

Stabilization in the braid groups I: MTWS

JOAN S BIRMAN
WILLIAM W MENASCO

Choose any oriented link type \mathcal{X} and closed braid representatives X_+, X_- of \mathcal{X} , where X_- has minimal braid index among all closed braid representatives of \mathcal{X} . The main result of this paper is a ‘Markov theorem without stabilization’. It asserts that there is a complexity function and a finite set of ‘templates’ such that (possibly after initial complexity-reducing modifications in the choice of X_+ and X_- which replace them with closed braids X'_+, X'_-) there is a sequence of closed braid representatives $X'_+ = X^1 \rightarrow X^2 \rightarrow \dots \rightarrow X^r = X'_-$ such that each passage $X^i \rightarrow X^{i+1}$ is strictly complexity reducing and non-increasing on braid index. The templates which define the passages $X^i \rightarrow X^{i+1}$ include 3 familiar ones, the destabilization, exchange move and flype templates, and in addition, for each braid index $m \geq 4$ a finite set $\mathcal{T}(m)$ of new ones. The number of templates in $\mathcal{T}(m)$ is a non-decreasing function of m . We give examples of members of $\mathcal{T}(m)$, $m \geq 4$, but not a complete listing. There are applications to contact geometry, which will be given in a separate paper [6].

[57M25](#), [57M50](#)

1 Introduction

1.1 The problem

Let \mathcal{X} be an oriented link type in the oriented 3–sphere S^3 or $\mathbb{R}^3 = S^3 \setminus \{\infty\}$. A representative $X \in \mathcal{X}$ is said to be a *closed braid* if there is an unknotted curve $\mathbf{A} \subset (S^3 \setminus X)$ (the *axis*) and a choice of *fibration* \mathbf{H} of the open solid torus $S^3 \setminus \mathbf{A}$ by meridian discs $\{H_\theta; \theta \in [0, 2\pi]\}$, such that whenever X meets a fiber H_θ the intersection is transverse. We call the pair (\mathbf{A}, \mathbf{H}) a *braid structure*. The fact that X is a closed braid with respect to \mathbf{H} implies that the number of points in $X \cap H_\theta$ is independent of θ . We call this number the *braid index* of X , and denote it by the symbol $b(X)$. The *braid index* of \mathcal{X} , denoted $b(\mathcal{X})$, is the minimum value of $b(X)$ over all closed braid representatives $X \in \mathcal{X}$.

Closed braid representations of \mathcal{X} are not unique, and Markov’s well-known theorem (see the book by Birman [3], and the papers by Birman and Menasco [12], Lambropoulou and Rourke [22], Markov [23], Morton [28] and Traczyk [31]) asserts that

any two are related by a finite sequence of elementary moves. One of the moves is *braid isotopy* by which we mean an isotopy of the pair $(X, \mathbb{R}^3 \setminus \mathbf{A})$ which preserves the condition that X is transverse to the fibers of \mathbf{H} . The other two moves are mutually inverse, and are illustrated in Figure 1. Both take closed braids to closed braids. We call them *destabilization* and *stabilization* where the former decreases braid index by one and the latter increases it by one. The ‘weight’ w denotes w parallel strands, relative to the given projection. The braid inside the box which is labeled P is an arbitrary $(w + 1)$ -braid. Later, it will be necessary to distinguish between positive and negative destabilizations, so we illustrate both now.




Figure 1: The two destabilization moves

Theorem 1 (Markov’s Theorem (MT) [23]) *Let X_+, X_- be closed braid representatives of the same oriented link type \mathcal{X} in oriented 3-space, with the same braid axis \mathbf{A} . Then there exists a sequence*

$$(1-1) \quad X_+ = X_1 \rightarrow X_2 \rightarrow \dots \rightarrow X_r = X_-$$

of closed braid representatives of \mathcal{X} such that, up to braid isotopy, each X_{i+1} is obtained from X_i by a single stabilization or destabilization.

It is easy to find examples of subsequences $X_j \rightarrow \dots \rightarrow X_{j+k}$ of (1-1) in Theorem 1 such that $b(X_j) = b(X_{j+k})$, but X_j and X_{j+k} are not braid isotopic. Call such a sequence a *Markov tower*. The stabilization and destabilization moves are very simple, but sequences of stabilizations, braid isotopies and destabilizations can have unexpected consequences. In the braid groups these moves are ‘site dependent’, unlike the stabilization–destabilization move in the Reidemeister–Singer Theorem. (For an example the reader should refer ahead to the specified site of the stabilization in the sequence in Figure 5.) Until now these moves have been predominately used to develop link invariants, but the Markov towers themselves have been ‘black boxes’. One of

the main motivating ideas of this work is to open up the black box and codify Markov towers.

Markov's Theorem is typical of an entire class of theorems in topology where some form of stabilization and destabilization play a central role. Other examples are:

- (1) The Reidemeister–Singer Theorem [30] relates any two Heegaard diagrams of the same 3–manifold, by a finite sequence of very simple elementary changes on Heegaard diagrams. The stabilization–destabilization move adds or deletes a pair of simple closed curves a, b in the defining Heegaard diagram, where $a \cap b = 1$ point and neither a nor b intersects any other curve a_i, b_j in the Heegaard diagram.
- (2) The Kirby Calculus [21] gives a finite number of moves which, when applied repeatedly, suffice to change any surgery presentation of a given 3–manifold into any other, at the same time keeping control of the topological type of a 4–manifold which the given 3–manifold bounds. The stabilization–destabilization move is the addition-deletion of an unknotted component with framing ± 1 to the defining framed link.
- (3) Reidemeister's Theorem (see Burde–Zieschang [15]) relates any two diagrams of the same knot or link, by a finite sequence of elementary moves which are known as RI, RII, RIII. The stabilization–destabilization move is RI. It is easy to see that Markov's Theorem implies Reidemeister's Theorem.

These theorems are all like Markov's Theorem in the sense that while the stabilization and destabilization moves are very simple, a sequence of these moves, combined with the appropriate isotopy, can have very non-trivial consequences. Here are other examples in which the stabilization move is not used, at the expense of restricting attention to a special example:

- (4) W Haken proved that any Heegaard diagram for a non-prime 3–manifold is equivalent to a Heegaard diagram which is the union of two separate Heegaard diagrams, one for each summand, supported on disjoint subsets of the given Heegaard surface. See Scharlemann–Thompson [29] for a very pleasant proof.
- (5) Waldhausen [32] proved that any two Heegaard diagrams of arbitrary but fixed genus g for the 3–sphere S^3 are equivalent.

In the course of an effort which we began in 1990 to discover the theorem which will be the main result of this paper (see [Theorem 2](#) below) the authors made several related contributions to the theory of closed braid representatives of knots and links:

- (4') A *split* (resp. *composite*) closed n -braid is an n -braid which factorizes as a product XY where the sub-braid X involves only strands $1, \dots, k$ and the sub-braid Y involves only strands $k + 1, \dots, n$ (resp. k, \dots, n). In the manuscript [7] the authors proved that if X is a closed n -braid representative of a split or composite link, then up to (braid-index preserving) isotopy and *exchange* moves, as in Figure 2, X may be assumed to be a split or composite closed braid.




Figure 2: The exchange move

- (5') In the manuscript [9] the authors proved that if X is a closed braid representative of the μ -component unlink \mathcal{X} , then a finite sequence of braid isotopies, exchange moves and destabilization can be found which change X to the closure of the identity braid in the braid group B_μ .
- (6) In the manuscript [11] the authors discovered that there is another move, the *3-braid flype* (see Figure 3) with the property that if X is a closed 3-braid representative of a knot or link type \mathcal{X} which cannot be represented by a 1-braid or 2-braid, then either X has a unique conjugacy class or X has exactly two conjugacy classes, and these two classes are related by a 3-braid flype. They also showed that the exchange move can be replaced by braid isotopy for prime links of braid index 3.




Figure 3: The two flype moves

The authors also established two fundamental facts which gave strong evidence that a more general result might be true:

- (7) In [8] the authors introduced a complexity function on closed braid representatives of \mathcal{X} and proved that, up to exchange moves, there are at most finitely many conjugacy classes of representatives of minimum complexity.
- (8) In [10] the authors proved that if a link type \mathcal{X} has infinitely many conjugacy classes of closed braid representatives of the same braid index, then, up to exchange moves, they fall into finitely many equivalence classes.

The goal of this paper is to generalize examples (4'), (5') and (6), taking into account (7) and (8), to arbitrary closed braid representatives of arbitrary oriented knots and links. We call our main theorem *Markov's Theorem Without Stabilization* (MTWS), because it is a direct modification of Markov's Theorem, but with his stabilization move replaced by other moves which allow one to jump from one isotopy class in the complement of \mathbf{A} to another, while keeping the braid index constant or decreasing it.

1.2 Block-strand diagrams and templates

Before we can state our main result, we need to introduce new concepts. Our moves will be described in terms of certain pairs of 'block-strand diagrams' which we call 'templates'. Examples are the block-strand diagram pairs which make up the templates in Figures 1, 2 and 3. The reader may wish to look ahead to the boxed pairs of block-strand diagrams in Figures 8, 9 and 10 for examples of more complicated templates.

A *block* B in $\mathbb{R}^3 - \mathbf{A}$ is a 3-ball having the structure of a 2-disk Δ crossed with an interval $[0, 1]$ such that (i) for any fiber $H_\theta \in \mathbf{H}$ the intersection $H_\theta \cap B$ is either \emptyset or $\Delta \times \{\theta\}$ for some $\theta \in [\theta_1, \theta_2]$, and (ii) there exists some $\theta \in (0, 2\pi)$ such that $H_\theta \cap B = \emptyset$. The disc $\mathbf{t} = B \cap H_{\theta_1}$ is the *top* of B and the disc $\mathbf{b} = B \cap H_{\theta_2}$ is the *bottom* of B . A *strand* l is homeomorphic to an interval $[0, 1]$ or a circle S^1 . It is oriented and transverse to each fiber of \mathbf{H} such that its orientation agrees with the forward direction of \mathbf{H} . When l is homeomorphic to an interval, $\partial l = l_0 \cup l_1$, where l_0 is the *beginning endpoint* of l and l_1 is the *ending endpoint* of l . A *block-strand diagram* \mathcal{D} is a collection of pairwise disjoint blocks $\{B^1, \dots, B^k\}$ and pairwise disjoint strands $\{l^1, \dots, l^1\}$ which together have the following structure:

- (1) If $l^i \cap B^j \neq \emptyset$ then $l^i \cap B^j = (l^i_1 \cap \mathbf{t}^j) \cup (l^i_0 \cap \mathbf{b}^j)$ where \mathbf{t}^j is the top of B^j and \mathbf{b}^j is the bottom of B^j . (We allow for the possibility that either $(l^i_1 \cap \mathbf{t}^j)$ or $(l^i_0 \cap \mathbf{b}^j)$ is empty.)
- (2) For each l^i_0 (resp. l^i_1) there is some $\mathbf{b}^j \subset B^j$ (resp. $\mathbf{t}^j \subset B^j$) such that $l^i_0 \subset \mathbf{b}^j$ (resp. $l^i_1 \subset \mathbf{t}^j$).
- (3) For each block B^j we have $|\mathbf{t}^j \cap (\cup_{1 \leq i \leq 1} l^i_1)| = |\mathbf{b}^j \cap (\cup_{1 \leq i \leq 1} l^i_0)| \geq 2$.

The fact that for each $j = 1, \dots, k$ there is a fiber which misses B^j shows that, by rescaling, we may find a distinguished fiber H_{θ_0} which does not intersect any block. We define the *braid index* $b(\mathcal{D})$ of the block-strand diagram \mathcal{D} to be the number of times the strands of \mathcal{D} intersect the distinguished fiber H_{θ_0} . Condition (3) above makes $b(\mathcal{D})$ well defined. For a specified block $B^j \subset \mathcal{D}$ we define its *braid index* $b(B^j) = |\mathfrak{t}^j \cap (\cup_{1 \leq i \leq l_1^j})|$. Continuing the definition of a block-strand diagram, we assume:

- (4) If $B^j \subset \mathcal{D}$ then the braid index $b(B^j)$ is strictly less than the braid index $b(\mathcal{D})$, with one exception. The exception occurs when \mathcal{D} is the block strand diagram that results after a destabilization has been performed, as in [Figure 1](#).

A *template* \mathcal{T} is a pair of block-strand diagrams $(\mathcal{D}_+, \mathcal{D}_-)$, both with blocks B^1, \dots, B^k and an isotopy which takes \mathcal{D}_+ to \mathcal{D}_- , in such a way that for every fixed choice of braiding assignments to the blocks B^1, \dots, B^k the resulting closed braids X, X' represent the same oriented link type \mathcal{X} . The diagrams \mathcal{D}_+ and \mathcal{D}_- are the *initial* and *final* block-strand diagrams in the pair. The *fixed blocks* and *fixed strands* in $\mathcal{T} = (\mathcal{D}_+, \mathcal{D}_-)$ are the blocks and strands where the isotopy is pointwise the identity. All other blocks and strands are *moving*. For example, in [Figure 3](#) the blocks P and Q are fixed blocks, whereas R is a moving block. A *braiding assignment* to a block-strand diagram \mathcal{D} is a choice of a braid on m_j strands for each $B^j \in \mathcal{D}$. That is, we replace B^j with the chosen braid, so that B^j with this braiding assignment becomes a braided tangle with m_j in-strands and m_j out-strands. In this way a block strand diagram gives us a closed braid representative of a link \mathcal{X} .

Let X be a closed m -braid. We say that X is *carried* by \mathcal{D} if there exists a braiding assignment for the blocks in \mathcal{D} such that the resulting closed braid is braid-isotopic to X .

When we first began to understand that templates were the appropriate settings for our work on the MTWS we wondered whether our definition was so broad (because the diagrams in question support so many knot and link types) as to be content-free! In this regard, the following fact is fundamental:

Proposition 1.2.1 *Let \mathcal{D} be a block-strand diagram of braid index n . Then there exist n -braids that \mathcal{D} does not carry.*

Proof Up to conjugation, a block-strand diagram may be described by a word $V_1 W_1 V_2 W_2 \dots V_k W_k$ in the standard elementary braid generators $\sigma_1, \dots, \sigma_{n-1}$ of the n -strand braid group, where each V_j represents a word which describes the braid carried by the j^{th} block (after making a braiding assignment to the block) and each W_j

is a braid word on n strands which describes the strands that connect the blocks. By hypothesis no block has more than $n - 1$ strands entering or leaving it, so by modifying the W_i 's we may assume without loss of generality that each V_j is a braid on the first q_j -strands, where $q_j < n$. After this modification, the only places where the elementary braid generator σ_{n-1} appears is in the braid words that describe the strands that join the blocks, ie the words W_1, \dots, W_k .

Now let $|W_i|$ be the number of times σ_{n-1} occurs in W_i . The $|W_i|$'s are fixed numbers since we were handed a block-strand diagram. For an arbitrary conjugacy class $\{C\}$ of n -braids, let $|C|$ be the minimum number of times the generator σ_{n-1} is used, in all possible words which represent $\{C\}$. Our block diagram can only carry closed n -braids C such that (up to conjugacy) $|C| \leq |W_1| + |W_2| + \dots + |W_k|$. But there are closed n -braids X such that $|Y|$ is arbitrarily large for all $Y \in \{X\}$. An example is $X = (\sigma_1 \sigma_2 \dots \sigma_{n-1})^{nN}$ for a large positive integer N . For, the braid word $(\sigma_1 \sigma_2 \dots \sigma_{n-1})^n$ is a full twist of the braid strands. It generates the center of the n -string braid group. It cannot be represented by any braid word that does not use all the elementary braid generators. Therefore it's N^{th} power uses the generator σ_{n-1} at least nN times. \square

Having [Proposition 1.2.1](#) on hand, we proceed to define the templates that we will use in the statement of the MTWS. Our main theorem begins with an arbitrary closed n -braid representative X_+ of an arbitrary oriented knot or link type \mathcal{X} in 3-space. Let X_- be a second such representative, where $b(X_-) = b(\mathcal{X})$. Our goal in the subsections which follow is to describe some of the templates that we need, and at the same time to describe the building blocks of all of them. Note that we regard braid isotopy as a trivial move, sometimes even forgetting to mention it. By a result of Morton [[26](#), [Theorem 1](#)] braids β, β' in the n -string braid group B_n are conjugate if and only if the associated closed braids are isotopic in the complement of the braid axis. In keeping with our motivating idea of codifying Markov towers, the names that we give some of our templates $\mathcal{T} = (D_+, D_-)$ correspond to the name of the isotopy that is used to move D_+ to D_- .

1.2.1 The two destabilization templates Our two *destabilization* templates were defined in [Figure 1](#). We distinguish the cases of positive and negative destabilization because the strands which join the fixed blocks are different, and so the templates are different. The destabilization templates do not have any moving blocks. They occur at every braid index.

1.2.2 The admissible flype templates Flypes first enter the picture when the braid index is 3, and we already illustrated the two 3-braid flype templates in [Figure 3](#). There

is an obvious way to generalize it to any braid index n , namely declare the strands to be weighted strands. See Figure 4(a), which shows the support of the flype with weighted strands. From now on, the term *flype* will always have this meaning. The *sign of a flype* is the sign of the single crossing (possibly weighted) which is not in the braid block. Both positive and negative flypes are illustrated in Figure 3. They have distinct templates.




Figure 4: (a) The support of a positive flype with weighted strands. (b) Example of an inadmissible flype. Notice the extra twists introduced because of the weighted strands.

There is a subtle point: Let X_+ and X_- be the closed braid before and after a flype, which we shall consider (for the purpose of describing our moves) as acting left to right. The flype motion is supported in a 3–ball B^3 . In Figure 4(a) observe that the fiber H_θ at $\theta = \pi/2$ intersects $X_+ \cap B^3$ in w' points, but intersects X_- in k points. Observe that $w + w' = k + k'$. We have shown that

$$b(X_+) - b(X_-) = w' - k = k' - w.$$

Thus flypes with weighted strands are non-increasing on braid index if and only if $w' - k = k' - w \geq 0$. We will refer to a flype which is non-increasing on braid index as an *admissible flype*. An example of an inadmissible flype is given in Figure 4(b). While we are obviously interested in the admissible flypes, it will turn out that the *inadmissible flypes* are important too, as they lead to additional, more complicated templates.

By Markov’s Theorem, the left and right braids in every admissible flype template must be related by a Markov tower. Figure 5 shows such a 2–step tower, in the case when the braid index is 3. The moves used in the sequence are (up to braid isotopy) a single stabilization and a single destabilization. Thus flypes arise in a very natural way in the study of stabilization in the braid groups: they replace a sequence stabilization,

destabilization by a single braid-index preserving (or possibly reducing) move. Notice that when R is a negative half-twist, the tower can be replaced by an exchange move. Observe that since flypes are replacements for Markov towers, we are now free to use

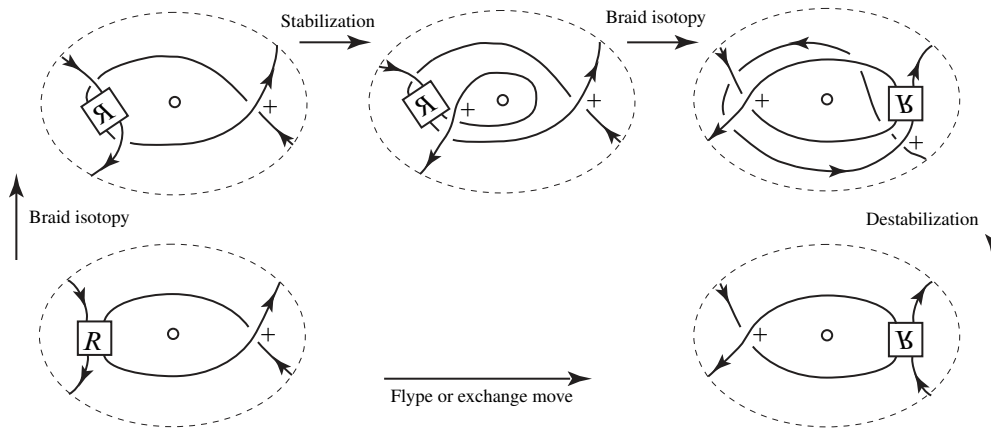


Figure 5: A very simple Markov tower

them to construct more general Markov towers.

1.2.3 The exchange move template and sequences of exchange moves The exchange move template was defined by the block-strand diagram in Figure 2. It was proved in our earlier paper [11] that for $n \leq 3$ it is equivalent to braid isotopy, and by Fiedler [18] that for $n \geq 4$, and generic choices of the braids P and Q , the exchange move cannot be replaced by braid isotopy.

Figure 6 shows how exchange moves, together with braid isotopy, can lead to infinitely many conjugacy classes of closed braid representatives of the same knot or link (see our earlier paper [9]). Indeed, in [10] the authors proved that if a link has infinitely many conjugacy classes of closed m -braid representatives for any fixed value of m then all but finitely many of them are related by exchange moves. This fact will shape the form of our main theorem. More precisely, our main theorem shows exactly how far one may go, using only exchange moves and destabilizations, and then identifies the finitely many moves which are needed *in addition to* exchange moves and destabilizations, to take one closed n -braid representative of a knot to another of the same braid index.

The sequences of exchange moves which we next define are very useful and important. Figure 7 shows how we unfold a piece of a closed braid to reveal that it has the structure of a ‘rooted block and strand tree’, and then ‘loop’ a distinguished subarc of the braid (always of weight 1) over the tree. (These concepts will be defined in Section 5.6. See




Figure 6: The exchange move can lead to arbitrarily many distinct braid isotopy classes of closed n -braid representatives of a single knot type, $n \geq 4$.




Figure 7: A sequence of exchange moves passes a distinguished strand (the thick black one) over a block-strand tree

Figure 59 and the nearby text. We hope the reader will be patient. Our initial goal is to state our main result.) Notice that, while the unlooping process does not preserve closed braids, we have retained the closed braid structure by the device of cutting the braid axis into 4 little ‘axis pieces’. Of course the fibers of \mathbf{H} are arranged radially

around these little axis pieces, in a sufficiently small neighborhood, so that when we open up the tree we can retain a local picture of the braid structure. During the looping motion the distinguished strand cuts each axis piece twice.

1.2.4 The cyclic templates Figure 8 gives an example of the cyclic template. We

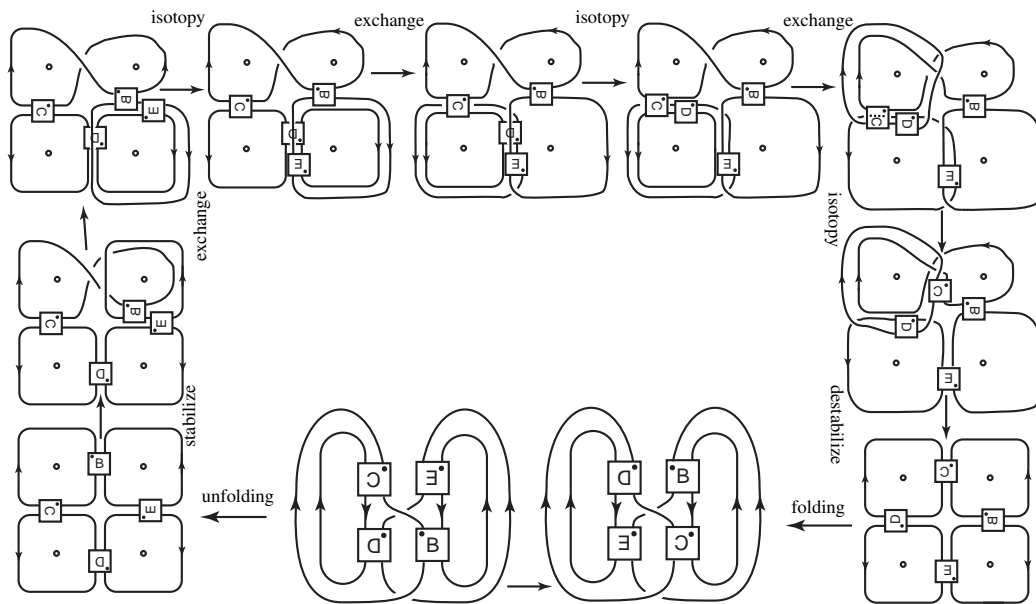


Figure 8: Example of a 4–braid cyclic template, and the Markov tower (stabilize, exchange, exchange, exchange, destabilize) that it replaces

have singled it out because it shows an interesting way in which stabilization introduces flexibility into the manipulation of closed braids, by allowing us to permute the blocks in a rather special and highly symmetric block-strand diagram. The associated closed braid diagrams have been unfolded to make it easier to follow the sequence of moves. A more general Markov tower for a cyclic template uses weighted strands, the entire tower being equivalent to permuting the blocks and weighted strands in a cycle. The resulting move on closed braids will be referred to as the *cyclic move*.

1.2.5 The G–flype and G–exchange templates The moves that we next describe are gathered together into the set of templates that we call $T(m)$, $m = 4, 5, \dots$, where m is the braid index $b(X_+)$ of X_+ .

A *G–flype* template (for ‘generalized’ flype template) is a block-strand diagram, the support of which is the replacement for a Markov tower in which (i) the first move

$X_1 \rightarrow X_2$ is an inadmissible flype; (ii) more generally adjacent terms in the sequence differ by destabilizations, exchange moves and not necessarily admissible flypes; (iii) $b(X_j) > b(X_1)$ for every $j = 2, \dots, k-1$; and (iv) $b(X_k) \leq b(X_1)$. An example is the boxed pair of 6–braid block-strand diagrams at the bottom of Figure 9. It can be understood by running around the diagram clockwise. As can be seen, the first step in that sequence is an inadmissible flype which increases the braid index by 1. The intermediate steps are exchange moves and the final step is an admissible flype that reduces braid index. In more general examples the final step could also be a destabilization. See Section 5.5 for a definition which shows precisely how G–flypes arise and describes the fixed and moving blocks that always occur.

A *G–exchange move* is the template that results from a sequence of k interrelated exchange moves, each of which moves a distinguished subarc $\{\alpha_1, \dots, \alpha_k\}$ of the closed braid across \mathbf{A} and has the property that the exchange move on α_i cannot be completed before at least part of the exchange move on some other α_j is started, for $1 \leq i \leq k$. See Section 5.6 for a definition which shows precisely how G–exchange moves arise and shows the fixed blocks that always occur.

The boxed pair of block-strand diagrams in the bottom row of Figure 10 are an example of a template for a G–exchange move on a 6–braid. There are 6 braid blocks: A, B, C, D, E, F . Running around the figure clockwise we show how a coordinated sequence of partial exchange moves, each of which can be completed as soon as enough of the other arcs are moved out of the way, achieves the same goal. In the first passage we have pushed strand a under the braid blocks A and B and across \mathbf{A} , to a position just to the right of braid block C . We have also lifted strand b above the braid blocks A and B and pulled it across \mathbf{A} to a position just to the left of braid block D . Then we begin our G–exchange move on arc c . In the fourth sketch we complete it. In the fifth sketch we begin the G–exchange move on arc d and complete the G–exchange move on arc a . In the final sketch we complete the G–exchange moves on arcs b and d . This example was discovered in the course of our proof. It illustrates the ideas developed in Section 5.6. Observe that, since exchange moves preserve link type and braid index, it follows that G–exchange moves do too.

1.3 Statement of results

We are finally ready to state our main result, the Markov theorem without stabilization. Let \mathcal{B} be the collection of all braid isotopy classes of closed braid representatives of oriented knot and link types in oriented 3–space. Among these, consider the subcollection $\mathcal{B}(\mathcal{X})$ of representatives of a fixed link type \mathcal{X} . Among these, let $\mathcal{B}_{\min}(\mathcal{X})$ be the subcollection of representatives whose braid index is equal to the braid index of \mathcal{X} .

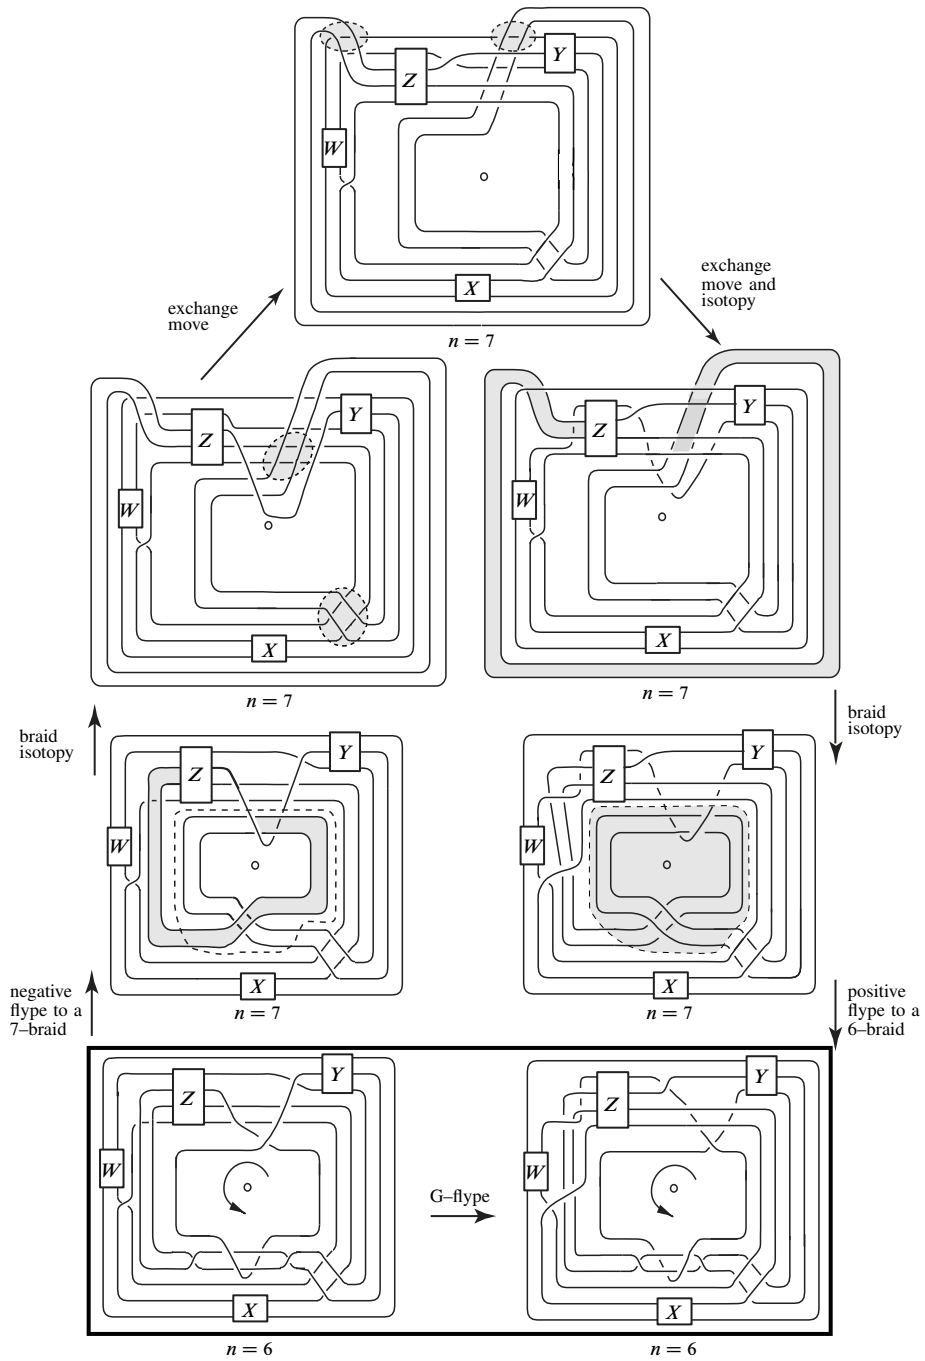


Figure 9: Example of a G-flype on a 6-braid




Figure 10: Example of a G-exchange move on a 6-braid

Theorem 2 (Markov’s Theorem Without Stabilization (MTWS)) Choose any $X_+ \in \mathcal{B}(\mathcal{X})$ and any $X_- \in \mathcal{B}_{\min}(\mathcal{X})$. Then

- there is a complexity function with values in $\mathbb{Z}_+ \times \mathbb{Z}_+ \times \mathbb{Z}_+$ which is associated to X_+, X_- , and
- for each braid index m there is a finite set $\mathcal{T}(m)$ of templates, each determining a move which is non-increasing on braid index,

such that the following hold:

First, there are initial modifications in the choices of X_+ and X_- , which replace them by X'_+ and X'_- :

$$(1-2) \quad X_- = X_-^1 \rightarrow \dots \rightarrow X_-^p = X'_-$$

$$(1-3) \quad X_+ = X_+^1 \rightarrow \dots \rightarrow X_+^q = X'_+$$

In (1-2) and (1-3) each passage $X_-^j \rightarrow X_-^{j+1}$ and $X_+^j \rightarrow X_+^{j+1}$ is strictly complexity-reducing. In (1-2) each passage is realized by an exchange move, so that $b(X_-^{j+1}) = b(X_-^j)$. In (1-3) each passage is realized by either an exchange move or a destabilization, so that $b(X_+^j) \geq b(X_+^{j+1})$.

After these initial modifications, there is another sequence taking X'_+ to X'_- :

$$(1-4) \quad X'_+ = X^q \rightarrow \dots \rightarrow X^r = X'_-$$

In (1-4) each passage $X^j \rightarrow X^{j+1}$ is strictly complexity-reducing, and is realized by either an exchange move, a destabilization, an admissible flype or a move defined by one of the templates $\mathcal{T} \in \mathcal{T}(m)$, where $m = b(X^j)$. The inequality $b(X^j) \geq b(X^{j+1})$ holds for each $j = q, \dots, r - 1$.

Remark 1.3.1 We explain why we refer to [Theorem 2](#) as *Markov's Theorem without Stabilization*. Setting $X^j = X_+^j$ when $j = 1, \dots, q - 1$, we have the combination of sequences (1-3) and (1-4):

$$X_+ = X^1 \rightarrow \dots \rightarrow X^q \rightarrow \dots \rightarrow X^r = X'_-$$

In this combined sequence $b(X^j) \geq b(X^{j+1})$ for each $j = 1, \dots, q, \dots, r$. On the other hand, $b(X_-) = b(X_-^1) = b(X_-^2) = \dots = b(X_-^p) = b(X'_-)$, so that the braid index of the second representative remains fixed during all of our modifications. Thus [Theorem 2](#) is, indeed, a version of [Theorem 1](#) which avoids stabilization. \diamond

Remark 1.3.2 The sequence (1-2) deals with the phenomenon which was exhibited in [Figure 6](#). It must be treated separately because if we only allowed modifications to X_+ then the complexity would be forced to increase as X_+ approached X_- , if X_- happened to be wound up as on the left in [Figure 6](#). Since (1-2) and (1-3) are inter-related, we treat (1-3) (which uses a limited subset of the moves in (1-4)) and (1-4) separately. \diamond

Remark 1.3.3 When a passage $X^j \rightarrow X^{j+1}$ is realized by a template \mathcal{T} , there are braiding assignments to the blocks in \mathcal{T} such that the initial and final diagrams of \mathcal{T}

carry the pair (X^j, X^{j+1}) . However, the template \mathcal{T} also carries infinitely many other knots and links, for other braiding assignments to the blocks. \diamond

Remark 1.3.4 The templates in the sets $\mathcal{T}(m), m \in \mathbb{Z}^+$, are precisely the additional moves which were not needed for the work in our earlier papers [8]–[10], but are needed for our particular proof of the MTWS. We discuss them briefly, starting with $m = 2$.

- (1) The 2–string braid group is an infinite cyclic group. Let σ_1 denote its generator. An arbitrary element is then $\sigma_1^k, k \in \mathbb{Z}$. It is easy to see that links which are closed 2–braids are either (i) the 2–component unlink ($k = 0$), or (ii) the unknot ($k = \pm 1$) or (iii) the type $(2, k)$ torus knots and links ($|k| \geq 2$). It is clear that the 2–component unlink and the type $(2, k)$ torus knots and links have unique closed 2–braid representatives. The unknot has exactly 2 closed 2–braid representatives, with σ_1 (resp. σ_1^{-1}) admitting a positive (resp. negative) destabilization. Since the set $\mathcal{T}(m)$ does not include the two destabilization templates, it follows that $\mathcal{T}(2) = \emptyset$. In our paper [6], which contains applications of [Theorem 2](#) to transverse knots, we will prove that, as a consequence of the main theorem in our paper [11], $\mathcal{T}(3) = \emptyset$.
- (2) It was proved by Fiedler [18] that closed 4–braids include infinitely many inequivalent 4–braid representatives of the unknot. His basic one is the example discovered by Morton [27], with the others obtained from it by the winding process which we illustrated in [Figure 6](#). Fifteen other families of 4–braid unknot examples were uncovered in the paper by Birman, Boldi, Rampichini and Vigna [13], in the course of a computer implementation of the unknot recognition algorithm of Birman and Hirsch [5]. All of them can be simplified to braids which admit a destabilization with the help of exchange moves. We do not know whether more general 4–braids are too complicated to be simplified with the use of the 2 destabilization templates, the exchange move template, the cyclic templates and the admissible flype templates.
- (3) Note that, given any template \mathcal{T} of braid index m , other templates for braid index $> m$ may be obtained from it by declaring the strands to be weighted, and also by replacing some of the blocks by other templates. From this it follows that the cardinality $|\mathcal{T}(m)|$ of $\mathcal{T}(m)$ is an increasing function of m . However, we do not have a precise description of $\mathcal{T}(m)$ for any $m > 3$, although we do not expect any fundamental difficulty in doing the actual enumeration for, say, $m = 4, 5$ and perhaps 6. For the special case $m = 6$ two examples were given in the boxed pairs of block-strand diagrams at the bottom of [Figures 9 and 10](#). The general picture seems to be quite complicated. \diamond

1.4 Plan of the paper

In [Section 2](#) we set up the topological construction which will be the basis for our work. We will show that there is a very special isotopy that takes us from X_+ to X_- . The trace of the isotopy will be seen to be an immersed annulus \mathcal{CA} whose double point set is the union of finitely many pairwise disjoint clasp arcs. We call it a ‘clasp annulus’.

The principle tool in our proof of [Theorem 2](#) is the study of certain ‘braid foliations’ of the immersed annulus \mathcal{CA} and its preimage \mathcal{PA} . Braid foliations were used by the authors in earlier work [\[8\]–\[10\]](#), but always in the setting of embedded surfaces. In [Section 3](#) we review the ideas that we need from the literature on braid foliations. Readers who are familiar with the literature will probably want to pass quickly over [Section 3](#), referring to it instead, as needed, later in the paper. In [Section 4](#) we study braid foliations of our immersed annulus. We will need to do hard technical work to arrange things so that the clasp arcs are close to or contained in a union of leaves and have nice neighborhoods (we call them normal neighborhoods) on the preimage annulus \mathcal{PA} .

