

BRAIDS AND COMBINATORIAL KNOT FLOER HOMOLOGY

PETER LAMBERT-COLE, MICHAELA STONE, AND DAVID SHEA VELA-VICK

ABSTRACT. We present a braid-theoretic approach to combinatorially computing knot Floer homology. To a knot or link K , which is braided about the standard disk open book decomposition for (S^3, ξ_{std}) , we associate a corresponding multi-pointed nice Heegaard diagram. We then describe an explicit algorithm for computing the associated knot Floer homology groups. We compute explicit bounds for the computational complexity of our algorithm and demonstrate that, in many cases, it is significantly faster than the previous approach using grid diagrams.

1. INTRODUCTION

Knot Floer homology is a powerful invariant of knots and links defined independently by Ozsváth and Szabó in [OS04b] and by Rasmussen in [Ras03]. It is part of the general Heegaard Floer package [OS04a], and has proven itself to be a tremendously useful and versatile invariant. Knot Floer homology contains a wealth of geometric information about knots in the 3–sphere and in arbitrary 3–manifolds. It is capable of detecting a knot’s genus [OS04c], determining if that knot is fibered [Ghi08, Ni07], and also contains a powerful concordance invariant which provides lower bounds for the slice genus [OS03, Ras03]. The knot Floer homology chain complex for a knot K also contains information about the Heegaard Floer homology of manifolds obtained via surgeries along K [OS08b, OS11].

In 2006, Manolescu, Ozsváth and Sarkar presented a combinatorial method for computing knot Floer homology via grid diagrams [MOS09]. Their construction is based on a result of Sarkar and Wang who described a method for combinatorially computing versions of Heegaard Floer homology in general via “nice” Heegaard diagrams [SW10]. The straightforward, combinatorial nature of Manolescu, Ozsváth and Sarkar’s construction has lead to many interesting and significant applications, including the discovery by Ozsváth, Szabó and Thurston [OST08] of a transverse invariant taking values in knot Floer homology.

For knots with many crossings, computing knot Floer homology via grid diagrams is impractical. Indeed, the computation grows factorially as a function of the grid-number (arc-index) and there exist torus knots for which the time complexity is $O(((c + 2)!)^3)$, where c is the crossing number. The present goal is to present an alternative method for combinatorially computing knot Floer homology which is often faster than via grid diagrams. Our approach goes by way of braid representations of topological knots, and is based on earlier work of

2010 *Mathematics Subject Classification.* 57M27; 57R58.

Key words and phrases. Heegaard Floer homology.

Baldwin, Vértesi and the third author equating the various transverse invariants defined in the context of knot Floer homology [BVV13]. Our main theorem is the following:

Theorem 1.1. *Given a braid word w representing a braid $\delta \in B_n$ in n -strands, the knot Floer homology $\text{HFK}^-(K(\delta))$ of the associated braid closure can be computed in $O((\frac{4^{l(w)}}{2n} + 4)^{6n} + 4^{l(w)} + n)$ time, where $l(w)$ is the length of the braid word w . Moreover, there exists some positive real constant c , not depending on n or $l(w)$, such that $\text{HFK}^-(K(\delta))$ can be computed in $O(c^{l(w)n}n^{-6n})$*

To prove Theorem 1.1, we begin by associating to any braid word $w \in B_n$ in n strands a natural multi-pointed Heegaard diagram for the corresponding braid closure that we call a “braid diagram”. Roughly speaking, a braid diagram is a Heegaard decomposition of the knot complement, modified slightly to incorporate information about the *transverse* link naturally associated to the braid word w (see [BVV13]). The multi-pointed Heegaard diagrams one obtains via this process are described in Section 2.3. They are not necessarily nice in the sense of Sarkar and Wang, but are nearly so – they can contain at most $n - 1$ six-sided bad regions, where n is the braid index. From here, the key step is to perform an appropriate collection of stabilizations and isotopies to connect the above bad regions to basepoint regions. The result is a nice diagram for the associated braid closure.

There are many examples for which computing knot Floer homology via our method is significantly faster than the corresponding computation in the grid setting. Generally speaking, this method is well-adapted to computing knot Floer homology for braid closures with large twist regions. As a simple example, via our method, the knot Floer complex of the $(2, n)$ -torus knot can be computed in polynomial time and has $(2n + 1)$ generators, whereas the corresponding computation in the grid setting grows factorially in n and requires $(n + 2)!$ generators. In Section 5, we compare the time required to compute knot Floer homology via our method and grid diagrams for general (p, q) -torus knots.

Finally, we remark that it should be possible to apply these methods to combinatorially compute the transverse braid invariant t , defined in [BVV13]. Doing so will require a careful analysis of the isomorphisms induced on knot Floer homology by curve isotopies described in Section 3.1. This is a problem we plan to return to in a future paper.

Organization. In Section 2, we provide some background on knot Floer homology, grid diagrams and braid diagrams. The algorithm and computing time required to produce a nice diagram associated to a braid is presented in Section 3. In Section 4, we prove Theorem 1.1, establishing an upper bound on the time required to compute the knot Floer homology complex via our algorithm. Finally, in Section 5, we compare and contrast our method with the traditional grid approach by focusing on the example of torus knots.

Acknowledgements. We would like to thank Samuel Connolly, Robert Lipshitz, Dylan Thurston and Ellen Weld for helpful conversations. We would also like to extend our sincere gratitude to Louisiana State University for sponsoring the 2012 LSU Research Experience for Undergraduates. Several of the ideas contained in this work originated as offshoots of this program. Finally, Lambert-Cole and Vela-Vick would like to acknowledge partial support from NSF Grant DMS-1249708.

2. PRELIMINARIES

In this section, we collect a few facts about knot Floer homology, Legendrian and transverse knots, and their invariants which will be useful in what follows.

2.1. Knot Floer homology. In this subsection, we review some basic definitions and results from knot Floer homology. Here and for the remainder of the paper we work with coefficients in $\mathbb{F} = \mathbb{Z}/2$. For a more in-depth and elementary treatment, we refer the interested reader to [OS04b, OS08a].

A multi-pointed Heegaard diagram for a (null-homologous) knot or link $K \subset S^3$ is a tuple $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{z}, \mathbf{w})$ consisting of the following:

- A genus g Riemann surface Σ
- Collections of disjoint simple, closed curves $\boldsymbol{\alpha} = \{\alpha_1 \cup \cdots \cup \alpha_{g+n-1}\}$ and $\boldsymbol{\beta} = \{\beta_1 \cup \cdots \cup \beta_{g+n-1}\}$, each of which span g -dimensional subspaces of $H_1(\Sigma; \mathbb{Z})$
- Collections of basepoints $\mathbf{z} = \{z_1, \dots, z_n\}$ and $\mathbf{w} = \{w_1, \dots, w_n\}$ such that each component of $\Sigma \setminus (\cup \alpha_i)$ and $\Sigma \setminus (\cup \beta_i)$ contains exactly one z_i and one w_i .

We require that the triple $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$ be a Heegaard diagram for the 3-sphere, and that the knot $K \subset S^3$ is specified by the collections of basepoints \mathbf{z} and \mathbf{w} as follows. First, choose oriented, embedded arcs $\gamma_1, \dots, \gamma_n$ in the complement $\Sigma \setminus (\cup \alpha_i)$ connecting the z to w -basepoints. Next, choose oriented, embedded arcs $\delta_1, \dots, \delta_n$ in the complement $\Sigma \setminus (\cup \beta_i)$ connecting the w to z -basepoints. Finally, depressing the interiors of these arcs into the α and β -handlebodies respectively, we require that their union $K = (\cup \gamma_i) \cup (\cup \beta_i)$ specify the knot K .

