

THE MURASUGI SUM AND SYMPLECTIC FLOER HOMOLOGY

By

Jared Able

A DISSERTATION

Submitted to
Michigan State University
in partial fulfillment of the requirements
for the degree of

Mathematics – Doctor of Philosophy

2022

ABSTRACT

This Ph.D. dissertation studies the operation of Murasugi sum on pairs of knots (or links) in the three-sphere S^3 . This operation, which produces a knot (or link) involves several choices which usually change the isotopy class of the produced object. This could lead one to believe the Murasugi sum has no hope of preserving properties of knots (or links). Indeed, without restricting the involved choices, we show in Chapter 2 that any knot is the Murasugi sum of any two knots. However, we also show that restricting the possible choices restricts the possible Murasugi sums. The contents of Chapter 2 are joint with Mikami Hirasawa [1].

Historically, the usual restriction has been to consider knots which are Murasugi summed along minimal genus Seifert surfaces. In this setting, knot genus is additive under Murasugi sums [11], and the rank of the “top” group of knot Floer homology \widehat{HFK} is multiplicative under Murasugi sums [4, 38]. Continuing this trend in Chapter 4, we study the symplectic Floer homology HF_* of a family of knots which are closures of a particular type of 3-string braid. These knots can be viewed as Murasugi sums performed along minimal genus Seifert surfaces, and we show that a large range of choices in these Murasugi sums all yield the same rank of HF_* . We carry out these calculations via the Bestvina-Handel algorithm [2] and the combinatorial formula for HF_* of pseudo-Anosov maps due to Cotton-Clay [5]. We hope that these calculations shed some light on the behavior of HF_* under Murasugi sums, as this group has garnered recent interest for its connection to the “next-to-top” group of knot Floer homology [15, 40].

Copyright by
JARED ABLE
2022

ACKNOWLEDGEMENTS

I would like to thank my advisor, Matt Hedden, for his willingness to listen to my ideas, and for his exceptional enthusiasm and breadth of knowledge. I would also like to thank Mikami Hirasawa, whose patience, as well as his concrete knowledge of Seifert surfaces and Murasugi sums, still astounds me. Finally, I would like to thank Emily, who continues to support and challenge me in this journey called life.

TABLE OF CONTENTS

CHAPTER 1	INTRODUCTION	1
CHAPTER 2	THE MURASUGI GRAPH OF KNOTS	5
CHAPTER 3	PSEUDO-ANOSOV MAPS AND THE BESTVINA-HANDEL ALGORITHM	20
CHAPTER 4	SYMPLECTIC FLOER HOMOLOGY	31
	BIBLIOGRAPHY	44
	APPENDIX	47

CHAPTER 1
INTRODUCTION

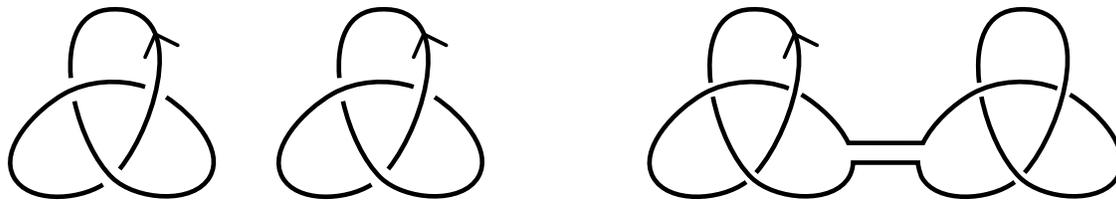


Figure 1.1. Two trefoil knots (left), and their connected sum (right).

A recurring theme in topology is to understand how manifolds behave under various operations. For example, a classical way to “add” two knots K_1, K_2 is the connected sum $K_1 \# K_2$, see Figure 1.1. This operation is well-behaved with respect to many geometric and algebraic invariants of knots such as the genus $g(K)$, the signature $\sigma(K)$, the Alexander polynomial $\Delta_K(t)$, and knot Floer homology $\widehat{\text{HF}}K(K)$ [30, 41]. For this reason, the connected sum and its various generalizations (e.g. satellite operations and band sums) have proven fruitful in constructing infinite families of distinct knots which all satisfy a given property, c.f. [6, 18, 19]. Historically, such constructions have motivated the search for increasingly refined invariants of knots.

We are mainly interested in the operation known as the Murasugi sum, which is a generalization of the connected sum of knots. Given two knots K_1, K_2 with Seifert surfaces F_1, F_2 satisfying $\partial F_1 = K_1, \partial F_2 = K_2$, the *Murasugi sum* $K_1 \star_m K_2$ identifies K_1, K_2 along an even-sided m -gon embedded in F_1, F_2 . See Figure 2.1, and see Definition 2.2.1 for a precise definition. This operation was introduced by Murasugi in [34] and was later coined as the Murasugi sum and popularized by Gabai in [11] via the slogan “The Murasugi sum is a natural geometric operation.” Indeed, Gabai showed that geometric invariants of Seifert surfaces such as fiberedness, incompressibility, and being of minimal genus are preserved under Murasugi sums.

Whereas the connected sum between two oriented knots is unique, “the” Murasugi sum of two knots is highly dependent on the choices of surfaces and embedded m -gons therein for

$m \geq 4$. In this thesis, we consider placing various restrictions on the choices involved in the Murasugi sum. In Chapter 2, we begin by studying the unrestricted behavior of Murasugi sums. As one might expect, allowing for all such choices yields interesting behavior. Indeed, we produce an efficient construction to justify the slogan “Any knot is the Murasugi sum of any two knots.” We then study the behavior of Murasugi sums when restricting the size of the m -gon. In particular, we obtain an obstruction to forming a knot as the Murasugi sum of two knots. In Chapter 3, we lay the groundwork necessary to understand pseudo-Anosov monodromies of fibered knots in terms of invariant train tracks obtained via the Bestvina-Handel algorithm [2]. This will allow us, in Chapter 4, to use work of Cotton-Clay [5] to calculate the rank of the symplectic Floer homology HF_* of the monodromies of a certain infinite family of fibered knots which we denote $K(j, |\beta_B| - j)$. A knot in this family is the closure of a particular homogeneous 3-string braid, and such a knot can be viewed as a Murasugi sum of a pair of closures of positive 2-string braids. One consequence of our calculation is that there are many ways to Murasugi sum such a pair of closures of 2-string braids which all yield the same rank of HF_* .

1.1 Knot theory background



Figure 1.2. A positive crossing (left), and a negative crossing (right).

The objects this thesis is concerned with are Seifert surfaces having knot boundaries. In this section, we recall some classical constructions and facts from knot theory. A *link* L of $|L|$ components is an embedding of $|L|$ disjoint oriented circles in the 3-sphere S^3 such that the image can be written as a finite union of line segments. A *knot* K is a link of one component. In a given diagram of L , we allocate to each crossing either $+1$ (a positive crossing) or -1 (a negative crossing) according to the convention indicated in Figure 1.2. While the sum of

all of these signs in some diagram of L fails to be an invariant of L , we have the following invariant of two-component links.

Definition 1.1.1. For a two-component link L with components L_1, L_2 , the *linking number* $lk(L_1, L_2)$ of L_1 and L_2 is half the sum of the signs, in a diagram for L , of the crossings at which one strand is from L_1 and the other is from L_2 .

The linking number has led to a wealth of calculable algebraic knot invariants, in part through its connection to Seifert surfaces. A *Seifert surface* F for a link L is an oriented, connected, compact surface with $\partial F = L$. Such surfaces always exist by Seifert's algorithm [30], see Figure 1.3, but they are never unique. Given a Seifert surface F for a knot K of genus $g = g(F)$, choose a representative basis of closed curves $\alpha_1, \dots, \alpha_{2g}$ for $H_1(F; \mathbb{Z}) \cong \mathbb{Z}^{2g}$. We form the *Seifert matrix* S as $S_{i,j} = lk(\alpha_i, \alpha_j^+)$, where $1 \leq i, j \leq 2g$ and α_j^+ denotes a small positive push-off of α_j . Up to the notion of so-called S -equivalence, the Seifert matrix is an invariant of K . A coarse knot invariant is given by the signature $\sigma(K)$ of K , which is the signature of $(S + S^T)$, i.e. the sum of the signs along the diagonal of a diagonalization of $(S + S^T)$. An often more refined invariant is the (symmetrized) Alexander polynomial $\Delta_K(t)$ of K , which is the determinant of $t^{-g}(S - tS^T)$ up to multiplication by $t^{\pm n}$. Even more refined is the knot Floer homology $\widehat{\text{HF}}K(K)$ of K , which categorifies $\Delta_K(t)$. This means that $\chi(\widehat{\text{HF}}K(K)) = \Delta_K(t)$ up to multiplication by $t^{\pm n}$, where χ denotes the graded Euler characteristic.

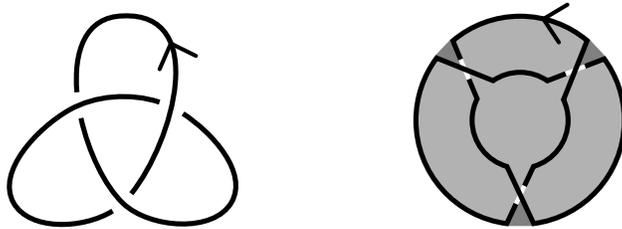


Figure 1.3. Forming a Seifert surface for a trefoil knot.

More geometrically, the *genus* $g(K)$ of a knot K is the minimum of the genera of all Seifert surfaces for K . While computing $g(K)$ from its definition alone is essentially impossible, a variety of invariants have been developed to calculate or bound $g(K)$, one such being $\widehat{\text{HF}}K$.

A minimal genus Seifert surface for a link is not usually unique, but it is unique (up to isotopy) if our link is fibered [9, 25, 39].

Definition 1.1.2. A link is called *fibered* if there is a fibration $\pi : S^3 - L \rightarrow S^1$ with fiber a Seifert surface F for L , the so-called *fiber surface* of L .

For L fibered with fiber surface F , the complement of a neighborhood of L in S^3 is diffeomorphic to the mapping torus M_ϕ of a diffeomorphism $\phi : F \rightarrow F$, i.e.

$$S^3 \setminus \nu(L) \cong M_\phi := \frac{F \times [0, 1]}{(x, 0) \sim (\phi(x), 1)}.$$

The isotopy class represented by ϕ is called the *monodromy* of L , and the monodromy is well-defined up to conjugation by a diffeomorphism of F .

We are particularly interested in knots which are Murasugi sums of Hopf bands, as these knots have monodromies which are usually easy to describe. We begin by noting that the monodromy of the positive Hopf band is a right Dehn twist about its core curve, and the monodromy of a negative Hopf band is a left Dehn twist about its core curve. See Figure 1.4. The behavior of monodromies under Murasugi sums is given as follows in [12].

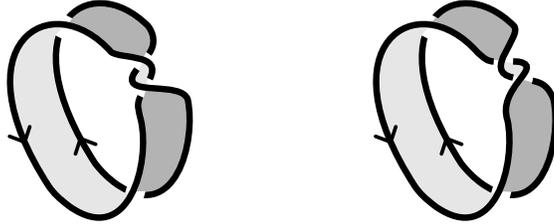


Figure 1.4. The positive Hopf band (left), and the negative Hopf band (right).

Theorem 1.1.1. *Suppose that F is a Murasugi sum of F_1 and F_2 , where $\partial F = L, \partial F_i = L_i$ for $i = 1, 2$, and L_i is a fibered link with monodromy ϕ_i such that $\phi_i|_{L_i} = \text{id}$. Then L is a fibered link with fiber surface F and monodromy $\phi = \phi'_2 \circ \phi'_1$, where $\phi'_i|_{F_i} = \phi_i$ and $\phi'_i|_{F-F_i} = \text{id}$ for $i = 1, 2$.*

The order of composition is determined by applying the monodromy of the “upper” surface F_1 first, and then the monodromy of the “lower” surface F_2 . See the discussion following Definition 2.2.1.

CHAPTER 2

THE MURASUGI GRAPH OF KNOTS

2.1 Introduction

One development in knot theory is to define and study the structure of the topological space \mathcal{S} composed of isotopy classes of knots. In [20], the *Gordian Complex* \mathcal{G} of knots was defined as follows: The vertex set of \mathcal{G} consists of isotopy classes of knots, and a set of $n + 1$ vertices K_0, \dots, K_n spans an n -simplex if and only if any pair of knots in it can be changed into each other by a single crossing change. Since then, many studies have been done by replacing the crossing change with other local operations on knots and on virtual knots as in [16, 21, 23, 36, 47].

In the following, we consider the operation of Murasugi sum along Seifert surfaces of links. An oriented, embedded surface F without a closed component is called a *Seifert surface* for an oriented link L if its boundary ∂F coincides with the oriented link L . An m -Murasugi sum is an operation to glue two Seifert surfaces F_1 and F_2 along an m -gon with m even (for the precise definition, see Definition 2.2.1).

We may regard a fixed knot K as an operation on the space of knots as follows. A knot K_1 is changed to a knot K_2 via K if a Seifert surface for K_2 is obtained from a Seifert surface F_1 for K_1 by Murasugi-summing a Seifert surface F for K to F_1 . Thus for each fixed K , the space of knots \mathcal{S} has the structure of directed graph, where an edge is an arrow or a double-headed arrow.

Definition 2.1.1. For a knot K , the Murasugi sum graph of knots $MSG(K)$ is a directed graph such that (1) the vertex set consists of all isotopy classes of knots, (2) two vertices K_1 and K_2 are connected by an edge with arrowhead on K_2 if there exist Seifert surfaces F, F_1, F_2 for K, K_1, K_2 such that F_2 is a Murasugi sum of F and F_1 . The *restricted Murasugi sum graph* $MSG(K, n)$ is considered by only allowing Murasugi sums along m -gons with $m \leq n$.

For $n = 2$, a 2-Murasugi sum is the connected sum operation. Hence for any knot K ,

$MSG(K, 2)$ has an obvious structure where the edges are arrows from each knot K' to the connected sum of K and K' . For $n = 4$, a 4-Murasugi sum is called a *plumbing*. It was shown in [11, 34, 42] that nice geometric properties of knots and surfaces (such as fiberedness and genus-minimality) were preserved under Murasugi sums and decompositions of minimal genus Seifert surfaces. On the other hand, Thompson [43] gave examples where the trefoil is obtained as a plumbing of two unknots, and the unknot is obtained as a plumbing of two figure-eight knots. Thus, expectations to generalize preservation results to Murasugi sums of non-minimal Seifert surfaces were negated.

In this chapter, we generalize Thompson's examples to show that given any three knots, we can produce one of them as a Murasugi sum of the other two.

Theorem A1. *For any three knots K_1, K_2, K_3 , there exist Seifert surfaces F_1, F_2, F_3 for them such that F_3 is a Murasugi sum of F_1 and F_2 .*

Therefore we have the following:

Corollary. *For a knot K , any set of knots $\{K_1, K_2, \dots, K_n\}$ in $MSG(K)$ composes a complete graph where all the edges are bi-directed.*

We refine the result of Theorem A1 by giving an algorithm to find a closed braid for K_3 which naturally splits into closed braids for K_1 and K_2 . See Figures 2.7, 2.8 following Example 2.2.6.

Theorem A2. *For any three knots K_1, K_2, K_3 , there are braids b_1, b_2, b_3 such that K_i is the closure of b_i ($i = 1, 2, 3$), satisfying the following:*

If the braid b_3 is expressed as a braid word W_3 with generators $\sigma_1, \sigma_2, \dots, \sigma_k, \dots, \sigma_n$, then W_1 (resp. W_2) is obtained from W_3 by deleting the generators $\sigma_1, \dots, \sigma_k$ (resp. $\sigma_{k+1}, \dots, \sigma_n$).

To further study the structure of $MSG(K, n)$ where the size of Murasugi sums is limited, we give, in Section 2.3, lower and upper bounds on the minimal m -gon required to form K_3 as a Murasugi sum of K_1 and K_2 . Our bounds are in terms of $d_{cb}(K, K')$ and $d_{bt}(K, K')$, which are, respectively, the minimal number of coherent band surgeries (resp. band-twists) required to transform K into K' . For the precise definition of these operations, see Definitions 2.2.2

and 2.3.2.

Theorem B. *If K_3 is an m -Murasugi sum of K_1 and K_2 such that m is minimal, then*

$$d_{cb}(K_1 \# K_2, K_3) + 2 \leq m \leq 2(d_{bt}(K_1, K_3) + d_{bt}(K_2, O) + 1),$$

where the roles of K_1 and K_2 can be switched to improve the upper bound.

2.2 Any knot is a Murasugi sum of any two knots

The original construction of the Murasugi sum was first introduced by Murasugi in [34] and was later coined as the Murasugi sum by Gabai in [11]. For simplicity, we define the Murasugi sum in terms of Murasugi decomposition. See Figure 2.1.

Definition 2.2.1. Let F be a Seifert surface in S^3 , let Σ be a 2-sphere such that $S^3 \setminus \Sigma$ is a union of two open 3-balls B_1, B_2 and such that $\Sigma \cap F$ is an m -gon Ω with m even. Denote $F \cap \overline{B_1} = F_1$ and $F \cap \overline{B_2} = F_2$. Then we say that F decomposes into F_1 and F_2 along Ω .

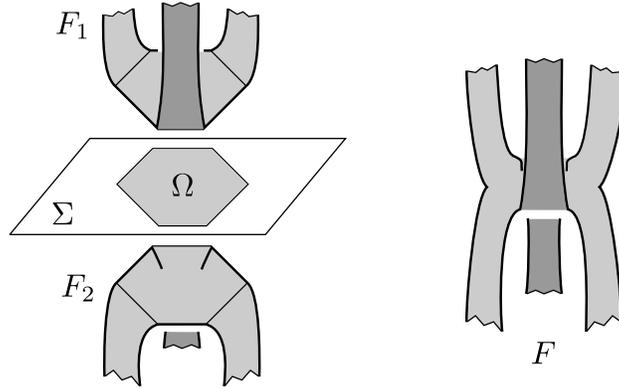


Figure 2.1. A local picture of a 6-Murasugi sum.

If F decomposes into F_1 and F_2 along an m -gon, then F is said to be an m -Murasugi sum of F_1 and F_2 , which we denote by $F_1 \star_m F_2$. Given two knots K_1, K_2 , we write $K_1 \star_m K_2$ to denote the boundary of $F_1 \star_m F_2$ for some Seifert surfaces F_1, F_2 for K_1, K_2 . Note that the summing disk Ω can initially appear stretched and twisted, but we can isotope the given surfaces to see Ω as flat. We can also isotope the surfaces so that Σ corresponds to the plane $z = 0$ in \mathbb{R}^3 , and in this situation, if F_1 lies above (resp. below) $z = 0$ and F_2 lies below (resp. above) $z = 0$, then we say that we Murasugi sum F_1 onto the positive (resp. negative) side of F_2 .

There are several operations one can perform on knot diagrams to obtain a new knot. One such operation is a crossing change, and more generally, an antiparallel full-twisting.

Definition 2.2.2. An *antiparallel full-twisting* on an oriented link is a local move where we select a pair of locally antiparallel strings $\begin{array}{c} \leftarrow \\ \rightarrow \end{array}$ and apply some number of full twists $\begin{array}{c} \curvearrowright \\ \curvearrowleft \end{array}$.

For convenience, we sometimes refer to antiparallel full-twisting as band twisting. We can realize this twisting operation along an arc α , where α is a short, unknotted arc connecting two antiparallel strings of a link L , which is contained within a small ball B such that $L \cap B$ is a trivial 2-string tangle. In this setting, we can span a Seifert surface F for L such that $F \cap B$ is a rectangle b containing the arc α . Then the twisting operation is realized by applying some full-twists to b .

Consider the two Seifert surfaces of the unknot in Figure 2.2. The following proposition states that a crossing change within a knot can be realized by either plumbing or deplumbing these surfaces. More generally, by increasing the number of full twists in R_+ or R_- , we can realize any antiparallel full-twisting by either plumbing or deplumbing Seifert surfaces for the unknot.

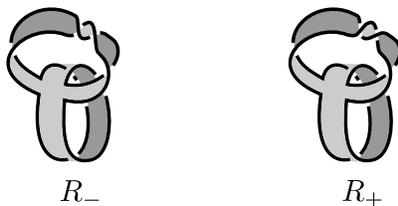


Figure 2.2. Two Seifert surfaces for the unknot.

Proposition 2.2.3. Let K_1, K_2 be knots such that K_2 is obtained by changing a positive crossing in K_1 . Then there exist Seifert surfaces F_1, F'_1 for K_1 , F_2, F'_2 for K_2 , and R_+, R_- for the unknot satisfying the following:

1. F'_1 is a plumbing of F_2 and R_+ .
2. F'_2 is a plumbing of F_1 and R_- .

Proof. We illustrate both statements in Figure 2.3. □

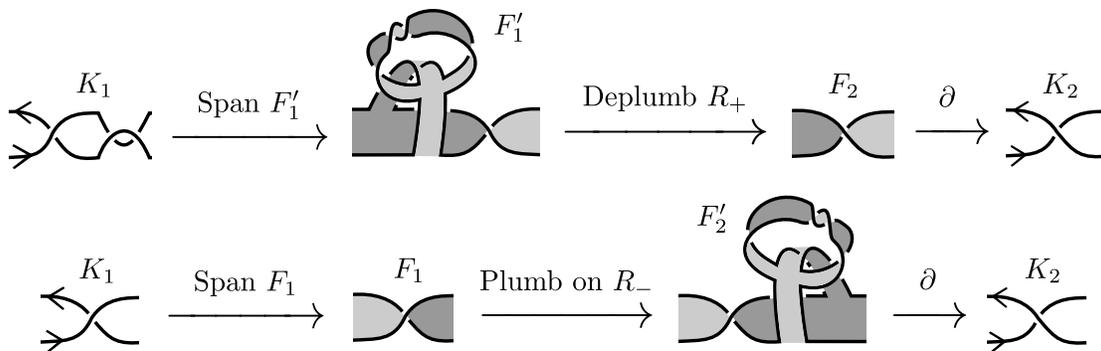


Figure 2.3. Changing a positive crossing.

Furthermore, we can perform any number of crossing changes simultaneously via a single Murasugi sum with a Seifert surface for the unknot, which allows us to prove the following lemma.