In [Section 5](#) we learn how to translate data in the braid foliation of \mathcal{CA} and the induced foliation of \mathcal{PA} into data about the passage from the closed braid X_+ to the closed braid X_- . The tools that are needed become increasingly complicated as we proceed. First, we ask how far we can get with exchange moves and destabilizations. Flypes enter the picture next, but in the form of very rudimentary examples which we call ‘microflypes’. A rather surprising use of stabilization becomes apparent in [Section 5.3](#). Briefly, we learn that stabilization is the tool for creating flypes with weighted strands and complicated braiding assignments in the moving blocks out of microflypes. G–exchange come into play next. There are hints in this part of the work about the need for G–flypes, however the reasons for needing G–flype templates will not become clear until we are part way through the proof of the MTWS.

The proof of [Theorem 2](#) is given in [Section 6](#). We will see how G–flype templates arise. The most difficult part of the argument will be the proof that for each fixed braid index $m = b(X_+)$ the cardinality $|\mathcal{T}(m)|$ is bounded. The finiteness can, perhaps, be understood by appreciating that the ‘infinite parts’ are pushed into the blocks in the block-strand diagrams of the templates in $\mathcal{T}(m)$. This is, perhaps, the key point about block-strand diagrams and templates: they are at the same time both very flexible and very inflexible. A given template supports a huge family of knots and links, because there are no restrictions on the braiding assignments in the blocks, but on the other hand a template always supports at most a special family of links.

The paper ends, in [Section 7](#) with a discussion of open problems suggested by the MTWS and by its proof. In a separate paper [\[6\]](#) several applications will be given to the study of transversal knot types in the standard contact structure on \mathbb{R}^3 .

Conventions Results which will be used explicitly in the proof of the MTWS are highlighted by calling them ‘propositions’ rather than ‘lemmas’. There are many remarks, scattered throughout the paper. Most of them are guides to the reader. The end of a remark is indicated by the symbol \diamond .

Acknowledgments The authors single out for special thanks John Etnyre, Keiko Kawamuro and the unknown G&T referee. All three studied our manuscript with care and attention to details in the proof, and asked the kind of questions that can only be asked by an expert, leading us to clarify many tricky points.

This work has been in progress for so many years that it’s difficult for us to even remember everyone else who helped. With apologies for oversights we also thank Tara Brendle, Oliver Dasbach, Ivan Dynnikov, Brian Mangum, Walter Neumann, Nancy Wrinkle, Joel Zablow, Xingru Zhang, Joseph Birman and Melissa Menasco.

The first author acknowledges partial support from the US National Science Foundation, under grants DMS-9402988, 9705019, 9973232 and 0405586. She also wishes to thank the Mathematics Department at the Technion (Israel Institute of Technology) for hospitality during several visits when this paper was the main focus of her work. The second author acknowledges partial support from the US National Science Foundations, under grants DMS-9200881, 9626884 and 0306062; and the Mathematical Sciences Research Institute, where he was a Visiting Member during winter/spring of 1997.

2 Getting started

In this section we develop the basic construction which will allow us to prove [Theorem 2](#). For ease in presentation, we give our construction first for the special case when \mathcal{X} is a knot. After that it will be easy to modify it in the more general case when we begin with a link. The section will end with a key example.

2.1 The basic construction for knots

Lemma 2.1.1 *Let X_+ , X_- be arbitrary disjoint closed braid representatives of the same knot type \mathcal{X} . Then there is an intermediate representative X_0 of \mathcal{X} such that the following hold:*

- (1) X_0 is the braid connected sum of X_+ and k closed braid representatives of the unknot, for some $k \geq 0$. These k representatives of the unknot bound pairwise disjoint discs.
- (2) $X_0 - X_-$ (resp. $X_+ - X_0$) is the boundary of an embedded annulus \mathcal{A}_- (resp. \mathcal{A}_+). It will be seen from the construction that X_0 is the boundary of a Seifert surface \mathbf{F}_0 , and \mathcal{A}_- is a collar neighborhood of X_0 on \mathbf{F}_0 .
- (3) The intersections $\mathcal{A}_+ \cap \mathcal{A}_-$ are precisely k clasp arcs. See [Figure 11](#).




Figure 11: Clasp intersections

Remark 2.1.1 A word is in order on the basic construction. There are very simple constructions which yield everything in [Lemma 2.1.1](#) except the fact that $\mathcal{A}_- \subset \mathbf{F}_0$. Astute readers will notice that \mathbf{F}_0 plays an almost invisible role in the pages that lie ahead, and ask why we needed it? Our reasoning was that in our earlier papers we had developed extensive machinery regarding braid foliations of Seifert surfaces for knots and links, including the case of discs bounded by unknots. In this paper we need related braid foliations of the immersed annulus of [Lemma 2.1.1](#). The easiest way to develop the tools that we need seemed to be to make one of our annuli (ie \mathcal{A}_-) a subset of a Seifert surface, and arrange that the other (ie \mathcal{A}_+) is a finite family of foliated discs banded together with narrow strips to form a second annulus that intersects the first in a controlled fashion. That will allow us to make full use of the existing braid foliation machinery for Seifert surfaces and discs, rather than to develop everything anew. This issue will become clear in [Section 3](#) and [Section 4](#) below. \diamond

Proof Without loss of generality we may assume that $X_+ \subset \mathbb{R}_+^3$ and $X_- \subset \mathbb{R}_-^3$, with X_+ far above $\mathbb{R}^2 = \mathbb{R}_-^3 \cap \mathbb{R}_+^3$, and X_- a little bit below \mathbb{R}^2 . We may further assume that X_+, X_- are closed braids with respect to the same braid structure (\mathbf{A}, \mathbf{H}) . Our first task is to construct a series of knots X'_-, X''_-, X_0 , all representing \mathcal{X} , with X_0 the braid-connected sum of X_+ and k pairwise disjoint and pairwise unlinked copies U_1, \dots, U_k of the unknot.

Choose a Seifert surface \mathbf{F}_- for X_- . Let $X'_- \subset \mathbf{F}_-$ be a preferred longitude for X_- , chosen close enough to X_- , so that the annulus that they cobound in \mathbf{F}_- does not intersect \mathbf{A} , and so that X'_- is also a closed braid. The knots X_- and X'_- will have algebraic linking number 0, but $X_- \cup X'_-$ will not be a split link unless \mathcal{X} is the unknot. Therefore, if we try to push X'_- out of \mathbb{R}^3 and into \mathbb{R}^3_+ it will get stuck, ie there will be a finite number of ‘undercrossing hooks’ where X'_- is forced to dip back into \mathbb{R}^3_- to pass under X_- , as in Figure 12(a).

Our first change is to modify X'_- (holding X_- fixed) to a new closed braid X''_- which has the same knot type as X'_- and is entirely in \mathbb{R}^3_+ . This can be accomplished by pushing X'_- across pairwise disjoint discs $D_1 \cup \dots \cup D_k$, as in Figure 12(b) to X''_- . By placing X_- very close to \mathbb{R}^2 and choosing the discs to be very ‘thin’ we may assume that each subarc β_i is in braid position and in \mathbb{R}^3_+ . By construction X''_- is a closed braid (because X'_- is a closed braid and the isotopy $X'_- \rightarrow X''_-$ is arbitrarily small), also it represents \mathcal{X} , and it is entirely in \mathbb{R}^3_+ .

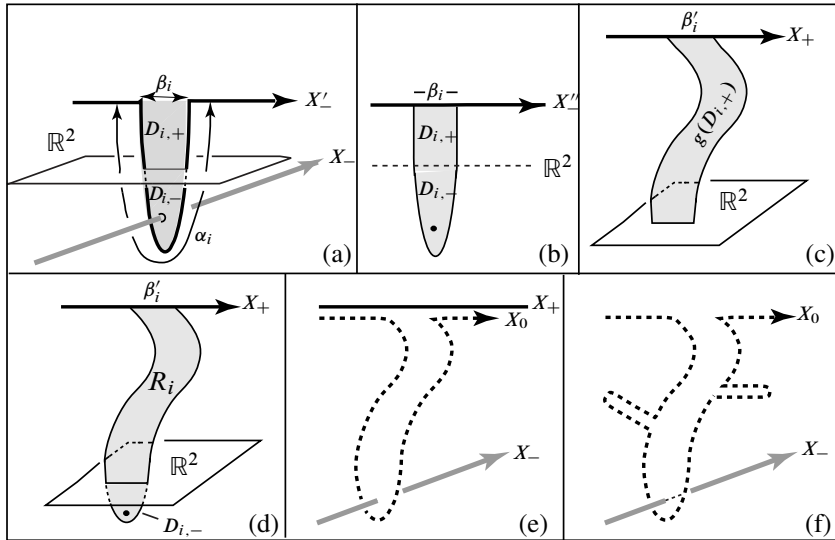


Figure 12: Constructing X_0

We are now ready to bring X_+ into the picture. The fact that X''_- and X_+ both represent \mathcal{X} and are both in the interior of \mathbb{R}^3_+ shows that we may find a homeomorphism $g: \mathbb{R}^3_+ \rightarrow \mathbb{R}^3_+$ which is the identity on \mathbb{R}^2 with $g(X''_-) = X_+$. Extend g by the identity on \mathbb{R}^3_- to a homeomorphism $G: \mathbb{R}^3 \rightarrow \mathbb{R}^3$. Let $R_i = G(D_{i,+} \cup D_{i,-}) = g(D_{i,+}) \cup D_{i,-}$. The facts that (i) G is a homeomorphism which is the identity in \mathbb{R}^3_- and (ii) if $i \neq j$ then $D_i \cap D_j = \emptyset$ tell us that the R_i 's are pairwise disjoint embedded discs, and also

that X_- pierces each R_i exactly once. The fact that X_+ was well above \mathbb{R}^2 shows that we may assume that each R_i intersects X_+ in a single arc $\beta'_i = g(\beta_i)$.

Let X'_0 be the knot which is obtained from X_+ by replacing each $\beta'_i \subset X_+$ by $\partial R_i \setminus \beta'_i$. Then X'_0 is the connected sum of X_+ and k copies of the unknot, the i^{th} copy being ∂R_i . By construction the k discs R_1, \dots, R_k are pairwise disjoint, so that our k unknots represent the k -component unlink. It may happen that X'_0 is not a closed braid. The only subarcs which might not be in braid position are the $2k$ ‘vertical’ arcs in each $\partial g(D_{i,+})$. To overcome this problem, recall that JW Alexander [1] introduced a very simple way to change an arc δ which is not in braid position to one that is. See Figure 13(a). Alexander showed that whenever a knot X is not transverse to the fibers of \mathbf{H} it can be divided into small segments which can then be pushed across the braid axis, one at a time, avoiding unwanted intersections with the rest of X , to change X to a closed braid. Using this construction, we change all the wrongly oriented subarcs of the $2k$ vertical arcs in the boundaries of the ‘ribbons’ $R_1 \cup \dots \cup R_k$ to subarcs that are everywhere transverse to the fibers of \mathbf{H} . We have proved (1).




Figure 13: (a) Alexander’s trick (b) A fragment of the preimage \mathcal{PA} of $\mathcal{A}_+ \cup \mathcal{A}_-$

There is an important aspect of our construction, which will give us the part of (2) that relates to \mathcal{A}_- :

- $X_- \cup X_0$ has the same link type as $X_- \cup X'_0$. For, by construction, the homeomorphism $G^{-1}: \mathbb{R}^3 \rightarrow \mathbb{R}^3$, being the identity on \mathbb{R}^3_- , sends $X_- \cup X_0$ to $X_- \cup X'_0$.

This simple fact gives us the annulus \mathcal{A}_- , in the following way: Since X'_- is a preferred longitude for X_- , and since $X_- \cup X_0$ has the same link type as $X_- \cup X'_-$, it follows that X_0 is also a preferred longitude for X_- . From this it follows that X_- is also a preferred longitude for X_0 . Choose a Seifert surface \mathbf{F}_0 for X_0 . Holding X_0 and X_- fixed, isotope the interior of \mathbf{F}_0 until X_- lies on \mathbf{F}_0 as a preferred longitude. Let $\mathcal{A}_- \subset \mathbf{F}_0$ be the annulus in \mathbf{F}_0 which X_0 and X_- bound. This annulus is embedded because \mathbf{F}_0 is embedded. Thus we have proved the part of (2) that relates to \mathcal{A}_- .

In fact, a small modification in X_0 also gives us \mathcal{A}_+ . The discs $R_1 \cup \cdots \cup R_k$ have a natural order which is determined by the order of the subarcs $\beta'_1 \cup \cdots \cup \beta'_k$ along X_+ . Using this order, and the framing provided by \mathbf{F}_0 , join the discs $R_1 \cup \cdots \cup R_k$ by k very narrow bands in \mathbf{F}_0 , each having one edge on X_+ . Modify X_0 by pushing it a little bit into \mathbf{F}_0 along the bands and along the R_i 's. (By an abuse of notation, we use the same names for the modified R_i 's and the modified X_0). The union of the new R_i 's and the bands is our annulus \mathcal{A}_+ . The annulus \mathcal{A}_+ is embedded because the R_i 's are disjoint and embedded, and the bands are too. We have proved (2). Since $X_- \subset \partial \mathcal{A}_-$ and since X_- pierces each $R_i \subset \mathcal{A}_+$ once, it follows that $\mathcal{A}_- \cap \mathcal{A}_+$ always contains k clasp arc intersections. Thus we have also proved (3), and the proof of [Lemma 2.1.1](#) is complete. \square

We establish conventions that will be used throughout this paper. The symbol \mathcal{CA} denotes the immersed annulus $\mathcal{A}_+ \cup \mathcal{A}_-$. We shall refer to it as a *clasp annulus*. We will also be interested in its preimage \mathcal{PA} under the immersion. [Figure 13\(b\)](#) is a schematic that illustrates our basic construction, as described in [Lemma 2.1.1](#). The knot X_0 will only be needed in the beginning of our argument, therefore we show it as a dotted curve. The closed braid X_+ will be the primary focus of our attention, and so we show it as a thick black line. We will eventually modify it to X_- , which we illustrate as a thick grey line (to suggest that it is a distant goal). To avoid clutter in our figures we will, whenever the meaning is unambiguous, suppress the labels X_+ and X_- . Most of the time the black-grey convention will enable us to recognize them without labels.

Remark 2.1.2 The reader may wonder why the isotopy which we constructed from X_+ to X_- in [Lemma 2.1.1](#) required us to increase the braid index of X_+ by connect-summing it with some number of closed braid representatives of the unknot, because one expects these to have braid index greater than 1. This does indeed sound wrong, in view of the fact that our goal is to have the braid index go down and not up, but in fact it is to be expected because we are attempting to prove the MTWS, and for that we first need to have in hand a proof of the MT. Indeed, in [\[12\]](#) we showed that a variation on the construction given here in [Lemma 2.1.1](#) and [Proposition 2.2.1](#) (but

without taking any note of the clasp arcs) can be used to give a new proof of Markov's classical result. When we began the work in this paper we started with that new proof of the Markov theorem and proceeded to modify the Markov tower that it gave to us. Later, we realized that it was not necessary to literally prove the Markov theorem, all we needed was the isotopy encoded by \mathcal{CA} . We hope that explains the logic. \diamond

2.2 The general case

Proposition 2.2.1 *Choose any μ -component oriented link type \mathcal{X} in oriented \mathbb{R}^3 . Let X_+, X_- be closed braid representatives sharing a braid structure (\mathbf{A}, \mathbf{H}) . Then an intermediate closed braid representative $X_0 \in \mathcal{X}$ exists such that the following hold:*

- (1) *Each component X_0^i of X_0 is the braid connected sum of X_+^i and k^i closed braid representatives of the unknot. These $k = k^1 + \dots + k^\mu$ representatives of the unknot bound k pairwise disjoint discs.*
- (2) *There is a pairing of the components of X_+ and X_- such that for each pair (X_+^j, X_-^j) there exists a closed braid X_0^j which represents the same component of \mathcal{X} . Moreover, for each $j = 1, \dots, \mu$ the following holds: $X_+^j - X_0^j$ (resp. $X_0^j - X_-^j$) is the boundary an embedded annulus \mathcal{A}_+^j (resp. \mathcal{A}_-^j). The union of these μ embedded annuli forms an embedded surface \mathcal{A}_+ (resp. \mathcal{A}_-).*
- (3) *The intersections $\mathcal{A}_+ \cap \mathcal{A}_-$ are a finite collection of pairwise disjoint clasp arcs. These intersections can be clasp intersections between the annuli \mathcal{A}_+^i and \mathcal{A}_-^i associated to a single component \mathcal{X}^i of \mathcal{X} and also clasp intersections between the annuli \mathcal{A}_+^i and \mathcal{A}_-^j associated to distinct components $\mathcal{X}^i, \mathcal{X}^j$ of \mathcal{X} .*
- (4) *Moreover, without loss of generality we may assume that \mathcal{X} is a non-split link. If, instead, \mathcal{X} is a split link, then the basic construction may be applied to each non-split component.*

Proof We leave it to the reader to check that, except for (4), the proof is a minor adaptation of the argument that we gave in the proof of [Lemma 2.1.1](#). As for (4), we note that in the manuscript [\[7\]](#) the authors used tools that are closely related to the tools that will be used in this paper to prove that if X is an arbitrary closed n -braid representative of a split (resp. composite) link, then there is a complexity function $C(X)$ with values in \mathbb{Z}_+ which is associated to X such that after a strictly complexity-reducing sequence of exchange moves, all of which preserve braid index, X can be changed to a split (resp. composite) closed n -braid representative of \mathcal{X} . Thus, for split (resp. composite) links, the proof of the MTWS can be preceded by applying the results in [\[7\]](#) to find the non-split (resp. prime) summands. However, we note that non-prime

knots and links do not require any special consideration in our work. Therefore there would not be any point in making the assumption that \mathcal{X} is prime. On the other hand, there are two places where the assumption that \mathcal{X} is not split simplifies our work a little bit, so in what follows we make the assumption (4) that \mathcal{X} is non-split. \square

2.3 A key example of the basic construction

The reader is referred to Figure 14. It is a key example, and we will study it in full detail during the course of this manuscript. We explain those features which can be understood at this time.

In our example \mathcal{X} is a knot, and the isotopy from X_+ to X_- is realized by a flype. The top sketch shows \mathcal{PA} as the union of two discs which are identified along bands which join the arcs ab and $a'b'$, and also cd and $c'd'$. Figure 14 depicts the disc neighborhood of one of its clasp arcs. The black (resp. grey) boundary is X_+ (resp. X_-). To avoid clutter we do not show X_0 . There is one clasp arc whose two preimages γ_- and γ_+ are shown. The clasp arc γ_ϵ , where $\epsilon = \pm$, has one endpoint on X_ϵ and the other in the interior of the annulus, at the point where $X_{-\epsilon}$ pierces the annulus. The bottom sketch illustrates the clasp annulus \mathcal{CA} . To visualize \mathcal{CA} in 3-space, first




Figure 14: A key example

give each of the discs in the top sketch a half-twist, as in the leftmost and rightmost bottom sketches. Then identify the two half-twisted discs along the clasp arc, as in the middle bottom sketch. The passage from X_+ (the black boundary) to X_- (the grey

boundary) is realized by a push of X_+ across the immersed annulus \mathcal{CA} to X_- . While the annulus is immersed, there are no self-intersections of the boundary braid as it is pushed across the annulus.

In fact, our example illustrates a (negative) 3–braid flype, as in the passage from the 3rd to the 4th sketch in [Figure 3](#), in the special case when

$$X_+ = \sigma_1^p \sigma_2^{-2} \sigma_1^q \sigma_2^{-1}, \quad X_- = \sigma_1^p \sigma_2^{-1} \sigma_1^q \sigma_2^{-2}$$

with p and q integers, also $p + q$ odd (so that we get a knot) and absolute value at least 3 (so that, by the work in our paper [\[11\]](#), we know that the flype cannot be replaced by braid isotopy). Notice that in the passage from the left to the right sketch in [Figure 3](#) the isotopy is supported in a 3–ball B^3 whose boundary intersects X_+ in 4 points. These are the endpoints of the 2 subarcs of X_+ in the bottom middle sketch in [Figure 14](#). The 3–ball B^3 contains in its interior the braid box R (which in our case is a single negative full twist), a little subarc of the braid axis and the single crossing to its right (which in our case is also negative). The signs of the singular points were chosen to correspond to the fact that the exponents of σ_2 in the braids which represent X_+ and X_- are both negative. The braiding in those blocks depends, of course, on the choices of the exponents p and q . They are examples of the ‘fixed blocks’ which are formed from the ‘braiding of bands of s –arcs’. (See the proof of [Proposition 5.3.2](#)).

It should be clear to the reader that (except in very special cases) the braiding of the long thin bands will lead to geometric linking between X_+ and X_- , and this implies that there will be additional ‘short clasp arcs’ in the bands. That matter will be discussed in [Section 4.2](#).

We remark that the puncture endpoints of the two clasp arcs in [Figure 14](#) are illustrated as being signed, but their signs have not been defined or discussed. We will need them later, so we define them now. The puncture point on a clasp arc is said to be a *positive* puncture point (resp. *negative* puncture point) according as the orientation of X_ϵ at the puncture agrees (resp. disagrees) with the orientation of the outward drawn normal, when the surface has the orientation induced by the orientation on X_+ . We will see later that the sign on one of the punctures determines the sign on the other. In [Figure 14](#) we illustrated the case when the left puncture is negative and the right puncture is positive, but the opposite choices are also possible.

3 Introducing braid foliations

[Theorem 2](#) is about the relationship between two closed braid diagrams which represent the same link. However, the method we use to prove it will not be in the setting of

link diagrams. Rather, we will be working with the immersed annulus \mathcal{CA} , and with certain ‘braid foliations’ of \mathcal{CA} . Foliated surfaces have been used before, in our earlier papers [8]–[10]. In this section we will review and describe the machinery which we use from these other papers. The reader who has seen these foliations before will be able to omit this section and go directly to Section 4, possibly pausing to refer back to this review to refresh his/her memory of details. A more detailed review may be found in the review article by Birman and Finkelstein [4].

3.1 Braid foliations of Seifert surfaces

We are given $X \in \mathcal{B}(\mathcal{X})$ with $b(X) = n$ and with braid structure (\mathbf{A}, \mathbf{H}) in \mathbb{R}^3 . To make this review as simple as possible, we assume that X is a knot. Choose a Seifert surface \mathbf{F} of minimal genus for X . After modifying \mathbf{F} we will show that it supports a special type of singular foliation which was studied and used by the authors in [8]–[10]. We call it a *braid foliation*.

There are choices of orientation which determine the sign conventions in braid foliations. First, we assign the standard orientation to \mathbb{R}^3 and choose the braid axis \mathbf{A} to be the positively oriented z -axis. Using cylindrical coordinates, this determines a sense of increase of the polar angle coordinate θ . Next, the closed braid X is oriented so that it points in the direction of increasing θ at each point of $X \cap H_\theta$. The orientation on X induces an orientation on \mathbf{F} , and so determines a positive normal at each interior point on \mathbf{F} .

Since $X = \partial\mathbf{F}$ is a closed braid, $\mathbf{A} \cap \mathbf{F}$ is non-empty. The intersections of \mathbf{H} with \mathbf{F} may be assumed to be (a) radial in a neighborhood of each point of $\mathbf{A} \cap \mathbf{F}$ and (b) transverse to the boundary in a neighborhood of $\partial\mathbf{F}$. See Figure 15.

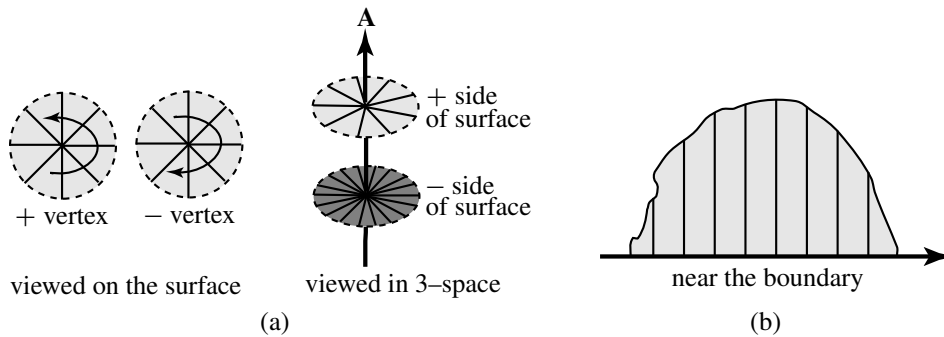


Figure 15: (a) Foliation of surface near a vertex (b) Foliation of surface near the boundary

Vertices in the foliation are points where \mathbf{A} pierces \mathbf{F} . We call a vertex positive or negative, according as \mathbf{A} intersects \mathbf{F} from the negative or positive side of \mathbf{F} respectively. The sketches in Figure 15 illustrate the positive side of \mathbf{F} , so that the flow induced by the fibration is anticlockwise (resp. clockwise) about a positive (resp. negative) vertex.

Singularities in the foliation occur at points where \mathbf{F} is tangent to one of the fibers of \mathbf{H} . The singularities may be assumed to be finite in number and to occur on distinct fibers of \mathbf{H} . By Morse's classification theorem every singularity may be assumed to result from a local maximum or minimum or a saddle point tangency between \mathbf{F} and a fiber of \mathbf{H} . Let s be a singular point of the foliation of \mathbf{F} , and let H_θ be the disc fiber which contains s . We say that the singularity s is positive if the outward-drawn oriented normal to the oriented surface \mathbf{F} coincides in direction with the normal to H_θ in the direction of increasing θ . Otherwise s is negative.

Leaves in the foliation are components of intersection of H_θ with the surface \mathbf{F} . A *singular leaf* is a leaf which contains a singularity of the foliation. Every other leaf is a *non-singular leaf*.

A very basic property of our braid foliations of Seifert surfaces is that non-singular leaves may all be assumed to be arcs. We review the reasons (which goes back to Bennequin [2]). Suppose that there is a simple closed curve $\alpha(\theta)$ in $\mathbf{F} \cap H_\theta$ for some non-singular polar angle θ . The fact that \mathbf{F} is pierced non-trivially by the braid axis, and that the foliation is transverse to the boundary, shows that \mathbf{F} cannot be foliated entirely by simple closed curves, so if we follow the sequence of arcs $\alpha(\theta)$ as θ increases or decreases we must arrive at a singularity. Let H_{θ_0} be the singular fiber. The singularity may be assumed to be either be a center or a saddle point, but if it is a center, then by following $\alpha(\theta)$ in the opposite direction we will arrive at another singularity, and that one cannot also be a center because \mathbf{F} is not a 2-sphere, so it must be a saddle point, and the singularity must be a homoclinic point (a singularity which is formed when a generic leaf has a saddle point singularity with itself), as illustrated in Figure 16(a). Note that the singular leaf $\alpha(\theta_0)$ lies in both \mathbf{F} and H_{θ_0} , and necessarily bounds a disc Δ in the latter. Assuming that Δ is innermost, we surger \mathbf{F} along Δ as in Figure 16.

The surgered surface has two components. By assumption \mathbf{F} has maximal Euler characteristic, which implies that one of the two components is a 2-sphere. Discarding it, and smoothing the new \mathbf{F} , we can eliminate the singularity.

It remains to consider the case when the interior of the disc Δ intersects \mathbf{F} . Since s_0 is the only singularity in H_θ , there are no singularities in the interior of Δ . But then each component of $\mathbf{F} \cap \text{int}(\Delta) \subset H_\theta$ must be a simple closed curve. Let c be an innermost such simple closed curve. Then we can surger \mathbf{F} along c , and then smooth the surgered

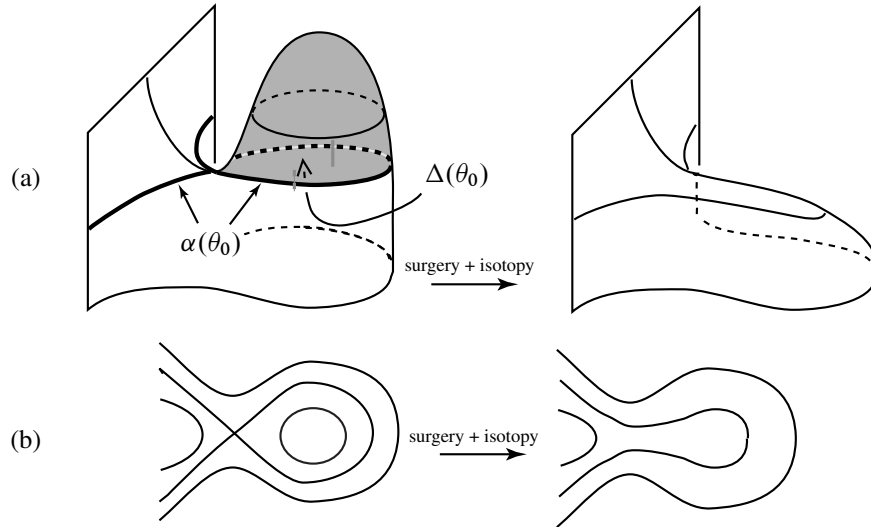


Figure 16: In this sketch the fibers of \mathbf{H} are to be thought of as horizontal planes. (a) Surgering \mathbf{F} along the disc Δ removes the singularity. (b) The change in foliation on \mathbf{F} .

surface by an isotopy. This will introduce center singularities but no additional saddles. Any S^2 's that are formed we discard. (We leave it to the reader to draw appropriate pictures). After a finite number of such surgeries we obtain a new surface \mathbf{F}' which has the same homoclinic point in its foliation, but does not meet the disc Δ . We then do the surgery which is illustrated in Figure 16, reducing the number of saddle point singularities. In this way all leaves which are simple closed curves can be eliminated.

Since each non-singular leaf is an arc, one of its endpoints could either be at a vertex of the foliation or a point on the boundary, however we now claim that non-singular leaves which have both endpoints on the boundary do not occur. For, suppose there is a leaf l which has both endpoints on the boundary. Let $N(l)$ be a neighborhood of the leaf on \mathbf{F} . Thinking of $N(l)$ as a rectangle, it will have a pair of opposite edges on the boundary, and the orientation on these edges will be consistent with the orientation on $N(l)$. However, l also lies in a fiber H_θ of \mathbf{H} , and the boundary of \mathbf{F} is a closed braid. But then, the orientation on one of the edges of $N(l)$ will agree with that of the normal to H_θ and (since l is by hypothesis non-singular) the orientation of the other must disagree with the that of the normal, but that is impossible because the boundary curve is a closed braid.

It follows from this that the non-singular leaves have two types: those which have one endpoint on the boundary and the other at a vertex (we call them *a-arcs*) and those

which have both endpoints being vertices of the foliation (we call them *b*-arcs). See Figure 17.




Figure 17: Non-singular leaves in the foliation of \mathbf{F}

Singularities fall into three types, which we call types *aa*, *ab* and *bb*, the notation indicating that just before an *aa*-singularity (resp. *ab*, *bb*-singularity) the non-singular leaves were both type *a* (resp. types *a* and *b*, resp. both type *b*). We shall refer to the 2-cells which are foliated neighborhoods of the singular leaves as ‘tiles’. See Figure 18.




Figure 18: Tiles of type *bb*, *ab*, *aa*

The foliation may be used to decompose the surface \mathbf{F} into a union of foliated 2-cells, each of which contains exactly one singularity of the foliation. Each 2-cell is a regular neighborhood on \mathbf{F} of a singular leaf. These foliated 2-cells are our *tiles* and the resulting decomposition of \mathbf{F} is a *tiling*. See the three sketches in Figure 18. The *tile vertices* are the points where the braid axis \mathbf{A} intersects the surface \mathbf{F} . (There are also other vertices on $\partial\mathbf{F}$, but we prefer to exclude them when we refer to tile vertices.) The *tile edges* are arbitrary choices of *a*-arcs or *b*-arcs. (There are also other tile edges which are subarcs of X , but it will be convenient to ignore those, just as we ignored the vertices which are on X .)

Assume that the vertices and singularities in the braid foliation of \mathbf{F} have been labeled. The *combinatorial data* which is associated to the foliation is a listing of the signed vertices, in the natural cyclic order in which they occur on \mathbf{A} , and a listing of the signed singularities, in the natural cyclic order in which they occur in the fibration. The following proposition is very natural, because leaves in the foliation of \mathbf{F} are level sets for the embedding of \mathbf{F} :

Proposition 3.1.1 (Birman–Menasco [10], Birman–Finkelstein [4, Theorem 4.1])
Let T be any tile in the braid foliation of \mathbf{F} . Label the vertices and singularities of T by their signs and their cyclic orders as above. Then this decorated braid foliation determines the embedding of T . More globally, the embedding of all of the aa and ab tiles in the foliation of \mathbf{F} determines the embedding of X in $\mathbb{R}^3 \setminus \mathbf{A}$.

We illustrate in Figure 19 the embeddings of the three tiles in Figure 18, for one of the finitely many choices of the combinatorial data on these tiles.




Figure 19: Embeddings of tiles of type bb , ab , aa in 3-space

3.2 Control over the foliations

Braid foliations are not unique, and in this section we describe some of the ways we have discovered to modify them. This is an essential part of the argument in the proof of the MTWS, because the vertices of valence 1 and 2 that we use to recognize destabilizations and exchange moves may not be present initially, but after a change in foliation they may be present. The existence of a vertex of valence 1 or 2 will be one of our ways to learn when the complexity can be reduced.

The foliation of \mathbf{F} depends upon the choice of half-planes H_θ in the fibration of $\mathbb{R}^3 \setminus \mathbf{A}$. A *change in fibration* is the choice of a new set of half-planes $\mathbf{H} = \{H_\theta : 0 \leq \theta \leq 2\pi\}$.

Equivalently, one could fix the fibers of \mathbf{H} and the link, and move the surface. This induces a *change in foliation*.

The changes which we shall make are always very controlled and very local in terms of changing the tiling. In particular they do not change the braid, and are supported in a neighborhood in \mathbb{R}^3 of subarcs of singular leaves. The question of when such changes are possible has been studied. They were used in earlier papers by the authors, and we use those results as needed here. We describe two changes in foliation. For details, see Birman–Finkelstein [4].




Figure 20: The first change in foliation, in the case when both tiles have type bb . There are two other cases, obtained from the one which is illustrated by deleting part of the bb -tile and adding one or both dotted arcs as boundary, to convert to an ab -tile. The 3 cases are: case (1), two bb -tiles; case (2), one ab and one bb -tile; case (3), two ab -tiles.

Lemma 3.2.1 (First change in foliation) *Let s_1 and s_2 be singularities of the same sign $\delta = \pm$ in tiles D_1 and D_2 , where D_1 and D_2 intersect in a common leaf vw of type b . For example, see the left sketch in Figure 20, which relates to the case when both tiles have type bb . Then after a change in foliation which is supported on a neighborhood in 3-space of an arc γ which joins the two singular points, the foliation of $D_1 \cup D_2$ changes in one of the two ways which are illustrated in Figure 20. In particular, the valence of the vertices v and w decreases as a result of the change.*

Proof See [4, Theorem 2.1] for a very detailed proof of the Lemma. We note the following feature of the proof. There are 6 vertices in $D_1 \cup D_2$, labeled v, a, y, w, z, b .

The proof in [4] shows that they remain fixed during the change in foliation. There are also singular leaves vy (resp. vz) with one endpoint at the vertex v in the region of interest and the other at y (resp. z). In the passage from the left to the middle (resp right) sketch in Figure 20 there is, at every point in the isotopy, a singular leaf vy (resp. vz) which contains one of the two singularities. This justifies our labeling that singularity s'_1 (resp. s'_2) in the middle (resp. right) sketch, because it evolves directly from s_1 (resp. s_2) during the isotopy which realizes the change. We call the other singularity s'_2 (resp. s'_1). In fact (see [4]) if s_1 and s_2 occur in an angular interval $[\theta_1, \theta_2] \subset [0, 2\pi]$ in the fibration \mathbf{H} , with $s_1 < s_2$ in this interval, then after the change in foliation the new singularities will still be in the same angular interval, only now, in both cases, we will have $s'_1 > s'_2$. \square

We will need one more change in foliation. It is similar to that of Lemma 3.2.1, except that it holds without restriction as to the signs of the two singularities which are involved. We call it the *second* change in foliation. The reader may find Figure 21 helpful in understanding what it says.




Figure 21: Sketch (a) shows the second change in foliation. Sketch (b) shows N , embedded in 3-space, before the change, illustrating the position of Δ . The move is a push downward of s_1 along Δ .

Lemma 3.2.2 (Second change in foliation) *Let γ be an arc which is located in a foliated disc N which is a subset of a foliated surface. Assume that the foliation of N contains exactly two singular points s_1 and s_2 . Let l_i be the singular leaf through*

$s_i, i = 1, 2$. Let H_{θ_i} be the fiber of \mathbf{H} which contains $s_i, i = 1, 2$. Suppose that there is a disc Δ in 3-space, such that:

- (i) The interior of Δ has empty intersection with \mathbf{F} .
- (ii) $\partial\Delta = \alpha \cup \gamma$, with $\Delta \cap N = \gamma$ and $\Delta \cap H_{\theta_2} = \alpha$.
- (iii) Δ is trivially foliated, ie there are no vertices or singularities in Δ .
- (iv) $\gamma \cap l_1 = s_1$ and $\gamma \cap l_2 = \{p\}$, where p is a point.

Then after a change in foliation which is induced by pushing N along Δ , changing the order of s_1 and s_2 , we may assume that $\gamma \cap l_1 = s'_2$ and $\gamma \cap l_2 = \emptyset$. Moreover, the change in the foliation of \mathbf{F} may be assumed to be supported on an arbitrarily small neighborhood on N of the subarcs $[s_1, p] \subset \gamma$ and $[p, s_2] \subset l_2$. There could also be several singularities s_1, s_2, \dots, s_k , with associated neighborhoods, and if the conditions are met for each of them in turn then the disc Δ may be used to push s_1 past many singularities, one at a time.