To each basepoint $w_i \in \mathbf{w}$, we associate a formal variable U_{w_i} . The knot Floer complex $\text{CFK}^-(\mathcal{H})$ is then the free $\mathbb{Z}[U_{w_1}, \dots, U_{w_n}]$ -module generated by the intersections of the two $(g+n-1)$ -dimensional tori $\mathbb{T}_\alpha = \alpha_1 \cup \cdots \cup \alpha_{g+n-1}$ and $\mathbb{T}_\beta = \beta_1 \cup \cdots \cup \beta_{g+n-1}$ inside the $(g+n-1)$ -fold symmetric product $\text{Sym}^{g+n-1}(\Sigma)$. The differential on $\text{CFK}^-(\mathcal{H})$ is obtained as follows. Let $\mathbf{x}, \mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ be a pair of generators for $\text{CFK}^-(\mathcal{H})$, $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ a Whitney disk connecting them, and J_t a generic path of almost complex structures on $\text{Sym}^{g+n-1}(\Sigma)$. Let $\mathcal{M}(\phi)$ be the moduli space of pseudo-holomorphic representatives of the disk ϕ , and denote by $\widehat{\mathcal{M}}(\phi)$ its quotient by the natural \mathbb{R} -action given by translation. The differential on $\text{CFK}^-(\mathcal{H})$ is then given by

$$\partial^- \mathbf{x} = \sum_{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} \sum_{\substack{\phi \in \pi_2(\mathbf{x}, \mathbf{y}), \\ \mu(\phi)=1, \\ n_{\mathbf{z}}(\phi)=0}} \# \widehat{\mathcal{M}}(\phi) \cdot U_{w_1}^{n_{w_1}(\phi)} \cdots U_{w_n}^{n_{w_n}(\phi)} \cdot \mathbf{y},$$

where $\mu(\phi)$ is the Maslov index of ϕ and $n_{w_i}(\phi)$ denotes the algebraic intersection number of the Whitney disk ϕ with the subvariety $\{w_i\} \times \text{Sym}^{g+n-2}(\Sigma)$.

The knot Floer complex possesses two natural types of gradings. The first is the Maslov (homological) grading, which is an absolute \mathbb{Q} -grading, specified up to an overall shift by the relation

$$M(\mathbf{x}) - M(\mathbf{y}) = \mu(\phi) - 2 \cdot \sum_i n_{w_i}(\phi),$$

for any pair of generators $\mathbf{x}, \mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ and Whitney disk connecting them $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$, and the requirement that multiplication by any of the formal variables U_{w_i} drop Maslov grading by 2. The second type of grading is known as the Alexander grading. If $K = K_1 \cup \dots \cup K_\ell$ is an ℓ -component link, then for each $i \in \{1, \dots, \ell\}$, the Alexander grading associated to K_i is an absolute \mathbb{Q} -grading, specified up to an overall shift by the relation

$$A_{K_i}(\mathbf{x}) - A_{K_i}(\mathbf{y}) = n_{\mathbf{z}_{K_i}}(\phi) - n_{\mathbf{w}_{K_i}}(\phi),$$

where \mathbf{x}, \mathbf{y} and ϕ are as above, \mathbf{z}_{K_i} and \mathbf{w}_{K_i} are the collections of basepoints corresponding to the link component K_i , and the requirement that multiplication by a formal variable U_{w_j} drop A_{K_i} by 1 if $w_j \in \mathbf{z}_{K_i}$ and otherwise preserve the grading.

The “minus” version of knot Floer homology is the homology of the complex $(\text{CFK}^-(\mathcal{H}), \partial^-)$:

$$\text{HFK}^-_*(K) := \text{H}_*(\text{CFK}^-(\mathcal{H}), \partial^-).$$

When w_i and w_j are basepoints corresponding to the same component of the link K , then their associated formal variables U_{w_i} and U_{w_j} act identically on $\text{HFK}^-(K)$. Choose for each component K_i of the link K , a formal variable U_i associated to some basepoint for K_i . Then the knot Floer homology $\text{HFK}^-(K)$ is an invariant of the link $K \subset S^3$, which is well-defined up to graded $\mathbb{F}[U_1, \dots, U_\ell]$ -module isomorphism.

There are two additional associated homology theories with which one commonly works. The first is known as the “hat” version of knot Floer homology and is obtained as follows. For each component $K_i \in K$, set exactly one of its associated formal variables $U_{w_i} = 0$ and denote by $\widehat{\partial}$ the associated differential on the quotient complex $(\widehat{\text{CFK}}(K), \widehat{\partial})$. It follows that the homology

$$\widehat{\text{HFK}}_*(K) := \text{H}_*(\widehat{\text{CFK}}(\mathcal{H}), \widehat{\partial}),$$

is an invariant of the link K up to \mathbb{F} -module isomorphism. Finally, it is often convenient to work with the further quotient of $(\widehat{\text{CFK}}(\mathcal{H}), \widehat{\partial})$ that is obtained by setting the remaining formal variables $U_{w_j} = 0$. The result is known as the “tilde” version of knot Floer homology and its complex is denoted $(\widetilde{\text{CFK}}(\mathcal{H}), \widetilde{\partial})$. The associated homology

$$\widetilde{\text{HFK}}_*(K) := \text{H}_*(\widetilde{\text{CFK}}(\mathcal{H}), \widetilde{\partial})$$

is an invariant of the link K together with the number of basepoints (n_1, \dots, n_ℓ) in \mathbf{w} corresponding to each component K_1, \dots, K_ℓ of K . As \mathbb{F} -modules, we have that

$$\widetilde{\text{HFK}}(K) \simeq \widehat{\text{HFK}}(K) \otimes V_1^{n_1-1} \otimes \dots \otimes V_\ell^{n_\ell-1},$$

where V_i is a rank 2 vector space spanned by vectors v and w with multi-gradings $M(v) = A_{K_j}(v) = 0$ for all j , and $M(w) = A_{K_i}(w) = -1$, $A_{K_j}(w) = 0$ for $j \neq i$.

2.2. Combinatorial computations and grid diagrams. In [SW10], Sarkar and Wang described a combinatorial method for computing Heegaard Floer invariants via so-called nice diagrams.

Definition 2.1. A multi-pointed Heegaard diagram $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{z}, \mathbf{w})$ is called *nice* if every region in $\Sigma \setminus (\boldsymbol{\alpha} \cup \boldsymbol{\beta})$ not containing a z -basepoint is topologically a disk with at most 4

corners. In other words, every region in the complement of the α and β -curves either contains a z -basepoint or is a bigon or square.

If a multi-pointed Heegaard diagram \mathcal{H} is nice, Sarkar and Wang showed that the differential on $\text{CFK}^-(\mathcal{H})$ can be computed combinatorially by counting embedded, empty rectangles and bigons connecting generators. They further showed algorithmically how, through a sequence of handleslides and isotopies in the complement of basepoints, any multi-pointed Heegaard diagram can be transformed into one which is nice.

If $K \subset S^3$ is a knot or link in the 3-sphere, then K can be represented combinatorially via a *grid diagram*. A grid diagram $G = (n; \mathbb{X}, \mathbb{O})$ consists of the following:

- (1) An $n \times n$ square planar grid,
- (2) Collections $\mathbb{X} = (X_1, \dots, X_n)$ and $\mathbb{O} = (O_1, \dots, O_n)$ of X 's and O 's in the squares of the grid such that each row and each column contains exactly one X and one O , and no square is occupied by both an X and an O .