Lemma 2.2.4. *Any knot K has a Seifert surface F which is a Murasugi sum of two Seifert surfaces F_1, F_2 for the unknot.*

Proof. For any diagram D of K , we can choose a subset \mathcal{C} of crossings such that we obtain the unknot by simultaneously changing the crossings in \mathcal{C} . To see this, start at some point in D and walk along the knot. Then can specify \mathcal{C} to consist of those crossings which we enter first along an under-path and later along an over-path. Near each crossing \curvearrowright in \mathcal{C} , apply a Reidemeister II move to introduce an antiparallel clasp $\curvearrowright\curvearrowleft$, where undoing the clasp results in changing that crossing as in \curvearrowleft . Thus we obtain a new diagram D' for K , and we obtain the unknot by simultaneously undoing all the clasps. Put a 3-ball B in the complement of K and isotope K so that all the clasps are in B , and span a Seifert surface F for K as in the left part of Figure 2.4. Then F decomposes into two Seifert surfaces F_1 and F_2 , where ∂F_2 is the unknot and ∂F_1 is the result of undoing all the clasps in K , and hence is the unknot. \square

Conversely, any knot is obtained from an unknot by simultaneously removing antiparallel clasps. Therefore, in the proof above, we may regard F and F_2 as surfaces for the unknot and F_1 as a surface for K . This gives the following.

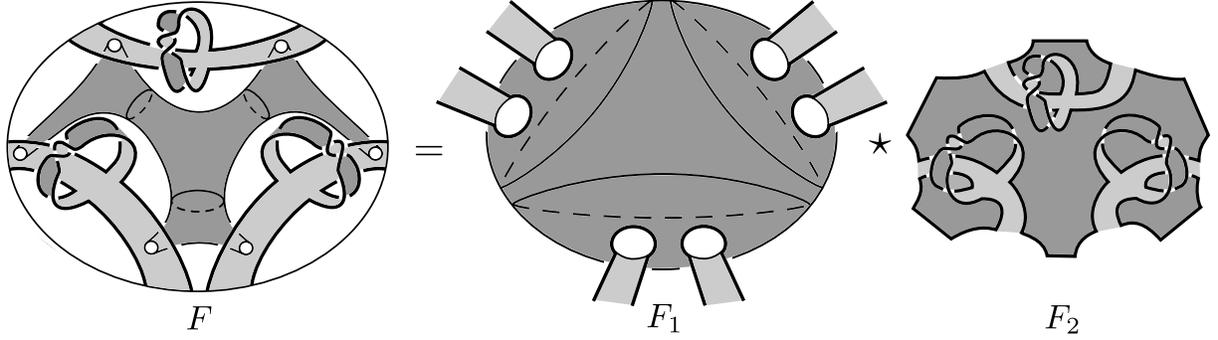


Figure 2.4. Producing $\partial F = K$ as a Murasugi sum of unknots.

Corollary 2.2.5. *Any knot K has a Seifert surface F_1 which becomes a Seifert surface for the unknot by Murasugi summing F_1 with some Seifert surface for the unknot.*

Proof of Theorem A1. By Corollary 2.2.5, there is a Seifert surface F_1 (resp. F_2) for K_1 (resp. K_2) that Murasugi sums with a Seifert surface F'_1 (resp. F'_2) of the unknot O to yield a Seifert surface for O . By Lemma 2.2.4, there exist two Seifert surfaces F_3, F'_3 for O that Murasugi sum to a Seifert surface for K_3 . The boundary connected sum of F_1, F'_2, F_3 (resp. F'_1, F_2, F'_3) is a Seifert surface for K_1 (resp. K_2), and we Murasugi sum these surfaces as in Figure 2.5 along the shaded n -gon. \square

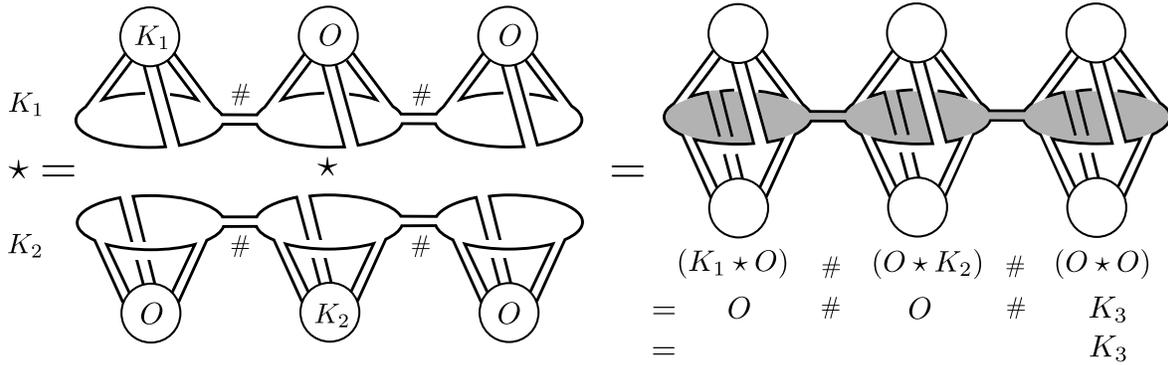


Figure 2.5. K_3 as a Murasugi sum of K_1 and K_2 .

Proof of Theorem A2. By Theorem A1, there are Seifert surfaces F_1, F_2, F_3 for K_1, K_2, K_3 such that F_3 is a Murasugi sum of F_1 and F_2 and the summing disk Ω is flat. Apply a trivial twist at each band attached to Ω as depicted in Figure 2.6. Then we have a diagram D' such that the summing disk is spanned by a Seifert circle C . Note that the canonical Seifert surface F'_3 is a Murasugi sum of F'_1 and F'_2 along the summing disk Ω ,

where $\partial F'_1 = \partial F_1, \partial F'_2 = \partial F_2, \partial F'_3 = \partial F_3$. Apply Yamada's braiding algorithm [46] to D' , independently inside and outside C . Then we have the desired braids. \square

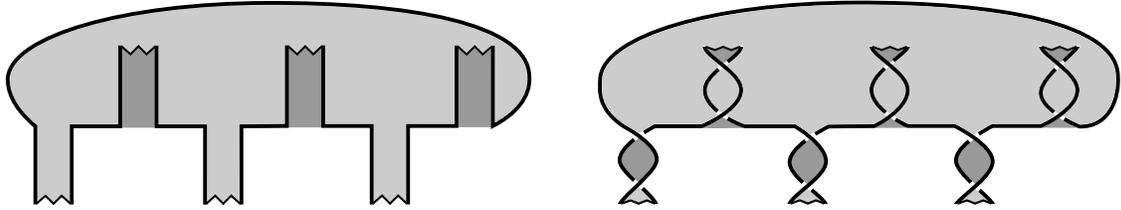


Figure 2.6. Making a canonical surface for the braid decomposition.

We use the notation $F \xrightarrow{\partial} K$ to mean that $\partial F = K$, and $F_1 \stackrel{\partial}{=} F_2$ to mean that $\partial F_1 = \partial F_2$.

Example 2.2.6. In Figure 2.7, we illustrate 7_5 as a Murasugi sum of two unknots, and in Figure 2.8, we illustrate the unknot as a Murasugi sum of 5_2 and 7_5 . In these figures, we also express the Murasugi sums in terms of braid decompositions, as described in Theorem A2. Note that we simplified some procedures in the proofs of Theorems A1 and A2. In particular, we already have a Seifert circle corresponding to the summing disk without twisting as in Figure 6, we eliminated some Seifert circles with two bands, and we used “long” bands to save us from depicting many generators.

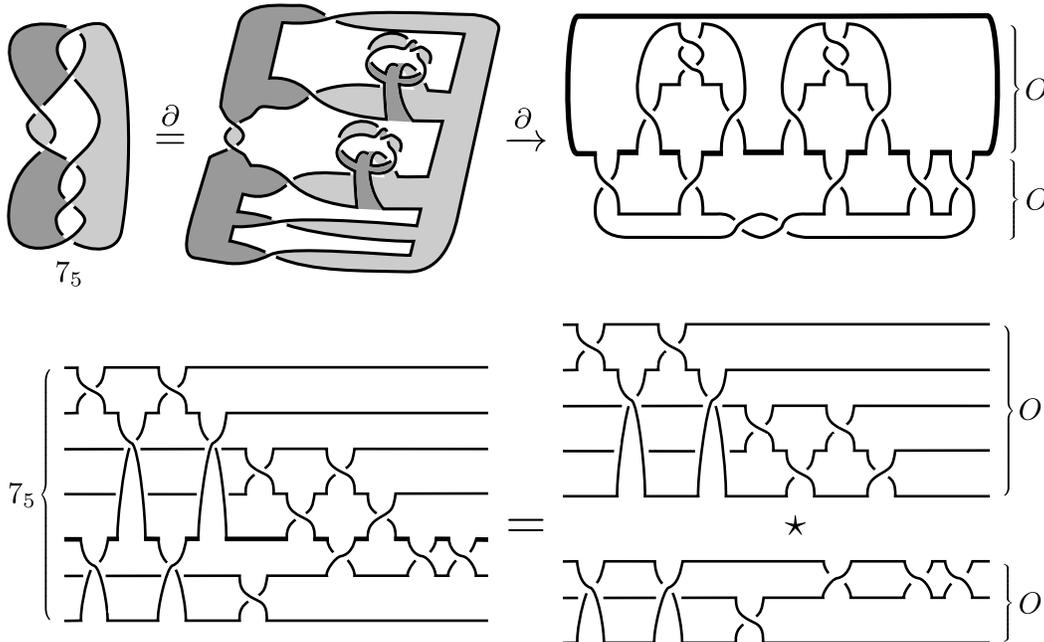


Figure 2.7. 7_5 as a Murasugi sum of two unknots, and its braid decomposition.

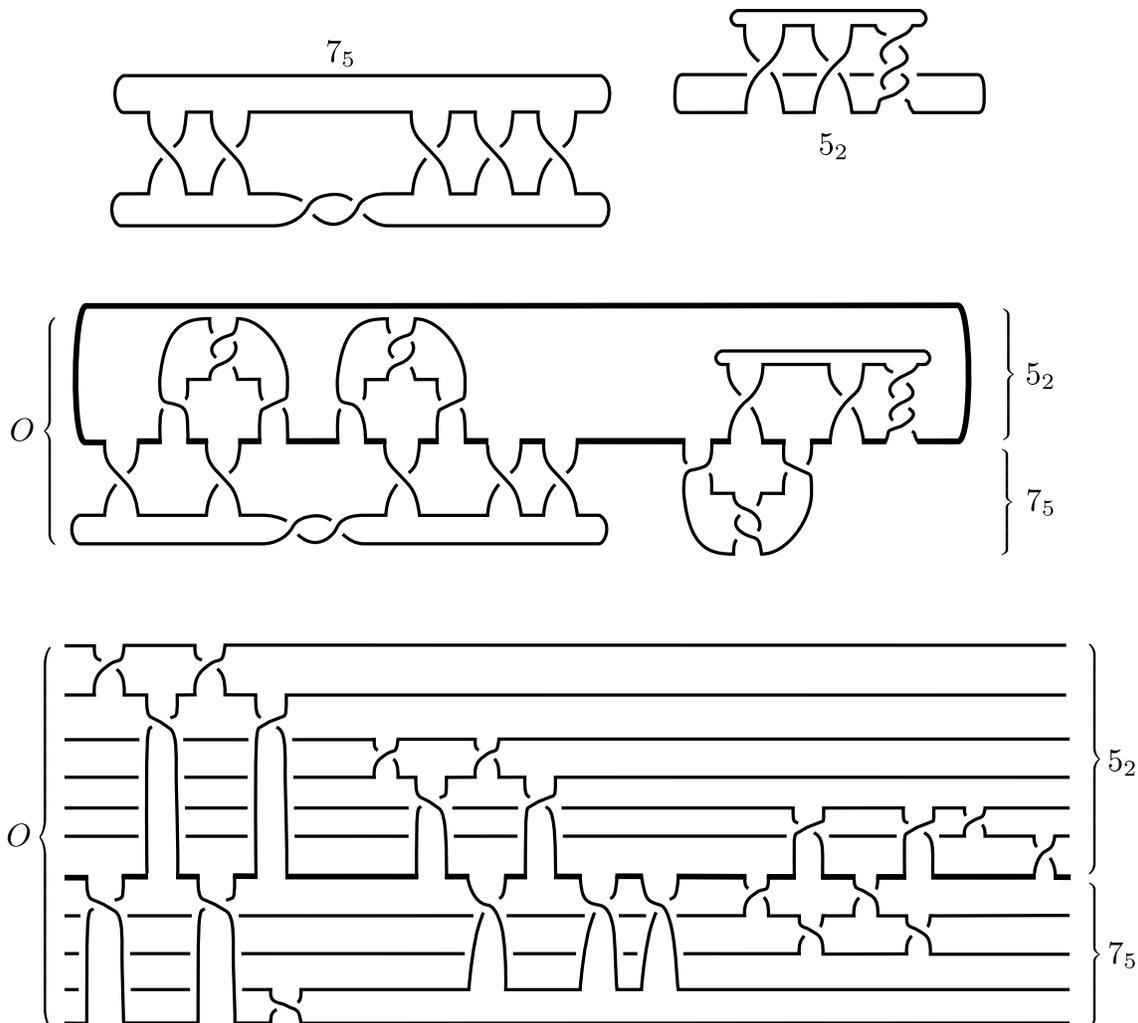


Figure 2.8. The unknot as a Murasugi sum of 5_2 and 7_5 , and its braid decomposition.

2.3 Constraining the minimal m -gon

In Theorem A1, it was shown that given three knots K_1, K_2, K_3 , we can form K_3 as a Murasugi sum of K_1 and K_2 by using a sufficiently large m -gon. We begin this section with Definition 2.3.1 in order to give lower and upper bounds for the minimal such m -gon, culminating in Theorem B from Section 2.1. In particular, if we restrict the size of our m -gons, Theorem B obstructs forming a knot as a Murasugi sum of two given knots. Previously, the only restriction on Murasugi sums was for plumbings of fiber surfaces, such as in [33], where Melvin and Morton showed that the Conway polynomials of fibered knots of genus 2 take a restricted form when the fiber surface is a plumbing of Hopf bands.

Definition 2.3.1. Given three oriented knots K_1, K_2, K_3 , define the *minimal size of Mura-*

sugi summation for $K_3 = K_1 \star K_2$ to be

$$d_M(K_1, K_2; K_3) = \min\{m \mid F_3 = F_1 \star_m F_2, \partial F_i = K_i \text{ for } i = 1, 2, 3\},$$

where the minimum is taken over F_1, F_2 .

Recall that m is even. From the definition, $d_M(K_1, K_2; K_3) = 2$ if and only if $K_3 = K_1 \# K_2$. Also, if $K_3 \neq K_1 \# K_2$ can be realized as a plumbing of K_1 and K_2 , then $d_M(K_1, K_2; K_3) = 4$. Now we recall the notion of a coherent band surgery. See Figure 2.9.



Figure 2.9. A band surgery along an unlink yielding $3_1 \# \overline{3_1}$.

Definition 2.3.2. Given a knot or link L and an embedded band $b : I \times I \rightarrow S^3$ with $L \cap b(I \times I) = b(I \times \partial I)$, we obtain a new link $L' = (L \setminus b(I \times \partial I)) \cup b(\partial I \times I)$, and we say that L' is obtained from L by a *band surgery*. For oriented L, L' , a *coherent* band surgery is a band surgery that respects the orientations of both links, that is, $L \setminus b(I \times \partial I) = L' \setminus b(\partial I \times I)$ as oriented spaces.

Given two oriented links L, L' , we denote by $d_{cb}(L, L')$ the minimal number of coherent band surgeries required to produce L' from L . This number is known as the coherent band-Gordian distance, and in the case of knots it is equal to twice the $SH(3)$ -Gordian distance [26]. In [27], d_{cb} is calculated for most pairs of knots up to seven crossings. For knots K, K' , note that $d_{cb}(K, K')$ is necessarily even because a coherent band surgery changes the number of components by one.

We have the following lower bounds for d_M in terms of d_{cb} and the signature σ .

Theorem 2.3.3. *For knots K_1, K_2, K_3 , we have*

$$d_{cb}(K_1 \# K_2, K_3) + 2 \leq d_M(K_1, K_2; K_3).$$

Consequently,

$$d_{cb}(K_1 \sqcup K_2, K_3) + 1 \leq d_M(K_1, K_2; K_3),$$

where $K_1 \sqcup K_2$ is a split link, and

$$|\sigma(K_1) + \sigma(K_2) - \sigma(K_3)| + 2 \leq d_M(K_1, K_2; K_3).$$

Proof. Suppose K_3 is an m -Murasugi sum of K_1 and K_2 , where $m = 2n$ is minimal. Then there is a sequence of $2(n - 1) = m - 2$ coherent band surgeries between K_3 and $K_1 \# K_2$, where each band lies within a Seifert surface for K_3 , so that

$$d_{cb}(K_1 \# K_2, K_3) + 2 \leq m.$$

See Figure 2.10. With one more band surgery, we have $m - 2 + 1$ coherent band surgeries between K_3 and the split link $K_1 \sqcup K_2$, so that

$$d_{cb}(K_1 \sqcup K_2, K_3) + 1 \leq m.$$

If an oriented link L' is obtained from L by a coherent band surgery, an estimate was given by Murasugi [35] on the difference of the signatures as $|\sigma(L) - \sigma(L')| \leq 1$. Since $\sigma(K_1 \# K_2) = \sigma(K_1) + \sigma(K_2)$, we obtain the third inequality. \square

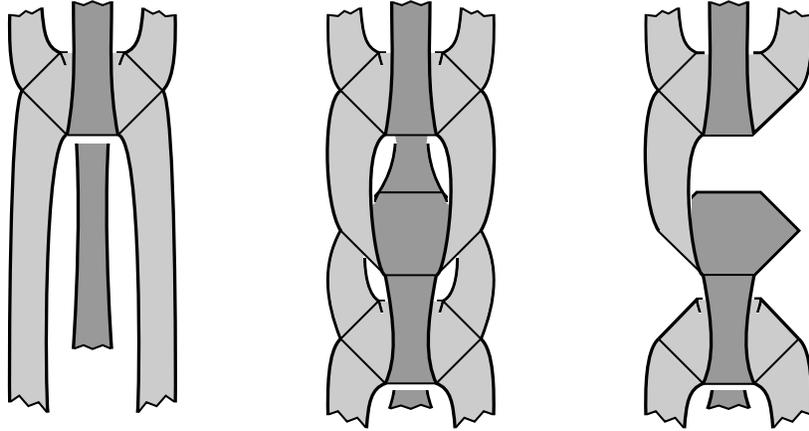


Figure 2.10. Performing $(6 - 2)$ band surgeries to recover $K_1 \# K_2$ from K_3 .

Example 2.3.4. By Theorem 2.3.3, the signature obstructs forming 9_1 as a plumbing of two copies of 3_1 . However, Figure 2.11 depicts how 9_1 desums into two copies of 3_1 along the shaded 6-gon, which is obtained by first merging two 4-gons into an 8-gon which is then reduced to the 6-gon. This process is explained in the proof of Lemma 2.3.8.

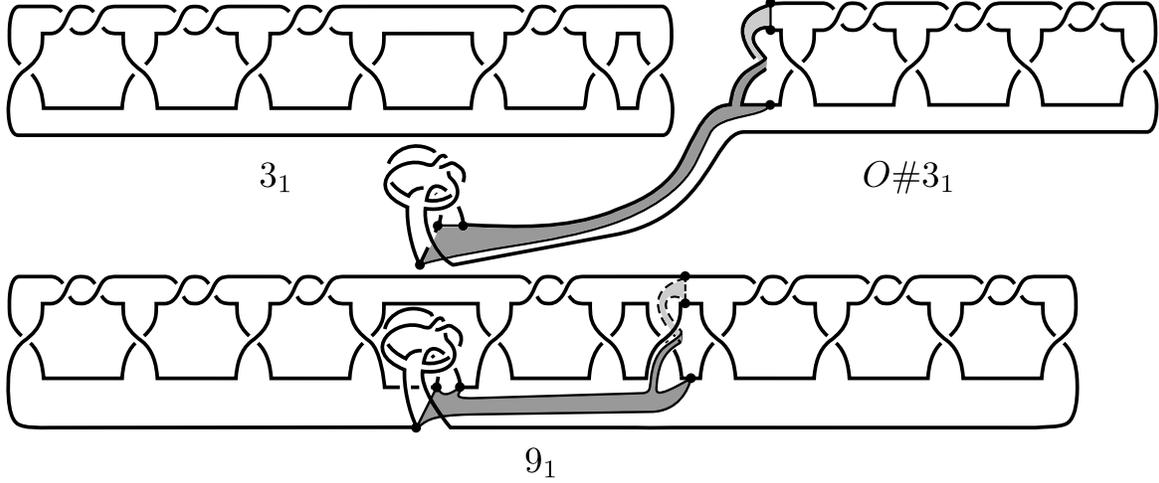


Figure 2.11. 9_1 as a 6-Murasugi sum of 3_1 and 3_1 .

Remark 2.3.5. For pairs of knots where d_{cb} has not been determined, we may apply lower bounds for d_{cb} from [26] in terms of the smooth four-ball genus g_4 , hence in terms of σ by [35], and in terms of the Nakanishi index e . Indeed, for knots K_1, K_2 we have

$$|\sigma(K_1) - \sigma(K_2)| \leq 2g_4(K_1\# - K_2) \leq d_{cb}(K_1, K_2),$$

where $-K_2$ is the reverse mirror of K_2 , and

$$|e(K_1) - e(K_2)| \leq d_{cb}(K_1, K_2).$$

We now move on to some upper bounds of d_M . Using Lemma 2.3.7 below, we can easily modify the proof of Theorem A1 to form K_3 as a $4(u(K_1) + u(K_2) + u(K_3))$ -Murasugi sum of K_1 and K_2 , where $u(K)$ is the unknotting number of K . This gives the upper bound

$$d_M(K_1, K_2; K) \leq 4(u(K_1) + u(K_2) + u(K_3)).$$

In what follows, we improve this upper bound into Theorem 2.3.9. Recall that the Gordian distance d_G between two knots K, K' is the minimal number of crossing change operations required to transform a diagram for K into a diagram for K' , where the minimum is taken over all diagrams for K . Since K can be transformed into K' by a sequence of crossing changes that passes through the unknot, we have $d_G(K, K') \leq u(K) + u(K')$. More generally, we consider the band-twist distance d_{bt} .

Definition 2.3.6. Define the *band-twist distance* between two oriented links L, L' , denoted by $d_{bt}(L, L')$, as the minimal n such that there exists a sequence of links $L = L_0, L_1, L_2, \dots, L_n = L'$, where L_{i+1} is obtained from L_i by an antiparallel full-twisting (a band-twist) as in Definition 2.2.2.