Proof Figures 21(a) shows the foliated disc N before and after the change which we propose to make. The arrows which are attached to the leaves indicate the direction of increasing θ . Using the foliation of N , and knowing the signs of the singularities, one may construct an embedding of N in 3-space, and we have done so in Figure 21(b) in the case when the signs of the singularities at s_1 and s_2 are different. (The other case is similar). In Figure 21(b) fibers of \mathbf{H} are to be thought of as horizontal planes. The auxiliary disc Δ , is also illustrated. The move which we make to realize the change in foliation in Figure 21(a) is to push N down along the disc Δ . To understand how this changes the foliation of N , we have labeled certain endpoints on ∂N with numbers 1,2,3,4,5,6. There are non-singular leaves which we call 12, 34, 56, each with arrows directed inward (to illustrate the direction of increasing θ) and joining 1 to 2, 3 to 4 and 5 to 6 respectively. In the left picture the first singularity occurs when leaf 12 approaches leaf 34, but in the right picture the first singularity occurs when leaf 12 approaches leaf 56. \square

3.3 Using braid foliations to detect destabilizations and exchange moves

In this section we will show how to recognize, from the foliation of \mathbf{F} , when a closed braid admits a destabilization or an exchange move.

Destabilizations are easy. It is shown in [4] that X_+ admits a destabilization if the foliation has a vertex of valence 1, as in Figure 22(a). The embedding of a tile which contains the vertex of valence 1, for one of the two possible choices of the sign of the singularity, is illustrated in sketch (b).




Figure 22: Destabilization along a valence 1 vertex, viewed (a) on the foliated surface and (b) in 3-space

Lemma 3.3.1 *Destabilizations which are predicated on the existence of a valence 1 vertex reduce the number of singularities (resp. vertices) in the foliation of \mathbf{F} by 1 (resp. 1).*

Proof Clear. See Figure 22(a). □

Before we can describe our exchange moves, a new concept is needed. We observe that, intuitively, b -arcs in the foliation of \mathbf{F} arise when there are ‘pockets’ in the surface, and we are now interested in the case when a ‘pocket is empty’ and so can be removed. We now make this precise. A b -arc β in the foliation is an *inessential* b -arc if it joins a pair of vertices v, w which are consecutive vertices in the natural cyclic ordering of vertices along \mathbf{A} . We use the term ‘inessential’ because, if we think of a fiber H_θ of \mathbf{H} as a disc with the braid axis \mathbf{A} as its boundary, then an inessential b -arc will cut off a disc $\Delta \subset H_\theta$ which has the property $\Delta \cap \mathbf{F} = \beta$. The disc Δ cannot be pieced by X , for if it were there would be an a -arc α which either intersected our inessential b -arc, which is impossible because \mathbf{F} is embedded, or α has a vertex endpoint between v and w , which is impossible because v and w are consecutive vertices on \mathbf{A} . The disc Δ can be used to push β (and nearby leaves in nearby disc fibers) across \mathbf{A} , reducing the number of vertices in the foliation of \mathbf{F} . Peek ahead to the right sketch in Figure 24 for examples of b -arcs which are inessential and also essential. An *essential* b -arc is one which is not inessential.

Exchange moves were defined in Figure 2 as a move on a block-strand diagram. Our task now is to detect them in the foliation. The complexity function mentioned in

Theorem 2 will include the number of singularities in the foliation of a clasp annulus which is bounded by the given braids. As will be seen shortly, exchange moves always reduce this complexity function. However, the exchange moves that are used in this paper come in two flavors, and it's necessary to check both.

The *ab*-exchange move: In **Figure 23(a)** we have illustrated v , a type *ab* valence




Figure 23: The type *ab*-exchange move

2 vertex. The left and right sketches in (a) (resp. (b)) show the configurations on the foliated surface \mathbf{F} (resp. in 3-space), the left sketches being before and the right sketches being after the exchange move. In both (a) and (b) the vertex v is adjacent to two type *ab* singularities, and these singularities have opposite signs. The *ab*-exchange move occurs in the passage from the left to the right sketch. One pushes the subarc $\alpha \subset X$ across the shaded disc to α' . If we choose α' so that it is everywhere transverse to leaves of the foliation (as is clearly possible) the move takes closed braids to closed braids. **Figure 23(b)** illustrates the embedding for one particular choice of signs and orderings.

Note that if the *b*-arcs which end at the vertex v are all essential, then X must encircle the subarc vw of the braid axis which is inside the pocket. The braid changes we make can then be understood by looking at the projection onto a plane orthogonal to the braid axis. After the exchange move the shaded disc will have vanished. Peering down the braid axis (as we did in **Figure 2**) we see that the projection has changed in the predicted manner. For full details, consult [4].

The *bb*-exchange move: In Figures 24 and 25 we have illustrated v , a type *bb* valence 2 vertex. The non-singular leaves which have an endpoint at v are all type *b*. The left and right sketches show the changes in 3-space. The changes in the braid projection can be understood by looking down the axis onto a plane orthogonal to the braid axis. The foliation does not change at all after a *bb*-exchange move, however there is a




Figure 24: The type *bb*-exchange move, followed by an isotopy of \mathbf{F}

change in the order of the vertices along the braid axis. The ‘pocket has been emptied’ and after the exchange move, the empty pocket can be collapsed by pushing every *b*-arc $\beta(\theta)$ across its disc $\Delta(\theta)$, as in Figure 25. Remark: The pocket could of course have much more complicated braiding inside it. The proof that it can always be emptied in this way is non-trivial; details may be found in our earlier paper [7] or the article by Birman and Finkelstein [4]. The passage from the left to right sketches in Figure 25 shows the change in the foliation, after the removal of all inessential *b*-arcs.




Figure 25: Changes in foliation after the removal of inessential *b*-arcs

Lemma 3.3.2 Suppose that the foliation of \mathbf{F} has a vertex v of valence 2 and type *ab* or *bb*. Assume that the adjacent singularities have opposite signs. Then the closed braid $X = \partial\mathbf{F}$ admits an exchange move. After the move, there is a surface \mathbf{F}' , isotopic to \mathbf{F} , with a decomposition containing two fewer vertices and two fewer singularities than the decomposition of \mathbf{F} .

Proof See Figures 22(a), 24 and 25. □

3.4 Using braid foliations to detect stabilizations

As elaborated on in the introduction to this paper, one of the questions that motivated our work was a desire to understand why stabilization played such an important role in the classical Markov Theorem. Having braid foliations of Seifert surfaces bounded by knots and links in hand, one answer to that question became clear: they allow one to simplify a Seifert surface by eliminating negative vertices and their associated singularities. The first sketch in Figure 26 shows how: When there is an ab -singularity




Figure 26: Stabilization along an ab -tile, viewed (a) as a move on the foliated surface, and (b) concentrating on how it alters the boundary

in the foliation, we may push X across the associated negative vertex and its singularity, in a neighborhood of the singular leaf, to a new position which is again everywhere transverse to the foliation. It follows that after we do this move we will have a new closed braid representative. This move simplifies \mathbf{F} because it eliminates a vertex and a singularity. The right sketch shows why the move is actually a stabilization. Figure 27 shows our stabilization move on the embedded surface in 3-space. If one looks carefully one can see the half-twist which has been introduced in the course of the push. We note that the pictures of ab -tiles in Figure 27 are deformations of the picture in Figure 18: we stretched out the top sheet to make visible a neighborhood of the singular leaf.

4 Braid foliations of the immersed annulus

We are now ready to investigate braid foliations of the clasp annulus \mathcal{CA} which we constructed in Section 2. We will need to confront two new aspects of the geometry:




Figure 27: Stabilization along an ab -tile, viewed in 3-space

- (1) There is a (basically trivial) new aspect to our geometry, which unfortunately will lead to new bookkeeping: The clasp annulus \mathcal{CA} has two boundary components, X_+ and X_- . By the basic construction, the X_- boundary is a curve in the interior of the chosen Seifert surface \mathbf{F}_0 for the closed braid X_0 . As was shown in the previous section, \mathbf{F}_0 admits a braid foliation. Clearly X_- can cut through the tiles in that foliation in any way as long as it is always transverse to the leaves of the foliation. Therefore we will allow for new tile types, to account for the partial tiles at the X_- boundary of \mathcal{CA} .
- (2) The second new aspect of the geometry is central to the work in this paper. Our annulus \mathcal{CA} is not embedded.

A preliminary modification in the clasp arcs will be helpful in what follows:

Lemma 4.0.1 *We may assume that the k clasp arcs are transverse to fibers of \mathbf{H} , and so also to the leaves in the braid foliation of \mathcal{CA} and \mathcal{PA} .*

Proof Let $\Gamma = \gamma_1 \cup \cdots \cup \gamma_k$ be the union of the clasp arcs. Then the graph $X_+ \cup X_0 \cup X_- \cup \Gamma$ is embedded in \mathcal{CA} , which is a subset of \mathbb{R}^3 . We focus now on that graph. By our earlier construction, its subsets X_+ , X_0 and X_- are in braid position, but in general Γ is not in braid position, ie the interior of some clasp arc may not be transverse to the fibers of \mathbf{H} . From the proof of [Lemma 2.1.1](#) we employ Alexander's braid trick to every wrongly oriented subarc of Γ , doing it so as to avoid intersections with X_+ , X_0 and X_- . The construction allows us to find an orientation-preserving PL homeomorphism $f: S^3 \rightarrow S^3$ which changes Γ to braid position. A classical result of Gugenheim (see Hempel [19, Theorem 1.5]) then tells us that we may assume that f is isotopic to the identity. Replacing \mathcal{CA} by $f(\mathcal{CA})$, and using the fact that f leaves X_+ , X_0 and X_- invariant, it follows that we may assume that every subarc of the graph $X_+ \cup X_0 \cup X_- \cup \Gamma \subset \mathcal{CA}$ is in braid position. \square

4.1 Tile types in \mathcal{PA}

Our work begins with the two closed braid representatives X_+ and X_- of \mathcal{X} . Our basic construction in Section 2 gave us the immersed annulus $\mathcal{CA} = \mathcal{A}_+ \cup \mathcal{A}_-$ and the clasp arcs. We begin by studying the braid foliations of the two annuli. The key point which will allow us to apply the machinery of Section 3 is that each is embedded. The really new feature is the presence of the clasp arcs. We first make the clasp arc as nice as possible (we change them to ‘tabs’). After that we will remove ‘short clasp arcs’, ie ones which do not pass through any singular leaves. Then we will create ‘normal neighborhoods’ about the doubly modified clasp arcs, which will isolate them from the rest of \mathcal{CA} . Our device for isolating them is to introduce lots of inessential b -arcs, which will give us the freedom we need to modify the clasp arcs when we need to do so.

In this section we are interested in the foliation of \mathcal{A}_+ and \mathcal{A}_- which are induced by intersections of these annuli with the half-planes of our braid structure. The closed braids X_+, X_0, X_- are all oriented so that they point in the direction of increasing θ at each point of intersection with an H_θ . We choose an orientation on the annulus \mathcal{A}_+ in such a way that it induces the given orientation on its boundary component X_+ . Notice that this means that the orientation on X_0 does not agree with that induced by the chosen orientation on \mathcal{A}_+ . Similarly, we choose an orientation on the annulus \mathcal{A}_- in such a way that it induces the given orientation on its boundary component X_0 , which implies that the orientation on X_- does not agree with that induced by the chosen orientation on \mathcal{A}_- .

As in the situation of Seifert surfaces, the foliation may be assumed to be radial in a neighborhood of each point of $\mathbf{A} \cap \mathcal{A}_\pm$ (see Figure 15(a)) and transverse to the boundary in a neighborhood of $\partial\mathcal{A}_\pm$ (Figure 15(b)). The braid axis \mathbf{A} pierces \mathcal{A}_\pm from either the negative or the positive side at each pierce point, and we have indicated this by attaching positive or negative signs to the pierce points on \mathcal{A}_\pm . As before, leaves in the foliation are singular if they contain a singularity of the foliation, otherwise they are non-singular. The singularities may be assumed to be finite in number and to occur on distinct fibers of \mathbf{H} .

We now show that, as in the situation of Seifert surfaces, we may assume that there are no leaves in the foliation of $\mathcal{CA} = \mathcal{A}_+ \cup \mathcal{A}_-$ which are simple closed curves. See Figure 28. There are new issues to settle. The first question we ask is whether a simple closed curve c could intersect both \mathcal{A}_+ and \mathcal{A}_- ? Assume that the homoclinic loop determined by c is innermost in H_θ , and also innermost in the foliation of \mathcal{CA} . Then c bounds a disc on $\mathcal{A}_+ \cup \mathcal{A}_-$ which is foliated by simple closed curves, so it would intersect X_0 twice if it intersects it at all. This would force a tangency between X_0

and a fiber of \mathbf{H} , but X_0 is in braid position, so this cannot happen. Therefore any leaf in the foliation which is a simple closed curve must lie in the interior of \mathcal{A}_ϵ , $\epsilon = +$ or $-$.

Next, for simplicity assume that the foliation of \mathcal{A}_+ contains the homoclinic loop c and refer to Figure 28. Our leaf c bounds a disc Δ_θ in the fiber H_θ which contains c . See Figure 28(a). (If c is homologically nontrivial in \mathcal{A}_+ then, since it bounds a disc in H_θ , X_+ would have a component that was the unknot.) If c is not intersected by any clasp arcs then we can apply the argument of Section 3 (see Figure 16). Since both Δ_θ and Δ_c are embedded their union $\Delta_\theta \cup_c \Delta_c$ is an embedded 2–sphere.

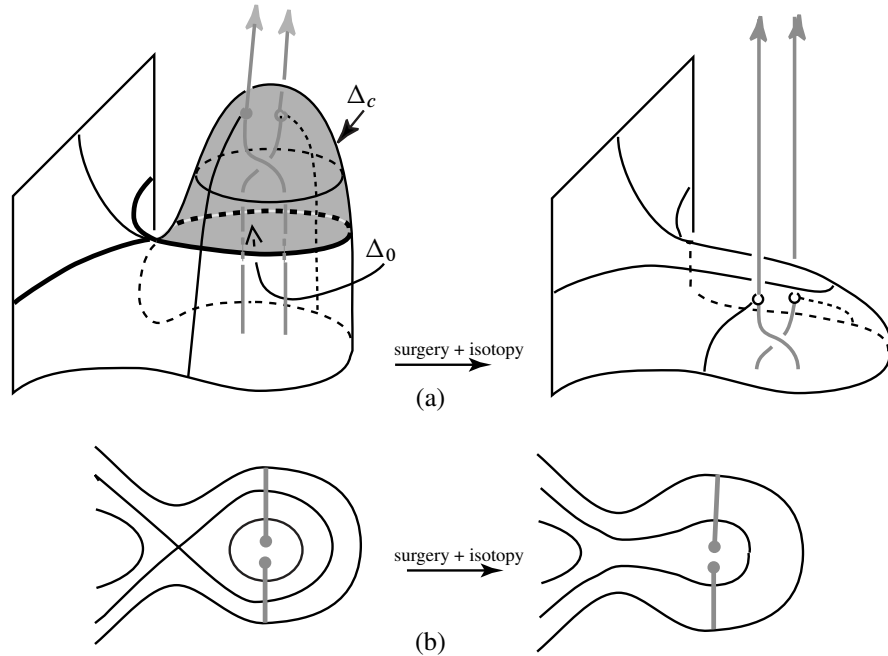


Figure 28: The elimination of a homoclinic singularity and associated simple closed curves, in the presence of clasp arcs. Sketch (a) is the geometric realization. Sketch (b) shows the preimage in $\mathcal{P}\mathcal{A}$.

Now consider a clasp arc that intersects c . Since clasp arcs are transverse to the foliation of $\mathcal{C}\mathcal{A}$, any clasp arc that intersects c intersects it exactly once. Thus, any clasp arc that intersects c must have a puncture endpoint in Δ_c , where X_- intersects \mathcal{A}_+ . (Referring to Figure 28(a), the gray braid strands represent X_- .) Since any braid that enters the 2–sphere $\Delta_\theta \cup_c \Delta_c$ must puncture Δ_θ , and can only exit by puncturing Δ_c , we conclude that only X_- can intersect Δ_θ .

Next, we consider the intersection of $\Delta_\theta \cap (\mathcal{A}_+ \cup \mathcal{A}_-) \subset \Delta_\theta$. In principle, this intersection set can contain three types of arcs/curves: arcs that have an endpoint on both X_- and c ; arcs that have both endpoints on c ; and simple closed curves. Notice that there can be no arcs that have both endpoints on X_- because this would violate the orientation of X_- . Since we are assuming that c is innermost for the moment we ignore the issue of simple closed curves. Referring to [Figure 28\(b\)](#) notice that for every puncture point in Δ_c there is a point on c that is an intersection with a clasp arc. But, also for every puncture point of X_- with Δ_c there must be a puncture point of X_- with Δ_θ . So the second type of arc of intersection—having both endpoints on c —cannot occur. If we then perform the surgery illustrated in [Figure 28\(a\)](#), we see that this corresponds to a truncation of the clasp arcs. That is, this surgery on a simple closed curve in the homoclinic leaf replaces the immersed annulus with a new immersed surface which is union of an annulus and a 2–sphere, where the latter can be discarded. It replaces a clasp intersection with a clasp intersection and ribbon intersection on the discarded 2–sphere. Afterwards, clasp intersections are still in braid position. This surgery eliminates at least one saddle singularity (and, possibly some points of $\mathcal{CA} \cap \mathbf{A}$). Reiterating this procedure we will arrive at a point where the foliation of \mathcal{CA} has no leaves that are simple closed curves, except in one situation.

The situation where we will not be able to perform the surgery illustrated in [Figure 28](#) is when the homoclinic point involves a singularity between an s –arc (see [Figure 29](#)) and a simple closed curve where the disc Δ_θ contains the endpoints of the resulting singular leaf which are in X_+ and X_- . However, we can alter the foliation of \mathcal{CA} in the following manner to eliminate the leaf c . Let γ_ϵ be any clasp arc that has an endpoint $e_{-\epsilon} \subset \Delta_c$. Let $\alpha \subset \mathcal{CA}$ be an arc that starts at $e_{-\epsilon}$ and ends at X_ϵ such that $\text{int}(\alpha)$ does not intersect any clasp arcs. It is convenient to take the support of α to be in a union of Δ_c and a regular neighborhood of the homoclinic singular leaf; and it can be assumed that α is transverse to the foliation of \mathcal{CA} except at a single point in Δ_c near $e_{-\epsilon}$. We now perform Alexander’s braid trick on α to make $\alpha \cup \gamma_\epsilon$ transverse to \mathbf{H} . In the resulting new foliation of \mathcal{CA} the path $\alpha \cup \gamma_\epsilon$ cannot intersect any leaf that is a simple closed curve for reasons of orientation. It is easily seen that once the homoclinic point is eliminated for γ_ϵ it will be eliminated for all clasp arcs that intersect Δ_c . No new homoclinic points are introduced.

Finally, if c is not innermost and $\Delta_\theta \cap (\mathcal{A}_+ \cup \mathcal{A}_-) \subset \Delta_\theta$ does contain other simple closed curves, we can achieve the assumption that c is innermost by first performing a surgery which is similar to that illustrated in [Figure 28](#) on all of the simple closed curves in Δ_θ , starting with the innermost. Such surgeries may or may not eliminate any saddle singularities, but they will create center points in the foliation. But, then we can perform the surgery to eliminate the homoclinic saddle singularity on c .

We have learned that we may assume that every leaf is an arc. The situation is a little bit more complicated than it was in the case of Seifert surfaces. Consult Figure 29. There are now 5 possible types of non-singular arcs in the foliation of \mathcal{A}_+ and \mathcal{A}_- : (i) arcs which have both endpoints at vertices (type b); (ii) arcs which have exactly one endpoint at a positive vertex (type a_+); (iii) arcs which have exactly one endpoint at a negative vertex (type a_-); (iv) arcs which have one endpoint on one boundary component and the other on the opposite boundary component (type s); and (v) arcs which have both endpoints on the same boundary component. As for type (v), the exact argument that we used in the case of Seifert surfaces applies, because X_+ , X_0 , X_- are all closed braids, so type (v) does not occur.




Figure 29: Non-singular leaves in the foliation of $\mathcal{A}_+ \cup \mathcal{A}_-$, as viewed in the preimage of \mathcal{CA}

Remark 4.1.1 With regard to Figure 29 we may need to reinterpret the boundary components, temporarily, as (X_+, X_0) or (X_0, X_-) instead of (X_+, X_-) . This ambiguity will be removed shortly. See Remark 4.3.1 below. \diamond

The annulus $\mathcal{A}_+ \cup \mathcal{A}_-$ is said to be *trivially foliated* or *trivial* if it is foliated without clasp arcs and the leaves in its foliation are all s -arcs.

An exhaustive list of the singularities which could, in principle, occur in the foliation of \mathcal{A}_+ and \mathcal{A}_- are types a_+a_+ , a_+b , a_+s , a_+a_- , a_-a_- , a_-b , a_-s , sb , ss , and bb , where the notation is consistent with that used in Section 3. However, we have:

Lemma 4.1.1 *Singularities of type ss do not occur in either \mathcal{A}_+ or \mathcal{A}_- .*

Proof There is an oriented flow on $\mathcal{A}_+ \cup \mathcal{A}_-$ that is transverse to the foliation. In particular, we have arranged that the foliation is transverse to the boundary in small neighborhoods of X_+ and X_- , so that the flow will be parallel to the boundary and oriented to agree with the orientations on X_+ and X_- , both of which are oriented closed braids. On the other hand, the fact that every singularity in the foliation is a saddle point implies that every singularity in the associated flow is hyperbolic. Suppose

that there was a singularity of type ss at polar angle 0. Then, for an arbitrarily small negative polar angle $-\epsilon$ there must be a pair of type s non-singular leaves s_1, s_2 moving toward one another on a collision course. However, a type s leaf separates the annulus and as we just noted the flow near the two endpoints of both s_1 and s_2 is in the direction of the orientation of X_+ and X_- . It is then impossible for the images of s_1 and s_2 in the flow to approach one-another, as the flow evolves. \square

Lemma 4.1.2 *If a singularity of type $a-a_-$ or $a-s$ occurs, then the associated singular leaf is always intersected by a clasp arc.*

Proof Suppose that a singularity of either type $a-a_-$ or $a-s$ occurs, and that no clasp arc intersects the associated singular leaf. See Figure 30. Notice that there is an arc,




Figure 30: Type $a-a_-$ and $a-s$ singularities do not occur

α , contained in the singular leaf which has both of its endpoints on X_- and which, together with one of the two subarcs β of $X_- \setminus (X_- \cap \partial\alpha)$ forms a simple closed curve C which bounds a disc $D \subset \mathcal{A}_-$. Since C bounds a disc it represents the unknot. Now observe that if we perturb the endpoints of α slightly along X_- we may change α to an arc α' which is transverse to fibers of \mathbf{H} . Thus, after an arbitrarily small perturbation, we may change C to $C' = \alpha' \cup \beta$, where C' is a closed braid and where $\beta \subset X_-$ is a subarc of the closed braid X_- . Let b be the braid index of C' . The fact that C' is everywhere transverse to fibers of \mathbf{H} guarantees that $b \geq 1$. But then, we may reduce the braid index of X_- by at least one without changing its knot type by replacing β by α' . However that is impossible because by hypothesis the braid index of X_- is minimal. \square

4.2 Preliminary modifications in the clasp arcs

We assume from now on that our clasp annulus \mathcal{CA} supports a braid foliation, so that its preimage \mathcal{PA} supports the lifted foliation. We continue to use the symbols \mathcal{CA} and \mathcal{PA} , but from now on \mathcal{CA} means the *foliated* clasp annulus and \mathcal{PA} means its *foliated* preimage. Each clasp arc γ in \mathcal{CA} will have two preimages γ_+ and γ_- in \mathcal{PA} , where




Figure 31: Possible tile types in the foliation of $\mathcal{A}_+ \cup \mathcal{A}_-$. The black (resp. grey) arcs always represent subarcs of X_+ (resp. X_-), with one exception: when we construct the tabs we will be working in \mathcal{A}_- and X_0 will enter the picture as a (dotted) grey arc.

γ_ϵ indicates the preimage of γ which begins on X_ϵ in the boundary of \mathcal{PA} and ends at its *puncture endpoint* on $X_{-\epsilon}$ in the interior of \mathcal{PA} .

We say that a clasp arc is *short* if γ_+ or γ_- does not cross any singular leaves. Since our ultimate goal is to push X_+ across \mathcal{CA} to X_- , every clasp arc will ultimately become short.

Lemma 4.2.1 *After a braid isotopy, we may assume that there are no short clasp arcs. Even more, suppose that γ_ϵ is a short clasp arc in \mathcal{PA} . Suppose further that the only non-singular leaves that γ_ϵ intersects are s -arcs and a_ϵ -arcs (without dividing the vertices of any $a_\epsilon a_\epsilon$ -singular leaves). Then we may eliminate γ_ϵ .*

Proof Since a short clasp arc does not cross any singular leaves and has one endpoint on X_+ it must be in the boundary of a region in \mathcal{PA} which is near X_+ and foliated entirely by a -arcs and s -arcs (call it an a/s -region). See Figure 32. Let e_+ and e_- be the endpoints of the clasp arc γ_+ , where $e_+ \in X_+$. Then e_- is on an a/s -arc, and running along that arc to X_+ we obtain a point $p \in X_+$. Modifying e_-p slightly, we make it transverse to the fibers of \mathbf{H} . The ‘triangular’ region e_+e_-p , shaded in Figure 32, is foliated entirely by subarcs of a/s -arcs. Rescale X_+ in such a way as to shrink the shaded region, pulling e_+ to p and γ_+ into the arc pe_- .




Figure 32: Eliminating short clasp arcs

The clasp arc is still in an a/s -region, as in the right sketch of Figure 32, which shows the local picture in 3-space. The black (resp. grey) boundary arc is X_+ (resp. X_-). We now push X_+ along e_+e_- to eliminate the short clasp arc, changing the geometric linking between X_+ and X_- in the process. Note that while it looks as if we are ‘unlinking’ the two curves, what we are really doing is to correct the geometric linking between X_+ (which started out by being geometrically unlinked from X_-) and X_- by putting it into the position of X_0 . \square

Remark 4.2.1 By our basic construction, the only way that \mathcal{CA} will fail to have clasp arcs is if X'_- , the pushoff of X_- onto a Seifert surface which was constructed during the proof of Lemma 2.1.1, is geometrically unlinked from X_- . The only way that can happen is if \mathcal{X} is the unlink. So, if \mathcal{X} is not the unlink, then clasp arcs occur. Short clasp arcs arise in the basic construction in situations where we could have arranged for X_+ to have the correct linking with X_- locally by a braid isotopy. If it happened that every clasp arc was short, then we would know that, after a braid isotopy of X_+ , it could be assumed that X_- is a preferred longitude for X_+ , ready to be pushed across \mathcal{A}_- , an embedded annulus. \diamond

4.3 Construction of the tabs

In this section we modify the discs R_1, \dots, R_k which we constructed in Section 2 to special foliated discs $\mathbf{T}_+^1, \dots, \mathbf{T}_+^k$ containing the clasp arcs.

We say that $\mathbf{T}_+^i \subset \mathcal{PA}$ is a *tab* associated to the clasp arc γ_+^i if the following hold:

- $\gamma_+^i \subset \mathbf{T}_+^i$.
- $\partial \mathbf{T}_+^i = \alpha \cup \beta$ where $\alpha \subset X_+$ and β is an arc that is transverse to the foliation of \mathcal{PA} .

- There is a simple path l_+ which is contained in a union of singular leaves in \mathbf{T}_+^i . The path l_+ starts on X_+ , ends at a negative vertex, and contains all the negative vertices in \mathbf{T}_+^i .
- The arc γ_+^i is the only clasp arc which intersects \mathbf{T}_+^i . It is everywhere transverse to the leaves in the foliation of \mathcal{PA} . The arc γ_+^i intersects each of the singular leaves in the induced foliation of \mathbf{T}_+^i exactly once.




Figure 33: An example of a tab

Lemma 4.3.1 We may assume that each R_i is a tab \mathbf{T}_+^i .

Proof We focus on a single $R_i \subset \mathcal{PA}$, where R_i is one of the discs that we constructed in Section 2. We construct the tab \mathbf{T}_+^i as a subset of the disc R_i in the basic construction. Going back to Section 3, recall that the boundary of R_i is a union of 4 arcs:

- a subarc of the modified X_0 ,
- a subarc of X_+ , and
- two modified s -arcs which join them.

It contains γ_+^i in its interior. Using the methods described in Section 3, we may assume that R_i supports a braid foliation. Initially, the foliated disc R_i will not look anything like a tab. What we wish to do is to rechoose the discs R_1, \dots, R_k so that they are as simple as possible. The new discs that we choose will be subsets of the old ones, with possibly modified foliations. In the course of the construction we will, in effect, be making a new choice of the separating curve X_0 which divides \mathcal{PA} into the embedded annuli \mathcal{A}_- and \mathcal{A}_+ . We note that right after the completion of this proof we will discard X_0 , as we will have no further use for it.

The arc γ_+^i has one of its endpoints at a point e_+^i on X_+ and its other endpoint $e_-^i \in X_-$ in the interior of R_i . In between, it winds around the vertices of the foliated

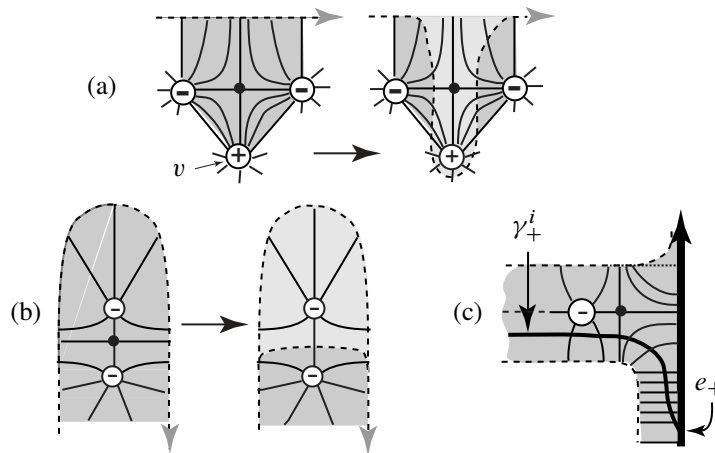


Figure 34: Rechoosing the disc R_i . In sketches (a) and (b) the dark area is in R'_i . The passage from left to right shows how we modify R'_i to eliminate positive vertices.

disc R_i , constrained to remain transverse to the leaves of the foliation. We may assume without loss of generality that the point e_-^i where X_- pierces the interior of R_i is not on a singular leaf. Therefore e_-^i belongs to a non-singular leaf of type s, a_+, a_- or b .

Let $\alpha \subset R_i \setminus \gamma_+^i$ be a simple path that starts at e_-^i and ends at a point $p \in X_+$. Then $\gamma_+^i \cup \alpha \subset R_i$ is a simple path that joins e_-^i to p in the interior of R_i . Without loss of generality we may assume that this simple path is in braid position. For, if it is not, apply Alexander's braiding trick to wrongly oriented subarcs. This will change the interior of R_i by an isotopy in 3-space. The changes will modify the foliation of R_i by introducing new vertices and singularities which allow α to avoid the points of non-transversality. After these modifications, let $R'_i \subset R_i$ be the foliated subdisc that $\gamma_+^i \cup \alpha$ splits off in R_i . Reapply the argument for eliminating leaves that are circles in the foliation of R'_i , so that R'_i supports a braid foliation. Therefore we may assume that R'_i supports a braid foliation.

We are now in position to re-choose γ_+^i, α and X_0 so that every vertex in R'_i is negative. See Figure 34. If the induced foliation of R'_i contains a positive vertex v then this sub-foliation must also contain a singular leaf that has its endpoint on $\gamma_+^i \cup \alpha$ and is adjacent to a (possibly different) positive vertex v' . We can then push $\gamma_+^i \cup \alpha$ along the singular leaf and across v , staying transverse to the foliation, as in Figure 34(a), moving v out of R'_i . Inducting on the number of positive vertices in the foliation of R'_i , we have arranged that the foliation of R'_i contains only negative vertices.

Now the graph of singular leaves in R'_i is either a linear tree, or it is a tree with branches. In the latter case, since e_-^i can be adjacent to only one negative vertex, either γ_+^i or α intersects a singular leaf in the foliation of R_i twice. We can then find a valence one vertex, as in [Figure 34\(b\)](#), with either γ_+^i or α in its boundary. After a ‘destabilization’ we can eliminate this negative vertex from R_i . Iterating this procedure, we alter $\gamma_+^i \cup \alpha$, until R'_i is changed to a tab.

There is only one more problem. It may happen that near the X_+ boundary of R'_i , either γ_+^i or α or both run along a band which is foliated without singularities by a and/or s -arcs, as depicted in [Figure 34\(c\)](#) in the case of γ_+^i . If so, we simply ‘rescale’ X_+ , as we did in the proof of [Lemma 4.2.1](#) to pull the long arcs back into the single a - s -tile which forms the base of the tab. \square

Remark 4.3.1 We make an important remark about simplified notation and simplified tile types. Having standardized the tabs, we will not have further use for X_0 . This eliminates the dual meanings of the boundary arcs in [Figure 31](#) and allows the following simplification. If $\epsilon = \pm$, an a_ϵ -arc has its interior endpoint on a vertex of sign ϵ and has its boundary endpoint on X_ϵ . At the same time, we will no longer need to depict X_0 in our figures, and will be free to use dotted arcs in other settings as we proceed through the proof of the MTWS. \diamond

Remark 4.3.2 Notice that the work in this section uses the braid foliation of R_i . If we had braid foliations on hand as a tool, during the basic construction, we could have arranged at that time for each R_i to be a tab. \diamond

Remark 4.3.3 In [Figure 14](#) we gave an example of a tab with the clasp arc on one side of the singular leaves. Soon we will develop the ability to move it so that it is transverse to the foliation and to one side or the other of a sequence of singular leaves, or alternatively to position it in a union of leaves (singular and/or non-singular), at the expense of introducing many new b -arcs into the foliation. See [Section 4.5](#) below. \diamond

4.4 The two finger moves

We need tools that will allow us to modify neighborhoods of the clasp arcs in \mathcal{CA} in a controlled manner, keeping track of the foliation on the two branches and making sure that no new self-intersections are introduced. The ‘finger moves’ will help us to do that. See [Figures 35](#) (and eventually [36](#)). To begin, we explain the pictures in [Figure 35](#). The first column shows foliated neighborhoods N of a subarc of one of the clasp arcs. We give separate pictures of the two foliated branches, N_1 (the *first branch*) and N_2 (the *second branch*). They intersect transversally. By hypothesis the clasp arc (and




Figure 35: Controlled local changes in the clasp arc in a neighborhood of a singularity after the first and second finger moves




Figure 36: The 4 possible choices of the signs of the pair (vertex, singularity)

so also its preimages γ_+ and γ_-) is transverse to the leaves of \mathbf{H} . We have oriented the clasp arc (arbitrarily).

We are interested in modifying the position of one of the clasp arcs in a neighborhood of a singular leaf on one of two branches N_1 or N_2 . Since singular leaves correspond to places where one of the two branches is tangent to a fiber of \mathbf{H} , and since the two branches intersect transversally along a clasp arc, we may assume that there is a neighborhood of the singular point in which the other branch is foliated without singularities. The neighborhood has been chosen so that N_2 is foliated without singularities, but N_1 contains a singularity, together with the vertex endpoint of one of the branches of the singular leaf. There are two sign choices: the sign of the vertex and the sign of the singularity. We have chosen these to be $(+, +)$, but in a moment we will consider the 4 possible sign choices. There are also little '+' signs next to the two components of the clasp arcs. They indicate which side of N_2 (resp. N_1) is the positive side, at γ_ϵ (resp. $\gamma_{-\epsilon}$).

Selected fibers H_{θ_i} have labels $i = 1, 2, 3$. The labels on the fibers and the way that γ intersects them and the little + signs next to γ_ϵ and $\gamma_{-\epsilon}$ completely determine the position of N in 3-space, relative to the coordinate system provided by the fibers of \mathbf{H} . We are ready to describe our two controlled changes in the clasp arc, and the corresponding changes in the foliation of N_1 and N_2 and in the way that γ_ϵ and $\gamma_{-\epsilon}$ intersect the leaves in the foliation.

We now define the moves precisely:

- (1) Our *first finger move*, illustrated in the middle column of [Figure 35](#), pushes γ across the horizontal branch of the singular leaf in the first branch, creating two points where it is not transverse to the fibers of \mathbf{H} , one on fiber 1 and the other on fiber 3. This move is always possible, because it occurs within an arbitrarily small neighborhood of the singularity. The corresponding change in the second branch can be understood by noticing that before the change γ_ϵ intersected fibers 1,2,3 transversally in that order. After the change there are two points of tangency with fibers of \mathbf{H} , the first with fiber 3 and the second with fiber 1. This explains the doubling back of $\gamma_{-\epsilon}$ after the first finger move.
- (2) Let α denote the subarc of the clasp arc which is between the two points of non-transversality, and let $p \in \alpha \cap \text{fiber 2}$ be the point which is closest to the singularity on γ_ϵ . The *second finger move*, illustrated in the right column of [Figure 35](#), pushes a neighborhood $N(p)$ of p on α across the singularity and across the vertex, staying within a neighborhood of the singular leaf. The foliation on the first branch is unchanged. The foliation of the second branch changes in an arbitrarily small neighborhood of p on the second branch. Two new vertices

of opposite parity and two new singularities of opposite parity are created, as illustrated in the bottom row sketch.

See [Figure 36](#) for the local changes with the four possible sign choices for the pair (vertex,singularity).

Remark 4.4.1 The two finger moves are always possible because of their local nature and because of our control over the geometry. We have illustrated the case when the singularity and vertex are both positive. The other three cases differ from this one by local symmetries. Observe that we have given a great deal of detailed data in [Figures 35](#) and [36](#), including the cyclic order of leaves in the fibration, the signs of the singularities and the signs of the vertices. Fibers of \mathbf{H} are level sets for positioning the two branches of the immersed surface in 3-space. The test for whether the finger moves are realizable in 3-space is to examine them on a sequence of fibers of \mathbf{H} , and the data in [Figures 35](#) and [36](#) suffices for that purpose.

If braid foliations had been available as a tool during the basic construction, we could have used the two finger moves then. We stress this because later we will use an inductive argument and we need to know that, after many changes, we are still in the situation of the basic construction. \diamond

4.5 Creating symmetric normal neighborhoods of the clasp arcs

In [Section 4.3](#) we modified the discs R_1, \dots, R_k which had been constructed as part of the basic construction to very special foliated discs which we called tabs. Thus we now know that each clasp arc γ_+^i is contained in a tab. In this section we study γ_+^i and γ_-^i and arrange that they are contained in much larger neighborhoods which support a canonical foliation. We will call them normal neighborhoods. The modifications will be made with the help of the finger moves of [Section 4.4](#), at the expense of adding new b -arcs (see [Section 3.3](#)). Intuitively, the new b -arcs add pouches to the surface, and so give additional ‘room to move around’. The normal neighborhoods which we will construct will give us choices, which can be made one way or another as it is convenient. For example we will be able to regard a clasp arcs as being on either side of its associated chain of singular leaves and in braid position, or as being contained in a union of leaves and so lying in a union of fibers of \mathbf{H} . Let $\gamma_\epsilon \subset \mathcal{PA}$ be a clasp arc with the property that γ_+ crosses k singular leaves. A neighborhood \mathbf{N}_ϵ of γ_ϵ is a *normal neighborhood* ([Figure 37](#), which gives an example, should be helpful) if the following conditions hold:

- (1) $\mathbf{N}_\epsilon \supset \gamma_\epsilon$. Also \mathbf{N}_ϵ intersects no other clasp arcs.