To such a grid diagram G , one associated a knot or link in S^3 as follows. First, draw (oriented) vertical line segments connecting the X 's and O 's in each column. Next, draw (oriented) horizontal line segments connecting the O 's and X 's so that the horizontal strands pass underneath the vertical strands. If we allow the X 's and O 's to play the roles of z 's and w 's respectively, then grid diagrams are clearly nice in the sense of Sarkar and Wang, and thus provide an avenue for combinatorially computing knot Floer homology.

In the context of Heegaard Floer theory, grid diagrams first appeared in the work of Manolescu, Ozsváth and Sarkar [MOS09], and Manolescu, Ozsváth, Szabó and Thurston [MOST07]. Since their initial appearance, grid diagrams have proven an essential tool for computing knot Floer homology and studying its applications to problems in topology, and in contact and symplectic geometry.

2.3. Braids and their associated Heegaard diagrams. Recall the well-known isomorphism between the braid group B_n and the mapping class group $\mathcal{M}(D, n)$ of the disk with n punctures.

The braid group B_n is the finitely presented group

$$B_n := \left\langle \sigma_1, \dots, \sigma_{n-1} \left| \begin{array}{l} \sigma_i \sigma_{i+1} \sigma_i = \sigma_{i+1} \sigma_i \sigma_{i+1} \\ \sigma_i \sigma_j = \sigma_j \sigma_i \quad \text{if } |i - j| \geq 2 \end{array} \right. \right\rangle$$

This presentation of B_n is the *Artin presentation* and the elements $\sigma_1, \dots, \sigma_{n-1}$ are *Artin letters*.

Let $Q = \{1, 2, \dots, n\}$ be a collection of n points in \mathbb{C} and let D be an embedded disk in \mathbb{C} containing Q . A self-homeomorphism ϕ of the pair (D, Q) is an orientation-preserving homeomorphism $\phi : D \rightarrow D$ that fixes Q setwise and ∂D pointwise. The *mapping class group* $\mathcal{M}(D, n)$ is the group of isotopy classes of self-homeomorphisms ϕ , with group multiplication given by composition of maps.

A *spanning arc* is a properly embedded arc γ in the interior of D such that the boundary of γ is two disjoint points in Q and the interior of γ is disjoint from Q . A *half-twist* τ_γ along a spanning arc γ is the following self-homeomorphism of the disk. Choose a neighborhood $\nu(\gamma)$

of γ in D disjoint from the other points in Q and an orientation-preserving identification ψ of $\nu(\gamma)$ with the disk $B = \{|z| \leq 3\} \subset \mathbb{C}$ that sends γ to the interval $[-1, 1]$. Choose a continuous function $g : \mathbb{R} \rightarrow \mathbb{R}$ with $g(x) = 0$ for $x \geq 3$ and $g(x) = -1$ for $x \leq 2$. Let $f : \mathbb{C} \rightarrow \mathbb{C}$ be the self-homeomorphism of \mathbb{C} defined by $f(z) = e^{g(|a|)i\pi}z$. Then the half-twist τ_γ along γ is defined to be $\tau_\gamma = \psi^{-1} \circ f \circ \psi$. Note that a half-twist is determined by the image of a curve that transversely intersects the spanning arc γ . See Figure 1.




FIGURE 1. A half-twist τ_γ along the spanning arc γ

Let $\gamma_i = [i, i+1] \times \{0\}$ denote the spanning arc connecting the i^{th} and $(i+1)^{\text{st}}$ basepoint in Q . Then there is a group isomorphism

$$B_n \simeq \mathcal{M}(D, n)$$

given by the map $\sigma_i \leftrightarrow \tau_{\gamma_i}$ that sends the i^{th} Artin generator to the half twist along the i^{th} spanning arc.

Now, let L be a link in S^3 , $\delta \in B_n$ a n -strand braid whose closure is L , and $\phi \in \mathcal{M}(D, n)$ the mapping class corresponding to δ .

In the disk D , label the basepoints $z_j = (j, 0)$ and choose a *basis* \mathbf{a} for (D, \mathbf{z}) , that is, a collection $\{a_1, \dots, a_{n-1}\}$ of properly embedded arcs in D such that each component of $D \setminus \mathbf{a}$ contains exactly one z basepoint. Index \mathbf{a} so that z_j lies in the component bounded by $\partial D, a_{j-1}, a_j$. Let \mathbf{b} be a second basis where each b_j is obtained by pushing off a_j along the orientation of ∂D and isotoping so that a_j and b_j intersect transversely in a unique point.


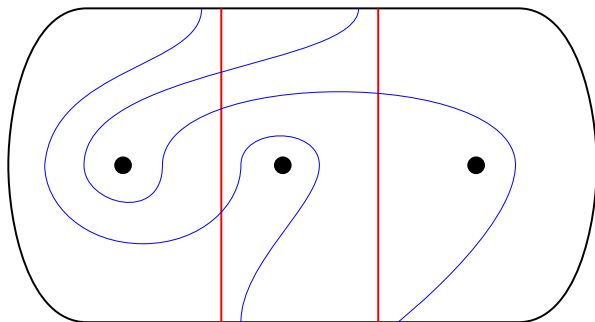
Let D' denote a second copy of this disk, with identical basepoints $\mathbf{w} = \{w_j\}$ and first basis $\mathbf{a}' = \{a'_j\}$. Let \mathbf{b}' be a second basis defined by setting $b'_j = \phi(b_j)$. We can assume, after possibly a perturbation of ϕ , that each pair a'_i and b'_i intersect transversely.

Let $\Sigma = D \cup -D'$ and for $i = 1, \dots, n$, let $\alpha_i = a_i \cup a'_i$ and $\beta_i = b_i \cup b'_i$.

Lemma 2.1. *For an oriented link L , let $\delta \in B_n$ be a braid whose closure is L and let $\phi \in \mathcal{M}(D, n)$ be the mapping class identified with δ . Then $\mathcal{H}_\phi = (\Sigma, \alpha, \beta, \mathbf{z}, \mathbf{w})$ is a multi-pointed Heegaard diagram for $-L$.¹*

Remark 2.2. It is often useful to consider the “dual” diagram obtained by using ϕ^{-1} to reparametrize Σ . Diagrammatically, this is accomplished by instead choosing bases $\mathbf{a}' = \{a'_i = \phi^{-1}(a_i)\}$ and $\mathbf{b}' = \{b'_i = b_i\}$ for D' before constructing \mathcal{H}_ϕ . In general, modifying the \mathbf{b}' by some ϕ is equivalent to instead modifying the \mathbf{a}' by ϕ^{-1} .

¹That the Heegaard diagram \mathcal{H}_ϕ specifies the oriented link $-L$ as opposed to L arises as a consequence of our choice to place the z -basepoints on D . This was done to allow for future applications to the study of transverse knot theory, a topic we hope to return to later.

FIGURE 2. The disk D with bases \mathbf{a}, \mathbf{b} FIGURE 3. The disk $-D'$ for $n = 3$ and $\delta = \sigma_2\sigma_1$

We call a multi-pointed Heegaard diagram \mathcal{H}_ϕ obtained in the above fashion a *braid diagram*. In the context of Heegaard Floer theory, such diagrams first appeared in the work of Baldwin, Vértesi and the third author in their work establishing an equivalence relating certain transverse invariants in knot Floer homology [BVV13]. There, it is shown that the distinguished collection of intersections \mathbf{x} lying on the disk $D \subset \Sigma$ together define a cycle in $\text{CFK}^-(\Sigma, \boldsymbol{\beta}, \boldsymbol{\alpha}, \mathbf{z}, \mathbf{w})$ whose associated homology class $[\mathbf{x}]$ is an invariant of the transverse knot of link type of the associated braid.

