Geometry & Topology 11 (2007) 2339–2412 **2339**

On combinatorial link Floer homology

CIPRIAN MANOLESCU
PETER OZSVÁTH
ZOLTÁN SZABÓ
DYLAN THURSTON

Link Floer homology is an invariant for links defined using a suitable version of Lagrangian Floer homology. In an earlier paper, this invariant was given a combinatorial description with mod 2 coefficients. In the present paper, we give a self-contained presentation of the basic properties of link Floer homology, including an elementary proof of its invariance. We also fix signs for the differentials, so that the theory is defined with integer coefficients.

57R58, 57M25

1 Introduction

Heegaard Floer homology [12] is an invariant for three-manifolds, defined using holomorphic disks and Heegaard diagrams. Ozsváth and Szabó [11] and Rasmussen [15] extended this construction to give an invariant, *knot Floer homology*, for null-homologous knots in a closed, oriented three-manifold. This construction is further generalized by Ozsváth and Szabó [9] to the case of oriented links. The definition of all these invariants involves counts of holomorphic disks in the symmetric product of a Riemann surface, which makes them rather challenging to calculate.

More recently, Sucharit Sarkar discovered a principle which ensures that for Heegaard diagrams with a certain property, the counts of holomorphic disks are combinatorial. Manolescu, Ozsváth and Sarkar [6] constructed Heegaard diagrams of the needed form from grid presentations of knots or links in S^3. This led to an explicit, combinatorial description of the knot or link Floer complex, taken with coefficients in $\mathbb{Z}/2\mathbb{Z}$, henceforth called $\mathbb{F}_2$. (See also Sarkar and Wang [16] for a different application of this principle.)

The purpose of the present paper is to develop knot (or link) Floer homology in purely elementary terms, starting from a grid presentation, and establish its topological invariance without appealing to the earlier theory. We also give a sign-refinement of this description, leading to a homology theory with coefficients in $\mathbb{Z}$.

Published: 17 December 2007 DOI: 10.2140/gt.2007.11.2339

We recall the chain complex from [6]; but first, we need to review some topological notions.

A *planar grid diagram* G lies on an $n \times n$ grid of squares in the plane. Each square is decorated either with an X, an O, or nothing. Moreover, the decorations are arranged so that:

- every row contains exactly one X and one O; and
- every column contains exactly one X and one O.

The number n is called the *grid number* of G. Sometimes we find it convenient to number the O's and X's by $\{O_i\}_{i=1}^n$ and $\{X_i\}_{i=1}^n$. We denote the set of all O's and X's by $\mathbb{O}$ and $\mathbb{X}$, respectively. As a point of comparison: the O_i correspond to the "white dots" of [6] and the w_i of [9], while the X_i to the "black dots" of [6] and the z_i of [9]. We find the current notation clearer for pictures.

Given a planar grid diagram G, we can place it in a standard position on the plane as follows: the bottom left corner is at the origin, and each cell is a square of edge length one. We then construct an oriented, planar link projection by drawing horizontal segments from the O's to the X's in each row, and vertical segments from the X's to the O's in each column. At every intersection point, we let the horizontal segment be the underpass and the vertical one the overpass. This produces a planar diagram for an oriented link $\vec{L}$ in S^3. We say that $\vec{L}$ has a grid presentation given by G. See Figure 1 for an example.

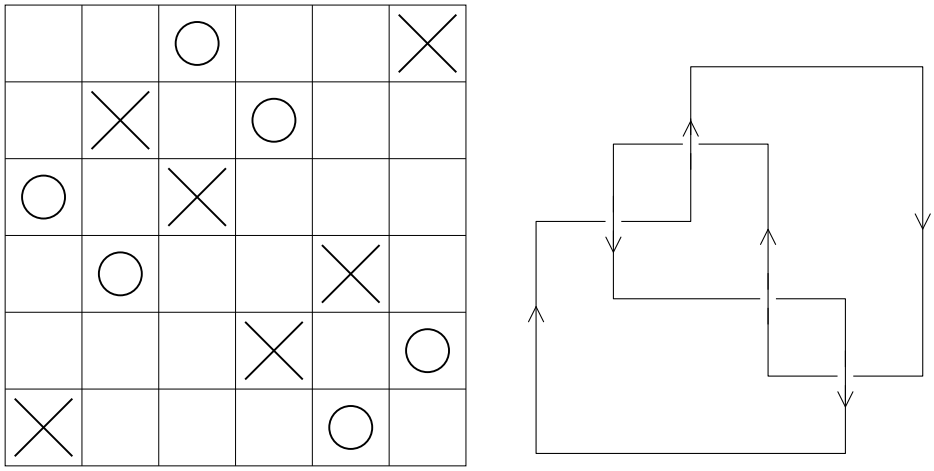

Figure 1: Grid presentation for the figure eight knot

We transfer our grid diagrams to the torus $\mathcal{T}$ obtained by gluing the topmost segment to the bottom-most one, and the leftmost segment to the rightmost one. In the torus, our horizontal and vertical arcs become horizontal and vertical circles. The torus inherits its orientation from the plane. We call the resulting object a *toroidal grid diagram*, or simply a grid diagram, for $\vec{L}$. We will again denote it by G.

Given a toroidal grid diagram, we associate to it a chain complex $(C^-(G), \partial^-)$ as follows. The set of generators of $C^-(G)$, denoted $\mathbf{S}$ or $\mathbf{S}(G)$, consists of one-to-one correspondences between the horizontal and vertical circles. More geometrically, we can think of the generators as n–tuples of intersection points between the horizontal and vertical circles, with the property that no intersection point appears on more than one horizontal (or vertical) circle.

Before defining the differentials, we turn to a grading and a filtration on the complex, determined by two functions $M\colon \mathbf{S} \longrightarrow \mathbb{Z}$ and $A\colon \mathbf{S} \longrightarrow (\frac{1}{2}\mathbb{Z})^\ell$.

The function M is defined as follows. Given two collections A, B of finitely many points in the plane, let $\mathcal{I}(A, B)$ be the number of pairs $(a_1, a_2) \in A$ and $(b_1, b_2) \in B$ with $a_1 < b_1$ and $a_2 < b_2$. Let $\mathcal{J}(A, B) = (\mathcal{I}(A, B) + \mathcal{I}(B, A))/2$. Take a fundamental domain $[0, n) \times [0, n)$ for the torus, cut along a horizontal and vertical circle, with the left and bottom edges included. Given a generator $\mathbf{x} \in \mathbf{S}$, we view $\mathbf{x}$ as a collection of points with integer coordinates in this fundamental domain. Similarly, we view $\mathbb{O} = \{O_i\}_{i=1}^n$ as a collection of points in the plane with half-integer coordinates. Define

$$M(\mathbf{x}) = \mathcal{J}(\mathbf{x}, \mathbf{x}) - 2\mathcal{J}(\mathbf{x}, \mathbb{O}) + \mathcal{J}(\mathbb{O}, \mathbb{O}) + 1.$$

We find it convenient to write this formula more succinctly as

$$M(\mathbf{x}) = \mathcal{J}(\mathbf{x} - \mathbb{O}, \mathbf{x} - \mathbb{O}) + 1, \tag{1}$$

where we extend $\mathcal{J}$ bilinearly over formal sums (or differences) of subsets. Note that the definition of M appears to depend on which circles we cut along to create a fundamental domain. In fact, it does not (see Lemma 2.4 below). Note also that this definition of the Maslov grading is not identical with that given in [6], but it is not difficult to see they agree. See Lemma 2.5 below, and the remarks following it.

For an ℓ–component link, we define an ℓ–tuple of Alexander gradings $A(\mathbf{x}) = (A_1(\mathbf{x}), \dots, A_\ell(\mathbf{x}))$ by the formula

$$A_i(\mathbf{x}) = \mathcal{J}(\mathbf{x} - \tfrac{1}{2}(\mathbb{X} + \mathbb{O}), \mathbb{X}_i - \mathbb{O}_i) - \Big(\frac{n_i - 1}{2}\Big), \tag{2}$$

where here $\mathbb{O}_i \subset \mathbb{O}$ is the subset corresponding to the ith component of the link, $\mathbb{X}_i \subset \mathbb{X}$ is the set of X's belonging to the ith component of the link, and where we

once again use the bilinear extension of $\mathcal{J}$. For links, the A_i may take half-integral values. Again, this quantity is independent of how the torus is cut up to form a planar rectangle (see Lemma 2.6 below).

Given a pair of generators $\mathbf{x}$ and $\mathbf{y}$, and an embedded rectangle r in $\mathcal{T}$ whose edges are arcs in the horizontal and vertical circles, we say that r *connects* $\mathbf{x}$ *to* $\mathbf{y}$ if $\mathbf{x}$ and $\mathbf{y}$ agree along all but two horizontal circles, if all four corners of r are intersection points in $\mathbf{x}\cup\mathbf{y}$, and if, as we traverse each horizontal boundary component of r in the direction dictated by the orientation that r inherits from $\mathcal{T}$, then the arc is oriented from a point in $\mathbf{x}$ to the point in $\mathbf{y}$. (See Figure 2 for an example.) Let $\mathrm{Rect}(\mathbf{x},\mathbf{y})$ denote the collection of rectangles connecting $\mathbf{x}$ to $\mathbf{y}$. If $\mathbf{x},\mathbf{y}\in\mathbf{S}$ agree along all but two horizontal circles, then there are exactly two rectangles in $\mathrm{Rect}(\mathbf{x},\mathbf{y})$; otherwise $\mathrm{Rect}(\mathbf{x},\mathbf{y})=\varnothing$. Let $\mathrm{Int}(r)$ denote the interior of the subset of $\mathcal{T}$ determined by r. A rectangle $r\in\mathrm{Rect}(\mathbf{x},\mathbf{y})$ is said to be *empty* if $\mathrm{Int}(r)\cap\mathbf{x}=\varnothing$, or equivalently if $\mathrm{Int}(r)\cap\mathbf{y}=\varnothing$. The space of empty rectangles connecting $\mathbf{x}$ and $\mathbf{y}$ is denoted $\mathrm{Rect}^{\circ}(\mathbf{x},\mathbf{y})$.

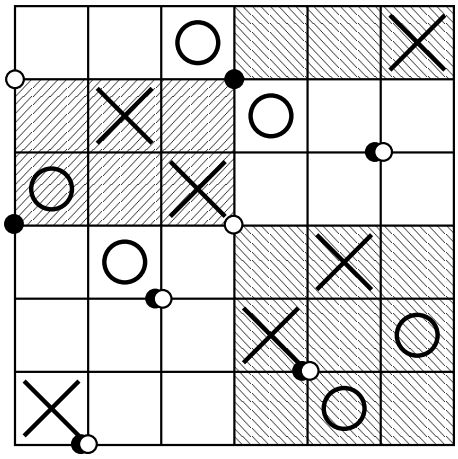

Figure 2: Rectangles. The small dark circles describe the generator $\mathbf{x}$ and the hollow ones describe $\mathbf{y}$. There are two rectangles in $\mathrm{Rect}(\mathbf{x},\mathbf{y})$, shown here shaded by two types of diagonal hatchings. The rectangle on the left is in $\mathrm{Rect}^{\circ}(\mathbf{x},\mathbf{y})$ while the other one is not, because it contains a dark circle in its interior.

Let R denote the polynomial algebra over $\mathbb{F}_2$ generated by variables which are in one-to-one correspondence between the elements of $\mathbb{O}$, and which we denote $\{U_i\}_{i=1}^{n}$. We think of this ring as endowed with a *Maslov grading*, defined so that the constant terms are in Maslov grading zero, and U_i are in grading -2. The ring is also endowed

with an *Alexander multifiltration*, defined so that constant terms are in filtration level zero, while the variables U_j corresponding to the i th component of the link drop the i th multifiltration level by one and preserve all others.

Let $C^-(G)$ be the free R–module with generating set $\mathbf{S}$.

We endow this module with an endomorphism $\partial^-\colon C^-(G) \longrightarrow C^-(G)$ defined by

$$\partial^-(\mathbf{x}) = \sum_{\mathbf{y}\in\mathbf{S}} \sum_{r\in \mathrm{Rect}^\circ(\mathbf{x},\mathbf{y})} U_1^{O_1(r)} \cdots U_n^{O_n(r)} \cdot \mathbf{y}, \tag{3}$$

where $O_i(r)$ denotes the number of times O_i appears in the interior of r (so $O_i(r)$ is either 0 or 1).

The results of [6] can be summarized by the following:

Theorem 1.1 (Manolescu–Ozsváth–Sarkar) *The data $(C^-(G), \partial^-)$ is a chain complex for the Heegaard-Floer homology $CF^-(S^3)$, with grading induced by M, and the filtration induced by A coincides with the link filtration of $CF^-(S^3)$.*[1]

In particular, appealing to the earlier theorem defined using holomorphic disks [9; 11; 15], the filtered quasi-isomorphism type of this chain complex C^- is a link invariant. Other knot and link invariants can be found by routine algebraic manipulations of C^- as well (for example, by taking the homology of the associated graded object).

Our main goal here is to prove the topological invariance of the filtered quasi-isomorphism type of the resulting chain complex $C^-(G)$, without resorting to any of the holomorphic disk theory, and in particular without resorting to Theorem 1.1. We prove the following:

Theorem 1.2 *Let $\vec{L}$ be an oriented, ℓ–component link. Number the elements of $\mathbb{O} = \{O_i\}_{i=1}^n$ so that $O_1, \dots, O_\ell$ correspond to different components of the link. Then the filtered quasi-isomorphism type of the complex $(C^-(G), \partial^-)$ over $\mathbb{Z}[U_1, \dots, U_\ell]$ is an invariant of the link.*

We also give independent verification of the basic algebraic properties of $C^-(G)$ which, with $\mathbb{F}_2$ (ie $\mathbb{Z}/2\mathbb{Z}$) coefficients, follow from Theorem 1.1, together with properties of the "Heegaard Floer homology package". Note that for technical reasons, for links with more than one component the chain complex in [9] was originally defined only with coefficients in $\mathbb{F}_2$.

[1] The reader should be warned: our conventions here on the Maslov grading are such that the total homology $H_*(CF^-(S^3))$ is isomorphic to a copy of the polynomial algebra in U, where the constants have grading equal to zero. In [12], the convention is that the constants have grading equal to -2.

There are some related constructions one could consider. In one of these, we set $U_1 = \cdots = U_\ell = 0$, and let $\widehat{C}(G)$ denote the resulting chain complex, equipped with its Alexander filtration. Taking the homology of the associated graded object, we get a group whose multigraded Euler characteristic is the multivariable Alexander polynomial of $\vec{L}$, times a suitable normalization factor (this is proved in Equation (1) of [9]; see also Theorem 6.1 below).

We have endeavored to separate the discussion of signs from the rest of the body of the paper, to underscore the simplicity of the $\mathbb{F}_2$ version which is sufficient for the knot-theoretic applications, and also simpler to calculate. In particular, in Section 3, we establish Theorem 1.2, working over coefficients in $\mathbb{F}_2$, where it could alternately be seen as an immediate consequence of Theorem 1.1. We hope, however, that the present combinatorial proof of invariance has value in its own simplicity; see also Ozsváth, Thurston and Szabó [13] for another application. The sign-refinements are dealt with in Section 4.

This paper is organized as follows. The algebraic properties are established in Section 2 and topological invariance with coefficients in $\mathbb{F}_2$ is established in Section 3. In Section 5, we describe some further properties of C^-. In Section 4, we describe the sign conventions and the modifications needed for the earlier discussion to establish Theorem 1.2 over $\mathbb{Z}$. Finally, in Section 6, we show that the Euler characteristic of the homology is the Alexander polynomial.

Acknowledgements We would like to thank Dror Bar-Natan, Sergei Duzhin, Sergey Fomin, John Morgan, and Sucharit Sarkar for helpful conversations.

CM was supported by a Clay Research Fellowship. PSO was supported by NSF grant number DMS-0505811 and FRG-0244663. ZSz was supported by NSF grant number DMS-0406155 and FRG-0244663. DPT was supported by a Sloan Research Fellowship.

2 Properties of the chain complex $C^-(G)$

2.1 Algebraic terminology

We recall some standard terminology from homological algebra.

For simplicity, we use coefficients in $\mathbb{F}_2 = \mathbb{Z}/2\mathbb{Z}$ for this section, and also the next two. The definitions from algebra can be made with $\mathbb{Z}$ coefficients with little change. Other aspects of $\mathbb{Z}$ coefficients will be handled in Section 4. (And in fact, the choices of signs in the formulas below which, of course, are immaterial over $\mathbb{F}_2$, have been chosen so as to work over $\mathbb{Z}$.)

Definition 2.1 We give $\mathbb{Q}^\ell$ its usual partial ordering, $(a_1,\dots,a_\ell) \le (b_1,\dots,b_\ell)$ if for all $i = 1,\dots,\ell$, $a_i \le b_i$. Let R be the ring $\mathbb{F}_2[U_1,\dots,U_n]$. A function $g\colon \{1,\dots,n\} \longrightarrow (\mathbb{Q}^{\ge 0})^\ell$ specifies a $\mathbb{Q}^\ell$ grading on R. Fix a grading on R. Let M be a module over R. A *$\mathbb{Q}^\ell$–filtration* on a module M is a collection of R–submodules $\{\mathcal{F}_s(M)\}_{s\in\mathbb{Q}^\ell}$ of M satisfying the following properties:

- $\mathcal{F}_s(M) \subset \mathcal{F}_t(M)$ if $s \le t$.
- Multiplication by U_i sends $\mathcal{F}_s(M)$ into $\mathcal{F}_{s-g(i)}(M)$.
- For all sufficiently large s (with respect to $\le$), $\mathcal{F}_s(M) = M$.

A *filtered R–module map* $\phi\colon M \longrightarrow N$ is an R–module map which carries $\mathcal{F}_s(M)$ into $\mathcal{F}_s(N)$. A *filtered chain complex* (C,∂) is a graded and filtered R–module, equipped with a filtered endomorphism ∂ which drops grading by one. Given filtered chain complexes A and B, a *filtered chain map* is a chain map $\phi\colon A \longrightarrow B$ which is a grading-preserving, filtered R–module map. Given two filtered chain maps $\phi_i\colon A \longrightarrow B$ for $i = 1,2$, a *filtered chain homotopy* is a filtered R–module map $H\colon A \longrightarrow B$ which raises grading by one and satisfies the formula

$$\partial_B \circ H + H \circ \partial_A = \phi_1 - \phi_2.$$

If a filtered chain homotopy exists between ϕ_1 and ϕ_2, then we say that ϕ_1 and ϕ_2 are *filtered chain homotopic*. Let $\phi\colon A \longrightarrow B$ be a filtered chain map. We say that ϕ is a *filtered chain homotopy equivalence* if there is a map $\psi\colon B \longrightarrow A$ with the property that $\phi\circ\psi$ and $\psi\circ\phi$ are filtered chain homotopic to the identity maps. A *filtered quasi-isomorphism* is a filtered map $\phi\colon A \longrightarrow B$ which induces an isomorphism from the homology groups $H_*(\mathcal{F}_s(A))$ to $H_*(\mathcal{F}_s(B))$. The *associated graded object* of a filtered chain complex C is the $\mathbb{Q}^\ell$–graded chain complex

$$\mathrm{gr}(C) = \bigoplus_{s\in\mathbb{Q}^\ell} \mathrm{gr}_s(C),$$

where $\mathrm{gr}_s(C)$ is the quotient of $\mathcal{F}_s(C)$ by the submodule generated by $\mathcal{F}_t(C)$ for all $t < s$, endowed with the differential induced from ∂.

A filtered chain homotopy equivalence is a filtered quasi-isomorphism. Moreover a map is a filtered quasi-isomorphism if and only if it induces an isomorphism on the homology of the associated graded object.

Definition 2.2 Given a filtered chain map $\phi\colon A \longrightarrow B$, we can form a new filtered chain complex, the *mapping cone* $M(\phi)$ whose underlying module is $A \oplus B$, and which is endowed with the differential $D(a,b) = (\partial a, \phi(a) - \partial b)$, where here ∂a and ∂b denotes the differentials of a and b within A and B, respectively.

The mapping cone fits into a short exact sequence of chain complexes (where the maps are all filtered chain maps)

$$0 \longrightarrow B \longrightarrow M(\phi) \longrightarrow A \longrightarrow 0,$$

and whose connecting homomorphism agrees with the map induced by ϕ.

Definition 2.3 Two filtered chain complexes A and B are *quasi-isomorphic* if there is a third filtered chain complex C and filtered quasi-isomorphisms from C to A and to B.

If $\phi_1\colon A \longrightarrow B$ and $\phi_2\colon A \longrightarrow B$ are chain homotopic, then their induced mapping cones are quasi-isomorphic.

Our chain complexes will always be finitely generated over $\mathbb{F}_2[U_1, \dots, U_n]$.

2.2 The chain complex C^-

We verify that $C^-(G)$ as defined in the introduction (using coefficients in $\mathbb{F}_2$) is a filtered chain complex in the above sense, with (Alexander) filtration induced from the function A and (Maslov) grading induced from the function M.

Lemma 2.4 *The function M is well-defined, ie it is independent of the manner in which a given generator $\mathbf{x} \in \mathbf{S}$ is drawn on the square.*

Proof Fix $\mathbf{x} \in \mathbf{S}$, thought of as drawn in the usual fundamental domain with the bottom and left edges included, so there is one component a with coordinates $(m, 0)$. Let $\mathbf{x}'$ denote the same generator in the fundamental domain with the top and left edges included, so there is now a component b with coordinates (m, n). For each i with $0 \le i < n, i \neq m$, there is one component c_i in $\mathbf{x}$ and $\mathbf{x}'$ with first coordinate i. For $m < i < n$, the pair (a, c_i) contributes 1 to the count of $\mathcal{J}(\mathbf{x}, \mathbf{x})$, whereas the corresponding pair (c_i, b) does not contribute to $\mathcal{J}(\mathbf{x}', \mathbf{x}')$. Symmetrically, for each i with $0 \le i < m$, the pair (c_i, a) does not contribute to $\mathcal{J}(\mathbf{x}, \mathbf{x})$, whereas (c_i, b) does contribute to $\mathcal{J}(\mathbf{x}', \mathbf{x}')$. It follows that $\mathcal{J}(\mathbf{x}, \mathbf{x}) + m = \mathcal{J}(\mathbf{x}', \mathbf{x}') + n - m - 1$. We can similarly analyze $\mathcal{J}(\mathbf{x}', \mathbb{O})$ to find

$$\mathcal{J}(\mathbf{x}', \mathbf{x}') = \mathcal{J}(\mathbf{x}, \mathbf{x}) + 2m - n + 1$$
$$2\mathcal{J}(\mathbf{x}', \mathbb{O}) = 2\mathcal{J}(\mathbf{x}, \mathbb{O}) + 2m - n.$$

In particular $M_{\mathbb{O}}(\mathbf{x}') = M_{\mathbb{O}}(\mathbf{x}) + 1$.

To complete the rotation, we have to change $\mathbb{O}$ to $\mathbb{O}'$ by moving the O in the bottom row, with coordinates $(l-\frac{1}{2},\frac{1}{2})$, to $(l-\frac{1}{2},n+\frac{1}{2})$. A similar analysis yields

$$2\mathcal{J}(\mathbf{x}',\mathbb{O}') = 2\mathcal{J}(\mathbf{x}',\mathbb{O}) + 2l - n$$
$$\mathcal{J}(\mathbb{O}',\mathbb{O}') = \mathcal{J}(\mathbb{O},\mathbb{O}) + 2l - n - 1.$$

Thus $M_{\mathbb{O}'}(\mathbf{x}') = M_{\mathbb{O}}(\mathbf{x}') - 1 = M_{\mathbb{O}}(\mathbf{x})$, which is the desired cyclic invariance.

The same reasoning also establishes invariance under horizontal rotation. □

The Maslov grading on R and the generating set $\mathbf{S}$ induces a Maslov grading on the chain complex C^-. Explicitly, the summand $C^-_d(G)$ is generated by expressions $U_1^{m_1}\cdots U_n^{m_n}\cdot\mathbf{x}$, with $\mathbf{x}\in\mathbf{S}$, where

$$d = M(\mathbf{x}) - 2\sum_{i=1}^{n} m_i.$$

Lemma 2.5 *Suppose that* $\mathbf{x},\mathbf{y}\in\mathbf{S}$, *and* $r\in\mathrm{Rect}(\mathbf{x},\mathbf{y})$ *is a rectangle with* $\mathbf{x}\cap\mathrm{Int}(r)=\varnothing$. *Then*

$$M(\mathbf{x}) = M(\mathbf{y}) + 1 - 2\sum_{i=1}^{n} O_i(r). \tag{4}$$

Proof Draw the torus $\mathcal{T}$ on a square in such a manner that the lower left corner of r coincides with the lower left corner of the square. Then it is clear that $\mathcal{J}(\mathbf{x},\mathbf{x}) = \mathcal{J}(\mathbf{y},\mathbf{y})+1$ (since the two new coordinates y_1 and y_2 in $\mathbf{y}$ are the only pair counted in $\mathcal{J}(\mathbf{x})$ which are not also counted in $\mathcal{J}(\mathbf{x})$), while $\mathcal{J}(\mathbb{O},\mathbf{x}) = \mathcal{J}(\mathbb{O},\mathbf{y}) + \#\{\mathbb{O}\cap r\}$, since each $O_i\in r$ gives rise to exactly one pair (x_1,O_i) counted in $\mathcal{J}(\mathbb{O},\mathbf{x})$ which is not also counted in $\mathcal{J}(\mathbb{O},\mathbf{x})$. Similarly, $\mathcal{J}(\mathbf{x},\mathbb{O}) = \mathcal{J}(\mathbf{y},\mathbb{O}) + \#\{\mathbb{O}\cap r\}$. Equation (4) now follows when M is calculated with respect to a particular manner of lifting the data on $\mathcal{T}$ to data on a square. But according to Lemma 2.4, the Maslov grading is independent of this data. □

The alert reader might notice that the definition of Maslov grading we give here does not identically agree with that given in [6], which we denote by M'. However, by connecting any two generators $\mathbf{x}\in\mathbf{S}$ by a sequence of rectangles satisfying Lemma 2.5 (the existence of which can be deduced from the fact that the symmetric group is generated by transpositions), we see at once that M is uniquely characterized, up to an additive constant, by Equation (4), which is also satisfied by M'. It now remains to show that $M(\mathbf{x}_0) = M'(\mathbf{x}_0)$ for some $\mathbf{x}_0\in\mathbf{S}$. To this end, we take $\mathbf{x}_0$ to be the generator for which x_i is on the lower left corner of the square marked with O_i.

