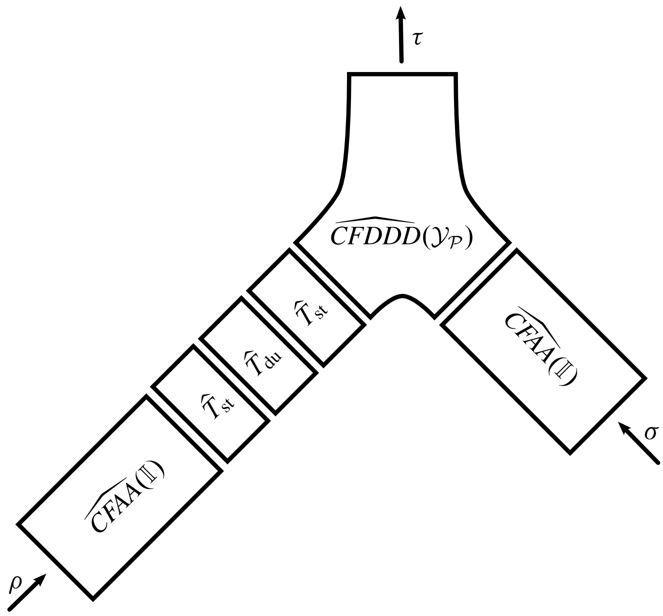

Figure 9: A schematic description of the trimodule $\widehat{CFDAA}(S^1 \times \mathcal{P})$ extracted from the calculation and conventions in [8].

The most involved is the merge operation: for a pair of bordered manifolds (M_i, α_i, β_i) for $i = 1, 2$, this produces a bordered manifold

$$\mathcal{M}_{1,2} = \mathcal{M}((M_1, \alpha_1 \beta_1), (M_2, \alpha_2, \beta_2))$$

using $S^1 \times \mathcal{P}$. As a bordered 3–manifold with three boundary components, $S^1 \times \mathcal{P}$ gives rise to a trimodule in the bordered theory (the object of study in work of the first author [8]). More precisely, we consider the type DAA trimodule $\widehat{CFDAA}(S^1 \times \mathcal{P})$.

We will extract $\widehat{CFDAA}(S^1 \times \mathcal{P})$ from the trimodule $\widehat{CFDDD}(\mathcal{Y}_{\mathcal{P}})$ computed in [8]. Note that $S^1 \times \mathcal{P}$ and $\mathcal{Y}_{\mathcal{P}}$ agree as manifolds, but have different bordered structure. The trimodule $\widehat{CFDDD}(\mathcal{Y}_{\mathcal{P}})$ has five generators, v, w, x, y and z, and the differential is given by (cf [8, Figure 10])

$$\begin{aligned}
\partial(v) &= \rho_3 \otimes x + \rho_1\sigma_3\tau_{123} \otimes y + \rho_{123}\sigma_{123}\tau_{123} \otimes y + \tau_3 \otimes z, \\
\partial(w) &= \rho_3\sigma_{12} \otimes x + \rho_1\sigma_3\tau_1 \otimes y + \rho_{123}\sigma_{123}\tau_1 \otimes y, \\
\partial(x) &= \rho_2\sigma_{12} \otimes v + \rho_2 \otimes w + \sigma_1\tau_3 \otimes y, \\
\partial(y) &= \sigma_2\tau_2 \otimes x + \rho_2\sigma_2 \otimes z, \\
\partial(z) &= \rho_3\sigma_1 \otimes y + \tau_2 \otimes w.
\end{aligned}$$

$m_2(x,\sigma_3)=z$	$m_3(v,\rho_1,\sigma_1)=\tau_1\otimes y$
$m_2(y,\sigma_2)=x$	$m_3(v,\rho_1,\sigma_{12})=\tau_1\otimes x$
$m_2(y,\sigma_{23})=z$	$m_3(v,\rho_1,\sigma_{123})=\tau_1\otimes z$
$m_2(w,\rho_3)=z$	$m_3(v,\rho_1,\sigma_{123})=\tau_{123}\otimes y$
$m_2(y,\rho_2)=w$	$m_3(v,\rho_{12},\sigma_1)=\tau_1\otimes w$
$m_2(y,\rho_{23})=z$	$m_3(v,\rho_{12},\sigma_{12})=\tau_{12}\otimes v$
$m_2(w,\sigma_{23})=\tau_{23}\otimes w$	$m_3(v,\rho_{12},\sigma_{123})=\tau_{123}\otimes w$
$m_2(x,\sigma_3)=\tau_{23}\otimes y$	$m_3(v,\rho_{123},\sigma_1)=\tau_1\otimes z$
$m_2(x,\rho_2)=\tau_2\otimes v$	$m_3(v,\rho_{123},\sigma_{12})=\tau_{123}\otimes x$
$m_2(x,\rho_{23})=\tau_{23}\otimes x$	$m_3(v,\rho_{123},\sigma_{123})=\tau_{123}\otimes z$
$m_2(v,\sigma_3)=\tau_3\otimes w$	$m_5(v,\rho_3,\rho_2,\rho_1,\sigma_1)=\tau_1\otimes z$
$m_2(y,\sigma_{23})=\tau_{23}\otimes y$	$m_5(v,\rho_3,\rho_2,\rho_1,\sigma_1)=\tau_{123}\otimes y$
$m_2(v,\rho_3)=\tau_3\otimes x$	$m_5(v,\rho_3,\rho_2,\rho_{12},\sigma_1)=\tau_{123}\otimes w$
$m_2(w,\sigma_2)=\tau_2\otimes v$	$m_5(v,\rho_3,\rho_2,\rho_{123},\sigma_1)=\tau_{123}\otimes z$
$m_2(z,\sigma_2)=\tau_{23}\otimes x$	
$m_2(z,\sigma_{23})=\tau_{23}\otimes z$	$m_7(v,\rho_3,\rho_2,\rho_3,\rho_2,\rho_1,\sigma_1)=\tau_{123}\otimes z$

Table 1: Operations for $\widehat{CFDAA}(\Gamma_{\mathcal{M}})$.

Since the trimodule has three commuting actions by three copies of the torus algebra (one corresponding to each boundary component), we use ρ, σ and τ to distinguish between algebra elements in each copy of $\mathcal{A}$. The σ–boundary of $\mathcal{Y}_{\mathcal{P}}$ is parametrized by a pair (α,β) such that β is a fibre of $\mathcal{P}\times S^1$ and α lies in the base $\mathcal{P}$, while the ρ– and τ–boundaries have the opposite parametrization. Thus the bordered manifold $S^1\times\mathcal{P}$ can be obtained from $\mathcal{Y}_{\mathcal{P}}$ by switching the role of α and β on (say) the ρ–boundary.

It is now clear how to compute $\widehat{CFDAA}(S^1\times\mathcal{P})$ from $\widehat{CFDDD}(\mathcal{Y}_{\mathcal{P}})$: we change the parametrization of the ρ–boundary by applying the bimodule $\widehat{T}_{\mathrm{st}}\boxtimes\widehat{T}_{\mathrm{du}}\boxtimes\widehat{T}_{\mathrm{st}}$ and then change the ρ– and σ–boundaries to type A by tensoring with $\widehat{CFAA}(\mathbb{I})$ (see Figure 9 for a schematic description). These are both straightforward computations; the resulting trimodule has five generators and the operations are listed in Table 1.

With this trimodule in hand, we have

$$\widehat{CFD}(\mathcal{M}_{1,2})\cong\widehat{CFDAA}(S^1\times\mathcal{P})\boxtimes((M_1,\alpha_1,\beta_1),(M_2,\alpha_2,\beta_2)),$$

where $\widehat{CFDAA}(S^1 \times \mathcal{P}) \boxtimes (\cdot, \cdot)$, by convention, tensors against the ρ–boundary in the first factor and against the σ–boundary in the second factor; compare Figure 9.

6.3 The effect of twist, extend and merge on loops

Restricting to the case of loop-type bordered invariants, the effect of the operations $\mathcal{T}^{\pm 1}$, $\mathcal{E}$ and $\mathcal{M}$ can be given simpler descriptions. Recall that $\mathrm{E} = \mathrm{T} \circ \mathrm{H}^{-1} \circ \mathrm{T}$, and that E can be easily calculated using Lemma 3.10.

Proposition 6.1 *If $\widehat{CFD}(M, \alpha, \beta)$ is of loop-type, and represented by the collection $\{\boldsymbol{\ell}_i\}_{i=1}^n$, then $\widehat{CFD}(\mathcal{T}^{\pm}(M, \alpha, \beta))$ is represented by $\{\mathrm{T}^{\pm 1}(\boldsymbol{\ell}_i)\}_{i=1}^n$ and $\widehat{CFD}(\mathcal{E}(M, \alpha, \beta))$ is represented by $\{\mathrm{E}(\boldsymbol{\ell}_i)\}_{i=1}^n$.*

Proof The proof is immediate from Proposition 3.11 with (1) and Proposition 3.12 with (2). $\square$

It follows easily that the operations $\mathcal{T}^{\pm}$ and $\mathcal{E}$ preserve the simple loop-type property.

Lemma 6.2 *Given a (simple) loop-type bordered three-manifold, the operations $\mathcal{T}^{\pm}$ and $\mathcal{E}$ produce (simple) loop-type manifolds.*

Proof It is an immediate consequence of Proposition 6.1 that if $\widehat{CFD}(M, \alpha, \beta)$ is represented by a collection of loops then the same is true for $\widehat{CFD}(\mathcal{T}(M, \alpha, \beta))$ and $\widehat{CFD}(\mathcal{E}(M, \alpha, \beta))$, since the operations $\mathrm{T}^{\pm 1}$ and E take loops to loops. Moreover, if the loops defining $\widehat{CFD}(M, \alpha, \beta)$ are simple then the loops resulting from $\mathcal{T}^{\pm}$ and $\mathcal{E}$ are simple, since changing framing by Dehn twists does not, by definition, change whether or not a loop is simple. It only remains to check that $\widehat{CFD}(\mathcal{T}(M, \alpha, \beta))$ and $\widehat{CFD}(\mathcal{E}(M, \alpha, \beta))$ have exactly one loop for each spin^c–structure on the corresponding manifolds. This is again clear since it is true for $\widehat{CFD}(M, \alpha, \beta)$ and the operations $\mathcal{T}^{\pm 1}$ and $\mathcal{E}$ amount to changing the parametrization on the boundary of M; changing the parametrization does not change the number of spin^c–structures, and the loop operations $\mathrm{T}^{\pm 1}$ and E do not change the number of loops. $\square$

For the merge operation, we will restrict further to the case that the loop(s) representing $\widehat{CFD}(M_1, \alpha_1, \beta_1)$ can all be written in standard notation with no stable chains. In this case, $\widehat{CFD}(\mathcal{M}_{1,2})$ is also a collection of loops.