3. THE ALGORITHM

In order to algorithmically compute the Knot Floer Homology of an oriented link $L \subset S^3$, we need to construct a multi-pointed Heegaard diagram for L that is nice, in the terminology of Sarkar-Wang.

Starting with a braid $B \in B_n$ whose closure is $-L$ and a braid word w in the Artin generators representing B , our approach is to

- (1) Use the braid word w to construct a specific, “efficient” self-homeomorphism of the disk with n punctures whose mapping class corresponds to B

- (2) Use this self-homeomorphism to construct an efficient multi-pointed Heegaard diagram for the original link L , as described in Subsection 2.3
- (3) Apply an appropriately modified version of the Sarkar-Wang algorithm to make the diagram nice.

3.1. Nice diagram. In this first subsection, we will describe the Heegaard diagram abstractly before describing how to algorithmically obtain the diagram in Subsection 3.2.

Let A_i and B_i denote the components of $\Sigma \setminus \alpha$ and $\Sigma \setminus \beta$, respectively, containing the i^{th} basepoint z_i . Thus, $\partial A_i = \alpha_i - \alpha_{i-1}$ and $\partial B_i = \beta_i - \beta_{i-1}$.

A *trivial bigon* is a region in the disk D' , disjoint from any basepoint w_i , whose boundary splits into two connected components, one a segment of an α -curve and the other a segment of a β -curve. Such a bigon can be eliminated by an isotopy of the α, β curves as it contains no basepoint to obstruct this isotopy. A monodromy ϕ that produces no trivial bigons in the Heegaard diagram is called *efficient*.

Choosing an efficient ϕ ensures that the diagram \mathcal{H}_ϕ is already very close to being nice.

Lemma 3.1. *Every elementary region of \mathcal{H}_ϕ not containing some z_j is a $2k$ -gon. Furthermore, suppose that ϕ is efficient and let $i = 1, \dots, n$. Then among the elementary regions of A_i , there is at most one 6-gon and no $2k$ -gons for $k \geq 4$. Similarly, among the elementary regions of B_i , there is at most one 6-gon and no $2k$ -gons for $k \geq 4$.*

Proof. By construction, A_1, B_1, A_n, B_n are topological disks and A_i, B_i are annuli for $i = 2, \dots, n-1$. Also, no α_i is contained within a single B_j and no β_j is contained within a single A_i .

Let R be an elementary region, contained in some A_i and in some B_j . Thus, segments of the boundary of R can only form parts of four curves: $\alpha_{i-1}, \alpha_i, \beta_{j-1}, \beta_j$. Moreover, each boundary component must have α and β edges.

Note that each α and β can appear as an edge in at most 1 boundary component. Since $\Sigma = S^2$, there is a 1-1 correspondence between the boundary components of R and connected components of $\Sigma \setminus R$. Therefore, a path in Σ between boundary components of R must be contained in R . Because each α, β is closed, if it forms part of two boundary components, it must lie entirely within R . But since R is elementary, that curve is actually the totality of a single boundary component.

Thus, R has either 1 or 2 boundary components and can be either a topological disk or an annulus.

Suppose that R is an annulus and let ϵ be an properly embedded arc in A_i connecting α_{i-1} to α_i . Then either ϵ is contained in R or ϵ intersects the β edges of ∂R at least once. There is clearly an arc in the disk D connecting α_{i-1} to α_i that passes through z_i and is disjoint from every β curve. Hence R must contain z_i .

Now suppose R is a $2k$ -gon. If some β edge of ∂R connects an α curve to itself, the union of this edge and some edge of that α forms a bigon. Since ϕ is efficient, that bigon must contain a basepoint and that basepoint must be w_i . Since w_i is unique, for each pair of bigons formed

in this way, one must either contain the other. If there are multiple such bigons, then R must be contained in one and R is a square.

If R is not a bigon or such a square, then exactly two β edges of ∂R connect α_{i-1} and α_i . There can be at most one more β edge connecting some α to itself, so R is either a square or a 6-gon.

Finally, suppose that there are two elementary 6-gons R_1, R_2 within A_i . Then one, say R_2 , must be contained in the nontrivial bigon associated to R_1 . However, this is a contradiction since R_2 must therefore be a square.

The exact same argument with A_i replaced by B_j and the α 's and β 's exchanged proves the statement for B_j . \square

Remark 3.2. In the terminology of Sarkar-Wang, this implies that the maximal badness of any region is 1 and that there are at most $n - 2$ bad regions.

Now, following the approach of [HKL07], we use the following *stabilization trick* to make the Heegaard diagram nice.

Definition 3.1. Let \mathcal{H}_ϕ be a Heegaard diagram with an empty, elementary 6-gon $R \subset A_i$. The *stabilization trick* consists of

- (1) Stabilize \mathcal{H}_ϕ by attaching a 1-handle to Σ with the attaching sphere given by two points, one in R and the other in a z -basepointed region
- (2) Apply the Sarkar-Wang algorithm

More specifically, choose a point x in R near some α boundary edge and a pushoff of that α curve containing x . Let $\widehat{\alpha}_i$ be an oriented subsegment of this pushoff disjoint from the bigon containing the basepoint w_i and whose boundary is $x \cup y$ for some point y in the region containing z_i . Then attach a 1-handle to Σ with attaching sphere $x \cup y$ and extend $\widehat{\alpha}_i$ across the 1-handle to a closed curve. Choose $\widehat{\beta}_i$ to be a meridian or belt-sphere of this 1-handle. It is isotopic to a pushoff of the boundary of the 6-gon.

Finally, perform “finger moves” by pushing $\widehat{\beta}_i$ across the three α edges and continue until reaching basepointed regions. This is easiest to see in the “dual” diagram in Figure 4b.

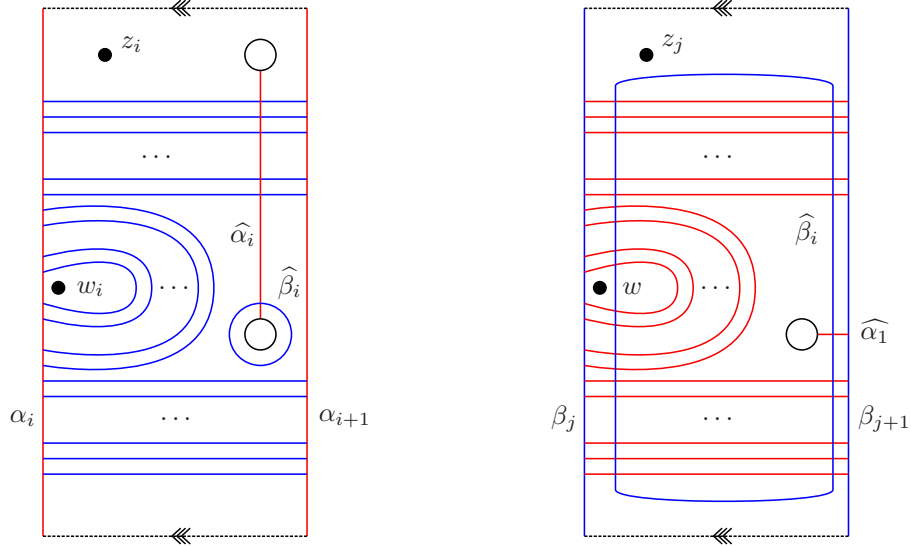
Suppose that \mathcal{H}_ϕ has g elementary 6-gons. Let $\mathcal{H}_\phi^{\text{nice}} = (\Sigma_g, \alpha \cup \widehat{\alpha}, \beta \cup \widehat{\beta}, \mathbf{z}, \mathbf{w})$ be the Heegaard diagram obtained by applying the stabilization trick to each 6-gon.