According to the conventions from [6], $M'(\mathbf{x}_0) = 1-n$; it is easy to verify that $M(\mathbf{x}_0) = 1-n$, as well.

For the Alexander gradings, we have the following analogue of Lemma 2.4:

Lemma 2.6 *For a given link component i, the function A_i is well-defined, ie it is independent of the manner in which a given generator $\mathbf{x} \in \mathbf{S}$ is drawn on the square.*

Proof For a point $p \in \mathbb{Z}^2$, the quantities $\mathcal{I}(p, \mathbb{X}_i - \mathbb{O}_i)$ and $\mathcal{I}(\mathbb{X}_i - \mathbb{O}_i, p)$ both compute the winding number of the ith component of the knot around the point p. This quantity is unchanged if p is moved from the very bottom to the very top of the diagram (since in that case the winding number is 0), and if $\mathbb{X}_i$ and $\mathbb{O}_i$ are rotated vertically once, it changes by ± 1 if p is in between the X and the O that are moved, and is unchanged otherwise. For a point p with half-integer coordinates, the inequalities used in the definition of $\mathcal{I}(p, \mathbb{X}_i - \mathbb{O}_i)$ effectively shift p up and to the right by $(\frac{1}{2}, \frac{1}{2})$ before computing the winding number. Similarly, $\mathcal{I}(\mathbb{X}_i - \mathbb{O}_i, p)$ computes the winding number around $p - (\frac{1}{2}, \frac{1}{2})$. Therefore $A_i(\mathbf{x})$, defined as $\mathcal{J}(\mathbf{x} - \frac{1}{2}(\mathbb{X} + \mathbb{O}), \mathbb{X}_i - \mathbb{O}_i)$, computes the winding number of the ith component around a weighted sum of points which has total weight 0 in each row and column. This combination is therefore invariant under cyclic rotation of the whole diagram. □

The function $A\colon \mathbf{S}(G) \longrightarrow (\frac{1}{2}\mathbb{Z})^\ell \subset \mathbb{Q}^\ell$ endows $C^-(G)$ with a $\mathbb{Q}^\ell$–filtration in the sense of Definition 2.1, for the function $g\colon \{1, \dots, n\} \longrightarrow \mathbb{Z}^\ell$ which associates to i the jth standard basis vector in $\mathbb{Z}^\ell$ if O_i belongs to the jth component of the link. The element $(U_1^{m_1} \cdots U_n^{m_n})\mathbf{x}$ has filtration level $a = (a_1, \dots, a_\ell)$, where

$$a = A(\mathbf{x}) - \sum_{i=1}^{n} m_i \cdot g(i).$$

It is sometimes useful to consider objects more general than rectangles, called domains. To define them, let us view the torus $\mathcal{T}$ as a two-dimensional cell complex, with the toroidal grid diagram inducing the cell decomposition with n^2 zero-cells, $2n^2$ one-cells and n^2 two-cells (the little squares). Let U_α be the one-dimensional subcomplex of $\mathcal{T}$ consisting of the union of the n horizontal circles.

Definition 2.7 Given $\mathbf{x}, \mathbf{y} \in \mathbf{S}$, a *path* from $\mathbf{x}$ to $\mathbf{y}$ is a 1–cycle γ on the cell complex $\mathcal{T}$, such that the boundary of the intersection of γ with U_α is $\mathbf{y} - \mathbf{x}$.

Definition 2.8 A *domain* p from $\mathbf{x}$ to $\mathbf{y}$ is a two-chain in $\mathcal{T}$ whose boundary ∂p is a path from $\mathbf{x}$ to $\mathbf{y}$. The *support* of p is the union of the closures of the two-cells appearing (with nonzero multiplicity) in the two-chain p.

Given $\mathbf{x}, \mathbf{y} \in \mathbf{S}$, let $\pi(\mathbf{x}, \mathbf{y})$ denote the space of domains from $\mathbf{x}$ to $\mathbf{y}$. There is a natural composition law

$$*\colon \pi(\mathbf{a}, \mathbf{b}) \times \pi(\mathbf{b}, \mathbf{c}) \longrightarrow \pi(\mathbf{a}, \mathbf{c}).$$

For a domain $p \in \pi(\mathbf{x}, \mathbf{y})$, we let $X_i(p)$ and $O_i(p)$ denote the multiplicity with which X_i and O_i, respectively, appear in p.

Proposition 2.9 *The differential ∂^- drops Maslov grading by one, and respects the Alexander filtration. Specifically, if $\mathbf{x} \in \mathbf{S}$ has $M(\mathbf{x}) = d$, then $\partial^-(\mathbf{x})$ is written as a sum of elements in Maslov grading $d-1$. Also, if $A(\mathbf{x}) = a$, then $\partial^-(\mathbf{x})$ is a sum of elements with Alexander filtrations $\leq a$.*

Proof The fact that ∂^- drops Maslov grading by one follows at once from Equation (4), together with the definition of ∂^-.

The fact that ∂^- respects the Alexander filtration follows from basic properties of winding numbers. Specifically, given $\mathbf{x}, \mathbf{y} \in \mathbf{S}$ and $r \in \mathrm{Rect}(\mathbf{x}, \mathbf{y})$, it is easy to see that

$$A(\mathbf{x}) - A(\mathbf{y}) = \sum_i (X_i(r) - O_i(r)) \cdot g(i).$$

Thus if $U_1^{m_1} \cdots U_n^{m_n} \cdot \mathbf{y}$ appears with nonzero coefficient in $\partial^-(\mathbf{x})$, then the Alexander filtration level of the corresponding term is smaller than the Alexander filtration level of $\mathbf{x}$ by $\sum_{i=1}^n X_i(r) \cdot g(i)$. □

With the terminology in place, we now verify that ∂^- is the differential of a chain complex.

Proposition 2.10 *The endomorphism ∂^- of $C^-(G)$ is a differential, ie $\partial^- \circ \partial^- = 0$.*

Proof Consider an element $\mathbf{x} \in \mathbf{S}$, viewed as a generator of $C^-(G)$. We can view $\partial^- \circ \partial^-(\mathbf{x})$ as a count

$$\partial^- \circ \partial^-(\mathbf{x}) = \sum_{\mathbf{z} \in \mathbf{S}} \sum_{\substack{p \in \pi(\mathbf{x}, \mathbf{y}) \\ \mathbf{x} \notin \mathrm{Int}\, p}} N(p) \cdot U_1^{O_1(p)} \cdots U_n^{O_n(p)} \cdot \mathbf{z},$$

where here $N(p)$ denotes the number of ways of decomposing a domain as a composite of two empty rectangles $p = r_1 * r_2$, where $r_1 \in \mathrm{Rect}^\circ(\mathbf{x}, \mathbf{y})$ and $r_2 \in \mathrm{Rect}^\circ(\mathbf{y}, \mathbf{z})$ for some $\mathbf{y} \in \mathbf{S}$.

If $\mathbf{z} \neq \mathbf{x}$, and if p has a decomposition $p = r_1 * r_2$, then we claim that there is a unique alternate decomposition $p = r_1' * r_2'$, where here $r_1' \in \mathrm{Rect}(\mathbf{x}, \mathbf{y}')$ and $r_2' \in \mathrm{Rect}(\mathbf{y}', \mathbf{z})$. In fact, if $p = r_1 * r_2$ is a domain obtained from two empty rectangles r_1 and r_2, then we claim that there are three possibilities for p:

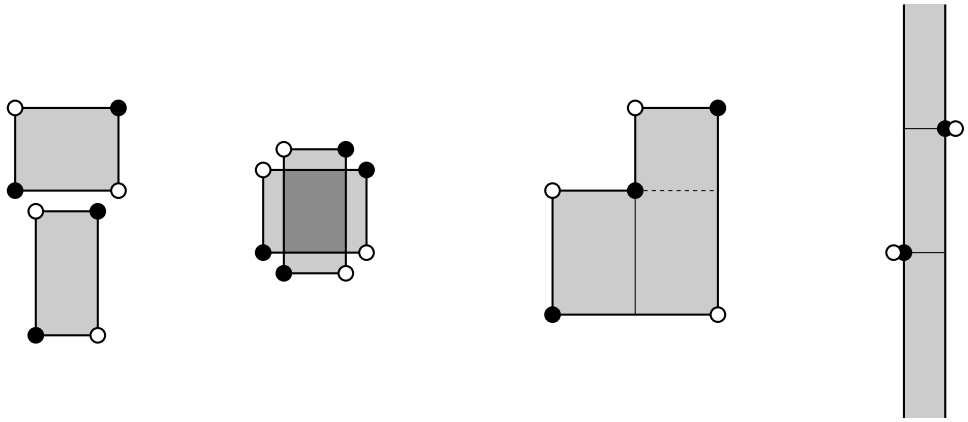

Figure 3: $\partial^- \circ \partial^- = 0$. The four combinatorially different ways the composite of two empty rectangles $r_1 * r_2$ can appear. The initial point is indicated by the dark circles, the final by the hollow ones.

- two disjoint rectangles;
- two rectangles with overlapping interiors (the darker region in Figure 3); and
- two rectangles which share a corner.

These three cases are illustrated in the first three diagrams in Figure 3. In each case, there are exactly two decompositions of the obtained domain as a juxtaposition of empty rectangles: in the first two cases by taking the rectangles in the two possible orders, and in third case by decomposing either along the thin or dotted lines, cf Figure 4. It follows at once that the $\mathbf{z}$ component of $\partial^- \circ \partial^-(\mathbf{x})$ vanishes for $\mathbf{z} \neq \mathbf{x}$.

When $\mathbf{z} = \mathbf{x}$, however, the only domains $p \in \pi(\mathbf{x}, \mathbf{x})$ which can be decomposed as a union of two empty rectangles are width one annuli, as in the fourth diagram in Figure 3, or height one annuli in the torus. There are $2n$ of these annuli. Each such annulus p has a unique decomposition $p = r_1 * r_2$ with $r_1 \in \mathrm{Rect}(\mathbf{x}, \mathbf{y})$ and $r_2 \in \mathrm{Rect}(\mathbf{y}, \mathbf{x})$ (for some uniquely specified $\mathbf{y}$). The row or column containing O_i contributes U_i in the formula for $\partial^- \circ \partial^-(\mathbf{x})$. Since O_i appears in exactly one row and exactly one column, it follows now that the $\mathbf{x}$ component of $\partial^- \circ \partial^-(\mathbf{x})$ vanishes, as well. □

The proof of the above proposition is elementary, depending on evident properties of rectangles in the torus. However, it does deserve a few extra words, since it is the starting point of this paper, and indeed a recurring theme throughout. Specifically, the alert reader will observe that the remarks concerning juxtapositions of pairs of rectangles is

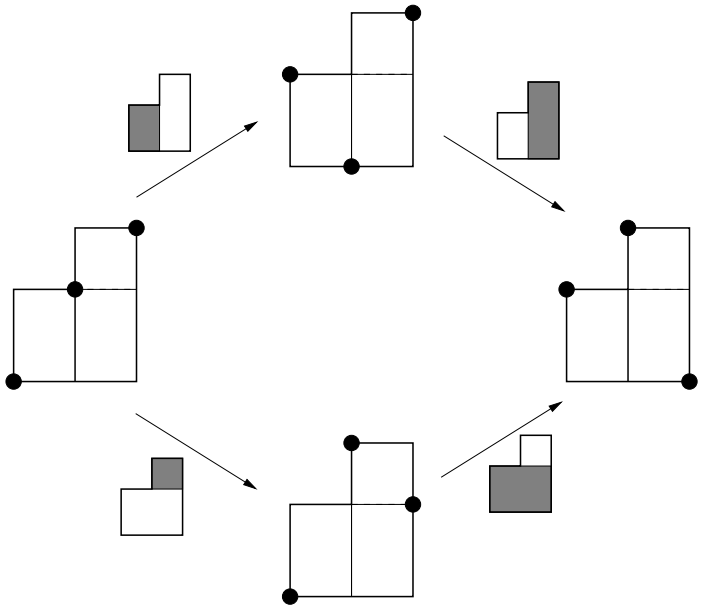

Figure 4: The third case of Figure 3. The three black dots are permuted to give four different generators. Each arrow represents a rectangle, which is shown shaded. There are two ways of connecting the initial generator $\mathbf{x}$ to the final generator $\mathbf{z}$: by following the top arrows, or the bottom ones. Each way gives a contribution to $\partial^- \circ \partial^-$, and in the final count these contributions cancel out.

one of the last vestiges of Gromov's compactness theorem, the foundation upon which Floer's theory of Lagrangian intersections is built [4] (and knot Floer homology can be viewed as a variant of that latter theory). The assertions about annuli can also be seen as remnants of Gromov's theory, as they are counting boundary degenerations.

In terms of combinatorics, we see a pattern that will be repeated throughout the paper: in order to prove an identity with differentials (e.g., that $(\partial^-)^2 = 0$, or that a map is a chain map) we consider the composites of two domains; generally the composite domain will have exactly two decompositions. In some cases we need to add or delete annuli of width or height one while taking care not to change the factors of U_i that appear.

2.3 Algebraic properties of C^-

We now turn to the basic algebraic properties of the chain complex.

In the following lemma, just as U_i is a chain map which drops filtration level by one, the filtered chain homotopy drops the filtration level by one.

Lemma 2.11 *Suppose that O_i and O_k correspond to the same component of $\vec{L}$. Then multiplication by U_i is filtered chain homotopic to multiplication by U_k.*

Proof Since filtered chain homotopies can be composed, it suffices to show that if O_i lies in the same row as some X_j which in turn is in the same column as O_k, then multiplication by U_i is filtered chain homotopic to multiplication by U_k. The filtered chain homotopy is furnished by counting rectangles which contain X_j.

Specifically, define

$$H\colon C^-(G) \longrightarrow C^-(G)$$

by the formula

$$H(\mathbf{x}) = \sum_{\mathbf{y}\in\mathbf{S}} \sum_{\substack{r\in \mathrm{Rect}^\circ(\mathbf{x},\mathbf{y}) \\ X_j \in r}} U_1^{O_1(r)} \cdots U_n^{O_n(r)} \cdot \mathbf{y}.$$

We claim that

$$\partial^- \circ H + H \circ \partial^- = U_i - U_k.$$

This follows from the same argument as Proposition 2.10: Most composite domains on the left hand side can be decomposed in exactly two ways. The exception are the horizontal and vertical annuli, necessarily containing X_j which contribute U_i and U_k, respectively. □

Proposition 2.12 *Suppose that the oriented link $\vec{L}$ has ℓ components. Choose an ordering of $\mathbb{O} = \{O_i\}_{i=1}^n$ so that for $i = 1, \dots, \ell$, O_i corresponds to the i th component of $\vec{L}$. Then the filtered chain homotopy type of $C^-(G)$, viewed as a chain complex over $\mathbb{F}_2[U_1, \dots, U_\ell]$, is independent of the ordering of $\mathbb{O}$.*

Proof Different numberings can be connected via the filtered chain homotopies of Lemma 2.11. □

The basic link invariant is the filtered quasi-isomorphism class of $C^-(\vec{L})$, thought of as a complex of $\mathbb{F}_2[U_1, \dots, U_\ell]$ modules. But there are some other natural constructions one can consider.

For example, we can consider the chain complex $\widehat{C}(G)$, which is a chain complex over $\mathbb{F}_2$, once again which is freely generated by elements of $\mathbf{S}$, by setting the $U_i = 0$ for $i = 1, \dots, \ell$. We let $\widehat{CL}(G)$ denote the graded object $\mathrm{gr}(\widehat{C}(G))$ associated to the Alexander filtration, and let $\widehat{HL}(G)$ denote its homology.

Lemma 2.13 *The group $\widehat{HL}(G)$ is a finitely-generated $\mathbb{F}_2$–module.*

Proof Clearly, $\widehat{CL}(G)$ is a finitely generated R–module. It follows from Lemma 2.11 that once we set $U_i = 0$ for $i = 1, \dots, \ell$, then multiplication by U_j is null-homotopic for all $j = 1, \dots, n$, and in particular it acts trivially on homology. It follows at once that $\widehat{HL}(G) = H_*(\widehat{CL}(G))$, which is clearly a finitely generated R–module, is in fact a finitely generated $\mathbb{F}_2$–module. □

There is another construction which is quite convenient to consider for calculations [1]. This is the chain complex $\widetilde{C}(G)$, which is obtained from $C^-(G)$ by setting all the $U_i = 0$, and then taking the associated graded object. (This complex is denoted simply $C(G)$ in [6], but we prefer to reserve this notation for later use.) Explicitly, this is the free $\mathbb{F}_2$–module generated by $\mathbf{S}$, endowed with the differential

$$\widetilde{\partial}(\mathbf{x}) = \sum_{\mathbf{y}\in\mathbf{S}} \#\left\{ r \in \mathrm{Rect}(\mathbf{x},\mathbf{y}) \,\middle|\, \begin{array}{l} \forall x \in \mathbf{x}, x \notin \mathrm{Int}(r), \\ \forall i, O_i \notin r \text{ and } X_i \notin r \end{array} \right\} \cdot \mathbf{y}.$$

It is easy to relate the homology of $\widetilde{CL}(G) = \mathrm{gr}(\widetilde{C}(G))$ with the homology of $\widehat{CL}(G)$, by some principles in homological algebra.

Lemma 2.14 *Let C be a filtered, graded chain complex of free modules over $\mathbb{F}_2[U_1, \dots, U_n]$, such that U_i decreases the homological grading by two and the filtration by one, and such that multiplication by U_i is chain homotopic to multiplication by U_j for any i, j. Then $H_*(C/\{U_i = 0\}_{i=1}^n) \cong H_*(C/U_1) \otimes V^{n-1}$, where V is the two-dimensional bigraded vector space spanned by one generator in bigrading $(-1,-1)$ and another in bigrading $(0,0)$.*

Proof Suppose for notational simplicity that $n = 2$. Consider the chain map from the mapping cone of the chain map $U_1\colon C \longrightarrow C$ to C/U_1 gotten by taking the quotient on the second summand. It follows easily from the five-lemma that this map is a quasi-isomorphism. Moreover, by iterating this observation, we see that $C/(U_1, U_2)$ is quasi-isomorphic to the mapping cone

$$\begin{array}{ccc} C & \xrightarrow{U_1} & C \\ {\scriptstyle U_2}\downarrow & & \downarrow{\scriptstyle U_2} \\ C & \xrightarrow{U_1} & C, \end{array}$$

which in turn is quasi-isomorphic to the mapping cone of

$$U_2\colon C/U_1 \longrightarrow C/U_1.$$

But since U_1 and U_2 are chain homotopic in C, we obtain an induced null-homotopy of the map induced by U_2 on C/U_1. Thus, this latter mapping cone is isomorphic to the mapping cone of zero, ie to the direct sum $C/U_1 \oplus C/U_1$, which in turn is quasi-isomorphic to $(C/U_1)\otimes V$.

We investigate now the filtrations and gradings. In order for the quasi-isomorphism from $U_1\colon C \longrightarrow C$ to C/U_1 to be a filtered and graded map, we must shift gradings an filtrations on the mapping cone $M(U_1)$ appropriately. Specifically, let $C[a,b]$ denote the graded and filtered chain complex with the property that $\mathcal{F}_s(C_d[a,b]) = \mathcal{F}_{s+b}(C_{d+a})$. Then the mapping cone $M(U_1)$ is $C[1,1]\oplus C$. Following through the above discussion, we see that the mapping cone $C/(U_1,U_2)$ is filtered and graded quasi-isomorphic to $C[1,1]/U_1 \oplus C/U_1 \cong (C/U_1)\otimes V$.

This discussion generalizes readily to the case where $n>2$. □

Proposition 2.15 *The homology groups $\widehat{HL}(G)$ determine $\widetilde{HL}(G)$; specifically,*

$$H_*(\widetilde{CL}(G)) \cong \widehat{HL}(G)\otimes \bigotimes_{i=1}^{\ell} V_i^{\otimes(n_i-1)},$$

where V_i is the two-dimensional vector space spanned by two generators, one in zero Maslov and Alexander multigradings, and the other in Maslov grading minus one and Alexander multigrading corresponding to minus the i th basis vector.

Proof This follows easily from Lemma 2.14, applied component by component. □

Notation Perhaps the reader will find it convenient if we collect our notational conventions here. The chain complex $C^-(G)$ refers to the full chain complex (and indeed, we soon drop the minus from the notation here), $CL^-(G)$ denotes its associated graded object, and $HL^-(G)$ is the homology of the associated graded object. $\widehat{C}(G)$ denotes the chain complex where we set one $U_i=0$ for each component of the link, $\widehat{CL}(G)$ is its associated graded object, and $\widehat{HL}(G)$ is the homology of the associated graded object. $\widetilde{C}(G)$ is the chain complex $C^-(G)$ modulo the relations that *every* $U_i=0$, $\widetilde{CL}(G)$ is the associated graded complex, and $\widetilde{HL}(G)$ is its homology. Most of these constructions have their analogues in Heegaard Floer homology; for example, according to [6], $HL^-(G)$ is identified with $HFL^-(L)$, and $\widehat{HL}(G)$ with $\widehat{HFL}(L)$. We find it useful to distinguish these objects, especially when establishing properties of the combinatorial complex which could alternatively be handled by appealing to [6], together with known properties of Heegaard Floer homology.

3 Invariance of combinatorial knot Floer homology

Our goal in this section is to use elementary methods to show that combinatorial knot Floer homology is independent of the grid diagram, proving Theorem 1.2 with coefficients in $\mathbb{F}_2$.

Following Cromwell [2] (compare also Dynnikov [3]), any two grid diagrams for the same link can be connected by a sequence of the following elementary moves:

(1) (*Cyclic permutation*) This corresponds to cyclically permuting the rows and then the columns of the grid diagram.

(2) (*Commutation*) Consider a pair of consecutive columns in the grid diagram G with the following property: if we think of the X and the O from one column as separating the vertical circle into two arcs, then the X and the O from the adjacent column occur both on one of those two arcs. Under these hypotheses, switching the decorations of these two columns is a commutation move, cf Figure 5. There is also a similar move where the roles of columns and rows are interchanged.

(3) (*Stabilization/destabilization*) Stabilization is gotten by adding two consecutive breaks in the link. More precisely, if G has arc index n, a stabilization H is an arc index $n+1$ grid diagram obtained by splitting a row in G in two and introducing a new column. For convenience, label the original diagram so it has decorations $\{X_i\}_{i=2}^{n+1}$, $\{O_i\}_{i=2}^{n+1}$. Let O_i and X_i denote the two decorations in the original row. We copy O_i onto one of the two new copies of the row it used to occupy, and copy X_i onto the other copy. We place decorations O_1 and X_1 in the new column so O_1 resp. X_1 occupy the same row as X_i resp. O_i in the new diagram, cf Figure 6. Destabilization is the inverse move to stabilization. Note that stabilization can be alternatively done by reversing the roles of rows and columns in the above description; however, such a stabilization can be reduced to the previous case, combined with a sequence of commutation moves. In fact, we can consider only certain restricted stabilization moves, where three of the four squares O_1, X_1, O_i, and X_i share a common vertex; ie the new column is introduced next to O_i or X_i. However, there are now different types of stabilizations corresponding to the different ways of dividing the O's and X's among the two new rows.

Of course, since our complex is associated not to the planar grid diagram, but rather to the induced picture on the torus, the fact that it is invariant under cyclic permutation is a tautology.

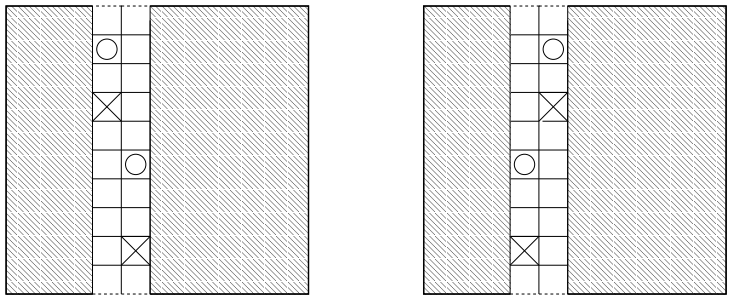

Figure 5: Commutation. The two grid diagrams differ from each other by interchanging the two columns, but correspond to the same link.

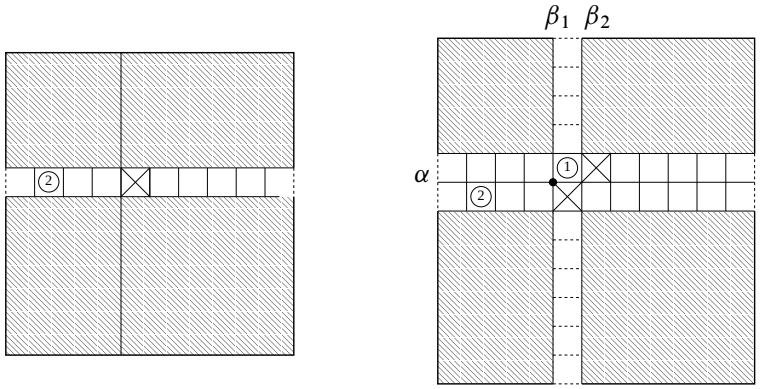

Figure 6: Stabilization. On the left, we have an initial grid diagram; on the right, a new diagram obtained by inserting the pictured row and column. Another stabilization is given by switching the roles of the new middle two rows.

We turn to commutation invariance next, and then stabilization invariance.

Note that all the chain complexes $\widehat{C}(G)$, $\widetilde{C}(G)$ depend on the quasi-isomorphism type of $C^-(G)$; thus, the latter is the most basic object. Thus, to streamline notation, we choose here to drop the superscript "$-$" from the notation of this chain complex and its differential.

3.1 Commutation invariance

Let G be a grid diagram for $\vec{L}$, and let H be a different grid diagram obtained by commuting two vertical edges. It is convenient to draw both diagrams on the same

torus, replacing a distinguished vertical circle β for G with a different one γ for H, as pictured in Figure 7. The circles β and γ meet each other transversally in two points a and b, which are not on a horizontal circle.