Since any crossing change may be realized as a band-twist, we have $d_{bt}(K, K') \leq d_G(K, K')$. Just as with crossing changes, we may perform any number of band-twist operations simultaneously.

Lemma 2.3.7. *For two links L and L' with $d_{bt}(L, L') = n$, there exists a Seifert surface F for L' with a set A of n mutually disjoint properly embedded arcs such that a Seifert surface F for L is obtained by applying a band-twist operation along each arc in A .*

Proof. Let $L = L_0, L_1, L_2, \dots, L_n = L'$ be a sequence of links related by the band-twisting operation. After obtaining L_1 from L_0 by twisting along an arc α_0 , instead of erasing α_0 , isotope it so that it is disjoint from the arc α_1 used to obtain L_2 . Repeating this, we obtain a set of arcs $\alpha_0, \alpha_1, \dots, \alpha_{n-1}$ attached to $L_n = L'$. By an isotopy, we can arrange the arcs to be short and contained in a ball B with $L' \cap B$ being a trivial $2n$ -string tangle. Splice L' along the arcs and push the resulting link L'' slightly off B . Span a Seifert surface for L'' disjoint from B . Then we obtain the desired Seifert surface for L' by attaching bands to L'' that pass through B . \square

If we wish to perform several Murasugi sums of several surfaces with a single surface, we can often combine these sums into a single sum as indicated by the following lemma. One implication of the below is that we may perform any number of band-twist operations simultaneously via a single Murasugi sum.

Lemma 2.3.8. *Suppose that a Seifert surface F' is obtained by Murasugi summing S_1, S_2, \dots, S_n on the same side of a connected Seifert surface F along mutually disjoint summing disks $\Omega_1, \Omega_2, \dots, \Omega_n$. Let Ω_i be an e_i -gon ($i = 1, 2, \dots, n$). Then F' is a Murasugi sum of F with a boundary connected sum of S_1, S_2, \dots, S_n along an m -gon, where $m = \sum_{i=1}^n e_i$. Moreover,*

if F is a Seifert surface for a knot, then we can merge the sums into an m' -gonal sum with $m' = m - 2(n - 1)$.

Proof. We may assume that each summand S_i is contained in a thin blister neighborhood of the summing disk Ω_i . Denote the edges in $\partial\Omega_i$ as $a_{i,1}, b_{i,1}, a_{i,2}, b_{i,2}, \dots$ for $i = 1, 2, \dots, n$, where the $a_{i,\cdot}$'s are sub-arcs of ∂F and the $b_{i,\cdot}$'s are properly embedded arcs in F . Since F is connected, we may find an embedded arc γ in F whose endpoints are the midpoints of $b_{1,j}$ and $b_{2,k}$ for some j and k as in Figure 2.12. We merge the two (dark shaded) summing disks Ω_1 and Ω_2 into an $(e_1 + e_2)$ -gon Ω' whose boundary consists of $((\partial\Omega_1 \cup \partial\Omega_2) \setminus (b_{1,j} \cup b_{2,k})) \cup (\gamma_1 \cup \gamma_2)$, where γ_1 and γ_2 are properly embedded in F and $b_{1,j} \cup \gamma_1 \cup b_{2,k} \cup \gamma_2$ is a rectangle $R = \gamma \times I$ such that $R \cap \text{int} \bigcup_{i=1}^n \Omega_i = \emptyset$. We see that the boundary connected sum of S_1 and S_2 is contained in a thin blister neighborhood of the new summing disk Ω' . By repeating this merging operation, we eventually combine all the summing disks into one and obtain the desired Murasugi sum.

For the last part of the assertion, the assumption that ∂F being connected ensures the existence of two arcs $a_{1,p}, a_{2,q}$ for some p, q such that one segment of ∂F between $a_{1,p}$ and $a_{2,q}$ does not pass through other summing disks. Then we can apply the previously mentioned merging of Ω_1 and Ω_2 so that γ_1 or γ_2 , say γ_1 , can be isotoped to a sub-arc of ∂F in $F \setminus \bigcup_{i=1}^n \Omega_i$. Then the three consecutive edges in Ω' , say, $a_{1,k}, \gamma_1, a_{2,j}$ for some k, j , can be regarded as one edge by merging the bands of S_1 and S_2 attached to $a_{1,k}$ and $a_{2,j}$ as in Figure 2.13. Then the new summing disk Ω'' is a p -gon, where $p = e_1 + e_2 - 2$. \square

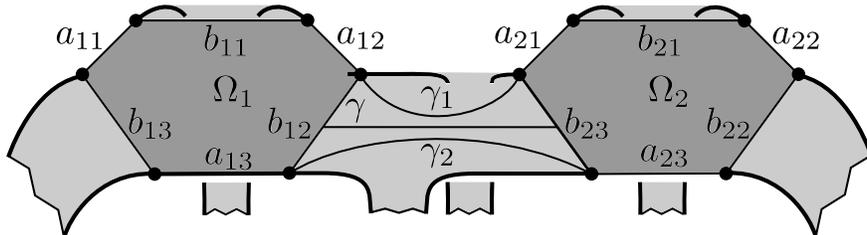


Figure 2.12. Merging disks Ω_1, Ω_2 in F along γ .

Combining Lemma 2.3.7 and Lemma 2.3.8, we arrive at the following improvement on our upper bound.

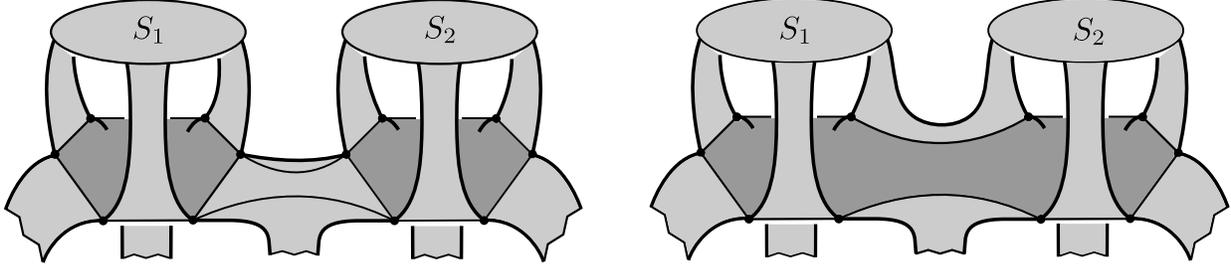


Figure 2.13. Merging two 6-gons into a $(6 + 6 - 2)$ -gon.

Theorem 2.3.9. *Let K_1, K_2, K_3 be knots. Then*

$$d_M(K_1, K_2; K_3) \leq 2(d_{bt}(K_1, K_3) + d_{bt}(K_2, O) + 1),$$

where the roles of K_1 and K_2 may be switched to improve the upper bound.

Proof. Suppose that $d_{bt}(K_1, K_3) = p$ and $d_{bt}(K_2, O) = q$. As guaranteed by Lemma 2.3.7, there is a Seifert surface F_1 for K_1 (resp. F_2 for K_2) with a collection \mathcal{A} of arcs $\alpha_1, \dots, \alpha_p$ (resp. \mathcal{B} of arcs β_1, \dots, β_q) such that performing band-twist operations along the arcs yields a Seifert surface F_3 for K_3 (resp. F_0 for O). Prepare Seifert surfaces A_1, \dots, A_p and B_1, \dots, B_q such that plumbing A_j along α_j (resp. B_k along β_k) results in applying the band-twist operation along α_j (resp. β_k). More precisely, each of the surfaces is a plumbing of a trivial annulus and an unknotted annulus with various numbers of full-twists (recall Figure 2).

Construct F'_1 from F_1 such that $\partial F'_1 = K_3$ by plumbing A_1, \dots, A_p along \mathcal{A} on the positive side of F_1 . Also, construct F'_0 from F_0 such that $\partial F'_0 = O$ by plumbing B_1, \dots, B_q along \mathcal{B} on the negative side of F_0 .

We merge the plumbed surfaces A_1, \dots, A_p so that F'_1 is a Murasugi sum of F_1 and A , where A is a boundary connected sum of A_1, \dots, A_p . Note that ∂A is an unknot O_1 . By Lemma 2.3.8, we may regard the Murasugi sum as along a $(4p - 2(p - 1))$ -gon and hence a $(2p + 2)$ -gon. Similarly, F'_0 is regarded as a $(2q + 2)$ -Murasugi sum of F_0 and a Seifert surface B for an unknot O_2 . Denote by F the 2-Murasugi sum (i.e., a boundary connected sum) of F'_1 and F'_0 , where F'_0 is summed on the positive side of F'_1 .

Then F is a $(2p + 2 + 2q + 2)$ -Murasugi sum of two summands, where one summand is the boundary connected sum of F_1 and B , and the other summand is the boundary

connected sum of F_0 and A . The boundary of the first (resp. second) summand is $K_1\#O_2$ (resp. $K_2\#O_1$). By Lemma 2.3.8, we can reduce the summing $(2p + 2q + 4)$ -gon into a $(2p + 2q + 2)$ -gon. Therefore, we have expressed K_3 as a $2(p + q + 1)$ -Murasugi sum of K_1 and K_2 . \square

By combining Theorems 2.3.3 and 2.3.9, we arrive at Theorem B. As an application of Theorem B, we determine $d_M(3_1, 3_1; K)$ for knots up to five crossings. Theorem B shows that

$$d_M(3_1, 3_1; 3_1) = 4, \quad d_M(3_1, 3_1; O) = d_M(3_1, 3_1; 4_1) = 6,$$

while it only gives the bounds

$$4 \leq d_M(3_1, 3_1; 5_1), \quad d_M(3_1, 3_1; 5_2) \leq 6.$$

We can show that $d_M(3_1, 3_1; 5_1) = d_M(3_1, 3_1; 5_2) = 4$ in the following way.

For $a_1, a_2, \dots, a_n \in 2\mathbb{Z}$, denote by $S[a_1, a_2, \dots, a_n]$ a linear plumbing of n unknotted annuli, where the i^{th} annulus has a_i half-twists. Note that all 2-bridge links have such a linear plumbing as a Seifert surface [17], which is of minimal genus if and only if $a_1 a_2 \cdots a_n \neq 0$.

Using the notation of Example 2.2.6, we have the following:

1. $S[a_1, a_2, \dots, a_n] \star_4 S[b_1, b_2, \dots, b_m] = S[a_1, \dots, a_n, b_1, \dots, b_m]$,
2. $S[a_1, a_2, \dots, a_n] \stackrel{\partial}{=} S[a_1, a_2, \dots, a_n, a_{n+1}, 0]$ for any a_{n+1} ,
3. $S[a_1, a_2, \dots, a_i, 0, a_{i+2}, \dots, a_n] \stackrel{\partial}{=} S[a_1, a_2, \dots, a_i + a_{i+2}, \dots, a_n]$.

Using the above, we have:

Example 2.3.10. 1. $3_1 \star_4 3_1 \stackrel{\partial}{\leftarrow} S[2, 2] \star_4 S[2, 2] = S[2, 2, 2, 2] \stackrel{\partial}{\rightarrow} 5_1$

2. $3_1 \star_4 3_1 \stackrel{\partial}{\leftarrow} S[2, 2, -2, 0] \star_4 S[2, 2] \stackrel{\partial}{=} S[2, 2, 0, 2] \stackrel{\partial}{=} S[2, 4] \stackrel{\partial}{\rightarrow} 5_2$.

Using this notation, we summarize what Thompson showed in [43] as follows:

1. $O \star_4 O \stackrel{\partial}{\leftarrow} S[2, 0] \star_4 S[0, 2] = S[2, 0, 0, 2] \stackrel{\partial}{=} S[2 + 0, 2] = S[2, 2] \stackrel{\partial}{\rightarrow} 3_1$

2. $4_1 \star_4 4_1 \stackrel{\partial}{\leftarrow} S[2, -2, 2, 0] \star_4 S[-2, 2] \stackrel{\partial}{=} S[2, -2, 0, 2] \stackrel{\partial}{=} S[2, 0] \stackrel{\partial}{\rightarrow} O$.

CHAPTER 3

PSEUDO-ANOSOV MAPS AND THE BESTVINA-HANDEL ALGORITHM

3.1 The classification of surface homeomorphisms

A celebrated result in low-dimensional topology is the classification of surfaces. This result states that two orientable surfaces are diffeomorphic if and only if they have the same number of boundary components b , the same genus g , and the same number of punctures n . Denote the diffeomorphism type of such a surface by $F = F_{b,g,n}$. Then the Euler characteristic of F is $\chi(F_{b,g,n}) = 2 - 2g - b - n$.

A natural object of study is the set of all self-diffeomorphisms of F . Let $\text{Diffeo}^+(F, \partial F)$ denote the group of orientation-preserving diffeomorphisms of F which restrict to the identity on ∂F (if $\partial F \neq \emptyset$). After endowing $\text{Diffeo}^+(F, \partial F)$ with the compact-open topology, we define the *mapping class group* of F as

$$\mathcal{MCG}(F) := \pi_0(\text{Diffeo}^+(F, \partial F)).$$

In other words, $\mathcal{MCG}(F)$ is the group of smooth isotopy class of elements of $\text{Diffeo}^+(F, \partial F)$, where isotopies are required to fix the boundary pointwise. Such a diffeomorphism can be represented as a composition of Dehn twists about some collection of essential simple closed curves in F [29].

Definition 3.1.1. A *right Dehn twist* about a simple closed curve $A \subset F$, denoted D_A^R , is defined as follows. Let $X = S^1 \times [0, 1]$ be the annulus with orientation induced by \mathbb{R}^2 , and identify a regular neighborhood $\nu(A)$ of A with X via an orientation-preserving diffeomorphism $\phi : X \rightarrow \nu(A)$. After defining the twist map $T : X \rightarrow X$ as

$$T(\theta, t) = (\theta - 2\pi t, t),$$

we define the homeomorphism $D_A^R : F \rightarrow F$ to be the identity outside of $\nu(A)$, and to be the composition $\phi \circ T \circ \phi^{-1}$ on $\nu(A)$. A *left Dehn twist* about A , denoted D_A^L , is defined similarly, where we instead consider the twist map $T(\theta, t) = (\theta + 2\pi t, t)$. See Figure 3.1.

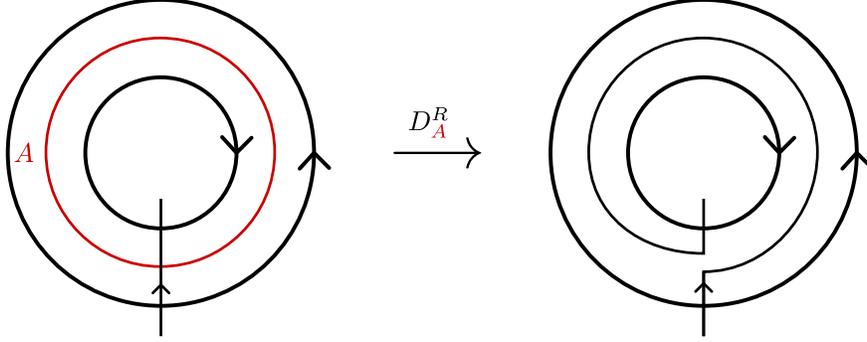


Figure 3.1. The action of a right Dehn twist about A on a line segment.

Elements of the mapping class group are classified according to the following theorem of Thurston [45]. Though we state this theorem for compact surfaces, a similar classification holds for surfaces with punctures [2].

Theorem 3.1.2. *A diffeomorphism ϕ of a compact oriented surface F is isotopic to a diffeomorphism ϕ' such that either*

1. ϕ' is periodic, i.e. there exists some positive integer m such that $(\phi')^m = \text{id}$; or
2. ϕ' is pseudo-Anosov, i.e. there is a number $\lambda > 1$ and a pair of transverse measured foliations $(\mathcal{F}_s, \mu_s), (\mathcal{F}_u, \mu_u)$ such that $\phi'(\mathcal{F}_s, \mu_s) = (\mathcal{F}_s, \frac{1}{\lambda}\mu_s)$ and $\phi'(\mathcal{F}_u, \mu_u) = (\mathcal{F}_u, \lambda\mu_u)$; or
3. ϕ' is reducible, i.e. there is a collection of disjoint simple closed curves C in the interior of F that is fixed by ϕ .

Thus, we can isotope the monodromy ϕ of the fiber surface F of a fibered knot K until it is either periodic, pseudo-Anosov, or reducible. We are particularly interested in fibered knots which have pseudo-Anosov monodromy, so-called *hyperbolic knots*. This class of knots consists of all knots which are not torus knots or satellite knots [44]. Conjecturally, most prime knots are hyperbolic knots, in the sense that the ratio of hyperbolic knots to prime knots approaches nearly 1 as the number of crossings goes to infinity [31].

While we called pseudo-Anosov maps diffeomorphisms in Theorem 3.1.2, it is more accurate to say that they are diffeomorphisms away from the finite set of singularities of the associated pair of measured foliations $\mathcal{F}_s, \mathcal{F}_u$. Such singularities in the interior of F are locally modelled on p -prongs, where $p \geq 3$. At a singularity which is a puncture of F , the

same p -pronged model applies, but we now allow $p \geq 1$. From a puncture, we can recover a boundary component by blowing up the puncture, e.g. as in §2.1 of [24]. See Figure 3.2. On a component of ∂F , singularities of \mathcal{F}_s and \mathcal{F}_u alternate, such that the boundary component is a cycle of leaves, and ϕ is the identity on the boundary.

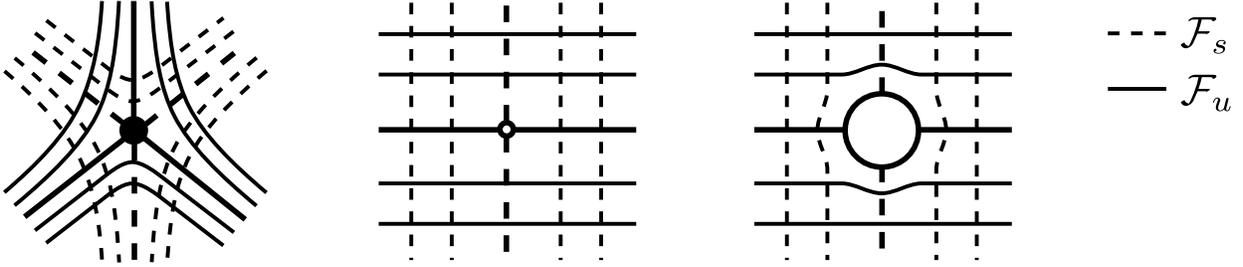


Figure 3.2. A 3-pronged interior singularity (left), a 2-pronged puncture (middle), and a blow-up of a 2-pronged puncture into a boundary component (right).

3.2 The Bestvina-Handel algorithm

While it has been established that a surface diffeomorphism *can* be isotoped into one of the classes of Thurston, it can be difficult to practically identify the particular class to which a given map belongs. To rectify this, Bestvina-Handel provide an algorithm to isotope a surface diffeomorphism into its Thurston class [2]. In what follows, we record only the parts of the algorithm which pertain to representatives of pseudo-Anosov maps on surfaces with one puncture. The algorithm as applied to such maps produces a so-called “invariant train-track” which combinatorially encodes the dynamics of the canonical pseudo-Anosov representative.

We begin with a spine of a punctured fiber surface F having monodromy ϕ , where the spine is viewed as a graph g with vertices and arbitrarily oriented edges. For an oriented edge e of this graph, we denote by \bar{e} the same edge with opposite orientation.

Definition 3.2.1. The collection of oriented edges emanating from a vertex v is called the *link* of v .

Continuing, we represent ϕ as some composition of Dehn twists within a regular neighborhood of the spine. After applying this composition to the edges of g , we realize the image

of each edge as a connected *edge-path* which travels through the edges and vertices of g . By abuse of notation, we refer to the induced graph map as ϕ .

Definition 3.2.2. Viewing ϕ as a graph map, we define the *derivative* D of ϕ at an edge e to be the first edge in the edge-path $\phi(e)$. The derivative D partitions the link of a vertex v into equivalence classes called *gates*, where two edges e_1, e_2 emanating from v are in the same gate if $D^m(e_1) = D^m(e_2)$ for some positive integer m .

Our initial edge-path calculation usually produces some redundancy which prevents the initial representative of ϕ from being pseudo-Anosov. This redundancy results from iterates of edge-paths which produce sub-paths of the form $e\bar{e}$, which we call *back-tracking*. However, the Bestvina-Handel algorithm says that we can perform a sequence of isotopies to remove such back-tracking. More precisely, the algorithm produces an efficient graph map.

Definition 3.2.3. The graph map $\phi : g \rightarrow g$ is called *efficient* if g has no vertices of valence 1 or 2. Additionally, every iterate of ϕ sends each edge to an edge-path which does not back-track. The second condition is equivalent to saying that for every edge e , the edge-path $\phi(e)$ contains no back-tracking, and $\phi(e)$ enters and exits any given vertex at different gates.

From an efficient graph map ϕ , we can construct an associated pair of transverse measured foliations, thus realizing ϕ as pseudo-Anosov. The six isotopies required to make a graph map efficient, called *moves*, are described combinatorially as follows. See Figure 3.3 for an example of how certain moves affect a graph. By abuse of notation, we denote the result of isotopies of ϕ as ϕ .

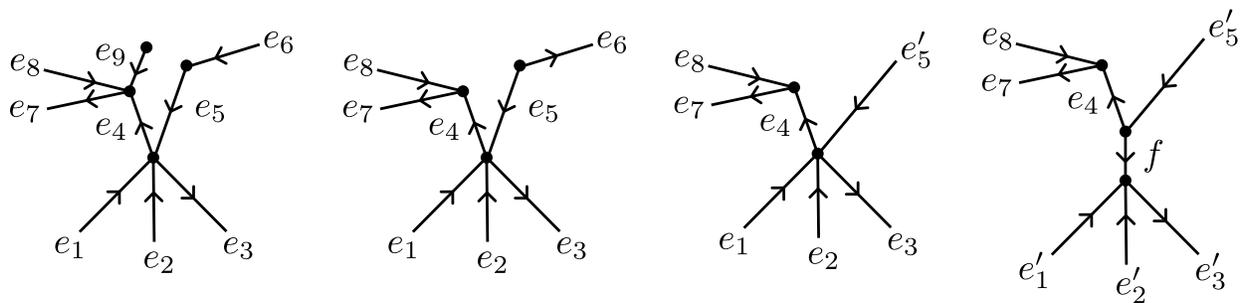


Figure 3.3. From left to right: starting with a graph, we perform a valence 1 isotopy, then a valence 2 isotopy, and then a fold.