Proposition 3.3. *The Heegaard diagram $\mathcal{H}_\phi^{\text{nice}}$ is nice.*

Proof. It is clear that all elementary regions in each B_i are either bigons or 4-gons, except possibly for a unique 6-gon. However, via the attached 1-handle, this 6-gon is identified with the elementary region containing the basepoint z_i .

Furthermore, note that unless w_i and z_i lie in the same elementary region, each region containing a w_i basepoint is either a bigon or a square. \square

3.2. Algorithm. In this subsection, we will explicitly describe the nice Heegaard diagram and compute the time required to obtain the diagram.



(A) The annulus A_i after stabilization and before finger moves (B) The “dual” diagram of B_j after finger moves

FIGURE 4. The Stabilization Trick

A *braid word* w for a braid $\delta \in B_n$ is a string $\sigma_{i_1}^{s_1} \sigma_{i_2}^{s_2} \dots \sigma_{i_k}^{s_k}$, where $\sigma_{i_j} \in \{\sigma_1, \dots, \sigma_{n-1}\}$ and $s_j \in \{\pm 1\}$, of Artin letters whose product equals δ in B_n . The integer $k = l(w)$ is called the *length* of the word w .

Throughout this section, the algorithm depends upon the explicit word w chosen for a braid δ , although the final Heegaard diagram does not. For this reason, we will use \mathcal{H}_w to denote Heegaard diagram.

In addition to the setup and choice of bases \mathbf{a}, \mathbf{b} from Subsection 2.3, we choose specific spanning arcs γ_i for $i = 0, \dots, n$. Recall that to identify B_n and $\mathcal{M}(D, n)$, we chose spanning arcs $\gamma_i = [i, i + 1]$ connecting the basepoints. We can assume that the boundary ∂D passes through the points $(0, 0)$ and $(n + 1, 0)$, and let $\gamma_0 = [0, 1]$ and $\gamma_n = [n, n + 1]$ be arcs that connect the 1st and (n) th basepoints to the boundary of the disk. Furthermore, we can choose \mathbf{a} such that $\alpha_i \cap \gamma_j$ is empty if $i \neq j$ and is exactly one point if $i = j$. Finally, we assume that there are no triple intersection points between α, β, γ curves.

The spanning arc γ_i cuts a'_i into two components, a_i^+ and a_i^- , which are contained in the upper-half and lower-half planes, respectively, and a_i cuts the spanning arc γ_i into two components γ_i^\pm . Let A_i^\pm denote the number of intersections $|a_i^\pm \cap \beta|$ between β curves and the components of $a'_i \setminus \gamma_i$ and define $A_i = A_i^+ + A_i^-$. Similarly, let Γ_i^\pm denote the number of intersections $|\Gamma_i^\pm \cap \beta|$ between β curves and the components of $\gamma_i \setminus \alpha_i$ and define $\Gamma_i = \Gamma_i^+ + \Gamma_i^-$.


Definition 3.2. The *complexity* of a multi-pointed Heegaard diagram \mathcal{H} is a tuple (n, v) , where $n - 1$ is the number of α curves and $v = |\alpha \cap \beta|$ is the total number of intersections between α and β curves.

Finally, we need to set a unit of time. A Heegaard diagram \mathcal{H} can be represented combinatorially by using $\alpha, \beta, \gamma, \mathbf{w}$ to determine a handle decomposition of Σ . Specifically, there is a unique 0-handle for each intersection point between some pair of α, β or γ curves and a unique 0-handle for each w -basepoint; a unique 1-handle for each segment of an α, β or γ curve between intersection points or w -basepoint; and a unique 2-handle for each elementary region. Each intersection point is 4-valent, so introducing or eliminating an intersection point requires a constant number of modifications of the handle decomposition. Thus, we declare that adding or removing a single intersection point between a pair of α and β curves or a pair of γ and β curves takes $O(1)$ time.

Given a braid word $\sigma_{i_1}^{s_1} \dots \sigma_{i_{l(w)}}^{s_{l(w)}}$ for a braid δ , it is straightforward to build the self-homeomorphism ϕ corresponding to δ by starting from a diagram \mathcal{H}_1 of the identity map and successively applying the half-twists $\tau_{i_j}^{s_j}$. After each step, we will also immediately remove any trivial bigons formed by pairs of α and β curves or pairs of γ and β curves.

Base case. The mapping class corresponding to the trivial braid 1 is the identity id , so D' is just a second copy of D (see Figure 2). The diagram \mathcal{H}_1 has complexity $(n, 2(n-1))$ and takes $O(n)$ time to construct. We choose it so that $A_i^+ = \gamma_i^- = 1$ and $A_i^- = \gamma_i^+ = 0$ for all $i = 1, \dots, n-1$.

Inductive step. Figure 5 describes the effect of the half-twist τ_i on the lower-half disk $-D'$ of \mathcal{H}_w . The half-twist is the identity outside of a neighborhood of the spanning arc γ_i and we can assume that in this neighborhood, there are no α, β intersections. For the following discussion, we restrict attention to the positive half-twist as the case of the inverse half-twist is identical up to mirror image.



(A) A neighborhood of γ_i before the half-twist (B) A neighborhood of γ_i after the half-twist

FIGURE 5. A half-twist γ_i

Consider what happens to a β strand that intersects γ_i^- under the half-twist. It now intersects first γ_{i-1}^+ , then a_i^- , then γ_i^+ , then γ_{i+1}^- , then finally a_i^- again before leaving $\nu(\gamma_i)$. On the other side, a β strand that intersected γ_i^+ now intersects first a_i^+ , then γ_{i-1}^+ , then γ_i^- , then a_i^+ , then finally γ_{i+1}^- before leaving $\nu(\gamma_i)$.

In order to see how new trivial bigons can be created, consider instead the “dual” diagram of a half-twist, in which we apply τ_i^{-1} to the α, γ curves instead in Figure 6. New trivial

bigons appear if prior to applying the half-twist, a β strand, after it leaves $\nu(\gamma_i)$, continues counterclockwise and intersects the next α and/or γ segments.




FIGURE 6. The possible trajectories of a β arc as it leaves $\nu(\gamma_i)$

For example, a β curve leaving $\nu(\gamma_i)$ at the top-left, between γ_{i-1} and a_i^+ , can next hit (in counter-clockwise order) a_i^+ , $\partial D'$, a_{i-1}^+ , γ_{i-1}^+ and if it crosses the spanning arc γ_{i-1} , it can next hit a_{i-1}^- , $\partial D'$, a_i^- . (We can assume, by choosing ν small, that this β segment does not enter $\nu(\gamma_i)$ again).

For each original β arc in $\nu(\gamma_i)$, the half-twist τ_{γ_i} can introduce at most two trivial bigons, one with α_i and one with $\gamma_{i\pm 1}$, in each direction it leaves $\nu(\gamma_i)$ for a total of at most four new trivial bigons.

For $w = w' * \sigma_i^*$, let $\mathcal{H}_{w'\sigma_i^*}$ denote the Heegaard diagram obtained by applying the half-twist $\tau_{\gamma_i}^*$ and then removing any trivial bigons.

We can now describe the complexity of the diagram and the time to obtain the diagram.