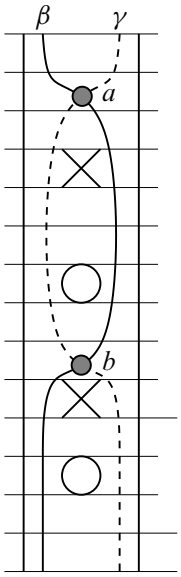

Figure 7: A commutation move, viewed as replacing one vertical circle (β, undashed) with another (γ, dashed)

We define a chain map $\Phi_{\beta\gamma}\colon C(G) \longrightarrow C(H)$ by counting pentagons. Given $\mathbf{x} \in \mathbf{S}(G)$ and $\mathbf{y} \in \mathbf{S}(H)$, we let $\mathrm{Pent}_{\beta\gamma}(\mathbf{x}, \mathbf{y})$ denote the space of embedded pentagons with the following properties. This space is empty unless $\mathbf{x}$ and $\mathbf{y}$ coincide at $n-2$ points. An element of $\mathrm{Pent}_{\beta\gamma}(\mathbf{x}, \mathbf{y})$ is an embedded disk in $\mathcal{T}$, whose boundary consists of five arcs, each contained in horizontal or vertical circles. Moreover, under the orientation induced on the boundary of p, we start at the β–component of $\mathbf{x}$, traverse the arc of a horizontal circle, meet its corresponding component of $\mathbf{y}$, proceed to an arc of a vertical circle, meet the corresponding component of $\mathbf{x}$, continue through another horizontal circle, meet the component of $\mathbf{y}$ contained in the distinguished circle γ, proceed to an arc in γ, meet an intersection point of β with γ, and finally, traverse an arc in β until we arrive back at the initial component of $\mathbf{x}$. Finally, all the angles here are required to be less than straight angles. These conditions imply that there is a particular intersection point, denoted a, between β and γ which appears as one of the corners of any pentagon in $\mathrm{Pent}_{\beta\gamma}(\mathbf{x}, \mathbf{y})$. The other intersection point b appears in

all of the pentagons in $\mathrm{Pent}_{\gamma\beta}(\mathbf{y},\mathbf{x})$. Examples are pictured in Figure 8. The space of empty pentagons $p \in \mathrm{Pent}_{\beta\gamma}(\mathbf{x},\mathbf{y})$ with $\mathbf{x} \cap \mathrm{Int}(p) = \varnothing$, is denoted $\mathrm{Pent}^{\circ}_{\beta\gamma}$.

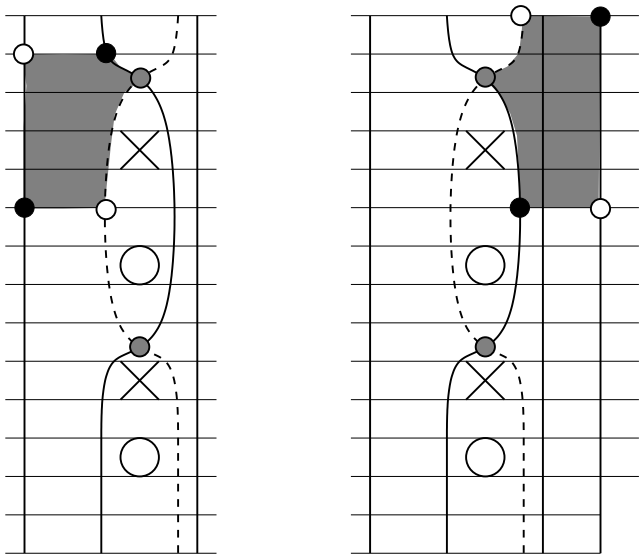

Figure 8: We have indicated here two allowed pentagons in $\mathrm{Pent}_{\beta\gamma}(\mathbf{x},\mathbf{y})$, where components of $\mathbf{x}$ are indicated by solid points, and those of $\mathbf{y}$ are indicated by hollow ones.

Given $\mathbf{x} \in \mathbf{S}(G)$, define

$$\Phi_{\beta\gamma}(\mathbf{x}) = \sum_{\mathbf{y}\in\mathbf{S}(H)} \sum_{p\in\mathrm{Pent}^{\circ}_{\beta\gamma}(\mathbf{x},\mathbf{y})} U_1^{O_1(p)} \cdots U_n^{O_n(p)} \cdot \mathbf{y} \in C(H).$$

Lemma 3.1 *The map $\Phi_{\beta\gamma}$ is a filtered chain map.*

Proof The fact that $\Phi_{\beta\gamma}$ preserves Alexander filtration and Maslov gradings is straightforward. Like the proof of Proposition 2.10, the proof that $\Phi_{\beta\gamma}$ is a chain map proceeds by considering domains which are obtained as a juxtaposition of a pentagon and a rectangle, representing terms in $\partial \circ \Phi_{\beta\gamma}$, and observing that such domains typically have an alternate decomposition to represent a term in $\Phi_{\beta\gamma} \circ \partial$. One example is illustrated in Figure 9. Other terms are more straightforward, consisting either of a disjoint rectangle and pentagon, a rectangle and pentagon with overlapping interior, or a rectangle and a pentagon which meet along a different edge; the pictures are similar to those in Figure 3. There is one special case, of a type of domain which has only one decomposition: these are the domains obtained as the union of a width one pentagon p

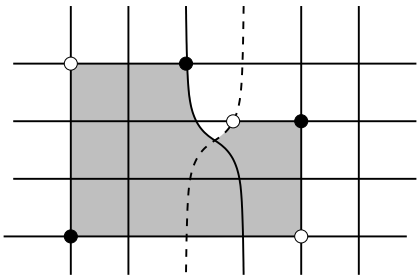

Figure 9: Chain map. The given domain can be decomposed either as a pentagon followed by a rectangle, or a rectangle followed by a pentagon. The first decomposition represents a term in $\partial\circ\Phi_{\beta\gamma}$, the second a term in $\Phi_{\beta\gamma}\circ\partial$.

and a width one rectangle r. In this case, if we let $\mathbf{x}\in\mathbf{S}(G)$, there is a canonical closest generator $c(\mathbf{x})\in\mathbf{S}(H)$ (with the property that $\mathbf{x}$ and $c(\mathbf{x})$ agree at all intersection points away from $\beta\cup\gamma$). It is easy to see, then, that our domain has the form $r*p$ or $p*r$ (depending on the local picture of $\mathbf{x}$), and it connects $\mathbf{x}$ to $c(\mathbf{x})$. But then, such domains are in one-to-one correspondence with domains of the form $r'*p'$ or $p'*r'$, where if p is a left pentagon, then p' is a right pentagon, and vice versa. See Figure 10. □

We can define chain homotopy operators analogously, only now counting hexagons.

More specifically, given $\mathbf{x},\mathbf{y}\in\mathbf{S}(G)$, we let $\mathrm{Hex}_{\beta\gamma\beta}(\mathbf{x},\mathbf{y})$ denote the space of embedded hexagons with the following property. This space, too, is empty unless $\mathbf{x}$ and $\mathbf{y}$ coincide at $n-2$ points. Moreover, an element of $\mathrm{Hex}_{\beta\gamma\beta}(\mathbf{x},\mathbf{y})$ is an embedded disk in $\mathcal{T}$, whose boundary consists of six arcs, each contained in horizontal or vertical circles. More specifically, under the orientation induced on the boundary of p, we start at the β–component of $\mathbf{x}$, traverse the arc of a horizontal circle, meet its corresponding component of $\mathbf{y}$, proceed to an arc of a vertical circle, meet its corresponding component of $\mathbf{x}$, continue through another horizontal circle, meet its component of $\mathbf{y}$, which contained in the distinguished circle β, continue along β until the intersection point b of β with γ, continue on γ to the intersection point a of β and γ, proceed again on β to the β–component of $\mathbf{x}$, which was also our initial point. Moreover, all corner points of our hexagon are again required to be less than straight angles. An example is given in Figure 11. We define the space of empty hexagons $\mathrm{Hex}^{\circ}_{\beta\gamma\beta}$, with interior disjoint from $\mathbf{x}$, as before. There is also a corresponding notion $\mathrm{Hex}_{\gamma\beta\gamma}$. We now define the

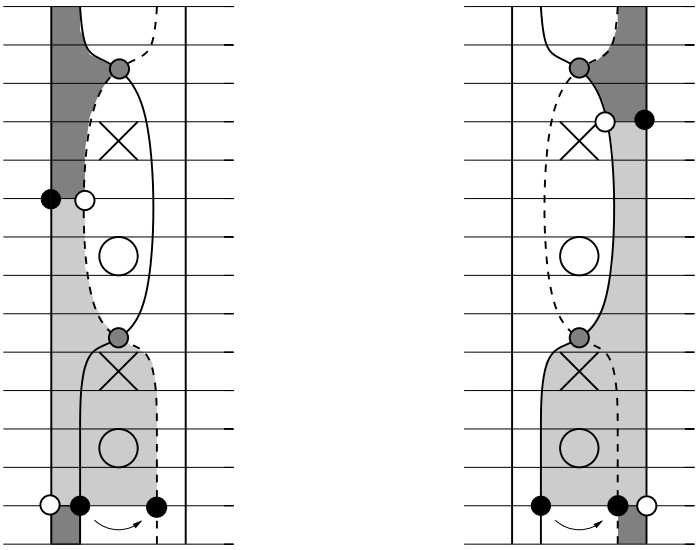

Figure 10: Special case of the chain map. The generators $\mathbf{x}$ and $c(\mathbf{x})$ are marked by dark circles; they differ from each other only on one row. The arrow indicates how the dark circle in $\mathbf{x}$ is replaced by a corresponding dark circle in $c(\mathbf{x})$. On the left we have a (darkly shaded) pentagon followed by a (lightly shaded) rectangle, and on the right we have a rectangle followed by a pentagon. The intermediate generators are marked by hollow circles.

function $H_{\beta\gamma\beta}\colon C(G) \longrightarrow C(G)$ by

$$H_{\beta\gamma\beta}(\mathbf{x}) = \sum_{\mathbf{y}\in\mathbf{S}(G)} \sum_{h\in \mathrm{Hex}^{\circ}_{\beta\gamma\beta}(\mathbf{x},\mathbf{y})} U_1^{O_1(h)} \cdots U_n^{O_n(h)} \cdot \mathbf{y}.$$

Proposition 3.2 *The map $\Phi_{\beta\gamma}\colon C(G) \longrightarrow C(H)$ is a chain homotopy equivalence; more precisely*

$$\mathbb{I} + \Phi_{\gamma\beta}\circ\Phi_{\beta\gamma} + \partial\circ H_{\beta\gamma\beta} + H_{\beta\gamma\beta}\circ\partial = 0$$
$$\mathbb{I} + \Phi_{\beta\gamma}\circ\Phi_{\gamma\beta} + \partial\circ H_{\gamma\beta\gamma} + H_{\gamma\beta\gamma}\circ\partial = 0.$$

Proof Juxtaposing two pentagons appearing in $\Phi_{\gamma\beta}\circ\Phi_{\beta\gamma}$, we generically obtain a composite domain which admits a unique alternative decomposition as a hexagon and a square, counted in $\partial\circ H_{\beta\gamma\beta}$ or $H_{\beta\gamma\beta}\circ\partial$. Typically, the remaining terms in $\partial\circ H_{\beta\gamma\beta}$ cancel with terms $H_{\beta\gamma\beta}\circ\partial$.

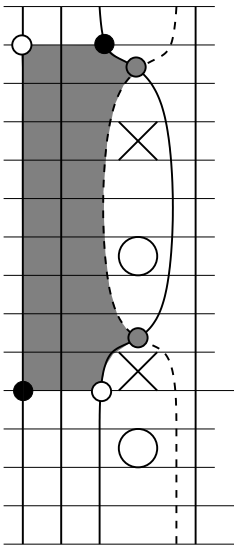

Figure 11: A hexagon in $\mathrm{Hex}_{\beta\gamma\beta}$

There is, however, one composite region which has a unique decomposition. Specifically, the vertical circles β_1, β_2, and γ divide up $\mathcal{T}$ into a collection of components, two of which are annuli and do not contain any X. Depending on the initial point $\mathbf{x}$, exactly one of these annuli can be thought of as a juxtaposition of two pentagons, or a hexagon and a rectangle which is counted once in $\Phi_{\gamma\beta}\circ\Phi_{\beta\gamma}+\partial\circ H_{\beta\gamma\beta}+H_{\beta\gamma\beta}\circ\partial$; but it is also counted in the identity map. See Figure 12. □

3.2 Stabilization invariance

Let G be a grid diagram and H denote a stabilization. We discuss in detail the case where we introduce a new column with O_1 immediately above X_1 (and X_2 is immediately to the left or to the right of O_1); the case where X_1 is immediately above O_1 can be treated symmetrically by a rotation of all diagrams by $180°$.

More specifically, given a horizontal arc from O_2 to X_2, we introduce a vertical segment (somewhere along the arc) consisting of a new pair O_1 and X_1, where O_1 is on the square right above X_1, which in turn is in the same row as the new copy of O_2, as in Figure 6. Indeed, do this in such a manner that three of the four squares marked O_1, O_2, X_1, and X_2 share a common vertex. Furthermore, by applying commutation, we can assume without loss of generality that these three squares are

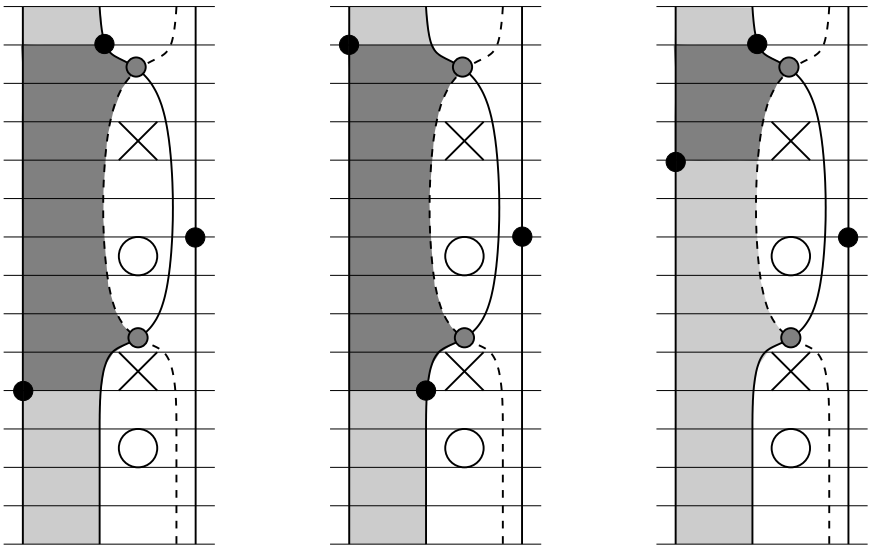

Figure 12: Decomposing the identity map. Consider the three configurations in $C(G)$, indicated by dark circles. The shaded region can be thought of as decomposed into a hexagon followed by a rectangle (as on the left), a rectangle followed by a hexagon (as in the middle), or a pair of pentagons as on the right. The first can be thought of as counting terms in $\partial \circ H_{\beta\gamma}$, the middle terms in $H_{\beta\gamma} \circ \partial$, and the right in $\Phi_{\gamma\beta} \circ \Phi_{\beta\gamma}$. There are three more cases, if the β–component of the configuration lies on the other arc in β; in this case, we must decompose the annulus on the right.

O_1, X_1, and X_2. Thus, the grid of H is gotten by inserting a new column of squares, where two consecutive squares are marked by O_1 and X_1. We let β_1 be the vertical circle on the left, and β_2 the one on the right. Let α denote the new horizontal circle in H which separates O_1 from X_1.

Let $B = C^-(G)$ and $C = C^-(H)$. Let C' be the mapping cone of

$$U_2 - U_1 \colon B[U_1] \longrightarrow B[U_1],$$

ie $C'[U_1] = B[U_1] \oplus B[U_1]$, endowed with the differential $\partial \colon C' \longrightarrow C'$ given by

$$\partial'(a, b) = (\partial a, (U_2 - U_1) \cdot a - \partial b)$$

where here ∂ denotes the differential within C (actually, in the sequel we drop the prime from the differential within C', as well, and hope that the differential is clear from the context). Note that B is a chain complex over $\mathbb{F}_2[U_2, \dots, U_n]$, so that $B[U_1]$

denotes the induced complex over $\mathbb{F}_2[U_1,\dots,U_n]$ gotten by introducing a new formal variable U_1. Let $\mathcal{L}$ and $\mathcal{R}\cong B[U_1]$ be the subgroups of C' of elements of the form $(c,0)$ and $(0,c)$ for $c\in B[U_1]$, respectively. The module $\mathcal{R}$ inherits Alexander and Maslov gradings from its identification with $B[U_1]$, while $\mathcal{L}$ is given the Alexander and Maslov gradings which are one less than those it inherits from its identification with $B[U_1]$. With respect to these conventions, the mapping cone is a filtered complex of R–modules.

Lemma 3.3 *The map from C' to B that takes (a,b) to $a/\{U_1=U_2\}$ is a quasi-isomorphism.*

Proof In general, the mapping cone C' of a map $f\colon C_1\to C_2$ fits into a short exact sequence on homology from C_2 to C' to C_1. The connecting homomorphism in the corresponding long exact sequence on homology is the map induced by f. In this case, f is U_1-U_2, which is injective on the homology of $B[U_1]$, so the map from C' to B is a quasi-isomorphism. □

It therefore suffices to define a filtered quasi-isomorphism

$$F\colon C\longrightarrow C'. \tag{5}$$

To do this, we introduce a little more notation.

Let $\mathbf{S}(G)$ be the generating set of B, and $\mathbf{S}(H)$ be the generating set of C. Let x_0 be the intersection point of α and β_1 (the dark dot in Figure 6). Let $\mathbf{I}\subset\mathbf{S}(H)$ be the set of $\mathbf{x}\in\mathbf{S}(H)$ which contain x_0. There is, of course, a natural (pointwise) identification between $\mathbf{S}(G)$ and $\mathbf{I}$, which drops Alexander and Maslov grading by one. More precisely, given $\mathbf{x}\in\mathbf{S}(G)$, let $\phi(\mathbf{x})\in\mathbf{S}(H)$ denote the induced generator in $\mathbf{I}$ which is gotten by inserting x_0. We then have

$$M_{C(G)}(\mathbf{x})=M_{C(H)}(\phi(\mathbf{x}))+1=M_{C'}(0,\phi(\mathbf{x}))=M_{C'}(\phi(\mathbf{x}),0)+1 \tag{6}$$

$$A_{C(G)}(\mathbf{x})=A_{C(H)}(\phi(\mathbf{x}))+g(1)=A_{C'}(0,\phi(\mathbf{x}))=A_{C'}(\phi(\mathbf{x}),0)+g(1) \tag{7}$$

where g is the function from Section 2.2, mapping from i to the basis vector corresponding to the component of the link containing O_i. With this said, we will henceforth suppress ϕ from the notation, thinking of $\mathcal{L}$ and $\mathcal{R}$ as generated by configurations in $\mathbf{I}\subset\mathbf{S}(H)$.

As such, the differentials within $\mathcal{L}$ and $\mathcal{R}$ count rectangles which do not contain x_0 on their boundary, although they may contain x_0 in their interior. Note however that the boundary operator (in $\mathcal{L}$ and $\mathcal{R}$) for rectangles containing x_0 does not involve the variable U_1.

Definition 3.4 For $\mathbf{x} \in \mathbf{S}(H)$ and $\mathbf{y} \in \mathbf{I} \subset \mathbf{S}(H)$, a domain $p \in \pi(\mathbf{x}, \mathbf{y})$ is said to be of type L or R if either it is trivial, in which case p has type L, or it satisfies the following conditions:

- p has only nonnegative local multiplicities.
- For each $c \in \mathbf{x} \cup \mathbf{y}$, other than x_0, at least three of the four adjoining squares have vanishing local multiplicities.
- In a neighborhood of x_0 the local multiplicities in three of the adjoining rectangles are the same number k. When p has type L, the lower left corner has local multiplicity $k-1$, while for p of type R the lower right corner has multiplicity $k+1$.
- ∂p is connected.

The *complexity* of the trivial domain is 1; the complexity of any other domain is the number of horizontal lines in its boundary. The set of type L (or R) domains from $\mathbf{x}$ to $\mathbf{y}$ is denoted $\pi^L(\mathbf{x}, \mathbf{y})$ (or $\pi^R(\mathbf{x}, \mathbf{y})$). We set $\pi^F(\mathbf{x}, \mathbf{y}) = \pi^L(\mathbf{x}, \mathbf{y}) \cup \pi^R(\mathbf{x}, \mathbf{y})$, and call its elements domains of *type* F; see Figure 13 for examples. We denote by π^F the union of the sets $\pi^F(\mathbf{x}, \mathbf{y})$, over all possible $\mathbf{x}$ and $\mathbf{y}$.

The *innermost height* (resp. *width*) of a domain in π^F is the vertical (resp. horizontal) distance from the corner adjacent horizontally (resp. vertically) to x_0 to the corner after that.

We now define maps

$$F^L\colon C \longrightarrow \mathcal{L}$$
$$F^R\colon C \longrightarrow \mathcal{R}$$

where F^L (resp. F^R) counts domains of type L (resp. R) without factors of U_1. Specifically, define

$$F^L(\mathbf{x}) = \sum_{\mathbf{y}\in\mathbf{S}} \sum_{p\in\pi^L(\mathbf{x},\mathbf{y})} U_2^{O_2(p)} \cdots U_n^{O_n(p)} \cdot \mathbf{y}$$
$$F^R(\mathbf{x}) = \sum_{\mathbf{y}\in\mathbf{S}} \sum_{p\in\pi^R(\mathbf{x},\mathbf{y})} U_2^{O_2(p)} \cdots U_n^{O_n(p)} \cdot \mathbf{y}.$$

We put these together to define a map

$$F = \begin{pmatrix} F^L \\ F^R \end{pmatrix}\colon C \longrightarrow C'.$$

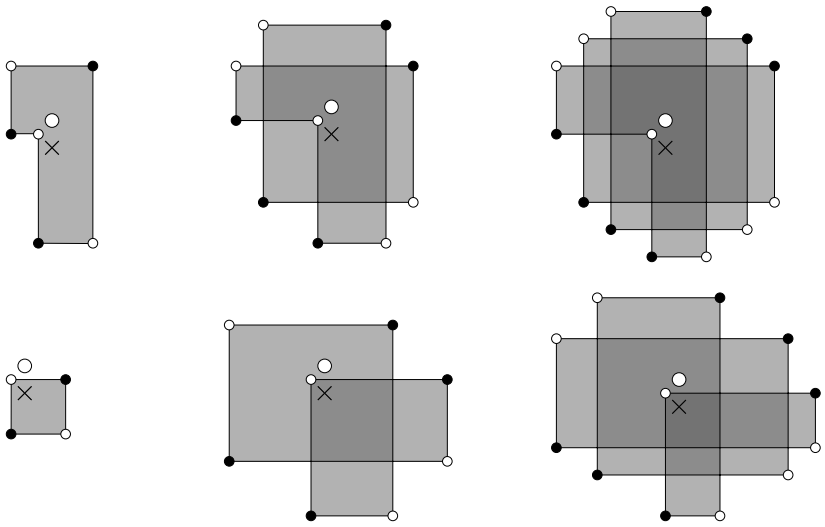

Figure 13: Types of domains. We have listed here domains in the stabilized diagram, labeling the initial points by dark circles, and terminal points by empty circles. The top row lists domains of type L, while the second row lists some of type R. The marked O and X are the new ones in the stabilized picture. Complexities from the left on the first row are 3, 5, and 7 respectively; on the second, they are 2, 4, and 6. Darker shading corresponds to higher local multiplicities. Not shown is the trivial domain of type L, which has complexity 1.

Lemma 3.5 *The map $F\colon C \longrightarrow C'$ preserves Maslov grading, respects Alexander filtrations, and is a chain map.*

Proof The fact that the gradings and filtrations are respected is straightforward. For instance, the Alexander filtration shift of a region p is given by counting the number of O's minus the number of X's contained in p. A region of type L contains O_1 and X_1 an equal number of times, and every other O_i comes with a cancelling factor of U_i, so the Alexander filtration shift is negative. The other shifts can be checked in a similar way.

To prove that F is a chain map, we consider all the terms in the expression $\partial \circ F$ or $F \circ \partial$. Most of these are counts of composite domains $p * r$ or $r * p$, where r is a rectangle and p is a type L or R domain. A rectangle $r \in \pi(\mathbf{x}, \mathbf{y})$ cannot contribute to this count if any component $x \in \mathbf{x}$ is in the interior of r, except in the special case where $x = x_0$, and the rectangle is thought of as connecting two intersection points in

$\mathcal{L}$ or $\mathcal{R}$, in which case we say it is of Type 2. All other empty rectangles are said to be of Type 1.

There are several cases of domains contributing to $\partial \circ F$ or $F \circ \partial$, which we group according to whether r is a Type 1 or Type 2 rectangle, and to how many corners p and r have in common. We list the cases below; verifying that these are the only cases is a straightforward exercise in planar geometry.

If r is of Type 1, we have the following possibilities:

I(0) A composition in either order of a domain $p \in \pi^F$ and an empty rectangle r of Type 1, with all corners distinct. This domain appears in both $\partial \circ F$ and $F \circ \partial$ as compositions in two different orders, $p * r$ and $r' * p'$, where r has the same support as r' and p has the same support as p'.

I(1) A composition in either order of a nontrivial domain $p \in \pi^F$ and an empty rectangle r, with p and r sharing one corner and r disjoint from x_0 (including the boundary). The union of these two domains has a unique concave corner not at x_0, and we can slice this into a domain in π^F and a rectangle of Type 1 in two ways by cutting in either way from this concave corner. This gives the domain as a composition in exactly two ways. An example is shown in Figure 14.

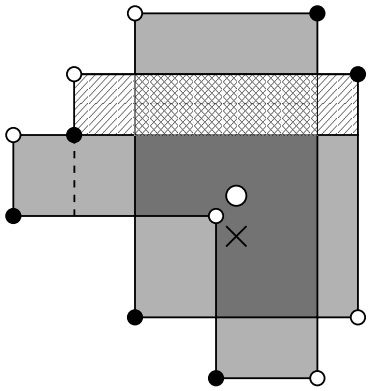

Figure 14: Case I(1). An example of a domain with two decompositions $r * p = r' * p'$, both accounted for in case I(1)

I($1'$) A composition $r * p$ with r and p sharing one corner and x_0 appearing on the horizontal or vertical boundary of r. The composite looks again like a domain in π^F or the rotation by $180°$ of such a domain. See Figure 15. A special case worth mentioning is when $r \in \mathrm{Rect}^\circ(\mathbf{x}, \mathbf{y})$ with $\mathbf{y} \in \mathbf{I}$; in this case p is trivial, with complexity 1, as in Figure 16.