Remark 6.3 By contrast, if $\widehat{CFD}(M_1, \alpha_1, \beta_1)$ and $\widehat{CFD}(M_2, \alpha_2, \beta_2)$ both contain stable chains in standard notation, $\widehat{CFD}(\mathcal{M}_{1,2})$ is not obviously of loop-type. However, in many cases it can be realized as a collection of loops after a homotopy equivalence.

It is enough to describe $\mathcal{M}$ on individual loops; we use $\textsc{m}$ to denote the corresponding operation on abstract loops. If $\widehat{CFD}(M_1,\alpha_1,\beta_1)$ is represented by a collection of loops $\{\boldsymbol{\ell}_i\}_{i=1}^n$ and $\widehat{CFD}(M_2,\alpha_2,\beta_2)$ is represented by a collection of loops $\{\boldsymbol{\ell}_j\}_{j=1}^m$, then $\widehat{CFD}(\mathcal{M}_{1,2})$ is given by $\bigcup_{1\le i\le n, 1\le j\le m}\textsc{m}(\boldsymbol{\ell}_i,\boldsymbol{\ell}_j)$. Determining $\textsc{m}(\boldsymbol{\ell}_i,\boldsymbol{\ell}_j)$ is a direct calculation using the trimodule and a key application of loop calculus.

Proposition 6.4 *Let $\boldsymbol{\ell}_1$ be a loop which can be written in standard notation with only type d_k unstable chains and let $\boldsymbol{\ell}_2$ be any loop. Then $\textsc{m}(\boldsymbol{\ell}_1,\boldsymbol{\ell}_2)$ is a collection of loops. If $\boldsymbol{\ell}_2$ cannot be written in standard notation, then $\textsc{m}(\boldsymbol{\ell}_1,\boldsymbol{\ell}_2)$ is one copy of $\boldsymbol{\ell}_2$ for each segment in $\boldsymbol{\ell}_1$. Otherwise, $\textsc{m}(\boldsymbol{\ell}_1,\boldsymbol{\ell}_2)$ is determined as follows: The ι_0–vertices of the $\mathcal{A}$–decorated graph correspond to pairs (u,v), where u is an ι_0–vertex of $\boldsymbol{\ell}_1$ and v is an ι_0–vertex of $\boldsymbol{\ell}_2$, and for each d_k segment from u_1 to u_2 in $\boldsymbol{\ell}_1$ we have:*

1. *For each a_l segment from v_1 to v_2 in $\boldsymbol{\ell}_2$, there is an a_l segment from (u_2,v_1) to (u_2,v_2).*
2. *For each b_l segment from v_1 to v_2 in $\boldsymbol{\ell}_2$, there is n b_l segment from (u_1,v_1) to (u_1,v_2).*
3. *For each c_l segment from v_1 to v_2 in $\boldsymbol{\ell}_2$, there is a c_{l-k} segment from (u_2,v_1) to (u_1,v_2).*
4. *For each d_l segment from v_1 to v_2 in $\boldsymbol{\ell}_2$, there is a d_{k+l} segment from (u_1,v_1) to (u_2,v_2).*

Proof The type D module represented by $\textsc{m}(\boldsymbol{\ell}_1,\boldsymbol{\ell}_2)$ is obtained by tensoring the type D modules corresponding to $\boldsymbol{\ell}_1$ and $\boldsymbol{\ell}_2$ with the ρ– and σ–boundaries, respectively, of the trimodule $\widehat{CFDAA}(\mathcal{Y}_{\mathcal{P}})$. We will denote this tensor by $\widehat{CFDAA}(\mathcal{Y}_{\mathcal{P}})\boxtimes(\boldsymbol{\ell}_1,\boldsymbol{\ell}_2)$. This trimodule has five generators; the idempotents associated with each generator on the ρ–, τ– and σ–boundaries are as follows:

generator	v	w	x	y	z
idempotents	$(\iota_0^\rho,\iota_0^\sigma,\iota_0^\tau)$	$(\iota_0^\rho,\iota_1^\sigma,\iota_1^\tau)$	$(\iota_1^\rho,\iota_0^\sigma,\iota_1^\tau)$	$(\iota_1^\rho,\iota_1^\sigma,\iota_1^\tau)$	$(\iota_1^\rho,\iota_1^\sigma,\iota_1^\tau)$

Note that the ι_0–generators in $\widehat{CFDAA}(S^1\times\mathcal{P})\boxtimes(\boldsymbol{\ell}_1,\boldsymbol{\ell}_2)$ arise precisely from the generator v tensored with ι_0–generators in $\boldsymbol{\ell}_1$ and $\boldsymbol{\ell}_2$.

First suppose that $\boldsymbol{\ell}_2$ is written in standard notation. Note also that, since $\boldsymbol{\ell}_1$ is assumed to have no a_k segments, we may ignore the m_5 and m_7 operations in Table 1. As a result, there are no operations in $\textsc{m}(\boldsymbol{\ell}_1,\boldsymbol{\ell}_2)$ that arise from more than one segment in either loop. This means that to compute $\mathrm{M}(\boldsymbol{\ell}_1, \boldsymbol{\ell}_2)$ we can feed $\boldsymbol{\ell}_1$ and $\boldsymbol{\ell}_2$ into the trimodule one segment at a time. For each combination of segment in $\boldsymbol{\ell}_1$ and segment in $\boldsymbol{\ell}_2$, the resulting portion of $\widehat{CFDAA}(S^1 \times \mathcal{P}) \boxtimes (\boldsymbol{\ell}_1, \boldsymbol{\ell}_2)$ is homotopy equivalent to the segment determined by (1)–(4) in the statement of the proposition. The proof is essentially contained in Figures 10, 11 and 12; we will describe one case in detail and leave the details of the other cases to the reader, with the figures as a guide.

Consider a segment d_k in $\boldsymbol{\ell}_1$ and a segment a_l in $\boldsymbol{\ell}_2$ with $k, l > 0$ (see Figure 10, top left, for the case of $k = l = 2$). Let the generators in these two segments be labelled as

$$d_k = \xrightarrow{\rho_{123}} \underset{x_1}{\circ} \xrightarrow{\rho_{23}} \cdots \xrightarrow{\rho_{23}} \underset{x_k}{\circ} \xrightarrow{\rho_2} \underset{u_2}{\bullet} \rightarrow \qquad a_l = \leftarrow \underset{v_1}{\bullet} \xrightarrow{\sigma_3} \underset{y_1}{\circ} \xrightarrow{\sigma_{23}} \cdots \xrightarrow{\sigma_{23}} \underset{y_l}{\circ} \xrightarrow{\sigma_2} \underset{v_2}{\bullet} \rightarrow$$

Note that the arrow adjacent to u_2 on the right is determined by the segment following d_k in $\boldsymbol{\ell}_1$, but, following Section 3.1 it must be an outgoing ρ_1, ρ_{12} or ρ_{123} arrow. Similarly, the arrows to the left of v_1 and the right of v_2 are outgoing σ_1, σ_{12} or σ_{123} arrows. The portion of $\widehat{CFDAA}(S^1 \times \mathcal{P}) \boxtimes (\boldsymbol{\ell}_1, \boldsymbol{\ell}_2)$ coming from d_k and a_l has generators

$$\begin{aligned}
x \otimes x_i \otimes v_j &\quad \text{for } 1 \le i \le k,\ j \in \{1, 2\},\\
y \otimes x_i \otimes y_j &\quad \text{for } 1 \le i \le k,\ 1 \le j \le l,\\
z \otimes x_i \otimes y_j &\quad \text{for } 1 \le i \le k,\ 1 \le j \le l,\\
v \otimes u_2 \otimes v_j &\quad \text{for } j \in \{1, 2\},\\
w \otimes u_2 \otimes y_j &\quad \text{for } 1 \le j \le l.
\end{aligned}$$

For each i in $\{1, \ldots, k\}$, the trimodule operations $m_2(x, \sigma_3)$, $m_2(y, \sigma_{23})$ and $m_2(y, \sigma_2)$ give rise to the unlabelled edges

$$\begin{aligned}
x \otimes x_i \otimes v_1 \to z \otimes x_i \otimes y_1 &\quad \text{for } 1 \le i \le k,\\
y \otimes x_i \otimes y_j \to z \otimes x_i \otimes y_{j+1} &\quad \text{for } 1 \le i \le k, 1 \le j \le l-1,\\
y \otimes x_i \otimes y_l \to x \otimes x_i \otimes v_2 &\quad \text{for } 1 \le i \le k.
\end{aligned}$$

We can cancel these unlabelled edges using the edge reduction algorithm described in Section 2.3. We cancel them in order of increasing i and, for fixed i, in the order above. It is not difficult to check that the only additional incoming arrows at $z \otimes x_i \otimes y_j$ and $z \otimes x_i \otimes v_2$ are given by

$$\begin{aligned}
z \otimes x_i \otimes y_j \xrightarrow{\tau_{23}} z \otimes x_i \otimes y_{j+1} &\quad \text{for } 1 \le i \le k, 1 \le j \le l-1,\\
z \otimes x_i \otimes y_l \xrightarrow{\tau_{23}} x \otimes x_i \otimes v_2 &\quad \text{for } 1 \le i \le k,\\
y \otimes x_{i-1} \otimes y_j \to z \otimes x_i \otimes y_j &\quad \text{for } 1 \le i \le k, 1 \le j \le l-1,\\
x \otimes x_{i-1} \otimes y_l \to x \otimes x_i \otimes v_2 &\quad \text{for } 1 \le i \le k.
\end{aligned}$$