Lemma 3.4. *Suppose that the complexity of \mathcal{H}_w is (n, v) . Then the complexity of $\mathcal{H}_{w\sigma_i^*}$ is at most $(n, v + 12(\Gamma_i))$ and it can be obtained from \mathcal{H}_w in $O(|\Gamma_i|)$. The complexity of \mathcal{H}_w is at most $(n, 4^{l(w)} + 2n)$ and it can be obtained in $O(4^{l(w)} + n)$.*

Proof. It's clear that performing this half twist introduces $2\Gamma_i$ new intersections between α_i and the β curves. Moreover, it introduces Γ_i new intersections between γ_{i-1}^+ , γ_{i+1}^- and the β curves. Removing the trivial bigons requires eliminating 2 vertices for at most $4\Gamma_i$ new trivial bigons.

The first pair of facts follow since at most $12\Gamma_i$ new vertices are introduced or eliminated, each of which takes $O(1)$ times. The second pair follows since at each step, $\Gamma_{i-1}, \Gamma_{i+1}$ each increase by $\Gamma_i < \max_j \Gamma_j$. The maximum among the Γ_j can at most double each iteration. \square

Finally, we need to make the diagram nice.

Proposition 3.5. *The nice diagram $\mathcal{H}_w^{\text{nice}}$ has complexity at most $(2n - 2, 4^{l(w)+1} + 8n)$ and can be obtained in $O(4^{l(w)})$.*

Proof. Attaching a handle in the i^{th} annulus introduces at most A_i intersections between $\widehat{\alpha}_i$ and the β curves and introduces at most $B_j + B_{j-1}$ intersections between $\widehat{\beta}_j$ and the α curves. In other words, we introduce at most 3 new intersection points for each original intersection point. Since the new number of intersection points is a constant multiple of the old number of intersection points, the time complexity is the same as in Lemma 3.4. \square

Remark 3.6. This bound on diagram complexity may overestimate the number of new intersections of β curves and spanning arc. It should be possible to choose a lower base in the exponential bound. However, the number of vertices can grow exponentially in the braid word length. For instance, the complexity of the diagram for the pseudo-Anosov braid $(\sigma_2^{-1}\sigma_1)^k$ grows exponentially in k .

4. COMPUTING HOMOLOGY

In this section, we will describe how to compute knot Floer homology from a nice, multi-pointed Heegaard diagram. The approach here is fairly straightforward and there are many techniques to improve the speed of computation. However, we are only interested in establishing the basic qualitative properties of our approach. The discussion draws from [HKL07] but is adapted to our particular Heegaard diagram.

Throughout this section, we consider Heegaard diagrams $\mathcal{H}_w^{\text{nice}}$ of complexity $(n+1, v)$.

First, we establish a bound on the size of the chain complex. Recall that the knot Floer complexes are generated by the intersection points $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$.

Lemma 4.1. *Let \mathcal{H} be a Heegaard diagram of complexity $(n+1, v)$. Then the number of generators of the knot Floer complex is at most $\left(\frac{v}{n}\right)^n$.*

Proof. The number of generators is the permanent of the matrix M whose $(i, j)^{\text{th}}$ entry is the number of intersection points between α_i and β_j . It is an easy exercise to prove that, restricted to the level set $\sum_{i,j} M_{i,j} = v$, the maximum of the permanent function is $\frac{v}{n} * \text{Id}_n$ and the minimum is $\frac{v}{n^2} * [1]$, where Id_n is the $n \times n$ identity matrix and $[1]$ is the $n \times n$ matrix with 1's in every entry. \square

In order to compute the Alexander and Maslov gradings of the generators and to identify the holomorphic disks in the differential, we need to find chains in the Heegaard diagram connecting pairs of generators. To accomplish this, we will use a linear map associated to the Heegaard diagram as follows.

Let R denote the \mathbb{R} -vector space generated by elementary regions $r \in \Sigma \setminus (\alpha \cup \beta)$. Let C denote the \mathbb{R} -vector space generated by the α curves $\{\alpha_1, \dots, \alpha_{n-1+g}\}$ and all but one of the β curves $\{\beta_1, \dots, \beta_{n-2+g}\}$. Let V denote the \mathbb{R} -vector space generated by the vertices of the Heegaard diagram. These spaces each have an inner product $\langle \cdot, \cdot \rangle$ defined by choosing the above bases to be orthonormal.

Let $\rho = \dim R$ be the number of regions, $v = \dim V$ the number of vertices, and e the number of edges in the Heegaard diagram. Then

$$2 - 2g = \rho - e + v$$

and since each vertex is 4-valent, the identity $2e = 4v$ implies that

$$(1) \quad 4 - 4g = 2\rho - 3v$$

Let \mathcal{D} be the linear map from $R \oplus C \rightarrow V$ that assigns to each region r the signed sum of vertices incident to r and to each α or β curve to the sum of vertices along it. The sign of a vertex v is positive if, traveling along the boundary of r in the direction induced by the orientation, one jumps from a β curve to an α curve at v , and the sign is negative otherwise.

By abuse of notation, let A_i and B_i denote the sum in R of elementary regions constituting the connected components A_i of $\Sigma \setminus \alpha$ and B_i of $\Sigma \setminus \beta$, respectively, containing the basepoint z_i .

A domain in Σ is *periodic* if its boundary is the union $\coprod_i k_i \alpha_i \coprod_i l_i \beta_i$ of some number of α and β curves. The domains A_i, B_j and sums of these domains are periodic.

Lemma 4.2. *\mathcal{D} is surjective and the kernel of \mathcal{D} is spanned by A_1, \dots, A_n and B_1, \dots, B_{n-1} .*

Proof. The statement about the kernel follows from two facts. First, $\mathcal{D}|_C$ is clearly injective and the image $\mathcal{D}(R)$ is perpendicular to the image $\mathcal{D}(C)$ with respect to the inner product since the signs of corners of a region r formed by some α or β cancel in pairs. Therefore, the kernel of \mathcal{D} lies in R .

Secondly, the kernel of \mathcal{D} in R exactly corresponds to periodic domains. It follows that since the ambient manifold is S^3 , the periodic domains are generated by the domains A_i, B_j subject to the relation $B_n = \sum_{i=1}^n A_i - \sum_{j=1}^{n-1} B_j$.

Since the nullity of \mathcal{D} is $2n - 1$ and the identity $4g - 4 = \rho - 3v$ in Equation 1 holds, this implies that \mathcal{D} is surjective. \square

Let \mathcal{W} denote the left inverse of \mathcal{D} . Then \mathcal{W} maps V bijectively onto $\ker(\mathcal{D})^\perp$ and $\mathcal{W}\mathcal{D}$ is the identity on $\ker(\mathcal{D})^\perp$.

Each generator \mathbf{x} of the knot Floer chain complex can be identified with a vector $x = x_1 + \dots + x_{n-1+g}$ in V . For a pair of generators, the map \mathcal{W} thus determines a fractional chain on the surface of the Heegaard diagram connecting the two generators.

Since \mathcal{D} is not injective, there is some ambiguity to recovering Whitney disks using the map \mathcal{W} . However, it follows from the proof of Lemma 4.2 that the only ambiguity is if \mathbf{x}, \mathbf{y} contain distinct vertices x', y' in the intersection of the same pair α_i, β_j of curves. Then the chain $P_{i,j} = \sum_{k \leq i} A_k - \sum_{k \leq j} B_k$ is the difference of two distinct chains connecting x' to y' . The inverse map \mathcal{W} will count each of these chains with coefficient $\frac{1}{2}$ and therefore the coefficients in $\mathcal{W}(x - y)$ lie in $\frac{1}{2}\mathbb{Z}$. To remove this ambiguity, we can consider the chains $\mathcal{W}(x - y) \pm \frac{1}{2}P_{i,j}$.