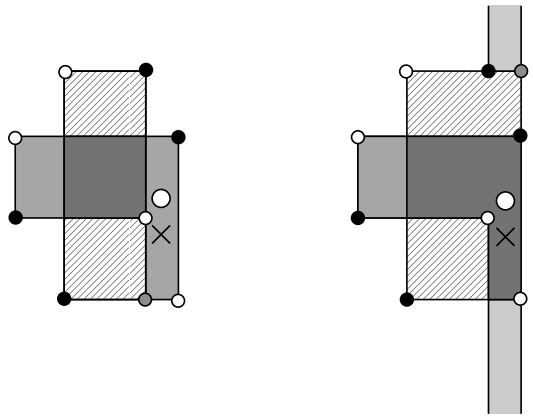

Figure 15: Cases I($1'$) and II(1). There are two terms in $\partial \circ F^L + F^L \circ \partial$ starting at the black dots and ending at the white dots, thought of as elements of $\mathcal{L}$. The term on the left is a juxtaposition $r * p$ (as in I($1'$)), while the second is $p' * r'$, where p' is of type L and r' is of Type 2 (as in II(1)).

I(2) A composition in either order of $p \in \pi^F$ and $r \in \mathrm{Rect}^\circ$, where p and r share two corners other than possibly x_0; see Figures 20 and 21. In this case p has complexity at least 3.

I(3) A domain that wraps around the torus with a decomposition as $p * r$ or $r * p$, where r is an empty rectangle of Type 1 and $p \in \pi^F$ has innermost height or width equal to 1, and r and p share three corners other than possibly x_0. This decomposition is unique. The total domain contains a unique vertical or horizontal annulus of height or width equal to 1. When the complexity m of p is equal to 2, the domain is just this annulus. Examples are shown in Figures 16 ($m = 3$, horizontal), 17 ($m = 5$, horizontal), 18 ($m = 5$, vertical), and 19 ($m = 4$, horizontal).

If r is of Type 2, the composition must be of the form $p * r$, because Type 2 rectangles only appear in the differential of the target complex C'. We only have two possibilities:

II(0) All the corners of p and r are disjoint.

II(1) A domain that wraps around the torus with a decomposition as $p * r$, where r is a rectangle of Type 2 that shares one corner with p. This decomposition is unique, and the total domain again contains a unique thin (ie width one or height one) annulus. See Figure 15.

Apart from these, there is one other special contribution to $F \circ \partial$, which does not come from a decomposition of a domain into $p * r$ or $r' * p'$:

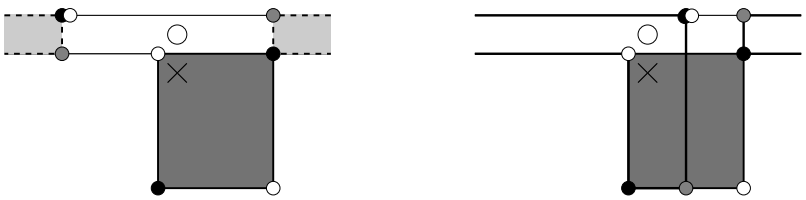

Figure 16: Cases I($1'$) and I(3), where $m = 1$. In both pictures, the darkly shaded rectangle represents a map from the black generator to the white one, followed by the natural map (induced by the trivial domain, which has complexity 1) to the white generator thought of as an element of in $\mathcal{L}$. This is accounted for in I($1'$). Depending on the placement of the black dot in the top row, we can cancel this either with a term in $F^L \circ \partial$ or $\partial \circ F^L$. In the first case (on the left), we have the domain $r * p$, where r is the height one (lightly shaded) rectangle in the row through O_1, to the intermediate generator (labelled by the shaded circle), thought of as a differential within $C(H)$, followed by a complexity 3 domain p with innermost height equal to one, which we trust the reader can spot. In the second case (on the right), we have the decomposition $p * r$, where p is the complexity 3 domain with innermost height equal to one from the black generator to the intermediate generator, which is bounded by the dark line, followed by a rectangle to the white generator, which again we leave to the reader to find. In both cases the alternate term is accounted for in case I(3).

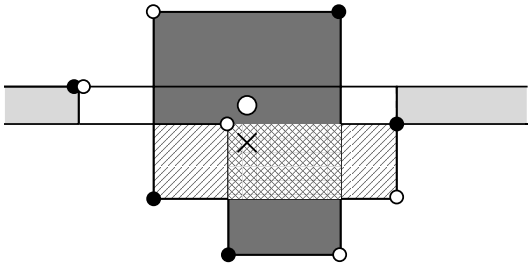

Figure 17: Cases I($1'$) and I(3), horizontal annulus. There are two terms in $\partial \circ F^L + F^L \circ \partial$ starting at the black dots and ending at the white dots. One of them counts the composite domain $r' * p'$ where r' is the hatched rectangle containing X, and p' is the darkly shaded complexity 3 domain (accounted for in I($1'$)); and the other is a count of $r * p$, where r is the height one, lightly shaded rectangle, followed by a complexity 5 domain with innermost height equal to one (accounted for in I(3)).

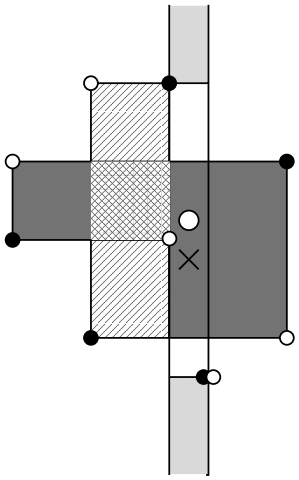

Figure 18: Cases I($1'$) and I(3), vertical annulus. There are two terms in $\partial \circ F^L + F^L \circ \partial$ starting at the black dots and ending at the white dots. One of them counts the composite domain $r' * p'$ where r' is the hatched rectangle containing the white dot x_0 in its boundary, and p' is the darkly shaded complexity 3 domain (accounted for in I($1'$)); and the other is a count of $r * p$, where r is the height one, lightly shaded rectangle, followed by a complexity 5 domain with innermost height equal to one (accounted for in I(3)).

(S) A domain $p \in \pi^L$ followed by the differential from $\mathcal{L}$ to $\mathcal{R}$, which multiplies by $U_2 - U_1$.

Contributions from case I(0) cancel each other out, and the same goes for those from case I(1). In fact, these cases are the exact analogs of the first three cases in Figure 3 for the proof of Proposition 2.10. See Figure 14 for an example.

We claim that contributions from case I($1'$) cancel with contributions from case II(1) or I(3), together with possibly a contribution from case (S). Indeed, for each domain of type I($1'$) made of a rectangle $r_1 \in \mathrm{Rect}^\circ(\mathbf{x}, \mathbf{y})$ and a domain $p_1 \in \pi^F(\mathbf{y}, \mathbf{z})$ of complexity m, let $p_0 = r_1 * p_1$. We can make a new domain p_0' by adding a thin annulus abutting x_0 on the opposite side of x_0 from r_1. (For instance, if the right side of r touches x_0, add a vertical annulus of width one whose left side touches x_0.) In the case when $m = 1$, when r_1 touches x_0 at a corner, we attach a horizontal annulus if r_1 contains X_1 and a vertical annulus otherwise, as in Figure 16. If the innermost height or width of p_1 is 1, then p_0 decomposes as $p_2 * r_2$, where $p_2 \in \pi^F$ has complexity m. This corresponds to a contribution from case II(1), as in Figure 15.

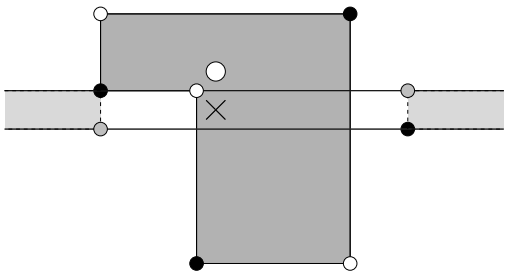

Figure 19: Cases I(1′), I(3), and (S). This case is similar to those in Figures 17 and 18, except that it also involves a domain of type (S). We count terms in $\partial \circ F + F \circ \partial$ starting at the black dots and ending at the white dots (thought of as representing an element of $\mathcal{R}$). The darkly shaded polygon represents a domain of type L from the black to the white generator. Postcomposing with the differential from $\mathcal{L}$ to $\mathcal{R}$, we get $(U_2 - U_1)$ times the white generator. Alternatively, the region can be decomposed as a rectangle containing O_1 (with a factor of U_1), composed with the rectangle containing X_1, thought of as a polygon of type R. Alternatively, there is a term induced by the height one (lightly shaded) rectangle, followed by a complexity 4 domain of type R, which the reader can easily spot. One of these two domains contains O_2, and hence the composite will count with a factor of U_2.

If, on the other hand, the innermost height or width (as appropriate) of p_1 is not 1, the new domain p_0' is of type I(3) and in turn decomposes as $p_2 * r_2$ or $r_2 * p_2$, depending on the placement of the generator on the new row or column, where $p_2 \in \pi^F$ has complexity $m + 2$. See Figures 16–19.

In these cases involving annuli, if $p_i \in \pi^R$ and the annulus is horizontal, the rectangle r_1 contains O_1 and so has a contribution which is multiplied by U_1, while the domain p_0' contains O_2 and so has a contribution which is multiplied by U_2. Thus these two terms contribute $U_1 - U_2$ to the composite map from $\mathbf{x}$ to $\mathbf{z}$. On the other hand, in this case the domain p_0 is itself in $\pi^L(\mathbf{x}, \mathbf{z})$, and so we get a cancelling contribution of type (S), as in Figure 19. In other cases the two domains p_0 and p_0' give the same contribution to the boundary map.

Compositions $r * p$ or $p * r$ from case I(2), with p of complexity $m \geq 3$, cancel out compositions $r' * p'$ from case II(0), with p' of complexity $m - 2$, as illustrated in Figures 20 ($m = 3$) and 21 ($m = 5$).

The only domains left to cancel are those of type I(3) with $m = 2$ and type (S) with $m = 1$. There are two kinds of domains of type I(3) with $m = 2$: a vertical and a horizontal annulus, containing O_1 and O_2, respectively, and in both cases

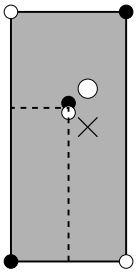

Figure 20: Cases I(2) and II(0), with complexity $m = 3$. The simplest case of the pairing between cases I(2) and II(0)

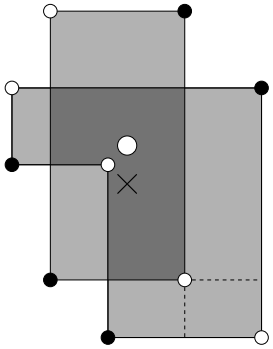

Figure 21: Cases I(2) and II(0). The illustrated domain can be decomposed as a complexity 3 domain of type L followed by a Type 2 rectangle (accounted for in II(0)), or alternatively a complexity 5 domain of type L followed by a Type 1 rectangle (accounted for in I(2)).

containing X_1. These domains map a generator $\mathbf{x} \in \mathbf{I}$ to itself, and so cancel the remaining contribution from the maps of type (S). □

In order to see that F is a quasi-isomorphism, we will introduce an appropriate filtration. Consider $\widetilde{C}(H)$. Let Q be a collection of $(n-1)^2$ dots, one placed in each square which do not appear in the row or column through O_1. Given $h \in (\frac{1}{2}\mathbb{Z})^\ell$, let $\widetilde{C}(H,h)$ denote the summand generated by generators $\mathbf{x}$ with Alexander gradings equal to h.

Note that for fixed $\mathbf{x}, \mathbf{y} \in \mathbf{S}_h$, for any two domains $p, p' \in \pi(\mathbf{x},\mathbf{y})$ with $O_i(p) = X_i(p) = O_i(p') = X_i(p')$ for all i, we have that $\#(Q \cap p) = \#(Q \cap p')$. Thus, we can find a function $\mathcal{F}$ so that for any $\mathbf{x}, \mathbf{y} \in \mathbf{S}$, if $p \in \pi(\mathbf{x},\mathbf{y})$ is a domain with

$O_i(p) = X_i(p) = 0$ for all i, then

$$\mathcal{F}(\mathbf{x}) - \mathcal{F}(\mathbf{y}) = \#(Q \cap p).$$

The function $\mathcal{F}$ determines a filtration on $\widetilde{C}(H,h)$, whose associated graded object counts only those rectangles which contain no O_i, X_i, or points in Q. Thus, these rectangles must be supported in the row or column through O_1. We let $\widetilde{C}_Q$ denote this associated graded object, and typically drop h from the notation.

We recall now a well-known principle from homological algebra (see for example McCleary [7, Theorem 3.2]).

Lemma 3.6 *Suppose that $F\colon C \longrightarrow C'$ is a filtered chain map which induces an isomorphism on the homology of the associated graded object. Then F is a filtered quasi-isomorphism.*

We decompose $\mathbf{S} = \mathbf{I} \cup (\mathbf{NI}) \cup (\mathbf{NN})$, where **NI** consists those configurations whose β_2 component is $\alpha \cap \beta_2$ and whose β_1 component is not in α, while **NN** consists of those whose β_2 component and β_1 component are not on α. We have corresponding decompositions of modules: $C = C^I \oplus C^{NI} \oplus C^{NN}$.

Lemma 3.7 *$H_*(\widetilde{C}_Q)$ is isomorphic to the free $\mathbb{F}_2$–module generated by elements of* **I** *and* **NI**.

Proof There are two cases, according to whether the X_2 marks the square to the left or the right of O_1.

Suppose X_2 is in the square just to the right of the square marked O_1. Then we have a direct sum splitting $\widetilde{C}_Q \cong \widetilde{C}_Q^{NI} \oplus B$, where the differentials in $\widetilde{C}_Q^{NI}$ are trivial, hence its homology is the free $\mathbb{F}_2$–module generated by elements of **NI**; and where B is a chain complex fitting into an exact sequence

$$0 \longrightarrow \widetilde{C}_Q^{I} \longrightarrow B \longrightarrow \widetilde{C}_Q^{NN} \longrightarrow 0.$$

Moreover, it is easy to see that $H_*(\widetilde{C}_Q^{NN}) = 0$. Finally, the differentials in $\widetilde{C}_Q^{I}$ are trivial, so its homology is the free $\mathbb{F}_2$–module generated by elements of **I**.

Suppose on the other hand that X_2 is just to the left of O_1. Then there is a direct sum splitting $\widetilde{C}_Q \cong \widetilde{C}_Q^{I} \oplus B'$, where once again the differentials on $\widetilde{C}_Q^{I}$ are trivial and B' fits into an exact sequence

$$0 \longrightarrow \widetilde{C}_Q^{NN} \longrightarrow B' \longrightarrow \widetilde{C}_Q^{NI} \longrightarrow 0,$$

where $H_*(\widetilde{C}_Q^{NN}) = 0$ and the differentials on $\widetilde{C}_Q^{NI}$ are trivial. □

Proposition 3.8 *The map F is a filtered quasi-isomorphism.*

Proof We consider the map induced by F:

$$\widetilde{F}_Q\colon \widetilde{C}_Q \longrightarrow \widetilde{C}'_Q.$$

$\widetilde{C}'_Q$ splits as a direct sum of chain complexes $\mathcal{L}_Q \oplus \mathcal{R}_Q$, both of which are freely generated by elements in $\mathbf{I}$.

There are two cases. First take the case where X_2 is in the square just to the right of the square marked O_1. Consider the subcomplex $\widetilde{C}^I_Q \oplus \widetilde{C}^{NI}_Q \subset \widetilde{C}_Q$. By Lemma 3.7, this subcomplex carries the homology, and hence it suffices to show that the restriction of $\widetilde{F}_Q$ to this subcomplex induces an isomorphism in homology.

To this end observe that $\widetilde{F}^L_Q$ restricted to $\widetilde{C}^I_Q$ is an isomorphism. Moreover, $\widetilde{F}^R_Q$ restricted to $\widetilde{C}^{NI}_Q$ counts rectangles supported in the row and column through O_1 and which contain X_1 in their interior and end up in $\mathbf{I}$ (since no other domains of type R is disjoint from Q). But for each element of $\mathbf{NI}$, there is a unique such rectangle. Thus $\widetilde{F}_Q$ is a quasi-isomorphism when X_2 is just to the right of O_1.

In the second case, where X_2 is just to the left of O_1, we proceed as follows. In this case $\widetilde{C}^I_Q$ is a direct summand of the complex $\widetilde{C}_Q$ (cf the proof of Lemma 3.7). Moreover, it is easy to see that $\widetilde{F}^L_Q$ restricted to $\widetilde{C}^I_Q$ is an isomorphism of chain complexes. It remains to show that the restriction of $\widetilde{F}^R_Q$ is a quasi-isomorphism. This is true because the only domains of type R which do not contain X_2 are rectangles, and those which are supported in the allowed region connect configurations of type $\mathbf{NI}$ to $\mathbf{I}$. Once again, the result now follows from the fact that there is a unique rectangle of type R connecting a given element of $\mathbf{NI}$ to an element of $\mathbf{I}$. This completes the verification that $\widetilde{F}_Q$ is a quasi-isomorphism.

We now appeal to Lemma 3.6 to conclude that $\widehat{F}$ is quasi-isomorphism; and another application of the same principle gives that F is a quasi-isomorphism, as well. □

Remark 3.9 The chain complex C' used in this stabilization proof can be viewed as the chain complex associated to the Heegaard diagram where the vertical circle β_1 is replaced by a small circle enclosing O_1 and X_1. In this Heegaard diagram it is straightforward to check that the counts of holomorphic disks are still combinatorial and equivalent to the boundary operator in C'.

3.3 Completion of topological invariance, without signs

We have now all the pieces needed to establish Theorem 1.2, with coefficients in $\mathbb{F}_2 = \mathbb{Z}/2\mathbb{Z}$.

Proof of Theorem 1.2 This result now is an immediate consequence of Cromwell's theorem, our earlier remarks on cyclic permutation, and Propositions 3.8 and 3.2. □

4 Signs

Definition 4.1 A *true sign assignment*, or simply a *sign assignment*, is a function

$$\mathcal{S}\colon \mathrm{Rect}^\circ \longrightarrow \{\pm 1\}$$

with the following properties:

(Sq) For any four distinct $r_1, r_2, r_1', r_2' \in \mathrm{Rect}^\circ$ with $r_1 * r_2 = r_1' * r_2'$, we have that

$$\mathcal{S}(r_1)\cdot\mathcal{S}(r_2) = -\mathcal{S}(r_1')\cdot\mathcal{S}(r_2').$$

(V) If $r_1, r_2 \in \mathrm{Rect}^\circ$ have the property that $r_1 * r_2$ is a vertical annulus, then

$$\mathcal{S}(r_1)\cdot\mathcal{S}(r_2) = -1.$$

(H) If $r_1, r_2 \in \mathrm{Rect}^\circ$ have the property that $r_1 * r_2$ is a horizontal annulus, then

$$\mathcal{S}(r_1)\cdot\mathcal{S}(r_2) = +1.$$

Theorem 4.2 *There is a sign assignment in the sense of Definition 4.1. Moreover, this sign assignment is essentially unique: if $\mathcal{S}_1$ and $\mathcal{S}_2$ are two sign assignments, then there is a function $f\colon \mathbf{S} \longrightarrow \{\pm 1\}$ so that for all $r \in \mathrm{Rect}^\circ(\mathbf{x},\mathbf{y})$, $\mathcal{S}_1(r) = f(\mathbf{x})\cdot f(\mathbf{y})\cdot\mathcal{S}_2(r)$.*

We turn to the proof of this theorem in Section 4.1. We can use the sign assignment from Theorem 4.2 to construct the chain complex over $\mathbb{Z}$ as follows. Fix a true sign assignment $\mathcal{S}$. Define $C^-(G)$ to be the free $\mathbb{Z}[U_1,\dots,U_n]$–module generated by $\mathbf{x} \in \mathbf{S}(G)$, endowed with Maslov grading and Alexander filtration as before. We endow this with the endomorphism

$$\partial^-_{\mathcal{S}}\colon C^-(G) \longrightarrow C^-(G)$$
$$\partial^-_{\mathcal{S}}(\mathbf{x}) = \sum_{\mathbf{y}\in\mathbf{S}} \sum_{r\in\mathrm{Rect}^\circ(\mathbf{x},\mathbf{y})} \mathcal{S}(r)\cdot U_1^{O_1(r)}\cdots U_n^{O_n(r)}\cdot\mathbf{y}.$$

We will check that this endomorphism gives the sign refinement of $C^-(G)$ needed in Theorem 1.2. In turn, the proof of that theorem involves reexamining the invariance proof from Section 3, and constructing sign refinements for the chain maps and homotopies used there. We turn to this task in Section 4.2. However, first we construct the sign assignments, proving Theorem 4.2.

4.1 The existence and uniqueness of sign assignments

Definition 4.3 A *thin rectangle* is a rectangle with width one. We denote the set of thin rectangles tRect; given $\mathbf{x}, \mathbf{y} \in \mathbf{S}$, we let $\mathrm{tRect}(\mathbf{x}, \mathbf{y}) = \mathrm{tRect} \cap \mathrm{Rect}(\mathbf{x}, \mathbf{y})$. For fixed $\mathbf{x}$ and $\mathbf{y}$ and $n > 2$, there can be at most one element in $\mathrm{tRect}(\mathbf{x}, \mathbf{y})$.

Sign assignments as in Theorem 4.2 are constructed in the following six steps.

(1) Define sign assignments in a more restricted sense, *sign assignments for the Cayley graph*. These are analogues of sign assignments defined only for thin rectangles supported in an $(n-1) \times (n-1)$ subsquare of the torus, satisfying a suitable restriction of Property (Sq) from Definition 4.1.

(2) Show that sign assignments for the Cayley graph satisfy a uniqueness property. Establish existence by giving an explicit formula. (It is also possible to give a more abstract existence argument, but the formula is needed in the next step.)

(3) Extend the formula to include all thin rectangles on the torus, and show that it satisfies, once again, axioms gotten by restricting Properties (Sq) and (V) to thin rectangles.

(4) Show that a sign assignment on thin rectangles can be extended uniquely to a function satisfying Properties (Sq) and (V), but not necessarily Property (H).

(5) Given the sign assignment on thin rectangles chosen in Step 3, establish a formula for the values of the function from Step 4 on empty rectangles supported in the $(n-1) \times (n-1)$ subsquare of the torus.

(6) With our choices of signs, use the explicit formulas from Step 5 to show that the function from Step 4 satisfies Property (H), thus giving a sign assignment in the sense of Definition 4.1.

Step 1: Define sign assignments on the Cayley graph We let $\Sigma = [0, n-1] \times [0, n-1]$ denote the $(n-1) \times (n-1)$–subsquare of the torus with the lower left corner at the origin.

Definition 4.4 Given $\mathbf{x}, \mathbf{y} \in \mathbf{S}$, a *thin rectangle in* Σ *from* $\mathbf{x}$ *to* $\mathbf{y}$ is a rectangle $r \in \mathrm{tRect}(\mathbf{x}, \mathbf{y})$ supported inside Σ. A thin rectangle in Σ *connects* $\mathbf{x}$ *and* $\mathbf{y}$ if it is a thin rectangle in Σ from $\mathbf{x}$ to $\mathbf{y}$ or from $\mathbf{y}$ to $\mathbf{x}$. The set of all thin rectangles in Σ is denoted tRect^*.

The set tRect^* has an interpretation in terms of a Cayley graph of the symmetric group in the following sense.

Consider the graph Γ whose vertices are elements in the symmetric group on n letters $\mathfrak{S}_n$, and whose edges are labeled by the $n-1$ adjacent transpositions $\{\tau_i\}_{i=1}^{n-1}$ in $\mathfrak{S}_n$, with an edge labelled τ_i connecting $\sigma_1, \sigma_2 \in \mathfrak{S}_n$ precisely when $\sigma_2 = \sigma_1 \cdot \tau_i$. When $\sigma_2 = \sigma_1 \cdot \tau_i$, we join σ_1 and σ_2 by exactly one edge, i.e. we do not draw an additional one for the relation $\sigma_1 = \sigma_2 \cdot \tau_i$. Γ is the Cayley graph of $\mathfrak{S}_n$ with respect to the generators τ_i.

There is a one-to-one correspondence between elements in $\mathfrak{S}_n$ and generators $\mathbf{S}$, which is obtained by viewing elements of $\mathbf{S}$ as graphs of permutations $\sigma_{\mathbf{x}}$. (To this end, we think of $\mathfrak{S}_n$ as permutations of the letters $\{0, \dots, n-1\}$.) This can be extended to a one-to-one correspondence between edges in the Cayley graph Γ and elements of tRect^*, sending a rectangle $r \in \mathrm{tRect}^*$ which connects $\mathbf{x}$ and $\mathbf{y}$ to the corresponding edge in the Cayley graph connecting $\sigma_{\mathbf{x}}$ and $\sigma_{\mathbf{y}}$:

$$\mathrm{tRect}^* \cong \mathrm{Edges}(\Gamma). \tag{8}$$

Definition 4.5 A *sign assignment on the Cayley graph* is a function

$$\mathcal{S}_0\colon \mathrm{Edges}(\Gamma) \longrightarrow \{\pm 1\}$$

with the following properties:

(Sq) If $\{e_1, \dots, e_4\}$ are four edges which form a square, then

$$\mathcal{S}_0(e_1) \cdots \mathcal{S}_0(e_4) = -1.$$

(Hex) If $\{e_1, \dots, e_6\}$ are six edges which form a hexagon, then

$$\mathcal{S}_0(e_1) \cdots \mathcal{S}_0(e_6) = 1.$$

Note that a square in the Cayley graph corresponds to two pairs of disjoint rectangles $r_i \in \mathrm{tRect}(\mathbf{x}_i, \mathbf{x}_{i+1})$ and $r_i' \in \mathrm{tRect}(\mathbf{x}_i', \mathbf{x}_{i+1}')$ for $i = 1, 2$, with $\mathbf{x}_1 = \mathbf{x}_1'$ and $\mathbf{x}_2 = \mathbf{x}_2'$, as pictured in Figure 22. Similarly, a hexagon in the Cayley graph corresponds to six thin rectangles $r_i \in \mathrm{tRect}(\mathbf{x}_i, \mathbf{x}_{i+1})$ and $r_i' \in \mathrm{tRect}(\mathbf{x}_i', \mathbf{x}_{i+1}')$ for $i = 1, 2, 3$ with $\mathbf{x}_1 = \mathbf{x}_1'$ and $\mathbf{x}_3 = \mathbf{x}_3'$, such that the union of the support of r_1, r_2, and r_3 is a rectangle (with width two), as pictured in Figure 23.