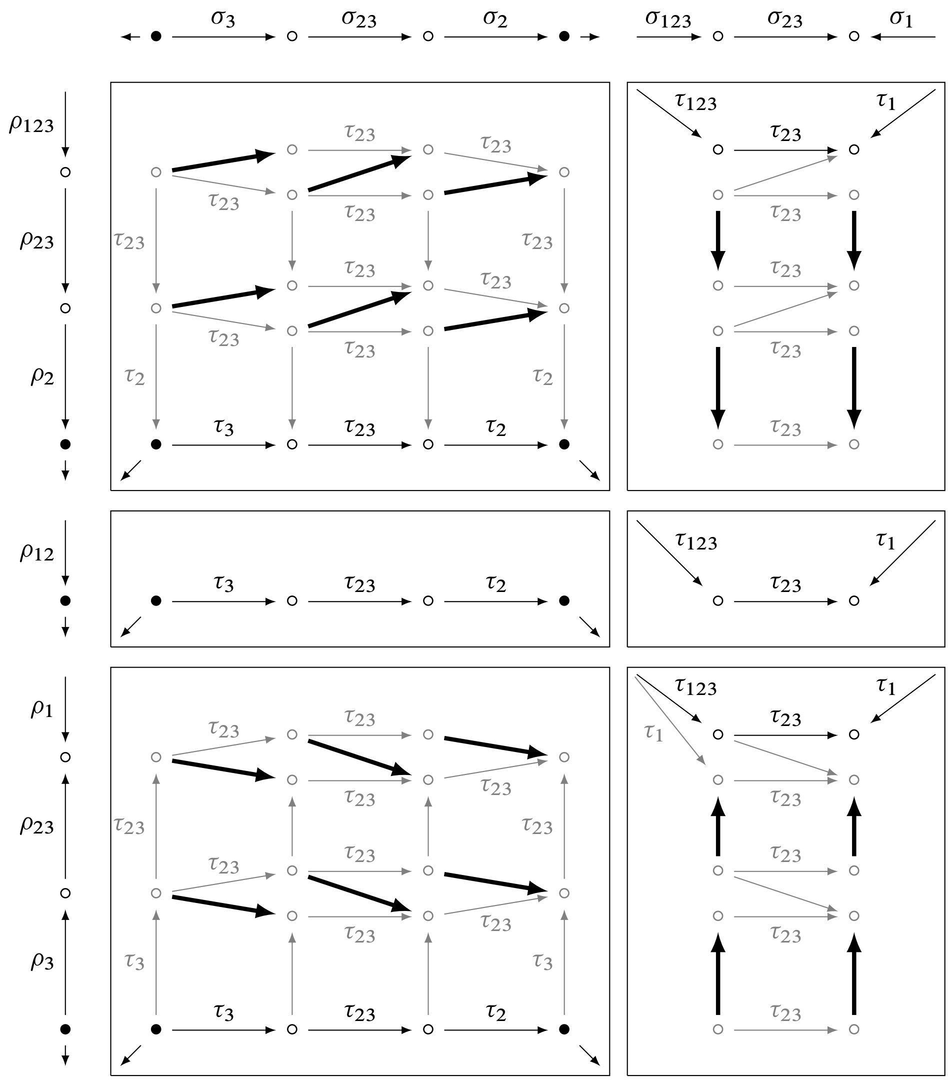

Figure 10: Each box contains the portion of $\widehat{CFDAA}(S^1\times\mathcal{P})\boxtimes(\boldsymbol{\ell}_1,\boldsymbol{\ell}_2)$ coming from a d_2, d_0 or d_{-2} segment in $\boldsymbol{\ell}_1$ (top, middle and bottom, respectively) and an a_2 or b_2 segment in $\boldsymbol{\ell}_2$ (left and right, respectively). Thick unlabelled arrows can be removed with the edge reduction algorithm; gray indicates generators and arrows that are eliminated when the differentials are cancelled.

It follows that each time we use the edge reduction algorithm on one of the unlabelled arrows mentioned above, there are no other incoming arrows at the terminal vertex that have not already been cancelled, and so cancelling the arrow produces no new arrows. After cancelling all of the unlabelled arrows, the only remaining generators are $v \otimes u_2 \otimes v_1$, $v \otimes u_2 \otimes v_2$, and $v \otimes u_2 \otimes y_j$ for $1 \le j \le l$. Arrows between these generators arise from trimodule operations involving only the generators v and w and with no ρ inputs; there are only three:

$$m_2(v, \sigma_3) = \tau_3 \otimes w, \quad m_2(w, \sigma_{23}) = \tau_{23} \otimes w \quad \text{and} \quad m_2(w, \sigma_2) = \tau_2 \otimes v.$$

It follows that the only arrows in the portion of $\widehat{CFDAA}(S^1 \times \mathcal{P}) \boxtimes (\boldsymbol{\ell}_1, \boldsymbol{\ell}_2)$ coming from the segments d_k and a_l are

$$\begin{aligned} v \otimes u_2 \otimes v_1 &\xrightarrow{\tau_3} w \otimes u_2 \otimes y_1, \\ w \otimes u_2 \otimes y_j &\xrightarrow{\tau_{23}} w \otimes u_2 \otimes y_{j+1} \quad \text{for } 1 \le j \le l-1, \\ w \otimes u_2 \otimes y_l &\xrightarrow{\tau_2} v \otimes u_2 \otimes v_2. \end{aligned}$$

That is, the portion of $\widehat{CFDAA}(S^1 \times \mathcal{P}) \boxtimes (\boldsymbol{\ell}_1, \boldsymbol{\ell}_2)$ coming from the segments d_k and a_l is an a_l segment from $v \otimes u_2 \otimes v_1$ to $v \otimes u_2 \otimes v_2$.

Finally, note that there can be no arrows connecting $w \otimes u_2 \otimes y_j$ to any other portions of $\widehat{CFDAA}(S^1 \times \mathcal{P}) \boxtimes (\boldsymbol{\ell}_1, \boldsymbol{\ell}_2)$ arising from different segments, since the only arrow connecting u_2 to a generator in a different segment of $\boldsymbol{\ell}_1$ is an outgoing ρ_1, ρ_{12} or ρ_{123} arrow, and the only trimodule operations involving these inputs also have σ_1, σ_{12} or σ_{123} as an input. The outgoing arrows from u_2 and from v_1 and v_2 do give rise to additional arrows out of $v \otimes u_2 \otimes v_1$ and $v \otimes u_2 \otimes v_2$; these show up in the portion of $\widehat{CFDAA}(\mathcal{Y}_\mathcal{P}) \boxtimes (\boldsymbol{\ell}_1, \boldsymbol{\ell}_2)$ coming from the segment following d_k in $\boldsymbol{\ell}_1$ and the segment following or preceding a_l in $\boldsymbol{\ell}_2$.

For other pairs of segments in $\boldsymbol{\ell}_1$ and $\boldsymbol{\ell}_2$, the proof is similar. Figure 10 depicts the relevant portion of $\widehat{CFDAA}(S^1 \times \mathcal{P}) \boxtimes (\boldsymbol{\ell}_1, \boldsymbol{\ell}_2)$ for d_2, d_0 or d_{-2} paired with a_2 or b_2. Any d_k paired with a_l or b_l behaves like one of these cases, depending on the sign of k. Note that, if $l < 0$, we simply take the mirror image of these diagrams, since $a_{-l} = \bar{a}_l$. This proves (1) and (2).

(3) can be deduced from (4) by observing that a c_l segment from v_1 to v_2 is the same as a d_{-l} segment from v_2 to v_1. To prove (4), consider pairing d_k in $\boldsymbol{\ell}_1$ with d_l in $\boldsymbol{\ell}_2$. The behaviour depends on the sign of k and l; Figures 11 and 12 depict the cases with k in $\{2, 0, -2\}$ and l in $\{3, 0, -3\}$. If $k = 0$, it is clear that the result is a segment of type d_l, and similarly, if $l = 0$, the result is a segment of type d_k. The case of k and l positive behaves like the top left box in Figure 11, and the case of k and l negative behaves like the bottom box in Figure 12. In each case all generators cancel except for those along the top and right edges, resulting in a segment of type d_{k+l}.

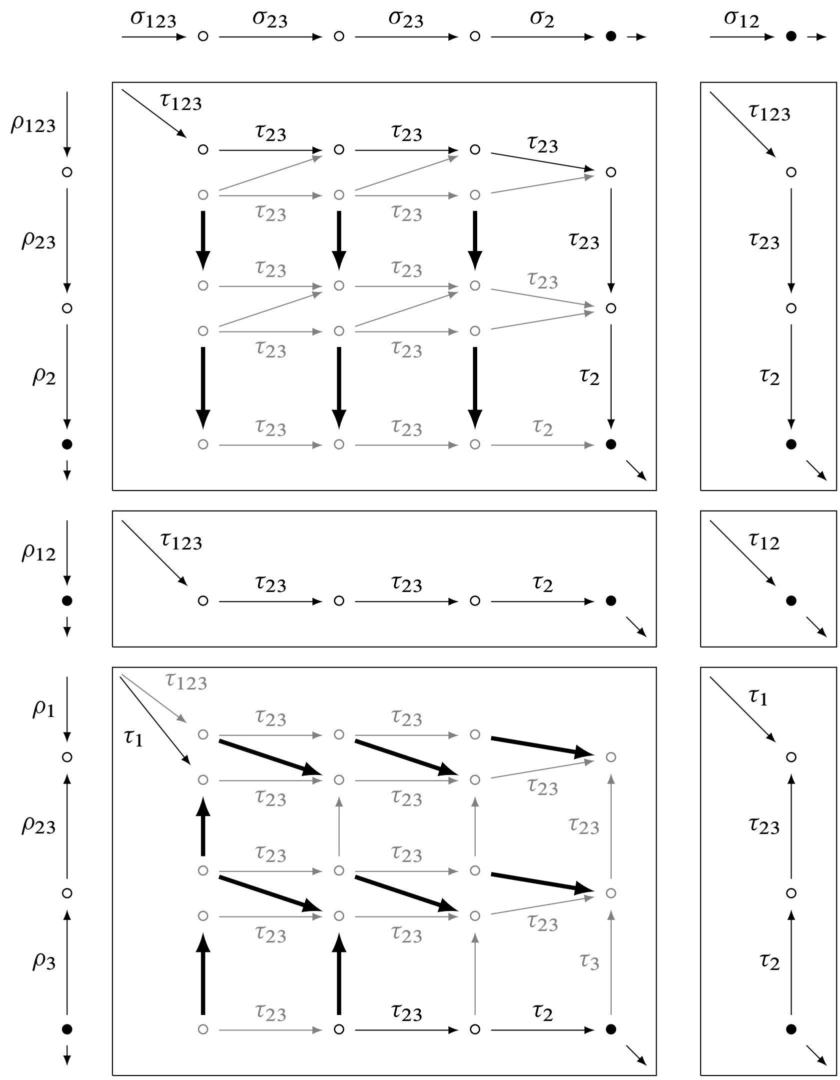

Figure 11: The portions of $\widehat{CFDAA}(S^1 \times \mathcal{P}) \boxtimes (\boldsymbol{\ell}_1, \boldsymbol{\ell}_2)$ coming from a d_2, d_0 or d_{-2} segment in $\boldsymbol{\ell}_1$ and a d_3 or d_0 segment in $\boldsymbol{\ell}_2$. Thick unlabelled arrows can be removed with the edge reduction algorithm; gray indicates generators and arrows that are eliminated when the differentials are cancelled.