Figure 37: Normal neighborhoods of γ_- and γ_+ in \mathcal{PA} . The left (resp. middle and right) sketch shows \mathbf{N}_+ (resp. $\mathbf{N}_-, \mathbf{N}_-$). The arc γ_- ends near X_+ in the middle sketch and in the interior of \mathcal{PA} in the right sketch. In all 3 sketches the foliation is defined up to reflection about a vertical axis which joins X_- and X_+ .

- (2) \mathbf{N}_ϵ contains $2k$ disjoint ‘horizontal’ paths $q_1^\epsilon, \dots, q_{2k}^\epsilon$, each contained in the singular leaves of \mathbf{N}_ϵ and each containing two vertices and two singularities. Traversing γ_ϵ , starting at the X_ϵ endpoint, each q_i^ϵ is crossed once. The vertices on q_{2i-1}^ϵ (resp. q_{2i}^ϵ) have sign ϵ (resp. $-\epsilon$). For $1 \leq i \leq k$, the singularities on q_{2i-1}^ϵ have parity $\epsilon\delta_i$ and the singularities on q_{2i}^ϵ have parity $-\epsilon\delta_i$, where $\delta_i = \pm$.
- (3) $\partial\mathbf{N}_\epsilon = \varphi_1 \cup \varphi_2 \cup \varphi_3 \cup \varphi_4$ where:
 - (a) $\varphi_1 \subset X_\epsilon$.
 - (b) φ_2 is a path contained in one arc of type a_ϵ and k singular leaves. It contains k vertices of sign ϵ .
 - (c) φ_3 is transverse to the foliation of \mathcal{PA} .
 - (d) φ_4 is a path contained in k singular leaves and one arc of type b (or type a_ϵ , in the special case when the puncture point on γ_ϵ is near $X_{-\epsilon}$). It contains k vertices of sign $-\epsilon$.
 - (e) Traversing $\varphi_1 \subset X_\epsilon \cap \partial\mathbf{N}_\epsilon$, beginning at the point $\varphi_1 \cap \varphi_4$, we pass through the endpoint of a type $a_\epsilon b$ singular leaf, and end at the X_ϵ endpoint of an a_ϵ leaf.

We say that $(\mathbf{N}_+, \mathbf{N}_-)$ is a *normal neighborhood pair* if each \mathbf{N}_ϵ is a normal neighborhood of γ_ϵ . Our main result in this section will be very important in the detailed work

we will need to do to push X_+ across \mathcal{CA} to X_- , and to prove the MTWS. We stress this by calling in a ‘proposition’.

Proposition 4.5.1 For each clasp arc pair $(\gamma_+, \gamma_-) = (\gamma_+^i, \gamma_-^i)$ we may assume that there is a normal neighborhood pair $(\mathbf{N}_+, \mathbf{N}_-)$. Moreover, within the normal neighborhoods, we may assume that instead of being transverse to the foliation each clasp arc is positioned in a finite union of leaves in its normal neighborhood pair.

Proof The proof of Proposition 4.5.1 will occupy the rest of this section. Our work begins with the tab neighborhoods of $\gamma_+^1, \dots, \gamma_+^k$ which we constructed earlier, when we simplified the discs R_1, \dots, R_k of the basic construction. The tab neighborhood construction told us nothing about the other sheet of \mathcal{CA} , and our first goal is to modify it so that we have related tab neighborhoods of $\gamma_-^1, \dots, \gamma_-^k$. See Figure 38 for an example. After that we will iterate the construction to produce normal neighborhoods.




Figure 38: Tab neighborhood pair in \mathcal{PA}

Let γ be a clasp arc and let γ_+, γ_- , be its preimages in \mathcal{PA} . We say that $(\mathbf{T}_+, \mathbf{T}_-)$ is a *tab neighborhood* of γ in \mathcal{PA} if the following hold:

- (i) $\partial \mathbf{T}_\epsilon = \alpha \cup \beta$ where $\alpha \subset X_\epsilon$ and β is an arc that is transverse to the foliation of \mathcal{PA} .
- (ii) There is a simple path $l_\epsilon \subset \mathbf{T}_\epsilon$, contained in singular leaves. It starts on X_ϵ , contains all k vertices and all k singularities. If the singularities on l_+ have sign $\delta_1, \delta_2, \dots, \delta_k$ on \mathbf{T}_+ , then the singularities on l_- have sign $-\delta_1, -\delta_2, \dots, -\delta_k$.
- (iii) The arc $\gamma_\epsilon \subset \mathbf{T}_\epsilon$ is the only clasp arc which intersects \mathbf{T}_ϵ . It intersects each of the k singular leaves in the induced foliation of \mathbf{T}_ϵ exactly once.

- (iv) Let $e_+e', e'e'', \dots, e' \cdots e_- \subset \gamma_+ \subset \mathbf{T}_+$ be a subdivision of γ_+ into k subarcs such that each subarc crosses one singular leaf. Then the corresponding induced subdivision of γ_- in \mathbf{T}_- given by the immersion \mathcal{CA} also has the property that each subarc crosses one singular leaf.

Lemma 4.5.1 *We may assume that each pair of clasp arcs γ has a tab neighborhood $(\mathbf{T}_+, \mathbf{T}_-)$ for its associated pair $(\gamma_+, \gamma_-) \subset \mathcal{PA}$.*

Proof To prove the lemma we will make repeated use of the first and second finger moves. Consult Figure 39, which shows the changes we will make as they appear on \mathbf{T}_+ , and Figure 40, which depicts the corresponding changes on \mathcal{PA} in a neighborhood of γ_- . After the sequence of changes γ_+ will have moved to the other side of the chain of singular leaves in \mathbf{T}_+ , but the key features will not have changed.




Figure 39: Local changes, as they appear on \mathbf{T}_+

By Lemma 4.3.1 we may assume that γ_+ is already contained in a tab neighborhood \mathbf{T}_+ . Since we have not changed the fact that γ is transverse to fibers of \mathbf{H} , it follows that in a sufficiently small foliated neighborhood of γ_- in \mathcal{PA} the arc γ_- will also be transverse to the leaves of the foliation of \mathcal{PA} . Some of the leaves intersecting γ_- may be singular. Label them l_1, \dots, l_t , as in Figure 40. The first change we introduce is to perform the first finger move on a small subarc of γ which is just below the singular leaf in the end-tilde of \mathbf{T}_+ , pushing γ across the singular leaf, as in the passage from sketch (a) to (b) in Figure 39(a) \rightarrow (b). This will induce a corresponding change in γ_- , as illustrated in Figure 40(a) \rightarrow (b).

We next perform the second finger move as shown in the passage Figure 39(b) \rightarrow 39(c) and the corresponding alteration 40(b) \rightarrow 40(c). It will be helpful to label the two new

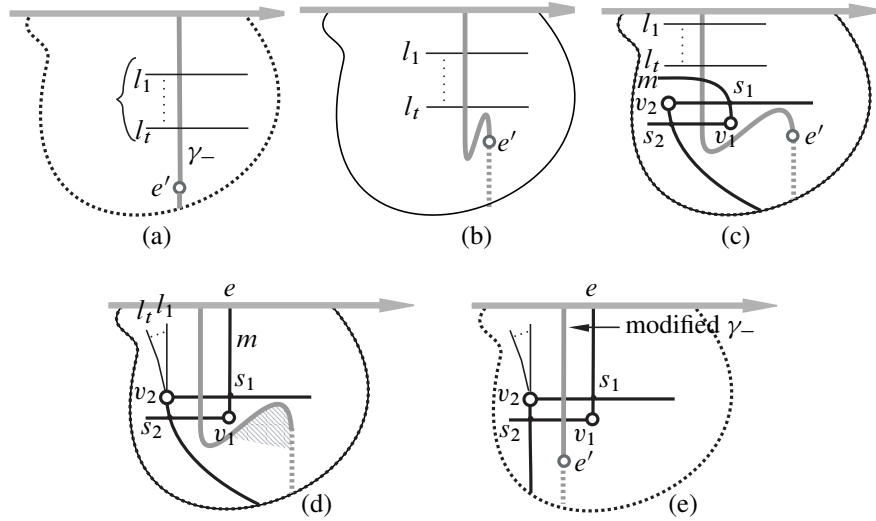


Figure 40: Local changes, depicted in a neighborhood of γ_- on \mathcal{PA}

vertices and singularities introduced by this second finger move as v_1 and v_2 (the vertices) and s_1 and s_2 (the singularities). This second finger move creates a disc Δ in the tab \mathbf{T}_+ whose interior is necessarily embedded because it is on a tab-neighborhood of γ_+ and there is only one clasp arc on each such tab-neighborhood. We use it to do the second change in foliation (see Lemma 3.2.2 and Figure 21). The notation has been chosen so that s_1 in Figure 40(c) and Δ in Figure 39(c) correspond to s_1 and Δ in Figure 21. The singularity s_2 of Figure 21 is not shown in Figure 40(c), also the singularity s_2 of Figure 40(c) is similarly not part of the geometry of Figure 21. After a series of such changes in foliation the singular leaf which is labeled m in Figure 40(c) will have exchanged order with the singularities in the leaves l_1, \dots, l_t as illustrated in Figure 40(d). (Note that this implies that each of the leaves l_1, \dots, l_t has one of its endpoints at v_2). After the change in foliation the clasp arc can be tightened, resulting in the picture we see in Figure 40(e).

From the tab neighborhood definition, condition (ii) on the signs of the vertices follows from the fact that the vertex endpoint of a singular leaf of the foliation of \mathcal{PA} which begins at X_ϵ necessarily has sign $-\epsilon$. But then, the tiles which make up \mathbf{T}_ϵ are all type $a_{-\epsilon}a_{-\epsilon}$ and never type $a_\epsilon a_\epsilon$.

The point e' in Figure 39(d) is defined to be the point where the deformed clasp arc is tangent to a fiber of \mathbf{H} . Figures 39(d) \rightarrow 39(e) \rightarrow 39(f) \rightarrow 39(g) show how the argument can be iterated. If there are k tiles on \mathbf{T}_+ , then after k iterations—finger move 1 followed by finger move 2 followed by the second change in foliation, or as

we shall call it, the *FFF* move, we will have created a tab neighborhood on \mathcal{PA} for γ_- . The subdivision of the tabs described in condition (iv) is achieved automatically via the iteration of the *FFF* sequence. Since we have not moved γ_+ outside of \mathbf{T}_+ and since after the final iteration of our second change in foliation γ_+ will again be transverse to the foliation, we will have created a tab neighborhood pair $(\mathbf{T}_+, \mathbf{T}_-)$ for γ as its associated pair (γ_+, γ_-) .

We refer the reader back to the discussion in Section 4.4, where we introduced the two finger moves, to verify that the parity information of condition (ii) is satisfied. The proof of Lemma 4.5.1 is complete. \square

We are ready to complete the proof of Proposition 4.5.1 By repeating the procedure of Lemma 4.5.1 we can replicate another \mathbf{T}_+ neighborhood inside the \mathbf{T}_+ neighborhoods of Figure 38. This places γ_+ inside a normal neighborhood \mathbf{N}_+ . Since \mathbf{N}_+ contains a tab neighborhood \mathbf{T}_+ which belongs to a tab neighborhood pair, we can interchange the roles of γ_+ and γ_- to produce a corresponding normal neighborhood for γ_- . The main point is that the application of our *FFF* procedure does not move γ_+ outside the tab neighborhood that is nested inside \mathbf{T}_+ . The key properties of the normal neighborhoods follow.




Figure 41: Positioning the clasp arc γ_- into a finite number of leaves (singular and/or non-singular)

We now claim that our procedure for creating $(\mathbf{N}_+, \mathbf{N}_-)$ also creates the conditions necessary for pushing the associated clasp intersection arc into a union of leaves. To

see this let us review our construction of the normal neighborhood pair. See [Figure 38](#). Recall the notation: $e^+e', e'e'', \dots, e' \cdots e^- \subset \gamma_- \subset \mathbf{T}_-$ is a subdivision of the clasp arc into k subarcs such that each subarc crosses one singular leaf. Observe that each of the subarcs $e^+e', e'e'', \dots, e' \cdots e^- \subset \gamma_- \subset \mathbf{T}_-$ has a neighborhood which has the foliation of one of the tiles in [Figure 31](#). So except for a neighborhood around each $e', e'', \dots, e' \cdots e^-$ in γ_ϵ , we can push γ_ϵ into a finite number of leaves. The parity of each of these singular leaves is the same as the parity of the pierce end-point of γ_ϵ . (See the end of [Section 2](#) for the definition of the sign). [Figure 41\(b\)](#) illustrates an example of a clasp arc having a positive pierce end-point that has been partially pushed into singular leaves. In the example which is illustrated the pierce point is positive, and since there is a negative singularity on the ‘vertical’ chain of singular leaves the clasp arc is forced to bend in the manner that is illustrated in sketch (b). [Figure 41\(c\)](#) illustrates that, after a change in fibration which introduces singular leaves that have two singular points ([Figure 41\(a\)](#)) it is possible to push our clasp arc into a finite union of leaves by pushing the remaining arc neighborhoods of the points $e', e', \dots, e' \cdots e^-$ into leaves of the foliation. These two-singularity tiles arise naturally during the change of fibration of [Figure 20](#) in the following way: as we perform the change of fibration in [Figure 20](#) we must pass through a tile having six sides and a two-singularity singular leaf. (This is explained carefully in Birman–Finkelstein [[4](#), [Figure 2.2](#)]). The two singular points will always have common parity. In the example, after the change we can reposition the clasp arc so that it only passes through positive singularities, as in sketch (c). Once we allow the use of such tiles in the foliation of \mathcal{PA} , the final assertion in the proof of [Proposition 4.5.1](#) follows. \square

Remark 4.5.1 When we push the clasp arcs into a union of leaves, we lose some of the symmetry of normal neighborhoods. In this regard we note the following. The symmetry will be needed in one place: for the construction of ‘thin annuli’ in [Section 5.3.1](#). The building blocks for those annuli are the regions which are illustrated in [Figure 43](#). To construct those regions, all we need is a normal neighborhood of that part of a positive clasp arc γ_+^i which begins on X_+ , follows an a_+ arc and along γ_+ to a little bit after it crosses the first singular leaf, together with the corresponding segment on γ_-^i . In such a subset of a normal neighborhood, the required symmetry will always be present, even when it does not extend to a more global symmetry.

5 Pushing across \mathcal{CA}

The machinery has been set up. In this section we learn how to use the foliation of \mathcal{CA} and \mathcal{PA} to push X_+ across the clasp annuli \mathcal{CA} to X_- . That is, we show how we use

the foliation to look inside the black box of Markov towers. Here is an overview of the work in this section:

In [Section 5.1](#) we learn how to produce the sequences (2) and (3) of [Theorem 2](#). Those sequences use the simplest moves of the MTWS, exchange moves and destabilization, to simplify the foliation outside the normal neighborhoods of the clasp arcs, to find the modified braids X'_- and X'_+ . We alert the reader to [Remark 5.1.1](#), where the connection between sequences (1) and (2) of the MTWS and the ‘winding’ phenomenon that was illustrated in [Figure 6](#) is explained.

In [Section 5.2](#) we introduce a new basic move, the ‘microflype’, which generalizes the example of a microflype which we gave earlier in [Section 2.3](#) and [Figure 14](#). In particular, we identify microflype ‘regions’ in the foliated annulus \mathcal{PA} . In [Section 5.3](#) we construct more complicated regions which we call ‘thin annuli’ and learn how to use sequences of microflypes, stabilizations and destabilizations, which can then be amalgamated into admissible flypes, to push X_+ across the thin annuli. At the same time we show how to construct a template that represents the push across a thin annulus. In [Section 5.4](#) we generalize this construction, setting up one of the main tools that will be used in the proof of the MTWS.

In [Section 5.5](#) we describe the foliation that \mathcal{CA} must have so that the Markov tower corresponds to a G–flype, and work out the foliation of \mathcal{CA} and \mathcal{PA} that gave us the G–flype template of [Figure 9](#). In [Section 5.6](#) we describe the foliation \mathcal{CA} and \mathcal{PA} must have so that the Markov tower corresponds to a G–exchange move, working out in detail the explicit foliation that yielded the G–exchange template of [Figure 10](#). Lastly, in [Section 5.7](#) we discuss the isotopy associated with cyclic templates, and in particular the example that we gave earlier in [Figure 8](#).

5.1 The complexity function $c(X_+, X_-, \mathcal{CA})$

By [Proposition 4.5.1](#) we may assume that each clasp arc pair is contained in a normal neighborhood pair, and can be pushed into the union of leaves. Assume from now on that has been done. Define the *complexity* $c(X_+, X_-, \mathcal{CA})$ to be the lexicographically ordered pair (c_1, c_2) , where c_1 is the number of singular points on the clasp arcs and c_2 is the number of singular points that are outside normal neighborhoods. Note that if the clasp arcs are transverse to the fibers of \mathbf{H} , then c_1 can be interpreted as the number of singular leaves crossed by γ_+ and γ_- .

The reader may wonder why we do not include a count of vertices in the foliation? The reason is simple: the vertex count is determined by the singularity count, using the Euler characteristic of the annulus. See the argument at the beginning of the proof of [Lemma 6.3.1](#).

The reader may also wonder why our complexity function ignores all the singularities that are in normal neighborhoods but not on the clasp arcs. The reason is again simple: when we constructed normal neighborhoods we created lots of inessential b -arcs. If we omit them, we will be forced to give up normal neighborhoods, and we don't want to do that now.

Finally, the reader may wonder why we are not including braid index in our complexity function. That is a more subtle matter. At this stage in the work it suffices to say that we will introduce it later (see the *augmented* complexity function of [Section 6](#).)

In the manuscript [\[9\]](#) the authors proved that when \mathcal{X} is the μ -component unlink, exchange moves and destabilization suffice to reduce any closed braid representative X_+ to the identity braid in the μ -strand braid group. In this section we see how far we can go in the simplification of our clasp annulus \mathcal{CA} with the help of exchange moves and destabilization. The moves that we use here will create the two subsequences (1) and (2) of the MTWS. The reason that we had to separate the sequences (1) and (2) in the statement of the MTWS from the sequence (3) is that it may happen that the given braid X_- does not have minimum complexity with respect to exchange moves, for example it may be wound up in the manner illustrated in [Figure 5](#). If so, then if we simply tried to modify the given X_+ to X_- , it might not be possible to do it without increasing complexity at some point.

We shall regard changes in foliation and braid isotopy to be 'trivial moves'. On the other hand, in general exchange moves modify the braid isotopy class and destabilizations change the braid index too. Our goal will be to minimize the complexity, using all four moves: braid isotopy, changes in foliation, exchange moves and destabilizations.

One expects the singular leaves to contain key information about a foliation. In the case of braid foliations more is true, because the singular leaves divide nicely into subsets which are characterized by the signs of their vertices and singularities. Let $G_{\epsilon,\delta}$, where ϵ and δ are \pm , be the set of all singular leaves which pass through only vertices of sign ϵ and singularities of sign δ . We consider the four subgraphs of the graph of the singular leaves $G_{+,+}, G_{-,-}, G_{+,-}, G_{-,+}$. By definitions $G_{\epsilon,\delta} \cap G_{-\epsilon,-\delta} = \emptyset$. See [Figure 42](#), which illustrates how the 4 graphs intersect a bb -tile. Similar graphs appeared in [Bennequin \[2\]](#), in connection with his studies of the characteristic foliation of surfaces bounded by knots.

The intersection of a clasp arc γ_ϵ with a subdisc Δ of \mathcal{PA} is *good* if $(\Delta \cap \gamma_\epsilon) \subset (\Delta \cap G_{\epsilon,\delta})$ for some $(\epsilon, \delta) \in \{(+, +), (+, -), (-, +), (-, -)\}$ and if no component of $\Delta \cap \gamma_\epsilon$ includes the puncture endpoint of γ_ϵ . A subdisc Δ of \mathcal{PA} is *good* if every component of intersection of Δ with a clasp arc γ_ϵ is good. A subdisc that is disjoint




Figure 42: The graph $G_{\epsilon,\delta}$ includes all singular leaves through vertices of sign ϵ and singularities of sign δ . The thick edges in this example illustrate the passage of $G_{+,+}, G_{-,-}, G_{+,-}, G_{-,+}$ through a bb -tile.

from the clasp arcs is, of course, good. A disc that contains a puncture endpoint of a clasp arc is not good.

A vertex is said to be *near* X_ϵ , where $\epsilon = \pm$, if it is the endpoint of a leaf of type a_ϵ in the foliation of \mathcal{PA} . A vertex v is said to be an *interior vertex* if it is not near either X_+ or X_- . In both cases we define $link(v)$ to be the closure of the union of all b -arcs and a_ϵ -arcs which meet v .

To begin our work, we study ab and bb -exchange moves and changes in foliation in the presence of clasp arcs, under the hypothesis that the moves are supported in good discs:

The $(ab)^*$ exchange move Let v be a vertex of valence 2 in the foliation of \mathcal{PA} which is near X_ϵ and is the endpoint of a b -arc. Assume that $\Delta = link(v)$ is a good disc which has non-empty intersection with a clasp arc γ_ϵ . Consult Figure 23(a), which illustrates the case when v is a positive vertex. The disc Δ is the closed disc bounded by the singular leaves wqp and wst in Figure 23(a) and by the subarc pt of X_+ . If $\epsilon = -$ all vertex signs are reversed and the roles of X_+ and X_- are interchanged, but the underlying phenomena are unchanged. We are interested in whether we can push X_+ across Δ in the presence of the clasp arc γ ?

Notice that all 4 graphs $G_{+,+}, G_{+,-}, G_{-,+}$ and $G_{-,-}$ intersect Δ in Figure 23(a). By our definition of a good disc, $\gamma \cap \Delta$ is a subset of some $G_{\epsilon,\delta}$, however the case $\epsilon = +$ is impossible. The reason is: $G_{+,\pm}$ ends at v in the interior of Δ (because the two singularities in Δ have opposite signs), also by hypothesis Δ is a good disc, so that γ cannot have its puncture endpoint in Δ . Therefore $(\gamma_\epsilon \cap \Delta) \subset G_{-,+}$ or $G_{-,-}$. Both cases can occur, and the two cases are essentially the same. In both cases Δ is embedded, because γ_ϵ lies in a normal neighborhood and so its partner $\gamma_{-\epsilon}$ cannot

intersect Δ . In fact normal neighborhoods tell us more: γ_ϵ is the only clasp arc that can intersect Δ . So γ begins at, say, $p \in X_+$ and passes through q to w , always in $\partial\Delta$. Just as in the embedded case, we can push X_+ across a neighborhood of Δ , removing two vertices and two singularities from the foliation. This is the ‘light bulb’ move of Figure 23(b). It shortens γ_ϵ and so also shortens its partner $\gamma_{-\epsilon}$, reducing the entry c_1 in $c(X_+, X_-, \mathcal{CA})$. The second sheet of \mathcal{PA} is unchanged by the move, except for the fact that its clasp arc is shortened by pulling in its puncture endpoint. The link X_+ remains embedded throughout the move. The argument is identical if X_+ is replaced by X_- . This is the move that we call the $(ab)^*$ exchange move. It is exactly the same as the ab -exchange move, but in the presence of clasp arcs and with the assumption of good discs.

The $(bb)^*$ exchange move Let v be an interior vertex of valence 2 in the foliation of \mathcal{PA} . Assume that $\Delta = \text{link}(v)$ is a good disc which has non-empty intersection with a clasp arc γ . Consult Figure 25, which illustrates the case when v is a negative vertex. The disc Δ is the closed disc bounded by the singular leaves w_1sw_2 and w_1qw_2 . All vertex signs could be reversed, it will not matter, so we assume they are as illustrated. We are interested in whether we can do the bb -exchange move and collapse of the pocket, as illustrated in the passage from the left to right sketches in Figure 25, in the presence of the clasp arc γ , when Δ is a good disc?

As in the case of the $(ab)^*$ move, the intersection of our clasp arc with Δ is assumed to be a subset of the intersection of one of the 4 graphs with Δ . Since the puncture point of γ cannot be in Δ the only possibilities are the graphs $G_{+,+}$ or $G_{+,-}$, because the other two possibilities lead to a puncture point in the interior of Δ . Assume without loss of generality that $\gamma \subset G_{+,+}$, for example γ might pass through the points pw_1sw_2p' in the left sketch in Figure 25. Using the hypothesis that γ is in a normal neighborhood, we can (by Proposition 4.5.1) push γ to a new position which we call $pw'_1s'w'_2p'$ where it is transverse to the leaves of the foliation. But then the bb -exchange move goes through as before, as described in full detail in sketches 2.15–2.19 and the accompanying text of Birman–Finkelstein [4]. Since that argument is long and technical, and since a complete reference is available, we do not repeat it here. After the move the ‘pocket may be collapsed’, as illustrated in the passage from the left to the middle sketch in Figure 24 of this paper, and also in the passage from the left to the right sketch in Figure 25. The clasp arc can then be pushed back into a union of leaves, ie to the position pw_2p' in the right sketch. The complexity $c(X_+, X_-, \mathcal{CA})$ has been reduced because the singularity s is no longer on γ . This is the move that we call the $(bb)^*$ exchange move. It is exactly the same as the bb -exchange move, but in the presence of clasp arcs and with the assumption of good discs.

The change of foliation of Lemma 3.2.1 in the presence of clasp arcs As in the case of the ab and bb -exchange moves, we will refer to the manuscript by Birman and Finkelstein [4] for all details, since the required changes in foliation were justified very carefully there, in the proof of Theorem 2.1 of that paper, under the hypothesis that the surface in question was embedded. But the immersed case is really no different from the embedded case because the changes in foliation that we need here can always be thought of as being induced by changes in the choice of disc fibers for the fibration \mathbf{H} of 3-space minus the braid axis, and not by any change in the position of \mathcal{CA} in 3-space. See Figure 20, which illustrates the case when we are interested in changing the foliation on two bb -tiles which meet in a common b -arc. The sketches show the changes when there are no clasp arcs. The fact that the clasp arcs have been positioned in normal neighborhoods shows us that we can change the foliation in one of the sheets without changing the other sheet. Assume that the union of our two bb -tiles is a good disc. There are three cases:

- (i) Both s_1 and s_2 are contained in a single normal neighborhood as consecutive singularities, on a clasp arc. After the change of foliation, the path of the clasp arc contains one fewer singularity, so $c(X_+, X_-, \mathcal{CA})$ is reduced. In particular, if $\epsilon = +$ and we have a symmetric clasp arc then the resulting change of foliation is the one that is labeled (ii) in Figure 20. In this case the clasp arc contains the two right-most positive vertices and the right-most singular point. If $\epsilon = -$, then the resulting change of foliation is the one that is labeled (i) in Figure 20, and the clasp arc contains the two left-most negative vertices and the left-most singular point.
- (ii) One of the two singularities, say s_1 , is contained in a normal neighborhood of a clasp arc, but not on the clasp arc. After the change in foliation, the path of the clasp arc is unchanged. For example, if $\epsilon = +$ and the clasp arc contains s_1 (resp. s_2) then we have a change of foliation that takes us to (i) (resp. (ii)).
- (iii) The two singularities are contained in disjoint normal neighborhoods of different clasp arcs. The path of the two clasp arcs is unchanged by the change in foliation. For example, if $\epsilon = +$ and the bottom singularity was contained in the clasp arc for the beginning tilings in the sequences in Figure 20 then the resulting change of foliation is (ii) in all three cases.

Armed with this knowledge, we are now able to adapt [4, Theorem 3.1] to our clasp annulus \mathcal{CA} and its foliated preimage \mathcal{PA} . We will use the symbol \mathbf{N}_ϵ for a normal neighborhood of the clasp arc γ_ϵ and the symbol \mathcal{N} for the union of all normal neighborhoods of all clasp arcs.

Proposition 5.1.1 *Each of the following holds for all four graphs $G_{+,+}$, $G_{-,-}$, $G_{-,+}$ and $G_{+,-}$:*

- (1) $G_{\epsilon,\delta} \cap G_{-\epsilon,-\delta} = \emptyset$.
- (2) *Every singular point and every vertex in the foliation of \mathcal{PA} is in $G_{+,+}$ or $G_{-,-}$ (and also in $G_{+,-}$ or $G_{-,+}$).*
- (3) $G_{\epsilon,\delta}$ has no interior isolated vertex v .
- (4) $G_{\epsilon,\delta} \setminus (G_{\epsilon,\delta} \cap \mathcal{N})$ has no interior endpoint vertex v .

Assume that all clasp arcs are positioned in normal neighborhoods, and that initially they have been pushed into unions of leaves of the foliation. Then after some number of exchange moves (combined with changes in foliation and isotopies in the complement of the axis) the following holds for all four graphs at once:

- (5) $G_{\epsilon,\delta}$ contains no closed loop l which bounds a good disk $\Delta \subset \mathcal{PA}$.

After some number of exchange moves and destabilizations, the following also holds for all four graphs at once:

- (6) *There is no closed loop l which is the union of an edgepath $\mathcal{E}_1 \subset G_{\epsilon,+}$ and an edgepath $\mathcal{E}_2 \subset G_{\epsilon,-}$ which bounds a good disc $\Delta \subset \mathcal{PA} \setminus \mathcal{N}$.*
- (7) *Let $\mathcal{E}_1, \mathcal{E}_2, \mathcal{E}_3$ be connected arcs, with $\mathcal{E}_1 \subset G_{\epsilon,\epsilon}$, $\mathcal{E}_2 \subset G_{\epsilon,-\epsilon}$, $\mathcal{E}_3 \subset X_\epsilon$. Then there is no closed loop $l = \mathcal{E}_1 \sqcup \mathcal{E}_2 \sqcup \mathcal{E}_3$ or $\mathcal{E}_2 \sqcup \mathcal{E}_3$ or $\mathcal{E}_1 \sqcup \mathcal{E}_3$ which bounds a good disc Δ on \mathcal{PA} .*

Proof The proof is almost identical with the proof of [4, Theorem 3.1]. Outside normal neighborhoods of the clasp arcs it is identical. Inside normal neighborhoods the key concepts which makes it possible to carry over arguments used in [4] are good discs. In the presence of clasp arcs we simply use the exchange moves $(ab)^*$ and $(bb)^*$ instead of ab and bb and the complexity function $c(X_+.X_-, \mathcal{CA})$. An application of $(ab)^*$ can shorten the length of a clasp arc γ_ϵ , reducing the number of singular leaves it intersects. It will necessarily shorten the length of $\gamma_{-\epsilon}$ and decrease the number of tiles of \mathcal{PA} that are in $N_{-\epsilon}$. Thus, there may be some inessential b -arcs that are now away from normal neighborhoods, and they can be eliminated.

Also, an application of $(bb)^*$ can shorten the length of a clasp arc γ_ϵ , reducing the number of singular leaves it intersects. But, it does not immediately reduce the number of singular leaves $\gamma_{-\epsilon}$ intersects. Moreover, neither γ_ϵ nor $\gamma_{-\epsilon}$ may be in normal neighborhoods anymore. To re-establish the symmetry between the the lengths of the two clasp arcs in \mathcal{PA} we take a tab neighborhood around γ_ϵ (since a tab neighborhood will have the shortest clasp length) and re-apply the normal neighborhood construction of Proposition 4.5.1. □

Our first corollary produces the modified braids X'_+ and X'_- in the statement of the MTWS. In the course of changing X_{\pm} to X'_{\pm} , it also produces the sequences (1) and (2) of the MTWS, which modify $X_{\pm} \rightarrow X'_{\pm}$:

Corollary 5.1.1 *Assume that all changes are made in the complement of normal neighborhoods of the clasp arcs. Then with the help of:*

- complexity-reducing exchange moves on X_+ and X_- ,
- complexity-reducing destabilizations of X_+ ,
- changes in foliation (which modify \mathcal{CA} but do not change X_+ or X_-),

the triplet (X_+, X_-, \mathcal{CA}) may be changed to a new triplet $(X'_+, X'_-, \mathcal{CA}')$ which has minimum complexity, and for which the foliation of \mathcal{PA} has the following properties:

- (1) *There are no $a_{\epsilon} s$ or $a_{\epsilon} a_{\epsilon}$ singularities in \mathcal{PA} for $\epsilon = +$ or $-$.*
- (2) *If α is a path in $G_{\epsilon, \delta}$ which begins and ends on X_{ϵ} , and Δ is the disc on \mathcal{PA} which is split off by α and a subarc of X_{ϵ} , then $\text{int}(\Delta)$ contains a puncture end-point of a $\gamma_{-\epsilon}$ clasp arc.*
- (3) *If $\alpha \subset G_{\epsilon, \delta}$ is a loop which bounds a disc Δ in \mathcal{PA} , then $\text{int}(\Delta)$ contains at least two puncture end-points of clasp arcs, one from some γ_{ϵ}^i and the other from $\gamma_{-\epsilon}^i$.*
- (4) *If $v \in G_{\epsilon, \delta}$ is an endpoint vertex, then either $\text{link}(v)$ contains a puncture endpoint of a clasp arc or v lies in a normal neighborhood of a clasp arc.*

Proof In the case of an embedded surface, details on how to pass from the analogue of [Proposition 5.1.1](#) to that of the Corollary are given in the proof of [\[4, Lemma 3.8\]](#). We have modified the statement of [\[4, Lemma 3.8\]](#), by saying that if a disc Δ exists with the stated undesirable properties, then Δ is not a good disk. In the case of assertion (3) of the Corollary we need the puncture endpoints of both γ_+^i and γ_-^i because if both are not in Δ then Δ will be embedded, but no such embedded disc exists. \square

Remark 5.1.1 Armed with [Corollary 5.1.1](#), the reader is advised to go back and look at [Figure 6](#) again. It was intuitively clear, when we first presented this sketch, that the unwinding process which is illustrated in the passage from left to right ought to be ‘complexity-reducing’. Now that we have finally defined an appropriate complexity function we can understand exactly what that means.

We need the modifications of [Corollary 5.1.1](#) for the following reason: If it should happen that the given braid X_- is wound up some number of times as on the left, and if

we were to modify X_+ but leave X_- unaltered, then we might be able to use exchange moves and destabilizations to modify X_+ to a minimum braid index representative, using only complexity-decreasing moves, but if we did not modify X_- too then we would have to increase complexity at the end. That is the reason that we need the initial complexity-reducing modifications of *both* X_+ and X_- . That is the content of [Corollary 5.1.1](#).

An annulus component of \mathcal{CA} is *standard* if every component $E \subset G_{\epsilon, \delta}$ on the annulus satisfies the following conditions:

- (1) E is homeomorphic to either S^1 or $[0, 1]$, ie either a circle or a line.
- (2) If E is a circle then it is a core circle of the annulus component.
- (3) If E is a line with $\partial E = p \cup p'$ then $p \in X_\epsilon$ and p' is near $X_{-\epsilon}$.

An annulus component of \mathcal{CA} is *trivially foliated* if it is foliated entirely by s -arcs. Examples of a standard annulus, and of a trivially foliated annulus, are given in [Figure 43](#).

Our second Corollary relates to modifications which may be needed after all the clasp arcs have been shortened, that is $c_1 = 0$.

Corollary 5.1.2 *Let (X_-, X_+, \mathcal{CA}) be of minimal complexity with $c_1 = 0$. Then each component of \mathcal{PA} is either an annulus which is foliated entirely with s -arcs or a standard annulus.*

Proof The statement follows directly from [Corollary 5.1.1](#), when we add the assumption that $c_1 = 0$. □

Remark 5.1.2 The problem of pushing X_+ across such a standard annulus will be treated in [Section 5.7](#).

Remark 5.1.3 We know exactly how the foliation of \mathcal{PA} looks inside a normal neighborhood of a clasp arc. We now ask what can happen outside the union of normal neighborhoods, when the puncture points are in the interior of \mathcal{PA} , as in the left and right sketches in [Figure 37](#). An example is illustrated in [Figure 44](#). The shaded regions are unions of overlapping normal neighborhoods of clasp arcs. The clasp arcs in question are negative. We show only one of them, to avoid clutter and enable us to focus on the features which are of interest now. It is to be compared with the right sketch in [Figure 37](#). The boundaries of the region we are studying are indicated with dotted black and grey arcs transverse to the foliation because in general they are in the interior of \mathcal{PA} (although in special cases there is no reason why they could not be X_+ and X_-). This sketch will be used in [Section 6](#) in the proof of the finiteness of $\mathcal{T}(m)$. ◇

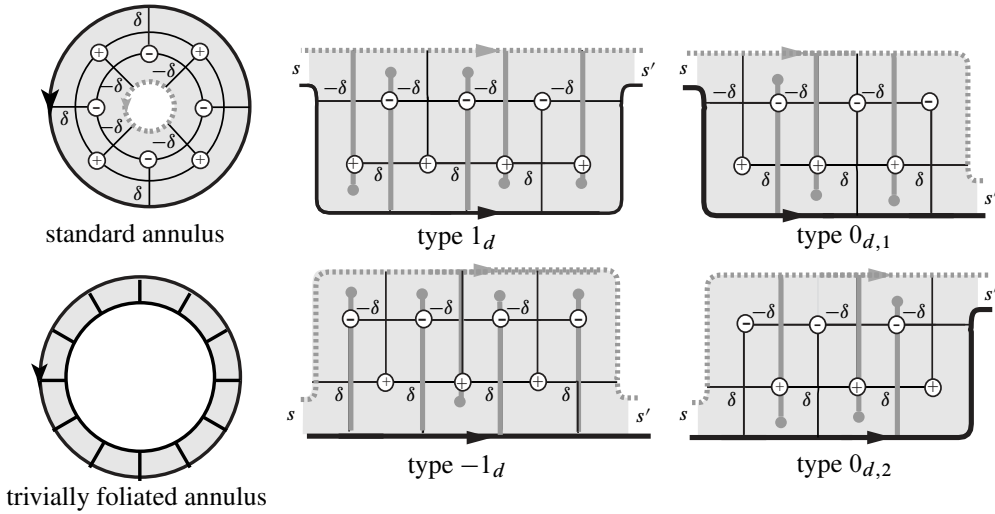


Figure 43: Examples of a standard annulus, a trivially foliated annulus, and regions in a thin annulus. Thin annuli will be defined in Section 5.3.