We can relate the above notion of a sign assignment to the earlier notion of sign assignments (Definition 4.1):

Lemma 4.6 *The restriction of a sign assignment in the sense of Definition 4.1 to the Cayley graph of $\mathfrak{S}_n$ is a sign assignment on the Cayley graph as defined in Definition 4.5.*

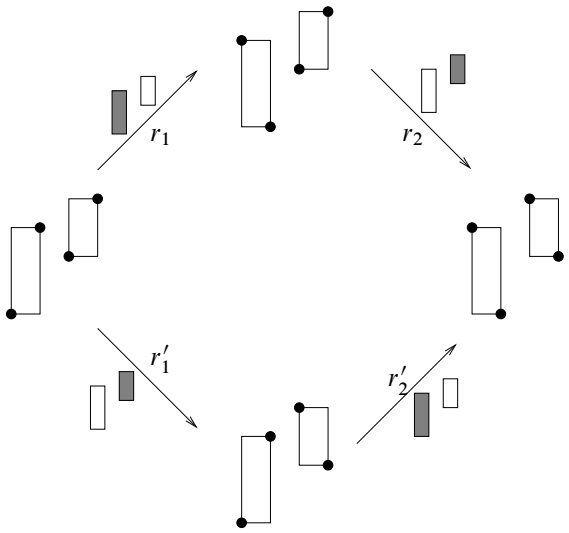

Figure 22: The rectangles in the square rule

Proof The first property follows from the corresponding property in Definition 4.1. For the second property, note that there is a rectangle r_4 of width 2 that cuts across a diagonal of the hexagon, cf Figure 23. Two applications of Property (Sq) (both involving r_4) now shows that the total number of sign changes around the hexagon must be even. □

Step 2: Signs assignments on the Cayley graph exist and are unique

Proposition 4.7 *A sign assignment on the Cayley graph exists and is unique up to equivalence given by changing the sign of the basis elements (as in Theorem 4.2).*

Proof Recall that for $A, B \subset \mathbb{R}^2$, we defined $\mathcal{I}(A, B)$ to be the number of pairs $(a_1, a_2) \in A$ and $(b_1, b_2) \in B$ with $a_1 < b_1$ and $a_2 < b_2$.

Given an edge of the Cayley graph, let $r \in \mathrm{tRect}^*(\mathbf{x}, \mathbf{y})$ denote its corresponding rectangle. Let $h = h(r)$ denote the height of the top edge of r (ie the four corners of r are (i, a), (i, h), $(i+1, a)$ and $(i+1, h)$, where (i, a) and $(i+1, h)$ belong to $\mathbf{x}$ and (i, h) and $(i+1, a)$ belong to $\mathbf{y}$).

We then define

$$\mathcal{S}(r) = (-1)^{\mathcal{I}(\mathbf{x}, \{(x_1, x_2) \in \mathbf{x} \mid x_2 \leq h(r)\})}. \tag{9}$$

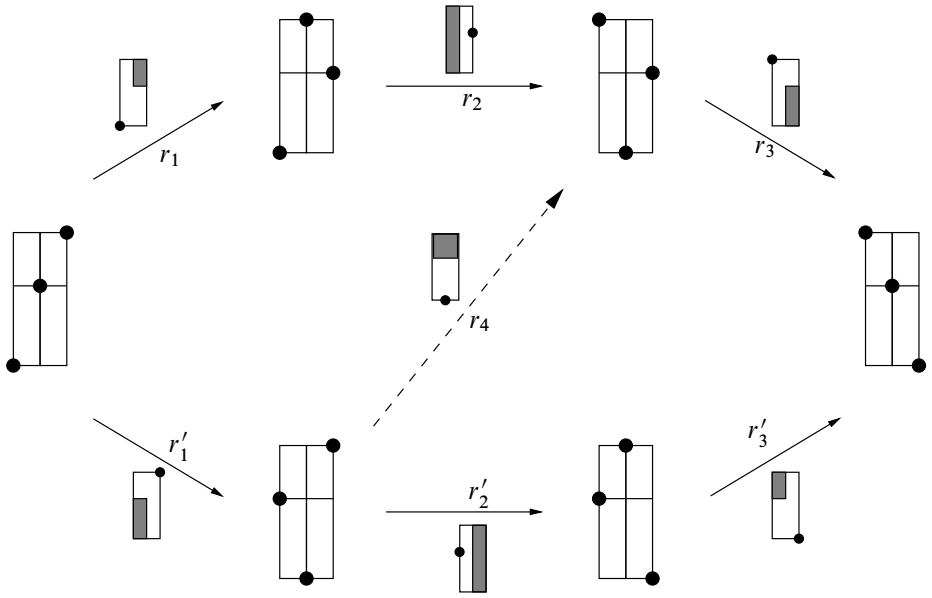

Figure 23: The rectangles in the hexagon rule. Two applications of the square rule (Sq) in Definition 4.1 give $\mathcal{S}(r_1)\cdot\mathcal{S}(r_2)=\mathcal{S}(r_1')\cdot\mathcal{S}(r_4)$ and $\mathcal{S}(r_4)\cdot\mathcal{S}(r_3)=\mathcal{S}(r_2')\cdot\mathcal{S}(r_3')$. These imply the hexagon rule $\mathcal{S}(r_1)\cdot\mathcal{S}(r_2)\cdot\mathcal{S}(r_3)=\mathcal{S}(r_1')\cdot\mathcal{S}(r_2')\cdot\mathcal{S}(r_3')$.

We check that each square anticommutes. To this end, observe that a square in the Cayley graph corresponds to four rectangles $r_1\in\mathrm{tRect}^*(\mathbf{x},\mathbf{y})$, $r_2\in\mathrm{tRect}^*(\mathbf{y},\mathbf{z})$, $r_1'\in\mathrm{tRect}^*(\mathbf{x},\mathbf{y}')$ and $r_2'\in\mathrm{tRect}^*(\mathbf{y}',\mathbf{z})$, where r_1 and r_2 have distinct corners and $r_1*r_2=r_1'*r_2'$ as in Figure 22. Number the rectangles so that $h(r_1)=h(r_2')<h(r_2)=h(r_1')$. It is easy to see that

$$\mathcal{S}(r_1)=\mathcal{S}(r_2'),\qquad \mathcal{S}(r_2)=-\mathcal{S}(r_1').$$

We can similarly check that a hexagon commutes. Consider the six thin rectangles corresponding to a hexagon in the Cayley graph $r_i\in\mathrm{tRect}(\mathbf{x}_i,\mathbf{x}_{i+1})$ and $r_i'\in\mathrm{tRect}(\mathbf{x}_i',\mathbf{x}_{i+1}')$ for $i=1,2,3$ with $\mathbf{x}_1=\mathbf{x}_1'$ and $\mathbf{x}_3=\mathbf{x}_3'$, as pictured in Figure 23. One can check that

$$\mathcal{S}(r_1)=\mathcal{S}(r_3'),\qquad \mathcal{S}(r_2)=\mathcal{S}(r_2'),\qquad \mathcal{S}(r_3)=\mathcal{S}(r_1'),$$

so that in particular

$$\mathcal{S}(r_1)\mathcal{S}(r_2)\mathcal{S}(r_3)=\mathcal{S}(r_1')\mathcal{S}(r_2')\mathcal{S}(r_3'),$$

as needed.

We will now prove uniqueness. Let $\mathcal{S}$ and $\mathcal{S}'$ be two sign assignments on $\mathfrak{S}_n$. Define a new function T on the Cayley graph by

$$T(\mathbf{x};\tau_i) = \mathcal{S}(\mathbf{x};\tau_i)\mathcal{S}'(\mathbf{x};\tau_i).$$

Then the product of T around any square or hexagon is equal to 1.

Let W be the Cayley complex of $\mathfrak{S}_n$: the 2–complex whose edges and vertices form the Cayley graph of $\mathfrak{S}_n$, and whose 2–cells are the squares connecting $\{\mathbf{x}, \mathbf{x}\tau_i, \mathbf{x}\tau_i\tau_j, \mathbf{x}\tau_j\}$ for $|i-j|>1$ and the hexagons connecting $\{\mathbf{x}, \mathbf{x}\tau_i, \mathbf{x}\tau_i\tau_{i+1}, \mathbf{x}\tau_i\tau_{i+1}\tau_i, \mathbf{x}\tau_{i+1}\tau_i, \mathbf{x}\tau_{i+1}\}$. Since these squares and hexagons (together with the relations $\tau_i^2 = 1$, which are suppressed in the definition of Γ) form a complete set of relations for $\mathfrak{S}_n$, the complex W is simply connected.

Now consider T as an element of $C^1(W;\{\pm 1\})$. The conditions on T are equivalent to saying that it is a cocycle: $\delta T = 0$. Since W is simply connected, there is therefore a function $f \in C^0(W;\{\pm 1\})$ so that $\delta f = T$. This function f gives the desired choice of signs on the basis. □

Remark 4.8 We could prove Proposition 4.7 without explicitly exhibiting the sign assignment: In general, suppose we have a 2–complex W and are looking for an assignment of ± 1 to the edges of the 2–complex so that the number of -1 signs is odd around a prescribed set of 2–cells. Such an assignment is unique (if it exists) if and only if $H^1(W;\{\pm 1\})$ is trivial, as in the proof of the Proposition. Furthermore, such an assignment exists if there is a 3–complex W' with W as its 2–skeleton so that $H^2(W';\{\pm 1\}) = 0$ and the set of faces with an odd number of -1 signs, considered as a 2–cocycle on W', is coclosed. In the case at hand, we can take W' to be the 3–skeleton of the *permutahedron* [17], which can be defined as the convex hull of the vectors obtained by permuting the coordinates of $(1, 2, \dots, n)$. This 3–skeleton is W with the following types of 3–cells attached:

- cubes corresponding to 3 disjoint transpositions, an $\mathfrak{S}_2 \times \mathfrak{S}_2 \times \mathfrak{S}_2 \subset \mathfrak{S}_n$;
- hexagonal prisms corresponding to $\mathfrak{S}_3 \times \mathfrak{S}_2 \subset \mathfrak{S}_n$; and
- truncated octahedra corresponding to $\mathfrak{S}_4 \subset \mathfrak{S}_n$.

(For the last case, note that the Cayley graph of $\mathfrak{S}_4$ is the boundary of a truncated octahedron.) In each case the number of squares on the boundary of the 3–cell is even, so an assignment of signs exists. The permutahedron is convex (hence contractible), so W' is 2–connected.

Step 3: Extend sign assignments to all thin rectangles in the torus

Definition 4.9 A *vertical sign assignment for thin rectangles* is a function

$$\mathcal{S}\colon \mathrm{tRect} \longrightarrow \{\pm 1\},$$

which satisfies the following properties:

(Sq) Given thin rectangles $r_1 \in \mathrm{tRect}(\mathbf{x}, \mathbf{y})$ and $r_2 \in \mathrm{tRect}(\mathbf{y}, \mathbf{z})$ with distinct corners, if we let $r_1' \in \mathrm{tRect}(\mathbf{x}, \mathbf{y}')$ and $r_2' \in \mathrm{tRect}(\mathbf{y}', \mathbf{z})$ be two other rectangles such that $r_1 * r_2 = r_1' * r_2'$, we have that

$$\mathcal{S}(r_1)\mathcal{S}(r_2) = -\mathcal{S}(r_1')\mathcal{S}(r_2').$$

(See Figure 22.)

(Hex) Given six thin rectangles $r_i \in \mathrm{tRect}(\mathbf{x}_i, \mathbf{x}_{i+1})$ and $r_i' \in \mathrm{tRect}(\mathbf{x}_i', \mathbf{x}_{i+1}')$ for $i = 1, 2, 3$ with $\mathbf{x}_1 = \mathbf{x}_1'$ and $\mathbf{x}_3 = \mathbf{x}_3'$, such that the union of the support of r_1, r_2, and r_3 is a rectangle (with width two), we have that

$$\mathcal{S}(r_1)\mathcal{S}(r_2)\mathcal{S}(r_3) = \mathcal{S}(r_1')\mathcal{S}(r_2')\mathcal{S}(r_3').$$

(See Figure 23.)

(V) If $r_1 \in \mathrm{tRect}(\mathbf{x}, \mathbf{y})$ and $r_2 \in \mathrm{tRect}(\mathbf{y}, \mathbf{x})$, then $\mathcal{S}(r_1) = -\mathcal{S}(r_2)$.

Proposition 4.10 *There is a vertical sign assignment for thin rectangles.*

Proof We extend Equation (9), as follows.

Note that $\mathcal{T}$ is obtained from Σ by adding one more row of squares, which are of the form

$$\{[i, i+1] \times [n-1, n]\}_{i=0}^{n-1},$$

and one more column of squares which are of the form

$$\{[n-1, n] \times [j, j+1]\}_{j=0}^{n-1}.$$

Consider a thin rectangle r in the torus. If r is contained in $\Sigma \subset \mathcal{T}$, then $\mathcal{S}(r)$ is as in Equation (9). If $r \in \mathrm{tRect}(\mathbf{x}, \mathbf{y})$ is a thin rectangle which is supported in the new column, but which is disjoint from the new row, so that it is of the form $[n-1, n] \times [a, b]$ with $0 \le a < b < n$ we define

$$\mathcal{S}(r) = (-1)^{\mathcal{I}\left(\mathbf{x}, \{(x_1, x_2) \in \mathbf{x} \mid x_2 \le a\}\right) + \mathcal{I}\left(\mathbf{x}, \{(x_1, x_2) \in \mathbf{x} \mid a < x_2 < b \text{ and } x_2 \text{ even}\}\right) + b}. \tag{10}$$

The thin rectangles not covered by the above two cases are those whose interiors meet the new row of squares. It is easy to see that for each such rectangle $r \in \mathrm{tRect}(\mathbf{x}, \mathbf{y})$,

there is a unique other thin rectangle $r' \in \mathrm{tRect}(\mathbf{y}, \mathbf{x})$, whose interior does not meet the row of squares, and hence whose sign is defined either by Equation (9) or by Equation (10). We then define

$$\mathcal{S}(r) = -\mathcal{S}(r').$$

This ensures that Property (V) in Definition 4.9 is satisfied. We must now verify Properties (Sq) and (Hex).

If all the thin rectangles involved are contained in Σ, then the two conditions were already checked in the proof of Proposition 4.7. Let us consider the cases when all the rectangles involved are disjoint from the new row, but at least one of them is supported in the new column, so that its sign is given by Equation (10).

Let us consider the square rule, with the support of r_1 (which is the same as the support of r_2') being in the last column, and the support of r_2 (the same as that of r_1') in Σ. Let $h(r_1)$ and $h(r_2)$ be the heights of the top edges of r_1 and r_2, respectively. If $h(r_1) \leq h(r_2)$, then, just as in the proof of Proposition 4.7, we have:

$$\mathcal{S}(r_1) = \mathcal{S}(r_2'), \qquad \mathcal{S}(r_2) = -\mathcal{S}(r_1').$$

If $h(r_1) \geq h = h(r_2)$, cf Figure 24, let $r_1 = [n-1, n] \times [a, b]$. If $h < a$, we obtain

$$\mathcal{S}(r_1) = -\mathcal{S}(r_2'), \qquad \mathcal{S}(r_2) = \mathcal{S}(r_1').$$

If $h \in (a, b)$ as in Figure 24 then, in comparing $\mathcal{S}(r_2)$ and $\mathcal{S}(r_1')$ using Equation (9),

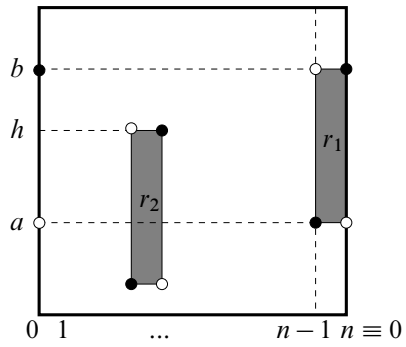

Figure 24: Anticommutation of width one rectangles. In computing the sign of the rectangle r_2 we use the black dot on the horizontal line of height a, while for the rectangle r_1', which has the same support, we use the white dot on the leftmost vertical edge.

there is a discrepancy coming from a pairs of points where the second point has coordinates $(n-1,a)$, and $h-a$ pairs of points where the first point has coordinates $(0,a)$. Therefore,

$$\mathcal{S}(r_2) = (-1)^h \mathcal{S}(r_1').$$

On the other hand, in comparing $\mathcal{S}(r_1)$ and $\mathcal{S}(r_2')$ using Equation (10), there can only be a discrepancy of one extra pair (two corners of r_2), which appears in case h is even. Thus,

$$\mathcal{S}(r_1) = (-1)^{h+1} \mathcal{S}(r_2'),$$

which implies that $\mathcal{S}(r_1)\mathcal{S}(r_2) = -\mathcal{S}(r_1')\mathcal{S}(r_2')$, as desired.

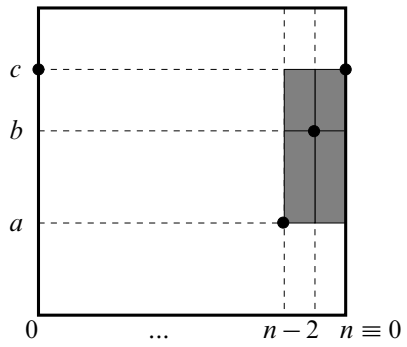

Figure 25: Hexagon rule in the last two columns. We compute the signs in terms of the generator $\mathbf{x}$; some of its components are the black dots shown here.

Let us now consider the hexagon rule, where the rectangles are supported in the last two columns, as in Figure 25. We denote by a, b, and c (with $a<b<c$) the three possible heights at which the relevant rectangles have their horizontal edges. We use the notation from Definition 4.9, where $\mathbf{x}$ is the initial generator of r_1 and r_1'. When applying formulas (9) and (10), we must keep in mind that the initial point of the relevant rectangle may differ from $\mathbf{x}$ (at heights a, b and c), cf Figure 25. After finding the relevant discrepancies, we can express everything in terms of $\mathbf{x}$:

$$\begin{aligned}
\mathcal{S}(r_1) &= (-1)^{\mathcal{I}\left(\mathbf{x},\{(x_1,x_2)\in\mathbf{x}\mid x_2\le b\}\right)+\mathcal{I}\left(\mathbf{x},\{(x_1,x_2)\in\mathbf{x}\mid b<x_2<c,\ x_2 \text{ even})\}\right)+c}\\
\mathcal{S}(r_1') &= (-1)^{\mathcal{I}\left(\mathbf{x},\{(x_1,x_2)\in\mathbf{x}\mid x_2\le b\}\right)}\\
\mathcal{S}(r_2) &= (-1)^{\mathcal{I}\left(\mathbf{x},\{(x_1,x_2)\in\mathbf{x}\mid x_2\le a\}\right)+\mathcal{I}\left(\mathbf{x},\{(x_1,x_2)\in\mathbf{x}\mid a<x_2<c\}\right)+1}\\
\mathcal{S}(r_2') &= (-1)^{\mathcal{I}\left(\mathbf{x},\{(x_1,x_2)\in\mathbf{x}\mid x_2\le a\}\right)+\mathcal{I}\left(\mathbf{x},\{(x_1,x_2)\in\mathbf{x}\mid a<x_2<c,\ x_2 \text{ even})\}\right)+b+c+1}\\
\mathcal{S}(r_3) &= (-1)^{\mathcal{I}\left(\mathbf{x},\{(x_1,x_2)\in\mathbf{x}\mid x_2\le a\}\right)+\mathcal{I}\left(\mathbf{x},\{(x_1,x_2)\in\mathbf{x}\mid a<x_2\le b,\ x_2 \text{ even})\}\right)+b}\\
\mathcal{S}(r_3') &= (-1)^{\mathcal{I}\left(\mathbf{x},\{(x_1,x_2)\in\mathbf{x}\mid x_2\le a\}\right)+\mathcal{I}\left(\mathbf{x},\{(x_1,x_2)\in\mathbf{x}\mid a<x_2<c\}\right)}
\end{aligned}$$

Putting these relations together, we obtain the required identity:

$$\mathcal{S}(r_1)\mathcal{S}(r_2)\mathcal{S}(r_3) = \mathcal{S}(r_1')\mathcal{S}(r_2')\mathcal{S}(r_3').$$

There is a similar computation that needs to be done for the hexagon rule when the rectangles are supported in the first and the last column. We leave this case as an exercise for the reader.

Finally, we need to check the square and the hexagon rule when some of the rectangles involved wrap vertically around the torus, ie their support has nontrivial intersection with the horizontal line l of height $n-\frac{1}{2}$. We call such rectangles *vertically wrapped*. For the square rule, either two or all four of the four rectangles involved are vertically wrapped. For the hexagon rule, exactly four out of the six rectangles involved are vertically wrapped. The square and the hexagon rules now follow from the corresponding ones when we replace the vertically wrapped rectangles r with their counterparts r' such that $r * r'$ are vertical annuli. Indeed, from equations of the form $\mathcal{S}(r) = -\mathcal{S}(r')$ we always pick up an even number of minus signs (either two or four), so the overall signs are unchanged. □

Step 4: Extend vertical sign assignment to all empty rectangles We weaken the notion of sign assignments from Definition 4.1 a little.

Definition 4.11 A *vertical sign assignment* is a function

$$\mathcal{S}\colon \mathrm{Rect}^\circ \longrightarrow \{\pm 1\}$$

which satisfies Properties (Sq) and (V) from Definition 4.1. Sometimes, we call this a *vertical sign assignment on all empty rectangles*, to distinguish it from the seemingly weaker vertical sign assignments on thin rectangles.

Our goal in this step is to show that a vertical sign assignment for thin rectangles can be uniquely extended to a vertical sign assignment on all empty rectangles.

Given a vertical sign assignment for thin rectangles $\mathcal{S}_0$, we define an extension

$$\mathcal{S}\colon \mathrm{Rect}^\circ \longrightarrow \{\pm 1\},$$

by extending the definition inductively on the width w of the rectangle. Explicitly, if $r \in \mathrm{Rect}^\circ(\mathbf{x}, \mathbf{y})$ is a rectangle with width 1, then $\mathcal{S}(r) = \mathcal{S}_0(r)$. Suppose next that $\mathcal{S}$ is defined for all rectangles of width less than w for some $w > 1$. Given $r \in \mathrm{Rect}^\circ(\mathbf{x}, \mathbf{y})$ of width w, there is exactly one rectangle r_1 ending at $\mathbf{x}$ with width one whose upper left corner coincides with the lower left corner of r, as in the first diagram in Figure 26.

Then $r_1 * r$ has an alternate decomposition as $r_2 * r_3$, where r_2 has width $w-1$, and r_3 has width one. We can then define

$$\mathcal{S}(r) = -\mathcal{S}_0(r_1)\mathcal{S}(r_2)\mathcal{S}_0(r_3). \tag{11}$$

(The right hand side is defined, since the width of r_2 is $w-1$.)

We verify that $\mathcal{S}$ is a vertical sign assignment in the sense of Definition 4.1, which we do in stages.

Lemma 4.12 *Suppose that we have four rectangles $r_1, r_2, r_1', r_2' \in \mathrm{Rect}^\circ$ with $r_1 * r_2 = r_1' * r_2'$, where r_1 and r_2' have width one, and r_1 and r_2 share exactly one corner. Then*

$$\mathcal{S}(r_1)\mathcal{S}(r_2) = -\mathcal{S}(r_1')\mathcal{S}(r_2'). \tag{12}$$

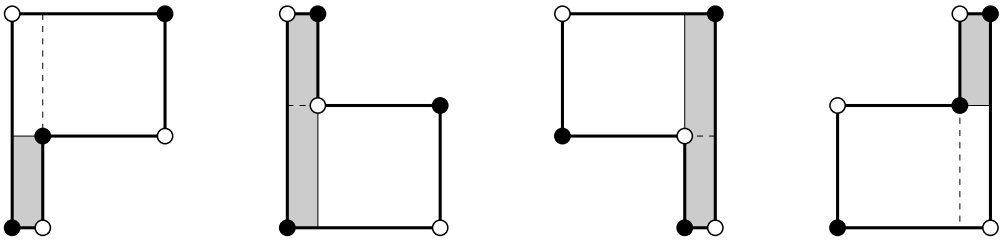

Figure 26: The four cases in Lemma 4.12. Each case is a decomposition $r_1 * r_2$, where r_1 is shown shaded. The dotted line gives the alternate decomposition $r_1' * r_2'$.

Proof There are four cases in the proof, as illustrated in Figure 26. In the first case, when the upper-left corner of r_1 is the lower-left corner of r_2, the conclusion holds by the definition of the extension of $\mathcal{S}$, Equation (11). In the second case, when the lower-right corner of r_1 is the lower-left corner of r_2, the conclusion follows from Property (V) of the vertical sign assignment: if we relate both thin rectangles, r_1 and r_2', to the other rectangle in the same vertical column, we get the same rectangles as in the previous case.

Otherwise, we will prove the result by induction on the maximum of the widths of r_2 and r_1'. We treat the base case first, where this maximum is equal to two. There are two cases:

- If the upper-left corner of r_1 is the upper-right corner of r_2, Equation (12) follows from Property (Sq) of the vertical sign assignment, the hexagon rule, and the definition of $\mathcal{S}$.

- If the lower-right corner of r_1 is the upper-right corner of r_2, Equation (12) follows from Property (V) applied to the case above.

We now treat the induction on the maximum of the widths of r_2 and r_1', which we may assume is greater than two. If the upper-left corner of r_1 is the upper-right corner of r_2 (the third case in the figure), we may again apply Property (V) to change to the last case. So we may assume that the lower right corner of r_1 is shared with the upper right corner of r_2. Since the width of r_2 is greater than two, we can find a thin rectangle r_0 to $\mathbf{x}$, the initial point of r_1, disjoint from r_1 with the property that r_0 and r_2 share a corner. We consider now the composite $p = r_0 * r_1 * r_2$. We organize the various decompositions of p into the graph in Figure 27. Each edge corresponds to some rectangle in some decomposition of p, and each vertex corresponds to one of the various elements of $\mathbf{S}$ which can be connected by rectangles in some decomposition of p. Each face corresponds to two decompositions of p which have some rectangle in common. Thus we have organized the decompositions of p in a cube. A face is said to *anticommute* if the product of the signs associated to its four edges is -1. Our goal is to show that the face belonging to the two decompositions $r_0 * r_1' * r_2'$ and $r_0 * r_1 * r_2$ anticommutes. To see this, we observe that the other five faces anticommute: two of them do, as they correspond to rearranging two disjoint thin rectangles (Property (Sq) of the vertical sign assignment for thin rectangles), two of the faces anticommute by the definition of the extension, Equation (11), and the fifth face anticommutes by induction on the width of r_2. □

Lemma 4.13 *Suppose that we have four distinct rectangles $r_1, r_2, r_1', r_2' \in \mathrm{Rect}^\circ$ with $r_1 * r_2 = r_1' * r_2'$, where r_1 and r_2' have width one, then*

$$\mathcal{S}(r_1)\mathcal{S}(r_2) = -\mathcal{S}(r_1')\mathcal{S}(r_2').$$

Proof We prove the result by induction on the width of r_2.