If the C components of $\mathcal{W}(x - y)$ are 0 and the R components are integers, then $\phi = \mathcal{W}(x - y)$ is a 2-chain and represents a Whitney disk in $\pi_2(\mathbf{x}, \mathbf{y})$. In particular, if the R components of the vector $\mathcal{W}(x - y)$ are either 1 or 0 and the C components are 0, then there is an embedded Whitney disk connecting \mathbf{x} to \mathbf{y} . If we correct $\mathcal{W}(x - y)$ by adding the appropriate periodic domains $\pm \frac{1}{2}P_{i,j}$ for all such pairs x', y' , the collection gives exactly the Whitney disks connecting \mathbf{x} to \mathbf{y} . Let $\pi_2^0(\mathbf{x}, \mathbf{y})$ denote the collection of all such domains and note that there are at most $2n$ such domains.

Alexander grading. The relative Alexander grading of a pair \mathbf{x}, \mathbf{y} is determined by the formula:

$$A(\mathbf{x}) - A(\mathbf{y}) = n_z(\phi) - n_w(\phi)$$

Let Z, W be the vectors in R that are the sums of the basis vectors corresponding to regions with a z or w basepoint, respectively. Then, the relative Alexander grading is given by

$$A(\mathbf{x}) - A(\mathbf{y}) = \langle \phi, Z - W \rangle$$

for any $\phi \in \pi_2^0(\mathbf{x}, \mathbf{y})$.

Maslov grading. The relative Maslov grading of a pair \mathbf{x}, \mathbf{y} is determined by the formula

$$M(\mathbf{x}) - M(\mathbf{y}) = \mu(\phi) - 2n_w(\phi)$$

The Maslov index of the Whitney disk ϕ can be computed according to Lipshitz's formula [Lip06]

$$\mu(\phi) = e(\phi) + \mu_{\mathbf{x}}(\phi) + \mu_{\mathbf{y}}(\phi)$$

where $e(\phi)$ is the Euler measure of the domain ϕ and $\mu_{\mathbf{x}}, \mu_{\mathbf{y}}$ are point measures. The *Euler measure* of a domain ϕ is $\frac{1}{2\pi} \int_{\phi} \omega_g$ where ω_g is the curvature of a metric g on Σ for which all α, β curves are geodesics and always intersect at right angles. The Euler measure is clearly additive and the Euler measure of a $2n$ -gon is $1 - \frac{n}{2}$. The *point measure* $\mu_x(\phi)$ of a domain at a vertex x is the average of the coefficients of ϕ for the four elementary regions incident to the vertex x . The point measure $\mu_{\mathbf{x}}(\phi) = \sum \mu_{x_i}(\phi)$ with respect to a generator is the sum of point measures of ϕ at each vertex x_i of the generator.

Define $E \in R$ to be the vector whose r^{th} entry is the Euler measure $e(r)$. Then the relative Maslov grading is given by

$$M(\mathbf{x}) - M(\mathbf{y}) = \langle \phi, E \rangle + \langle |\mathcal{D}|(\phi), x + y \rangle - 2\langle \phi, W \rangle$$

for any $\phi \in \pi_2^0(\mathbf{x}, \mathbf{y})$.

Lemma 4.3. *The relative Alexander and Maslov gradings can be computed from $\mathcal{H}_w^{\text{nice}}$ in $O(v^3 + v^2n + n^2 + (v^2 + n) \left(\frac{v}{n}\right)^n)$.*

Proof. Gauss-Jordan elimination on a $k \times l$ matrix takes $O(k^2 + l^2k + l^3)$ and so computing the pseudoinverse \mathcal{W} takes $O(v^3 + v^2n + n^2)$. Obtaining $\mathcal{W}(x - y)$ takes $O(vn + n^2)$ and checking all $\mathcal{W}(x - y) \pm \frac{1}{2}P_{i,j}$ to see if they are domains takes a further $O(vn)$.

To compute the inner product $\langle \phi, Z - W \rangle$ takes $O(n)$ time since $Z - W$ has at most $2n$ nonzero entries. Thus it takes $O(n \left(\frac{v}{n}\right)^n)$ to compute all Alexander gradings.

Computing $\langle \phi, E \rangle$ takes $O(v)$, computing $|\mathcal{D}|(\phi)$ via matrix multiplication takes $O(v^2)$ and computing $\langle |\mathcal{D}|(\phi), x + y \rangle$ and $\langle \phi, W \rangle$ each take $O(n)$. Thus computing all Maslov gradings takes $O((v^2 + n) \left(\frac{v}{n}\right)^n)$. \square

Differential. Given a nice diagram, the Sarkar-Wang result [SW10] guarantees a unique holomorphic representative for each empty, embedded square or bigon on the Heegaard diagram connecting two generators. Moreover, these are the only possible holomorphic disks and so the differential can be computed in terms of these squares or bigons.

Thus, to compute the differential we only need to consider pairs of generators \mathbf{x}, \mathbf{y} that differ in at most two vertices. Specifically,

$$\begin{aligned} \mathbf{x} = (x_1, x_2, x_3, \dots, x_{n-1+g}) \quad \text{and} \quad \mathbf{y} = (y_1, y_2, x_3, \dots, x_{n-1+g}) \quad \text{or} \\ = (y_1, x_2, x_3, \dots, x_{n-1+g}) \end{aligned}$$

In the first case, no square region contains a \mathbf{w} basepoint, so the relative gradings satisfy $A(\mathbf{x}) = A(\mathbf{y})$ and $M(\mathbf{x}) = M(\mathbf{y}) + 1$ and there is no ambiguity in $\mathcal{W}(x - y)$. A Whitney disk $\phi \in \pi_2^0(\mathbf{x}, \mathbf{y})$ is unique and if it is embedded then it is empty. The U coefficients always vanish.

In the second case, there can be at most two disks $\phi \in \pi_2^0(\mathbf{x}, \mathbf{y})$ connecting \mathbf{x} to \mathbf{y} . Again, if \mathbf{x}, \mathbf{y} have the appropriate relative gradings then an embedded disk is empty. The U coefficients can be computed by taking the inner product $\langle \phi, W \rangle$ which takes $O(n)$ time.

Lemma 4.4. *The differential can be computed in $O(v \left(\frac{v}{n}\right)^{2n})$.*

Proof. It takes $O(v)$ to check whether a domain between a given pair of generators is embedded. \square

Homology. Having computed the relative gradings for all generators as well as the differential ∂^- , we obtain the chain complex CFK^- and can compute the homology.

Tabulating the combined time required to perform each of the above steps, we have the following proposition:

Proposition 4.5. *The homology HFK^- can be computed from a diagram $\mathcal{H}_w^{\text{nice}}$ of complexity (n, v) in $O\left(\left(\frac{v}{n}\right)^{3n}\right)$ time.*

Proof. Given the differential ∂^- , the homology HFK^- can be computed via Gauss-Jordan elimination in $O\left(\left(\frac{v}{n}\right)^{3n}\right)$. The proposition now follows by combining this fact with Lemmata 4.3 and 4.4 and the fact that $n < v \leq \left(\frac{v}{n}\right)^n$ \square

Theorem 1.1 now follows by combining Propositions 3.5 and 4.5.

5. EXAMPLE

As an example, we compare the computational complexity of this braid algorithm to the GRID algorithm for torus knots. The actual homology is known and is fairly straightforward to obtain. However, this class provides a good comparison to understand the algorithmic strengths of this braid approach.