Figure 44: The foliation of $\mathcal{P}\mathcal{A}$, near the puncture point of a clasp arc γ_- but outside its normal neighborhood. It is assumed here that the puncture point of γ_- is in the interior of $\mathcal{P}\mathcal{A}$.

5.2 Pushing across a microtype region

Having simplified the foliation of $\mathcal{C}\mathcal{A}$ as much as possible with the use of exchange moves and destabilizations, we begin to investigate new features of the foliation of $\mathcal{C}\mathcal{A}$.

A clasp arc γ in $\mathcal{C}\mathcal{A}$ has two preimages in $\mathcal{P}\mathcal{A}$, namely γ_+ and γ_- . Recall that γ_ϵ , $\epsilon = \pm$, begins on X_ϵ and ends at an interior point of $\mathcal{P}\mathcal{A}$. Recall also that γ_ϵ was defined to be short if it does not intersect any singular leaves. We now define γ_ϵ to be long if the puncture endpoint of γ_ϵ is on an $a_{-\epsilon}$ -arc. We say that γ is doubly long if both γ_+ and γ_- are long. An example of a doubly long clasp arc was given in Figure 14. On the other hand, in Figure 56 there are 4 clasp arcs, and among the eight preimages $\gamma_-^1, \gamma_-^2, \gamma_+^3$ and γ_+^4 are long, but their partners $\gamma_+^1, \gamma_+^2, \gamma_-^3$ and γ_-^4 are not,

so that there are no doubly long clasp arcs. (Remark: the foliated annulus in [Figure 56](#) gave rise to the 6–braid G–flype template which we described in the introduction to this paper, in [Figure 9](#)).

A clasp arc γ_ϵ in \mathcal{PA} is *intermediate* if it is neither short nor long. This implies that its puncture endpoint is on a b –arc. Look ahead to [Figure 57](#) for examples. Eight clasp arcs are depicted there, and the 16 preimages are all intermediate. (Remark: the foliated annulus [Figure 57](#) gave rise to the 6–braid G–exchange template which we described in the introduction to this paper, in [Figure 10](#)). The *length* of the clasp arc γ_+ is the number of singular leaves which γ_+ crosses. We will not be concerned with the length of γ_- . Note that if long clasp arcs occur, then $c_1 > 0$.

In this subsection we consider the case when there is exactly one clasp arc γ and it is doubly long, so that $c(X_+, X_-, \mathcal{CA}) = (1, 0)$. Since $c_2 = 0$, we may assume that outside normal neighborhoods of γ_+ and γ_- , the foliation of \mathcal{PA} consists entirely of bands which are foliated by s –arcs. But then, since there are no other clasp arcs in \mathcal{PA} , we may just as well simplify the normal neighborhoods by the deletion of inessential b –arcs. Thus \mathcal{PA} is a union of an a_+s –tile and an a_-s –tile, joined up by bands of s –arcs. This very simple foliated annulus was introduced long ago, as the ‘key example’ at the end of [Section 2](#), in [Figure 14](#). Understanding microflypes will allow us, later, to consider very much more complicated annuli \mathcal{PA} . Recall that microflypes are braid-index preserving moves which replace very simple Markov towers, as can be seen from [Figure 5](#), so that they might be expected to be basic to our work.

Of course, when we first encountered the microflype region in \mathcal{PA} in [Section 2](#), as an example of the basic construction, we did not have available to us the machinery of braid foliations. By our work in [Section 4.5](#) we know that the signs of the singularities in the two tiles are opposite.

The leaves in the braid foliation of \mathcal{PA} are level sets for the embedding of the two tiles which make up \mathcal{PA} . Referring back to [Figure 14](#), it should now be completely clear that the bottom row of sketches represents the immersion in 3–space of the two foliated tiles in the top row of sketches, and that the motion of X_+ across the two illustrated discs is indeed realized by a flype which takes X_+ to a new position which is separated from X_- by a band of s –arcs. The associated block strand diagram is clearly that for a flype, which in the case which is illustrated is a negative flype. The braid inside the braid block is a single full twist of two strands. The sign of the full twist that is illustrated is negative.

Summarizing: A *microflype region* in \mathcal{PA} is a subset of \mathcal{PA} which is a union of an a_+s –tile and an a_-s –tile, intersected by paired clasp arcs γ_+ and γ_- , each of which intersects a single singular leaf. A *microflype* on a block-strand diagram is a flype in

which the braid in the braid block R consists of exactly one full twist of either sign on two strands. Such a braid block R will be called a *microblock*. All strands have weight 1. The *sign of a microflype* is the pair (α, β) where α is sign of the half-twist which is outside the braid block and β is the sign of the full twist which is inside the braid block. In [Section 5.3](#), below, we will show that microflypes with their associated microblocks are the building blocks of the most general flypes, with arbitrary braids in the braid block and arbitrary weights on the strands.

By the construction in [Proposition 4.5.1](#), it is possible to push the clasp arcs into the associated singular leaves, giving a more symmetric embedding in 3-space.

5.3 Pushing across thin annuli

While we have been able, up to now, to make the tacit assumption that we are working with knots, we now return to the general case of links.

Our task in this section is to learn how to use flypes to shorten the length of long (but not necessarily doubly long) clasp arcs in the foliation of \mathcal{PA} , thereby reducing the integer c_1 in the complexity pair (c_1, c_2) . We have already seen that in the situation where there is exactly one clasp arc, we may use a microflype to push X_+ across the clasp annulus \mathcal{CA} . The situation which we face now has two factors which make it significantly more complicated. The first is that we must allow for the possibility that there are $k \geq 1$ clasp arcs. The second is that if X_+ has μ components, then \mathcal{CA} will be the image of μ annuli under an immersion, and we must allow for the possibility of clasp intersections between distinct annuli.

5.3.1 Constructing the thin annuli Preparing for the shortening of long clasp arcs, we will construct a family \mathcal{S} of ‘thin annuli’ which is a subset of \mathcal{PA} . Normal neighborhoods will play a key role in the construction. We will prove:

Lemma 5.3.1 *Assume that \mathcal{X} has μ components. Suppose that \mathcal{PA} contains long clasp arcs $\gamma_+^i, \gamma_+^j, i \neq j$. Then there exists a family \mathcal{S} of μ annuli, each a subannulus of \mathcal{PA} , and each with a component of X_+ as one of its boundaries, such that every annulus in \mathcal{S} is either trivially foliated by s -arcs, or is a standard annulus, or has a foliation satisfying the following (see [Figure 43](#) for examples):*

- (i) *Each non-trivially foliated annulus contains at least one long clasp arc. Moreover, all of its clasp arcs are doubly long and of length 1, with respect to the induced foliation of \mathcal{S} .*
- (ii) *After an isotopy of \mathcal{PA} which leaves X_+ and the other components of $\partial\mathcal{S}$ fixed, each b -arc in \mathcal{S} may be assumed to have at least one of its endpoints on a clasp arc.*

Proof The case when there is exactly one long clasp arc was just discussed in Section 5.2, so we assume that there are at least two long clasp arcs. Let γ_+^i be a long clasp arc, so that γ_-^i begins on X_- and ends near X_+ . Assume that γ_+^i and γ_-^i have been pushed into a union of leaves in their normal neighborhoods. This is possible, by the construction in Proposition 4.5.1. Let \mathbf{N}_-^i be the normal neighborhood of γ_-^i . We focus our attention on a rectangle which we call $\bar{\mathbf{N}}_-^i$. See sketches (1) and (2) in Figure 45. It is a subset of \mathbf{N}_-^i , and it coincides with \mathbf{N}_-^i when the clasp arc has length 2. It has three edges which are in $\partial\mathbf{N}_-^i$. We choose an arc which is transverse to the foliation as its other ‘horizontal’ boundary.




Figure 45: Sketches (1) and (2) illustrate the two possibilities for $\bar{\mathbf{N}}_-^i$. Sketches (3) and (4) illustrate the two possibilities for $\bar{\mathbf{N}}_+^i$.

There are corresponding subrectangles $\bar{\mathbf{N}}_+^i$ and we show the two possible arrangements in sketches (3) and (4) of Figure 45. As in the case of $\bar{\mathbf{N}}_-^i$, the rectangle $\bar{\mathbf{N}}_+^i$ has 3 boundary edges which are in $\partial\mathbf{N}_+^i$. We choose its 4th boundary edge to be an arc which is transverse to the foliation, so that the rectangle has 4 vertices and 4 singularities.

In the special case when γ_+^i has length 2 the puncture endpoint of the arc γ_+^i will be an interior point of $\bar{\mathbf{N}}_+^i$, but if γ_+^i has length ≥ 3 then $\gamma_+^i \cap \bar{\mathbf{N}}_+^i$ will have both of its endpoints on the boundary. We correct this by modifying $\gamma_+^i \cap \bar{\mathbf{N}}_+^i$ to a subarc of γ_+^i which has the same image in \mathcal{CA} as the intersection of γ_-^i with its induced normal neighborhood. This will give us a shortened induced arc which (by an abuse of notation) we continue to refer to as γ_+ . It begins on X_+ and ends at a point in the interior of $\bar{\mathbf{N}}_+^i$. By construction, γ_-^i and the new γ_+ have the same image in \mathcal{CA} , and determine the clasp intersection between the rectangles $\bar{\mathbf{N}}_-^i$ and $\bar{\mathbf{N}}_+^i$ which is induced by the clasp intersection in \mathcal{CA} corresponding to γ^i . The four cases which are illustrated in Figure 45 will be referred to as types (1),(2),(3),(4).

Let \mathcal{N}_S be the union of all of the $\bar{\mathbf{N}}_-^i$ and $\bar{\mathbf{N}}_+^i$. Note that, while $\bar{\mathbf{N}}_\pm^i \cap (\gamma_-^j \cup \gamma_+^j) = \emptyset$ for all $j \neq i$, there is no reason why $\bar{\mathbf{N}}_\pm^i \cap \bar{\mathbf{N}}_\pm^j$ should be empty. This leads us to the following preliminary definition of a connected collection of normal neighborhoods: Choose $\bar{\mathbf{N}}, \bar{\mathbf{N}}' \in \mathcal{N}_S$. We say that $\bar{\mathbf{N}}$ and $\bar{\mathbf{N}}'$ are *connected*, and write $\bar{\mathbf{N}} \longleftrightarrow \bar{\mathbf{N}}'$, if $\bar{\mathbf{N}} \cap \bar{\mathbf{N}}' \neq \emptyset$ in \mathcal{PA} . Two examples are given in Figure 46. In the top row, $\bar{\mathbf{N}}$ has

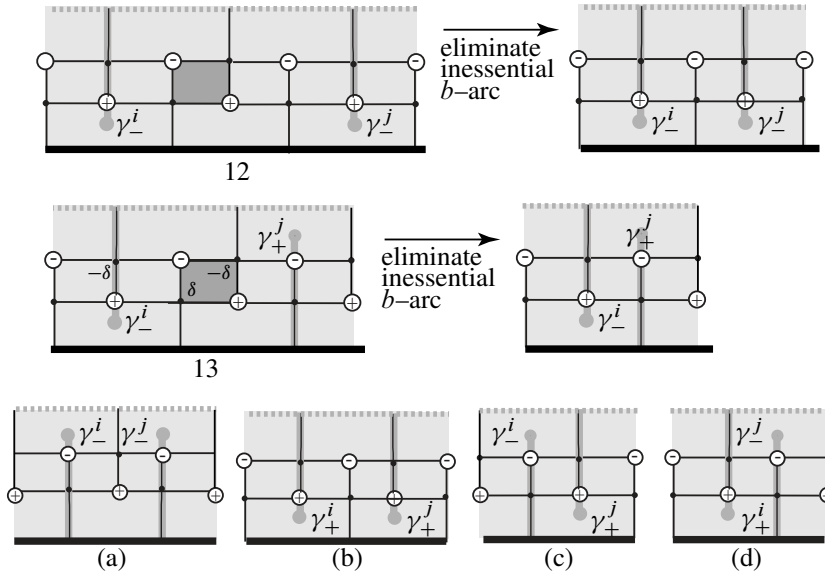


Figure 46: Adjacent pairs of foliated connected normal neighborhoods. Sketches a,b,c,d in the bottom row show all possibilities, after the elimination of inessential b -arcs.

type 1 and \bar{N}' has type 2 and $\bar{N} \cap \bar{N}' = \partial \bar{N} \cap \partial \bar{N}'$. In the middle row \bar{N} has type 1, \bar{N}' has type 3 and they intersect along two singular leaves and the disc between them. Observe that this is the maximal possible intersection, because the normal neighborhood of a clasp arc never intersects another clasp arc. It follows that the possible sequences in a connected set are 12, 13, 14, 23, 24, 34, 21, 31, 41, 32, 42, 43. A collection of normal neighborhoods $\{\bar{N}^1, \dots, \bar{N}^p\} \subset \mathcal{N}_S$ is connected if there is a connecting path between any two neighborhoods in the collection, ie if $\bar{N}', \bar{N}'' \in \{\bar{N}^1, \dots, \bar{N}^p\}$ then there exists a subcollection $\{\bar{N}^{i_1}, \dots, \bar{N}^{i_a}\} \subset \{\bar{N}^1, \dots, \bar{N}^p\}$ such that $\bar{N}' \longleftrightarrow \bar{N}^{i_1} \longleftrightarrow \bar{N}^{i_2} \longleftrightarrow \dots \longleftrightarrow \bar{N}^{i_a} \longleftrightarrow \bar{N}''$. A connected component of \mathcal{N}_S is called a *region* and is denoted by the symbol \mathcal{R} .

We now observe that each b -arc in each \mathcal{R} is in a normal neighborhood of some clasp arc. We distinguish between two types of b -arcs: those whose endpoints are vertices which do not meet a clasp arc, and those which have at least one vertex endpoint which is on a clasp arc. Let's look first at the former. Examples can be seen in the darkly shaded subrectangles in the left sketches in rows 1 and 2 of Figure 46. Recall that our normal neighborhoods were created by the repeated use of finger moves, which necessarily created some inessential b -arcs. But all of the b -arcs which do not intersect clasp arcs are inessential, and may be deleted by an isotopy of $\mathcal{P}\mathcal{A}$ which is supported on a disc in the interior of the connected region, as in the passage from the left to the

right in Figure 46. Therefore we may assume that \mathcal{R} contains no such b -arcs. Notice that we have chosen our definition of complexity so that this modification does not alter the complexity. With the modification, it is easy to see that there are precisely 4 possible sequences of two modified normal neighborhoods, as illustrated in the bottom row in Figure 46, sketches (a),(b),(c),(d).

A region \mathcal{R} is either an annulus or a rectangle, as illustrated in Figure 43. If it is an annulus, then it satisfies properties (i)–(ii) of Lemma 5.3.1. Assume it's a rectangle. The lower horizontal boundary of \mathcal{R} is a subarc of a component of X_+ , and so the connected components $\mathcal{R}_1, \dots, \mathcal{R}_q$ associated to any given component of X_+ have a natural cyclic order on X_+ . We would like to use this natural order to join them by bands of s -arcs to obtain annuli. However the vertical edges of the rectangles are not s -arcs. The following observation saves the day: the grey dotted horizontal boundary of each rectangle was chosen in a rather arbitrary way as an arc which is transverse to the foliation and in the interior of \mathcal{PA} , and if we now modify our choices by replacing the vertical edges of the rectangles in (a),(b),(c),(d) by the sketches in (e),(f),(g),(h) of Figure 47 we will be in business.

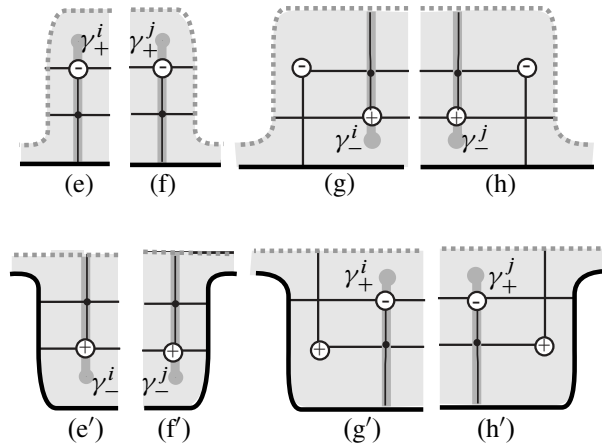


Figure 47: Possible arrangements of clasp arcs in adjacent intersecting normal neighborhoods in \mathcal{PA} , near the right and left boundaries of a connected collection of normal neighborhoods

Sketches (e) and (f) are obtained from sketch (a) of Figure 47 by modifying the grey boundary on the left and right respectively. Note that the modified grey boundary is everywhere transverse to the leaves of the foliation. On the other hand, if we attempt to do the same thing in the situation of (b), choosing the grey boundary to be close to the left (resp. right) clasp arc, a point of tangency with leaves of the foliation will

be introduced, so it is necessary to include the singular leaf which is on the left (resp. right) in the modified connected region, as illustrated in (g) (resp. (h)). We leave it to the reader to check that (e) and (h) are modifications of (c), and that (f) and (g) are modifications of (d). In (e)–(h), which are based directly upon the sketches in [Figure 45](#), all have their rightmost singular leaf in $G_{-, \delta}$. But there is a second possibility: we are given the foliation near X_+ , and we do not know much about it outside the normal neighborhoods. In fact there are 4 other possible patterns for the right and left boundaries, also restricted by the fact that X_+ must be transverse to fibers of \mathbf{H} . They are illustrated in sketches (e'), (f'), (g'), (h') of [Figure 47](#), and were obtained from (e), (f), (g), (h) by interchanging the roles of grey and black.

Let's examine the possibilities for the regions. [Figure 43](#) shows the foliated annulus and the four possible foliated disc regions in \mathcal{S} , up to the number of pairs of components of $G_{-, \delta}$ and $G_{+, -\delta}$. In each, we give examples of how the clasp arcs might be placed. To construct \mathcal{S} from the various non-trivially foliated regions we first join disc regions which are consecutive as one travels along a component of X_+ with bands of s -arcs which run between X_+ and X_- . We may also need some number of standard annuli which have long clasp arcs. (See [Figure 43](#) again.) Finally we may need some number of annuli which are foliated entirely by s -arcs. This completes the proof of [Lemma 5.3.1](#). \square

Let \mathcal{S} be an annular subset of \mathcal{PA} which contains X_+ (or a component of X_+ in the case when \mathcal{X} is a link) as one of its boundary components. We consider the foliation of \mathcal{S} which is induced by the foliation on \mathcal{PA} . A (possibly empty) family of s -arcs $\mathbf{S} = \{s_1, \dots, s_l : s_i \subset \mathcal{S}\}$ is a *complete collection of s -arcs* in \mathcal{S} if (i) no two s -arcs in the collection split off a sub-band of \mathcal{S} that is foliated entirely by s -arcs, and (ii) for any other s -arc $s \subset \mathcal{PA}$ there exists an $s_i \in \mathbf{S}$ such that $s \cup s_i$ splits off a sub-band of \mathcal{S} that is foliated entirely by s -arcs. It is immediate that cutting \mathcal{S} open along a complete collection \mathbf{S} of s -arcs decomposes \mathcal{S} into a disjoint union of *thin regions* and *bands of s -arcs*. This construction will be used in what follows.

The grey boundary of the thin annuli will in general be in the interior of \mathcal{PA} , although in special cases it will coincide with X_- . Observe that the difference in braid index $b(X_+) - b(X_-)$ across the union of the annuli which make up \mathcal{S} is the total number of positive vertices in \mathcal{S} minus the total number of negative vertices in \mathcal{S} . From [Figure 43](#) we see that this difference is always either 0, 1 or -1 in a single non-trivially foliated region. Thus the motion of X_+ across \mathcal{S} to X_- increases braid index if and only if there are more regions of type -1 than of type $+1$ in \mathcal{S} . The regions of type 0 do not affect the count. We call our regions types $1_d, -1_d, 0_{d,1}$ and $0_{d,2}$, the subscript indicating that the region is a disc. (The two cases of type 0_d are distinguished by the

placement of their clasp arcs.) There is also the special case of the standard annulus. The standard annulus first appeared in [Corollary 5.1.2](#) as an embedded annulus. We are now allowing for the occurrence of clasp arcs.

The previous construction of a thin annular subset \mathcal{S} in \mathcal{PA} has X_+ as one of its boundary components. But, we can also use the same construction to produce a thin annular subset which has X_- as one of its boundary components. Let \mathcal{S}_+ (resp. \mathcal{S}_-) be the thin annular subset of \mathcal{PA} having X_+ (resp. X_-) as one of its boundary components. Notice that \mathcal{S}_- , like \mathcal{S}_+ , will have type $0_d, 1_d, -1_d$ regions that are connected by s -bands. However, the dotted grey boundary will be replaced by a solid grey boundary since it will now be X_- ; and the solid black boundary will be replaced by a dotted black boundary since it may or may not be X_+ .

Referring to the regions in [Figure 43](#) we consider the cyclic ordering of the singularities in \mathbf{H} . We say a type $0_{d,1}$ or 1_d (resp. $0_{d,2}$ or -1_d) region is a *fan* if all of its singularities of parity $-\delta$ (resp. δ) occur in sequence in the fibration, followed by all of its singularities of parity δ (resp. $-\delta$).

We now have a proposition that allows us to use $b(X_+)$ and $b(X_-)$ to limit the occurrences of these regions in \mathcal{S}_+ and \mathcal{S}_- . It will play a key role, later, when we prove the finiteness of the set of templates. (Look ahead to part (6) of [Lemma 6.3.2](#) in [Section 6](#).)

Let $\rho: \mathcal{PA} \rightarrow \mathcal{CA}$ be the natural immersion.

Proposition 5.3.1 *The annular regions $\mathcal{S}_+, \mathcal{S}_- \subset \mathcal{PA}$ contribute to $b(X_+)$ and $b(X_-)$ in the following ways.*

- (1) *The number of type 1_d regions in \mathcal{S}_+ is bounded by $b(X_+)$.*
- (1') *The number of type -1_d regions in \mathcal{S}_- is bounded by $b(X_-)$.*
- (2) *If $\mathcal{R}_1 \subset \mathcal{S}_+$ is a type $0_{d,1}$ region and $\mathcal{R}_2 \subset \mathcal{S}_+$ is a type $0_{d,2}$ region with $\rho(\mathcal{R}_1) \cap \rho(\mathcal{R}_2) \neq \emptyset$, then the pair $(\mathcal{R}_1, \mathcal{R}_2)$ contributes at least $+1$ to $b(X_+)$.*
- (2') *If $\mathcal{R}_1 \subset \mathcal{S}_-$ is a type $0_{d,1}$ region and $\mathcal{R}_2 \subset \mathcal{S}_-$ is a type $0_{d,2}$ region with $\rho(\mathcal{R}_1) \cap \rho(\mathcal{R}_2) \neq \emptyset$, then the pair $(\mathcal{R}_1, \mathcal{R}_2)$ contributes at least $+1$ to $b(X_-)$.*
- (3) *Suppose $\mathcal{R}_1, \mathcal{R}_2 \subset \mathcal{S}_+$ are regions of type $0_{d,1}$ and $0_{d,2}$ and that $\mathcal{R}_3 \subset \mathcal{S}_+$ is a region of any type that is a fan. Assume that*
 - (a) $\rho(\mathcal{R}_1) \cap \rho(\mathcal{R}_3) \neq \emptyset$,
 - (b) $\rho(\mathcal{R}_2) \cap \rho(\mathcal{R}_3) \neq \emptyset$,
 - (c) \mathcal{R}_1 and \mathcal{R}_2 are adjacent to a common s -band \mathcal{R}_s .

Then the triple $(\mathcal{R}_1, \mathcal{R}_2, \mathcal{R}_s)$ contributes at least $+1$ to $b(X_+)$.

(3') Suppose $\mathcal{R}_1, \mathcal{R}_2 \subset \mathcal{S}_-$ are regions of type $0_{d,1}$ and $0_{d,2}$ and that $\mathcal{R}_3 \subset \mathcal{S}_-$ is a region of any type that is a fan. Assume that

- (a) $\rho(\mathcal{R}_1) \cap \rho(\mathcal{R}_3) \neq \emptyset$,
- (b) $\rho(\mathcal{R}_2) \cap \rho(\mathcal{R}_3) \neq \emptyset$,
- (c) \mathcal{R}_1 and \mathcal{R}_2 are adjacent to a common s -band \mathcal{R}_s .

Then the triple $(\mathcal{R}_1, \mathcal{R}_2, \mathcal{R}_s)$ contributes at least $+1$ to $b(X_-)$.

Proof of (1) Consider a type 1_d region. Its boundary is a union of two s -arcs, labeled s and s' , and subarcs x_+, x_- of X_+, X_- . Let $H_\theta, H_{\theta'}$ be the fibers which contain s, s' . Then x_+ (resp. x_-) begins on H_θ , travels around the braid axis some number of times, staying transverse to fibers, and ends at $H_{\theta'}$. Since there are more positive than negative vertices in a type 1_d region, it follows that x_+ makes at least one more circuit about the braid axis than x_- . In particular, x_+ makes at least one full circuit about \mathbf{A} and so contributes at least 1 to $b(X_+)$.

Proof of (2) Since $\rho(\mathcal{R}_1)$ and $\rho(\mathcal{R}_2)$ intersect along a common clasp arc, we can cut open their preimages in the thin annulus subset S of \mathcal{PA} along the two components of the clasp arc. Paste them together in the unique new way to obtain two new regions which resemble type 1_d and -1_d regions in S . Observe that, while the new regions look like they are type $\pm 1_d$, their images in \mathcal{CA} will have the property: the X_ϵ boundary, $\epsilon = \pm$, say x_ϵ will be a union of two arcs, $x_{\epsilon,1} \cup x_{\epsilon,2}$, where the final endpoint of $x_{\epsilon,1}$ and the initial endpoint of $x_{\epsilon,2}$ are on the same fiber, but are at different points of that fiber. We call this a 'surgery'. After the surgery, we may use the same argument that we used to prove (1), to show that the angular support of $x_{+,1} \cup x_{+,2}$ is 2π greater than the angular support of $x_{-,1} \cup x_{-,2}$. Therefore $x_{+,1} \cup x_{+,2}$ contributes at least 1 to $b(X_+)$.

Proof of (3) We have the following list of possibilities:

Case (i) \mathcal{R}_1 and \mathcal{R}_2 have different types, say \mathcal{R}_1 is type $0_{d,1}$ and \mathcal{R}_2 is type $0_{d,2}$.

- (ia) As we traverse X_+ on \mathcal{R}_3 we encounter an a_+ -arc that is adjacent to a clasp arc that \mathcal{R}_3 shares with \mathcal{R}_1 before we encounter an a_+ -arc that is adjacent to a clasp arc that \mathcal{R}_3 shares with \mathcal{R}_2 .

We surger our three regions $\mathcal{R}_1, \mathcal{R}_2$ and \mathcal{R}_3 , in \mathcal{PA} , along the pre-images of the clasp arcs in $\rho(\mathcal{R}_1) \cap \rho(\mathcal{R}_3)$ and $\rho(\mathcal{R}_2) \cap \rho(\mathcal{R}_3)$. Since the angular support for \mathcal{R}_s is either contained within the angular support of $X_+ \cap \mathcal{R}_3$, or is outside the angular support of $X_+ \cap \mathcal{R}_3$, the angular length of $X_+ \cap (\mathcal{R}_1 \cup \mathcal{R}_s \cup \mathcal{R}_2)$ must be greater than 2π . (Conceptually, \mathcal{R}_3 creates an 'adequate amount of spacing' between \mathcal{R}_1 and \mathcal{R}_2 so that the triple $(\mathcal{R}_1, \mathcal{R}_2, \mathcal{R}_s)$ contributes $+1$ to $b(X_+)$.)

- (ib) As we traverse X_+ on \mathcal{R}_3 we encounter an a_+ -arc that is adjacent to a clasp arc that \mathcal{R}_3 shares with \mathcal{R}_2 before we encounter an a_+ -arc that is adjacent to a clasp arc that \mathcal{R}_3 shares with \mathcal{R}_1 .

We need to consider the positioning of the s -band, \mathcal{R}_s , that runs from \mathcal{R}_1 to \mathcal{R}_2 . If the angular length of \mathcal{R}_s is greater than 2π then we are done. So assume that the angular length of \mathcal{R}_s is less than 2π . Next, we can assume that the angular length of $X_+ \cap \mathcal{R}_3$ is also less than 2π .

Now, a convenient way of accounting for braid index is to line up the foliations of \mathcal{R}_1 , \mathcal{R}_2 and \mathcal{R}_3 (which are still thought of as being in \mathcal{PA}) so that corresponding clasp arcs line up. (For examples the reader should jump ahead to Figures 48 and 50). Since our argument is concerned with the measurement of angular length we need only focus on where we start and end this measurement. It is convenient to use the endpoint of a clasp arc in \mathcal{R}_1 (or \mathcal{R}_2) as our starting point and the endpoint a clasp arc in \mathcal{R}_2 (or, reversing the interplay, \mathcal{R}_1) as our ending point. Thus, as a simplifying measure in drawing such a illustration we can focus in on the portion of the foliation of \mathcal{R}_1 and \mathcal{R}_2 where these clasp arcs are placed, for each such portion of the foliation has just two vertices, two singularities and a single clasp arc. Consider the position of \mathcal{R}_s in this superimposed picture. If \mathcal{R}_s is not next to \mathcal{R}_3 then we can again do a cut-and-paste to produce a topological annulus that contributes $+1$ to the braid index of $b(X_+)$. If \mathcal{R}_s is next to the foliation of \mathcal{R}_3 then it is easy to see that $X_+ \cap (\mathcal{R}_1 \cup \mathcal{R}_s \cup \mathcal{R}_2)$ contributes $+1$ to $b(X_+)$. In the superimposed image of $\mathcal{R}_1 \cup \mathcal{R}_s \cup \mathcal{R}_2$ in \mathcal{R}_3 it can be seen that $X_+ \cap (\mathcal{R}_1 \cup \mathcal{R}_s \cup \mathcal{R}_2)$ bounds a type 1_d region. Again, the assumption that \mathcal{R}_3 is a fan is not needed. (In the right sketch of Figure 48(ii), the short s -band that runs from dot g to dot c' is an example of an s -band that would be superimposed in the foliation of the right sketch in Figure 48(ii).

Case (ii) \mathcal{R}_1 and \mathcal{R}_2 are both type $0_{d,1}$ (or both type $0_{d,2}$). Suppose regions $\mathcal{R}_1, \mathcal{R}_2$ are of type $0_{d,1}$ and \mathcal{R}_3 is a region that is a fan. We now employ this procedure of superimposing the foliations of \mathcal{R}_1 and \mathcal{R}_2 onto the foliation of \mathcal{R}_3 . We again consider the positioning of \mathcal{R}_s . As before we assume that the angular support of \mathcal{R}_s and $X_+ \cap \mathcal{R}_3$ is less than 2π . When we superimpose $\mathcal{R}_1 \cup \mathcal{R}_s \cup \mathcal{R}_2$ onto \mathcal{R}_3 we will see two copies of a type $0_{d,1}$ region side-by-side. (Please refer to Figure 43 for a understanding of the labels s and s' .) We have two sub-possibilities:

- (iia) \mathcal{R}_s runs from the s' label of the left region (which will be \mathcal{R}_1 in our side-by-side image) to the s label of the right region, which is \mathcal{R}_2 .

Notice that the cyclic ordering of the singularities in \mathcal{R}_3 imposes an order on the singularities of \mathcal{R}_1 and \mathcal{R}_2 where they intersect \mathcal{R}_3 . Thus, since \mathcal{R}_3 is a fan the $-\delta$ singularity in \mathcal{R}_1 must occur after the δ singularity in \mathcal{R}_2 . (Recall we are

making the simplifying assumption that \mathcal{R}_1 and \mathcal{R}_2 only have two singularities.) So the angular support of \mathcal{R}_s must overlap with the angular support of $\partial\mathcal{R}_1$ forcing $X_+ \cap (\mathcal{R}_1 \cup \mathcal{R}_s)$ to have angular length greater than 2π . Thus, the triple $(\mathcal{R}_1, \mathcal{R}_2, \mathcal{R}_s)$ contributes $+1$ to $b(X_+)$.

(iib) \mathcal{R}_s runs from the s label of the left region, which is \mathcal{R}_1 , to the s' label of the right region, which is \mathcal{R}_2 .

We cut-and-paste $\mathcal{R}_1, \mathcal{R}_2, \mathcal{R}_3$ to produce a type $0_{d,1}$ region having a portion of X_+ in its boundary but sheared in two places. If we adjoin \mathcal{R}_s to this new $0_{d,1}$ region we see that the resulting X_+ boundary contributes $+1$ to $b(X_+)$.

The proofs of (1'), (2') and (3') are almost identical to the proofs of (1),(2),(3), and we leave them to the reader. \square

5.3.2 Using flypes to push X_+ across \mathcal{S} In this subsection we assume that \mathcal{S} is a collection of thin (but not standard) annuli. The case of standard annuli will be considered separately, in [Section 5.7](#).

To simplify the notation, we use ‘**bold-faced type**’ for the black boundary and all its auxiliary structures; and ‘Roman type’ for the grey boundary and all its auxiliary structure. In particular, \mathcal{S} is foliated by s -arcs, \mathbf{a} -arcs (adjacent to \mathbf{X}), b -arcs and \mathbf{a} -arcs (adjacent to \mathbf{X}). (We will specify further auxiliary structure in due course.) Thus, if $\mathcal{S} \subset \mathcal{CA}$ is thin, then $\partial\mathcal{S} = \mathbf{X} \cup X$ where $\mathbf{X} \subset \mathcal{CA}$ is on the X_+ side of \mathcal{S} in \mathcal{CA} , and $X \subset \mathcal{CA}$ is on the X_- side of \mathcal{S} in \mathcal{CA} .

The main goal of this section is to establish that the isotopy of \mathbf{X} across \mathcal{S} to X is the result of a collection of flypes, not necessarily all admissible. Specifically, given \mathcal{S} we will produce a recipe for constructing a template (\mathbf{D}, D) where, via a common braiding assignment to the blocks of \mathbf{D} and D , we have that \mathbf{D} carries \mathbf{X} and D carries X . Our recipe will tell us how to designate blocks of both flavors (fixed and moving), and how to designate the strands connecting the blocks. Moreover, we will show how the thin structure of \mathcal{S} gives rise to a collection of flypes that carries \mathbf{D} to D .

To motivate our work, we begin with examples. The key to understanding the moving blocks will be a ‘block amalgamation’ process. Here we give two examples which illustrate: (i) how the clasp arcs of \mathcal{S} are used to designate the microblocks first mentioned in [Section 5.2](#); (ii) how the foliation of \mathcal{S} is used to amalgamate the microblocks into larger moving blocks; and how the foliation of \mathcal{S} is used to flype these moving blocks.

Example 1 is illustrated in [Figures 48 and 49](#). [Figure 48\(i\)](#) depicts two regions of the preimage of \mathcal{S} sharing two clasp intersections, γ^1 and γ^2 . For each clasp intersection




Figure 48: The isotopy of X_+ across \mathcal{S} in a simple example

in \mathcal{S} there is a readily identifiable pair of tab neighborhoods in the preimage (just look at what remains after stabilizing \mathbf{X} (resp. X) along the singular leaf θ_4 (resp. θ_3)). So for each microflype we have an associated microblock that contains a positive full twist. Figure 48(i) coupled with the left sketch in Figure 49 illustrates the strands in the two microblocks: strands ab and df are associated with the microblock for γ^1 ; and strands $a'b'$ and $d'f'$ are associated with the microblock for γ^2 .




Figure 49: The 3-space embedding of X_+ and X_- in the Example in Figure 48

Next, notice that the foliation of \mathcal{S} yields the following sequence of isotopies for moving \mathbf{X} to X : in Figure 48(i)–(ii) we stabilize \mathbf{X} along the θ_4 singular leaf; in

(ii)–(iii) we perform two microflypes; and in (iii)–(iv) we destabilize \mathbf{X} . The geometric realization of Figure 48(ii) is depicted in the left sketch of Figure 49. (The alphabetic labeling of points on \mathbf{X} and \mathbf{X} are meant to correspond between the two figures and the reader is encouraged to check the details of this correspondence.)

Now observe that the right two sketches in Figure 49 illustrates an amalgamation between the two full twists of the microblocks, allowing us to consolidate this stabilization, microflypes, destabilization sequence into a single flype. Also, observe that the two regions in Figure 48(i) will be fans. This can be verified by checking the ordering of the angular support of the edgepaths $cdefd'e'f'g'$ and $aba'b'$ that are in \mathbf{X} , along with the corresponding angular support of the analogous edgepaths in \mathbf{X} , since all singular leaves have an endpoint on the boundary of these regions.

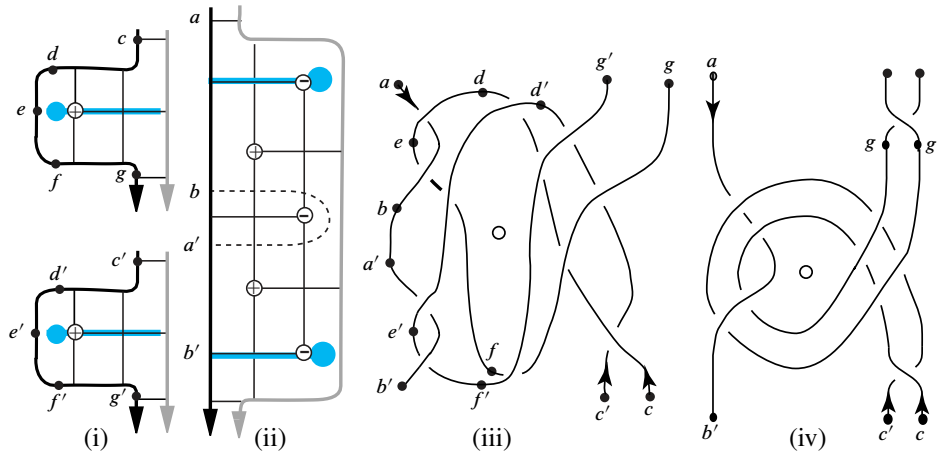


Figure 50: Amalgamating microflypes can result in flypes with weighted strands

The three regions of \mathcal{S} for Example 2 are illustrated in Figure 50(i) and (ii). Again, there are two clasp intersections in \mathcal{S} along with two pairs of tab neighborhoods which clearly delineate the microflypes and microblocks associated with each clasp intersection. Figure 50(iii) shows the geometric realization of \mathbf{X} : strands ab and df properly contains a full twist in one microblock; and strands $a'b'$ and $d'f'$ properly contains a full twist in the another microblock. And, again the foliation of \mathcal{S} supplies us with a sequence of stabilizations, microflypes, destabilization for moving \mathbf{X} to \mathbf{X} : the starting stabilization is along the segment $ba' \subset \mathbf{X}$; perform the two microflypes; and then do the two remaining destabilizations. The question is, can we amalgamate the two microblocks to achieve a three strand block that would contain the braiding depicted in Figure 50(iv)? The answer is that the information in the foliation is ambiguous; we

have not assigned any angular information to the singularities in Figure 50(i) and (ii). In Figure 50(iii) we have illustrated the point f (resp. d') as occurring after (resp. before) the point b' (resp. d) in the fibration, however there is no information in Figure 50(i) and (ii) that forces this choice. The occurrences might have been reversed. If they were reversed then the amalgamation of Figure 50(iv) would not have been possible. Thus, the foliation of Figure 50(i) and (ii) could depict either one braid -index-decreasing flype with a three strand block, or two elementary flypes followed by a destabilization.