The case where r_2 shares exactly one corner with r_1 was handled in Lemma 4.12, so we may assume the corners of r_1 and r_2 are distinct. According to Property (V), we can assume without loss of generality that the supports of r_1 and r_2 are disjoint. If r_2 also has width one, we are done by Property (Sq).

Otherwise, we proceed by induction in a way similar to the last case of Lemma 4.12. Specifically, we can find another width one rectangle r_0 ending at the initial point of r_1 which shares one corner with r_2. Consider the polygon $p = r_0 * r_1 * r_2$. We can once again organize the various decompositions of this polygon into a cube. In the case where r_0 and r_1 are disjoint, two of the faces anticommute according to Property (Sq)

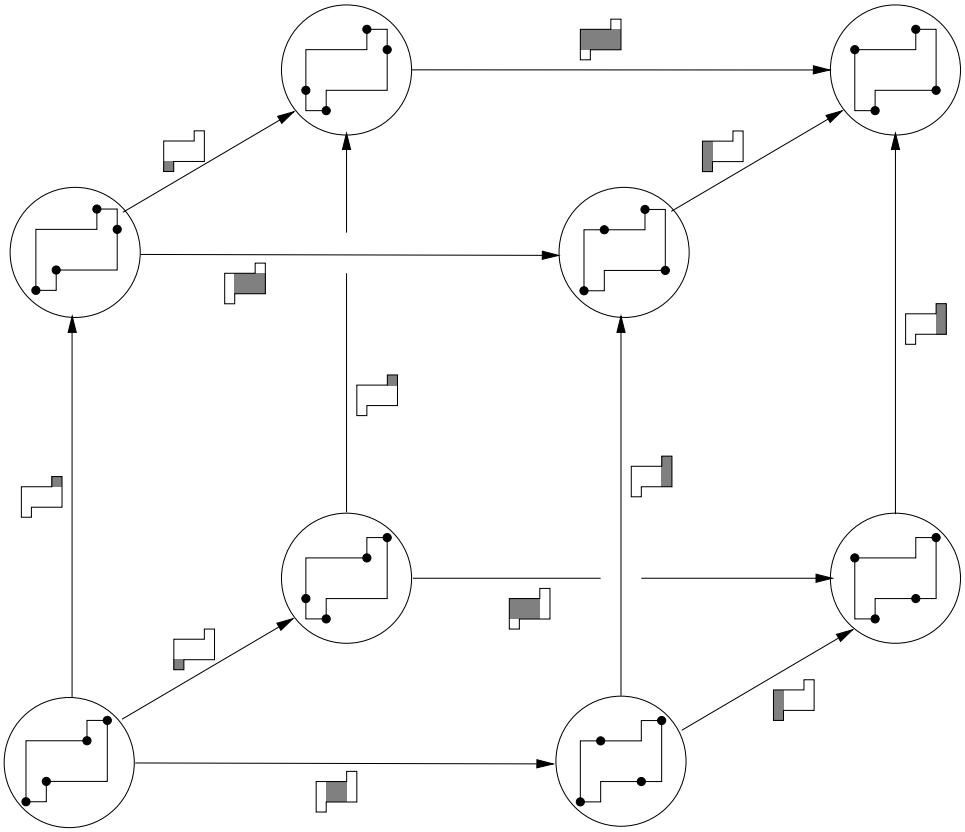

Figure 27: The cube of decompositions in Lemma 4.12. Each vertex is a configuration (shown by black dots); each edge is a rectangle (shown shaded). The different decompositions are the different ways to go from the lower-left corner to the upper-right corner following three edges. We are trying to show that the back face anticommutes; the front face is the inductive case.

of $\mathcal{S}_0$, two anticommute by Lemma 4.12, and a fifth anticommutes by induction on the width of r_2. Thus, the sixth must anticommute, as well.

In the other case, where r_0 and r_1 share a corner (which we need to consider only when the width of r_2 is three), four of the squares in the cube anticommute according to Lemma 4.12. The fifth involves a domain $r_1' * r_2'$ where r_1' and r_2' are both rectangles of width two, with r_1' containing r_1 and r_2' contained in r_2. To show that this last face anticommutes, we can find another rectangle r_0' that can be precomposed with these two to give a domain $r_0' * r_1' * r_2'$. The various decompositions of this domain can again be arranged into a cube, in which four of the squares anticommute by Lemma 4.12, the

fifth anticommutes by Property (Sq), and the sixth is the face with domain $r_1' * r_2'$, as desired. □

Proposition 4.14 *Given any four empty rectangles $r_1, r_2, r_1', r_2' \in \mathrm{Rect}^\circ$ with $r_1 * r_2 = r_1' * r_2'$, we have that*

$$\mathcal{S}(r_1)\mathcal{S}(r_2) = -\mathcal{S}(r_1')\mathcal{S}(r_2').$$

Proof Using suitably placed width one rectangles as before, we can narrow r_1 and/or r_2 by induction, until one or the other has width one and hence is covered by Lemma 4.13. □

In effect, the above proposition shows that a vertical sign assignment for thin rectangles $\mathcal{S}_0$ can be canonically extended to a vertical sign assignment for arbitrary rectangles $\mathcal{S}$ in the sense of Definition 4.11.

Step 5: Signs for rectangles supported in Σ In Proposition 4.10 we constructed a vertical sign assignment for thin rectangles. According to Step 4 above, this gives a vertical sign assignment $\mathcal{S}\colon \mathrm{Rect}^\circ \to \{\pm 1\}$ on all empty rectangles. We aim to give an explicit formula for $\mathcal{S}(r)$ in the case when r is supported in the subsquare $\Sigma = [0, n-1] \times [0, n-1]$.

Proposition 4.15 *Let $r = [a, b] \times [c, d] \in \mathrm{Rect}^\circ(\mathbf{x}, \mathbf{y})$, with $0 \le a < b \le n-1$ and $0 \le c < d \le n-1$. Denote by $\mathcal{D}(r)$ the number of points $(x_1, x_2) \in \mathbf{x}$ which lie strictly below r, ie $x_1 \in (a, b)$ and $x_2 \in [0, c)$, cf Figure 28. Then*

$$\mathcal{S}(r) = (-1)^{\mathcal{I}\left(\mathbf{x}, \{(x_1,x_2)\in\mathbf{x} \mid x_2 \le d\}\right) + \mathcal{D}(r)\cdot\left(\mathcal{I}\left(\mathbf{x}, \{(x_1,x_2)\in\mathbf{x} \mid c < x_2 \le d\}\right)+1\right)}. \tag{13}$$

Proof We use induction on the width w of r. When r has width one, (13) is just the formula (9).

Assume $w > 1$. We distinguish two cases, according to the position of the point $(a+1, l) \in \mathbf{x}$ in the vertical line $x_1 = a+1$. When it is below the support of r, the inductive step follows by using the anticommutation rule for a decomposition of the form $r * r_1 = r_2 * r'$, where the rectangles r_1 and r_2 have width one and lie in the ath column, so that their signs can be computed using (9), and r' has width $w-1$. The latter case is similar, but the decompositions are of the form $r_1 * r = r' * r_2$. In both cases the computations are straightforward. □

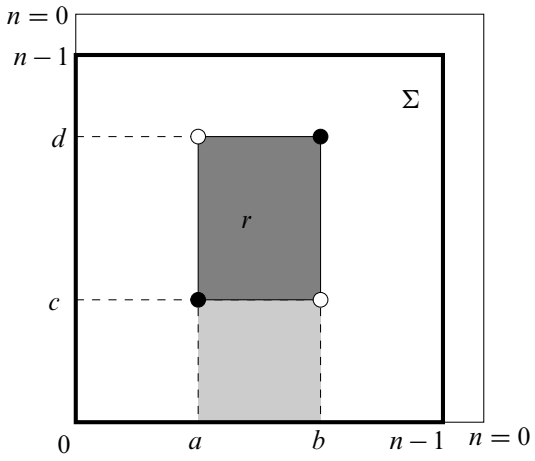

Figure 28: Signs for rectangles on the square. The subsquare Σ of the torus is shown bordered by the thick lines. Inside we have a rectangle r. The quantity $\mathcal{D}(r)$ counts the number of points in $\mathbf{x}$ which lie in the lightly shaded region below r.

Step 6: From vertical sign assignments to true sign assignments

Lemma 4.16 *Suppose $\mathcal{S}$ is a vertical sign assignment. Then there is some function $\rho\colon \{1,\dots,n\}\to\{\pm 1\}$ with the property that if r_1 and r_2 are two height one rectangles, with $r_1 * r_2$ connecting some $\mathbf{x}\in\mathbf{S}$ to itself, then $\mathcal{S}(r_1)\mathcal{S}(r_2)=\rho(i)$, where the support of $r_1 * r_2$ consists of the i th row.*

Proof We claim first that when the support of $r_1 * r_2$ is a horizontal annulus $[0,n)\times[i,i+1]$, $\mathcal{S}(r_1)\mathcal{S}(r_2)$ depends only on i and the components of $\mathbf{x}$ in the ith and $(i+1)$st rows. This follows readily from the square rule; if $\mathbf{x}$ and $\mathbf{x}'$ are two generators whose components on the ith and $(i+1)$st agree and there is some $r\in\mathrm{Rect}^\circ(\mathbf{x},\mathbf{x}')$, then two applications of the square rule establish the claim. More generally, we can always get between two generators $\mathbf{x}$ and $\mathbf{x}'$ that agree on these two rows by a sequence of squares whose corners are disjoint from the ith row, verifying the claim.

We next verify that if $r_1\in\mathrm{Rect}^\circ(\mathbf{x},\mathbf{y})$ are $r_2\in\mathrm{Rect}^\circ(\mathbf{y},\mathbf{x})$ are two rectangles which, together, form the ith row, then $\mathcal{S}(r_1)\mathcal{S}(r_2)$ depends only on the row in which r_1 and r_2 are supported. To this end, observe that there is another rectangle $r_3\in\mathrm{Rect}(\mathbf{x},\mathbf{z})$ with height one supported in the row $i+1$, and $r_4\in\mathrm{Rect}(\mathbf{z},\mathbf{x})$. We now claim that $r_1*r_2*r_3*r_4$ differs from another decomposition $r_1'*r_2'*r_3'*r_4'$ in two steps, in such a way that the supports of r_3 and r_4 agree with those of r_3' and r_4' (hence together

they occupy the $(i+1)$st row), and the supports of r_1' and r_2' also occupy the ith row, but the support of r_1' and r_2' are different from the supports of r_1 and r_2. Thus, it follows from the square rule that $\mathcal{S}(r_1)\mathcal{S}(r_2) = \mathcal{S}(r_1')\mathcal{S}(r_2')$. It is easy to see that any two $r_1 \in \mathrm{Rect}^\circ(\mathbf{x},\mathbf{y})$ and $r_2 \in \mathrm{Rect}^\circ(\mathbf{y},\mathbf{x})$ which occupy the ith row can be connected by a finite sequence of such moves. □

We now specialize to the vertical sign assignment $\mathcal{S}$ constructed (in Step 3) from the vertical sign assignment on thin rectangles exhibited in Proposition 4.10, which was based on the formulas (9) and (10). We claim that the function ρ from Lemma 4.16 is identically 1.

Lemma 4.17 *$\rho(i) = 1$ for $i = 1, \dots, n-1$.*

Proof It suffices to check that $\mathcal{S}(r_1)\mathcal{S}(r_2) = 1$ when $r_1 \in \mathrm{Rect}(\mathbf{x},\mathbf{y})$ is of the form $[0, n-1]\times[i-1, i]$, and $r_2 \in \mathrm{Rect}(\mathbf{y},\mathbf{x})$ is the square $[n-1, n]\times[i-1, i]$ in the last column, cf Figure 29. Proposition 4.15 gives

$$\mathcal{S}(r_1) = (-1)^{\mathcal{I}\left(\mathbf{x},\{(x_1,x_2)\in\mathbf{x} \mid x_2\le i\}\right)+i+1} = (-1)^{\mathcal{I}\left(\mathbf{x},\{(x_1,x_2)\in\mathbf{x} \mid x_2\le i-1\}\right)+1}.$$

On the other hand, from Equation (10) we get

$$\mathcal{S}(r_2) = (-1)^{\mathcal{I}\left(\mathbf{y},\{(y_1,y_2)\in\mathbf{y} \mid y_2\le i-1\}\right)+i} = (-1)^{\mathcal{I}\left(\mathbf{x},\{(x_1,x_2)\in\mathbf{x} \mid x_2\le i-1\}\right)+1}. \qquad \square$$

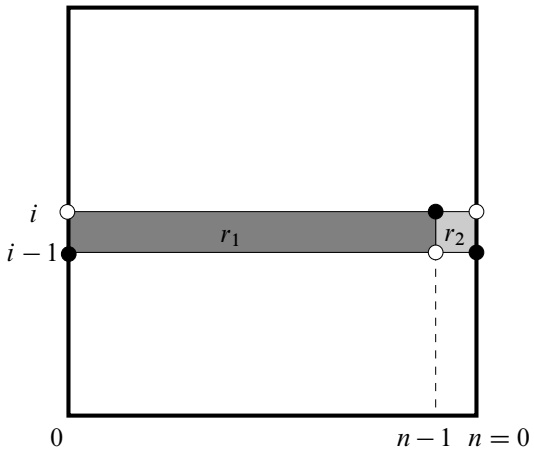

Figure 29: Computing $\rho(i)$ for $i \le n$. The generator $\mathbf{x}$ is shown in black dots, and $\mathbf{y}$ in hollow dots. When computing the sign of r_2 with formula (10), we need to use $\mathbf{y}$ in the place of $\mathbf{x}$.

To check that $\rho(n) = 1$, we first prove the following:

Lemma 4.18 *Let $k \in 1, 2, \dots, n-1$. Denote by $\mathbf{x}_k, \mathbf{y}_k \in \mathbf{S}$ the configurations*

$$\mathbf{x}_k = \{(i, n-1-i) | 0 \le i < k\} \cup \{(k, 0)\} \cup \{(i, n-i) | k < i < n\}$$
$$\mathbf{y}_k = \{(0, 0)\} \cup \{(i, n-1-i) | 1 \le i < k\} \cup \{(k, n-1)\} \cup \{(i, n-i) | k < i < n\}.$$

Let $r_k \in \mathrm{Rect}(\mathbf{x}_k, \mathbf{y}_k)$ be the rectangle of width k and height 1 supported in the last row, cf Figure 30. Then, its sign is given by

$$\mathcal{S}(r_k) = (-1)^n.$$

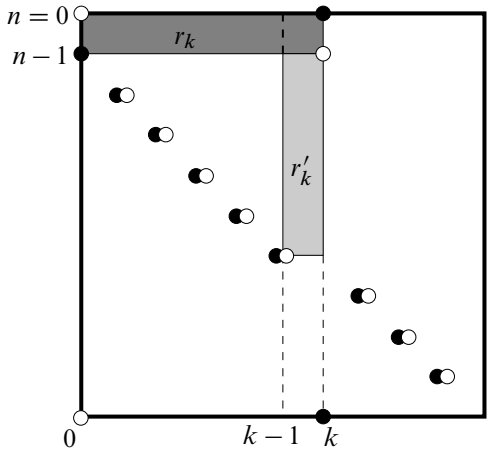

Figure 30: Computing $\rho(n)$. The rectangles r_k is darkly shaded, and r'_k lightly shaded. The generators $\mathbf{x}_k$ and $\mathbf{y}_k$ are represented by the black and white dots, respectively. There is an alternate decomposition of $r_k * r'_k$ given by cutting along the dashed line in the top row.

Proof Induction on k. When $k = 1$, the sign of r_1 is minus the sign of a thin rectangle of width one and height $n-1$ supported in the first column. The latter can be computed using Equation (9), which gives the answer $(-1)^{n-1}$; therefore, $\mathcal{S}(r_1) = (-1)^n$.

For $k > 1$, let $r'_k \in \mathrm{Rect}(\mathbf{y}_k, \mathbf{z}_k)$ be the rectangle of width one supported in the kth column, cf Figure 30. Its sign is $(-1)^{n+k}$ by formula (9). The domain $r_k * r'_k$ has an alternate decomposition as $p_k * r_{k-1}$, where the rectangle p_k is again supported in the kth column and has a counterpart p'_k such that $p_k * p'_k$ is a vertical annulus. Formula (9) gives $\mathcal{S}(p'_k) = (-1)^{n+k}$, so that $\mathcal{S}(p_k) = (-1)^{n+k+1}$. The inductive step now follows form the anticommutation relation $\mathcal{S}(r_k)\mathcal{S}(r'_k) = -\mathcal{S}(r_{k-1})\mathcal{S}(p_k)$. □

Lemma 4.19 $\rho(n) = 1$.

Proof By the previous lemma, the sign of $r_{n-1} \in \mathrm{Rect}(\mathbf{x}_{n-1}, \mathbf{y}_{n-1})$ is $(-1)^n$. There is also a thin vertical rectangle $r' \in \mathrm{Rect}(\mathbf{x}_{n-1}, \mathbf{y}_{n-1})$ supported in the last column, whose sign is $(-1)^{n-1}$ by (10). For the little square $r'' \in \mathrm{Rect}(\mathbf{y}_{n-1}, \mathbf{x}_{n-1})$ supported in the top right corner we have $\mathcal{S}(r'') = -\mathcal{S}(r')$. Since $r_{n-1} * r''$ is a horizontal annulus and $\mathcal{S}(r'') = \mathcal{S}(r_{n-1})$, we get $\rho(n) = 1$, as desired. □

We can now complete Step 6:

Proof of Theorem 4.2 We proved existence of vertical sign assignments in Step 4. According to Lemma 4.16, the resulting signs give a true sign assignment provided that the function ρ defined there is identically one. This was checked in Lemmas 4.17 and 4.19.

To see uniqueness, let $\mathcal{S}_1$ and $\mathcal{S}_2$ be two true sign assignments. We can restrict them to the square Σ and get sign assignments on the Cayley graph (cf Lemma 4.6). Using the uniqueness part of result of Proposition 4.7, we obtain a function $f\colon \mathbf{S} \longrightarrow \{\pm 1\}$ such that the composite

$$\mathcal{B}\colon \mathrm{Rect}^\circ \longrightarrow \{\pm 1\}, \quad \mathcal{B}(r) = f(\mathbf{x}) \cdot f(\mathbf{y}) \cdot \mathcal{S}_1(r) \cdot \mathcal{S}_2(r) \text{ for } r \in \mathrm{Rect}^\circ(\mathbf{x}, \mathbf{y})$$

satisfies the following properties:

- $\mathcal{B}(r) = 1$ for any rectangle of width one supported in Σ.
- (*Commutation rule*) $\mathcal{B}(r_1) \cdot \mathcal{B}(r_2) = \mathcal{B}(r_1') \cdot \mathcal{B}(r_2')$ whenever $r_1, r_2, r_1', r_2' \in \mathrm{Rect}^\circ$ are distinct and satisfy $r_1 * r_2 = r_1' * r_2'$.
- $\mathcal{B}(r_1) = \mathcal{B}(r_2)$ if $r_1 * r_2$ is a vertical annulus.
- $\mathcal{B}(r_1) = \mathcal{B}(r_2)$ if $r_1 * r_2$ is a horizontal annulus.

We claim that $\mathcal{B}$ is identically 1. Indeed, the third property above implies that $\mathcal{B}(r) = 1$ whenever r has width one and is not supported in the last column. The fourth property implies that $\mathcal{B}(r) = 1$ when r is a square of side length one supported in the last column; the same must be true for all width one rectangles in the last column by induction on height, applying the commutation rule. Thus $\mathcal{B}$ takes the value one on all vertical thin rectangles. The fact that $\mathcal{B}(r) = 1$ for all r follows by induction on width, again using the commutation rule. This shows that f satisfies the property required in the statement of Theorem 4.2. □

4.2 Properties of the sign-refined chain complex and the proof of Theorem 1.2

Proposition 4.20 *Let $\mathcal{S}$ be a sign assignment. The $\mathbb{Z}[U_1,\dots,U_n]$–module $C^-(G)$, endowed with the endomorphism $\partial^-_{\mathcal{S}}$, is a chain complex. Moreover, if $\mathcal{S}_1$ and $\mathcal{S}_2$ are two different sign assignments, then there is an isomorphism of chain complexes $(C^-,\partial^-_{\mathcal{S}_1})\cong(C^-,\partial^-_{\mathcal{S}_2})$.*

Proof In the expression $\partial_{\mathcal{S}}\circ\partial_{\mathcal{S}}(\mathbf{x})$, terms can be paired off as in the proof of Proposition 2.10. These terms cancel, according to the axioms on $\mathcal{S}$.

Suppose we are given sign assignments $\mathcal{S}_1$ and $\mathcal{S}_2$. Consider the map

$$\Phi\colon (C^-(G),\partial_{\mathcal{S}_1})\longrightarrow(C^-(G),\partial_{\mathcal{S}_2})$$

defined by $\Phi(\mathbf{x})=f(\mathbf{x})\cdot\mathbf{x}$, where f is the function from Theorem 4.2. It is straightforward to see that Φ is an isomorphism of chain complexes. □

Other algebraic properties from Section 2 have straightforward generalizations to this context. For example:

Lemma 4.21 *Suppose that O_i and O_k correspond to the same component of $\vec{L}$. Then multiplication by U_i is filtered chain homotopic to multiplication by U_k.*

Proof The chain homotopy from Lemma 4.21 works to establish the present lemma. It is important here that for $\mathbf{x}\in\mathbf{S}$, and r_1 and r_2 are the decompositions of the horizontal annulus containing X_j, and r_3 and r_4 are the analogous decompositions of the vertical annulus, then $\mathcal{S}(r_1)\mathcal{S}(r_2)=-\mathcal{S}(r_3)\mathcal{S}(r_4)$, but this is ensured by the axioms of the sign assignment. This ensures that U_i is chain homotopic to U_k (rather than, say, $-U_k$). □

Again, we view the chain complex $(C^-(G),\partial^-)$ as a module over $\mathbb{Z}[U_1,\dots,U_\ell]$ where the $\{U_i\}_{i=1}^{\ell}$ correspond to the ℓ components of our link $\vec{L}$. As before, we have the following:

Proposition 4.22 *Suppose that the oriented link $\vec{L}$ has ℓ components. Choose an ordering of $\mathbb{O}=\{O_i\}_{i=1}^{n}$ so that for $i=1,\dots,\ell$, O_i corresponds to the i th component of $\vec{L}$. Then the filtered chain homotopy type of $C^-(G)$, viewed as a chain complex over $\mathbb{Z}[U_1,\dots,U_\ell]$, is independent of the ordering of $\mathbb{O}$. Then $\widehat{HL}(G)$ and $\widetilde{HL}(G)$ are finitely generated Abelian groups. Moreover,*

$$H_*(\widetilde{CL}(G))\cong\widehat{HL}(G)\otimes\bigotimes_{i=1}^{\ell}V_i^{\otimes(n_i-1)},$$

where V_i is the two-dimensional vector space spanned by two generators, one in zero Maslov and Alexander multigradings, and the other in Maslov grading minus one and Alexander multigrading corresponding to minus the i th basis vector.

Proof This is a routine adaptation of Proposition 2.12, Lemma 2.13, and Proposition 2.15 from Section 2. □

We now turn to the proof of the main theorem, Theorem 1.2, with signs.

We adopt the strategy from Section 3; however, we must specify the signs used in defining our various chain maps and chain homotopies, and verify that they are indeed chain maps and chain homotopies with appropriate signs. As in Section 3, we begin with the case of commutation.

We adopt notation from Section 3.1. Consider the pentagons $\mathrm{Pent}_{\beta\gamma}(\mathbf{x},\mathbf{y})$ used there. Straightening out the $\beta\cap\gamma$–corner of the pentagons naturally induces rectangles in G. (We could have in fact defined the map in Section 3.1 as counts of rectangles, where the O's and X's in the central column are moving, but then it would have been a little confusing to write down exactly when they are counted with powers of the U's.)

Formally, we obtain a "straightening map"

$$e\colon \mathrm{Pent}_{\beta\gamma}(\mathbf{x},\mathbf{y}) \longrightarrow \mathrm{Rect}(\mathbf{x},\mathbf{y}'),$$

where $\mathbf{y}'$ is the generator corresponding to $\mathbf{y}$, where we slide horizontally from the γ back to the β component. Clearly, the image of e consists of rectangles with a vertical segment along β. There are two possibilities: either the rectangle lies to the right of this vertical segment (ie the segment is a left edge of the rectangle), or it lies to the left of this vertical segment. In the first case, we say the pentagon is a *right pentagon*, in the latter, we say it is a *left pentagon*. In Figure 8, the one pictured on the left is a left pentagon, and the one on the right is a right pentagon.

We define

$$\Phi_{\beta\gamma}(\mathbf{x}) = \sum_{\mathbf{y}\in\mathbf{S}(H)} \sum_{\substack{p\in\mathrm{Pent}_{\beta\gamma}(\mathbf{x},\mathbf{y})\\ \mathbf{x}\cap\mathrm{Int}(p)=\varnothing}} \epsilon(p)\cdot U_1^{O_1(p)}\cdots U_n^{O_n(p)}\cdot\mathbf{y},$$

where
$$\epsilon(p) = \begin{cases} \mathcal{S}(e(p)) & \text{if } p \text{ is a left pentagon}\\ -\mathcal{S}(e(p)) & \text{if } p \text{ is a right pentagon.}\end{cases}$$

We obtain the following analogue of Lemma 3.1:

Lemma 4.23 *The map $\Phi_{\beta\gamma}$ is a filtered antichain map, ie*

$$\partial \circ \Phi_{\beta\gamma} + \Phi_{\beta\gamma} \circ \partial = 0.$$

Proof Again, the proof follows from the proof of Lemma 3.1. In fact, the fact that the terms cancel in pairs typically follows from the same pairing which we see in Proposition 4.20. There are two cases which look different, though. One of these corresponds to the two decompositions pictured as in Figure 9. After projecting via e, both decompositions of the composite region in this case correspond to the same decomposition of the composite region into two rectangles. However, in one case, the rectangle corresponding to the pentagon is on the left, in the other, it is on the right. Thus, the signs given by ϵ are opposite. The other case, the rotation by 180° of Figure 9, works similarly. □

It seemed more natural in the above proposition to consider antichain maps, rather than the more traditional chain maps. Just as chain maps induce maps on homology, so do antichain maps. One could alternatively consider the chain map $\widetilde{\Phi}$ defined by

$$\widetilde{\Phi}_{\beta\gamma}(\mathbf{x}) = (-1)^{M(\mathbf{x})} \cdot \Phi_{\beta\gamma}.$$

We now turn to the chain homotopies gotten by counting hexagons.