Move 1 (Pulling tight)

If we encounter an edge-path containing a back-tracking subpath $e\bar{e}$, we remove this subpath. This has no effect on the graph.

Move 2 (Collapsing an invariant forest)

Suppose we encounter an invariant forest of g , i.e. a subgraph h of g such that $\phi(h) \subseteq h$ and each component of h is contractible. Then we collapse each component of h to a vertex, and we remove all occurrences of edges in the invariant forest from all edge-paths.

Move 3 (Valence 1 isotopy)

Suppose we encounter a valence 1 vertex. Then we remove this vertex and the emanating edge, as well as removing all occurrences of the emanating edge from all edge-paths.

Move 4 (Valence 2 isotopy)

Suppose we encounter a valence 2 vertex with emanating edges e_1, e_2 . Then we remove this vertex and amalgamate e_1, e_2 into an edge labeled, say, e'_1 . The edge-path $\phi(e'_1)$ is the concatenation $\phi(\bar{e}_2) \cdot \phi(e_1)$. Furthermore, we delete e_2 , and we replace e_1 with e'_1 in all edge-paths.

Move 5 (Folding)

Suppose we encounter a pairwise adjacent collection of emanating edges e_1, \dots, e_m in the link of v having common starting image the edge-path P . Then we identify (fold) the starting portion of this collection of edges to a single edge f with $\phi(f) = P$, and we relabel $e_1 = e'_1, \dots, e_m = e'_m$. Furthermore, we replace all occurrences of e'_1, \dots, e'_m in all edge-paths with $f e'_1, \dots, f e'_m$. We call the operation of a folding which involves the start of \bar{e} : folding the end of e .

Move 6 (Subdivision)

Suppose we encounter an edge e with interior point v mapping to a vertex of g . Then we can subdivide $e = e_1 e_2$, where \bar{e}_1, e_2 emanate from the valence 2 vertex v . As a result, we replace all occurrences of e with $e_1 e_2$, and we have $\phi(e) = \phi(e_1) \cdot \phi(e_2)$.

The order and way in which we perform these six moves is sometimes governed by the following matrix.

Definition 3.2.4. For a graph g with edges e_1, \dots, e_m , the *transition matrix* of a graph map $\phi : g \rightarrow g$ is the matrix M whose entry m_{ij} is the number of times that e_j and \bar{e}_j appear in the edge-path $\phi(e_i)$. In other words, M records how many times e_i runs over e_j under ϕ , regardless of orientation.

For ϕ to be in a pseudo-Anosov class, it is necessary that M is Perron-Frobenius and hence has a unique positive eigenvector. The i^{th} entry of this eigenvector corresponds to a width w assigned to the edge e_i . Similarly, the transpose M^T is Perron-Frobenius and hence has a unique positive eigenvector. The i^{th} entry of this eigenvector corresponds to a length l assigned to the edge e_i . Furthermore, the positive eigenvalue associated to either such eigenvector equals the spectral radius of M (and of M^T).

3.2.1 Producing an efficient graph map

As previously mentioned, we begin with a spine of a punctured fiber surface F having monodromy ϕ , where the spine is viewed as a graph g with vertices and arbitrarily oriented edges. We represent ϕ as some composition of Dehn twists within a regular neighborhood of the spine and apply this composition to the edges of g . Viewing ϕ as a graph map, we perform the following sequence of steps involving Moves 1-6.

Step 1

Pull ϕ tight wherever possible.

Step 2

If ϕ has a nontrivial invariant forest, collapse it. If this results in back-tracking, perform Step 1 again.

Step 3

Remove all valence 1 vertices by performing valence 1 isotopies. If this results in back-tracking, perform Step 1 again. If this results in a nontrivial invariant forest, perform Step

2 again.

Step 4

Remove all valence 2 vertices by performing valence 2 isotopies. When performing a given valence 2 isotopy in this step, we delete the incident edge having larger width w as determined by M . If both edges have the same width, we can delete either edge. If this results in backtracking, perform Step 1 again. If this results in a nontrivial invariant forest, perform Step 2 again.

Step 5

At this point, if ϕ is efficient, the algorithm stops. Otherwise, there exists some edge e and some positive integer m such that $\phi^m(e)$ backtracks. Choose m to be minimal. If $m = 2$, so that $e \mapsto e_1e_2 \mapsto \cdots f\bar{f}\cdots$, we fold to e' all edges in the link of $\{\bar{e}_1, e_2\}$ having common starting image the edge-path P , where $\phi(e') = P$. This operations allows us to pull $\phi(e)$ tight, thus removing the back-tracking resulting from the edge-path e_1e_2 . After performing this fold, we return to Step 1.

If $m > 2$, so that $e \mapsto e_1e_2 \mapsto \cdots \mapsto h_1h_2 \mapsto i_1i_2 \mapsto j_1j_2 \mapsto k\bar{k}$, we morally fold at the link of $\{\bar{j}_1, j_2\}$, then at the link of $\{\bar{i}_1, i_2\}$, then at the link of $\{\bar{h}_1, h_2\}$, and so on until we fold at the link of $\{\bar{e}_1, e_2\}$. However, this sequence of folds may prevent us from pulling $\phi(e)$ tight. Let p be the point in the interior of e which maps to the point $\phi(p)$ that separates the edge-path e_1e_2 . If at any step in our folding sequence, a folded edge-path from the folded edge e starts (or ends) at $\phi(p)$, we subdivide e (possibly by first subdividing the edges appearing in the iterates of $\phi(e)$). This allows us to carry on our folding sequence without folding the point p at which some iterate of ϕ fails to be locally injective, so that we can ultimately pull $\phi(e)$ tight. After performing this sequence of folds, we return to Step 1. Note that this algorithm terminates because the pulling-tight portion of Step 6 always reduces the eigenvalue λ of M , and no other steps increase λ , until we have arrived at the dilatation λ_{\min} of ϕ .

3.2.2 Producing an invariant train track from an efficient graph map

After applying the algorithm described in the previous section to the spine graph g of a punctured fiber surface F having monodromy ϕ , we have an efficient graph map $\phi : g \rightarrow g$. We now describe the local behavior of this graph map near vertices. At a given vertex v , replace v with a small circle C having points on its boundary corresponding to the gates of the link of v . Outside of C , an edge e of g in the link of v meets a point on the circle corresponding to the gate of e , such that e is orthogonal to C . These are the so-called “real” edges of our train track. Within C , we connect two points p_1, p_2 on C corresponding to distinct gates with an “infinitesimal” edge f if there is some real edge e of g and some positive integer k such that $\phi^k(e)$ enters C via p_1 and exits C via p_2 . Additionally, we require that f is orthogonal to C . See Figure 3.5. By inserting infinitesimal edges at each vertex of g , we obtain an invariant train track τ for ϕ .

3.2.3 Producing a foliation from an invariant train track

Given an invariant train track τ for ϕ , we form a Markov partition of a regular neighborhood of τ as follows. Cover each edge e (both real and infinitesimal) by a rectangle of width $w(e)$ and length $l(e)$ as determined by the unique positive eigenvectors of, respectively, M and M^T . We do so in such a way that each rectangle meets each circle C in an arc of C , and such that rectangles covering adjacent edges within the same gate intersect at a single point. Rectangles have a canonical pair of transverse measured foliations consisting of vertical strips and horizontal strips, so our Markov partition yields a foliation of a regular neighborhood of τ .

We extend this foliation to the components of $F - \tau$ as follows. Given a k -sided component X having cusp vertices, some iterate of ϕ maps each side of ∂X over itself. This yields a fixed point of some iterate of ϕ associated to each of the k sides of ∂X , and any vertex of ∂X is equidistant from the fixed points on adjacent sides in terms of the length l . We can now collapse all sides of ∂X onto a spine of X in a length-preserving fashion by identifying all fixed points to a singular point (which is removed if X is the component containing the

puncture). See Figure 3.4.

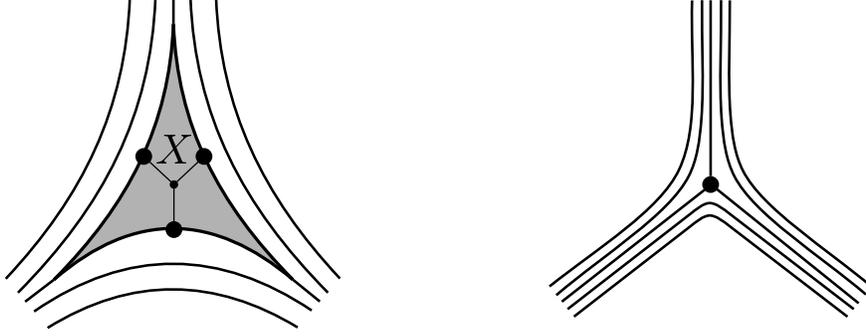


Figure 3.4. Extending a foliation to a face of $F - \tau$ by collapsing the face onto a spine.

Example 3.2.5. We now carry out the process described in this section to produce an efficient graph map, an invariant train track, and a foliation for a simple pseudo-Anosov map ϕ on the punctured torus. The initial graph and Dehn curves are pictured in Figure 3.5, where ϕ is a right Dehn twist about A , followed by a left Dehn twist about B . The resulting graph map

$$a \xrightarrow{D_A^R} a \xrightarrow{D_B^L} ba, \quad b \xrightarrow{D_A^R} ba \xrightarrow{D_B^L} bba$$

is already efficient since we have

$$\begin{aligned} a &\mapsto ba \mapsto \cdots ab \cdots \mapsto \cdots ab \cdots \mapsto \cdots \\ b &\mapsto bb \cdots \mapsto \cdots ab \cdots \mapsto \cdots ab \cdots \mapsto \cdots \end{aligned}$$

We now construct the invariant train track τ . Replacing the vertex with a small circle C , we partition the real edges a, \bar{a}, b, \bar{b} into their gates at points on C . To calculate the gates, we iterate D on the link $\{a, \bar{a}, b, \bar{b}\}$ until the image becomes periodic. Once the image of D on this link becomes periodic, we classify two edges as being in the same gate if their image under this periodic iterate of D is the same, and as being in distinct gates otherwise. We calculate

$$\{a, \bar{a}, b, \bar{b}\} \xrightarrow{D} \{b, \bar{a}, b, \bar{a}\} \mapsto \{b, \bar{a}, b, \bar{a}\}$$

From this, we conclude that a, b are in the same gate, and \bar{a}, \bar{b} are in the same gate. Each of these two gates corresponds to a cusp. Now we insert the infinitesimal edges within C .

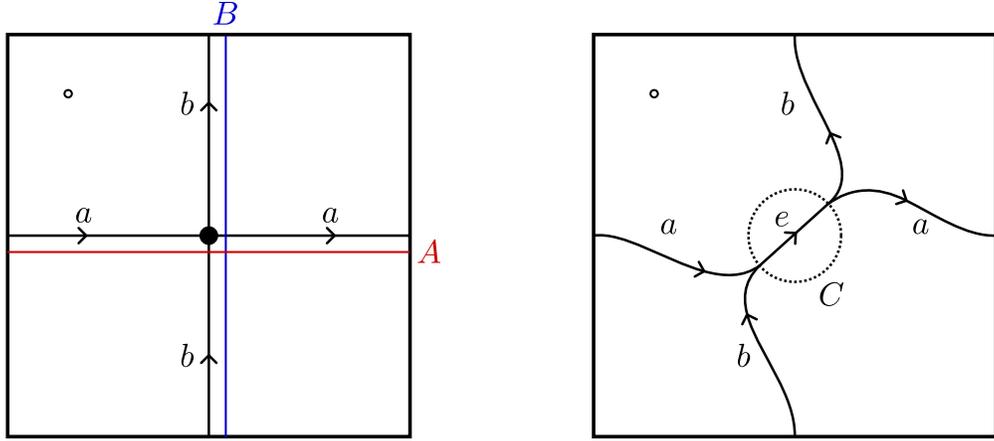


Figure 3.5. Our starting spine and Dehn curves (left), and adding the infinitesimal edge (right).

The edge-path $a \mapsto ba$ connects the distinct gates $\{\bar{b}, a\}$, so we insert an infinitesimal edge e connecting these two gates. Now our graph map is written as $a \mapsto bea$, $b \mapsto bebea$, $e \mapsto e$. The result is an invariant train track where the component containing the puncture is a disk with two cusps.

We now construct the pair of transverse foliations. The transition matrix is

$$M = \begin{pmatrix} 1 & 1 & 0 \\ 1 & 2 & 0 \\ 1 & 2 & 1 \end{pmatrix}$$

which has unique positive eigenvector $\left(\frac{3-\sqrt{5}}{2}, \frac{\sqrt{5}-1}{2}, 1\right)$, so that $w(a) = \frac{3-\sqrt{5}}{2}$, $w(b) = \frac{\sqrt{5}-1}{2}$, $w(e) = 1$. The transpose M^T has unique positive eigenvector $\left(\frac{\sqrt{5}-1}{2}, 1, 0\right)$, so that $l(a) = \frac{\sqrt{5}-1}{2}$, $l(b) = 1$, $l(e) = 0$. We cover each edge with a rectangle having length and width corresponding to that edge's l, w weight. See Figure 3.6 for the Markov partition and its image under ϕ .

To obtain the associated unstable foliation, we foliate the rectangles of the Markov partition, and then we collapse the 2-gon containing the puncture onto a spine of the 2-gon. See Figure 3.7.

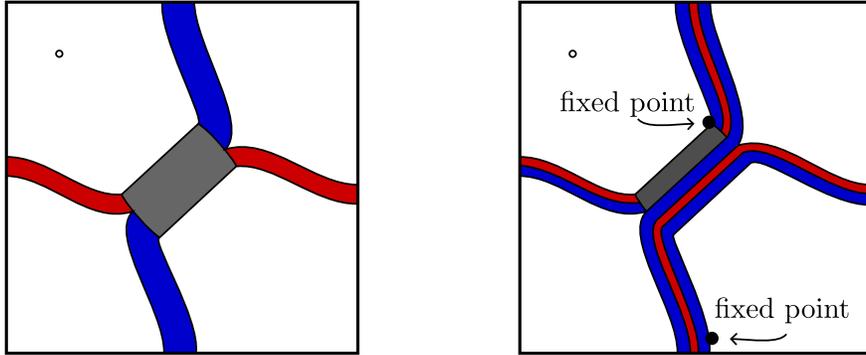


Figure 3.6. The initial Markov partition (left), and its image under ϕ (right).

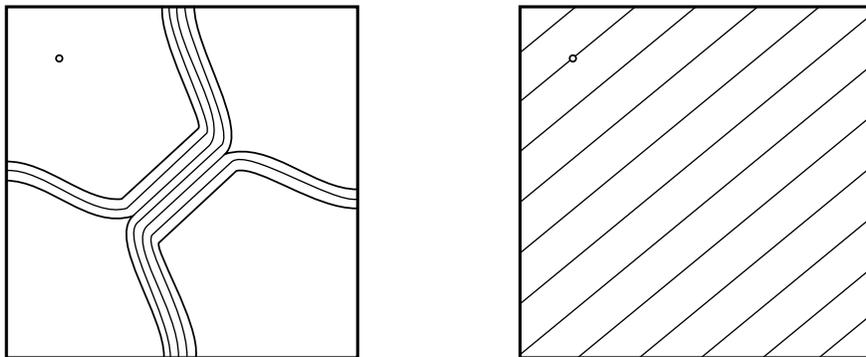


Figure 3.7. The unstable foliation of a regular neighborhood of τ (left), and the extension of this foliation to F (right).

CHAPTER 4

SYMPLECTIC FLOER HOMOLOGY

4.1 Cotton-Clay's calculation of symplectic Floer homology

In [5], Cotton-Clay calculated the symplectic Floer homology HF_* of mapping classes of pseudo-Anosov maps of compact surfaces. For such a mapping class ϕ , the rank of HF_* is simply a count of the (nondegenerate) fixed points of a *symplectic* representative of ϕ . Consequently, Cotton-Clay was able to give a combinatorial formula for the rank of HF_* from an invariant train track for a *singular* representative of ϕ . However, this formula is stated for closed surfaces, while our present application requires such a formula for surfaces with one boundary component. At the end of this section, we give such a formula based on Cotton-Clay's work.

We now outline the basic construction of HF_* , which is also known as fixed point Floer homology. For more details, see [5, 7, 10, 32]. Throughout the rest of this section, we consider a compact symplectic surface (F, ω) of negative Euler characteristic and an area-preserving symplectomorphism $\phi : F \rightarrow F$ with nondegenerate fixed points.

Definition 4.1.1. A fixed point x is *nondegenerate* if the linear differential $d\phi_x$ does not have 1 as an eigenvalue. A map ϕ is *nondegenerate* if all of its fixed points are nondegenerate.

In particular, such a ϕ has a finite fixed point set, denoted $\text{Fix}(\phi)$, consisting of isolated fixed points. After enforcing certain technical conditions on ϕ and ω , we can take the homology of the chain complex

$$CF_*(\phi) := \mathbb{Z}_2\langle x \mid x \in \text{Fix}(\phi) \rangle,$$

where the differential ∂ counts mod 2 the flow lines from x to y in a suitable moduli space. More precisely, let $J = \{J_t\}_{t \in \mathbb{R}}$ be a smooth path of ω -compatible almost-complex structures on F such that $J_{t+1} = \phi^* J_t$. We define the moduli space $M_1(x, y, J)$ as the set of maps $u : \mathbb{R}^2 \rightarrow F$ such that

1. $u(s, t) = \phi(u(s, t + 1))$,
2. $\lim_{s \rightarrow +\infty} u(s, \cdot) = x$, $\lim_{s \rightarrow -\infty} u(s, \cdot) = y$,

3. $\partial_s u + J_t(u)\partial_t u = 0,$

4. The Maslov index satisfies $\mu(u) = 1.$

We call such a u an index one *flow line* from x to y . After perturbing J to be “generic,” we realize $M_1(x, y, J)$ as a smooth 1-dimensional manifold. The quotient of $\mathcal{M}_1(x, y, J)$ by the free \mathbb{R} -action of translation in the s -variable yields a compact 0-dimensional manifold, so we can define

$$\partial x := \sum_{y \in \text{Fix}(\phi)} \left(\left| \frac{\mathcal{M}_1(x, y, J)}{\mathbb{R}} \right| \text{ mod } 2 \right) y.$$

In Proposition 3.1.1 of [5], it is shown that $HF_*(\phi)$ is well-defined for a nondegenerate symplectomorphism ϕ in a pseudo-Anosov mapping class h . In particular, $HF_*(\phi) = HF_*(h)$ is independent of choice of $\omega, J,$ and ϕ .

There are two necessary conditions for a flow line to exist between two fixed points x, y . The first condition is that x, y belong to the same Nielsen class.

Definition 4.1.2. Two fixed points x, y of ϕ are in the same *Nielsen class* if there exists a path $\gamma : [0, 1] \rightarrow F$ with $\gamma(0) = x, \gamma(1) = y$ such that $\phi(\gamma)$ is homotopic (rel. endpoints) to γ .

Indeed, given a flow line $u(s, t)$ from x to y , let $\gamma_t : [0, 1] \rightarrow F$ be a reparametrization of $u(\cdot, t)$ (for each t). Then γ_1 is the desired Nielsen equivalence of x and y as described in [5]. The second condition is that the Lefschetz signs of x, y are different.

Definition 4.1.3. The *Lefschetz sign* of a nondegenerate fixed point x is the quantity defined by $(-1)^{\varepsilon(x)} = \text{sign}(\det(\text{id} - d\phi_x)) \in \{\pm 1\}.$

Indeed, this condition follows since we require that the Maslov index satisfies $\mu(u) = 1,$ and we also have $\mu(u) = \varepsilon(x) - \varepsilon(y) \text{ mod } 2$ as in [10, 22]. Since the linear map $d\phi_x$ is symplectic, its two eigenvalues are either real with product 1, or imaginary with product 1. Thus, we can partition $\text{Fix}(\phi)$ into the following three classes.

Definition 4.1.4. A nondegenerate fixed point x is called *elliptic* if the eigenvalues of $d\phi_x$ lie on the unit circle. A nondegenerate fixed point x is called *positive hyperbolic* if the

eigenvalues of $d\phi_x$ are real and positive. A nondegenerate fixed point x is called *negative hyperbolic* if the eigenvalues of $d\phi_x$ are real and negative.

Note that the Lefschetz sign of elliptic and negative hyperbolic fixed points is 1, and the Lefschetz sign of positive hyperbolic fixed points is -1 . In particular, there can only be nontrivial flow lines between elliptic fixed points and positive hyperbolic fixed points, or between positive and negative hyperbolic fixed points.

4.2 Combinatorial formulae for HF_*

A (singular) pseudo-Anosov map $\phi : F \rightarrow F$ is smooth away from the singularities of its associated foliations. Furthermore, with respect to an appropriate area form on F , ϕ is symplectic away from these singularities. In Sections 3.2 and 4.2 of [5], Cotton-Clay uses Hamiltonian vector fields to give a symplectic smoothing of ϕ near these singularities and near ∂F (if $\partial F \neq \emptyset$). The result is a symplectic representative ϕ_{sm} of ϕ having nondegenerate fixed points.

In obtaining ϕ_{sm} , Cotton-Clay also determines the number and type of fixed points that result near each singularity and each puncture (i.e. each boundary component). At a rotated interior singularity which is a p -prong, we have one elliptic fixed point corresponding to the singularity, a so-called Type IIIc fixed point. Near a fixed, unrotated interior singularity which is a p -prong ($p \geq 3$), we have $(p - 1)$ positive hyperbolic fixed points, which are so-called Type IIIb- p fixed points. Near a rotated puncture, we have no fixed points. Near a fixed, unrotated puncture which is a p -prong ($p \geq 1$), we have p positive hyperbolic fixed points, which are so-called Type IIId- p fixed points.

In [3], it was shown that no two fixed points of a singular pseudo-Anosov map ϕ on a closed surface are in the same Nielsen class. Furthermore, Cotton-Clay shows that fixed points of ϕ_{sm} associated to a given interior singularity are in their own Nielsen class, and fixed points of ϕ_{sm} associated to a given puncture are in their own Nielsen class. Since the fixed points of ϕ_{sm} associated to a given interior singularity or puncture are all of the same Lefschetz sign by the preceding discussion, we conclude that all differentials vanish in

$HF_*(\phi_{sm})$. Hence, the rank of $HF_*(\phi_{sm})$ is given by $\text{Fix}(\phi_{sm})$, which equals the minimum number of fixed points of a nondegenerate symplectic representative of ϕ . In particular, $\text{Fix}(\phi_{sm})$ is well-defined.