The case that k and l have opposite signs is slightly more complicated. Assume first that k is negative and l is positive. If $k \geq -l$, the resulting complex looks like the bottom left box in Figure 11. Note that starting from the top left corner, there is a path to the bottom right corner consisting of a τ_1 arrow, an odd length "zigzag" sequence

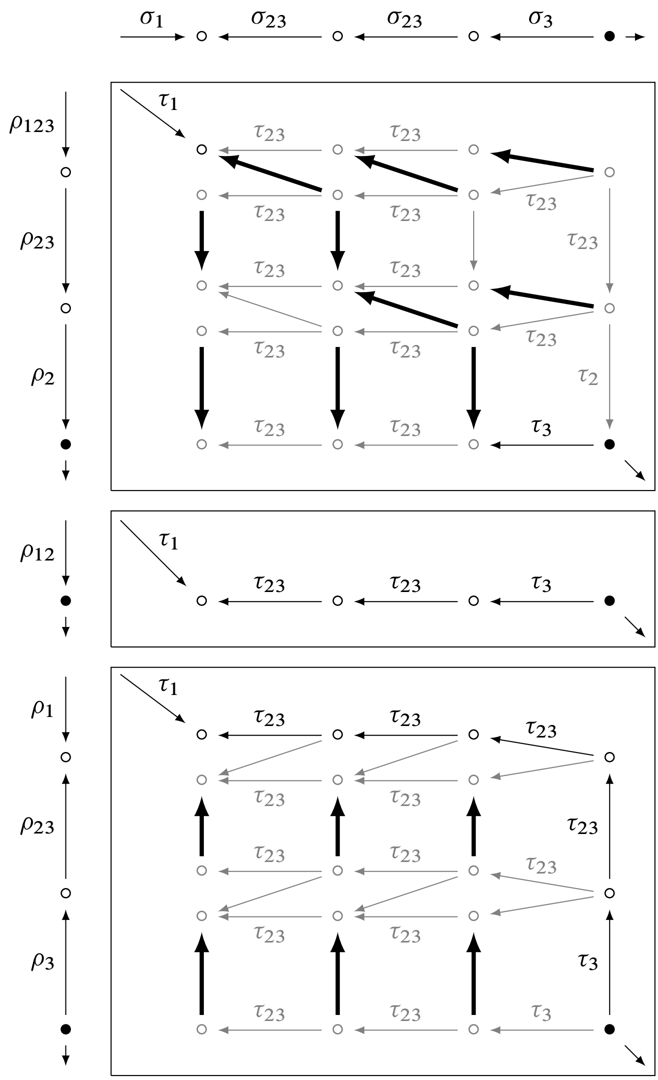

Figure 12: Each box contains the portion of $\widehat{CFDAA}(S^1 \times \mathcal{P}) \boxtimes (\boldsymbol{\ell}_1, \boldsymbol{\ell}_2)$ coming from a d_2, d_0 or d_{-2} segment in $\boldsymbol{\ell}_1$ and a d_{-3} segment in $\boldsymbol{\ell}_2$. Thick unlabelled arrows can be removed with the edge reduction algorithm; gray indicates generators and arrows that are eliminated when the differentials are cancelled.

of unlabelled arrows, $k+l$ τ_{23} arrows and a τ_2 arrow. Everything in the diagram not involved in this sequence can be cancelled without adding new arrows. Cancelling the remaining unlabelled arrows turns the τ_1 arrow and the first τ_{23} arrow into a τ_{123} arrow if $k+l>0$, or it turns the τ_1 arrow and τ_2 arrow into a τ_{12} arrow if $k+l=0$. The

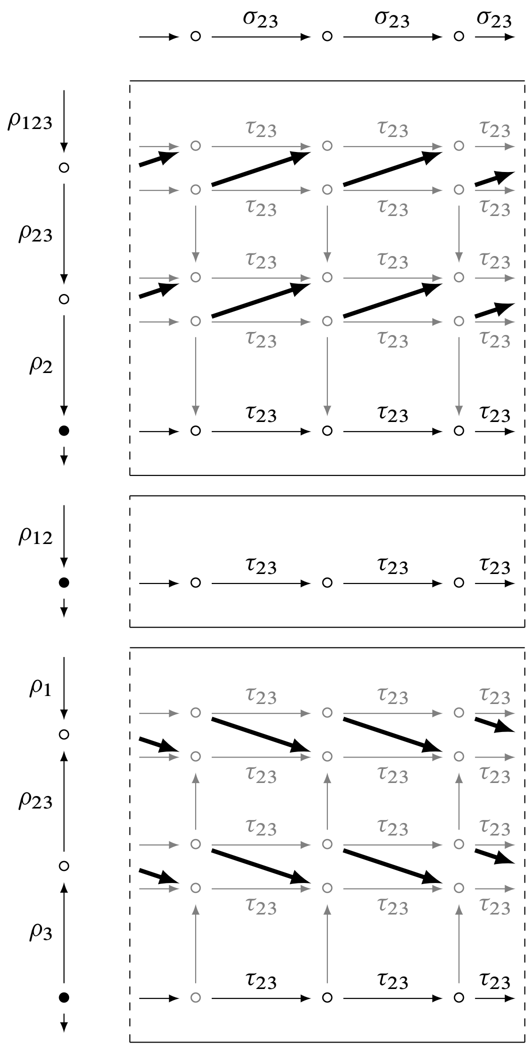

Figure 13: The portions of $\widehat{CFDAA}(S^1 \times \mathcal{P}) \boxtimes (\boldsymbol{\ell}_1, \boldsymbol{\ell}_2)$ coming from a d_2, d_0 or d_{-2} segment in $\boldsymbol{\ell}_1$ when $\boldsymbol{\ell}_2$ is a collection of e^* segments. The left and right edges of each box should be identified. Thick unlabelled arrows can be removed with the edge reduction algorithm; gray indicates generators and arrows that are eliminated when the differentials are cancelled. The result is a copy of $\boldsymbol{\ell}_2$ for each ι_0–generator in $\boldsymbol{\ell}_1$.

result is a segment of type d_{k+l}. The case that $k < -l$ is slightly different. It is not pictured separately, but the main difference is that the zigzag sequence of unlabelled arrows starting at the end of the τ_1 arrow has even length and ends on the right side of the diagram instead of the bottom. It is then followed by $-k-l-1$ backwards τ_{23} arrows and a backwards τ_3 arrow. Everything not involved in this sequence cancels, and removing the unlabelled arrows produces a segment of type d_{k+l} from the top left corner to the bottom right corner.

To complete the proof of (4), consider the case that k is positive and l is negative. If $l<-k$, the resulting complex looks like the top box in Figure 12. The complex reduces to a τ_1 arrow followed by a chain of unlabelled arrows, $-k-l-1$ backwards τ_{23} arrows and a backwards τ_3 arrow. This further reduces to a chain of type d_{k+l}. If instead $l\geq -k$, the chain of unlabelled arrows following the τ_1 arrow ends on the right side of the diagram instead of the bottom. Again, the complex reduces to a chain of type d_{k+l}.

Finally, we must consider the case the case that $\boldsymbol{\ell}_2$ cannot be written in standard notation. In this case, $\boldsymbol{\ell}_2$ is a collection of e^* segments. Note that in this case we can ignore all operations in Table 1 that involve v, x, or σ inputs other than σ_{23} (this only leaves seven operations). It is easy to see that the complex $\widehat{CFDAA}(S^1\times\mathcal{P})\boxtimes(\boldsymbol{\ell}_1,\boldsymbol{\ell}_2)$ collapses to a copy of $\boldsymbol{\ell}_2$ for each d_k segment in $\boldsymbol{\ell}_1$ (see Figure 13). □

In practice, it is helpful to compute $\mathrm{M}(\boldsymbol{\ell}_1,\boldsymbol{\ell}_2)$ by creating an i by j grid, where i is the length of $\boldsymbol{\ell}_1$ and j is the length of $\boldsymbol{\ell}_2$. The (i,j) entry of this grid is a square containing a single segment connecting two of its corners, as dictated by Proposition 6.4 (see, for example, Figure 14). The collection of loops $\mathrm{M}(\boldsymbol{\ell}_1,\boldsymbol{\ell}_2)$ can now be read off the grid by identifying the top and bottom edges and the left and right edges. Note that the number of disjoint loops in $\mathrm{M}(\boldsymbol{\ell}_1,\boldsymbol{\ell}_2)$ is given by $\gcd(i,j')$, where i is the number of d_k segments in $\boldsymbol{\ell}_1$ (by assumption, this is the length of $\boldsymbol{\ell}_1$) and j' is the number of d_k segments in $\boldsymbol{\ell}_2$ minus the number of c_k segments in $\boldsymbol{\ell}_2$ (if $j'=0$, we use the convention that $\gcd(i,0)=i$; our assumptions rule out the case that $i=0$). Note that i and j' can be given in terms of the $(\mathbb{Z}/2\mathbb{Z})$–grading on $\boldsymbol{\ell}_1$ and $\boldsymbol{\ell}_2$: up to sign, we have $i=\chi_\bullet(\boldsymbol{\ell}_1)$ and $j'=\chi_\bullet(\boldsymbol{\ell}_2)$. To see this, observe that, for an appropriate choice of relative grading, each d_k segment contributes 1 to $\chi_\bullet$, each c_k segment contributes -1 and each a_k or b_k segment contributes 0.

Lemma 6.5 *Consider a pair of loop-type, bordered, rational homology solid tori (M_1,α_1,β_1) and (M_2,α_2,β_2). If the loops representing $\widehat{CFD}(M_1,\alpha_1,\beta_1)$ contain only standard unstable chains then*

$$\mathcal{M}_{1,2}=\mathcal{M}((M_1,\alpha_1,\beta_1),(M_2,\alpha_2,\beta_2))$$

is of loop-type. If, in addition, the loops representing $\widehat{CFD}(M_2,\alpha_2,\beta_2)$ contain only standard unstable chains, then $\mathcal{M}_{1,2}$ is of simple loop-type.

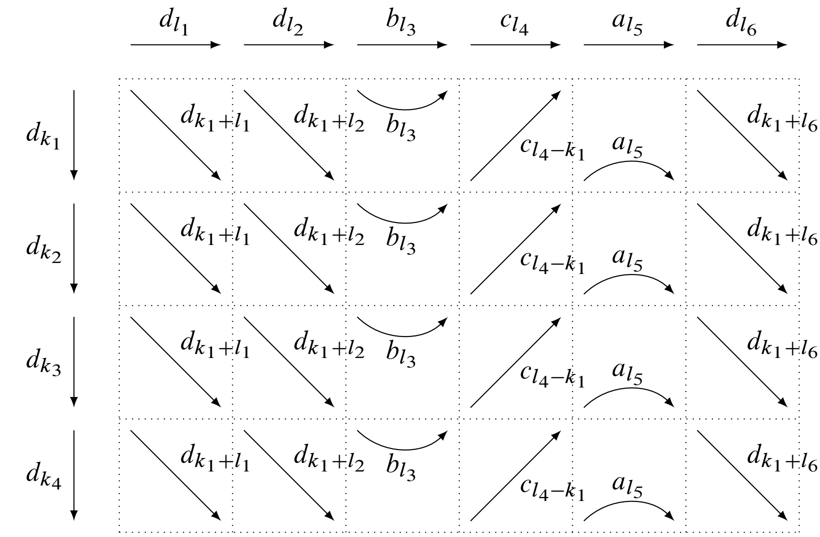

Figure 14: Computing $\mathrm{M}(\boldsymbol{\ell}_1, \boldsymbol{\ell}_2)$ for two loops $\boldsymbol{\ell}_1 = (d_{k_1} d_{k_2} d_{k_3} d_{k_4})$ and $\boldsymbol{\ell}_2 = (d_{l_1} d_{l_2} b_{l_3} c_{l_4} a_{l_5} d_{l_6})$.