As in Example 1 of Figure 48, it is easy to see that the region in Figure 50(ii) will be a fan. This will be independent of how the amalgamation of blocks occurs.

Example 3 (see Figure 51) shows a situation where successive flypes cannot be amalgamated. We see four flypes, with associated braid blocks X,Y,Z,W. In the initial diagram it looks as if it might be possible to amalgamate X and Y, but if we study the final diagram we see that this is impossible.




Figure 51: Four independent flypes which cannot be amalgamated

Remark 5.3.1 Observations based on these examples lead us to the following remarks about thin annuli \mathcal{S} :

- (i) The foliation of \mathcal{S} enables us to readily identify tab neighborhoods of clasp arcs.
- (ii) For each pair of tab neighborhoods (where the pairing is via the pairing of the clasp arcs) we have an associated microflype and microblock. By our parity conditions on the singularities of a region (refer to Figure 43) we know that all of the microflypes associated with a particular region are of the same

parity, ie either all positive flypes or all negative flypes. In particular, any type $0_{d,1}, 0_{d,2}, 1_d, -1_d$ region used in a flype will be a fan.

- (iii) We can move \mathbf{X} across the regions of \mathcal{S} to \mathbf{X} through a sequence of stabilizations of \mathbf{X} , microflypes and destabilizations of the resulting new \mathbf{X} . \diamond

Recall that our main goal in this subsection is to establish that the isotopy of \mathbf{X} across \mathcal{S} to \mathbf{X} is the result of a collection of flypes, not necessarily all admissible. For that we need to understand when it is possible to amalgamate a collection of microblocks into a larger block that is moved by a flype. So let us finally formalize the definition of ‘amalgamation’. To do this we first need to re-characterize microblocks in term of the foliation of \mathcal{S} .

By hypothesis, all clasp arcs in \mathcal{S} are doubly long. In particular, there are no short clasp arcs. For any pair γ_+^i, γ_-^i we have pushed both arcs into a chains of leaves. (For thin annuli these chains will always have exactly two leaves.) Since γ_+^i begins on \mathbf{X} , one of the endpoints of γ_+^i is a point $\mathbf{q}_+^i \in \mathbf{X}$. The fact that γ^i is doubly long shows that one of the endpoints of γ_-^i is also near \mathbf{X} . The singular leaf which contains it ends at a point $\mathbf{q}_-^i \in \mathbf{X}$. Since our clasp arc is doubly long, there is an analogous picture near \mathbf{X} . Both are illustrated in Figure 52. The part of \mathcal{CA} that belongs to the two normal neighborhoods runs all the way from the black boundary to the grey boundary. The two normal neighborhoods intersect along the clasp arc.

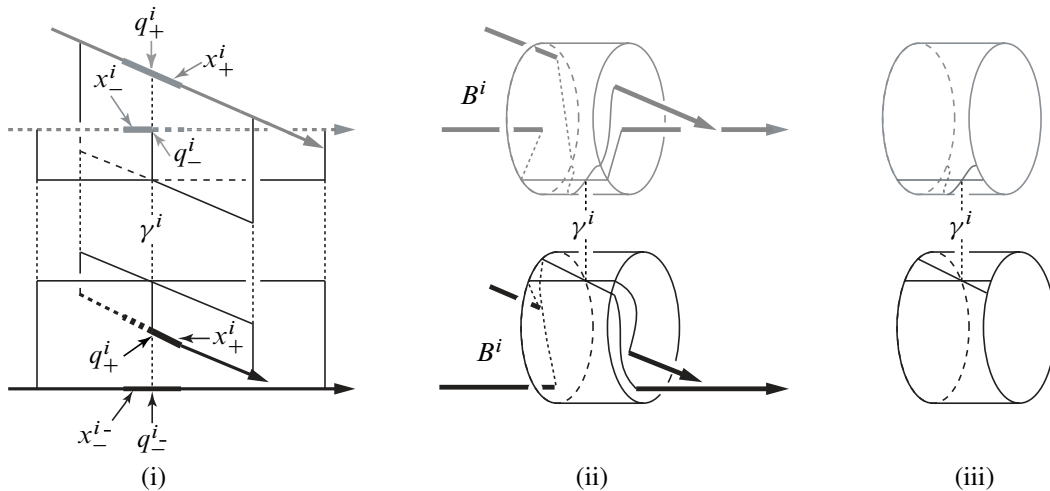


Figure 52: (i) The four microstrands associated to a doubly long clasp arc. (ii) The associated microblocks. (iii) The braid projection of the microstrands on the cylinder walls.

We fix subarcs \mathbf{x}_+^i and \mathbf{x}_-^i of \mathbf{X} , where \mathbf{x}_+^i (resp. \mathbf{x}_-^i) is a closed neighborhood of \mathbf{q}_+^i (resp. \mathbf{q}_-^i) on \mathbf{X} . The subarc \mathbf{x}^i is chosen so that it does not intersect the subarc associated to adjacent singular leaves on \mathbf{X} , also so that the arcs \mathbf{x}_+^i and \mathbf{x}_-^i have the same angular support $[\theta_{i_1}, \theta_{i_2}] \subset \mathbf{H}$. We call these arcs the *black microstrands*. There are, of course, similar thick *grey microstrands* \mathbf{x}_+^i and \mathbf{x}_-^i of \mathbf{X} . Notice that the isotopy of \mathbf{X} across \mathcal{S} pushes \mathbf{X} across \mathbf{N}_+^i (resp. \mathbf{N}_-^i) to \mathbf{X} , so the two black microstrands are mapped to the two grey microstrands.

Construct disjoint solid cylinders $\mathbf{B}^i, \mathbf{B}^i$ in 3-space which have the structure of blocks, as defined in Section 1 of this paper. These cylinders are, of course, microblocks, as defined in Section 5.2. The microblock \mathbf{B}^i is foliated by discs and contains the braided arcs \mathbf{x}_+^i and \mathbf{x}_-^i , which meet the disc fibers transversally inside \mathbf{B}^i . It intersects the fibers of \mathbf{H} in the interval $[\theta_{i_1}, \theta_{i_2}]$. We also have, without further work, a grey microblock. (See Figure 52(ii).)

There is an additional and very important feature of the geometry: The fact that our clasp arcs are doubly long tells us, immediately, that the two tab neighborhoods of the clasp arcs reach all the way from \mathbf{X} to \mathbf{X} , joining the black and grey microblocks. In fact, they intersect the side boundaries of the black (resp. grey) microblocks, and the intersection is a pair of intersecting arcs. We label the double points of the black and grey projection with the index of γ^i . (See Figure 52(iii).) We can think of these crossed arcs as ‘local projections’ of the microblocks. The very interesting feature of this projection is that the black and grey projections are joined to one-another by the tab neighborhoods of γ_+ and γ_- . (This is easy to see in Figure 52(ii), even though we omitted it to keep the picture as simple as possible).

Construct the black and grey microblocks, $\{\mathbf{B}^1, \dots, \mathbf{B}^r\}$ and $\{\mathbf{B}^1, \dots, \mathbf{B}^r\}$, one for each clasp arc in \mathcal{S} . Let $\mathbf{B} \subset S^3 \setminus \mathbf{A}$ be a 3-ball having the structure of a 2-disc cross an interval, $\Delta^2 \times [0, 1]$. Decompose $\partial\mathbf{B}$ as $\mathbf{c} \cup \mathbf{t} \cup \mathbf{b}$, where $\mathbf{c} = \partial\Delta^2 \times [0, 1]$; $\mathbf{t} = \Delta^2 \times \{0\}$; and $\mathbf{b} = \Delta^2 \times \{1\}$. (Our notation was chosen to suggest \mathbf{t} for ‘top’; \mathbf{c} for ‘cylinder’; and \mathbf{b} for ‘bottom’.) Then \mathbf{B} is an *amalgamating block* for the microblocks $\{\mathbf{B}^1, \dots, \mathbf{B}^r\}$ associated to the microstrands $\{\mathbf{x}_+^1, \mathbf{x}_-^1, \dots, \mathbf{x}_+^r, \mathbf{x}_-^r\}$ of \mathcal{S} if the following hold:

- (1) Each $\mathbf{B}^i \subset \text{int}(\mathbf{B})$ for $i \leq i \leq r$.
- (2) Its support, ie $\{\theta \mid \mathbf{B} \cap H_\theta \neq \emptyset\}$, is a closed interval having length less than 2π . Every component of $\mathbf{X} \cap \mathbf{B}$ contains a microstrand.
- (3) The top \mathbf{t} and bottom \mathbf{b} of $\partial\mathbf{B}$ are contained in generic disc fibers of \mathbf{H} . Moreover, each component of $\mathbf{t} \cap \mathcal{S}$ (resp. $\mathbf{b} \cap \mathcal{S}$) is a subarc of an \mathbf{a} -arc having one endpoint on $\partial\mathbf{t}$ (resp. $\partial\mathbf{b}$) and one endpoint on \mathbf{X} . The leaves of the induced foliation on

$\mathbf{c} \subset \partial \mathbf{B}$ are circles. The only non-singular leaves that \mathbf{c} intersect are \mathbf{a} -arcs and $\mathbf{c} \cap \mathbf{X} = \emptyset$.

(4) If \mathbf{x}_+^i (resp. \mathbf{x}_-^i) $\subset \text{int}(\mathbf{B})$, then its partner \mathbf{x}_-^i (resp. \mathbf{x}_+^i) is also a subset of $\text{int}(\mathbf{B})$. Also, the dotted segment between each \mathbf{q}_+^i and \mathbf{q}_-^i is a subset of $\text{int}(\mathbf{B})$.

(5) Replacing

$$\{\mathbf{B}, \mathbf{t}, \mathbf{c}, \mathbf{b}, \mathbf{a}, \{\mathbf{B}^1, \dots, \mathbf{B}^r\}, \{\mathbf{x}_+^1, \mathbf{x}_-^1, \dots, \mathbf{x}_+^r, \mathbf{x}_-^r\}\}$$

by

$$\{\mathbf{B}, \mathbf{t}, \mathbf{c}, \mathbf{b}, \mathbf{a}, \{\mathbf{B}^1, \dots, \mathbf{B}^r\}, \{\mathbf{x}_+^1, \mathbf{x}_-^1, \dots, \mathbf{x}_+^r, \mathbf{x}_-^r\}\}$$

we can also define the amalgamation of microblocks $\{\mathbf{B}^1, \dots, \mathbf{B}^r\}$.

Our next question is: what conditions need to be met so that we can say the isotopy across \mathcal{S} takes \mathbf{B} to \mathbf{B} ? A hint comes from Figure 52(iii). Consider the projection of the microstrands in \mathbf{B}^i onto its cylinder boundary and the projection of the microstrands in \mathbf{B}^i onto its cylinder boundary. Notice that these two projections are reflections of each other. This observation leads us to the following two definitions.

First, let \mathbf{B} be an amalgamating block. The *braid projection* $\pi(\mathbf{B})$ of X_+ onto $\mathbf{c} \subset \partial \mathbf{B}$ is the graph $\mathbf{c} \cap \mathcal{S}$, regarded as a subset of \mathbf{c} together with the clasp arc index labeling of this graph's double points. We define $\pi(\mathbf{B})$ in a similar fashion.

Second, amalgamating blocks \mathbf{B} and \mathbf{B} are a pair of *related* amalgamating blocks if the following hold:

- (a) $\{\mathbf{B}^1, \dots, \mathbf{B}^r\}$ and $\{\mathbf{B}^1, \dots, \mathbf{B}^r\}$ are microblocks associated with clasp intersections $\{\gamma^1, \dots, \gamma^r\}$.
- (b) \mathbf{B} is a block amalgamation of microblocks $\{\mathbf{B}^1, \dots, \mathbf{B}^r\}$.
- (c) \mathbf{B} is a block amalgamation of microblocks $\{\mathbf{B}^1, \dots, \mathbf{B}^r\}$.
- (d) $\pi(\mathbf{B})$ and $\pi(\mathbf{B})$ are reflections of each other.

Recall Figure 51 and see Figure 53 for an example.

Proposition 5.3.2 *Let \mathcal{S} be a collection of thin annuli with boundary components \mathbf{X} and \mathbf{X} . Then the motion of \mathbf{X} to \mathbf{X} may be realized by a collection of flypes and destabilizations.*

Proof We will construct a pair of candidates for block-strand diagrams (\mathbf{D}, \mathbf{D}) such that \mathbf{D} (resp. \mathbf{D}) carries \mathbf{X} (resp. \mathbf{X}) and the motion $\mathbf{D} \rightarrow \mathbf{D}$ is a sequence of flypes. We say that our diagrams (\mathbf{D}, \mathbf{D}) are merely *candidates* for block-strand diagrams is that




Figure 53: The projection of related microstrands onto the cylindrical walls in the boundaries of amalgamating blocks

condition (4) in the definition of a block-strand diagram may not hold. Indeed, it may be difficult to verify whether (4) does or does not hold at this stage of the argument. For example, it's possible that we could amalgamate blocks, and so produce new blocks which have full braid index. That is a matter which we are not ready to discuss at this time. (Look ahead to [Section 6.2](#) for more on this. See, in particular, [Corollary 6.2.1](#).) However, if we can construct \mathbf{D} and D so that they carry \mathbf{X} and X , and prove that the motion $\mathbf{D} \rightarrow D$ is a sequence of flypes and destabilizations, we will have accomplished our task.

To construct \mathbf{D} and D we need to understand several aspects of their structure. Observe that the strands of \mathbf{X} which change position during the passage to X are precisely those which bound the 'tiled' part of \mathcal{CA} , that is the part which is away from the bands of s -arcs. On the other hand, the part that stays fixed is the part that intersects the bands of s -arcs. So we can consider, separately, the 'moving' and 'fixed' parts of (\mathbf{D}, D) :

Moving blocks An amalgamating block \mathbf{B} will be moved to an amalgamating block B if and only if they are related amalgamating blocks, as defined above. Note that there may be some choices involved when we select the amalgamating blocks. We make those choices in such a way that the set of all moving blocks has minimal cardinality.

Moving Strands The moving strands in \mathbf{D} are all subarcs of \mathbf{X} which are not amalgamated into moving blocks and are not in the bands of s -arcs.

Fixed blocks Here is the intuitive idea: The fixed blocks are associated to braiding between the strands of \mathbf{X} (and so also of X) in the part of \mathcal{S} which is foliated by bands

of s -arcs. The braiding between bands of s -arcs is not part of our geometry. \mathcal{S} only detects the places where \mathbf{X} and \mathbf{X} differ in a non-trivial way, but in regions where there are bands of s -arcs they essentially coincide. However, what might happen is that there is braiding between the bands of s -arcs, but that the braiding is interrupted by, for example, strands of \mathbf{X} which separate two potential blocks. In order to determine the angular regions where the blocks occur, we therefore look for singularities in the tiled part of \mathcal{S} . The block subdivision so-obtained will be too extensive, and we will then need to amalgamate fixed blocks.

We consider the H_θ -sequence for \mathcal{S} . Let $\mathcal{S}_{\text{tiled}} \subset \mathcal{S}$ be the portion of \mathcal{S} that is *not* foliated by s -arcs. Let $\{\theta_0, \theta_1, \dots, \theta_\xi\} \subset [0, 2\pi)$ be a cyclic listing of all the angles at which the corresponding $H_{\theta_i} \in \mathbf{H}$ contains a singularity, ordered according to their natural cyclic order in \mathbf{H} . Since the bands of s -arcs are foliated without singularities, they must be located in

$$H' = \bigcup_{\theta \in (\theta_i, \theta_{i+1}), i=0}^{i=\xi} \{H_\theta \setminus H_\theta \cap \mathcal{S}_{\text{tiled}}\}.$$

Each component of H' is a disc, and it contains only s -arcs. Let C be a connected component of H' . Then C has a $\Delta^2 \times [\theta_i, \theta_{i+1}]$ structure, and it may contain some number of bands of s -arcs which braid with one-another inside C . If so, we amalgamate this braiding of s -arc bands in C into a single block $B(C)$, which will be a fixed block in the template (\mathbf{D}, \mathbf{D}) . If C has no s -arcs then $B(C)$ is vacuous. If C has a single s -arc then $B(c)$ is a single fixed strand.

Suppose that there is another connected component C' in H' , with its fixed block $B(C')$, and suppose further that $\{s\text{-arcs}\} \cap C' \cap H_{\theta_{i+1}} \subset \{s\text{-arcs}\} \cap C \cap H_{\theta_{i+1}}$. If this happens, the singularity at θ_{i+1} could have been ignored. Another way to say this is that we can amalgamate the blocks $B(C)$ and $B(C')$ into a larger fixed block of the template $B(C \cup C')$ of (\mathbf{D}, \mathbf{D}) . There may be some choices involved, and we make them so that the set of all fixed blocks has minimal cardinality.

Fixed strands It may happen that we have two blocks, $B(C)$ and $B(C')$, for which $(C \cap H_{\theta_{i+1}}) \cap (C' \cap H_{\theta_{i+1}}) \neq \emptyset$, but $(C \cap H_{\theta_{i+1}}) \neq (C' \cap H_{\theta_{i+1}})$, so that the amalgamating condition fails. In this situation there are strands that run between $B(C)$ and $B(C')$, however not all of the s -bands coming out of $B(C)$ go into $B(C')$, and/or not all of the strands going into $B(C')$ are strands that emerge from $B(C)$. Any braiding between the s bands that run between $B(C)$ and $B(C')$ can obviously be pushed into either $B(C)$ or $B(C')$ (eliminating $H_{\theta_{i+1}}$ as a singular fiber) and they will then become fixed strands.

It is clear from the construction of (\mathbf{D}, \mathbf{D}) that the template (\mathbf{D}, \mathbf{D}) carries (\mathbf{X}, \mathbf{X}) .

We still need to show that for each moving block pair (\mathbf{B}, \mathbf{B}) we have a flype (possibly with weighted strands) taking \mathbf{B} to \mathbf{B} as in Figure 4(a). To do this we first need to isolate each moving black block. Specifically, we wish to subdivide the foliation of \mathcal{S} so that there is a subannular region $\mathcal{S}' \subset \mathcal{S}$ such that if $\mathbf{y} \subset \mathbf{X}$ is a strand of \mathbf{D} then \mathbf{y} intersects s -arcs in the foliation of \mathcal{S}' . We refer to Figure 54 to see how \mathcal{S}' can be obtained through the stabilization of X . Suppose we have two clasp intersections γ^i and γ^j where the black microblock associated with γ^i is in an amalgamated black block $\mathbf{B}(i)$ and the black microblock associated with γ^j is in an amalgamated black block $\mathbf{B}(j)$. Let $\mathbf{y} \subset \mathbf{X}$ be a strand of \mathbf{D} that has its endpoints on $\mathbf{B}(i)$ and $\mathbf{B}(j)$.




Figure 54: Subdividing non-trivially foliated regions in \mathcal{S} . The black (resp. dotted grey) horizontal arcs are the arcs we refer to in the text as ‘ \mathbf{y} ’ (resp. ‘ \mathbf{y} ’). The intersections of the clasp arcs with the regions of interest are labeled g_{\pm}^i and g_{\pm}^j .

Since \mathbf{y} is a strand it cannot intersect any microstrands of \mathbf{X} . Thus, it either intersects s -arcs in the foliation of \mathcal{S} (in which case no subdivision is necessary), or the only non-singular leaves it intersects are \mathbf{a} -arcs. To describe the needed subdivision in the latter situation, we look at the corresponding grey microblocks associated with our two clasp intersections; the corresponding grey amalgamating blocks $\mathbf{B}(i)$ and $\mathbf{B}(j)$; and the corresponding grey strand $\mathbf{y} \subset \mathbf{X}$ which has endpoints on $\mathbf{B}(i)$ and $\mathbf{B}(j)$. By assumption, the only non-singular leaves intersected by the strand \mathbf{y} are \mathbf{a} -arcs in the foliation of \mathcal{S} . If \mathbf{y} intersects an ab -singular leaf (as illustrated in Figure 54(a)) then we can stabilize X along this singular leaf to produce a subannulus of \mathcal{S} that

has y intersecting s -arcs. If the strand y does not intersect an ab -singular leaf then we can enlarge the foliation of \mathcal{S} thru the addition of two vertices and singularities, as illustrated in the three remaining sequences, Figure 54(b)–(d). (The new b -arc will necessarily be inessential.) The strand y will now intersect the endpoint of an ab -singular leaf and a stabilization of X is possible. The corresponding black strand y will then intersect s -arcs in the resulting \mathcal{S}' , isolating the block $\mathbf{B}(i)$ from the block $\mathbf{B}(j)$.




Figure 55: An example that shows how a stabilization, two microfypes and a destabilization are amalgamated into a flype. The bottom left and right sketches show the corresponding subsets of \mathbf{D} and D . The top right sketch illustrates the part of the diagram that we refer to as D' in the text.

Let $\mathcal{S}' \subset \mathcal{S}$ be thin subannuli that isolate all the moving blocks. Then there is a pair (\mathbf{D}, D') where D' is obtained from D by the stabilizations in the subdivision of Figure 54. Part of it is illustrated in Figure 55. Saying the moving blocks of (\mathbf{D}, D') are isolated is saying that every strand of \mathbf{D} or D intersects s -arcs in the foliation of \mathcal{S}' . Now, if we can show that moving across \mathcal{S}' is a collection of flypes, then it will follow that the destabilizations come from going from D' to D .

To show that the isotopy across \mathcal{S}' takes \mathbf{D} to D' , using a collection of flypes, we proceed as follows: since the blocks are isolated, we need only consider what the isotopy is for one block. This means we are back in a situation similar to that of Example 1 and Example 2 at the beginning of this section. We have a subcollection of regions $\{\mathcal{R}_1, \dots, \mathcal{R}_h\} \subset \mathcal{S}'$ in the foliation of \mathcal{S}' (as defined in Section 5.3.1) that are grouped together by clasp intersections. These regions are all either type $0_{d,1}$, $0_{d,2}$, 1_d or -1_d . All of the associated microblocks of the clasp intersections of this collection

of regions can be amalgamated into a block pair (\mathbf{B}, \mathbf{B}) . Our claim is that the isotopy of \mathbf{B} across the regions $\{\mathcal{R}_1, \dots, \mathcal{R}_h\}$ to \mathbf{B} is a flype. By statement (ii) of [Remark 5.3.1](#) we know that we can subdivide the microblocks in \mathbf{B} and flype them across to the microblocks in \mathbf{B} . Also, the fact $\pi(\mathbf{B})$ and $\pi(\mathbf{B})$ are reflections tells us that we can re-amalgamate them into \mathbf{B} . In the simplest case of a subdivision into two blocks ([Figure 55](#)) we see that the entire isotopy of statement (ii) of [Remark 5.3.1](#) can be consolidated into a single flype. This situation can obviously be generalized into a larger number of blocks in the subdivision by an iterated nesting, the first iteration being: before doing ‘flype 1’ we subdivide block \mathbf{P} into two new blocks and, since block \mathbf{P} in the upper-left sketch could have been flyped, perform the entire sequence on these two new blocks first. Thus, \mathbf{B} is flyped to \mathbf{B} and the proof of [Proposition 5.3.2](#) is complete. \square

Example 5.3.1 The projection criterion is essential to choosing block amalgamations that maintain their integrity during a flype isotopy. This is clear from the example that we gave earlier, in [Figure 51](#). Before the flype it looks as if X and W (and also Y and Z) can be amalgamated, but after the flype that is impossible. Indeed, the requirement that the projections be reflections of each other will be violated if we attempt to amalgamate X and W , or Y and Z . \diamond

5.4 Constructing a candidate for a template (D_+, D_-) from \mathcal{CA}

We interrupt the flow of the argument, briefly. The procedure that we gave in [Section 5.3.2](#) for constructing a candidate for a template from thin annuli is almost identical to the procedure for constructing a candidate for a template (D_+, D_-) from the foliation of \mathcal{CA} . We will need to know how to do it when we reach the end of [Section 6](#), and since the alert reader will have the essential ideas in mind right now, this seems like a good moment to go through the steps in the construction. The argument will be repetitious, because it is modeled on that in [Section 5.3.2](#), however it seems better to be repetitious than to keep asking the reader to turn back.

Assume that we are given \mathcal{CA} , with boundary $X_+ - X_-$. Assume that the complexity (c_1, c_2) has been minimized by the use of exchange moves. Assume further that we have removed all short clasp arcs, and that all non-short clasp arcs are in normal neighborhoods. We will show that we can construct a pair of diagrams $(\mathcal{D}_+, \mathcal{D}_-)$, and an isotopy $\mathcal{D}_+ \rightarrow \mathcal{D}_-$ such that $(\mathcal{D}_+, \mathcal{D}_-)$ carries (X_+, X_-) , and the motion $\mathcal{D}_+ \rightarrow \mathcal{D}_-$ sends $X_+ \rightarrow X_-$. As in [Proposition 5.3.2](#) we do not know whether \mathcal{D}_+ and \mathcal{D}_- are block strand diagrams because we do not know whether condition (4) for a block-strand diagram holds. That matter will not concern us now. (Look ahead to [Section 6.2](#) for more on this. See, in particular, [Corollary 6.2.1](#)).

As in [Section 5.3.2](#) there are four ingredients in the structure of (D_+, D_-) : the moving blocks, the fixed blocks, the moving strands, and the fixed strands. The precise description of how we construct each is very close to that in [Section 5.3.2](#), except that:

- \mathcal{CA} replaces \mathcal{S} ,
- $\mathcal{CA}_{\text{tiled}}$ (the portion of \mathcal{CA} not foliated by s -arcs) replaces $\mathcal{S}_{\text{tiled}}$. As a consequence, we are constructing the pair $(\mathcal{D}_+, \mathcal{D}_-)$ (instead of the pair (\mathbf{D}, \mathbf{D})). Also the motion that takes $\mathcal{D}_+ \rightarrow \mathcal{D}_-$ now takes $X_+ \rightarrow X_-$ (instead of the motion that takes $\mathbf{D} \rightarrow \mathbf{D}$ taking $\mathbf{X} \rightarrow \mathbf{X}$).

(i) **The Moving Blocks** If there are no doubly long clasp arcs, then there also are no moving blocks, and this part of the construction ends. So we assume that there are doubly long clasp arcs. Therefore the two tab neighborhoods of the clasp arcs reach all the way from X_+ to X_- . We are in the situation of [Figure 52](#), and we may construct black and grey microstrands, exactly as we did in [Section 5.3.2](#), only now the black and grey microstrands are subarcs of X_+ and X_- . Notice that the isotopy of X_+ across \mathcal{CA} pushes X_+ across \mathbf{N}_+^i to X_- , so the two X_+ microstrands are always mapped to the two X_- microstrands. Exactly as in [Section 5.3.2](#) we have cylinders $\mathbf{B}^i, \mathbf{B}^i$ for our microblocks. The local projection of the microstrands in \mathbf{B}^i (resp. \mathbf{B}^i) are determined by the intersections of \mathcal{S}_+ (resp. \mathcal{S}_-) with the walls of \mathbf{B}^i (resp. \mathbf{B}^i). As in [Section 5.3.2](#), the double points of the local projections are labeled with the index of g^i .

Recall the amalgamation of microblocks, and the example of a template in [Figure 51](#). Before the two flypes it looks as if X and W (and also Y and Z) can be amalgamated, but after the flype we see that is impossible. The requirement that we introduced, that the projections be reflections of each other, is violated if we attempt to amalgamate X and W , or Y and Z . We developed a ‘projection criterion’ for when this is possible. The essential idea was that the block amalgamations that we used maintained their integrity during the passage from the black side to the grey side.

Let’s begin the construction of the black microblocks. The blocks \mathbf{B} and \mathbf{B} have the same structure as in [Section 5.3.2](#). We continue to use the terms ‘top’, ‘cylinder’, ‘bottom’, with the same meaning as in [Section 5.3.2](#). Then \mathbf{B} is an amalgamation of microblocks $\{\mathbf{B}^1, \dots, \mathbf{B}^r\}$ for microstrands $\{\mathbf{x}_+^1, \mathbf{x}_-^1, \dots, \mathbf{x}_+^r, \mathbf{x}_-^r\}$ of \mathcal{CA} if the following hold:

- (1) As in [Section 5.3.2](#): Each $\mathbf{B}^i \subset \text{int}(\mathbf{B})$ for $i \leq i \leq r$.
- (2) As in [Section 5.3.2](#): Its support, ie $\{\theta \mid \mathbf{B} \cap H_\theta \neq \emptyset\}$, is a closed interval having length less than 2π . Every component of $\mathbf{X} \cap \mathbf{B}$ contains a microstrand.

- (3) A subtle change from (3) in Section 5.3.2: note that $\mathbf{t} \cap \mathcal{CA}$ and $\mathbf{b} \cap \mathcal{CA}$ are each still a collection of subarcs of a_+ -arcs and, possibly, subarcs in b -arcs. This necessarily implies that \mathbf{t} and \mathbf{b} are contained in a generic disc fiber of \mathbf{H} . We can assume that neither \mathbf{t} nor \mathbf{b} intersect any microstrands. Continuing, we have the immersion $\rho: \mathcal{PA} \rightarrow \mathcal{CA}$. Each component of $\mathbf{R} \subset \rho^{-1}(\mathbf{B} \cap \mathcal{CA}) \subset \mathcal{PA}$ is a rectangular region that does not contain any singular points. Specifically:
- (a) If $\mathbf{R} \cap X_+ \neq \emptyset$ then \mathbf{R} is trivially foliated with $\partial\mathbf{R} = \beta_{+1} \cup \beta_{+2} \cup \beta_{+3} \cup \beta_{+4}$, where (i) $\beta_{+1} \subset X_+$, and β_{+1} intersects some microstrands in our specified set. (ii) β_{+2} and β_{+4} are subarcs of a_+ -arcs. (iii) β_{+3} is transverse to the foliation of \mathcal{PA} . Moreover the only non-singular leaves it intersects are a_+ -arcs.
 - (b) If $\mathbf{R} \cap X = \emptyset$ then \mathbf{R} is trivially foliated, with $\partial\mathbf{R} = \beta_{+1} \cup \beta_{+2} \cup \beta_{+3} \cup \beta_{+4}$ where β_{+1} and β_{+3} are in leaves; also β_{+2} and β_{+4} are transverse to the foliation of \mathcal{PA} .
- (4) As in Section 5.3.2: If \mathbf{x}_+^i (resp. \mathbf{x}_-^i) $\subset \text{int}(\mathbf{B})$, then its partner \mathbf{x}_-^i (resp. \mathbf{x}_+^i) is also a subset of $\text{int}(\mathbf{B})$. Also, the dotted segment between each \mathbf{q}_+^i and \mathbf{q}_-^i is a subset of $\text{int}(\mathbf{B})$.
- (5) As in Section 5.3.2: Replacing

$$\{\mathbf{B}, \mathbf{t}, \mathbf{c}, \mathbf{b}, \mathbf{a}, \{\mathbf{B}^1, \dots, \mathbf{B}^r\}, \{\mathbf{x}_+^1, \mathbf{x}_-^1, \dots, \mathbf{x}_+^r, \mathbf{x}_-^r\}\}$$

by

$$\{\mathbf{B}, \mathbf{t}, \mathbf{c}, \mathbf{b}, \mathbf{a}, \{\mathbf{B}^1, \dots, \mathbf{B}^r\}, \{\mathbf{x}_+^1, \mathbf{x}_-^1, \dots, \mathbf{x}_+^r, \mathbf{x}_-^r\}\},$$

we can also define the amalgamation of microblocks $\{\mathbf{B}^1, \dots, \mathbf{B}^r\}$.

As before, we also have related concepts for \mathbf{B} , the amalgamation of microblocks $\{\mathbf{B}^1, \dots, \mathbf{B}^r\}$.

Our definition of the braid projection of X_+ onto $\mathcal{B} \subset \partial\mathbf{B}$ has changed a little bit because of the new conditions in (3) above. For a given amalgamating block \mathbf{B} we consider $\{\mathbf{R}^1, \dots, \mathbf{R}^l\} \subset \rho^{-1}(\mathbf{B} \cap \mathcal{CA})$. These are the regions referred to in condition (3)(a). Let $\{\beta_{+3}^1, \dots, \beta_{+3}^l\}$ be the corresponding β_{+3} boundary sides of these regions. Then the graph $\rho(\cup_{1 \leq i \leq l} \beta_{+3}^i) \subset \mathbf{c}$, along with the clasp arc index labeling of the double points, are in $\pi(\mathbf{B})$. Similarly, $\pi(\mathbf{B})$ is just the labeled graph $\rho(\cup_{1 \leq i \leq l} \beta_{-3}^i) \subset \mathbf{c}$. Using this definition of the braid projection of X_+ onto $\mathcal{B} \subset \partial\mathbf{B}$, our definition of ‘related amalgamating blocks’ is the same as before, viz:

Let \mathbf{B} be an amalgamating block. The braid projection $\pi(\mathbf{B})$ of X_+ onto $\mathbf{c} \subset \partial\mathbf{B}$ is the graph $\mathbf{c} \cap \mathcal{S}$, regarded as a subset of \mathbf{c} together with the clasp arc index labeling of this graph’s double points. We define $\pi(\mathbf{B})$ in a similar fashion. Amalgamating blocks \mathbf{B} and \mathbf{B} are a pair of related amalgamating blocks if the following hold:

- (a) $\{\mathbf{B}^1, \dots, \mathbf{B}^r\}$ and $\{B^1, \dots, B^r\}$ are microblocks associated with clasp intersections $\{\gamma^1, \dots, \gamma^r\}$.
- (b) \mathbf{B} is a block amalgamation of microblocks $\{\mathbf{B}^1, \dots, \mathbf{B}^r\}$.
- (c) B is a block amalgamation of microblocks $\{B^1, \dots, B^r\}$.
- (d) $\pi(\mathbf{B})$ and $\pi(B)$ are reflections of each other.

Look back to [Figure 53](#) for an example.

(ii) Fixed Blocks The fixed blocks are associated to braiding between the strands of \mathbf{X} (and so also of X) in the part of \mathcal{S} which is foliated by bands of s -arcs. The braiding between bands of s -arcs is not part of our geometry. \mathcal{S} only detects the places where \mathbf{X} and X differ in a non-trivial way, but in regions where there are bands of s -arcs they essentially coincide. However, what might happen is that there is braiding between the bands of s -arcs, but that the braiding is interrupted by, for example, strands of \mathbf{X} which separate two potential blocks. In order to determine the angular regions where the blocks occur, we therefore look for singularities in the tiled part of \mathcal{S} . The block subdivision so-obtained will be too extensive, and we will then need to amalgamate fixed blocks. For further details, see the proof of [Proposition 5.3.2](#), which applies without any changes at all.

(iii) Moving Strands The moving strands in \mathbf{D} are all subarcs of \mathbf{X} which are not amalgamated into moving blocks and are not in the bands of s -arcs. We will give details on how they are placed in in the proof of [Proposition 6.3.1](#).

(iv) Fixed Strands It may happen that we have two blocks, $B(C)$ and $B(C')$, for which $(C \cap H_{\theta_{i+1}}) \cap (C' \cap H_{\theta_{i+1}}) \neq \emptyset$, but $(C \cap H_{\theta_{i+1}}) \neq (C' \cap H_{\theta_{i+1}})$, so that the amalgamating condition fails. In this situation there are strands that run between $B(C)$ and $B(C')$, however not all of the s -bands coming out of $B(C)$ go into $B(C')$, and/or not all of the strands going into $B(C')$ are strands that emerge from $B(C)$. Any braiding between the s bands that run between $B(C)$ and $B(C')$ can obviously be pushed into either $B(C)$ or $B(C')$ (eliminating $H_{\theta_{i+1}}$ as a singular fiber) and they will then become fixed strands.

We are now in a position to take \mathcal{CA} and construct an associated candidate for a template (D_+, D_-) . It is clear from the construction that $(D_+$ and $D_-)$ will carry X_+ and X_- .

Remark 5.4.1 Go back to [Figure 9](#). The bottom two sketches show the actual pair of block-strand diagrams (D_+, D_-) that make up the template in [Figure 9](#), whereas the other 5 sketches give it meaning by showing the actual combination of flypes (not

necessarily admissible), exchange moves and destabilizations that explain the isotopy from \mathcal{D}_+ to \mathcal{D}_- .

The construction that we just gave falls short of doing what we did in [Figure 9](#) in two ways. The first has been noted before this, and we stress it again, but the second is new:

- (1) As noted earlier, the blocks may not either be ‘consolidated’ or ‘optimal’, two matters that will be discussed in [Section 6](#). That’s why we have said, repeatedly, that our \mathcal{D}_+ and \mathcal{D}_- are *candidates* for block-strand diagrams, and the pair $(\mathcal{D}_+, \mathcal{D}_-)$ is a *candidate* for a template.
- (2) We have not constructed the intermediate stages that explain the isotopy. That is clear, because all we dealt with is the long clasp arcs, but to explain the intermediate isotopy we would have had to also look at clasp arcs which are not, initially, doubly long but become doubly long after part of \mathcal{CA} has been crossed. That construction remains for future work, for readers who may be interested in, for example, in working on our ‘open problem (3)’ in [Section 7](#) of this paper. \diamond

5.5 Pushing across regions with a G–fype foliation

We have seen that if \mathcal{PA} contains a long clasp arc, then there is a family of thin annuli $\mathcal{S}^1 \subset \mathcal{PA}$ having $\partial\mathcal{S}^1 = \mathbf{X}^0 \cup \mathbf{X}^1$, where $\mathbf{X}^0 = X_+$, also \mathbf{X}^0 can be pushed across \mathcal{S}^1 to \mathbf{X}^1 using a sequence of flypes. The complement of \mathcal{S}^1 in \mathcal{PA} will be a new family of annuli which has boundary $\mathbf{X}^1 \cup X_-$. It may happen that we can then apply [Proposition 5.1.1](#) to simplify \mathcal{S}^1 via $(ab)^*$ exchange moves; or it may happen that $\mathcal{PA} \setminus \mathcal{S}^1$ has long clasp arcs, in which case we can iterate the construction. Putting the two cases together we can construct a second annulus \mathcal{S}^2 with $\partial\mathcal{S}^2 = \mathbf{X}^1 \cup \mathbf{X}^2$ such that the movement across \mathcal{S}^2 is either an $(ab)^*$ exchange move or a sequence of flypes. This procedure can be iterated until either we produce thin annuli that have X_- as a boundary component, or there are no more $(ab)^*$ exchange moves, or there are no more long clasp arcs. If we do have thin annuli with X_- in the boundary then we will have decomposed \mathcal{CA} into a sequence of thin annuli. And, we have moved across \mathcal{CA} using flypes and ab or ab^* exchange moves. The difficulty in moving across \mathcal{PA} this way is that it may happen that we needed to use inadmissible flypes.