So we can calculate the rank of $HF_*(\phi_{sm})$ by counting the number of normal fixed points of ϕ and then by accounting for how $\text{Fix}(\phi)$ is affected when we replace ϕ with ϕ_{sm} . An invariant train track τ for ϕ yields the count of normal fixed points. Indeed, as described in Section 3.2.3, the transition matrix M of an efficient graph map ϕ yields a Markov partition of a regular neighborhood of τ consisting of rectangles. Under ϕ , the length of these rectangles is stretched by the dilatation λ , and the width of these rectangles is shrunk by λ . By the contraction mapping theorem, ϕ has a fixed point each time a rectangle maps over itself. The total number of times that rectangles map over themselves is given by $\text{tr}(M)$, so that $\text{tr}(M)$ is the fixed point count of ϕ , up to some correction terms.

There are over-counted fixed points resulting from a given edge b_i that, without loss of generality, has edge-path image $\phi(b_i) = b_i \cdots$. This would mean that there is another edge b_j in the link of b_i that, without loss of generality, has edge-path image $\phi(b_j) = b_j \cdots$. Cotton-Clay contends that such a fixed point contributes 2 to $\text{tr}(M)$, so that we have double-counted this fixed point.

However, there appears to be a small error in Cotton-Clay's description of these over-counted fixed points (which ultimately does not affect our calculation in Section 4.3). Indeed, Cotton-Clay captures such over-counting as $\sum s_i$, where s_i equals one-half the number of times that $\phi(b_i)$ starts or ends with b_i which has the correct orientation. However, it is possible that $\sum s_i$ is half-integer valued for once-punctured surfaces when the punctured region of $F - \tau$ is unrotated and has two cusps or has one cusp. For example, using the spine and monodromy from Example 3.2.3, we can isotope the curve A slightly upwards, so that the image of ϕ is given by $a \mapsto ba$, $b \mapsto bab$. This also specifies an invariant train track, and this would give $\sum s_i = 3/2$. When this phenomenon occurs for a punctured region with two cusps, two fixed points are over-counted in $\text{tr}(M)$ as 3. When this phenomenon occurs

for a punctured region with one cusp, one fixed point is over-counted in $\text{tr}(M)$ as 3. So we modify Cotton-Clay's $\sum s_i$ count to be

$$\tilde{s}_i := \begin{cases} \sum s_i, & \text{if } \sum s_i \in \mathbb{Z}, \\ \lfloor \sum s_i \rfloor, & \text{if } \sum s_i \in \mathbb{Z} + \frac{1}{2} \text{ and the punctured region has two cusps,} \\ \lfloor \sum s_i \rfloor + 1, & \text{if } \sum s_i \in \mathbb{Z} + \frac{1}{2} \text{ and the punctured region has one cusp.} \end{cases}$$

This definition of \tilde{s}_i would not accurately describe over-counting for all multi-punctured surfaces, so we emphasize that we are only considering surfaces with at most one puncture.

Finally, there is an under-counting of fixed points resulting from two distinct edges b_i, b_j that, without loss of generality, share an initial switch and have edge-path images $\phi(b_i) = b_j \cdots$ and $\phi(b_j) = b_i \cdots$. Such a fixed point does not contribute to $\text{tr}(M)$. This is called a *flip*; let f be the total number of flips.

From the preceding discussion, the number of normal fixed points of ϕ is

$$\text{Fix}(\phi) = \text{tr}(M) - \tilde{s}_i + f.$$

Assume, for now, that F is a closed surface. We describe the effect on the fixed points of ϕ when we replace ϕ with ϕ_{sm} . As mentioned above, the symplectic map ϕ_{sm} agrees with ϕ away from the singularities. A fixed interior p -pronged singular point ($p \geq 3$) that is nontrivially rotated has one elliptic fixed point which is not accounted for in $\text{tr}(M)$. Denote by r the number of rotated fixed interior singular points. A fixed interior p -pronged singular point ($p \geq 3$) that is unrotated has $(p - 1)$ positive hyperbolic fixed points, which are over-counted in $\text{tr}(M)$ as p , corresponding to the p sides of the associated component of $F - \tau$ that are collapsed to the singularity. Denote by u the number of unrotated fixed interior singular points. This yields the following formula of Cotton-Clay.

Theorem 4.2.1. *For a closed surface F and a pseudo-Anosov map $\phi : F \rightarrow F$,*

$$\text{rk } HF_*(\phi) = \text{tr}(M) - \sum_i s_i + f + r - u.$$

Now we assume that F has one boundary component. The only difference from the closed case is that one component of $F - \tau$ contains a puncture. The contribution of fixed points of the faces which do not contain the puncture remains the same, so we must now consider how the face containing the puncture affects our calculation. Say this face is a p -gon, $p \geq 1$. If this face is rotated, it contributes no normal fixed points to ϕ (and nothing to $\text{tr}(M)$). After replacing ϕ with ϕ_{sm} , we obtain no additional fixed points. If this face is unrotated, it contributes a normal fixed point for each side of the p -gon, namely p fixed points, and these fixed points are already accounted for in the trace. If this face has one cusp or two cusps, there may be some over-counting as described above, but this is accounted for in \tilde{s}_i . After replacing ϕ with ϕ_{sm} , these p positive hyperbolic fixed points are retained and are in the same Nielsen class. Consequently, we have the following theorem.

Theorem 4.2.2. *For a surface F with one boundary component and a pseudo-Anosov map $\phi : F \rightarrow F$, the rank of $HF_*(\phi)$ is*

$$\text{rk } HF_*(\phi) = \text{tr}(M) - \tilde{s}_i + f + r - u.$$

4.3 Invariant train tracks for $K(j, |\beta_B| - j)$

In this section, we record invariant train tracks for the monodromies of a particular family of 3-string braids. Let L_A be the closure of the braid $\beta_A = \sigma_1 \sigma_1$, and let L_B be the closure of the braid $\beta_B = \sigma_1^{-1} \cdots \sigma_1^{-1}$, where $|\beta_B| \geq 2$ is even. We consider those knots which can be written as $K = L_A \star_4 L_B$ such that K is the closure of the braid

$$\sigma_1 \underbrace{\sigma_2^{-1} \cdots \sigma_2^{-1}}_j \sigma_1 \underbrace{\sigma_2^{-1} \cdots \sigma_2^{-1}}_{|\beta_B| - j}.$$

Denote such a knot as $K(j, |\beta_B| - j)$, where j is odd.

Before we present our train tracks for $K(j, |\beta_B| - j)$, we make some remarks. Since the images of infinitesimal edges can be reconstructed from the images of real edges, we only label the real edges in our graphs and calculate the images of the real edges. Furthermore, since we are only interested in the behavior of our train tracks at switches, i.e. where edges

meet other edges, we illustrate the behavior of our train tracks only near switches. Finally, by noting that $K(j, |\beta_B| - j)$ is isotopic to $K(|\beta_B| - j, j)$, one can verify that the five train tracks that follow are sufficient to give a train track for the monodromy of any knot of the form $K(j, |\beta_B| - j)$. These calculations are explicitly performed in the Appendix.

An invariant train track for the monodromy of the closure of $K(1, 1) = 4_1$ is given by Figure 4.1. An invariant train track for the monodromy of the closure of $K(3, 3) = \overline{8}_5$ is given by Figure 4.2. An invariant train track for the monodromy of $K(|\beta_B| - 1, 1)$, where $|\beta_B| \geq 3$, is given by Figure 4.3. An invariant train track for the monodromy of $K(j, 3)$, where $j \geq 5$, is given by Figure 4.4. An invariant train track for the monodromy of $K(j, |\beta_B| - j)$, where $j \geq 5$ and $|\beta_B| - j \geq 5$ is given by Figure 4.5.

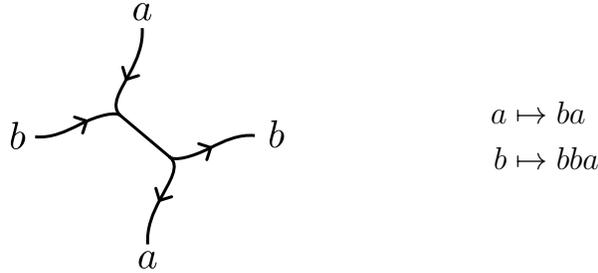


Figure 4.1. An invariant train track for the monodromy of $K(1, 1) = 4_1$.

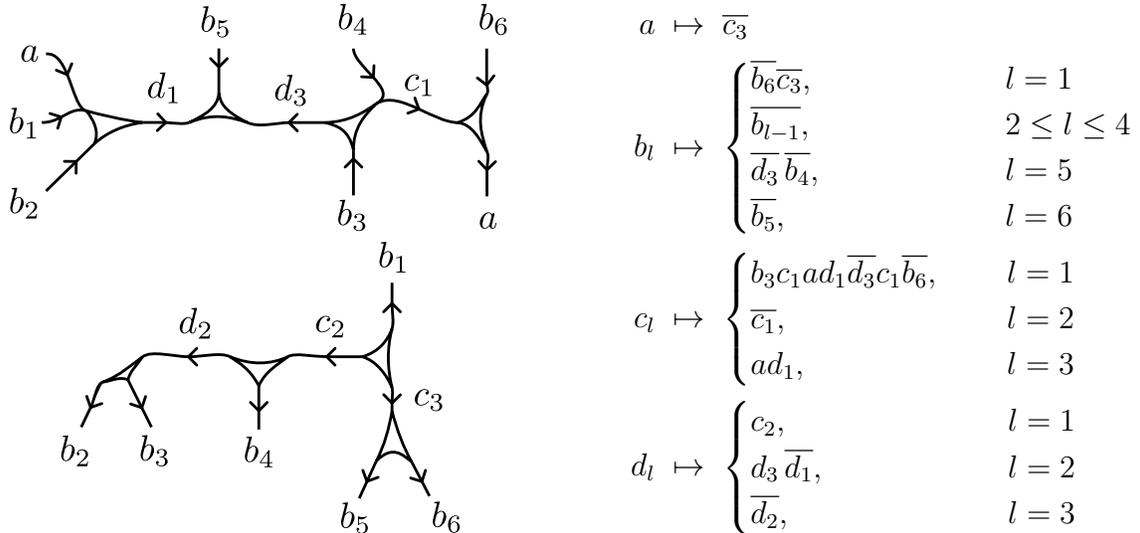


Figure 4.2. An invariant train track for the monodromy of $K(3, 3) = \overline{8}_5$.

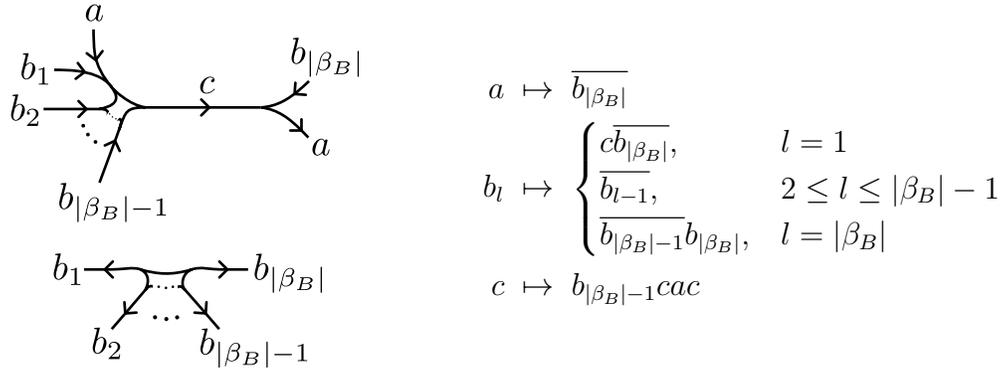


Figure 4.3. An invariant train track for the monodromy of $K(|\beta_B| - 1, 1)$, where $|\beta_B| \geq 3$.

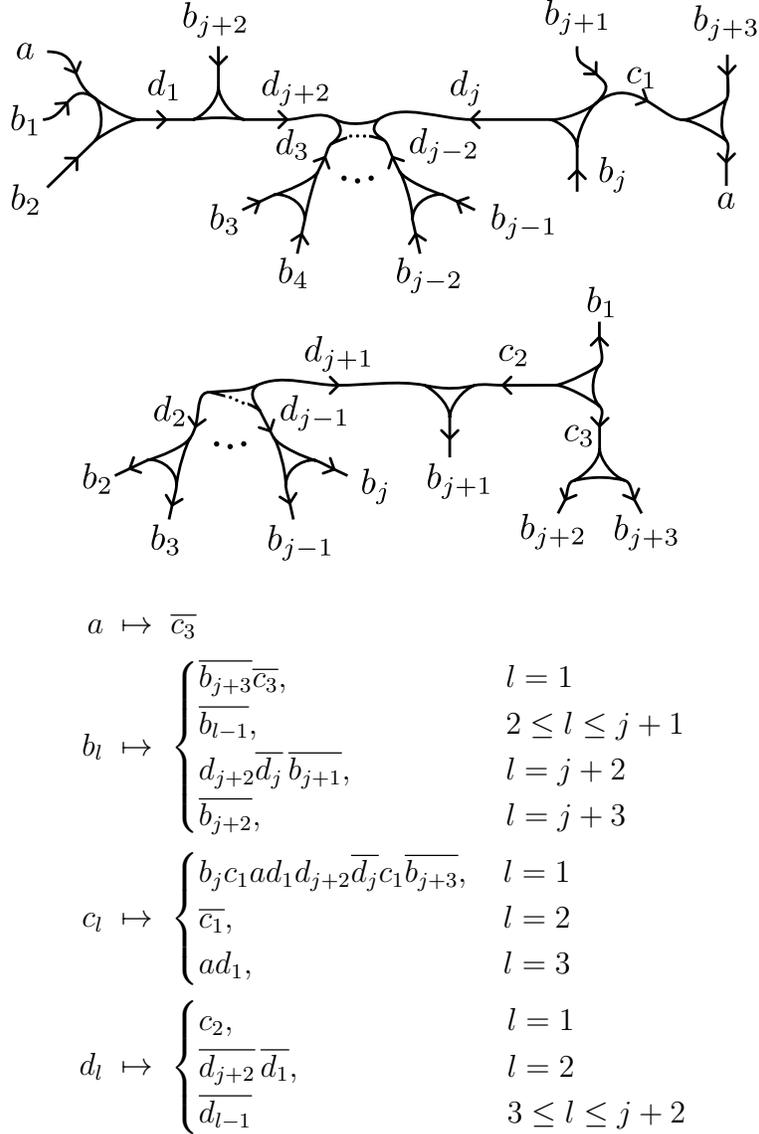


Figure 4.4. An invariant train track for the monodromy of $K(j, 3)$, where $j \geq 5$.

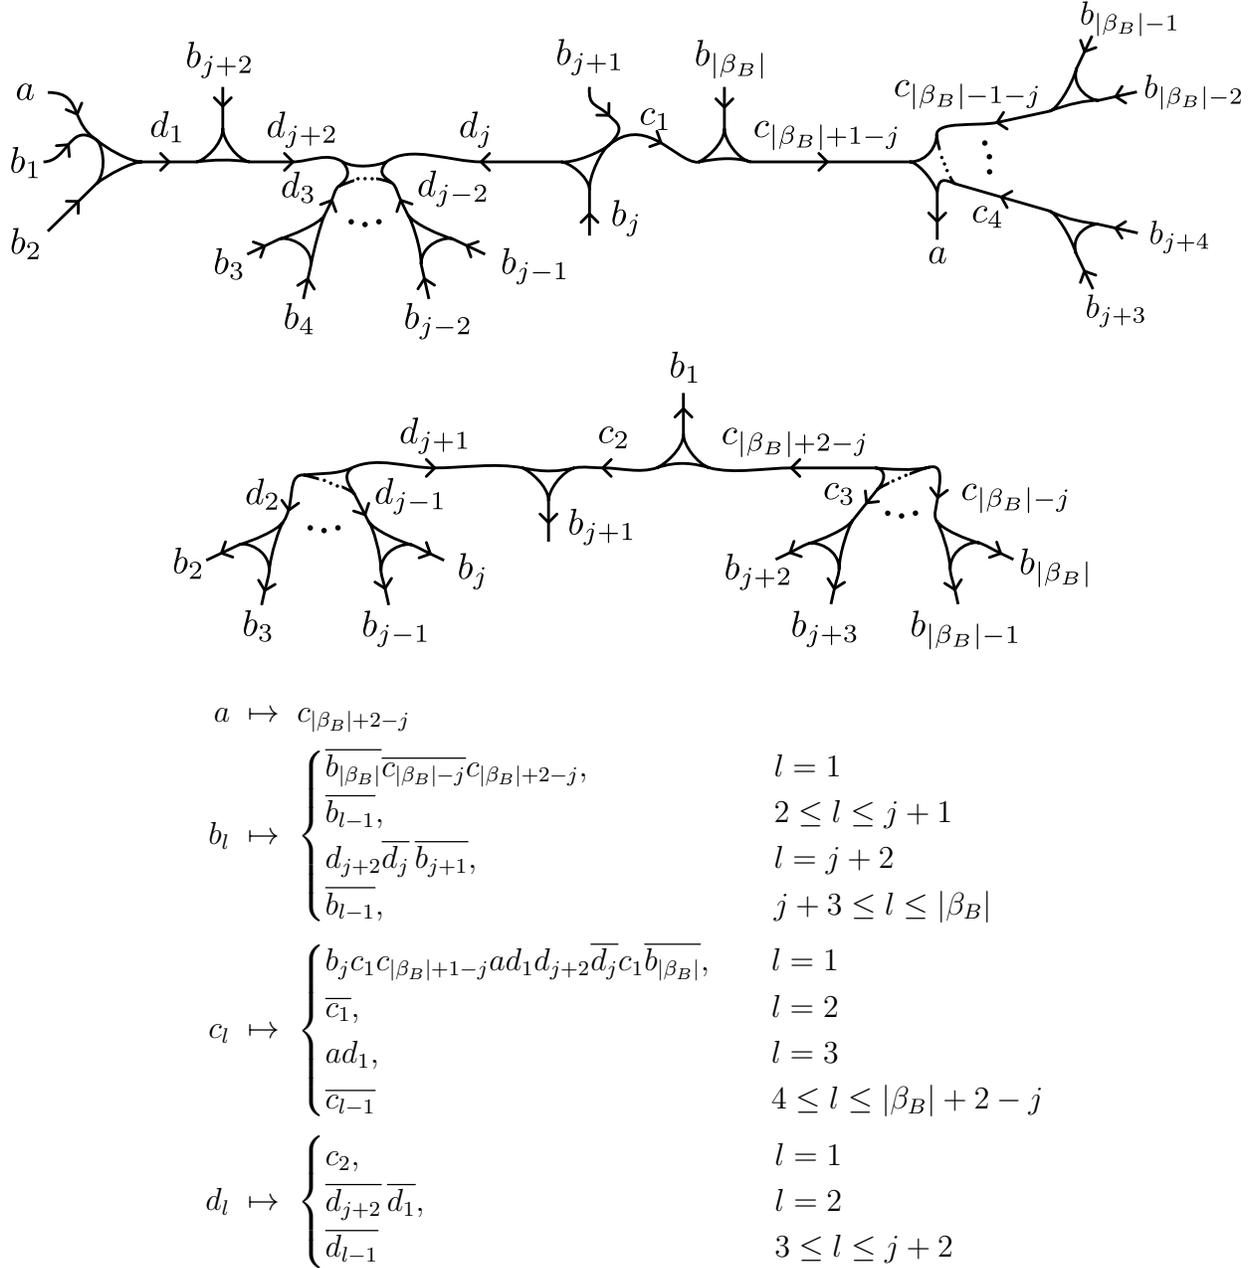


Figure 4.5. An invariant train track for the monodromy of $K(j, |\beta_B| - j)$, where $j \geq 5$ and $|\beta_B| - j \geq 5$.

4.4 Calculation of $HF_*(\phi_K)$ for $K = K(j, |\beta_B| - j)$

As described in Section 4.3, we can calculate the rank of $HF_*(\phi)$ from an invariant train track for ϕ . In particular, we apply Theorem 4.2.2 to each of the invariant train tracks just presented. As described in the Appendix, in each such train track, every interior p -pronged singular point ($p \geq 3$) is permuted with other singular points, so the contribution of $\sum_p (r_p - u_p)$ is 0. In each train track, there are also no flips, so the contribution of f is 0. Indeed, one can check that in each train track, there are no edges e_1, e_2 which emanate from the same vertex such that $\phi(e_1) = e_2 \cdots$ and $\phi(e_2) = e_1 \cdots$. We have an over-counting of fixed points in the train track for $K(1, 1) = 4_1$ corresponding to the pair $a \mapsto \cdots a, b \mapsto b \cdots$. The trace of the transition matrix of this train track is 3. Similarly, we have an over-counting of fixed points in the train track for $K(|\beta_B| - 1, 1)$, where $|\beta_B| \geq 3$, corresponding to the pair $b_{|\beta_B|} \mapsto \cdots b_{|\beta_B|}, c \mapsto \cdots c$. The trace of the transition matrix of this train track is 3. All of the other train tracks have no over-counting and have transition matrices with trace equal to 2. So the rank of HF_* is 2 in every case.

Theorem 4.4.1. *For $K = K(j, |\beta_B| - j)$ as defined above,*

$$\text{rk } HF_*(\phi_K) = 2.$$

Such calculations of $HF_*(\phi)$, where ϕ is some composition of Dehn twists along a nicely chosen collection of simple closed curves, have been performed via other methods [8]. However, to the author's knowledge, the maps ϕ considered here are not present in the literature. It should be noted that an invariant train track for the pseudo-Anosov monodromy of an iterated plumbing of Hopf bands is given in [13], but this is done without the Bestvina-Handel algorithm and without graphs.

One reason that such calculations are relevant is the recently established correspondence for a fibered knot K between the “next-to-top” group of $\widehat{HF\bar{K}}(K)$ and $HF_*(\phi_K)$ [15]. Indeed, while the “top” group of $\widehat{HF\bar{K}}$ yields the genus and fibered status of a knot [14, 37, 39], the “next-to-top” group of $\widehat{HF\bar{K}}$ of a nontrivial fibered knot K yields the number of nondegenerate

fixed points (plus one) of a symplectic representative of the monodromy of ϕ_K . Consequently, we have the following corollary.