Proof If $\widehat{CFD}(M_1, \alpha_1, \beta_1)$ and $\widehat{CFD}(M_2, \alpha_2, \beta_2)$ are collections of loops and the loops in $\widehat{CFD}(M_1, \alpha_1, \beta_1)$ contain only standard unstable chains, then, by Proposition 6.4, $\widehat{CFD}(\mathcal{M}_{1,2})$ is a collection of loops. Moreover, if the loops in $\widehat{CFD}(M_2, \alpha_2, \beta_2)$ also contain only standard unstable chains, Proposition 6.4 implies that the resulting loops in $\widehat{CFD}(\mathcal{M}_{1,2})$ contain only standard unstable chains, and in particular are simple. It only remains to check that $\widehat{CFD}(\mathcal{M}_{1,2})$ has exactly one loop for each spin^c–structure on $\mathcal{M}_{1,2}$.

Recall that the operation $\mathcal{M}$ corresponds to gluing two bordered manifolds to two of the three boundary components of $\mathcal{P} \times S^1$, where $\mathcal{P}$ is S^2 with three disks removed. Thus $\partial(S^1 \times \mathcal{P})$ has three connected components; denote the i^{th} connected component by $\partial(S^1 \times \mathcal{P})_i$. For each $i \in \{1, 2, 3\}$ let f_i denote a curve in $\partial(S^1 \times \mathcal{P})_i$ which is a fibre $\{\text{pt}\} \times S^1$ and let b_i denote the relevant component of $\partial\mathcal{P} \times \{\text{pt}\}$. According to the conventions introduced in Section 6.1, applying the operation $\mathcal{M}$ to two bordered manifolds (M_1, α_1, β_1) and (M_2, α_2, β_2) corresponds to gluing M_1 and M_2 to $\mathcal{P} \times S^1$ by identifying α_1 with f_1, β_1 with b_1, α_2 with f_2 and β_2 with b_2. The result is the manifold $M_3 = \mathcal{M}_{1,2}$.

We consider homology groups with coefficients in $\mathbb{Z}$. The spin^c–structures on a 3–manifold M with boundary are indexed by $H^2(M) \cong H_1(M, \partial M)$. Using the appropriate Mayer–Vietoris sequences, we have that $H_1(M_3, \partial M_3)$ is homomorphic to the quotient

$$H_1(S^1 \times \mathcal{P}, \partial(S^1 \times \mathcal{P})_3) \oplus H_1(M_1) \oplus H_1(M_2)/\{\alpha_1 \sim f_1, \beta_1 \sim b_1, \alpha_2 \sim f_2, \beta_2 \sim b_2\}.$$

Note that $H_1(S^1 \times \mathcal{P}, \partial(S^1 \times \mathcal{P})_3)$ is generated by f_i and b_i for $i \in \{1,2,3\}$ with the relations $f_1 = f_2 = f_3$ and $b_1 + b_2 = b_3 = 0$. For $i = 1, 2$, since M_i is a rational solid torus there is a unique (possibly disconnected) curve in ∂M_i that bounds a surface in M_i, to which we associate p_i/q_i (where p_i and q_i may not be relatively prime), so that $p_i\alpha_i + q_i\beta_i$ generates the kernel of the inclusion of $H_1(\partial M_i)$ into $H_1(M_i)$. The long exact sequence for relative homology gives

$$H_1(M_i) \cong H_1(M_i, \partial M_i) \oplus H_1(\partial M_i)/\langle p_i\alpha_i + q_i\beta_i \rangle.$$

In $H_1(M_3, \partial M_3)$, the relation $p_i\alpha_i + q_i\beta_i = 0$ translates to $p_i f_i + q_i b_i = q_i b_i = 0$. It follows that

$$H_1(M_3, \partial M_3) \cong H_1(M_1, \partial M_1) \oplus H_1(M_2, \partial M_2) \oplus \langle b_1 \rangle/(q_1 b_1 = q_2 b_1 = 0).$$

Thus, for each spin^c–structure on M_1 and each spin^c–structure on M_2, there are $\gcd(q_1, q_2)$ spin^c–structures on M_3. Note that q_2 may be 0; in this case we use the convention that $\gcd(q, 0) = q$. The assumption that $\widehat{CFD}(M_1, \alpha_1, \beta_1)$ contains only standard unstable chains implies that $q_1 \neq 0$.

Recall that the rational longitude of a rational homology solid torus can be read off of the bordered invariants. More precisely, by Proposition 2.10, the curves in ∂M_i which are nullhomologous in M_i are integer multiples of

$$\chi_\circ(\widehat{CFD}(M_i, \alpha_i, \beta_i; \mathfrak{s}_i))\alpha_i + \chi_\bullet(\widehat{CFD}(M_i, \alpha_i, \beta_i; \mathfrak{s}_i))\beta_i,$$

where $\mathfrak{s}_i$ is any spin^c–structure on M_i. Thus $q_i = \chi_\bullet(\widehat{CFD}(M_i, \alpha_i, \beta_i; \mathfrak{s}_i))$. By assumption, $\widehat{CFD}(M_i, \alpha_i, \beta_i; \mathfrak{s}_i)$ is a loop for each $\mathfrak{s}_i$. It was observed above that, given two loops $\boldsymbol{\ell}_1$ and $\boldsymbol{\ell}_2$ with $\boldsymbol{\ell}_1$ consisting only of standard unstable chains, $\mathrm{M}(\boldsymbol{\ell}_1, \boldsymbol{\ell}_2)$ is a collection of $\gcd(\chi_\bullet(\boldsymbol{\ell}_1), \chi_\bullet(\boldsymbol{\ell}_2))$ loops. Thus, for each loop in $\widehat{CFD}(M_1, \alpha_1, \beta_1)$ and for each loop in $\widehat{CFD}(M_2, \alpha_2, \beta_2)$, there are $\gcd(q_1, q_2)$ loops in $\widehat{CFD}(\mathcal{M}_{1,2})$. It follows that $\widehat{CFD}(\mathcal{M}_{1,2})$ has exactly one loop for each spin^c–structure on M_3 (where $\mathcal{M}_{1,2} \cong (M_3, \alpha_3, \beta_3)$ as bordered manifolds). $\square$

6.4 Bordered invariants of graph manifolds

With this description of $\mathcal{T}$, $\mathcal{E}$ and $\mathcal{M}$ on bordered invariants in hand, we may now return our focus to graph manifolds. The first author has described (and implemented) an algorithm for computing the (bordered) Heegaard Floer invariants of graph manifolds [8]; we will now outline a version of this algorithm for graph manifolds with a single boundary component and adapt it to the loops setup. Recall that, given a graph Γ with associated bordered manifold $(M_\Gamma, \alpha, \beta)$ (as described in Section 6.1), we write $\widehat{CFD}(\Gamma)$ for $\widehat{CFD}(M_\Gamma, \alpha, \beta)$.

In order to compute $\widehat{CFD}(\Gamma)$, we inductively build up the plumbing tree Γ using the three plumbing tree operations $\mathcal{T}^{\pm 1}$, $\mathcal{E}$ and $\mathcal{M}$ depicted in Figure 8, starting from the plumbing tree

$$\Gamma_0 = \overset{0}{\bullet} \text{----}$$

Note that M_{Γ_0} is a solid torus, with bordered structure $(M_{\Gamma_0}, \alpha, \beta)$, where α is a fibre of the S^1–bundle over D^2 (ie a longitude of the solid torus) and β is a curve in the base surface (ie a meridian of the solid torus, identified by $\partial D^2 \times \{\mathrm{pt}\}$). Thus $\widehat{CFD}(\Gamma_0)$ is represented by the loop $\boldsymbol{\ell}_\bullet \sim (d_0)$. As we apply the operations $\mathcal{T}$, $\mathcal{E}$ and $\mathcal{M}$, we keep track of the bordered invariants of the relevant manifolds. Let Γ_1 and Γ_2 be single-boundary plumbing trees. As shown in the previous section,

$$\begin{aligned}
\widehat{CFD}(\mathcal{T}(\Gamma_1)) &\cong \widehat{T}_{\mathrm{st}} \boxtimes \widehat{CFD}(\Gamma_1),\\
\widehat{CFD}(\mathcal{E}(\Gamma_1)) &\cong \widehat{T}_{\mathrm{st}} \boxtimes \widehat{T}_{\mathrm{du}} \boxtimes \widehat{T}_{\mathrm{st}} \boxtimes \widehat{CFD}(\Gamma_1),\\
\widehat{CFD}(\mathcal{M}(\Gamma_1, \Gamma_2)) &\cong \widehat{CFDAA}(\Gamma_{\mathcal{M}}) \boxtimes (\widehat{CFD}(\Gamma_1), \widehat{CFD}(\Gamma_2)).
\end{aligned}$$

Specializing Lemmas 6.2 and 6.5 to graph manifold leads to the following observation:

Lemma 6.6 *Let Γ be a single boundary plumbing tree constructed as described above from copies of Γ_0 using the operations $\mathcal{T}^\pm$, $\mathcal{E}$, and $\mathcal{M}$. $\widehat{CFD}(\Gamma)$ has simple loop-type as long as, each time the operation $\mathcal{M}$ is applied, the two input plumbing trees have simple loop-type bordered invariants with only unstable chains in standard notation.*

Proof This follows from Lemma 6.2, which says that $\mathcal{T}^\pm$ and $\mathcal{E}$ take simple loop-type manifolds to simple loop-type manifolds, and Lemma 6.5, which says that $\mathcal{M}$ takes two simple loop-type manifolds to a simple loop-type manifold provided the simple loops corresponding to both inputs consist only of unstable chains. □

With this observation, we can describe a large family of simple loop-type manifolds.

For each vertex v of a plumbing tree Γ, let $w(v)$ denote the Euler weight associated to v and let $n_+(v)$ and $n_-(v)$ denote the number of neighbouring vertices v' for which $w(v') \geq 0$ and $w(v') \leq 0$, respectively. We will say that v is a *bad vertex* if

$-n_-(v) < w(v) < n_+(v)$, and otherwise v is a *good vertex* (this should be viewed as a generalization of the notion of bad vertices defined for negative definite plumbing trees in [23]).

Proposition 6.7 *Let Γ be a plumbing tree with a single boundary edge at the vertex v_0, and suppose that every vertex other than v_0 is good. Then $\widehat{CFD}(\Gamma)$ is a collection of loops consisting only of standard unstable chains; up to reversal we can assume these unstable chains are of type d_k. Moreover, if $w(v_0)$ is (strictly) greater than $n_+(v_0)$, then these unstable chains all have subscripts (strictly) greater than 0. If $w(v_0)$ is (strictly) less than $n_-(v_0)$ then the unstable chains all have subscript (strictly) less than 0.*

Proof We proceed by induction on the number of vertices of Γ. The base case, where Γ has only one vertex v_0, is trivial; $\widehat{CFD}(\Gamma)$ in this case is given by the loop $(d_{w(v_0)})$.