With that difficulty in mind, we say that a family of thin annuli $\mathcal{S}^k \cup \mathcal{S}^{k+1} \cup \dots \cup \mathcal{S}^{k+l}$ supports a *G–fype foliation* if (i) it is possible to cross the region using only flypes and ab or $(ab)^*$ exchange moves, and (ii) there exist integers k, q such that $b(\mathbf{X}^{k+i}) \geq b(\mathbf{X}^k)$ for $1 \leq i \leq (q-1)$ and $b(\mathbf{X}^{k+q}) \leq b(\mathbf{X}^k)$. In this situation the combined flyping and ab or $(ab)^*$ motion across the region $\mathcal{S}^k \cup \mathcal{S}^{k+1} \cup \dots \cup \mathcal{S}^{k+q}$ is a *G–fype*. It is readily seen that if (\mathbf{B}, \mathbf{B}) is a pair of related amalgamating blocks for the pair

(X_+, X_-) then \mathbf{B} maintains its integrity through successive flypes and is isotoped to \mathbf{B} .

Example 5.5.1 The template in Figure 56 is an illustration of a G-flype.




Figure 56: The foliation of \mathcal{PA} , in the situation of the 6-braid template of Figure 9

It shows the foliated annulus \mathcal{PA} which was used to construct the template which we saw in Section 1 in the boxed sketches at the bottom of Figure 9. The labels ‘block Y’ and ‘block Z’ refer to the blocks in Figure 9. In the interest of keeping the foliation as simple and understandable as possible, we decided to position the clasp arcs only in tab neighborhoods, and not in normal neighborhoods. Thus, the clasp arcs are *not* contained in the finite union of leaves. (The addition of 12 vertices and 12 singularities would be needed in order to position the clasps in normal neighborhoods.) In order to remove distracting details from the picture, we have eliminated inessential b -arcs, collapsing the normal neighborhoods of the clasp arcs γ_+^1 and γ_+^2 . This is always possible, and we did it to save space. The arcs γ^1 and γ^2 in sketch (d) are the clasp arcs used to construct the first thinly foliated annuli that are split off by the dotted curves in sketches (b) and (c). After we have flyped across this initial thin annulus (the first negative flype of Figure 9) we will have four resulting vertices where we can perform (ab) exchange move. These are the two double-strand exchange moves in the isotopy sequence of Figure 9. Finally, the remaining portion of \mathcal{CA} will just

be the region in sketch (a) and tab neighborhoods around clasp arcs γ^3 and γ^4 in sketches (b) and (c). All of these individual motions were across thin annuli, giving us a decomposition of \mathcal{CA} into thin annuli. Since the braid index starts at $b(X_+) = 6$, goes to 7 after the first flype, remains at 7 for all of the exchange moves and only returns to $6 = b(X_-)$ after the last flype, all of the thin annuli in this decomposition go to make up the motion across \mathcal{PA} . The combination of all of these isotopies is the 6–braid G–flype which we illustrated in [Figure 9](#) and [Figure 56](#).

Notice that there are no moving blocks in [Figure 9](#) because there are no doubly long clasp arcs in [Figure 56](#). The following features of the template in [Figure 9](#) are outside the support of the isotopy: the four blocks X, Y, Z, W and the following strands that join them: all strands that enter and leave W and X , also strand 1 entering Y , strand 1 leaving Y , strands 3 and 4 entering Z and strands 1 and 4 leaving Z . We need to account for the changes in strands 2 and 3 entering Y , strands 2 and 3 leaving Y , strands 1 and 2 entering Z and strands 2 and 3 leaving Z . These are all described completely by the data in the foliated annulus.

5.6 Pushing across regions with a G–exchange foliation

In this subsection all clasp arcs have intermediate length, that is the pierce points of both γ_+ and γ_- are on b –arcs.

We begin with an example which illustrates how G–exchange moves arise. The example is the foliated subsurface of \mathcal{PA} that supports the G–exchange move of [Figure 10](#). It will not be difficult to understand this figure, now that the main tools in this paper, ie the foliated immersed annulus \mathcal{CA} and its foliated preimage \mathcal{PA} , are in place.

The foliated subsurface of \mathcal{PA} that supports the G–exchange move of [Figure 10](#) is illustrated in [Figure 57](#), which shows four discs on \mathcal{PA} . Label the discs \mathcal{R}_1 (top left), \mathcal{R}_2 (top right), \mathcal{R}_3 (bottom left) and \mathcal{R}_4 (bottom right). No \mathcal{R}_i is good, because each contains the puncture endpoints of clasp arcs. On the other hand, each is an embedded subset of \mathcal{PA} (because no \mathcal{R}_i contains both preimages of a clasp arc pair). The clasp arcs are all doubly-intermediate (because all of the puncture endpoints are on b –arcs).

Ignoring the clasp arcs momentarily, we see that each disc is topologically equivalent to the shaded disc in the upper left sketch in [Figure 23](#), ie it contains a positive vertex of valence 2 and type ab , and singularities of opposite sign. Each of our 4 regions contains in its boundary a subarc of X_- and a subarc of X_+ , colored dark and light respectively. These arcs are labeled $a_{\pm}, b_{\pm}, c_{\pm}$ and d_{\pm} . If the clasp arcs were not there we could use four ab –exchange moves to push a_+ (resp. b_+, c_+, d_+) across its disc to a_- (resp. b_-, c_-, d_-).

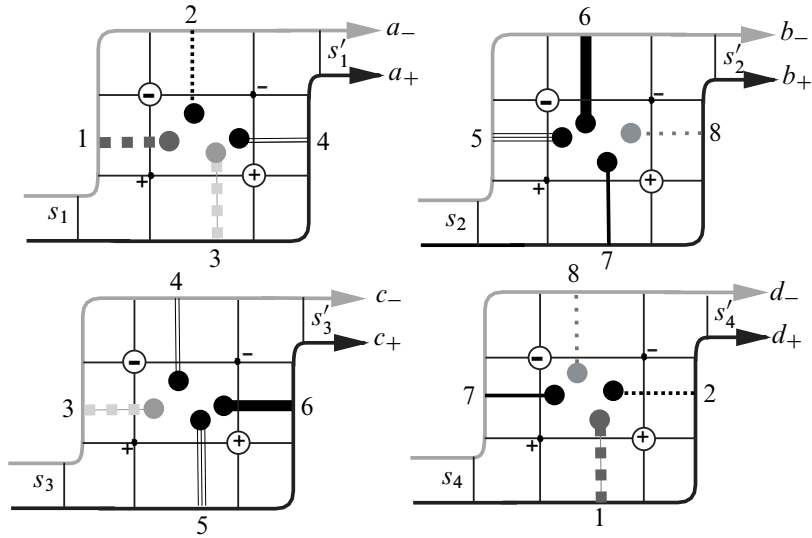


Figure 57: Foliated subsurface of \mathcal{PA} which are the support of the G–exchange move of Figure 10

Since \mathcal{R}_1 intersects \mathcal{R}_4 along clasp arcs 1 and 2, and \mathcal{R}_3 along clasp arcs 3 and 4, we cannot complete the move on strand a_+ until we begin the moves on strands c_+ and d_+ . But then, the motions of strands c_+ and d_+ across discs \mathcal{R}_3 and \mathcal{R}_4 cannot be completed because those discs have clasp intersections with \mathcal{R}_2 , and in fact no one of the motions can be completed until all of the others are completed too. That is, we have a G–exchange move. This particular very simple G–exchange move was discussed in Section 1 to this paper, and was depicted in the 6–braid example in Figure 10.

Remark 5.6.1 While we have gone to some pains to insure that all clasp arcs are in normal neighborhoods, in the example just given of a G–exchange move the normal neighborhoods are ignored for reasons of space, as they would enlarge the pictures to the point where they would obscure the features that are of interest. Since normal neighborhoods were created by adding many many inessential b –arcs; going the other way, they can also be deleted by an isotopy of the embedded part of \mathcal{CA} . \diamond

To generalize this example we first need to understand the foliated subregions of \mathcal{PA} which lead to sequences of exchange moves that carry a subarc of X_+ over a ‘rooted block and strand tree’. It will be easiest to study them first without the clasp arcs, and then add the clasp arcs later. Then we will need to understand the associated block and strand diagrams. Since the definitions are somewhat detailed, it may be helpful to see them worked out in a few special cases first.

In Figure 58(a) a subarc of X_+ can be pushed across the shaded regions by two exchange moves: the first across the darkly shaded disc (containing the vertices v and w and the singularities s and r), the second across the lighter shaded disc, which contains the vertices v', w' and the singularities s', r' . The support of the first exchange

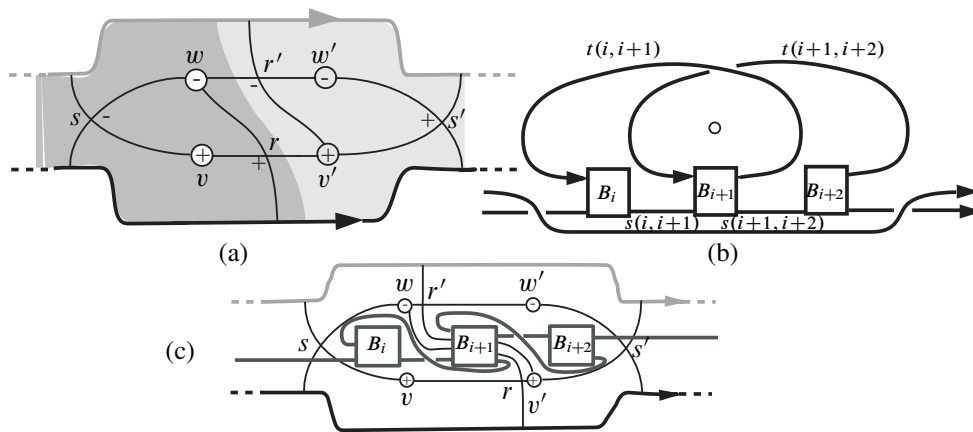


Figure 58: (a) The root foliation on $\mathcal{P}\mathcal{A}$. (b) The root diagram. (c) The two viewpoints are put together.

move, ie the dark shaded region, is a ‘pouch’ which is pierced twice by the axis, at v and w . If the foliation of $\mathcal{P}\mathcal{A}$ contains only essential b -arcs, then the closed braid must wrap around the braid axis (perhaps with many strands traveling along together, say t strands in all) in between the two pierce-points v and w . The first exchange move is a push of a subarc of X_+ across the pouch, crossing the axis twice as it does so at v and w . Then there is a second exchange move across a second pouch, crossing the axis twice at v' and w' . There will also be t' braid strands wrapping about the axis in between v' and w' , and perhaps a braid in between the weighted strands t and t' .

Keeping all this in mind we turn to Figure 58(b). It shows a picture of a root diagram which is part of a closed braid diagram. There are 3 blocks, labeled B_i, B_{i+1}, B_{i+2} and two weighted strands, labeled $t(i, i+1)$ and $t(i+1, i+2)$, traveling around the braid axis. These blocks and strands are there because if not the b -arcs which foliate the pouch would not be essential. The pouch is not shown in this picture.

In Figure 58(c) we put together the information in Figure 58(a) and (b). The region in (a) is to be thought of as a very flexible disc with two pouches. We are looking through the pouches to the block and strand tree which is visible inside them. The braid axis \mathbf{A} pierces the pouches in axis pieces vw and $v'w'$. The darker pouch (we called it C) is the support of the first ab -exchange move in sketch (a). It covers the darker Δ -disc.

The lighter shaded pouch C' covers the lighter Δ -disc. The motion of our subarc of X_+ is the sequence of two exchange moves over the two pouches. The subarc is like a 'handle' which moves over the braid blocks B_i, B_{i+1}, B_{i+2} and the weighted strands.

Now for a formal definition of roots, branches and a block and strand tree. [Figure 59](#) should be helpful. A collection of braid blocks $\{B_1, \dots, B_l\}$ and weighted strands $\{s(1, 2), \dots, s(i, i+1), \dots\}$, $1 \leq i \leq l$, and additional weighted strands $\{t^1(1, 2), \dots, t^{k_1}(1, 2), t^1(2, 3), \dots, t^{k_2}(2, 3), \dots, t^1(i, i+1), \dots, t^{k_i}(i, i+1), \dots\}$ is a *root* if $s(i, i+1)$, $1 \leq i \leq l$ has endpoints at the bottom of B_i and at the top of B_{i+1} ; $t^j(i, i+1)$, $1 \leq j \leq k_i$ has endpoints at the top of B_i and at the bottom of B_{i+1} , and if there exist embedded discs $\Delta(i, i+1, j) \subset S^3$ satisfying the following further conditions:

- The Δ -discs have disjoint interiors. Also, for each $t^j(i, i+1)$, $1 \leq j \leq k_i$ there is one associated disc.
- The braid axis \mathbf{A} intersects each Δ -disc transversally in a single point.
- $\partial\Delta(i, i+1, j) \subset t^j(i, i+1) \cup B_i \cup s(i, i+1) \cup B_{i+1}$. Also $\Delta(i, i+1, j) \cap s(i, i+1) = s(i, i+1)$ and $\Delta(i, i+1, j) \cap t^j(i, i+1) = t^j(i, i+1)$.
- If $\Delta(i, i+1, j)$ and $\Delta(i, i+1, j')$ are distinct Δ -discs which are intersected in succession by some meridian loop of $s(i, i+1)$ then there exist $\Delta(i-1, i, m)$ and $\Delta(i+1, i+2, q)$ such that the unoriented \mathbf{A} intersects first $\Delta(i, i+1, j)$, then $\Delta(i-1, i, m)$ and $\Delta(i+1, i+2, q)$ (in either order); then $\Delta(i, i+1, j')$; then all other Δ -discs.
- For each $\Delta(i, i+1, j)$ there exist a $\Delta(x, y, z)$ with either $(x, y, z) = (i-1, i, m)$ or $(x, y, z) = (i+1, i+2, q)$. Moreover, the unoriented \mathbf{A} intersects in succession $\Delta(i, i+1, j)$, $\Delta(x, y, z)$; then all remaining Δ -discs.

Given a braid structure (\mathbf{H}, \mathbf{A}) , a *radial sphere* is a 2-sphere that is transversally intersected by \mathbf{A} twice and is transverse to all of the disc fibers of \mathbf{H} . An *axis piece* α in a radial sphere S is a closed arc whose interior is transverse to the disc fibers of \mathbf{H} and which has empty intersection with at least one disc fiber. Axis pieces $\alpha, \beta \subset S$, are loop equivalent if $\partial\alpha = \partial\beta = \alpha \cap \beta$ and if $\alpha \cup \beta$ bounds a 2-disc $\delta \subset S$ such that $\delta \cap \mathbf{A} = \emptyset$.

A *branch* is a block B with associated weighted strands t , along with a 2-disc Δ such that:

- Δ is transversally intersected by \mathbf{A} at one point.
- $\partial\Delta = t \cup a$ where $a \subset \partial B$. Specifically, a is an arc made up of three segments, $a = a_T \cup a_S \cup a_B$ where: a_T is on the top of B ; a_S is on the side of B ; and a_B is on the bottom of B .

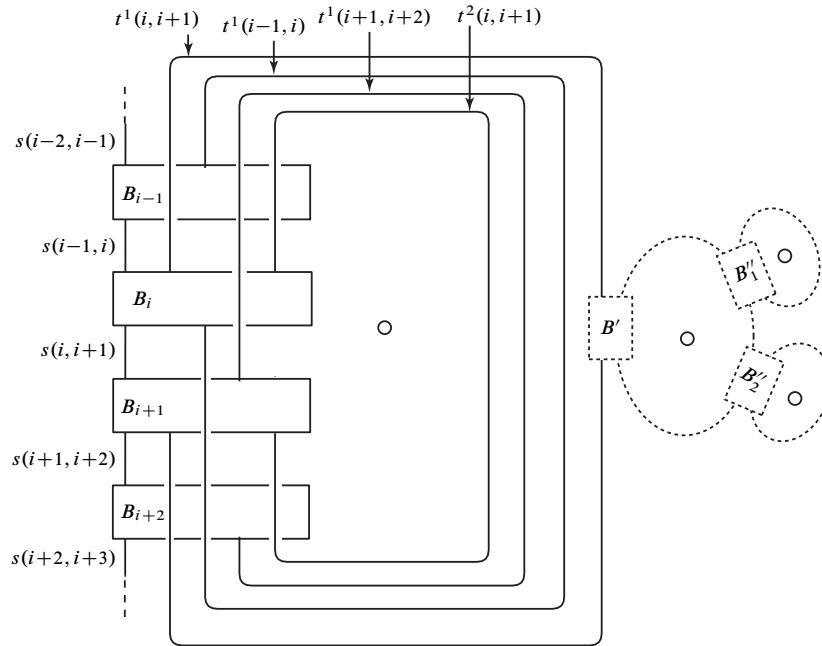


Figure 59: A block and strand tree. The dotted parts will be discussed later.

- $\text{int}(\Delta) \cap B = \emptyset$

The braid block B''_1 and the weighted strand which emerges from it and loops around the axis is an example.

Iterating the construction of attaching a branch to a root, we obtain a *rooted block-strand tree*.

Next we need to understand the foliated subregions of \mathcal{PA} which lead to sequences of exchange moves that carry a handle subarc of X_+ over a rooted block and strand tree. Again it will be helpful to see examples before we give the general definition. We first illustrate how the very simple pouch \mathcal{R}_1 in Figure 57 might itself develop a pouch. As before, we ignore the clasp arcs. Figure 60(a) shows the double pouch foliation of the region \mathcal{R} of Figure 60(a). In the expanded foliation the handle arc (the arc α of Figures 23) cannot be pushed across \mathcal{R} by ab exchanges. Fortunately bb -exchange moves come to the rescue. (See Figures 24 and 25.) The new vertex x has valence 2 and type (b, b) . Lemma 3.3.2, part (2), applies. We can do a bb exchange, and then remove the resulting inessential b -arcs. Now an ab -exchange is possible. As for the corresponding braid picture, the presence of the new vertices x, y means that the dotted root block B' of sketch Figure 59 has grown new branches B''_1




Figure 60: Expanding root foliations to tree foliations

and B''_2 . The move over the tree in Figure 60(a) will be a bb -exchange followed by an ab exchange. In the presence of clasp arcs, the G-exchange move of Figures 10 and 57 will of course become much more complicated when new branches are added. In this regard we note: an important feature is that every time new branches are added the braid index increases. Thus very complicated block and strand trees will only be encountered at very high braid index.

A slightly more complicated example, in Figure 60(b), shows the changes in foliation when we grow new branches in two different ways in the foliation of the region \mathcal{R} of Figure 58(a). The changes are supported inside the region $wsvv'r'w$. The branch associated to the new vertices x and y (resp. x' and y') is attached to the strand joining blocks B_i and B_{i+1} (resp. blocks B_{i+1} and B_{i+2}). The root diagram of Figure 58(b) has changed to a block-strand tree. In the foliation, the growth has all been 'inward'. This time two bb exchanges and two ab exchanges are needed to realize the G-exchange move over the block-strand tree. The reason G-exchange moves can be hard to visualize is because the part of the surface that undergoes the change in foliation is always far away from the block-strand tree in the closed braid. Putting this is another way, the foliated surface points out the way to organize very big sequences of exchange moves, some of which can be quite difficult to see in the closed braid diagram.

Finally, we come to the general definition. Let $\mathbf{S} \subset \mathcal{PA}$ be a complete collection of s -arcs. For present purposes a region $\mathcal{R} \subset \mathcal{PA} \setminus \mathbf{S}$ is either a rectangular shaped subdisc or a subannular region of \mathcal{PA} . Thus, as before, if \mathcal{R} is a subdisc then $\partial\mathcal{R} = s \cup Y_+ \cup s' \cup Y_-$, where: s and s' are subarcs of leaves in the foliation of \mathcal{PA} ; and Y_+ and Y_- are oriented arcs transverse to the foliation in the positive direction. If \mathcal{R} is a subannulus then $\partial\mathcal{R} = Y_+ \cup Y_-$ where Y_{\pm} are oriented curves transverse to

the foliation in the positive direction. Inside these regions we will have the induced foliation. While our regions may be intersected by clasp arcs, we are not concerned with them at this time.

Let $\mathcal{R} \subset \mathcal{PA} \setminus \mathbf{S}$ be as above. Let C be a component of $\mathcal{R} \cap G_{\epsilon, \delta}$, where the graph $G_{\epsilon, \delta}$ was defined pictorially in [Figure 42](#). We say that \mathcal{R} has a *root foliation* if:

- (1) C is homeomorphic to either S^1 or $[0, 1]$.
- (2) If C is homeomorphic to S^1 then \mathcal{R} is an annulus and C is homotopically equivalent to a core circle of \mathcal{R} .
- (3) If C is homeomorphic to $[0, 1]$ then C has an endpoint on Y_{\pm} and an endpoint near Y_{\mp} .

Proposition 5.6.1 *Let $\mathcal{R} \subset \mathcal{PA} \setminus \mathbf{S}$ be a component which is embedded. Assume that all b -arcs are essential. Assume that \mathcal{R} has a root foliation. Then the isotopy which corresponds to pushing a component of \mathbf{Y}_+ across \mathcal{R} is a sequence of ab -exchange moves over a root.*

Proof The isotopy of the braid across a region which has a root foliation can be realized by a sequence of ab exchanges. To see this, notice that at least one pouch \mathcal{P}_i is associated to each such region, and since the regions are crossed in a definite order the pouches can be joined in the same order. The assumption that each b -arc is essential implies that \mathcal{P}_i cannot be removed by isotopy, and is associated to a new braid block or blocks. The union of all of the \mathcal{P}'_i s gives a disc region \mathcal{R} with a subarc of \mathbf{Y}_+ in its boundary. The union of all of the blocks is a *root*. The isotopy of the subarc of \mathbf{Y}_+ across \mathcal{P} is a sequence of exchange moves across this root. This completes the proof. □

Remark 5.6.2 It might happen, in the situation of [Proposition 5.6.1](#), that \mathbf{S} is empty on some component of \mathcal{R} . If that occurs, the component in question will be a standard annulus. Pushing across a standard annulus will be treated in [Section 5.7](#). ◇

A region \mathcal{R} has a *tree foliation* if, after a sequence of bb -exchanges the foliation is reduced to a root foliation.

The next proposition shows how a tree foliation imposes a block decomposition on our two braids:

Proposition 5.6.2 *Let $\mathcal{R} \subset \mathcal{PA} \setminus \mathbf{S}$ be a component whose image is embedded. Assume that $\mathbf{S} \neq \emptyset$. Assume that all b -arcs are essential, and that \mathcal{R} has a tree foliation. Then the isotopy which corresponds to pushing \mathbf{Y}_+ across \mathcal{R} is a sequence of exchange moves over a block-strand tree.*

Proof The isotopy of the braid through the sequence of bb -exchanges corresponds to collapsing a tree to its root. The collapsing is realized by bb -exchange moves followed by braid isotopy to remove inessential b -arcs. After that, an isotopy of \mathbf{Y}_+ across the new \mathcal{R} corresponds to a sequence of aa and ab -exchange moves over this root. \square

Remark 5.6.3 Two remarks are in order. They relate to Propositions 5.6.1 and 5.6.2.

- (1) In both propositions the basic assumption is that the region \mathcal{R} is embedded. Of course this will be the case if no clasp arcs intersect \mathcal{R} . But it will also be the case if (a) clasp arcs intersect \mathcal{R} but no clasp arc pair intersects \mathcal{R} , or (b) there are no long clasp arcs which intersect \mathcal{R} .
- (2) After a G -exchange move across \mathcal{R} , valence 1 vertices may be revealed. See Figure 22. Later, when we get to the final steps in the proof, we will remove them by destabilization. \diamond

We are finally ready to introduce clasp arcs into the picture. In the most general case there will be several related regions $\mathcal{R}_1, \dots, \mathcal{R}_k$ which are intersected by paired clasp arcs. Thus the exchange moves across one \mathcal{R}_i will have to be interrupted midway to do part of an exchange move along an associated \mathcal{R}_j . Let $\{\mathcal{R}_1, \dots, \mathcal{R}_k\} \subset \mathcal{PA}$ be a collection of regions such that each $\text{int}(\mathcal{R}_i)$, $1 \leq i \leq k$, is embedded in $S^3 \setminus (X_+ \cup X_-)$ and each \mathcal{R}_i has a tree foliation. Assume each \mathcal{R}_i has at least one clasp arc with an endpoint on X_+ and at least one clasp arc with an endpoint on X_- , and with the puncture endpoints on b -arcs in \mathcal{R}_i , so that in particular no clasp arc is long. Moreover, assume that the image of $\{\mathcal{R}_1, \dots, \mathcal{R}_k\}$ in \mathcal{CA} is connected. Then the collection of regions is said to have a G -exchange foliation. The motion across a region with a G -exchange foliation is a G -exchange move. As we have shown, in the absence of clasp arcs it would be a sequence of exchange moves which carries a subarc of X_+ over a rooted block and strand tree.

Proposition 5.6.3 Let $\{\mathcal{R}_1, \dots, \mathcal{R}_k\} \subset \mathcal{PA}$ be a collection of regions which, taken together, have a G -exchange foliation. Then an isotopy of X_+ to X_- across the regions $\{\mathcal{R}_1, \dots, \mathcal{R}_k\}$ is realized by a G -exchange move.

Proof Since each region \mathcal{R}_i is embedded, the isotopy of any one arc $\mathcal{R}_i \cap X_+$ across \mathcal{R}_i to $\mathcal{R}_i \cap X_-$ corresponds to a G -exchange move. Moreover, since each region contains clasp arcs having endpoints on both X_+ and X_- and since their image in \mathcal{CA} is connected, we have an interdependence of the isotopies across all the regions. Thus, an isotopy across the collection $\{\mathcal{R}_1, \dots, \mathcal{R}_k\}$ is a G -exchange move. \square

5.7 Pushing across a standard annulus

In constructing the thin annuli, we omitted the case of a standard annulus (with or without clasp arcs). Motions across a standard annulus will in general occur at the very end of the isotopy from $X_+ \rightarrow X_-$, after all singularities that are on clasp arcs have been eliminated, and all short clasp arcs have been removed. To give an example of such a motion, let us suppose that \mathcal{X} is a link of $\mu \geq 2$ components and that the first component, X_+^1 cobounds with X_-^1 a standard annulus with $2k$ vertices, with or without clasp arcs. Figure 61 shows an example without clasp arcs, when $k = 4$.

Figure 61 shows that a stabilization along one of the singular leaves in $G_{-, \epsilon}$, followed by a sequence of $k - 1$ exchange moves, followed by a destabilization, suffices to move X_+^1 to X_-^1 . The total change in braid index is zero. We have a Markov tower similar to the one in Figure 5. If there are clasp arcs, we may assume that they have been pushed into singular leaves. Keeping in mind the stabilization move in Figure 61, where X_+ is pushed across a singular leaf, we note that when there is a pair of clasp arcs a microflype pushes both the black boundary arc and the dotted grey boundary arc, simultaneously, across a pair of singular leaves. After all of the clasp arc pairs have been removed with the help of microflypes, the rest of the standard annulus can be crossed with the help of complexity-reducing destabilizations and exchange moves. In




Figure 61: Pushing across a standard annulus

Section 1 of this paper, in Figure 8, we constructed the Markov tower and the associated template for the crossing of a standard annulus without clasp arcs. If there are clasp arcs, then as was just shown the annulus can be crossed using microflypes, exchange moves and destabilizations, and no other template is needed. The blocks shown in Figure 8 allow for the possibility that other components of \mathcal{X} braid with X^1 . If there were no other components, then the entire move in Figure 8 can be realized by braid isotopy. If there are other components, or if $k > 4$, this may not be the case. The cyclic move exists for every k . The associated templates are the cyclic templates of the MTWS.

6 The proof of the MTWS

The machinery which we need to prove [Theorem 2](#), the MTWS, has been set up. In this section we give our proof. The reader is referred to [Section 1.3](#) for the statement of the theorem.

6.1 Constructing the sequences (1–2), (1–3) and (1–4) and the templates in $\mathcal{T}(m)$

We are given a μ -component oriented link type \mathcal{X} in oriented S^3 or \mathbb{R}^3 . We are also given closed braid representatives $X_+ \in \mathcal{B}(\mathcal{X})$ and $X_- \in \mathcal{B}_{\min}(\mathcal{X})$. Indexing the components of \mathcal{X} as $\mathcal{X}^1, \dots, \mathcal{X}^\mu$, we choose corresponding indices for the components of \mathcal{CA} and X_+ and X_- , so that each annulus \mathcal{A}^j in \mathcal{CA} has $\partial\mathcal{A}^j = X_+^j - X_-^j$.

By [Proposition 2.2.1](#), which describes the ‘basic construction’ for links, we know how to construct the clasp annulus \mathcal{CA} . By the results in [Section 4](#) we may assume that \mathcal{CA} supports a braid foliation. In particular, by [Proposition 4.5.1](#) we may assume that each clasp arc has a normal neighborhood, and that it has been pushed into a union of leaves. Let $c(X_+, X_-, \mathcal{CA}) = (c_1, c_2)$ be the complexity of the triple (X_+, X_-, \mathcal{CA}) , as defined in [Section 5.1](#). Thus c_1 is the number of singularities that are on clasp arcs and c_2 is the number of singularities that are outside normal neighborhoods of the clasp arcs. The first step is:

Construction of the sequences (1–2) and (1–3) We make as many modifications as are possible, using only exchange moves on X_- and only exchange moves and destabilizations on X_+ . With these restrictions, [Corollary 5.1.1](#) tells us that we may find sequences $X_- = X_-^1 \rightarrow \dots \rightarrow X_-^p = X'_-$ as in (1–2) and $X_+ = X_+^1 \rightarrow \dots \rightarrow X_+^q = X'_+$ as in (1–3) of the MTWS such that $c(X'_+, X'_-, \mathcal{CA}')$ is minimal up to exchange moves and destabilizations. Exchange moves preserve braid index and destabilizations reduce it. Exchange moves and destabilizations are both strictly reducing on c_2 . Exchange moves in the presence of clasp arcs (ie the moves ab^* and bb^*) preserve or reduce c_1 . So our sequences are strictly complexity reducing with respect to $c(X_+, X_-, \mathcal{CA})$.

We next turn our attention to the construction of the sequence (1–4) of the MTWS. We begin with a weak version.

Claim For all triplets $(X'_+, X'_-, \mathcal{CA}')$ as above, we may find a sequence

$$(6-1) \quad X'_+ = X^1 \rightarrow X^2 \rightarrow \dots \rightarrow X^r = X'_-$$

which is strictly complexity reducing with respect to $c(X'_+, X'_-, \mathcal{CA}')$ such that every adjacent pair $X^i \rightarrow X^{i+1}$ is related by a single destabilization, exchange move, cyclic

move, G–exchange move or flype. Notice that we do *not* require that the flypes be admissible, so that the sequence may not be non-increasing on braid index.

Remark 6.1.1 (a) We draw the reader’s attention to the similarities between this claim and the discussion at the beginning of [Section 5.5](#) where we dealt with the special case when \mathcal{CA} admits a decomposition into thin annuli. Conceptually, subannuli of \mathcal{CA} that are thin annuli can be thought of as clasp annuli in their own right, because the boundary components represent \mathcal{X} and satisfy the conclusions of [Lemma 2.1.1](#) and [Proposition 4.5.1](#). What this claim is asserting is that \mathcal{CA}' admits a decomposition into clasp subannuli. That is, for each pair (X^i, X^{i+1}) , the link $X^i - X^{i+1}$ bounds a clasp annulus that has its foliation satisfying [Proposition 4.5.1](#). In the same fashion as in [Section 5.5](#) we use the notation \mathcal{S}^i for the clasp subannulus associated with (X^i, X^{i+1}) . The claim then asserts that \mathcal{S}^i corresponds to one of the following isotopies: destabilization, exchange move, cyclic move, G–exchange move or flype.

We interrupt the remark to prove the claim.

Proof of the claim The proof is by induction on $c(X'_+, X'_-, \mathcal{CA}') = (c_1, c_2)$, using lexicographical ordering.

To begin the induction, assume that $(c_1, c_2) = (0, 0)$. Since $c_1 = 0$, it follows that if there are clasp arcs they must all be short, in which case [Lemma 4.2.1](#) says that we may eliminate them by braid isotopy. Therefore \mathcal{CA}' is embedded and foliated without singularities. The foliation then consists entirely of s –arcs. But then X'_+ and X'_- represent the same braid isotopy class, and the MTWS is trivially true.

Inductive hypothesis Assume that the complexity of the family of foliated annuli \mathcal{PA} is (c_1, c_2) . Assume that the claim is true whenever the complexity is less than (c_1, c_2) . There are several cases:

Case (a) $c_1 > 0$, and there are long clasp arcs By the construction which was given in [Lemma 5.3.1](#), we find a family \mathcal{S} of subannuli of \mathcal{PA} such that each component supports a thin foliation or is a standard annulus. Some of these components may be trivially foliated, but since there is at least one long clasp arc they are not all trivially foliated. If \mathcal{S} is the union of standard annuli then we can push across \mathcal{S} by a cyclic move as shown in [Section 5.7](#). Otherwise, \mathcal{S} is not trivially foliated but does contain s –arcs. Then, by [Proposition 5.3.2](#) we may use a sequence of flypes to push X'_+ across \mathcal{S} . (Note: the flypes may not be admissible.) This reduces complexity. By the induction hypothesis, the claim is true.

Case (b) $c_1 > 0$, but there are no long clasp arcs In the presence of clasp arcs, but the absence of long clasp arcs, [Corollary 5.1.1](#) tells us that there must be at least one G–exchange region. Pushing across it reduces complexity.

Case (c) $c_1 = 0$, but c_2 is arbitrary We have already removed all short clasp arcs, so there are none. But then there also are no normal neighborhoods. The annulus \mathcal{CA} is embedded. By [Corollary 5.1.2](#), each component of \mathcal{PA} will either be either an annulus which is trivially foliated (in which case we are done as before) or there is a component which is a standard annulus without clasp arcs. In the latter case the motion is realized by a cyclic template. Again, the claim is true.

While we have found a strictly monotonic complexity-reducing sequence of closed braids (sequence [\(6–1\)](#)) joining X'_+ to X'_- , we have not established sequence [\(1–4\)](#) of the MTWS because we have not established that the braid index is non-increasing.

To deal with this deficiency we introduce the *augmented complexity function*

$$C(X'_+, X'_-, \mathcal{CA}) = (c_0, c_1, c_2),$$

where $c_0 = b(X'_+)$ and c_1 and c_2 are as before. This is the complexity function which is referred to in the statement of [Theorem 2](#). Notice that:

- Since the sequence [\(1–2\)](#) is complexity-reducing with respect to $c(X_+, X_-, \mathcal{CA})$ and since exchange moves preserve braid index, it is complexity-reducing with respect to $C(X_+, X_-, \mathcal{CA})$.
- Since the sequence [\(1–3\)](#) is complexity-reducing with respect to $c(X_+, X_-, \mathcal{CA})$ and since destabilizations reduce braid index, it is complexity-reducing with respect to $C(X_+, X_-, \mathcal{CA})$.

Now notice that we can use the braid index entry of our augmented complexity function to pick out a subsequence of [\(6–1\)](#):

$$(6-2) \quad X_+ = X^{i_1} \rightarrow X^{i_2} \rightarrow \dots \rightarrow X^{i_r} = X_-$$

such that $0 < i_1 < i_2 < \dots < i_r$ and $b(X^{i_j}) \geq b(X^{i_{j+1}})$, and for any X^k of [\(6–1\)](#) with $i_j < k < i_{j+1}$ we have $b(X^k) > b(X^{i_j})$. Since $X^{i_j} - X^{i_{j+1}}$ cobound a clasp annulus \mathcal{S}^{i_j} which is the union of the clasp subannuli mentioned in [Remark 6.1.1](#) our augmented complexity function is non-increasing on the triple $(X^{i_j}, X^{i_{j+1}}, \mathcal{S}^{i_j})$. The properties of sequence [\(1–4\)](#) follow. \square

Remark 6.1.1 (b) Our decomposition of \mathcal{CA} into clasp subannuli $\mathcal{S}^{i_1} \cup \mathcal{S}^{i_2} \cup \dots \cup \mathcal{S}^{i_r}$ also shows us the origin of templates in $\mathcal{T}(m)$. Specifically, an \mathcal{S}^{i_j} may be a clasp subannulus of \mathcal{CA} that is the union of any of the types of foliations of [Section 5.3.1](#), [5.5](#), [5.6](#) or [5.7](#) along with destabilization. This gives us a natural decomposition for $\mathcal{T}(m)$. Let $\mathcal{T}(m, n)$ be the subset of all templates in $\mathcal{T}(m)$ whose initial braid has

braid index m and whose final braid has braid index $n \leq m$. The union of the subsets $\mathcal{T}(m, n)$, $n \leq m$ determines $\mathcal{T}(m)$ because:

$$(6-3) \quad \mathcal{T}(m) = \mathcal{T}(m, m) \sqcup \mathcal{T}(m, m-1) \sqcup \cdots \sqcup \mathcal{T}(m, 1)$$

Thus, employing the procedure given in [Section 5.4](#) for constructing a template, we can use each triple $(X^{i_j}, X^{i_{j+1}}, \mathcal{S}^{i_j})$ to construct a template in $\mathcal{T}(b(X^{i_j}), b(X^{i_{j+1}}))$. Moreover, using the triple $(X_+, X^{i_{j+1}}, \mathcal{S}^{i_1} \cup \cdots \cup \mathcal{S}^{i_j})$ we can construct a template in $\mathcal{T}(m, b(X^{i_{j+1}}))$. \diamond

Our proof of the MTWS is almost complete, except for two missing facts: the proof that no block in any template in $\mathcal{T}(m)$ has full braid index, and the proof that $\mathcal{T}(m)$ is finite. We will deal with the former in [Corollary 6.2.1](#) and with the latter in [Proposition 6.3.1](#).