Let $T(p, q)$ denote a torus knot, with $q > 0$. Then it follows from work of Etnyre and Honda [EH01] that the arc index, and equivalently the minimal grid size, of $T(p, q)$ is $|p| + q$. Using grid diagrams will therefore take $O((p + q)!^3)$ to compute the knot Floer homology.

The knot can be represented as the closure of the q -strand braid

$$B_{p,q} = (\sigma_1 \sigma_2 \cdots \sigma_{q-1})^p$$

which has length $(q - 1)^p$ in the Artin letters.

The upper bound in Lemma 3.4 significantly overestimates the time required to obtain the Heegaard diagram.

From now on assume that $p > 0$; the case of negative p follows similarly. Since each twist is positive, the number of α, β intersections and the number of generators of the chain complex grow monotonically. Note that if $p = kq$ for some integer k , then the braid is just k full twists. Hence, we can bound the complexity of the diagram and time complexity of the algorithm by studying full twists.

A full twist is equivalent to the following. Separate the q strands into two collections: the first i strands and the final $q - i$ strands. Perform a full twist on the first i strands then perform a full twist on the final $q - i$ strands. Finally, perform a full twist exchanging the two collections of strands.

Since the Heegaard diagram depends only on the braid, this makes it easy to see what happens to β_i .

Twists among the first i strands and final $q - i - 1$ strands do not affect β_i , but the full twist of the two groups of strands does. Each half twist draws β_i across each α curve exactly twice and so a full twist introduces 4 new intersections between β_i and each α curve. This is true for all i and so each full twist adds $4(q - i)^2$ new vertices to the diagram.

If p is not a multiple of q , then there will be some hexagonal regions that will need to be stabilized away, adding at most $q - 2$ new pairs of α, β curves and at most sextupling the number of vertices.

Therefore, it follows that the complexity of the Heegaard diagram $\mathcal{H}_{T(p,q)}^{nice}$ is at most $(2q - 3, 24(\frac{p}{q} + 1)(q - 1)^2)$ and it can be obtained in $O(pq + q^2)$ time.

We can also bound the number of generators fairly easily by using the fact that the permanent is multilinear and that each quantity $|\alpha_i \cap \beta_j|$ grows linearly in $\frac{p}{q}$. Thus the number of generators of the knot Floer complexes obtained from $\mathcal{H}_{T(p,q)}^{nice}$ is bounded by

$$\left(4\frac{p}{q}\right)^{q-1} \left(8\frac{p}{q}\right)^{q-2} (2q - 3)! \leq c^q p^{2q-3}$$

for some constant c not depending on p or q .

Combining these facts with Lemmata 4.3 and 4.4 and applying Gauss-Jordan elimination shows that $\text{HFK}^-(T(p, q))$ can be computed in $O((cp^2)^{3q}q^4)$ time.

Thus, for fixed q , the computational complexity grows polynomially in p , with degree determined by q . For $p \gg q$ the braid approach becomes significantly faster. However, for $p \sim q$, the grid diagram approach is faster as $(2q)! \leq q^{2q}$.

REFERENCES

- [Bal10] John A. Baldwin, *Comultiplication in link Floer homology and transversely nonsimple links*, *Algebr. Geom. Topol.* **10** (2010), no. 3, 1417–1436. MR 2661532 (2012d:57014)
- [BVV13] John A Baldwin, David Shea Vela-Vick, and Vera Vértési, *On the equivalence of Legendrian and transverse invariants in knot Floer homology*, *Geom. Topol.* **17** (2013), no. 2, 925–974. MR 3070518
- [EH01] John B. Etnyre and Ko Honda, *Knots and contact geometry. I. Torus knots and the figure eight knot*, *J. Symplectic Geom.* **1** (2001), no. 1, 63–120. MR 1959579 (2004d:57032)

- [Ghi08] Paolo Ghiggini, *Knot Floer homology detects genus-one fibred knots*, Amer. J. Math. **130** (2008), no. 5, 1151–1169. MR 2450204 (2010f:57013)
- [HKL07] Jonathan Hales, Dmytro Karabash, and Michael T. Lock, *A modification of the Sarkar-Wang algorithm and an analysis of its computational complexity*, Preprint, [arXiv:0711.4405](https://arxiv.org/abs/0711.4405) [math.GT], 2007.
- [Lip06] Robert Lipshitz, *A cylindrical reformulation of Heegaard Floer homology*, Geom. Topol. **10** (2006), 955–1097. MR 2240908 (2007h:57040)
- [MOS09] Ciprian Manolescu, Peter Ozsváth, and Sucharit Sarkar, *A combinatorial description of knot Floer homology*, Ann. of Math. (2) **169** (2009), no. 2, 633–660. MR 2480614 (2009k:57047)
- [MOST07] Ciprian Manolescu, Peter Ozsváth, Zoltán Szabó, and Dylan Thurston, *On combinatorial link Floer homology*, Geom. Topol. **11** (2007), 2339–2412. MR 2372850 (2009c:57053)
- [Ni07] Yi Ni, *Knot Floer homology detects fibred knots*, Invent. Math. **170** (2007), no. 3, 577–608. MR 2357503 (2008j:57053)
- [OS03] Peter Ozsváth and Zoltán Szabó, *Knot Floer homology and the four-ball genus*, Geom. Topol. **7** (2003), 615–639. MR 2026543 (2004i:57036)
- [OS04a] ———, *Holomorphic disks and topological invariants for closed three-manifolds*, Ann. of Math. (2) **159** (2004), no. 3, 1027–1158. MR 2113019 (2006b:57016)
- [OS04b] ———, *Holomorphic disks and knot invariants*, Adv. Math. **186** (2004), no. 1, 58–116. MR 2065507 (2005e:57044)
- [OS04c] ———, *Holomorphic disks and genus bounds*, Geom. Topol. **8** (2004), 311–334. MR 2023281 (2004m:57024)
- [OS08a] ———, *Holomorphic disks, link invariants and the multi-variable Alexander polynomial*, Algebr. Geom. Topol. **8** (2008), no. 2, 615–692. MR 2443092 (2010h:57023)
- [OS08b] ———, *Knot Floer homology and integer surgeries*, Algebr. Geom. Topol. **8** (2008), no. 1, 101–153. MR 2377279 (2008m:57075)
- [OS11] ———, *Knot Floer homology and rational surgeries*, Algebr. Geom. Topol. **11** (2011), no. 1, 1–68. MR 2764036 (2012h:57056)
- [OST08] Peter Ozsváth, Zoltán Szabó, and Dylan Thurston, *Legendrian knots, transverse knots and combinatorial Floer homology*, Geom. Topol. **12** (2008), no. 2, 941–980. MR 2403802 (2009f:57051)
- [Ras03] Jacob Rasmussen, *Floer homology and knot complements*, Ph.D. thesis, Harvard University, 2003, [arXiv:math/0306378](https://arxiv.org/abs/math/0306378) [math.GT].
- [SW10] Sucharit Sarkar and Jiajun Wang, *An algorithm for computing some Heegaard Floer homologies*, Ann. of Math. (2) **171** (2010), no. 2, 1213–1236. MR 2630063 (2012f:57032)

DEPARTMENT OF MATHEMATICS, LOUISIANA STATE UNIVERSITY

E-mail address: plambe7@lsu.edu

URL: <https://www.math.lsu.edu/~plambe7>

DEPARTMENT OF MATHEMATICS, LOUISIANA STATE UNIVERSITY

E-mail address: mston16@lsu.edu

URL: <https://www.math.lsu.edu/~mstone16>

DEPARTMENT OF MATHEMATICS, LOUISIANA STATE UNIVERSITY

E-mail address: shea@math.lsu.edu

URL: <https://www.math.lsu.edu/~shea>