For the inductive step, first suppose the boundary vertex v_0 of Γ has valence two (including the boundary edge). That is, Γ has the form

$\Gamma =$ 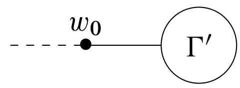

where $w_0 = w(v_0)$. Let v_1 be the boundary vertex of Γ', and let w_1 be the corresponding weight. Let $n_\pm(v_1)$ denote the counts of neighbouring vertices defined above for v_1 as a vertex in Γ, and let $n'_\pm(v_1)$ denote these counts for v_1 as a vertex in Γ' (ie ignoring the vertex v_0). Since v_1 is a good vertex, we have one of the following two cases:

(1) $w_1 \le -n_-(v_1) \le -n'_-(v_1) \le 0$.

(2) $w_1 \ge n_+(v_1) \ge n'_+(v_1) \ge 0$.

Note that $\Gamma = \mathcal{T}^{w_0}(\mathcal{E}(\Gamma'))$, and so $\widehat{CFD}(\Gamma)$ is given by $\mathrm{T}^{w_0}\big(\mathrm{E}(\widehat{CFD}(\Gamma'))\big)$. Furthermore, we assume by induction that the proposition holds for Γ'.

In case (1), we have that $\widehat{CFD}(\Gamma')$ consists of d_k segments with nonpositive subscripts. Thus in dual notation $\mathrm{E}(\widehat{CFD}(\Gamma'))$ is a loop consisting of d_k^* segments with nonnegative subscripts, and in standard notation $\mathrm{E}(\widehat{CFD}(\Gamma'))$ consists of d_k segments with nonnegative subscripts. Moreover, if $w_0 \le 0$, then w_1 is strictly less than $-n'_-(v_1) = -n_-(v_1) + 1$, so the d_k segments in $\widehat{CFD}(\Gamma)$ have strictly negative subscripts, the d_k^* segments in $\mathrm{E}(\widehat{CFD}(\Gamma'))$ have strictly positive subscripts, and in standard notation $\mathrm{E}(\widehat{CFD}(\Gamma'))$ consists only of d_0 and d_1 segments. It follows that

$\widehat{CFD}(\Gamma)$ is a collection of loops consisting of d_k segments with subscripts satisfying $k<0$ if $w_0<-n_-(v_0)=-1$, $k\le 0$ if $w_0=n_-(v_0)=-1$, $k\ge 0$ if $w_0=n_+(v_0)\ge 0$ and $k>0$ if $w_0>n_+(v_0)\ge 0$. Thus the proposition holds for Γ.

In case (2), $\widehat{CFD}(\Gamma')$ consists of d_k segments with nonnegative subscripts. Thus $\mathrm{E}(\widehat{CFD}(\Gamma'))$ consists of d_k^* with nonpositive subscripts in dual notation and of d_k with nonpositive subscripts in standard notation. Moreover, if $w_0\ge 0$ then $n_+(v_1)>n'_+(v_1)$, so the d_k segments in $\widehat{CFD}(\Gamma')$ have strictly positive subscripts, the d_k^* segments in $\mathrm{E}(\widehat{CFD}(\Gamma'))$ have strictly negative subscripts and the d_k segments in $\mathrm{E}(\widehat{CFD}(\Gamma'))$ have subscripts in $\{0,-1\}$. It follows that $\widehat{CFD}(\Gamma)$ is a collection of loops consisting of d_k segments with subscripts satisfying $k<0$ if $w_0<-n_-(v_0)\le 0$, $k\le 0$ if $w_0=n_-(v_0)\le 0$, $k\ge 0$ if $w_0=n_+(v_0)=1$ and $k>0$ if $w_0>n_+(v_0)=1$. Thus the proposition holds for Γ.

Now suppose that v_0 has valence higher than two. This means that Γ can be obtained as $\mathcal{M}(\Gamma_1,\Gamma_2)$, where Γ_1 and Γ_2 have fewer vertices than Γ. By induction, the proposition holds for Γ_1 and Γ_2, and so $\widehat{CFD}(\Gamma_1)$ and $\widehat{CFD}(\Gamma_2)$ may be represented by a collection of loops consisting only of standard type d_k chains. By Proposition 6.4, merging two (collections of) loops with only d_k chains produces a new collection of loops with only d_k chains. Moreover, each chain in the $\widehat{CFD}(\Gamma)$ is of the form d_{k+l} for some chains d_k in $\widehat{CFD}(\Gamma_1)$ and d_l in $\widehat{CFD}(\Gamma_2)$, so the maximum (resp. minimum) subscript in $\widehat{CFD}(\Gamma)$ is the sum of the maximum (resp. minimum) subscripts in $\widehat{CFD}(\Gamma_1)$ and $\widehat{CFD}(\Gamma_2)$.

For $i\in\{0,1\}$, let v_i be the boundary vertex of Γ_i and let $w_i=w(v_i)$ be the corresponding weight. We can choose any values for w_i provided $w_1+w_2=w_0$. If $w_0\ge n_+(v_0)=n_+(v_1)+n_+(v_2)$, then we can choose $w_1=n_+(v_1)$ and $w_2\ge n_+(v_2)$. By the inductive assumption, $\widehat{CFD}(\Gamma_1)$ and $\widehat{CFD}(\Gamma_2)$ both have only nonnegative subscripts, so the same is true of $\widehat{CFD}(\Gamma_2)$. Furthermore, if $w_0>0$ then $w_2>n_+(v_2)$, so $\widehat{CFD}(\Gamma_2)$ has only strictly positive subscripts and the same follows for $\widehat{CFD}(\Gamma)$. Similarly, if w_0 is (strictly) less than $n_-(v_0)=n_-(v_1)+n_-(v_2)$, we can choose $w_1=n_-(v_1)$ and w_2 (strictly) less than $n_-(v_2)$ and conclude that $\widehat{CFD}(\Gamma)$ only has subscripts (strictly) less than 0. $\square$

Note that if Γ is a plumbing tree with at most one bad vertex, we can compute $\widehat{HF}(M_\Gamma)$ by taking the dual filling of $\widehat{CFD}(\Gamma')$, where Γ' is obtained from Γ by adding a boundary edge to the bad vertex, or to any vertex if there are no bad vertices. $\widehat{CFD}(\Gamma')$, which has simple loop-type by Proposition 6.7, can be computed using the operations $\mathrm{T}^{\pm 1}$, E and M. In particular, we have the following generalization of [23, Lemma 2.6]:

Corollary 6.8 *If Γ is a closed plumbing tree with no bad vertices, then the manifold M_Γ is an L–space.*

Proof Add a boundary edge to any vertex of Γ to produce a single boundary plumbing tree Γ'. By Proposition 6.7, $\widehat{CFD}(\Gamma')$ is a collection of loops consisting only of type d_k segments with either all subscripts greater than or equal to zero or all subscripts less than or equal to zero. By Observation 3.9, it follows that in dual notation the loops representing $\widehat{CFD}(\Gamma')$ have no stable chains. M_Γ is obtained from $M_{\Gamma'}$ by dual filling, and dual filling is an L–space if there are no stable chains in dual notation. □

Recall that in the context of Theorem 1.3 it is important to distinguish solid torus-like manifolds from other manifolds of simple loop-type. Toward that end, we check the following:

Proposition 6.9 *Let Γ be a plumbing tree with a single boundary edge at the vertex v_0, and suppose that every vertex other than v_0 is good. Then M_Γ is solid torus-like if and only if it is a solid torus.*

Proof We proceed by induction on the size of Γ as in the proof of Proposition 6.7. Applying Dehn twists does not change whether or not a manifold is a solid torus, nor does it change whether or not the corresponding bordered invariants are solid torus-like. It follows that the proposition holds for $\mathcal{T}^{\pm}(\Gamma)$ and $\mathcal{E}(\Gamma)$ if it holds for Γ. We only need to check that it holds for $\mathcal{M}(\Gamma_1, \Gamma_2)$, assuming by induction that it holds for Γ_1 and Γ_2. By Proposition 6.7, we may assume that the loops representing $\widehat{CFD}(\Gamma_1)$ and $\widehat{CFD}(\Gamma_2)$ consist only of d_k chains.

Suppose $\widehat{CFD}(\mathcal{M}(\Gamma_1, \Gamma_2))$ is solid torus-like. Recall that, by the appropriate generalization of Lemma 3.15, the loops representing $\widehat{CFD}(\mathcal{M}(\Gamma_1, \Gamma_2))$ consist only of d_k segments such that the maximum and minimum subscripts appearing differ by at most one. Note also that the difference between maximum and minimum subscripts is additive under $\mathcal{M}$. It follows that either Γ_1 or Γ_2 must have bordered invariant consisting only of d_n segments for a fixed $n \in \mathbb{Z}$; say Γ_2 has this property. M_{Γ_2} is solid torus-like; by induction, M_{Γ_2} is a solid torus, and $\widehat{CFD}(\Gamma_2)$ is given by the single loop (d_n). Proposition 6.4 then implies that $\widehat{CFD}(\mathcal{M}(\Gamma_1, \Gamma_2))$ is obtained from $\widehat{CFD}(\Gamma_1)$ by adding n to the subscript of each segment. In other words, $\widehat{CFD}(\mathcal{M}(\Gamma_1, \Gamma_2)) \cong \widehat{CFD}(\mathcal{T}^n(\Gamma_1))$.

This equivalence does not only hold on the level of bordered invariants. Indeed, $\mathcal{M}(\Gamma_1, \Gamma_2)$ is equivalent to $\mathcal{T}^n(\Gamma_1)$ up to the graph moves in [22]; in particular, the corresponding graph manifolds are diffeomorphic. To see this, note that Γ_2 must be equivalent to the plumbing tree

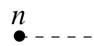

since $\widehat{CFD}$ of this tree is (d_n), and it is clear from Figure 8 that merging with this tree has the same effect as applying $\mathcal{T}^n$. Since $\widehat{CFD}(\mathcal{T}^n(\Gamma_1))$ is solid torus-like, $\widehat{CFD}(\Gamma_1)$ is solid torus-like and, by the inductive hypothesis, M_{Γ_1} is a solid torus. It follows that $M_{\mathcal{M}(\Gamma_1,\Gamma_2)} = M_{\mathcal{T}^n(\Gamma_1)} = \mathcal{T}^n(M_{\Gamma_1})$ is a solid torus. □

6.5 An explicit example: the Poincaré homology sphere

As an example of the algorithm and loop operations described above, we will compute $\widehat{HF}$ of the Poincaré homology sphere using the plumbing tree

$\Gamma =$ 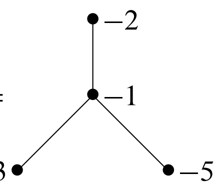

We start with the loop (d_0) representing $\widehat{CFD}(\Gamma_0)$. For this example, by abuse of notation, we will equate the bordered invariants with their loop representatives; thus, $\widehat{CFD}(\Gamma_0) = (d_0)$. We use the twist and extend operations to compute invariants for the plumbing trees