Corollary 4.4.1. *For $K = K(j, |\beta_B| - j)$ as defined above,*

$$\text{rk } \widehat{\text{HFK}}(K, g(K) - 1) = 3.$$

Furthermore, the Alexander polynomial takes the form

$$\Delta_K(t) = t^{-g(K)} - 3t^{1-g(K)} \dots - 3t^{g(K)-1} + t^{g(K)}$$

4.5 Further research

In Chapter 2, we have shown that given a knot K and any two knots K_1, K_2 , we can form K as a Murasugi sum of K_1 and K_2 , and that this situation can be illustrated in a closed braid form. So anything can happen when we use Murasugi sums of general complexity, but as we have also shown, Murasugi sums are more well-behaved when the size of Murasugi sums is restricted. For further study, we may also put other restrictions on the Murasugi sums, for example by restricting the genera of the Seifert surfaces involved in the Murasugi sums. So we can ask the following decomposition question about the result of the Murasugi sum:

Question. *For a nontrivial knot K with genus $g(K)$, what is the minimal genus of a Seifert surface F for K which is a Murasugi sum of two unknots?*

In this situation, if K has an alternating diagram which can be unknotted with $u(K)$ crossing changes, then Seifert's algorithm yields a minimal genus Seifert surface, so our constructions give that $1 \leq g(F) - g(K) \leq u(K)$. Note that this inequality is also true for knots with $u(K) = 1$, because by [28], there exists a minimal genus Seifert surface for K on which the unknotting crossing change can be done by twisting a band. Similarly, we can ask the following composition question about the summands of the Murasugi sum:

Question. *For two nontrivial knots K_1, K_2 , what are the minimal genera of Seifert surfaces F_1, F_2 of K_1, K_2 that Murasugi sum to the unknot?*

Answering these questions in general seems to be more involved than what we treat in this thesis, where we typically form some Seifert surfaces, manipulate them, then dissolve

them to obtain their knot boundaries. In particular, one would need to be more explicit with how the original Seifert surfaces are formed. Studying these questions via band surgery and band twisting (or perhaps other moves) should yield insight into a collection of subgraphs of $MSG(K)$ similar to $MSG(K, n)$.

In Chapter 4, we calculated the rank of HF_* of the monodromies of a family of knots which are the closures of particular homogeneous 3-string braids. It is the hope of the author that similar techniques can be used to calculate HF_* of the knot closure of any homogeneous braid with pseudo-Anosov monodromy. Such knots can be viewed as iterated Murasugi sums of closures of positive braids. Noting that such a statement for closures of positive braids (which include torus knots having periodic monodromy) might be proven by using the Skein exact sequence of \widehat{HFK} , we present the following conjecture.

Conjecture 4.5.1. *Suppose β is a homogeneous braid in the generators $\sigma_1, \dots, \sigma_n$ such that each σ_i appears at least twice, and suppose that the closure of β is a knot K . For $1 \leq i \leq n - 1$, consider a subword of β which begins with $\sigma_i^{\pm 1}$, ends with $\sigma_{i+1}^{\mp 1}$, and contains no other $\sigma_i^{\pm 1}$ or $\sigma_{i+1}^{\mp 1}$. Similarly, consider a subword of β which begins with $\sigma_{i+1}^{\mp 1}$, ends with $\sigma_i^{\pm 1}$, and contains no other $\sigma_i^{\pm 1}$ or $\sigma_{i+1}^{\mp 1}$. Let a_i be the number of times that a subword of one of these two types appears in β . Then*

$$\text{rk } \widehat{HFK}(K, g(K) - 1) = 1 + \sum_{i=1}^{n-1} \left\lceil \frac{a_i}{2} \right\rceil.$$

For example, the knot $12n_{22}$ is the closure of the homogenous braid

$$\sigma_1^{-1} \sigma_2 \sigma_3 \sigma_1^{-1} \sigma_3 \sigma_4^{-1} \sigma_2 \sigma_4^{-1} \sigma_3 \sigma_4^{-1} \sigma_3 \sigma_4^{-1}.$$

Here, $a_1 = 3, a_2 = 0, a_3 = 5$, so our conjecture (accurately) predicts that

$$\text{rk } \widehat{HFK}(12n_{22}, g(12n_{22}) - 1) = 6.$$

More generally, it would be interesting to describe the behavior of the next-to-top group of \widehat{HFK} of all fibered knots under Murasugi sum along minimal genus Seifert surfaces. While the author has not succeeded in proving a statement of this type using Heegaard diagrams, we nevertheless state the following generalization of Conjecture 4.5.1.

Conjecture 4.5.2. *For fibered knots K, K_1, K_2 , suppose that $K = K_1 \star_m K_2$, where the Murasugi sum is performed along minimal genus Seifert surfaces. Then*

$$\left| \text{rk } \widehat{\text{HFK}}(K, g(K) - 1) - \text{rk } \widehat{\text{HFK}}(K_1, g(K_1) - 1) - \text{rk } \widehat{\text{HFK}}(K_2, g(K_2) - 1) \right| \leq \frac{m}{2} - 1.$$

Clearly, this conjecture is true for $m = 2$, as this agrees with the Künneth formula for $\widehat{\text{HFK}}$ of the connected sum of knots due to Ozsváth-Szabó [41]. In proving this conjecture, one would give a deeper understanding of how symplectic fixed points behave under Murasugi sums, and one would also give another obstruction to forming Murasugi sums.

BIBLIOGRAPHY

- [1] Jared Able and Mikami Hirasawa. Construction and decomposition of knots as Murasugi sums of Seifert surfaces, December 2021. arXiv:2112.00297.
- [2] M. Bestvina and M. Handel. Train-tracks for surface homeomorphisms. *Topology*, 34(1):109–140, 1995.
- [3] Joan S. Birman and Mark E. Kidwell. Fixed points of pseudo-Anosov diffeomorphisms of surfaces. *Adv. in Math.*, 46(2):217–220, 1982.
- [4] Zhechi Cheng, Matthew Hedden, and Sucharit Sarkar. Murasugi sum and extremal knot floor homology, February 2022. arXiv:2202.09041.
- [5] Andrew Walker Cotton-Clay. *Symplectic Floer homology of area-preserving surface diffeomorphisms and sharp fixed point bounds*. ProQuest LLC, Ann Arbor, MI, 2009. Thesis (Ph.D.)—University of California, Berkeley.
- [6] P. R. Cromwell. Some infinite families of satellite knots with given Alexander polynomial. *Mathematika*, 38(1):156–169, 1991.
- [7] Stamatis Dostoglou and Dietmar A. Salamon. Self-dual instantons and holomorphic curves. *Ann. of Math. (2)*, 139(3):581–640, 1994.
- [8] Eaman Eftekhary. Floer homology of certain pseudo-Anosov maps. *J. Symplectic Geom.*, 2(3):357–375, 2004.
- [9] David Eisenbud and Walter Neumann. *Three-dimensional link theory and invariants of plane curve singularities*, volume 110 of *Annals of Mathematics Studies*. Princeton University Press, Princeton, NJ, 1985.
- [10] Alexander Fel’shtyn. The growth rate of symplectic Floer homology. *J. Fixed Point Theory Appl.*, 12(1-2):93–119, 2012.
- [11] David Gabai. The Murasugi sum is a natural geometric operation. In *Low-dimensional topology (San Francisco, Calif., 1981)*, volume 20 of *Contemp. Math.*, pages 131–143. Amer. Math. Soc., Providence, RI, 1983.
- [12] David Gabai. The Murasugi sum is a natural geometric operation. II. In *Combinatorial methods in topology and algebraic geometry (Rochester, N.Y., 1982)*, volume 44 of *Contemp. Math.*, pages 93–100. Amer. Math. Soc., Providence, RI, 1985.
- [13] David Gabai and William H. Kazez. Pseudo-Anosov maps and surgery on fibred 2-bridge knots. *Topology Appl.*, 37(1):93–100, 1990.
- [14] Paolo Ghiggini. Knot Floer homology detects genus-one fibred knots. *Amer. J. Math.*, 130(5):1151–1169, 2008.

- [15] Paolo Ghiggini and Gilberto Spano. Knot Floer homology of fibred knots and Floer homology of surface diffeomorphisms, January 2022. arXiv:2201.12411.
- [16] Amrenda Gill, Madeti Prabhakar, and Andrei Vesnin. Gordian complexes of knots and virtual knots given by region crossing changes and arc shift moves. *J. Knot Theory Ramifications*, 29(10):2042008, 24, 2020.
- [17] Allen Hatcher and William Thurston. Incompressible surfaces in 2-bridge knot complements. *Invent. Math.*, 79(2):225–246, 1985.
- [18] Matthew Hedden, Paul Kirk, and Charles Livingston. Non-slice linear combinations of algebraic knots. *J. Eur. Math. Soc. (JEMS)*, 14(4):1181–1208, 2012.
- [19] Matthew Hedden and Liam Watson. On the geography and botany of knot Floer homology. *Selecta Math. (N.S.)*, 24(2):997–1037, 2018.
- [20] Mikami Hirasawa and Yoshiaki Uchida. The Gordian complex of knots. *J. Knot Theory Ramifications*, 11(3):363–368, 2002.
- [21] Sumiko Horiuchi, Kasumi Komura, Yoshiyuki Ohyama, and Masafumi Shimozawa. The Gordian complex of virtual knots. *J. Knot Theory Ramifications*, 21(14):1250122, 11, 2012.
- [22] Michael Hutchings. An index inequality for embedded pseudoholomorphic curves in symplectizations. *J. Eur. Math. Soc. (JEMS)*, 4(4):313–361, 2002.
- [23] Atsushi Ishii and Kengo Kishimoto. The IH-complex of spatial trivalent graphs. *Tokyo J. Math.*, 33(2):523–535, 2010.
- [24] Bo Ju Jiang and Jian Han Guo. Fixed points of surface diffeomorphisms. *Pacific J. Math.*, 160(1):67–89, 1993.
- [25] Andras Juhasz. Knot Floer homology and Seifert surfaces. *Algebr. Geom. Topol.*, 8(1):603–608, 2008.
- [26] Taizo Kanenobu. $SH(3)$ -move and other local moves on knots. *Osaka J. Math.*, 51(2):439–457, 2014.
- [27] Taizo Kanenobu and Hiromasa Moriuchi. $SH(3)$ -Gordian distances between knots with up to seven crossings. *Topology Appl.*, 196(part B):537–550, 2015.
- [28] Tsuyoshi Kobayashi. Minimal genus Seifert surfaces for unknotting number 1 knots. *Kobe J. Math.*, 6(1):53–62, 1989.
- [29] W. B. R. Lickorish. A finite set of generators for the homeotopy group of a 2-manifold. *Proc. Cambridge Philos. Soc.*, 60:769–778, 1964.
- [30] W. B. Raymond Lickorish. *An introduction to knot theory*, volume 175 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1997.

- [31] Andrei V. Mal'yunin. On the question of genericity of hyperbolic knots. *Int. Math. Res. Not. IMRN*, (21):7792–7828, 2020.
- [32] Dusa McDuff and Dietmar Salamon. *Introduction to symplectic topology*. Oxford Graduate Texts in Mathematics. Oxford University Press, Oxford, third edition, 2017.
- [33] Paul Melvin and Hugh Morton. Fibred knots of genus 2 formed by plumbing Hopf bands. *J. Lond. Math. Soc. (2)*, 34(1):159–168, 1986.
- [34] Kunio Murasugi. On a certain subgroup of the group of an alternating link. *Amer. J. Math.*, 85:544–550, 1963.
- [35] Kunio Murasugi. On a certain numerical invariant of link types. *Trans. Amer. Math. Soc.*, 117:387–422, 1965.
- [36] Yasutaka Nakanishi and Yoshiyuki Ohyama. The Gordian complex with pass moves is not homogeneous with respect to Conway polynomials. *Hiroshima Math. J.*, 39(3):443–450, 2009.
- [37] Yi Ni. A note on knot Floer homology of links. *Geom. Topol.*, 10:695–713, 2006.
- [38] Yi Ni. Sutured Heegaard diagrams for knots. *Algebr. Geom. Topol.*, 6:513–537, 2006.
- [39] Yi Ni. Knot Floer homology detects fibred knots. *Invent. Math.*, 170(3):577–608, 2007.
- [40] Yi Ni. Knot floer homology and fixed points, January 2022. arXiv:2201.10546.
- [41] Peter Ozsváth and Zoltán Szabó. Holomorphic disks and knot invariants. *Adv. Math.*, 186(1):58–116, 2004.
- [42] John Stallings. Constructions of fibred knots and links. In *Algebraic and geometric topology (Proc. Sympos. Pure Math., Stanford Univ., Stanford, Calif., 1976), Part 2*, Proc. Sympos. Pure Math., XXXII, pages 55–60. Amer. Math. Soc., Providence, R.I., 1978.
- [43] Abigail Thompson. A note on Murasugi sums. *Pacific J. Math.*, 163(2):393–395, 1994.
- [44] William P. Thurston. Three-dimensional manifolds, Kleinian groups and hyperbolic geometry. *Bull. Amer. Math. Soc. (N.S.)*, 6(3):357–381, 1982.
- [45] William P. Thurston. On the geometry and dynamics of diffeomorphisms of surfaces. *Bull. Amer. Math. Soc. (N.S.)*, 19(2):417–431, 1988.
- [46] Shuji Yamada. The minimal number of Seifert circles equals the braid index of a link. *Invent. Math.*, 89(2):347–356, 1987.
- [47] Kai Zhang, Zhiqing Yang, and Fengchun Lei. The $H(n)$ -Gordian complex of knots. *J. Knot Theory Ramifications*, 26(13):1750088, 7, 2017.

APPENDIX: CALCULATION OF INVARIANT TRAIN TRACKS FOR
 $K(j, |\beta_B| - j)$

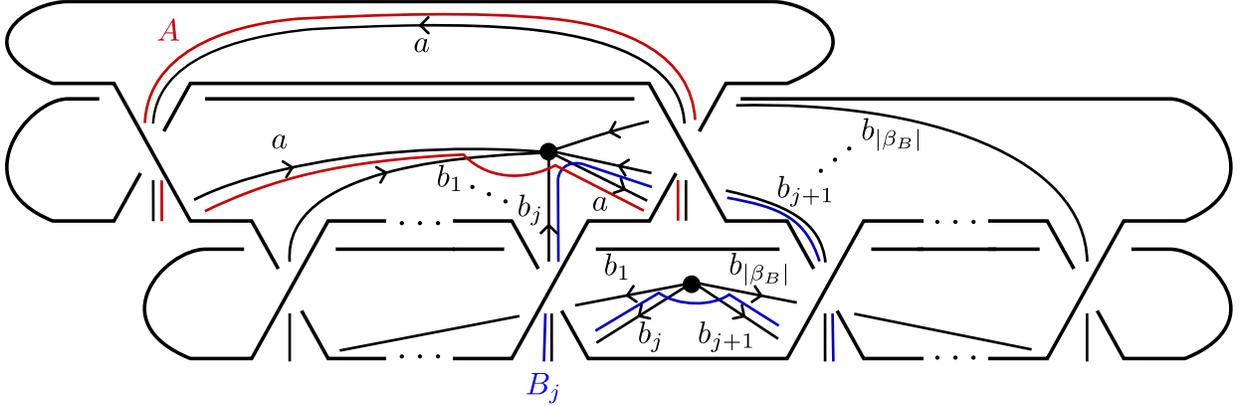


Figure 4.6. A spine of our surface, and the core curves A, B_j .

Here we explicitly calculate invariant train tracks of the knots $K(j, |\beta_B| - j)$ defined in Section 4.3. We begin with a spine of our surface, which consists of two vertices and the oriented edges $a, b_1, \dots, b_{|\beta_B|}$. See Figure 4.6. We choose the canonical homology basis A for the Hopf band which is the closure of $\beta_A = \sigma_1 \sigma_1$, so that $A \simeq a$, and the canonical homology basis $B_1, \dots, B_{|\beta_B|-1}$ for the closure of $\beta_B = \sigma_1^{-1} \dots \sigma_1^{-1}$, so that $B_l \simeq b_l \overline{b_{l+1}}$ for $1 \leq l \leq |\beta_B| - 1$. We isotope these curves so that the intersection between A and B_l is empty except for some fixed $l = j$. We further isotope $A, B_1, \dots, B_{|\beta_B|-1}$ so that A intersects b_1, \dots, b_j each geometrically once, B_l intersects b_l, b_{l+1} each geometrically once for $1 \leq l \leq |\beta_B|$, B_j intersects a geometrically once, and there are no other intersections.

As described in Section 1.1, the monodromy of K is given as the composition of Dehn twists $\phi = D_{B_{|\beta_B|-1}}^L \circ \dots \circ D_{B_1}^L \circ D_A^R$, where we first apply D_A^R , then $D_{B_1}^L$, and so on. In this setup, we calculate the nontrivial action of each Dehn twist as: $D_{B_j}^L(a) = \overline{b_{j+1}} b_j a$, and

$$\text{For } 1 \leq l \leq |\beta_B| - 1 : D_{B_l}^L(b_l) = b_{l+1}, \quad D_{B_l}^L(b_{l+1}) = b_{l+1} \overline{b_l} b_{l+1}$$

$$\text{For } 1 \leq l \leq j : D_A^R(b_l) = b_l a$$

We construct our invariant train tracks for five different cases. In each case, we assume that $|\beta_B|, j + 1$ are even so that $K(j, |\beta_B| - j)$ is a knot. For Case 1, we additionally assume that $|\beta_B| - j = j = 1$. For Case 2, we instead assume that $|\beta_B| = j + 1$ and $j \geq 3$. For Case

3, which involves the most moves out of all of the cases, we instead assume that $|\beta_B| \geq j + 5$ and $j \geq 5$. From the invariant train track for Case 3, we recover invariant tracks for Case 4, where we instead assume that $j \geq 5$ and $|\beta_B| = j + 3$, and for Case 5, where we instead assume that $|\beta_B| = j + 3$ and $j = 3$. Throughout our calculations, we only draw local pictures of our graphs around the vertices.

Case 1

Assume that $|\beta_B| - j = j = 1$. The Bestvina-Handel algorithm is illustrated in Figure 4.7, where we begin on the left with an initial spine graph. The initial image of the edges under ϕ is given by, after we pull tight,

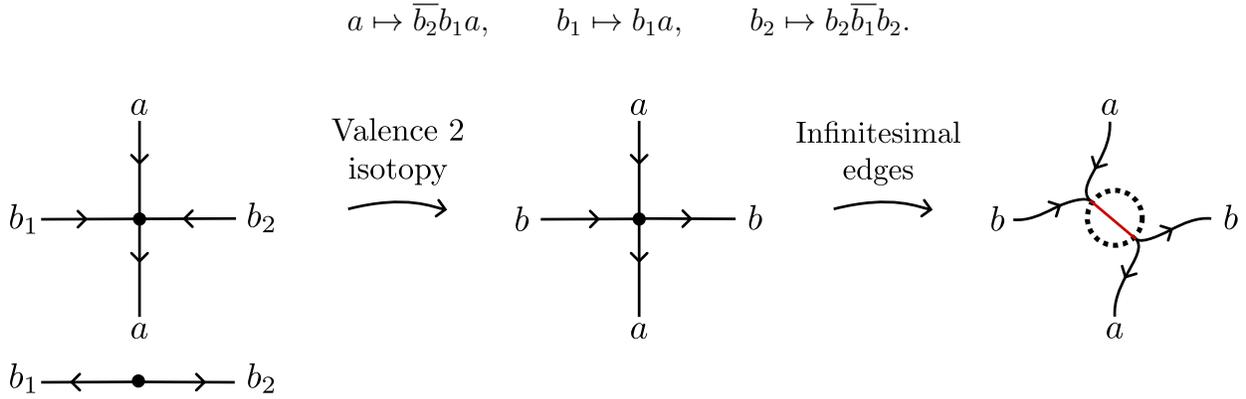


Figure 4.7. A sequence of moves to produce an invariant train track for Case 1.

At this point, we have a valence 2 vertex with link $\{b_1, b_2\}$, so we perform a valence 2 isotopy to remove b_2 . After dropping the subscript on b_1 , we obtain $a \mapsto ba$, $b \mapsto bba$. We claim that this specifies an invariant train track. To verify this claim, we must calculate the gates occurring at the vertex specified by $\{a, \bar{a}, b, \bar{b}\}$. To do so, we must iterate the graph derivative D at this vertex until the image becomes periodic. The image of an edge emanating from a vertex under D is the first edge in the edge-path of the emanating edge. Once the image of D at a vertex becomes periodic, we classify two edges as being in the same gate if their image under this periodic iterate of D is the same, and as being in distinct gates otherwise. We calculate

$$\{a, \bar{a}, b, \bar{b}\} \xrightarrow{D} \{b, \bar{a}, b, \bar{a}\} \mapsto \{b, \bar{a}, b, \bar{a}\}$$

From this, we conclude that a, b are in the same gate, and \bar{a}, \bar{b} are in the same gate. Each of these two gates corresponds to a cusp. Now we insert the infinitesimal edges within our vertex. The edge-path $a \mapsto ba$ connects the distinct gates $\{\bar{b}, a\}$, so we insert an infinitesimal edge connecting these two gates. See the right of Figure 4.7 for the result of distinguishing the gates and adding the infinitesimal edge (in red). The result is an invariant train track where the component containing the puncture is a disk with two cusps.

Case 2

Assume that $|\beta_B| = j + 1$ and $j \geq 3$. The Bestvina-Handel algorithm is illustrated in Figure 4.8, where we begin on the left with an initial spine graph. The initial image of the edges under ϕ is given by, after we pull tight,

$$a \mapsto \overline{b_{|\beta_B|}} b_{|\beta_B|-1} a$$

$$b_l \mapsto \begin{cases} b_{|\beta_B|-1} a, & l = 1 \\ b_{|\beta_B|} \overline{b_{l-1}} b_{|\beta_B|-1} a, & 2 \leq l \leq |\beta_B| - 1 \\ b_{|\beta_B|} \overline{b_{|\beta_B|-1}} b_{|\beta_B|}, & l = |\beta_B| \end{cases}$$

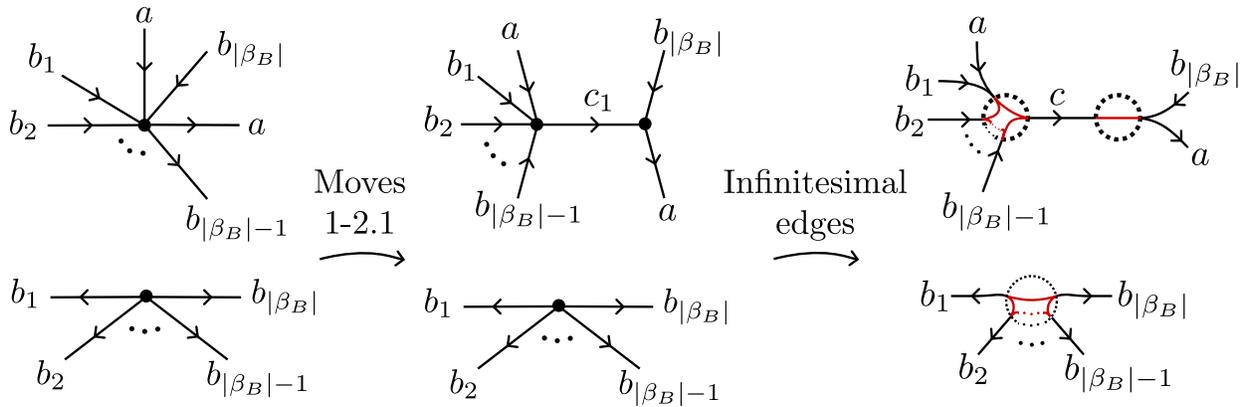


Figure 4.8. A sequence of moves to produce an invariant train track for Case 2.