Once again, there is a straightening map

$$e' \colon \mathrm{Hex}_{\beta\gamma\beta}(\mathbf{x}, \mathbf{y}) \longrightarrow \mathrm{Rect}(\mathbf{x}, \mathbf{y}),$$

and we can define a homotopy operator $H_{\beta\gamma\beta} \colon C(G) \longrightarrow C(G)$ by

$$H_{\beta\gamma\beta}(\mathbf{x}) = \sum_{\mathbf{y} \in \mathbf{S}(G)} \sum_{\substack{h \in \mathrm{Hex}_{\beta\gamma\beta}(\mathbf{x},\mathbf{y}) \\ \mathbf{x} \cap \mathrm{Int}(h) = \varnothing}} \epsilon'(h) \cdot U_1^{O_1(h)} \cdots U_n^{O_n(h)} \cdot \mathbf{y},$$

where $$\epsilon'(h) = \mathcal{S}(e'(h)).$$

Similarly define $H_{\gamma\beta\gamma}$.

Proposition 4.24 *With respect to the sign refinements, the map $\Phi_{\beta\gamma}$ induced by commuting two columns induces an isomorphism in homology.*

Proof The proof of Proposition 3.2 adapts readily to show that

$$\mathbb{I} + \Phi_{\gamma\beta} \circ \Phi_{\beta\gamma} + \partial \circ H_{\beta\gamma\beta} + H_{\beta\gamma\beta} \circ \partial = 0$$
$$\mathbb{I} + \Phi_{\beta\gamma} \circ \Phi_{\gamma\beta} + \partial \circ H_{\gamma\beta\gamma} + H_{\gamma\beta\gamma} \circ \partial = 0.$$

Note that in the terms $\Phi_{\gamma\beta}\circ\Phi_{\beta\gamma}$ and $\Phi_{\beta\gamma}\circ\Phi_{\gamma\beta}$, the two pentagons that appear are either both right pentagons or both left pentagons, so the extra minus sign for right pentagons has no effect. The proposition now follows. □

With commutation invariance in hand, we now turn to stabilization invariance, following the steps in Section 3.2.

As a first step, we need a sign in the definition of C', the mapping cone of the chain map

$$U_1 - U_2\colon B[U_1] \longrightarrow B[U_1],$$

to ensure it is a chain complex. One way of doing this is to define

$$\partial'(a,b) = (\partial a, (U_2 - U_1)\cdot a - \partial b).$$

We will find the following terminology useful.

Definition 4.25 Suppose $p \in \pi(\mathbf{x},\mathbf{y})$ can be decomposed as $p = r_1 * \cdots * r_m$ for some n and $r_i \in \mathrm{Rect}^\circ$. Suppose moreover that for some i we have $r_i * r_{i+1} = r_i' * r_{i+1}'$, for some $r_i', r_{i+1}' \in \mathrm{Rect}^\circ$. Then we say that the decompositions $r_1 * \cdots * r_m$ and $r_1 * \cdots * r_i' * r_{i+1}' * \cdots * r_m$ differ by an *elementary move*.

Recall that in Section 3.2, we identified $C(G)$ with a chain complex whose generators are $\mathbf{I} \subset \mathbf{S}(H)$, which contain the distinguished point x_0. The differentials count either empty rectangles which do not contain x_0 as corner points, or those of "Type 2", ie those which include x_0 (and hence also O_1 and X_1) in their interior. These rectangles are counted, but not with a power of U_1.

Starting with a sign assignment $\mathcal{S}$ on rectangles in H, we can induce one on $\mathbf{I}$ with the above two types of differential. To do this, we must give an explicit decomposition of each Type 2 rectangle r' as a product of three empty rectangles. There are four ways of doing this. We choose the following one: the initial rectangle is lower left (and involves x_0), and the third (final) one uses the upper left corner, cf Figure 31, and call it the *standard decomposition* $D_0(r')$ of the rectangle r'. For consistency, if r is a Type 1 rectangle in G, and r' is the corresponding rectangle viewed as a rectangle in H, we let $D_0(r')$ denote the length 1 decomposition $D_0(r') = r'$.

Lemma 4.26 *Fix $\mathbf{x},\mathbf{y} \in \mathbf{S}(G)$ and let $r \in \mathrm{Rect}^\circ(\mathbf{x},\mathbf{y})$ correspond to the rectangle r' connecting $\mathbf{x}',\mathbf{y}' \in \mathbf{S}(H)$ under the correspondence above between $\mathbf{S}(G)$ with $\mathbf{I} \subset S(H)$. For any sign assignment $\mathcal{S}$ for H, define*

$$\mathcal{S}_0(r) = \begin{cases} \mathcal{S}(r') & \text{if } r' \text{ has Type 1} \\ \mathcal{S}(r_1)\mathcal{S}(r_2)\mathcal{S}(r_3) & \text{if } r' \text{ has Type 2 and } D_0(r') = r_1 * r_2 * r_3. \end{cases}$$

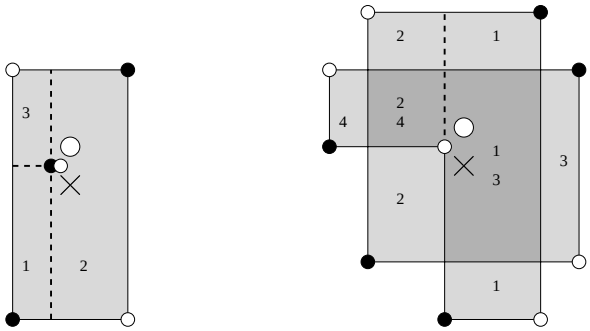

Figure 31: Decomposing polygons. On the left, we a have a standard decomposition of a Type 2 rectangle: more precisely, the decomposition consists of $r_1 * r_2 * r_3$, and the number on a region indicates which rectangle it belongs to. On the right, we have indicated the standard decomposition of a complexity 5 polygon of type L, $r_1 * \cdots * r_4$. Some regions (which have local multiplicity 2) are contained in the support of more than one rectangle, and hence are labelled with more than one integer.

Then $\mathcal{S}_0$ induces a sign assignment in the sense of Definition 4.1 for rectangles in G.

Proof We must show that if $r_1 * r_2 = r_1' * r_2'$ in G, then the decompositions $D_0(r_1) * D_0(r_2)$ and $D_0(r_1') * D_0(r_2')$ can be connected by an odd number of elementary moves. This follows from a routine case analysis of how r_1 and r_2 can interact. When one rectangle is Type 1 and the other is not, it takes 3, 5, or 7 elementary moves to connect the two decompositions. The most complicated case is the one where both are Type 2 rectangles; in that case, we can connect the two decompositions by nine moves: Write the standard decomposition $D_0(r_1) = s_1 * s_2 * s_3$ and $D_0(r_2) = t_1 * t_2 * t_3$, then we have a decomposition $D_0(r_1) * D_0(r_2) = s_1 * s_2 * s_3 * t_1 * t_2 * t_3$. In three moves, we commute t_1 to the beginning, then in three more moves we commute t_2 to the second place, and in three more moves we commute t_3 to the third spot. The resulting decomposition can be easily seen to agree with $D_0(r_1') * D_0(r_2')$. The other two axioms of a sign assignment are also easily verified. □

Remark 4.27 By "commuting the t_1 to the beginning of $s_1 * s_2 * s_3 * t_1 * t_2 * t_3$", we mean the following string of operations: apply three elementary moves, the first of which replaces the consecutive terms $s_3 * t_1$ by an alternative pair $t_1' * s_3'$, then apply another elementary move to the consecutive pair $s_2 * t_1'$, to get $t_1'' * s_2'$, and finally apply an elementary move to the pair $s_1 * t_1''$ to get $t_1''' * s_1'$. We will use this shorthand in several future proofs, as well.

We need now to introduce signs in the definition of the stabilization map F of Equation (5) to ensure that it is, in fact, a chain map.

As a first step, we define a function

$$\mu\colon \pi^F \longrightarrow \{\pm 1\}.$$

For this, we give a specific decomposition of $p \in \pi^F(\mathbf{x}, \mathbf{y})$ as an ordered juxtaposition of rectangles. Specifically, recall that ∂p can be thought of as an oriented, connected, curve. Order now the β–arcs $\{v_i\}_{i=1}^m$ so that they inherit the cyclic ordering from the orientation of ∂p, and so that v_m contains the stabilization point x_0 (which in turn is a component of $\mathbf{y}$). We can decompose

$$p = r_1 * \cdots * r_{m-1},$$

where r_i is a rectangle containing v_i, compare Figure 31, and define

$$\mu(p) = S(r_1)\cdots S(r_{m-1}).$$

Note that the left edge of each odd rectangle is contained in the circle β_1 containing x_0, while the right edge of each even rectangle is contained in β_1. We call this decomposition the *standard decomposition* $D(p)$. For polygons with complexity m, there are $m-1$ rectangles in this decomposition; m is odd if the polygon is of type L and even if the polygon is of type R.

We will analyze the signs according to the cases in the proof of Lemma 3.5.

Lemma 4.28 *Fix a complexity m domain $p \in \pi^F$ and a rectangle r as in cases I(0) or I(1); that is, they are either disjoint or share one corner, with r disjoint from x_0. Then we can either compose $r * p$ or $p * r$, and this composite has an alternate decomposition which is either of the form $r' * p'$ or $p' * r'$ where r' is an empty rectangle distinct from r and p' is a domain of type F distinct from p. We have the following cases:*

- *If $p * r = r' * p'$ (or $r * p = p' * r'$), then $\mu(p)\mathcal{S}_0(r) + (-1)^m \mathcal{S}_0(r')\mu(p') = 0$.*
- *If $r * p = r' * p'$ or $p * r = p' * r'$, then $\mathcal{S}_0(r)\mu(p) + \mathcal{S}_0(r')\mu(p') = 0$.*

Proof If $r * p = p' * r'$ and p has complexity m, then the decomposition $r * D(p)$ can be obtained from $D(p') * r'$ by $m-1$ elementary moves: we successively commute the rectangles in $D(p)$ past the rectangle r. The case where $r' * p' = p * r$ is symmetric.

Otherwise, r shares an edge with some rectangle r_i contained in $D(p)$. With some number k of elementary moves, we can change to a composition series where some rectangle s, with the same support as r, appears next to r_i. Then we can perform one elementary move to change these two rectangles (s and r_i) to s' and r_i', respectively,

where s has the same support as r'; then k more elementary moves returns us to the composition series for r' and $D(p')$, for a total of $2k+1$ elementary moves, which is odd, as desired. □

Lemma 4.29 *Suppose that $p \in \pi^F(\mathbf{x}, \mathbf{y}')$ is a domain of complexity m, $r \in \mathrm{Rect}(\mathbf{y}, \mathbf{z})$ is a Type 2 rectangle, and the corners of p and r are distinct; that is, they are in case II(0). This case matches with case I(2), so either $p * r$ has an alternate decomposition $r' * p'$, in which case*

$$\mu(p)\mathcal{S}_0(r) + (-1)^m \mu(p')\mathcal{S}_0(r') = 0, \tag{14}$$

or $p * r$ has an alternate decomposition as $p' * r'$, where r' is a Type 1 rectangle, in which case

$$\mu(p')\mathcal{S}_0(r') + \mu(p)\mathcal{S}_0(r) = 0. \tag{15}$$

Proof In the first case, we find it convenient to start with the decomposition $r' * p'$, and write the standard decomposition

$$D(p') = r'_1 * \cdots * r'_{m+1}.$$

Recall that r' shares two corner points with p'. There are two cases, according to whether these two corner points are upper right or lower left corners of p'. We consider the upper right case. In this case, the boundary of r' meets the boundary of two consecutive odd rectangles, as in Figure 32. Write the first as r'_{2i-1}. Starting from

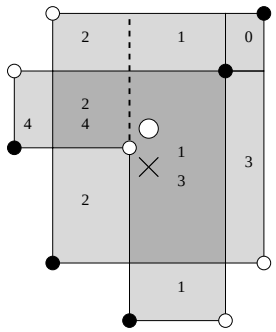

Figure 32: $p * r = r' * p'$, where r has Type 2. The rectangle labeled by 0 is the rectangle r', and p' is a polygon with complexity 5, whose standard decomposition is indicated by the numbers. This decomposition can be transformed into a decomposition $D(p) * D(r)$, where r has Type 2, in an odd number of steps.

$r' * D(p')$, we perform $2i-1$ elementary moves, to commute r' past r'_{2i-1}, to obtain a new decomposition $r_1 * \cdots * r_{2i-1} * r'' * r'_{2i} * r'_{2i+1} * \cdots * r'_{m-1}$. The support of the union of the three consecutive rectangles $r'' * r'_{2i} * r'_{2i+1}$ is a rectangle, which is decomposed so that the upper right rectangle r'' comes first.

We now need some terminology, and then a simple observation. Suppose that $s_1 * s_2 * s_3$ are three rectangles, whose union is a rectangle (so that some point x in its interior is a corner of two of the three original rectangles). Suppose also that s_1, s_2, and s_3 are ordered so that the lower left corner of the total rectangle is the lower left corner of s_1, while the upper left corner in the total rectangle is the upper left corner of s_3. We then say that $s_1 * s_2 * s_3$ is the standard decomposition of a rectangle. (This notion coincides with the earlier standard decomposition of a Type 2 rectangle, when the central point $x = x_0$.) Suppose now that r is some rectangle which can be postcomposed with $s_1 * s_2 * s_3$, and which has two corners inside the support of $s_1 * s_2 * s_3$. Then after an odd number of elementary moves (actually, 3 or 5), we can transform $s_1 * s_2 * s_3 * r$ to $r' * s'_1 * s'_2 * s'_3$, so that the supports of r and r' coincide, and $s'_1 * s'_2 * s'_3$ is a standard decomposition of the rectangle.

Starting from the composition $r'' * r'_{2i} * r'_{2i+1}$, we can apply two elementary moves to transform it into the standard decomposition $s_1 * s_2 * s_3$ of a rectangle. Then, applying the principle in the previous paragraph $m-2i$ times, we can commute $s_1 * s_2 * s_3$ to the end of the decomposition

$$D' = r_1 * \cdots * r_{2i-1} * s_1 * s_2 * s_3 * r'_{2i+1} * \cdots * r'_{m-1},$$

turning it into the desired decomposition $D(p) * D_0(r)$, where r is a Type 2 rectangle. In all, the number of elementary moves has the same parity as $m+1$, verifying Equation (14) in the case where r' is on the upper right side of the support of p'. The case when r' meets p' in lower left corners of p' is similar.

For Equation (15), again there are two cases, according to whether r' shares two lower right or two upper left corner points of p'. Assume they are lower right, and write the standard decomposition of p', $D(p') = r'_1 * \cdots * r'_{m+1}$. Now, one edge of r' is contained in r'_{2i-1}, while the other is contained in r'_{2i+1}. Consider now the decomposition $D(p') * r'$. In $m-2i$ elementary moves, we commute r' before r'_{2i-1}, to obtain a new decomposition

$$r'_1 * \cdots * r'_{2i-2} * r_{2i-1} * s_1 * s_2 * s_3 * t_{2i+2} * \cdots * t_{m+2},$$

where $s_1 * s_2 * s_3$ is a decomposition of a rectangle. After two elementary moves, we can change it to a standard decomposition of the rectangle. Applying the observation about commuting rectangles discussed above, we can now commute this decomposition

of the rectangle to the end of our discussion; the number of steps is congruent to $m+1$ modulo two. The new decomposition is the decomposition $D(p)*r$, and it was obtained from $D(p')*r$ by an odd number of elementary moves. Once again, the case when r' meets p' in upper left corners of p' works similarly. □

Lemma 4.30 *Let $r \in \mathrm{Rect}^\circ(\mathbf{x},\mathbf{y})$ and $p \in \pi^F(\mathbf{y},\mathbf{z})$ a domain with complexity m. Suppose that r and p share one corner, and suppose that x_0 appears in the interior of the boundary of r (Case I(1′)). Then there is a horizontal or vertical annulus so that the domain p_0' obtained by adding the annulus to $r*p$ has an alternate decomposition. We have the following cases:*

(1) *If there is a $\mathbf{y}'\neq\mathbf{y}$ and $r'\in\mathrm{Rect}^\circ(\mathbf{x},\mathbf{y}')$ and $p'\in\pi^F(\mathbf{y}',\mathbf{z})$ so that $r'*p'=p_0'$ (Case I(3)), then $\mathcal{S}_0(r)\mu(p)+\mathcal{S}_0(r')\mu'(p)=0$.*

(2) *If there is a $\mathbf{y}'\neq\mathbf{y}$ and $p'\in\pi^F(\mathbf{x},\mathbf{y}')$ and $r'\in\mathrm{Rect}^\circ(\mathbf{y}',\mathbf{z})$ so that $p'*r'=p_0'$ (Case I(3) if r' has Type 1, Case II(1) if r' has Type 2) then $\mathcal{S}_0(r)\mu(p)+(-1)^m\mu(p')\mathcal{S}_0(r')=0$.*

Proof Note that in the alternate decomposition, in the case where r' is of Type 1, the rectangles r' and p' meet in three points, and the composite domain of r' and p' contains an annulus.

Consider Case (1). This can be subdivided into two subcases: either the annulus is vertical or horizontal. Suppose the annulus is horizontal. In this case, write $D(p')=r_1'*\cdots*r_{m+1}'$. Consider the decomposition $r'*D(p')$. After m elementary moves (commuting r' so that it is next-to-last), we obtain an alternate decomposition, where the last two rectangles compose to the row through O_1. Cancelling these last two, and performing $m-1$ more elementary moves (bringing the last term to the first), we obtain the decomposition $r*D(p)$. The total change in sign is $(-1)^{2m-1}=-1$, so $\mathcal{S}_0(r)\mu(p)+\mathcal{S}_0(r')\mu(p')=0$ as claimed.

Consider next the case that the annulus is vertical. Write $D(p')=r_1'*\cdots*r_{m+1}'$, and consider the decomposition $r'*r_1'*\cdots*r_{m+1}'$. The rectangles r' and r_1' together form the column through O_1. Thus, cancelling the first two terms (and introducing a minus sign by Property (V)), we obtain a decomposition $r_2'*\cdots*r_{m+1}'$. Now $r_2'=r$, and $r_3'*\cdots*r_{m+1}'$ is a decomposition of p. We have changed signs only once, so again $\mathcal{S}_0(r)\mu(p)+\mathcal{S}_0(r')\mu(p')=0$.

Examples of both possibilities for Case (1) are given in Figure 33.

Consider Case (2), which we divide into subcases: either r' has Type 1 or Type 2. Consider first the case where r' has Type 1. This again can be divided into two

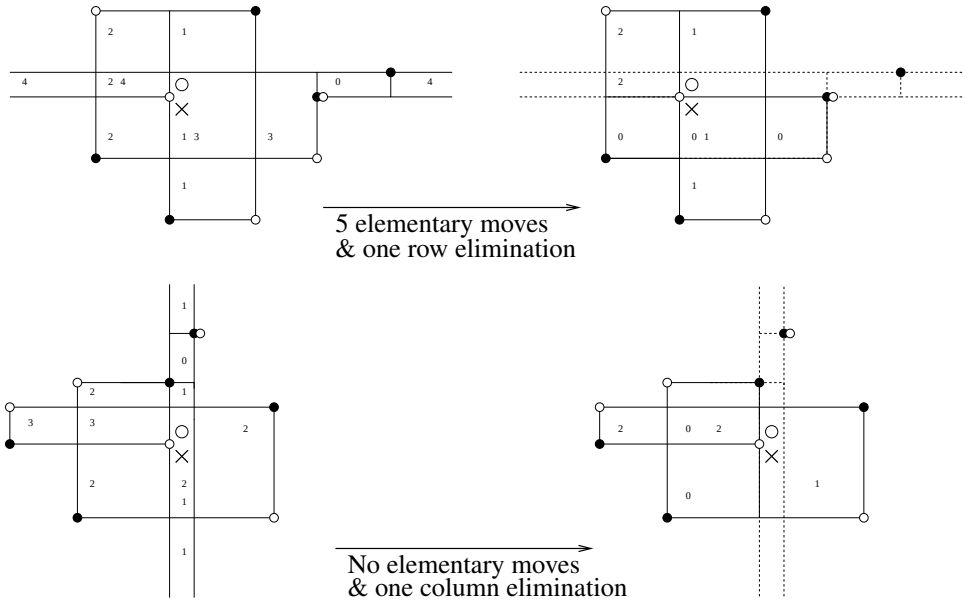

Figure 33: Case (1) of Lemma 4.30. We have illustrated examples of Case (1) in Lemma 4.30. The domains on the right column correspond to decompositions of the form $r * p$, where x_0 is contained in a boundary of r. On the left, we have corresponding alternate composite domains $r' * p'$. The domains here are decomposed into ordered rectangles; the integers in a region give the number of the rectangle the given region belongs to.

subcases, according to whether the annulus in the decomposition $r * p$ is horizontal or vertical. Suppose first that it is horizontal. Write $p' = r'_1 * \cdots * r'_{m+1}$, and consider the decomposition $D(p') * r'$. Performing an elementary move on the last two rectangles, and then on the next-to-last two, we obtain a new decomposition $r'_1 * \cdots * r'_{m-1} * s_m * s_{m+1} * s_{m+2}$, with the property that s_m and s_{m+1} form the row through O_1. Thus, they can be cancelled; performing $m-1$ elementary moves (commuting s_{m+2} to the beginning of the decomposition), we obtain the decomposition $r * D(p)$ with total sign change $(-1)^{m+1}$, verifying the claim. In the case where the annulus is vertical, write $D(p') = r'_1 * \cdots * r'_{m+1}$ and consider the decomposition $D(p') * r'$. We commute the last rectangle r' to the second place in m moves, then cancel the first two rectangles (which, since they form a vertical annulus, introduces the sign -1) to obtain the alternate decomposition $r * D(p)$ (with total sign change $(-1)^{m+1}$). This verifies the stated relation when r' has Type 1.

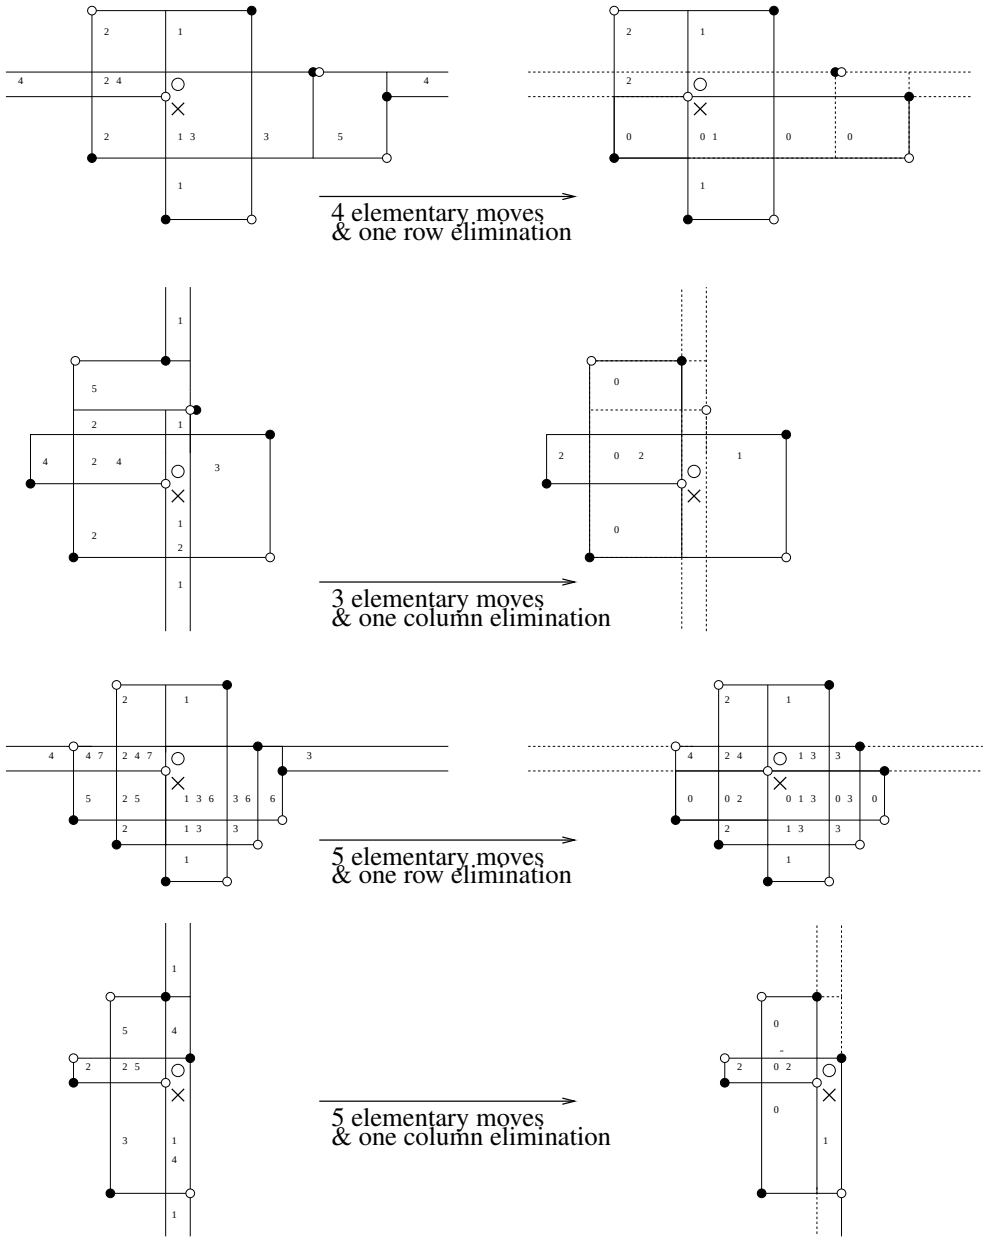

Figure 34: Case (2) of Lemma 4.30. We have illustrated examples of Case (2) in Lemma 4.30. The conventions are the same is an Figure 33.

We turn to the case where r' has Type 2. Again, we have two cases, according to whether the annulus is vertical or horizontal. Assume the annulus is horizontal. Write the standard decompositions $D(p') = r'_1 * \cdots * r'_{m-1}$ and $D_0(r') = r'_m * r'_{m+1} * r'_{m+2}$, and consider the decomposition $D(p') * D_0(r') = r'_1 * \cdots * r'_{m+2}$. Performing one elementary move, replacing $r'_m * r'_{m+1}$ by $s_m * s_{m+1}$, we obtain a new decomposition in which the rectangles r'_{m-1} and s_m form a row, and hence can be cancelled. Finally, in $m-2$ elementary moves (commuting s_{m+1} to the beginning of the decomposition), we obtain the decomposition $r * D(p)$, with total sign change $(-1)^{m-1}$.