$\Gamma_1 =$ 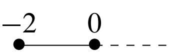

$\Gamma_2 =$ 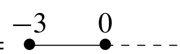

$\Gamma_3 =$ 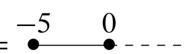

Note that $\Gamma_1 = \mathcal{E}(\mathcal{T}^{-2}(\Gamma_0))$, so

$$\widehat{CFD}(\Gamma_1) = \mathrm{E}\big(\mathrm{T}^{-2}((d_0))\big) = \mathrm{E}((d_{-2})) = (d_2^*) \sim (d_1 d_0).$$

Similarly, we find that

$$\widehat{CFD}(\Gamma_2) = \mathrm{E}\big(\mathrm{T}^{-3}((d_0))\big) = \mathrm{E}((d_{-3})) = (d_3^*) \sim (d_1 d_0 d_0),$$
$$\widehat{CFD}(\Gamma_3) = \mathrm{E}\big(\mathrm{T}^{-5}((d_0))\big) = \mathrm{E}((d_{-5})) = (d_5^*) \sim (d_1 d_0 d_0 d_0 d_0).$$

Now let

$\Gamma_4 =$ 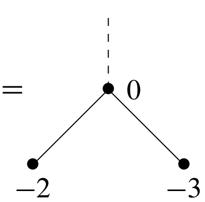

$\Gamma_5 =$ 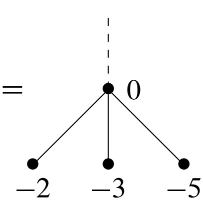

$\Gamma_6 =$ 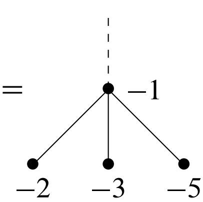

We have that $\Gamma_4 = \mathcal{M}(\Gamma_1, \Gamma_2)$, $\Gamma_5 = \mathcal{M}(\Gamma_4, \Gamma_3)$ and $\Gamma_6 = \mathcal{T}^{-1}(\Gamma_5)$. Reading diagonally from the first grid below, we see that $\widehat{CFD}(\Gamma_4) = \textsc{m}((d_1 d_0), (d_1 d_0 d_0)) = (d_2 d_0 d_1 d_1 d_1 d_0)$:

	d_1	d_0	d_0
d_1	d_2	d_1	d_1
d_0	d_1	d_0	d_0

	d_1	d_0	d_0	d_0	d_0
d_2	d_3	d_2	d_2	d_2	d_2
d_0	d_1	d_0	d_0	d_0	d_0
d_1	d_2	d_1	d_1	d_1	d_1
d_1	d_2	d_1	d_1	d_1	d_1
d_1	d_2	d_1	d_1	d_1	d_1
d_0	d_1	d_0	d_0	d_0	d_0

The second grid tells us that

$$\widehat{CFD}(\Gamma_5) = \\ (d_3 d_0 d_1 d_1 d_1 d_1 d_2 d_0 d_1 d_1 d_2 d_0 d_2 d_0 d_1 d_2 d_1 d_0 d_2 d_0 d_2 d_1 d_1 d_0 d_2 d_1 d_1 d_1 d_1 d_0).$$

Applying the operation $\textsc{t}^{-1}$, we find that

$$\widehat{CFD}(\Gamma_6) = \big(d_2 d_{-1} d_0 d_0 d_0 d_0 d_1 d_{-1} d_0 d_0 d_1 d_{-1} d_1 d_{-1} d_0 d_1 d_0 d_{-1} d_1 d_{-1} d_1 d_0 d_0 \\ d_{-1} d_1 d_0 d_0 d_0 d_0 d_{-1}\big).$$

Finally, to compute $\widehat{CF}$ of the closed manifold M_Γ we must fill in the boundary of $(M_{\Gamma_6}, \alpha, \beta)$ with a $D^2 \times S^1$ such that the meridian $\partial D^2 \times \{\mathrm{pt}\}$ glues to β and the longitude $\{\mathrm{pt}\} \times S^1$ glues to α. In other words, M_Γ is the 0–filling of $(M_{\Gamma_6}, \alpha, \beta)$, so $\widehat{CF}(M_\Gamma)$ is obtained from $\widehat{CFD}(\Gamma_6)$ by tensoring with $\ell_\circ^A$. To do this, we first write $\widehat{CFD}(\Gamma_6)$ in dual notation (using the procedure described in Section 3.2),

$$\widehat{CFD}(\Gamma_6) = (d_0^* b_1^* a_5^* b_1^* a_3^* b_1^* a_1^* b_1^* a_2^* b_2^* a_1^* b_1^* a_1^* b_3^* a_1^* b_5^* a_1^*).$$

Tensoring with $\ell_\circ$ produces one generator for each segment in dual notation, but also one differential for each type a^* segment. Since $\widehat{CFD}(\Gamma_6)$ has 17 dual segments and 8 a^* segments, all but one generator in $\ell_\circ^A \boxtimes \widehat{CFD}(\Gamma_6) \cong \widehat{CF}(M_\Gamma)$ cancels in homology. Thus $\widehat{HF}(M_\Gamma)$ has dimension 1, as is now well known (this was first calculated in [24, Section 3.2]).

Remark 6.10 While computer computation is not our primary motivation, it is worth pointing out that using loop calculus as described in this section instead of taking box tensor products of modules, bimodules and trimodules greatly improves the efficiency of the algorithm in [8] for rational homology sphere graph manifolds. This is illustrated by the fact that the example above can easily be done by hand, while computing the relevant tensor products would be tedious without a computer. The largest computation in [8] (for which the dimension of $\widehat{HF}$ is 213 312) took roughly 12 hours; when the computer implementation is adapted to use loop calculus, the same computation runs in 30 seconds.[3] The caveat is that the purely loop calculus algorithm may not work for all rational homology sphere graph manifolds, since there may be graph manifold rational homology solid tori that are not of loop-type, but in practice it works for most examples.

7 L–spaces and non-left-orderability

We conclude by proving the remaining results quoted in the introduction: Theorem 1.4, Theorem 1.5 and, finally, Theorem 1.1.

To begin, we observe that Seifert-fibred rational homology tori have simple loop-type bordered invariants. Any Seifert-fibred space over S^2 can be given a star-shaped plumbing tree in which all vertices of valence one or two have weight at most -2; in particular, there are no bad vertices except for the central vertex. Adding a boundary edge to the central vertex corresponds to removing a neighbourhood of a regular fibre, creating a Seifert-fibred space over D^2. By Proposition 6.7, such a plumbing tree has simple loop-type $\widehat{CFD}$. Moreover, by Proposition 6.9, such a plumbing tree is nonsolid torus-like unless the corresponding manifold is a solid torus.

The only other option to consider is a Seifert-fibred space over the Möbius band, since a Seifert-fibred space over any other base orbifold has $b_1 > 0$. Such a manifold can be obtained from a Seifert-fibred space over D^2 by removing a neighbourhood of a regular fibre and gluing in the Euler number 0 bundle over the Möbius band, fibre to fibre and base to base. On the plumbing tree, this corresponds to adding

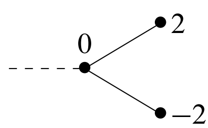

to the central vertex, or equivalently to merging with the plumbing tree

$\Gamma_N =$ 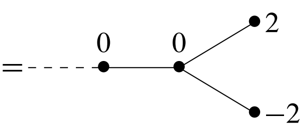

Proposition 7.1 *$\widehat{CFD}$ for the plumbing tree Γ_N above consists of two loops, one $(a_1 b_1) \sim (d_1^* d_{-1}^*)$ and the other $(e^* e^*)$.*

[3]An implementation is available from the authors upon request.

Proof Using the loop operations described in the preceding section, this is a simple computation. Note that $\Gamma_N = \mathcal{E}\big(\mathcal{M}\big(\mathcal{E}(\mathcal{T}^2(\Gamma_0)), \mathcal{E}(\mathcal{T}^{-2}(\Gamma_0))\big)\big)$. Thus $\widehat{CFD}(\Gamma_N)$ is given by

$$\begin{aligned}\mathrm{E}\big(\mathrm{M}\big(\mathrm{E}\big(\mathrm{T}^2((d_0))\big), \mathrm{E}\big(\mathrm{T}^{-2}((d_0))\big)\big)\big) &= \mathrm{E}\big(\mathrm{M}\big(\mathrm{E}((d_2)), \mathrm{E}((d_{-2}))\big)\big)\\ &= \mathrm{E}\big(\mathrm{M}((d_{-1}d_0), (d_1d_0))\big)\\ &= \mathrm{E}((d_0d_0) \amalg (d_1d_{-1}))\\ &= (d_0^*d_0^*) \amalg (d_1^*d_{-1}^*).\end{aligned}$$ □

Remark 7.2 This result was first established by a direct calculation by Boyer, Gordon and Watson [3]; this calculation is greatly simplified by appealing to loop calculus.

We now complete the proof that Seifert-fibred rational homology tori have simple loop-type.

Proof of Theorem 1.4 As observed above, the case of Seifert-fibred manifolds over D^2 is a special case of Proposition 6.7. For a Seifert-fibred manifold over the Möbius band, a plumbing tree Γ is given by $\mathcal{M}(\Gamma', \Gamma_N)$, where Γ' is a star shaped plumbing tree for a Seifert-fibred manifold over D^2. By Proposition 6.7, the loops in $\widehat{CFD}(\Gamma')$ contain only unstable chains in standard notation. By Propositions 6.4 and 7.1, we find that $\widehat{CFD}(\Gamma)$ is a collection of disjoint copies of $\widehat{CFD}(\Gamma_N)$, and in particular is a collection of simple loops. Moreover, by Lemma 6.5 there is one loop for each spin^c–structure, and so $\widehat{CFD}(\Gamma)$ is of simple loop-type. Finally, with the foregoing in place (in particular Proposition 6.9) one checks that the only solid torus-like manifold in this class is the solid torus itself. □

We are now in a position to assemble the pieces and give the proof of Theorem 1.1. A key observation is the following consequence of our gluing theorem, which provides some alternative characterizations of the set of strict L–space slopes of a given simple loop-type manifold:

Theorem 7.3 *Let M be a simple loop-type manifold that is not solid torus-like. The following are equivalent*:

(i) *γ is a strict L–space slope for M, that is, $\gamma \in \mathcal{L}_M^\circ$.*

(ii) *$M \cup_h M'$ is an L–space, where $h(\gamma) = \lambda$ and M' is a nonsolid torus-like, simple, loop-type manifold with rational longitude λ for which $\mathcal{L}_{M'}$ includes every slope other than λ.*

(iii) *$M \cup_h N$ is an L–space, where $h(\gamma) = \lambda$ and N is the twisted I–bundle over the Klein bottle with rational longitude λ.*

Proof Given $\gamma \in \mathcal{L}_M^\circ$ and a simple, loop-type manifold M' for which $\mathcal{L}_{M'}$ includes every slope other than λ, Theorem 5.7 ensures that $M \cup_h M'$ is an L–space. Indeed, for any $\gamma' \neq \gamma$ we have that $\lambda \neq h(\gamma') \in \mathcal{L}_{M'}^\circ$. This proves that (i) implies (ii).