Move 1

Fold the start of $b_2, \dots, b_{|\beta_B|}$ to $c_1 \mapsto b_{|\beta_B|}$, so that the new edges $b'_2, \dots, b'_{|\beta_B|}$ are specified by

$$b_2 = c_1 b'_2, \dots, b_{|\beta_B|} = c_1 b'_{|\beta_B|}.$$

This folding leaves a valence 2 vertex with link $\{b_1, c_1\}$, so we also perform a valence 2 isotopy to remove c_1 . This move does not change the appearance of the graph. After pulling tight and dropping primes, we obtain

$$\begin{aligned}
a &\mapsto \overline{b_{|\beta_B|}} b_{|\beta_B|-1} a \\
b_l &\mapsto \begin{cases} \overline{b_{|\beta_B|}} b_{|\beta_B|-1} a, & l = 1 \\ \overline{b_{l-1}} b_{|\beta_B|-1} a, & 2 \leq l \leq |\beta_B| - 1 \\ \overline{b_{|\beta_B|-1}} b_{|\beta_B|}, & l = |\beta_B| \end{cases}
\end{aligned}$$

Move 2

We now perform a sequence of folds to eliminate the backtracking

$$b_{|\beta_B|} \mapsto \overline{b_{|\beta_B|-1}} b_{|\beta_B|} \mapsto \cdots b_{|\beta_B|-2} \overline{b_{|\beta_B|-1}} \cdots \mapsto \cdots a \bar{a} \cdots .$$

First, we fold the end of $a, b_1, \dots, b_{|\beta_B|-1}$ to $c_1 \mapsto b_{|\beta_B|-1} a$, so that the new edges are specified by

$$a = a' c_1, b_1 = b'_1 c_1, \dots, b_{|\beta_B|-1} = b'_{|\beta_B|-1} c_1.$$

After pulling tight and dropping primes, we obtain

$$\begin{aligned}
a &\mapsto \overline{b_{|\beta_B|}} \\
b_l &\mapsto \begin{cases} \overline{b_{|\beta_B|}}, & l = 1 \\ \overline{c_1 b_{l-1}}, & 2 \leq l \leq |\beta_B| - 1 \\ \overline{c_1 b_{|\beta_B|-1}} b_{|\beta_B|}, & l = |\beta_B| \end{cases} \\
c_1 &\mapsto b_{|\beta_B|-1} c_1 a c_1
\end{aligned}$$

Move 2.1

As a result of Move 2, one more move is needed to remove the given backtracking since

$$b_{|\beta_B|} \mapsto \overline{b_{|\beta_B|-1}} b_{|\beta_B|} \mapsto \cdots c_1 \bar{c}_1 \cdots .$$

Next, we fold the start of $b_2, \dots, b_{|\beta_B|}$ to $c_2 \mapsto \bar{c}_1$, so that the new edges are specified by

$$b_2 = c_2 b'_2, \dots, b_{|\beta_B|} = c_2 b'_{|\beta_B|}.$$

This folding leaves a valence 2 vertex with link $\{b_1, c_2\}$, so we also perform a valence 2 isotopy to remove c_2 . This move does not change the appearance of the graph. After pulling tight, dropping primes, and dropping the subscript on c_1 , we obtain

$$\begin{aligned} a &\mapsto \overline{b_{|\beta_B|}} \\ b_l &\mapsto \begin{cases} \overline{c b_{|\beta_B|}}, & l = 1 \\ \overline{b_{l-1}}, & 2 \leq l \leq |\beta_B| - 1 \\ \overline{b_{|\beta_B|-1} b_{|\beta_B|}}, & l = |\beta_B| \end{cases} \\ c &\mapsto b_{|\beta_B|-1} c a c \end{aligned}$$

We claim that this specifies an invariant train track. To verify this claim, we must calculate the gates occurring at the three vertices $\{\bar{c}, a, \overline{b_{|\beta_B|}}\}$, $\{\bar{a}, \bar{b}_1, \dots, \overline{b_{|\beta_B|-1}}\}$, $\{b_1, \dots, b_{|\beta_B|}\}$. We carry out this calculation for the vertex $\{\bar{c}, a, \overline{b_{|\beta_B|}}\}$:

$$\{\bar{c}, a, \overline{b_{|\beta_B|}}\} \xrightarrow{D} \{\bar{c}, \overline{b_{|\beta_B|}}, \overline{b_{|\beta_B|}}\} \mapsto \{\bar{c}, \overline{b_{|\beta_B|}}, \overline{b_{|\beta_B|}}\}$$

From this, we conclude that $a, \overline{b_{|\beta_B|}}$ are in the same gate, which corresponds to a cusp, while \bar{c} is in its own gate. Similarly, the emanating vertices at the other vertices are all in distinct gates except for the gate $\{\bar{a}, \bar{b}_1\}$, which corresponds to another cusp.

Now we insert the infinitesimal edges within each vertex between some number of distinct gates. For instance, the edge-path $b_1 \mapsto \overline{c b_{|\beta_B|}}$ connects the distinct gates $\{\bar{c}, \overline{b_{|\beta_B|}}\}$, so we insert an infinitesimal edge e connecting these two gates. Furthermore, we calculate the image of e under iterates of ϕ and insert infinitesimal edges between the two gates which $\phi^k(e)$ connects for some $k \in \mathbb{Z}_{\geq 0}$. This is done similarly to how we calculated the gates, where we now iterate the graph derivative D on the gate $\{\bar{c}, \overline{b_{|\beta_B|}}\}$ until it becomes periodic.

We calculate

$$\{\bar{c}, \overline{b_{|\beta_B|}}\} \xrightarrow{D} \{\bar{c}, \overline{b_{|\beta_B|}}\}$$

We conclude that no additional infinitesimal edges are introduced from e . By performing a similar calculation for all other edge-paths, we conclude that within a vertex, each pair of adjacent gates is connected by an infinitesimal edge. See Figure 4.8 for the result of distinguishing the gates and adding infinitesimal edges (in red). The result is an invariant train track where the component containing the puncture is a disk with two cusps. Furthermore, there are two $|\beta_B|$ -prongs which are permuted.

Case 3

Assume that $j \geq 5$ and $|\beta_B| \geq j + 5$. Our initial spine is pictured in Figure 4.9. The image of the edges under ϕ is given by, after we pull tight,

$$a \mapsto \overline{b_{|\beta_B|}} b_j a$$

$$b_l \mapsto \begin{cases} b_j a, & l = 1 \\ b_{|\beta_B|} \overline{b_{l-1}} b_j a, & 2 \leq l \leq j \\ b_{|\beta_B|} \overline{b_{l-1}} b_{|\beta_B|}, & j + 1 \leq l \leq |\beta_B| \end{cases}$$

Move 1

Fold the start of $b_2, \dots, b_{|\beta_B|}$ to $c_1 \mapsto b_{|\beta_B|}$. This leaves a valence 2 vertex with link $\{b_1, c_1\}$, so we also perform a valence 2 isotopy to remove c_1 . This does not change the appearance of the graph. After pulling tight and dropping primes, we obtain

$$a \mapsto \overline{b_{|\beta_B|}} b_j a$$

$$b_l \mapsto \begin{cases} \overline{b_{|\beta_B|}} b_j a, & l = 1 \\ \overline{b_{l-1}} b_j a, & 2 \leq l \leq j \\ \overline{b_{l-1}} b_{|\beta_B|}, & j + 1 \leq l \leq |\beta_B| \end{cases}$$

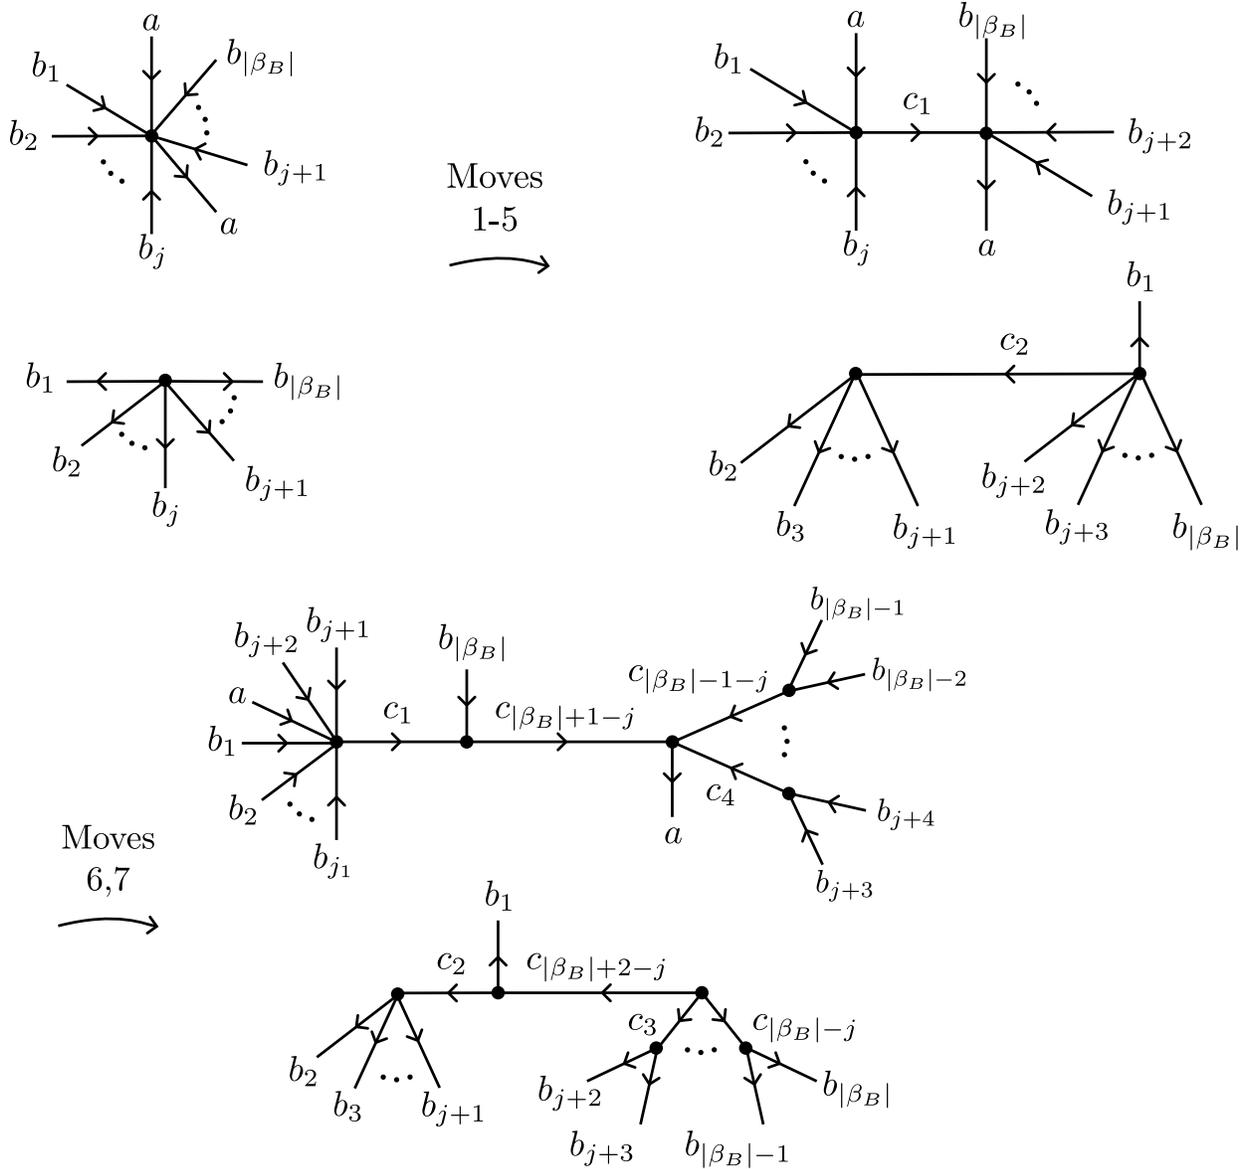


Figure 4.9. The initial spine graph of Case 3, and the results of Moves 1-7.

Move 2

At this point, we would like to eliminate the backtracking

$$b_j \mapsto \overline{b_{j-1}b_j} \cdots \mapsto \cdots b_{j-2}\overline{b_{j-1}} \cdots \mapsto \cdots a\overline{a} \cdots$$

However, this would involve a folding which increases the valence of the point p within b_j where $\phi^3(p)$ fails to be locally injective. To correct this, we subdivide $b_j = b_j^1 b_j^2$, where

$$b_j^1 \mapsto \overline{b_{j-1}}, \quad b_j^2 \mapsto b_j^1 b_j^2 a.$$

Now we can fold the end of $a, b_1, \dots, b_{j-1}, b_j^2$ to $c_1 \mapsto b_j^2 a$, so that the new edges are specified by

$$a = a'c_1, b_1 = b'_1c_1, \dots, b_{j-1} = b'_{j-1}c_1, b_j^2 = (b_j^2)'c_1.$$

After pulling tight and dropping primes, we obtain

$$\begin{aligned} a &\mapsto \overline{b_{|\beta_B|}}b_j^1 \\ b_l &\mapsto \begin{cases} \overline{b_{|\beta_B|}}b_j^1, & l = 1 \\ \overline{c_1 b_{l-1}}b_j^1, & 2 \leq l \leq j-1 \\ \overline{c_1 b_j^2} \overline{b_j^1} b_{|\beta_B|}, & l = j+1 \\ \overline{b_{l-1}}b_{|\beta_B|}, & j+2 \leq l \leq |\beta_B| \end{cases} \\ b_j^1 &\mapsto \overline{c_1 b_{j-1}} \\ b_j^2 &\mapsto b_j^1 \\ c_1 &\mapsto b_j^2 c_1 a c_1 \end{aligned}$$

Move 2.1

As a result of Move 2, one more move is needed to remove the given backtracking since

$$b_j \mapsto \overline{b_{j-1}}b_j^1 \mapsto c_1 \overline{c_1}$$

We fold the start of $b_2, \dots, b_{j-1}, b_j^1, b_{j+1}$ to $c_2 \mapsto \overline{c_1}$, so that the new edges are specified by

$$b_2 = c_2 b'_2, \dots, b_{j-1} = c_2 b'_{j-1}, b_j^1 = c_2 (b_j^1)', b_{j+1} = c_2 b'_{j+1}.$$

After pulling tight and dropping primes, we obtain

$$\begin{aligned}
a &\mapsto \overline{b_{|\beta_B|}} c_2 b_j^1 \\
b_l &\mapsto \begin{cases} \overline{b_{|\beta_B|}} c_2 b_j^1, & l = 1 \\ \overline{b_1} c_2 b_j^1, & l = 2 \\ \overline{b_{l-1}} b_j^1, & 3 \leq l \leq j-1 \\ \overline{b_j^2} \overline{b_j^1} \overline{c_2} b_{|\beta_B|}, & l = j+1 \\ \overline{b_{j+1}} \overline{c_2} b_{|\beta_B|}, & l = j+2 \\ \overline{b_{l-1}} b_{|\beta_B|}, & j+3 \leq l \leq |\beta_B| \end{cases} \\
b_j^1 &\mapsto \overline{b_{j-1}} \overline{c_2} \\
b_j^2 &\mapsto c_2 b_j^1 \\
c_1 &\mapsto b_j^2 c_1 a c_1 \\
c_2 &\mapsto \overline{c_1}
\end{aligned}$$

Move 3

At this point, we have a valence 2 vertex with link $\{\overline{b_j^1}, b_j^2\}$, so we perform a valence 2 isotopy to remove b_j^1 . After pulling tight, dropping primes, and dropping the superscript on b_j^2 , we have

$$\begin{aligned}
a &\mapsto \overline{b_{|\beta_B|}c_2} \\
b_l &\mapsto \begin{cases} \overline{b_{|\beta_B|}c_2}, & l = 1 \\ \overline{b_1c_2}, & l = 2 \\ \overline{b_{l-1}}, & 3 \leq l \leq j \\ \overline{b_j\overline{c_2}b_{|\beta_B|}}, & l = j + 1 \\ \overline{b_{j+1}\overline{c_2}b_{|\beta_B|}}, & l = j + 2 \\ \overline{b_{l-1}b_{|\beta_B|}}, & j + 3 \leq l \leq |\beta_B| \end{cases} \\
c_1 &\mapsto b_jc_1ac_1 \\
c_2 &\mapsto \overline{c_1}
\end{aligned}$$

Move 4

At this point, we would like to eliminate the backtracking

$$b_{|\beta_B|} \mapsto \overline{b_{|\beta_B|-1}b_{|\beta_B|}} \mapsto \cdots b_{|\beta_B|-2}\overline{b_{|\beta_B|-1}} \cdots \mapsto \cdots b_{|\beta_B|}\overline{b_{|\beta_B|}} \cdots$$

However, this would involve a folding which increases the valence of the point p within $b_{|\beta_B|}$ where $\phi^3(p)$ fails to be locally injective. To correct this, we subdivide $b_{|\beta_B|} = b_{|\beta_B|}^1 b_{|\beta_B|}^2$, where

$$b_{|\beta_B|}^1 \mapsto \overline{b_{|\beta_B|-1}}, \quad b_{|\beta_B|}^2 \mapsto b_{|\beta_B|}^1 b_{|\beta_B|}^2.$$

Now we can fold the start of a and the end of $b_{j+1}, \dots, b_{|\beta_B|-1}, b_{|\beta_B|}^2$ to $c_3 \mapsto \overline{b_{|\beta_B|}^2}$, so that the new edges are specified by

$$a = c_3 a', \quad b_{j+1} = b'_{j+1} \overline{c_3}, \quad \dots, \quad b_{|\beta_B|-1} = b'_{|\beta_B|-1} \overline{c_3}, \quad b_{|\beta_B|}^2 = (b_{|\beta_B|}^2)' \overline{c_3}.$$

Together, Moves 4, 4.1, 5 do not change the appearance of the graph. After pulling tight and dropping primes, we obtain

$$\begin{aligned}
a &\mapsto \overline{b_{|\beta_B|}^1} c_2 \\
b_l &\mapsto \begin{cases} c_3 \overline{b_{|\beta_B|}^2} \overline{b_{|\beta_B|}^1} c_2, & l = 1 \\ \overline{b_1} c_2, & l = 2 \\ \overline{b_{l-1}}, & 3 \leq l \leq j \\ \overline{b_j} \overline{c_2} b_{|\beta_B|}^1, & l = j + 1 \\ c_3 \overline{b_{j+1}} \overline{c_2} b_{|\beta_B|}^1, & l = j + 2 \\ c_3 \overline{b_{l-1}} b_{|\beta_B|}^1, & j + 3 \leq l \leq |\beta_B| - 1 \end{cases} \\
b_{|\beta_B|}^1 &\mapsto c_3 \overline{b_{|\beta_B|-1}} \\
b_{|\beta_B|}^2 &\mapsto b_{|\beta_B|}^1 \\
c_1 &\mapsto b_j c_1 c_3 a c_1 \\
c_2 &\mapsto \overline{c_1} \\
c_3 &\mapsto c_3 \overline{b_{|\beta_B|}^2}
\end{aligned}$$

Move 4.1

As a result of Move 4, one more move is needed to remove the given backtracking since

$$b_{|\beta_B|} \mapsto \overline{b_{|\beta_B|-1}} b_{|\beta_B|}^1 \mapsto \overline{c_3} c_3$$

Now we can fold the start of $b_1, b_{j+2}, \dots, b_{|\beta_B|-1}, b_{|\beta_B|}^1$ to $c_4 \mapsto c_3$ so that the new edges are specified by

$$b_1 = c_4 b'_1, b_{j+2} = c_4 b'_{j+2}, \dots, b_{|\beta_B|-1} = c_4 b'_{|\beta_B|-1}, b_{|\beta_B|}^1 = c_4 (b_{|\beta_B|}^1)'$$

After pulling tight and dropping primes, we obtain

$$\begin{aligned}
a &\mapsto \overline{b_{|\beta_B|}^1 \overline{c_4} c_2} \\
b_l &\mapsto \begin{cases} \overline{b_{|\beta_B|}^2 \overline{b_{|\beta_B|}^1 \overline{c_4} c_2}}, & l = 1 \\ \overline{b_1 \overline{c_4} c_2}, & l = 2 \\ \overline{b_{l-1}}, & 3 \leq l \leq j \\ \overline{b_j \overline{c_2} c_4 b_{|\beta_B|}^1}, & l = j + 1 \\ \overline{b_{j+1} \overline{c_2} c_4 b_{|\beta_B|}^1}, & l = j + 2 \\ \overline{b_{l-1} b_{|\beta_B|}^1}, & j + 3 \leq l \leq |\beta_B| - 1 \end{cases} \\
b_{|\beta_B|}^1 &\mapsto \overline{b_{|\beta_B|-1} \overline{c_4}} \\
b_{|\beta_B|}^2 &\mapsto c_4 b_{|\beta_B|}^1 \\
c_1 &\mapsto b_j c_1 c_3 a c_1 \\
c_2 &\mapsto \overline{c_1} \\
c_3 &\mapsto c_3 \overline{b_{|\beta_B|}^2} \\
c_4 &\mapsto c_3
\end{aligned}$$