Consider the final case, where r' has Type 2 and the annulus is vertical. Again, write the decomposition $D(p') = r'_1 * \cdots * r'_{m-1}$, and $D(r') = r'_m * r'_{m+1} * r'_{m+2}$. Moving r'_{m+1} to the second place in $m-1$ steps, we obtain a new decomposition whose first two terms make up a column. Cancel this, at the cost of introducing one more -1. Next, commute the last two rectangles, and then move the pair to the first two spots in an even number of steps. In this way, we end up with the decomposition $D(r) * D(p)$, with a change in sign of $(-1)^{m+1}$, as needed.

Examples for all possibilities of Case (2) are shown in Figure 34. □

We can now define the stabilization map. In the same way as in Section 3.2, define

$$F^L(\mathbf{x}) = \sum_{\mathbf{y}\in\mathbf{S}} \sum_{p\in\pi^L(\mathbf{x},\mathbf{y})} \mu(p)\cdot U_2^{O_2(p)}\cdots U_n^{O_n(p)}\cdot\mathbf{y}$$
$$F^R(\mathbf{x}) = \sum_{\mathbf{y}\in\mathbf{S}} \sum_{p\in\pi^R(\mathbf{x},\mathbf{y})} \mu(p)\cdot U_2^{O_2(p)}\cdots U_n^{O_n(p)}\cdot\mathbf{y},$$

and put them together to give

$$F = \begin{pmatrix} F^L \\ F^R \end{pmatrix} : C \longrightarrow C'.$$

We have the following sign-refinement of Lemma 3.5:

Lemma 4.31 *The map $F\colon C \longrightarrow C'$ preserves Maslov grading, respects Alexander filtrations, and is a chain map with coefficients in $\mathbb{Z}$.*

Proof Our goal is to show that

$$F\circ\partial_C - \partial_{C'}\circ F = 0.$$

Recall that ∂'_C has three terms: rectangles in $\mathcal{L}$, rectangles in $\mathcal{R}$ (counted with the opposite sign) and the differential from $\mathcal{L}$ to $\mathcal{R}$, multiplication by $(U_2 - U_1)$. Again,

we can collect the terms in $F\circ\partial_C-\partial_{C'}\circ F$ into terms of Types I(0), I(1), I(1′), I(2), I(3), II(0), II(1), and (S). In the proof of Lemma 3.5, we have seen that these terms can be grouped into pairs. We must show that in each pair, the associated signs cancel.

Lemma 4.28 ensures that the terms in Case I(0) and Case I(1) drop out in cancelling pairs. The interesting case is when we have alternate decompositions $p*r=r'*p'$; in this case, if p is of type L, m is odd and the differential in $\mathcal{L}$ corresponding to r is taken with the usual sign $\mathcal{S}_0(r)$, while if p is of type R, m is even and the differential in $\mathcal{R}$ is taken with sign $-\mathcal{S}_0(r)$. In both cases the terms cancel. Similarly, Lemma 4.29 ensures that all terms with complexity $m\geq 3$ in Case I(2) drop out with their corresponding terms in II(0).

Lemma 4.30 ensures that all terms in I(1′) cancel with their corresponding terms of types I(3) or II(1), leaving possible terms of Type (S). Specifically, a term of type I(1′) corresponds to a decomposition $p_0=r_1*p_1$, where r_1 is an empty rectangle and p_1 is a term of type F and complexity m. Adding an annulus to p_0 as in the proof of Lemma 3.5, we obtain a new domain p_0', which in turn decomposes as r_2*p_2 or p_2*r_2 as in case I(3) or II(1). In cases where $O_1\notin r_1$, these terms appear in cancelling pairs according to Lemma 4.30.

In the cases where $O_1\in r_1$, the decomposition r_1*p_1 contributes once counted with a multiple of U_1 (as r_1, which contains O_1, is thought of as a differential in C), the decomposition p_2*r_2 or r_2*p_2 contributes with a multiple of U_2 (as it contains the row through X_1), but there is also a contribution coming from the composite domain p_0, thought of as a domain of type L, times (U_2-U_1), the differential from $\mathcal{L}$ to $\mathcal{R}$ within C'. Cancellation of the terms involving U_1 follows from the observation that $r_1*D(p_1)$ differs from the standard decomposition of p_0 by a sequence of $m-1$ elementary moves (commuting r_1 to the very end), where here m denotes the complexity of p_1 (which is necessarily even). Cancellation of the terms involving U_2 follows since they have the opposite sign from the terms involving U_1.

It is straightforward to see that the remaining possible $m=2$ terms in II(2) cancel the remaining possible two $m=1$ terms of type (S). □

Putting everything together, we have the following:

Proof of Theorem 1.2 This result now is an immediate consequence of Cromwell's theorem and the sign refinements discussed above. Specifically, independence of the choice of sign assignment is established in the uniqueness statement of Theorem 4.2. Commutation invariance follows from Proposition 4.24. Stabilization invariance follows from Lemma 4.31, together with a straightforward adaptation of the proof of Proposition 3.8. □

It is sometimes convenient to consider the chain complex $CL^-(G)$ which is the graded object associated to the Alexander filtration of $C^-(G)$. Explicitly, it is the group with the same underlying chain complex, endowed with a differential as in Equation (3). It is a formal consequence of Theorem 1.2 that the homology $HL^-(G)$ of $CL^-(G)$, thought of as a module over $\mathbb{Z}[U_1,\dots,U_\ell]$, is a link invariant.

5 More properties

We next give a few of the basic properties of knot and link Floer homology. Again, most of these properties are well-known [11; 15; 9]; but again, we can give a self-contained derivation here. Let $\widehat{HL}_d(\vec{L},\mathbf{s})$ be the part of $\widehat{HL}(\vec{L})$ with Alexander grading $\mathbf{s}$ and Maslov grading d.

Proposition 5.1 *The total homology groups of the chain complex $C^-(G)$ are isomorphic to the module $\mathbb{Z}[U]$, where all the U_i act as multiplication by U. The homology groups of $\widehat{C}(G)$ are isomorphic to $\mathbb{Z}$.*

Proof The chain complex $C^-(G)$ refers to the $\{X_i\}_{i=1}^n$ only through its Alexander filtration; in particular, the homology of $C^-(G)$ makes no reference to this placement, and it is unchanged by a rearrangement of these decorations (though it does appear to depend on the placement of the $\{O_i\}_{i=1}^n$). Now, given any grid diagram G, we can consider instead the alternate grid diagram H gotten by placing X_i in the square immediately under each O_i. This new diagram clearly represents a suitably stabilized diagram for the unknot. Indeed, after destabilizing sufficiently many times, we can reduce to the 2×2–grid diagram J for the unknot. A direct calculation in this case gives that $H_*(C^-(J))\cong\mathbb{Z}[U]$ (or $\mathbb{F}_2[U]$ with coefficients modulo 2).

The analogous statement for $\widehat{C}(G)$ follows similarly. □

Proposition 5.1 allows us to define the invariant $\tau(K)$ for a knot K (see Ozsváth and Szabó [10] and Rasmussen [15]; compare also Rasmussen [14]): If we consider the natural inclusion map $\iota_m\colon \mathcal{F}_m(\widehat{C}(G))\longrightarrow\widehat{C}(G)$, then $\tau(K)$ is the smallest integer m for which the map induced on homology by ι_m is nontrivial, as a map to $H_*(\widehat{C}(G))\cong\mathbb{Z}$.

The Alexander polynomial of a link remains unchanged under overall orientation reversal, it is a symmetric polynomial, and it is invariant under mirror. These three facts are reflected in Propositions 5.2, 5.3, and 5.5, respectively.

Proposition 5.2 *The filtered quasi-isomorphism type of the complex $C^-(G)$ does not change if we reverse the orientation of all components of the link $\vec{L}$.*

Proof Consider the diagram G' obtained by switching the x and y coordinates, thus flipping G along the diagonal from the bottom left to upper right corner. Switching the x and y coordinates also gives a map from the original set of generators $\mathbf{S}$ to the new set of generators $\mathbf{S}'$ which preserves both degrees and is a chain map. The new diagram G' is a diagram for $\vec{L}$ with the orientation of each component reversed.

A few more remarks are needed when working over $\mathbb{Z}$, since the precomposition of a sign assignment with reflection through the diagonal is not quite a sign assignment, in the sense of Definition 4.1: the roles of rows and columns are reversed. However, this is remedied by substituting $-U_i$ in place of U_i. □

Proposition 5.3 *Given $\mathbf{s} = (s_1, \dots, s_\ell) \in (\frac{1}{2}\mathbb{Z})^\ell$, we have that*

$$\widehat{HL}_d(\vec{L}, \mathbf{s}) \cong \widehat{HL}_{d-2S}(\vec{L}, -\mathbf{s}),$$

where $S = \sum_{i=1}^{\ell} s_i$.

Proof Fix a grid diagram for $\vec{L}$, and let A^1 and M^1 denote its total Alexander filtration and Maslov grading. (By total Alexander filtration, we mean the sum of the components of the Alexander multifiltration.) Switching the roles of $\mathbb{O}$ and $\mathbb{X}$, we obtain a grid diagram for $-\vec{L}$. Differentials within $\widetilde{C}$ are the same for the two diagrams, but the Alexander and Maslov gradings are different. We let A^2 and M^2 denote the Alexander and Maslov gradings of the new diagram. We find it convenient to symmetrize A^i, defining

$$\widetilde{A}^i(\mathbf{x}) = A^i(\mathbf{x}) + \Big(\frac{n_1 - 1}{2}, \dots, \frac{n_\ell - 1}{2}\Big)$$

for $i = 1, 2$. It is a straightforward calculation from Equations (1) and (2) that

$$\begin{aligned} M^1 - 2\sum_{i=1}^{\ell} \widetilde{A}^1_i &= M^2 \\ -\widetilde{A}^1 &= \widetilde{A}^2. \end{aligned}$$

The result now follows from Proposition 5.2 together with Proposition 4.22. □

Indeed, we have the following more general version:

Proposition 5.4 *Let $\vec{L}$ be an oriented, ℓ–component link, and let $\vec{L}'$ be the oriented link obtained from $\vec{L}$ by reversing the orientation of its i th component. Then, writing $\mathbf{s} = (s_1, \dots, s_\ell)$,*

$$\widehat{HL}_d(\vec{L}, (s_1, \dots, s_\ell)) \cong \widehat{HL}_{d-2s_i+\ell_i}(\vec{L}', (s_1, \dots, s_{i-1}, -s_i, s_i, \dots, s_\ell)),$$

where here ℓ_i denotes the total linking number of the i th component of $\vec{L}$ with the remaining components.

Proof From a grid diagram for $\vec{L}$ we can obtain a grid diagram for $\vec{L}'$ by switching the roles of $\mathbb{O}_i$ and $\mathbb{X}_i$, ie those markings which correspond to the i th component of the link. As in the proof of Proposition 5.3, the complexes $\widetilde{C}$ agree, but the Maslov and Alexander functions change, as follows. Let $\widetilde{A}^1_j$ and M^1 be the j th symmetrized Alexander filtration and Maslov grading for $\vec{L}$, respectively, and let $\widetilde{A}^2_j$ and M^2 be the corresponding functions for $\vec{L}'$. Let $\widetilde{\mathbb{O}}_i = \mathbb{O} \setminus \mathbb{O}_i$ and $\widetilde{\mathbb{X}}_i = \mathbb{X} \setminus \mathbb{X}_i$.

By a direct application of Equation (2),

$$\widetilde{A}^2_j(\mathbf{x}) = \begin{cases} \widetilde{A}^1_j(\mathbf{x}) & i \neq j \\ -\widetilde{A}^1_i(\mathbf{x}) & i = j, \end{cases}$$

and, using Equations (1) and (2),

$$\begin{aligned} M^2(\mathbf{x}) - M^1(\mathbf{x}) &= -2\mathcal{J}(\mathbf{x} - \mathbb{O}, \mathbb{X}_i - \mathbb{O}_i) + \mathcal{J}(\mathbb{X}_i - \mathbb{O}_i, \mathbb{X}_i - \mathbb{O}_i) \\ &= -2\widetilde{A}_i(\mathbf{x}) - \mathcal{J}(\mathbb{X} - \mathbb{O}, \mathbb{X}_i - \mathbb{O}_i) + \mathcal{J}(\mathbb{X}_i - \mathbb{O}_i, \mathbb{X}_i - \mathbb{O}_i) \\ &= -2\widetilde{A}^1_i(\mathbf{x}) + \mathcal{J}(\widetilde{\mathbb{X}}_i - \widetilde{\mathbb{O}}_i, \mathbb{X}_i - \mathbb{O}_i). \end{aligned}$$

Moreover, it is straightforward to see that $\mathcal{J}(\widetilde{\mathbb{X}}_i - \widetilde{\mathbb{O}}_i, \mathbb{X}_i - \mathbb{O}_i) = \ell_i$. □

Proposition 5.5 *Let $\vec{L}$ be a link, and let $r(\vec{L})$ denote its mirror. In this case, we have an identification*

$$\widehat{HL}_d(\vec{L}, \mathbf{s}) \cong \widehat{HL}^{2S-d}(\widehat{C}(r(\vec{L}), \mathbf{s}));$$

note the right-hand side denotes cohomology.

Proof Rotating the grid diagram G ninety degrees to get a new diagram G' corresponds to passing from the knot to its mirror. There is an induced map ϕ from $\mathbf{S}(G)$ to $\mathbf{S}(G')$. Letting $\widetilde{A}$, $\widetilde{M}$ and $\widetilde{A}'$, $\widetilde{M}'$ be the Alexander and Maslov gradings of G and G' respectively, it is clear that $\widetilde{A}(\mathbf{x}) = -\widetilde{A}'(\phi(\mathbf{x}))$, $\widetilde{M}(\mathbf{x}) = -\widetilde{M}'(\phi(\mathbf{x}))$. Indeed, if we think of ϕ as taking $\mathbf{x}$ to $\mathbf{x}^*$, the dual basis element of $\widetilde{C}(H)$ which is one on $\mathbf{x} \in \mathbf{S}(H)$ and zero on all other $\mathbf{y} \in \mathbf{S}(H)$, then ϕ induces an isomorphism of chain complexes. The shift in absolute grading now follows from Proposition 5.3. □

6 Relation to the Alexander polynomial

In this section we will show that the Euler characteristic of the multigraded complex $\widehat{CL}$ with respect to the Maslov grading is the Alexander polynomial. Precisely, fix a grid

diagram G of a link $\vec{L}$ with ℓ components. Given $\mathbf{s} \in (\frac{1}{2}\mathbb{Z})^\ell$, let $t = (t_1, \dots, t_\ell)$ be a collection of variables, and for $\mathbf{s} = (s_1, \dots, s_\ell)$ an element of $(\frac{1}{2}\mathbb{Z})^\ell$, define

$$t^{\mathbf{s}} = t_1^{s_1} \cdots t_\ell^{s_\ell}.$$

For any multigraded groups $C_i(\mathbf{s})$ with Maslov grading i and Alexander grading s, define

$$\chi(C; t) = \sum_{i,\mathbf{s}} (-1)^i t^{\mathbf{s}} \operatorname{rank}(C_i(\mathbf{s})).$$

Theorem 6.1 *For any link $\vec{L}$, the Euler characteristic of $\widehat{HL}$ determines the multivariable Alexander polynomial up to sign. Precisely,*

$$\chi(\widehat{HL}(\vec{L})) = \begin{cases} \pm \prod_{i=1}^{\ell} (t_i^{1/2} - t_i^{-1/2}) \Delta_A(\vec{L}; t) & \ell > 1 \\ \pm \Delta_A(\vec{L}; t) & \ell = 1, \end{cases}$$

where $\Delta_A(\vec{L}; t)$ is the multivariable Alexander polynomial, normalized so that it is symmetric up to sign under the involution of sending all t_i to their inverses.

We will prove this by taking the Euler characteristic of the alternate complex $\widetilde{C}(\vec{L})$. The Maslov grading of a generator $\mathbf{x} \in \mathbf{S}$ is, up to an overall sign, the sign of $\mathbf{x}$ considered as a permutation. The Alexander grading is, up to an overall shift, minus the sum of the winding numbers around the generators. Summing over all generators $\mathbf{x}$, we get a "minesweeper determinant" as illustrated below:

$$\begin{vmatrix} 1 & 1 & 1 & t & t & t \\ 1 & 1 & t^{-1} & 1 & t & t \\ 1 & t & 1 & 1 & t & t \\ 1 & t & t & t & t^2 & t \\ 1 & t & t & t & t & 1 \\ 1 & 1 & 1 & 1 & 1 & 1 \end{vmatrix}$$

This turns out to give the Alexander polynomial, up to a sign, powers of the variables t_i, and factors of $(1 - t_i)$.

More formally, given a grid diagram G of size n, define an $n \times n$ matrix $M(G)$ by

$$M(G)_{ij} = t^{a(i,j)}$$

where $a(i,j) \in (\frac{1}{2}\mathbb{Z})^\ell$ is the vector whose kth component is the minus the winding number of the kth component of the link around the point (i,j). (Here we use the convention that the links runs between the O's and the X's, which have half-integer coordinates, so this winding number is well-defined.) Then we have:

Proposition 6.2 *For any grid diagram G of a link $\vec{L}$ with ℓ components, let n_i be the number of vertical segments corresponding to component i. Then*

$$\det M(g) = \begin{cases} \pm t^k (1-t)^{n-1} \Delta(\vec{L}; t) & \ell = 1 \\ \pm \left(\prod_{i=1}^{\ell} t_i^{k_i} (1-t_i)^{n_i} \right) \Delta(\vec{L}; t) & \ell > 1 \end{cases}$$

for some integers k_i. In the case $\ell = 1$, we write t, n, k for t_1, n_1, k_1 for convenience.

This proposition implies Theorem 6.1:

Proof of Theorem 6.1 It follows from Proposition 4.22 that

$$\left(\prod_i (1-t_i)^{n_i - 1} \right) \chi(\widehat{HL}(\vec{L})) = \chi(\widetilde{HL}(G)) = \chi(\tilde{C}(G)).$$

The theorem follows by Proposition 6.2 up to an overall sign and powers of the t_i. The powers of t_i are determined by Proposition 5.3 and the chosen normalization of the Alexander polynomial. □

Proof of Proposition 6.2 We will use Fox's free differential calculus [5] with respect to a presentation associated to a grid diagram for a link. This presentation was described by Neuwirth [8], who also proved that it is actually a presentation of the knot group.

The presentation, as shown in Figure 35, has one generator for each vertical segment of the link, starting from the basepoint (outside of the page at the position of the reader), coming down to the left of a vertical segment, going under the segment, and coming back out of the page. There are n generators, one for each vertical segment. There are $n-1$ relations, one for each position between horizontal segments. The path of the relation comes down to the left of the diagram, runs across the diagram at a constant horizontal level, and comes back up out of the page. On the one hand this loop is contractible (pull it beneath the diagram); on the other hand it is equal to the product of the generators corresponding to the vertical segments that we cross.

In the example in Figure 35, there are six generators, x_0 through x_5, and five relations:

$$r_1 = x_2 x_5 \qquad r_2 = x_1 x_2 x_3 x_5$$
$$r_3 = x_0 x_1 x_3 x_5 \qquad r_4 = x_0 x_3 x_4 x_5$$
$$r_5 = x_0 x_4.$$

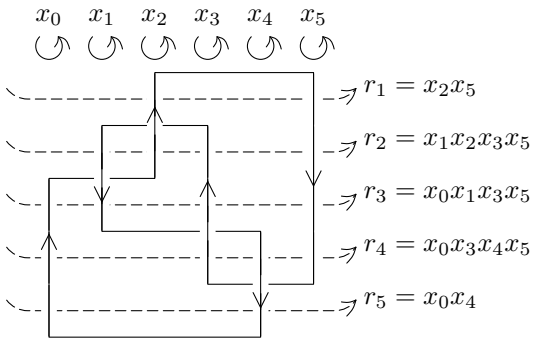

Figure 35: Generators and relations from a gridlink presentation

The Fox derivative matrix is particularly easy to compute in this case, since the generators appear only positively in the relations and at most once in each relation. The entries are the initial portions of the relations, and they appear in the positions given by the vertical strands. In our running example, we get

$$\left(\frac{\partial r_i}{\partial x_j}\right)_{ij} = \begin{pmatrix} 0 & 0 & 1 & 0 & 0 & x_2 \\ 0 & 1 & x_1 & x_1x_2 & 0 & x_1x_2x_3 \\ 1 & x_0 & 0 & x_0x_1 & 0 & x_0x_1x_3 \\ 1 & 0 & 0 & x_0 & x_0x_3 & x_0x_3x_4 \\ 1 & 0 & 0 & 0 & x_0 & 0 \end{pmatrix}.$$

To find the Alexander polynomial, map this matrix to the abelianization of the knot group, mapping each x_j to $t_i^{\pm 1}$, depending on which component the vertical segment belongs to and whether the corresponding vertical strand is oriented upwards or downwards. For knots, the Alexander polynomial is the determinant of a maximal minor of the resulting n by $n-1$ matrix, up to a factor of $\pm t^k$. For links, the multivariable Alexander polynomial is, up to the same factor, the determinant of a maximal minor divided by $(1-t_i)$ for each component i that is not the component of the deleted column.

In the example, we get

$$\Delta(t) = \pm t^k \begin{vmatrix} 0 & 0 & 1 & 0 & 0 \\ 0 & 1 & t^{-1} & 1 & 0 \\ 1 & t & 0 & 1 & 0 \\ 1 & 0 & 0 & t & t^2 \\ 1 & 0 & 0 & 0 & t \end{vmatrix}.$$

Now let us turn to computing $\det(M)$ where M is the minesweeper matrix defined earlier. Subtract each column from the next one. The winding numbers change by at most one when we move from one square to a neighbor. Therefore in every column but the first we have zero entries where the vertical segment does not intervene, and where a vertical segment of component i does intervene every entry is divisible by $t_i^{\pm 1}-1$. Thus, for each column but the first, we can factor out t_i-1 if the column is oriented upwards or $t_i^{-1}-1$ if it is oriented downwards. Furthermore, after this operation the last row contains only a single nonzero entry, 1 in the first column, so we can delete the first column and last row without changing the determinant (up to sign).

In the example, we get

$$\begin{vmatrix} 1 & 1 & 1 & t & t & t \\ 1 & 1 & t^{-1} & 1 & t & t \\ 1 & t & 1 & 1 & t & t \\ 1 & t & t & t & t^2 & t \\ 1 & t & t & t & t & 1 \\ 1 & 1 & 1 & 1 & 1 & 1 \end{vmatrix} = \begin{vmatrix} 1 & 0 & 0 & t-1 & 0 & 0 \\ 1 & 0 & t^{-1}-1 & 1-t^{-1} & t-1 & 0 \\ 1 & t-1 & 1-t & 0 & t-1 & 0 \\ 1 & t-1 & 0 & 0 & t^2-t & t-t^2 \\ 1 & t-1 & 0 & 0 & 0 & 1-t \\ 1 & 0 & 0 & 0 & 0 & 0 \end{vmatrix}$$

$$= \pm(t-1)^3(t^{-1}-1)^2 \begin{vmatrix} 0 & 0 & 1 & 0 & 0 \\ 0 & 1 & t^{-1} & 1 & 0 \\ 1 & t & 0 & 1 & 0 \\ 1 & 0 & 0 & t & t^2 \\ 1 & 0 & 0 & 0 & t \end{vmatrix}.$$

Up to the expected factors of $1-t_i$ and $\pm t^k$, this is the same determinant we got from the Fox derivative calculations. □

References

[1] **J A Baldwin**, **W D Gillam**, *Computations of Heegaard-Floer knot homology* `arXiv:math.GT/0610167`

[2] **P R Cromwell**, *Embedding knots and links in an open book. I. Basic properties*, Topology Appl. 64 (1995) 37–58 MR1339757

[3] **I A Dynnikov**, *Arc-presentations of links: monotonic simplification*, Fund. Math. 190 (2006) 29–76 MR2232855

[4] **A Floer**, *Morse theory for Lagrangian intersections*, J. Differential Geom. 28 (1988) 513–547 MR965228

[5] **R H Fox**, *Free differential calculus. I. Derivation in the free group ring*, Ann. of Math. (2) 57 (1953) 547–560 MR0053938

[6] **C Manolescu**, **P S Ozsváth**, **S Sarkar**, *A combinatorial description of knot Floer homology* arXiv:math.GT/060769

[7] **J McCleary**, *User's guide to spectral sequences*, Mathematics Lecture Series 12, Publish or Perish, Wilmington, DE (1985) MR820463

[8] **L P Neuwirth**, $*$ *projections of knots*, from: "Algebraic and differential topology—global differential geometry", Teubner-Texte Math. 70, Teubner, Leipzig (1984) 198–205 MR792695

[9] **P Ozsváth**, **Z Szabó**, *Holomorphic disks and link invariants* arXiv:math.GT/0512286

[10] **P Ozsváth**, **Z Szabó**, *Knot Floer homology and the four-ball genus*, Geom. Topol. 7 (2003) 615–639 MR2026543

[11] **P Ozsváth**, **Z Szabó**, *Holomorphic disks and knot invariants*, Adv. Math. 186 (2004) 58–116 MR2065507

[12] **P Ozsváth**, **Z Szabó**, *Holomorphic disks and topological invariants for closed three-manifolds*, Ann. of Math. (2) 159 (2004) 1027–1158 MR2113019

[13] **P Ozsváth**, **D P Thurston**, **Z Szabó**, *Transverse knots and combinatorial knot Floer homology*, in preparation

[14] **J A Rasmussen**, *Khovanov homology and the slice genus* arXiv:math.GT/0402131

[15] **J A Rasmussen**, *Floer homology and knot complements*, PhD thesis, Harvard University (2003) arXiv:math.GT/0306378

[16] **S Sarkar**, **J Wang**, *A combinatorial description of some Heegaard Floer homologies* arXiv:math.GT/0607777

[17] **P H Schoute**, *Analytic treatment of the polytopes regularly derived from the regular polytopes*, Verhandelingen der Koninklijke Akademie van Wetenschappen te Amsterdam 11 (1911) 1–87

Department of Mathematics, Columbia University
New York NY 10027, USA
Department of Mathematics, Columbia University
New York NY 10027, USA
Department of Mathematics, Princeton University
Princeton NJ 08544, USA
Department of Mathematics, Barnard College, Columbia University
New York NY 10027, USA
cm@math.columbia.edu, petero@math.columbia.edu,
szabo@math.princeton.edu, dthurston@barnard.edu

Proposed: Rob Kirby Received: 2 November 2006
Seconded: Yasha Eliashberg, Tom Mrowka Accepted: 12 June 2007