To see that (ii) implies (iii), is suffices to observe that $N(\gamma)$ is an L–space for all γ other than the rational longitude; this is indeed the case, as observed, for example, in [3, Proposition 5] (alternatively, this fact is an exercise in loop calculus).

Finally, suppose that $M \cup_h N$ is an L–space and consider the slope γ in ∂M determined by $h^{-1}(\lambda)$. Since $N(\lambda)$ is not an L–space, by Theorem 5.7 it must be that $\gamma \in \mathcal{L}_M^\circ$, as required, so that (iii) implies (i). □

Boyer and Clay consider a collection of Seifert-fibred rational homology solid tori $\{N_t\}$ for integers $t > 1$. In this collection, $N_2 = N$, the twisted I–bundle over the Klein bottle. More generally, the N_t are examples of Heegaard Floer homology solid tori (see [12, Section 1.5] for a expanded discussion on this class of manifolds). These manifolds are easily described by the plumbing tree

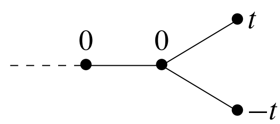

Translated into loop notation, the invariant $\widehat{CFD}(N_t, \varphi, \lambda)$ is simple, described by

$$(d_0^*)^t \amalg \left[\coprod_{i=1}^{t-1} (d_i^* d_{i-t}^*) \right].$$

This calculation is similar to that of Proposition 7.1; these are members of a much larger class of manifolds that are interesting in their own right. Note, in particular, that $N_t(\gamma)$ is an L–space for all γ other than the rational longitude; see Section 4 but compare also [3]. Therefore, the N_t satisfy the conditions of (ii) in Theorem 7.3 and we have the following:

Corollary 7.4 *Let M be a simple loop-type manifold that is not solid torus-like. The following are equivalent*:

(i) *γ is a strict L–space slope for M, that is, $\gamma \in \mathcal{L}_M^\circ$.*

(ii) *$M \cup_h N_t$ is an L–space, where $h(\gamma) = \lambda$ is the rational longitude, for any integer $t > 1$.*

(iii) *$M \cup_h N_2$ is an L–space, where $h(\gamma) = \lambda$ is the rational longitude.*

Note that Theorem 1.5 follows from the equivalence between (i) and (iii) in Corollary 7.4. This answers [2, Question 1.8] and considerably simplifies [2, Theorem 1.6] when restricting to Seifert-fibred rational homology solid tori. Indeed, we have shown:

Theorem 7.5 *Suppose $M \ncong D^2 \times S^1$ is a Seifert-fibred rational homology solid torus. The following are equivalent*:

(i) $\gamma \in (\mathcal{L}_M^\circ)^c$.

(ii) γ *is detected by a left-order* (*in the sense of Boyer and Clay* [2]).

(iii) γ *is detected by a taut foliation* (*in the sense of Boyer and Clay* [2]).

Proof This follows immediately from [2, Theorem 1.6] combined with Corollary 7.4. □

Proof of Theorem 1.1 The equivalence between (ii) and (iii) is due to Boyer and Clay [2]. To see that (i) is equivalent to either of these we first note that, if M is one of the two Seifert-fibred pieces in Y, then, according to Theorem 4.1, the set of all slopes $\widehat{\mathbb{Q}}$ is divided into (the restriction to $\widehat{\mathbb{Q}}$ of) two disconnected intervals $\mathcal{L}_M^\circ$ and $(\mathcal{L}_M^\circ)^c$. The latter is precisely the set of *NLS detected slopes* in the sense of Boyer and Clay [2, Definition 7.16], according to Theorem 7.5. (In particular, this observation should be compared with [2, Theorem 8.1].) Thus the desired equivalence follows from Theorem 1.3, on comparison with [2, Theorem 1.7] restricted to rational homology solid tori. □

References

[1] **M Boileau**, **S Boyer**, *Graph manifold $\mathbb{Z}$–homology 3–spheres and taut foliations*, J. Topol. 8 (2015) 571–585 MR Zbl

[2] **S Boyer**, **A Clay**, *Foliations, orders, representations, L-spaces and graph manifolds*, Adv. Math. 310 (2017) 159–234 MR Zbl

[3] **S Boyer**, **C M Gordon**, **L Watson**, *On L-spaces and left-orderable fundamental groups*, Math. Ann. 356 (2013) 1213–1245 MR Zbl

[4] **S Boyer**, **D Rolfsen**, **B Wiest**, *Orderable 3–manifold groups*, Ann. Inst. Fourier (Grenoble) 55 (2005) 243–288 MR Zbl

[5] **A Clay**, **T Lidman**, **L Watson**, *Graph manifolds, left-orderability and amalgamation*, Algebr. Geom. Topol. 13 (2013) 2347–2368 MR Zbl

[6] **N M Dunfield**, *Floer homology, group orderability, and taut foliations of hyperbolic 3–manifolds*, Geom. Topol. 24 (2020) 2075–2125 MR Zbl

[7] **F Haiden**, **L Katzarkov**, **M Kontsevich**, *Flat surfaces and stability structures*, Publ. Math. Inst. Hautes Études Sci. 126 (2017) 247–318 MR Zbl

[8] **J Hanselman**, *Bordered Heegaard Floer homology and graph manifolds*, Algebr. Geom. Topol. 16 (2016) 3103–3166 MR Zbl

[9] **J Hanselman**, *Splicing integer framed knot complements and bordered Heegaard Floer homology*, Quantum Topol. 8 (2017) 715–748 MR Zbl

[10] **J Hanselman**, **J Rasmussen**, **S D Rasmussen**, **L Watson**, *L-spaces, taut foliations, and graph manifolds*, Compos. Math. 156 (2020) 604–612 MR Zbl

[11] **J Hanselman**, **J Rasmussen**, **L Watson**, *Bordered Floer homology for manifolds with torus boundary via immersed curves*, preprint (2016) arXiv 1604.03466 To appear in J. Amer. Math. Soc.

[12] **J Hanselman**, **J Rasmussen**, **L Watson**, *Heegaard Floer homology for manifolds with torus boundary: properties and examples*, Proc. Lond. Math. Soc. (3) 125 (2022) 879–967 MR

[13] **J Hanselman**, **L Watson**, *Cabling in terms of immersed curves*, Geom. Topol. 27 (2023) 925–952

[14] **M Hedden**, **A S Levine**, *Splicing knot complements and bordered Floer homology*, J. Reine Angew. Math. 720 (2016) 129–154 MR Zbl

[15] **J Hom**, **T Lidman**, **L Watson**, *The Alexander module, Seifert forms, and categorification*, J. Topol. 10 (2017) 22–100 MR Zbl

[16] **A Kotelskiy**, **L Watson**, **C Zibrowius**, *Immersed curves in Khovanov homology*, preprint (2019) arXiv 1910.14584

[17] **A S Levine**, *Knot doubling operators and bordered Heegaard Floer homology*, J. Topol. 5 (2012) 651–712 MR Zbl

[18] **T Li**, **R Roberts**, *Taut foliations in knot complements*, Pacific J. Math. 269 (2014) 149–168 MR Zbl

[19] **R Lipshitz**, **P S Ozsváth**, **D P Thurston**, *Bimodules in bordered Heegaard Floer homology*, Geom. Topol. 19 (2015) 525–724 MR Zbl

[20] **R Lipshitz**, **P S Ozsvath**, **D P Thurston**, *Bordered Heegaard Floer homology*, Mem. Amer. Math. Soc. 1216, Amer. Math. Soc., Providence, RI (2018) MR Zbl

[21] **M Mauricio**, *On lattice cohomology and left-orderability*, preprint (2013) arXiv 1308.1890

[22] **W D Neumann**, *A calculus for plumbing applied to the topology of complex surface singularities and degenerating complex curves*, Trans. Amer. Math. Soc. 268 (1981) 299–344 MR Zbl

[23] **P Ozsváth**, **Z Szabó**, *On the Floer homology of plumbed three-manifolds*, Geom. Topol. 7 (2003) 185–224 MR Zbl

[24] **P Ozsváth**, **Z Szabó**, *Holomorphic disks and three-manifold invariants: properties and applications*, Ann. of Math. 159 (2004) 1159–1245 MR Zbl

[25] **I Petkova**, *Cables of thin knots and bordered Heegaard Floer homology*, Quantum Topol. 4 (2013) 377–409 MR Zbl

[26] **I Petkova**, *The decategorification of bordered Heegaard Floer homology*, J. Symplectic Geom. 16 (2018) 227–277 MR Zbl

[27] **J Rasmussen**, **S D Rasmussen**, *Floer simple manifolds and L-space intervals*, Adv. Math. 322 (2017) 738–805 MR Zbl

[28] **S D Rasmussen**, *L-space intervals for graph manifolds and cables*, Compos. Math. 153 (2017) 1008–1049 MR Zbl

[29] **L Watson**, *Surgery obstructions from Khovanov homology*, Selecta Math. 18 (2012) 417–472 MR Zbl

Department of Mathematics, University of Texas at Austin
Austin, TX, United States

Current address: *Department of Mathematics, Princeton University*
Princeton, NJ, United States

School of Mathematics and Statistics, University of Glasgow
Glasgow, United Kingdom

Current address: *Department of Mathematics, University of British Columbia*
Vancouver BC, Canada

jh66@math.princeton.edu, liam@math.ubc.ca

Proposed: Peter Ozsváth Received: 24 September 2015
Seconded: Ciprian Manolescu, Cameron Gordon Revised: 10 July 2020

GEOMETRY & TOPOLOGY

msp.org/gt

See inside back cover or msp.org/gt for submission instructions.

The subscription price for 2023 is US \$740/year for the electronic version, and \$1030/year (+\$70, if shipping outside the US) for print and electronic. Subscriptions, requests for back issues and changes of subscriber address should be sent to MSP. Geometry & Topology is indexed by Mathematical Reviews, Zentralblatt MATH, Current Mathematical Publications and the Science Citation Index.

Geometry & Topology (ISSN 1465-3060 printed, 1364-0380 electronic) is published 9 times per year and continuously online, by Mathematical Sciences Publishers, c/o Department of Mathematics, University of California, 798 Evans Hall #3840, Berkeley, CA 94720-3840. Periodical rate postage paid at Oakland, CA 94615-9651, and additional mailing offices. POSTMASTER: send address changes to Mathematical Sciences Publishers, c/o Department of Mathematics, University of California, 798 Evans Hall #3840, Berkeley, CA 94720-3840.

GT peer review and production are managed by EditFLOW® from MSP.

PUBLISHED BY

mathematical sciences publishers

nonprofit scientific publishing

http://msp.org/

GEOMETRY & TOPOLOGY

Volume 27 Issue 3 (pages 823–1272) 2023