Move 5

At this point, we have three valence 2 vertices with respective links $\{\overline{b_{|\beta_B|}^1}, b_{|\beta_B|}^2\}$, $\{\overline{c_1}, c_3\}$, $\{c_2, c_4\}$. So we perform three valence 2 isotopies to remove $b_{|\beta_B|}^1, c_3, c_4$. After pulling tight and dropping the superscript on $b_{|\beta_B|}^2$, we obtain

$$\begin{aligned}
a &\mapsto c_2 \\
b_l &\mapsto \begin{cases} \overline{b_{|\beta_B|}c_2}, & l = 1 \\ \overline{b_1c_2}, & l = 2 \\ \overline{b_{l-1}}, & 3 \leq l \leq j \\ \overline{b_jc_2} & l = j + 1 \\ \overline{b_{j+1}c_2} & l = j + 2 \\ \overline{b_{l-1}}, & j + 3 \leq l \leq |\beta_B| \end{cases} \\
c_1 &\mapsto b_jc_1ac_1\overline{b_{|\beta_B|}} \\
c_2 &\mapsto \overline{c_1}
\end{aligned}$$

Move 6

We would like to remove the backtracking

$$\begin{aligned}
c_1 &\mapsto \cdots c_1\overline{b_{|\beta_B|}} \mapsto \cdots \overline{b_{|\beta_B|}b_{|\beta_B|-1}} \cdots \\
&\xrightarrow{\phi^{|\beta_B|-j-3}} \cdots \overline{b_{j+3}b_{j+2}} \cdots \mapsto \cdots b_{j+2}\overline{b_{j+1}} \cdots \mapsto \cdots \overline{c_2}c_2 \cdots
\end{aligned}$$

through a sequence of pairs of folds. We begin by sliding the end of b_{j+1}, b_{j+2} over a , and then we fold the start of b_{j+2}, b_{j+3} to $c_3 \mapsto a$. Then for even k satisfying $j+3 \leq k \leq |\beta_B|-2$, perform the following pairs of folds in increasing k -order. First, fold the end of b_k, b_{k+1} to $c_{k+1-j} \mapsto \overline{c_{k-j}}$, then fold the start of b_{k+1}, b_{k+2} to $c_{k+2-j} \mapsto \overline{c_{k+1-j}}$. After performing this sequence of folds, we then fold the end of $c_1, b_{|\beta_B|}$ to $c_{|\beta_B|+1-j} \mapsto \overline{c_{|\beta_B|-j}}$. From all of this folding, we obtain

$$\begin{aligned}
a &\mapsto c_2 \\
b_l &\mapsto \begin{cases} \overline{c_{|\beta_B|+1-j} b_{|\beta_B|} \overline{c_{|\beta_B|-j} c_2}}, & l = 1 \\ \overline{b_1 c_2}, & l = 2 \\ \overline{b_{l-1}}, & 3 \leq l \leq |\beta_B| \end{cases} \\
c_l &\mapsto \begin{cases} b_j c_1 c_{|\beta_B|+1-j} a c_1 \overline{b_{|\beta_B|}}, & l = 1 \\ \overline{c_{|\beta_B|+1-j} c_1}, & l = 2 \\ a, & l = 3 \\ \overline{c_{l-1}}, & 4 \leq l \leq |\beta_B| + 1 - j \end{cases}
\end{aligned}$$

Move 7

We now have the backtracking

$$b_2 \mapsto \overline{b_1 c_2} \mapsto \cdots c_{|\beta_B|+1-j} \overline{c_{|\beta_B|+1-j}} \cdots,$$

so we fold the start of b_1, c_2 to $c_{|\beta_B|+2-j} \mapsto \overline{c_{|\beta_B|+1-j}}$. After pulling tight, we have

$$\begin{aligned}
a &\mapsto c_{|\beta_B|+2-j} c_2 \\
b_l &\mapsto \begin{cases} \overline{b_{|\beta_B|} \overline{c_{|\beta_B|-j} c_{|\beta_B|+2-j} c_2}}, & l = 1 \\ \overline{b_1 c_2}, & l = 2 \\ \overline{b_{l-1}}, & 3 \leq l \leq |\beta_B| \end{cases} \\
c_l &\mapsto \begin{cases} b_j c_1 c_{|\beta_B|+1-j} a c_1 \overline{b_{|\beta_B|}}, & l = 1 \\ \overline{c_1}, & l = 2 \\ a, & l = 3 \\ \overline{c_{l-1}}, & 4 \leq l \leq |\beta_B| + 2 - j \end{cases}
\end{aligned}$$

Move 8

We now have the backtracking

$$c_1 \mapsto b_j c_1 \cdots \mapsto \cdots \overline{b_{j-1} b_j} \cdots \xrightarrow{\phi^{j-3}} \cdots \overline{b_2 b_3} \cdots \mapsto \cdots b_1 \overline{b_2} \cdots \mapsto \cdots c_2 \overline{c_2} \cdots .$$

In a way similar to Move 6, we must perform a sequence of pairs of folds. The first such pair of folds is to fold the end of a, b_1, b_2 to $d_1 \mapsto c_2$, and then to fold the start of b_2, b_3 to $d_2 \mapsto \overline{d_1}$. Then for odd k satisfying $3 \leq k \leq j - 2$, perform the following pairs of folds in increasing k -order. First, fold the end of b_k, b_{k+1} to $d_k \mapsto \overline{d_{k-1}}$, then fold the start of b_{k+1}, b_{k+2} to $d_{k+1} \mapsto \overline{d_k}$. After performing this sequence of folds, we then fold the start of c_1 and the end of b_j, b_{j+1} to $d_j \mapsto \overline{d_{j-1}}$. From all of this folding, we obtain

$$\begin{aligned}
 a &\mapsto c_{|\beta_B|+2-j} \\
 b_l &\mapsto \begin{cases} \overline{b_{|\beta_B|} \overline{c_{|\beta_B|-j} c_{|\beta_B|+2-j}},} & l = 1 \\ \overline{b_{l-1}}, & 2 \leq l \leq j \\ \overline{d_j b_{l-1}}, & j+1 \leq l \leq j+2 \\ \overline{b_{l-1}}, & j+3 \leq l \leq |\beta_B| \end{cases} \\
 c_l &\mapsto \begin{cases} b_j c_1 c_{|\beta_B|+1-j} a d_1 \overline{d_j c_1 \overline{b_{|\beta_B|}}}, & l = 1 \\ \overline{c_1 d_j}, & l = 2 \\ a d_1, & l = 3 \\ \overline{c_{l-1}}, & 4 \leq l \leq |\beta_B| + 2 - j \end{cases} \\
 d_l &\mapsto \begin{cases} c_2, & l = 1 \\ \overline{d_{l-1}}, & 2 \leq l \leq j \end{cases}
 \end{aligned}$$

Move 9

We now have the backtracking

$$b_1 \mapsto \overline{b_{|\beta_B|} c_{|\beta_B|-j}} \cdots$$

$$\xrightarrow{\phi^{|\beta_B|-j-3}} \cdots \overline{b_{j+3} c_3} \cdots \mapsto \cdots b_{j+2} \overline{d_1} \cdots \mapsto \cdots \overline{b_{j+1} c_2} \cdots \mapsto \cdots d_j \overline{d_j} \cdots .$$

Again, we must perform a sequence of pairs of folds. The first such pair of folds is to fold the start of b_{j+1} and the end of c_2 to $d_{j+1} \mapsto \overline{d_j}$, and then to fold the end of b_{j+2}, d_1 to $d_{j+2} \mapsto \overline{d_{j+1}}$, and then to fold the end of c_3 and start of b_{j+3} to $d_{j+3} \mapsto \overline{d_{j+2}}$. Then for even k satisfying $4 \leq k \leq |\beta_B| - j - 1$, perform the following pairs of folds in increasing k -order. First, fold the start of c_k and end of b_{j+k} to $d_{j+k} \mapsto \overline{d_{j+k-1}}$, then fold the start of b_{j+k+1} and end of c_{k+1} to $d_{j+k+1} \mapsto \overline{d_{j+k}}$. After performing this sequence of folds, we then have $|\beta_B| - j - 2$ valence 2 vertices, so we perform valence 2 isotopies to remove $d_{j+3}, \dots, d_{|\beta_B|}$. As a result, we obtain

$$\begin{aligned}
a &\mapsto c_{|\beta_B|+2-j} \\
b_l &\mapsto \begin{cases} \overline{b_{|\beta_B|} c_{|\beta_B|-j} c_{|\beta_B|+2-j}}, & l = 1 \\ \overline{b_{l-1}}, & 2 \leq l \leq j+1 \\ d_{j+2} \overline{d_j} \overline{b_{j+1}}, & l = j+2 \\ \overline{b_{l-1}}, & j+3 \leq l \leq |\beta_B| \end{cases} \\
c_l &\mapsto \begin{cases} b_j c_1 c_{|\beta_B|+1-j} a d_1 d_{j+2} \overline{d_j} c_1 \overline{b_{|\beta_B|}}, & l = 1 \\ \overline{c_1}, & l = 2 \\ a d_1, & l = 3 \\ \overline{c_{l-1}}, & 4 \leq l \leq |\beta_B| + 2 - j \end{cases} \\
d_l &\mapsto \begin{cases} c_2, & l = 1 \\ \overline{d_{j+2}} \overline{d_1}, & l = 2 \\ \overline{d_{l-1}}, & 3 \leq l \leq j+2 \end{cases}
\end{aligned}$$

We claim that this specifies an invariant train track. To verify this claim, we must calculate the gates. We give a sample calculation for the gates corresponding to the vertex with emanating edges $\{\overline{a}, \overline{b_1}, \overline{b_2}, d_1\}$. In this example,

$$\{\overline{a}, \overline{b_1}, \overline{b_2}, d_1\} \xrightarrow{D} \{\overline{c_{|\beta_B|+2-j}}, \overline{c_{|\beta_B|+2-j}}, b_1, c_2\}.$$

Since $D(\overline{a}) = D(\overline{b_1})$, the edges $\overline{a}, \overline{b_1}$ emanate from the same gate, and this corresponds to a cusp in the train track. Continuing to iterate D , we obtain:

$$\begin{aligned}
\{\overline{c_{|\beta_B|+2-j}}, \overline{c_{|\beta_B|+2-j}}, b_1, c_2\} &\mapsto \{c_{|\beta_B|+1-j}, c_{|\beta_B|+1-j}, \overline{b_{|\beta_B|}}, \overline{c_1}\} \\
&\mapsto \{\overline{c_{|\beta_B|-j}}, \overline{c_{|\beta_B|-j}}, b_{|\beta_B|-1}, b_{|\beta_B|}\} \\
&\xrightarrow{D^{|\beta_B|-j-3}} \{\overline{c_3}, \overline{c_3}, b_{j+2}, b_{j+3}\} \\
&\mapsto \{\overline{d_1}, \overline{d_1}, d_{j+2}, \overline{b_{j+2}}\} \\
&\mapsto \{\overline{c_2}, \overline{c_2}, \overline{d_{j+1}}, b_{j+1}\} \\
&\mapsto \{c_1, c_1, d_j, \overline{b_j}\} \\
&\mapsto \{b_j, b_j, \overline{d_{j-1}}, b_{j-1}\} \\
&\xrightarrow{D^{j-3}} \{b_3, b_3, \overline{d_2}, b_2\} \\
&\mapsto \{\overline{b_2}, \overline{b_2}, d_1, \overline{b_1}\} \\
&\mapsto \{b_1, b_1, c_2, \overline{c_{|\beta_B|+2-j}}\}
\end{aligned}$$

By repeating this iteration of D three more times, we return to $\{\overline{c_{|\beta_B|+2-j}}, \overline{c_{|\beta_B|+2-j}}, b_1, c_2\}$, so the edges $\overline{a}, \overline{b_1}$ lie in the same gate, and the edges $\overline{b_2}, d_1$ lie in distinct gates. By performing a similar calculation at all vertices, we conclude that all other edges are in distinct gates except for the edges $\{c_1, \overline{b_{j+1}}\}$, which correspond to another cusp.

Finally, we must calculate the infinitesimal edges. We produce a sample calculation corresponding to the infinitesimal edge which is added from the image of $c_1 \mapsto b_j c_1$. This edge-path connects the distinct gates $\{\overline{b_j}, c_1\}$, and its image under iterations of D is specified by

$$\begin{aligned}
\{\overline{b}_j, c_1\} &\mapsto \{b_{j-1}, b_j\} \\
&\mapsto \{\overline{b}_{j-2}, \overline{b}_{j-1}\} \\
&\xrightarrow{D^{j-3}} \{\overline{b}_1, \overline{b}_2\} \\
&\mapsto \{\overline{c}_{|\beta_B|+2-j}, b_1\} \\
&\mapsto \{c_{|\beta_B|+1-j}, \overline{b}_{|\beta_B|}\} \\
&\xrightarrow{D^{|\beta_B|-j-2}} \{\overline{c}_3, b_{j+2}\} \\
&\mapsto \{\overline{d}_1, d_{j+2}\} \\
&\mapsto \{\overline{c}_2, \overline{d}_{j+1}\} \\
&\mapsto \{c_1, d_{j+1}\}
\end{aligned}$$

The final gates in this calculation correspond to the same vertex and a different infinitesimal edge from our starting gates $\{\overline{b}_j, c_1\}$. By repeating this iteration of D two more times, we return to $\{\overline{b}_j, c_1\}$, so we conclude that all vertexes with three gates have adjacent gates connected by an infinitesimal edge. Similarly, the adjacent gates within each of $\{a, \overline{c}_4, \dots, \overline{c}_{|\beta_B|-1-j}, \overline{c}_{|\beta_B|+1-j}\}$, $\{c_3, \dots, c_{|\beta_B|-j}, c_{|\beta_B|+2-j}\}$, $\{\overline{d}_3, \dots, \overline{d}_{j-2}, \overline{d}_j, \overline{d}_{j+2}\}$, $\{d_2, \dots, d_{j-1}, d_{j+1}\}$ are connected by infinitesimal edges. See Figure 4.11 for the result of distinguishing the gates and adding infinitesimal edges. The result is an invariant train track where the component containing the puncture is a disk with two cusps. Furthermore, there are two $\frac{|\beta_B|+1-j}{2}$ -prongs which are permuted (colored blue), there are two $\frac{j+1}{2}$ -prongs which are permuted (colored gray), and there are $(|\beta_B| + 2)$ 3-prongs which are permuted (colored red).

Case 4

Now we assume that $j \geq 5$, $|\beta_B| = j + 3$. Our initial spine graph is pictured in the top left of Figure 4.9, where we relabel $|\beta_B| = j + 3$. The moves we perform are essentially the same as those performed in Case 3, except for several key differences. After Move 4, there are

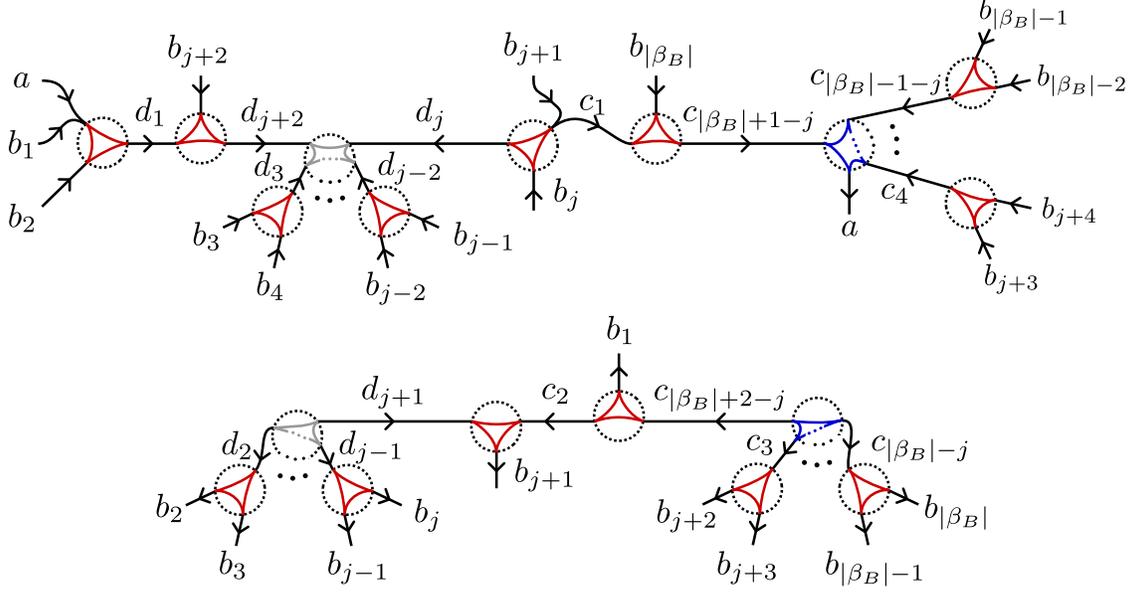


Figure 4.11. Adding infinitesimal edges at vertices to obtain the train track of Case 3.

images of b_l for $j + 3 \leq l \leq |\beta_B| - 1$, which would be empty in our current case. One can check that this does not affect how the images are written after Move 5.

In Move 6, our sequences of folds for $j + 3 \leq k \leq |\beta_B| - 2$ is now empty. Consequently, we delete the following labeled edges from the bottom of Figure 4.9. From the upper graph, we delete the edges labeled $b_{j+3}, \dots, b_{|\beta_B|-1}, c_4, \dots, c_{|\beta_B|-1-j}$, and from the lower graph, we delete the edges labeled $b_{j+4}, \dots, b_{|\beta_B|}, c_5, \dots, c_{|\beta_B|-j}$. Then, we relabel $|\beta_B| = j_1 + 3$. This results in two valence 2 vertices with respective links $\{a, \overline{c_4}\}, \{c_3, c_5\}$. The image of our current case's edges after Move 7 is obtained from the previous case's image by relabeling $|\beta_B| = j + 3$, and by performing two valence 2 isotopies to remove c_4, c_5 .

In Move 9, our sequence of folds for $4 \leq k \leq |\beta_B| - j - 1$ is now empty, but in the same move, we removed all such folds by valence 2 isotopy. Consequently, this has no effect on our current case's graph or the image of edges. See Figure 4.12 for the final graph. The image of our edges is given by

$$\begin{aligned}
a &\mapsto \overline{c_3} \\
b_l &\mapsto \begin{cases} \overline{b_{j+3}c_3}, & l = 1 \\ \overline{b_{l-1}}, & 2 \leq l \leq j+1 \\ d_{j+2}\overline{d_j b_{j+1}}, & l = j+2 \\ \overline{b_{j+2}}, & l = j+3 \end{cases} \\
c_l &\mapsto \begin{cases} b_j c_1 a d_1 d_{j+2} \overline{d_j c_1 b_{j+3}}, & l = 1 \\ \overline{c_1}, & l = 2 \\ a d_1, & l = 3 \end{cases} \\
d_l &\mapsto \begin{cases} c_2, & l = 1 \\ \overline{d_{j+2} d_1}, & l = 2 \\ \overline{d_{l-1}}, & 3 \leq l \leq j+2 \end{cases}
\end{aligned}$$

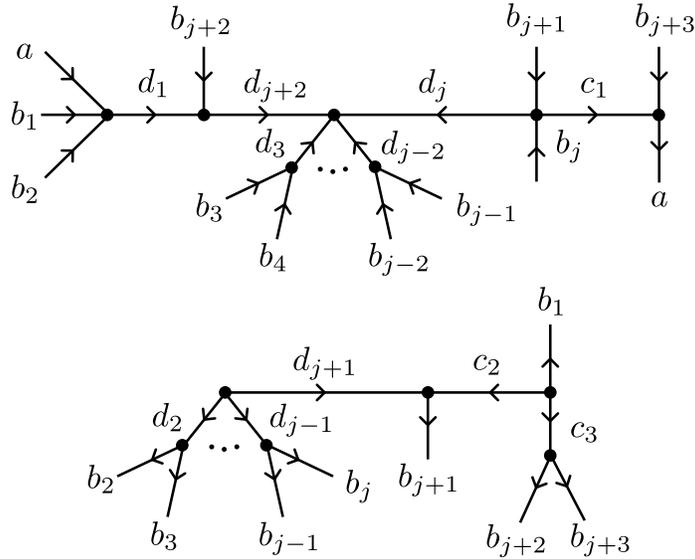


Figure 4.12. The final Case 4 graph.

We claim that this specifies an invariant train track. By calculating the gates in a way similar to the previous case, we conclude that at each vertex, every emanating edge

is in a distinct gate except for the gates $\{\bar{a}, \bar{b}_1\}, \{\bar{b}_{j+1}, c_1\}$ which correspond to cusps. By another infinitesimal edge calculation, we conclude that within each vertex, adjacent gates are connected by an infinitesimal edge. See Figure 4.4. The result is an invariant train track where the component containing the puncture is a disk with two cusps. Furthermore, there are two $\frac{j+1}{2}$ -prongs which are permuted, and there are $(j+5)$ 3-prongs which are permuted.

Case 5

Now we assume that $j = 3$, $|\beta_B| = j + 3$. Similarly to Case 4, we trace through the moves of Case 3 and alter our graph and the image of our edges as required. See Figure 4.13. As a result, we obtain

$$\begin{aligned}
 a &\mapsto \bar{c}_3 \\
 b_l &\mapsto \begin{cases} \bar{b}_6 \bar{c}_3, & l = 1 \\ \bar{b}_{l-1}, & 2 \leq l \leq 4 \\ \bar{d}_3 \bar{b}_4, & l = 5 \\ \bar{b}_5, & l = 6 \end{cases} \\
 c_l &\mapsto \begin{cases} b_3 c_1 a d_1 \bar{d}_3 c_1 \bar{b}_6, & l = 1 \\ \bar{c}_1, & l = 2 \\ a d_1, & l = 3 \end{cases} \\
 d_l &\mapsto \begin{cases} c_2, & l = 1 \\ d_3 \bar{d}_1, & l = 2 \\ \bar{d}_2, & l = 3 \end{cases}
 \end{aligned}$$

We claim that this specifies an invariant train track where the component containing the puncture is a disk with two cusps. Furthermore, there are eight 3-prongs which are permuted. See Figure 4.2.

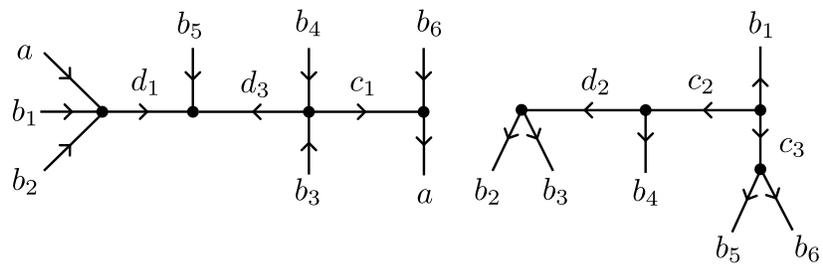


Figure 4.13. The final Case 5 graph.