We emphasize that the destabilization template is *not* contained in $\mathcal{T}(m)$. On the other hand, for the remainder of the paper, it will be convenient to think of exchange moves, admissible flypes and G-exchange moves as templates in $\mathcal{T}(m)$, rather than as separate moves.

6.2 Cleaning up the templates in $\mathcal{T}(m)$

Our goal is to prove that no block in any template in $\mathcal{T}(m)$ has full braid index, however we will do that in a setting that yields additional information for later use. The concept of a block-strand diagram is, at this moment, rather loose. One can imagine that blocks could be slid around and amalgamated with one-another, and also subdivided. We would like them to be more canonical. But even if we succeed to make them canonical, a direct attack on the problem seems difficult. We use a more roundabout approach. When we organized the moves of the MTWS into three separate subsequences, we separated the subsequence (1–2), which requires only exchange moves, and the subsequence (1–3), which requires only exchange moves and destabilizations, from subsequence (1–4), which requires admissible flypes as well as the more general moves in $\mathcal{T}(m)$. We will see that, as a consequence, the templates in $\mathcal{T}(m)$ have no blocks whose braid index is equal to the braid index of the diagram. In fact, we will be able to show more: that the templates of $\mathcal{T}(m)$ are ‘consolidated’ in a sense that will be made precise below, with the braid index condition on the blocks being just one of the nice properties of consolidated templates.

Let \mathcal{D} be a block-strand diagram having blocks B and B' . We use the notation $t, b \subset \partial B$ (resp. $t', b' \subset \partial B'$) for the top and bottom of B (resp. top and bottom of B'). Recall our use of the term ‘amalgamating blocks’ in [Section 5.3.2](#). We now say that B'

(resp. B) can be amalgamated into B (resp. B') if via a braid isotopy we can move $t' \subset b$ (resp. $b \subset t'$). (These are two ways of looking at the same phenomenon.) We say the blocks of \mathcal{D} are *consolidated* if no two blocks of \mathcal{D} can be amalgamated. Next, let s and s' be strands in \mathcal{D} having endpoints on common blocks (still called) B and B' . We say that strands s and s' are *topologically parallel* if there exists a rectangular disc R such that $R \cap \mathcal{D} = \partial R = s \cup \beta_b \cup s' \cup \beta_{t'}$ where $\beta_b \subset b$ and $\beta_{t'} \subset t'$. Note that if strands s, s' are topologically parallel, it could happen, for example, that s travels from B to B' without winding around the axis, but s' winds about the axis as it does so. To rule out this sort of complication, we say that s and s' are *braid parallel* if they are topologically parallel and $R \cap \mathbf{A} = \emptyset$. We say that \mathcal{D} is *consolidated* if the blocks of \mathcal{D} are consolidated and all of its topologically parallel strands are braid parallel.

Proposition 6.2.1 *Given a template (D_+, D_-) with D_- having minimal braid index there exists a new template (D'_+, D'_-) , where D'_+ is obtained from D_+ via a sequence of exchange moves and destabilizations; also D'_- is obtained from D_- via a sequence of exchange moves; also D'_+ and D'_- are both consolidated block-strand diagrams. In particular, we can assume that the block-strand diagrams in every template in $\mathcal{T}(m)$ are consolidated.*

Proof Let R be a rectangular disc that demonstrates that strands s and s' are topologically parallel for, say, D_+ . We look at the induced foliation of \mathbf{H} on R . We can make the standard assumptions about R being transverse to \mathbf{A} and all but finitely many disc fibers of \mathbf{H} being transverse to R . For the finitely many non-transverse disc fibers, we can assume that each one contains a single saddle singularity. We can then argue that there are no leaves in the foliation that are circles. The foliation of R can then be seen as a union of $aa-$, $ab-$, $bb-$ and $sb-$ tiles. In particular, since the two sides β_b and $\beta_{t'}$ can be assumed to be s -arcs, R will either be trivially foliated by s -arcs, or will contain sb -singularities.

In the case where R is not trivially foliated we assign an arbitrary orientation to R . We can then talk about the $G_{\epsilon, \delta}$ graphs of R . We wish to invoke the statements of [Proposition 5.1.1](#) to simplify the graph of R through a sequence of exchange moves and destabilizations. Since R is embedded all discs corresponding to those satisfying statements (5), (6) and (7) of [Proposition 5.1.1](#) will automatically be good. It is clear that after applying [Proposition 5.1.1](#) repeatedly, R will be trivially foliated, thus our resulting strands will be braid parallel.

We have a similar argument in the case where our strands are in D_- . The only change is that our application of [Proposition 5.1.1](#) cannot yield any destabilizations because D_- has minimal braid index. This establishes that we can replace a template (D_+, D_-) with

(D'_+, D'_-) via the sequence mentioned in statements 1 and 2, so that any topologically parallel strands in D'_+ or D'_- are braid parallel.

Next, suppose there are two blocks, $B(1)$ and $B(2)$, in the template (D'_+, D'_-) such that we can amalgamate $B(2)$ into $B(1)$ in, say D'_- . (These blocks may be moving or fixed.) Such an amalgamation occurs because all of the strands entering the top of $B(2)$ are braid parallel and start at the bottom of $B(1)$. If we could also amalgamate them in D'_+ , then there would be nothing to stop us from doing the amalgamation in both block-strand diagrams—our template would still be a template and would be simplified. The obstruction to amalgamating $B(2)$ into $B(1)$ in D'_- is that there is no similar amalgamation of $B(2)$ into $B(1)$ in D'_+ . But, since the strands entering the top of $B(2)$ are all braid parallel in D'_- , they must be topologically parallel in D'_+ . But, we are assuming that all topologically parallel strand in D'_+ are braid parallel. So there is no obstruction. Thus, the block amalgamation is possible. After an iteration of amalgamations we may assume that (D'_+, D'_-) is consolidated. \square

Corollary 6.2.1 *Let $\mathcal{T} = (D_+, D_-)$ be a template in $\mathcal{T}(m)$. Then there is no block $B \subset D_-$ such that $b(B) = b(D_-)$.*

Proof Suppose that D_- does contain a block B such that the number of strands entering the top of B or leaving the bottom of B is equal to the braid index of D_- . Then we can amalgamate all other blocks in D_- with B . That is, we can push all braiding of strands immediately before and after B into B , thus making all such braiding parallel. This allows us to amalgamate any block immediately before or after B in D_- with B . We then iterate this procedure to amalgamate all blocks in D_- with B . After this comprehensive amalgamation the consolidated block-strand diagram D'_- will have a single block and will still carry X_- . By [Proposition 6.2.1](#) we will thus have gone from X_+ (carried by D_+) to X_- (carried by the consolidated D'_-) through a sequence of braid isotopies, exchange moves and destabilizations. But note that these are the moves that made up the sequences (1–3) and (1–2). This means that our X'_+ of sequence (1–3) is braid isotopic to our X'_- of sequence (1–2). Thus, sequence (1–4) is vacuous and there is no template to construct. We conclude that we can assume that the templates of $\mathcal{T}(m)$, which are all constructed using sequence (1–4), have no blocks of full braid index. \square

Our main goal for this section, the proof that no block in a template in $\mathcal{T}(m)$ has full braid index, has been achieved. In so-doing, we learned a little bit more: we may assume that the templates of $\mathcal{T}(M)$ are consolidated. In fact, there is still is one more step we can take to make them more canonical. Referring back to the discussion of *Fixed Blocks* in [Section 5.3](#), in the first two paragraphs we used the occurrence of singularities

in the fibration \mathbf{H} to designate fixed blocks. In the third paragraph we then specified how some of these just-constructed fixed blocks can be amalgamated. Essentially, we were observing the existence of “missing blocks”. We make this concept and its implications rigorous below. While this is not necessary for pushing our finiteness argument forward, it does tie up this loose end coming from [Section 5.3](#). Our brief detour will end at the end of this subsection.

Let \mathcal{D} be a block-strand diagram having braid structure (\mathbf{A}, \mathbf{H}) . Suppose there exist 3-balls B in the braid structure having a $\Delta \times [0, 1]$ structure with the discs $\Delta \times p$ contained in fibers of \mathbf{H} , $0 \leq p \leq 1$. Assume that $\text{int}(B)$ contains at least one block of \mathcal{D} and intersects only the strands of \mathcal{D} , but that $B \setminus (B \cap \mathcal{D})$ is not homeomorphic to an interval cross a $2k$ -punctured 2-sphere. We call such a 3-ball B a *missing block* in \mathcal{D} .

Let (D_+, D_-) be a template that is consolidated. Then (D_+, D_-) is an *optimal template* if for every missing block of D_- , the ambient isotopy of S^3 that takes D_- to D_+ does not result in an missing block of D_+ . In fact, we claim that we may assume that all templates of $\mathcal{T}(M)$ are optimal.

To see this, let $(D_+, D_-) = \mathcal{T} \in \mathcal{T}(m)$ and let $\{B^1, \dots, B^k\}$ be a complete listing of all the blocks in \mathcal{T} . As before, $b(B^i)$ is the braid index of the block B^i . We define the complexity of \mathcal{T} to be a lexicographically ordered 2-tuple $(k, \sum_1^k (n - b(B^i)))$. Now, if B_- is a missing block of D_- that is taken to a missing block B_+ of D_+ by the ambient isotopy of S^3 that relates D_- and D_+ then by assumption B_+ must contain at least one of the B^i blocks plus something else. That something else could be either additional blocks or strands of D_+ that do not intersect B^i in the set $B_+ \cap D_+ \subset B_+$. We now replace $\mathcal{T} \in \mathcal{T}(m)$ with a new template

$$\mathcal{T}' = ((D_+ \setminus \text{int}(B_+)) \cup B_+, [(D_- \setminus \text{int}(B_-)) \cup B_-]).$$

Notice that the complexity of \mathcal{T}' is less than the complexity of \mathcal{T} : if B_+ contains more than one block then \mathcal{T}' has fewer blocks than \mathcal{T} ; and, if B_+ contains just one block, $B^i \subset D_+$, along with some number of extra strands then $m - b(B_+) < m - b(B^i)$. In both situations the complexity is reduced.

It is not hard to see that employing the recipe given in [Section 5.4](#) for producing a (D_+, D_-) from an arbitrary triplet (X_+, X_-, \mathcal{CA}) does not necessarily produce an optimal template. We emphasize this point by ending this subsection with a useful definition which we will need in our finiteness argument.

Let $\{l_+^1, \dots, l_+^h\}$ be a listing of the components of $X_+ \cap \mathcal{CA}_{\text{tiled}}$ and, similarly, $\{l_-^1, \dots, l_-^h\}$ be a listing of the components of $X_- \cap \mathcal{CA}_{\text{tiled}}$ such that the isotopy

across \mathcal{CA} has l_+^i being taken to l_-^i for $1 \leq i \leq h$. Specifically, there is a component $\mathcal{R}^i \subset \mathcal{CA}_{\text{tiled}}$ that has $l_{\pm}^i \subset \partial \mathcal{R}^i$, $1 \leq i \leq h$. Now, if we collapse each of the s -arcs of $\mathcal{CA} - \mathcal{CA}_{\text{tiled}}$ to a point we can conceptually think of X_- as being obtained from X_+ by replacing every l_+^i in X_+ with l_-^i , or to abuse notation $[(X_+ \setminus (l_+^1 \cup \dots \cup l_+^h)) \cup (l_-^1 \cup \dots \cup l_-^h)] = X_-$. (We will continue this abuse of notation below.)

Recall that by Morton [26, Theorem 1] the conjugacy class of X_+ is determined by the link type $X_+ \sqcup \mathbf{A}$. We say that a triplet has *unnecessary motion* if its complexity can be reduced in the following manner. Suppose now that there exists a proper subset $\{l_+^{i_1}, \dots, l_+^{i_r}\} \subset \{l_+^1, \dots, l_+^h\}$ such that $[(X_+ \setminus (l_+^{i_1} \cup \dots \cup l_+^{i_r})) \cup (l_-^{i_1} \cup \dots \cup l_-^{i_r})] \sqcup \mathbf{A}$ has the same link type as $X_+ \sqcup \mathbf{A}$. Then let \mathcal{CA}' be the clasp annulus that is obtained from \mathcal{CA} by replacing each component $\{\mathcal{R}^{i_1}, \dots, \mathcal{R}^{i_r}\} \subset \mathcal{CA}_{\text{tiled}}$ with an s -band that is parallel to the strands $\{l_-^{i_1}, \dots, l_-^{i_r}\}$. Our new clasp annulus is still cobounded by $X_+ - X_-$ and $C(X_+, X_-, \mathcal{CA}') < C(X_+, X_-, \mathcal{CA})$.

6.3 The set $\mathcal{T}(m)$ is finite

We return to the main thread of the argument. There is still one very big unanswered question: how do we know that $\mathcal{T}(m)$ is finite? Proving that it is, in fact, finite, is the main result in this final subsection. See Proposition 6.3.1. The proof of the MTWS will follow immediately.

We begin with an investigation of restrictions on the tiling of \mathcal{CA} which follow from the fact that it is a topological annulus. Let W be a vertex in the foliation of \mathcal{PA} . The *valence* v of W is the number of singular leaves which have an endpoint on W . Traveling around W in either direction one will encounter a cyclically ordered sequence (t_1, t_2, \dots, t_v) of non-singular leaf types, where each t_i is either a_{\pm} or b . If this sequence includes α edges of type a_{\pm} and $v - \alpha$ edges of type b , then we say that W has *type* $(\alpha, v - \alpha)$. Let $V(\alpha, v - \alpha)$ be the number of vertices of valence v and type $(\alpha, v - \alpha)$ in the foliation of \mathcal{PA} . Let $E(s)$ be the number of s -arcs which are boundary edges of a band of s -arcs in the foliation of \mathcal{PA} . For example, in the foliated annulus of Figure 43 we have $E(s) = 4$.

Lemma 6.3.1 *The vertices in the foliation of \mathcal{CA} satisfy the following restriction:*

$$(6-4) \quad V(1, 1) + 2V(1, 0) + 2V(0, 2) + V(0, 3) = 2E(s) + V(2, 1) + 2V(3, 0) + \sum_{v=4}^{\infty} \sum_{\alpha=0}^v (v + \alpha - 4)V(\alpha, v - \alpha),$$

with every term on both the LHS and the RHS non-negative. Notice that every vertex type appears and is counted, with the following 3 exceptions: vertices of type $(1, 2)$, $(2, 0)$ and $(0, 4)$ do not appear in this equation because in all three cases the coefficient $(v + \alpha - 4) = 0$.

Proof On each annular component of \mathcal{PA} the foliation determines a cellular decomposition which goes over to a cellular decomposition of S^2 on shrinking the 2 boundary components to points. Letting V , E and F be the number of vertices, edges and tiles, the fact that $\chi(S^2) = 2$ shows that on each component of the foliated surface \mathcal{CA} we have $V + 2 - E + F = 2$. Each tile has four edges and each edge is an edge of exactly 2 tiles, so that $E = 2F$. Combining this with the previous equation we learn that $2V = E$. Let $E(a)$, $E(b)$, and $E(s)$ be the number of a -, b -, and s -edges, where we count both a_+ and a_- edges as being type a . Then

$$(6-5) \quad 2V = E(a) + E(b) + E(s).$$

Since

$$(6-6) \quad V = \sum_{v=1}^{\infty} \sum_{\alpha=0}^v V(\alpha, v-\alpha)$$

$$(6-7) \quad E(a) = \sum_{v=1}^{\infty} \sum_{\alpha=0}^v \alpha V(\alpha, v-\alpha)$$

$$(6-8) \quad 2E(b) = \sum_{v=1}^{\infty} \sum_{\alpha=0}^v (v-\alpha) V(\alpha, v-\alpha)$$

we may combine equations (6-5) through (6-8) to obtain

$$(6-9) \quad \sum_{v=1}^{\infty} \sum_{\alpha=0}^v (4-v-\alpha) V(\alpha, v-\alpha) = 2E(s).$$

Rearranging terms, we have proved the lemma. \square

Our next goal is to learn which of the terms in Equation (6-4) are bounded, and which terms can grow without bound, when we fix the braid indices $b(X_+)$ and $b(X_-)$. For a triple (X_+, X_-, \mathcal{CA}) of minimal complexity with $(b(X_+), b(X_-)) = (m, n)$, let $\mathcal{N} \subset \mathcal{PA}$ be the union of all normal neighborhoods of the clasp arcs and let \mathcal{N}' be its complement in \mathcal{PA} . As before, let $V(\alpha, \beta)$ be the number of vertices of valence v and type (α, β) in the foliation of \mathcal{PA} . Let $V'(\alpha, \beta)$ denote the number of vertices which are in \mathcal{N}' and contribute to $V(\alpha, \beta)$.

Our primary method for accounting for the braid index of the braid $X_+ \sqcup X_-$ is by counting the number of b -arcs in non-singular disc fibers of \mathbf{H} . This number is independent of the choice of the fiber. Let \mathbf{b} be a b -arc in $\mathcal{CA} \cap H_\theta$. Then \mathbf{b} has 2 endpoints on \mathbf{A} and divides H_θ into 2 subdiscs, Δ_1 and Δ_2 . In Section 4.5 we created normal neighborhoods of the clasp arcs, at the expense of introducing lots of b -arcs into the foliation of \mathcal{CA} . Those b -arcs are ones that can always be removed, if we wish to do so, by reversing the isotopy. We need to make sure that they do not mess up our counting process. For that reason, we now sharpen the concept of an essential b -arc. (Up to this point, a b -arc was defined to be essential if it was not inessential.)

- A b -arc \mathbf{b} is *strongly essential* if the interiors of Δ_1 and Δ_2 each contain either an a_\pm -arc or an s -arc. Note that in the creation of normal neighborhoods we never introduced new strongly essential b -arcs.
- A b -arc \mathbf{b} is *weakly essential* if the interiors of Δ_1 and Δ_2 do not contain a_\pm or s arcs. The reason these are only ‘weakly essential’ is because, if their vertex endpoints are not adjacent on \mathbf{A} , then one of discs Δ or Δ' contains a family of b -arcs and the innermost member of that family will be inessential. Thus a weakly essential b -arc can always be changed to an inessential b -arc by an appropriate sequence of isotopies.
- A b -arc \mathbf{b} is *inessential* if v and w are adjacent on \mathbf{A} . This implies that neither Δ_1 nor Δ_2 contains an a_ϵ -arcs, an s -arc or a b -arc.

The construction of normal neighborhoods introduces numerous weakly essential b -arcs in \mathcal{N} . There may also be weakly essential b -arcs in \mathcal{N}' , e.g. \mathbf{b} may be in \mathcal{N}' but Δ and/or Δ' may contain b -arcs that are in \mathcal{N} . For this reason, we refine our count. We already defined $V(\alpha, \beta)$. Recall that $V'(\alpha, \beta)$ is the number of vertices which are in \mathcal{N}' and contribute to $V(\alpha, \beta)$. Let $V^e(\alpha, \beta^e)$ denote the number of vertices which are adjacent to α a -arcs and β^e strongly essential b -arcs. We will obtain the bounds that we need from $V(\alpha, \beta)$, $V'(\alpha, \beta)$ and $V^e(\alpha, \beta^e)$

Lemma 6.3.2 *The following hold for the individual terms in Equation (6–4), when the braid indices m and n are fixed and (X_+, X_-, \mathcal{CA}) has minimal complexity through the application of exchange moves (applied to X_+ and X_-) and destabilizations (applied to X_+):*

- (1) $V(1, 0)$ is zero.
- (2) $V(\alpha, v - \alpha)$ is zero for $\alpha \geq 2$.
- (3) If the complexity is minimal, then $V'(0, 2) = V'(0, 3) = 0$.
- (4) $V^e(1, 2)$ and $V^e(0, 4)$ are bounded.

- (5) $E(s)$ is bounded.
- (6) $V(1, 1)$ is bounded.
- (7) $\sum_{\beta=3}^{\infty} V^e(1, \beta^e)$ is bounded.
- (8) $\sum_{\beta=5}^{\infty} V^e(0, \beta^e)$ is bounded.

Remark 6.3.1 We do not obtain bounds for $V(0, 2)$, $V(0, 3)$, $V(0, 4)$, $V(1, 2)$, $V'(1, 2)$ and $V'(0, 4)$, and indeed they cannot be bounded. With this observation, we have accounted for every possible term $V(\alpha, \beta)$ and $V'(\alpha, \beta)$ for $\alpha, \beta \geq 0$. \diamond

Proof We consider the various inequalities in order:

Proof of (1) and (2) These two statements follow from [Corollary 5.1.1](#). Specifically, statement 1. of [Corollary 5.1.1](#) forces $V(1, 0)$ to be zero. If there is a vertex v_ϵ that is adjacent to the boundary of \mathcal{PA} that contributes to the count of $V(\alpha, v - \alpha)$ for $\alpha \leq 2$ then there are two a_ϵ -arcs, $a_\epsilon^1, a_\epsilon^2$, adjacent to v_ϵ that are not isotopic to each other in the foliation. The subdisc, $\Delta(v_\epsilon)$, that $a_\epsilon^1 \cup a_\epsilon^2$ splits off in \mathcal{PA} can only contain γ_ϵ clasp arcs. The absence of $\gamma_{-\epsilon}$ makes $\Delta(v_\epsilon)$ embedded when considered in \mathcal{CA} . This implies that all components of $G_{\epsilon, \delta} \cap \Delta(v_\epsilon) \subset \Delta(V_\epsilon)$ are simply connected, ie by statement 3. of [Corollary 5.1.1](#) there are no loops in the graphs $G_{\epsilon, \delta} \cap \Delta(v_\epsilon)$. This further implies that there exists a path α in either $G_{-\epsilon, +}$ or $G_{-\epsilon, -}$ satisfying the assumptions in statement 2. of [Corollary 5.1.1](#). Since our subdisc does not contain a $\gamma_{-\epsilon}$ clasp arc, we get a contradiction of statement 2. We conclude that $V(\alpha, v - \alpha)$ is zero.

Proof of (3) By [Proposition 5.1.1](#), if a vertex W has type $(0, 2)$ or $(0, 3)$, and if its link is a good disc, then that vertex can be removed by changes of foliation followed by exchange moves, reducing complexity. However, we are assuming minimum complexity. Therefore no such W exists, unless $\text{link}(W)$ is not good, ie it intersects \mathcal{N} .

Proof of (4) Suppose that W is a vertex that contributes to the count of $V(1, 2)$. It is near X_+ or X_- , so assume that it is adjacent to a strongly essential b -arc. An example was given in [Figure 58\(a\)](#). In this situation we showed in [Figure 58\(b\)](#) that the corresponding embedding is a root. Each new block in the root contributes at least 1 to the braid index. This forces the braid index of $X_+ \sqcup X_-$ to grow, contradicting our assumption that it is fixed at $m + n$. A similar argument applies to $V(0, 4)$.

Proof of (5) Suppose that $E(s)$ is unbounded. Then there will be an unbounded number of singularities of type $a_\epsilon s$ and/or sb . Notice that there is a direct correspondence between the number of $a_\epsilon s$ -singularities and the number of vertices contributing to $V(1, 0)$ and $V(\alpha, v - \alpha)$ for $\alpha \geq 2$. But by statements (1) and (2) we know that $V(1, 0)$

and $V(\alpha, v - \alpha)$, $\alpha \geq 2$ are bounded. Thus, the only way that $E(s)$ can grow is if there is a growth in the number of singularities of type sb .

Previously in Section 5.3.1 we defined the notion of a complete collection of s -arcs for thin annuli. We have a similar notion for \mathcal{CA} and \mathcal{PA} . A (possibly empty) family of s -arcs $\mathbf{S} = \{s_1, \dots, s_l : s_i \subset \mathcal{CA}\}$ is a *complete collection of s -arcs* in \mathcal{CA} if: (i) no two s -arcs in the collection split off a sub-band of \mathcal{CA} that is foliated entirely by s -arcs; and, (ii) for any other s -arc $s \subset \mathcal{CA}$ there exists an $s_i \in \mathbf{S}$ such that $s \cup s_i$ splits off a sub-band of \mathcal{CA} that is foliated entirely by s -arcs. It is immediate that cutting \mathcal{CA} open along a complete collection \mathbf{S} of s -arcs decomposes \mathcal{CA} into a disjoint union of components that contain in one-to-one correspondence the components of $\mathcal{CA}_{\text{tiled}}$ and bands of s -arc.

Now fixing \mathcal{CA} , let \mathbf{S} to be a complete collection of s -arcs in the foliation. Assume that $E(s)$ can be arbitrarily large. We know that the number of components of $E(s)$ which have angular length $\geq 2\pi$ is bounded, because each time that a band of s -arcs travels completely around \mathbf{A} it contributes 1 to $b(X_+)$ and 1 to $b(X_-)$. This means that all of the growth in the cardinality $|\mathcal{CA} \setminus \mathbf{S}|$ (which is the same as $|\mathcal{CA}_{\text{tiled}}|$) comes from components \mathcal{C} with the angular length of $X_+ \cap \mathcal{C}$ and $X_- \cap \mathcal{C}$ being strictly less than 2π . In particular, for such a component the set of fibers H_θ for which $H_\theta \cap (X_+ \cap \mathcal{C}) \neq \emptyset$ coincides with the set of fibers for which $H_\theta \cap (X_- \cap \mathcal{C}) \neq \emptyset$, since $X_+ \cap \mathcal{C}$ and $X_- \cap \mathcal{C}$ have their endpoints on the same two bs -singularities. We are thus seeing a growth in the components of $\mathcal{CA}_{\text{tiled}}$ that are characterized by the condition:

- There exists an $H_\theta \in \mathbf{H}$ with $H_\theta \cap (X_\epsilon \cap \mathcal{C}) = \emptyset$, $\epsilon = +$ and $-$.

We conclude that the following holds: Let $\{\mathcal{C}_1, \dots, \mathcal{C}_j\} \subset \mathcal{CA} \setminus \mathbf{S}$ be the set of components for which there exists an $H_\theta \in \mathbf{H}$ such that $H_\theta \cap (X_\pm \cap \mathcal{C}_j) = \emptyset$ for all $1 \leq j \leq J$. (Note that there may be a different H_θ for each \mathcal{C}_j .) If $E(s)$ grows then the index J must also grow.

Now suppose we have a subcollection of such components $\{\mathcal{C}_{i_1}, \dots, \mathcal{C}_{i_L}\} \subset \{\mathcal{C}_1, \dots, \mathcal{C}_J\}$ such that for every $H_\theta \in \mathbf{H}$ we have that $H_\theta \cap [\bigcup_{1 \leq j \leq L} \mathcal{C}_{i_j}]$ contains a strongly essential b -arc. Then this subcollection contributes to the braid index of $X_+ \sqcup X_-$. Since our braid index is fixed such subcollections cannot grow in number. Thus, as the index J grows we can only have growth in a subset $\{\mathcal{C}'_{i_1}, \dots, \mathcal{C}'_{i_{L'}}\} \subset \{\mathcal{C}_1, \dots, \mathcal{C}_J\}$ with the property that there exists a fixed disc fiber H'_θ with $H'_\theta \cap [\bigcup_{1 \leq j \leq L'} \mathcal{C}'_{i_j}]$ being a union of weakly inessential b -arcs. Pushing these weakly essential b -arcs off of H'_θ we see that the isotopy across the components $\{\mathcal{C}'_{i_1}, \dots, \mathcal{C}'_{i_{L'}}\}$ can be achieved in the complement of \mathbf{A} . Thus, our original triplet (X_+, X_-, \mathcal{CA}) has unnecessary motion

and was not of minimal complexity. But, we are assuming that we started with minimal complexity. Thus, $E(s)$ cannot grow.

Proof of (6) We refer back to [Proposition 5.3.1](#) to establish the boundedness of $V(1, 1)$. If we consider the construction of \mathcal{S}_+ and \mathcal{S}_- we know from (1) and (2) of [Proposition 5.3.1](#) that there can be at most m type 1_d regions in \mathcal{S}_+ and n type -1_d regions in \mathcal{S}_- . Each type 1_d (resp. -1_d) region in \mathcal{S}_+ (resp. \mathcal{S}_-) accounts for two vertices that contribute to the count of $V(1, 1)$. So if $V(1, 1)$ is able to increase it must be through either the occurrence of type $0_{d,1}$ and $0_{d,2}$ regions in \mathcal{S}_+ and \mathcal{S}_- , or through the occurrence of a subregion like those in [Figure 57](#). (The distinguishing feature between the two cases is whether the clasp arc intersecting the region is long or not.) An increase in the latter will not increase $b(X_+) + b(X_-)$ since it will be associated with a G-exchange move. So we need only consider growth in \mathcal{S}_+ and \mathcal{S}_- .

If $V(1, 1)$ is allowed to grow arbitrarily large then there will be one annular component of \mathcal{PA} that will contribute an arbitrarily large number of vertices to the count of $V(1, 1)$. Thus, we will have a single component of \mathcal{PA} which will contribute an arbitrarily large number of type $0_{d,1}$ or $0_{d,2}$ regions to the construction of either \mathcal{S}_+ or \mathcal{S}_- . Since this growth occurs on a single component of \mathcal{PA} , any two type 0_d regions on that \mathcal{PA} component will be adjacent to a common s -band, (an assumption which is needed to apply (3) and (3') of [Proposition 5.3.1](#)).

Focusing on \mathcal{S}_+ , we know from (2) of [Proposition 5.3.1](#) that we cannot have growth in the number of pairs of type $0_{d,1}$ regions and type $0_{d,2}$ regions that intersect each other. From (5) of [Proposition 5.3.1](#) we know that we cannot have a single region (see \mathcal{R}_3 in [Proposition 5.3.1](#)) which is intersected by a growing number of type 0_d regions. (By [Remark 5.3.1](#) any \mathcal{R}_3 region in \mathcal{S}_+ or \mathcal{S}_- will be a fan, which is an assumption needed for the application of (3) and (3') of [Proposition 5.3.1](#).)

Thus, we can only have an increase in $V(1, 1)$ if it comes from a pair of intersecting regions. Dealing with the growth of $V(1, 1)$ in \mathcal{S}_+ , we list the possibilities: (i) a type $0_{d,1}$ could intersect another type $0_{d,1}$; (ii) a type $0_{d,1}$ could intersect a type $0_{d,2}$; (iii) a type $0_{d,1}$ (or $0_{d,2}$) could intersect a type 1_d ; or (iv) a type $0_{d,1}$ (or $0_{d,2}$) could intersect a type -1_d . If we have possibility (i), for one of the $(1, 1)$ vertices the link(v) will be a good disc and we could have eliminated it by (7) of [Proposition 5.1.1](#). This violates our minimal complexity assumption, so possibility (i) does not occur. If we have possibilities (ii) or (iii) then by (2) and (2') of [Proposition 5.3.1](#) there will be a contribution of $+1$ to $b(X_+)$. So these occurrences are bounded.

Finally, we consider the growth of $V(1, 1)$ in \mathcal{S}_+ through an unbounded number of occurrences of possibility (iv). Suppose there is growth in pairs of regions $(\mathcal{R}_0, \mathcal{R}_{-1_d})$ such that: $\mathcal{R}_0 \subset \mathcal{S}_+$ is a type $0_{d,2}$ region; $\mathcal{R}_{-1_d} \subset \mathcal{S}_+$ is a type -1_d region; and

$\rho(\mathcal{R}_0) \cap \rho(\mathcal{R}_{-1}) \neq \emptyset$. We refer the reader back to Figures 46 and 43, and adapt them to our purpose at hand. In Figure 43 we see an illustration of a type -1_d region. Given such a region we can perform the inverse of the operation illustrated in the top sketch of Figure 46 to introduce an inessential b -arc that is positioned as the ‘left most’ b -arc in \mathcal{R}_{-1_d} . (Referring to the type -1_d region in Figure 43, as we traverse the black boundary in the direction of its orientation, with this introduction of an inessential b -arc, the first singularity of parity δ will no longer be associated with a clasp intersection.) Now, if we stabilize X_+ along this first singularity we will eat into \mathcal{R}_{-1_d} and the remaining portion of \mathcal{R}_{-1_d} will be a type $0_{d,1}$ region which we call \mathcal{R}'_{-1_d} . The new stabilized X_+ we call X'_+ , and we will have $b(X'_+) = b(X_+) + 1$. But, since \mathcal{R}_0 is type $0_{d,2}$ and $\rho(\mathcal{R}_0) \cap \rho(\mathcal{R}'_{-1_d}) \neq \emptyset$, by (3) of Proposition 5.3.1 we know that this intersection contributes $+1$ to $b(X'_+)$. Since X'_+ came from X_+ by a single stabilization we know $2 \leq b(X'_+) = b(X_+) + 1$. Thus, if we had x such $(\mathcal{R}_0, \mathcal{R}_{-1_d})$ pairs, for each pair we could have performed a similar stabilization on X_+ to produce a braid X'_+ ; and we would know that $2x \leq b(X'_+) = b(X_+) + x$. This implies $x \leq b(X_+)$. So we have bounded $V(1, 1)$.

Proof of (7) and (8) We study Equation (6–4) and ask which terms can grow without bound on both sides? By statements (1) and (6) of this lemma we know that the terms $V(1, 0)$ and $V(1, 1)$ on the LHS cannot grow without bound for fixed m and n . By statement (3) of this Lemma we know that if $V(0, 2)$ and/or $V(0, 3)$ on the LHS grow without bound, then the growth must occur inside the union \mathcal{N} of all normal neighborhoods of clasp arcs. By equation (7) we know that a growth in $V(1, 1)$ will force a growth in $\sum_{\beta=3}^{\infty} V(1, \beta)$ or $E(s)$.

Passing to the RHS of Equation (6–4), we know from statement (5) that $E(s)$ cannot grow without bound, for fixed m, n . By statement (2) we know that $V(2, 1)$ and $V(3, 0)$ are bounded. But then, the only terms which might not be bounded, on the RHS of Equation (6–4), are those in the double sum. However, of the terms in the double sum we know from statement (2) that $V(\alpha, v - \alpha)$ is bounded if $\alpha \geq 2$.

Suppose that β is bounded but that $V(1, \beta)$ increases without bound. This means that there is some fixed value of β for which there are arbitrarily many vertices of type $(1, \beta)$. An example was illustrated earlier, in Figure 44. In this illustration vertices U, V, Y, Z are vertices ‘at the bottom’ of a region that is composed of normal neighborhoods (see the right sketch of Figure 37), and the vertices W, W'' can be thought of as ‘coning’ these vertices and their associated singular leaves. (The vertex W'' should be thought of in a similar manner.) Then W and W'' contribute to the count of $V(0, \beta)$ or $V(1, \beta)$ and we are able to see the interplay between these terms and $V(0, 2)$ and $V(0, 3)$ in equation (7). The vertices U, V, Y, Z are in \mathcal{N} , but are adjacent to a vertex outside of \mathcal{N} . They have valence 2.

Such a coning idea can be iterated. Referring back to [Figure 62](#), if we imagine an additional vertex W''' lying below the black dotted arc we could conceivably cone W, W', W'' and all of the associated singular leaves to such a W''' . Since the number of vertices in the shaded region (vertices like U, V, Y, Z) can possibly grow, the number of vertices coning the bottom (or top) of a region composed of normal neighborhoods can also grow. And, the number of vertices coning the W -flavored vertices can then also grow, etc. So we see that there is no inherent reason why $\sum_{\beta=5}^{\infty} V(0, \beta)$ or $\sum_{\beta=3}^{\infty} V(1, \beta)$ should be bounded. However, we need only establish that $\sum_{\beta=5}^{\infty} V^e(0, \beta^e)$ or $\sum_{\beta=3}^{\infty} V^e(1, \beta^e)$ are bounded.

By construction the vertices that contribute to $V^e(0, \beta^e)$ and $V^e(1, \beta^e)$ count are outside normal neighborhoods. We suppose that a growth in them is balanced in [Equation \(6-4\)](#) by a growth in $V(0, 2)$ and $V(0, 3)$ that are associated with the normal neighborhoods of γ_+ arcs.

We need to go back to our original construction of \mathcal{CA} and extract an embedded annulus from \mathcal{CA} that contains the vertices that contribute to the $\sum_{\beta=5}^{\infty} V^e(0, \beta^e)$ or $\sum_{\beta=3}^{\infty} V^e(1, \beta^e)$. We do this by taking a tab neighborhood for each $\gamma_- \subset \mathcal{PA}$ (see the left sketch of [Figure 38](#)) and removing it from \mathcal{PA} and its image from \mathcal{CA} . (This is equivalent to stabilizing X_- along all of the singular leaves that are in the tab neighborhoods of the preimages γ_- of the clasp arcs.) Through an abuse of notation (and in keeping with [Section 2](#)) we call this embedded annulus \mathcal{A}_+ .

Recall the notation $X_-, X_0, X_+, \mathcal{A}_-, \mathcal{A}_+$ from the basic construction in [Section 2](#). Choose an annular neighborhood \mathcal{A}_0 of X_0 in $\mathcal{A}_+ \cup \mathcal{A}_-$ which does not intersect the clasp arcs. Then \mathcal{A}_0 is embedded and has X_0 as its core circle. Since \mathcal{A}_+ and \mathcal{A}_- are both embedded, we may extend them to embedded annuli $\mathcal{A}'_+ = \mathcal{A}_+ \cup \mathcal{A}_0$ and $\mathcal{A}'_- = \mathcal{A}_- \cup \mathcal{A}_0$ which have a common framing, also both are embedded and both have X_0 as a core circle. From the construction in [Section 2](#) we know that the algebraic linking number $Lk(X_-, X_0) = 0$. It then follows that the linking number $Lk(X_+, X_0)$ is also 0, and so \mathcal{A}'_+ can be extended to a minimal genus Seifert surface \mathbf{F}_+ (different from the one which we used in [Section 2](#)) having X_+ as its boundary. Observe that all of the vertices that contribute to $\sum_{\beta=5}^{\infty} V(0, \beta)$ or $\sum_{\beta=3}^{\infty} V(1, \beta)$ are in \mathcal{A}_+ , and so also are in \mathbf{F}_+ .

Now consider the count $\sum_{\beta=5}^{\infty} V^e(0, \beta^e)$ on $\mathcal{A}_+ \subset \mathbf{F}_+$, and suppose we have a type $(0, 2)$ vertex, $v \in \mathcal{A}_+$. The possibilities are: that both of the b -arcs that are adjacent to v are strongly essential; or one is strongly essential and the other is weakly essential; or both are weakly essential. (Our notion of weakly essential and strongly essential are now with respect to the surface \mathbf{F}_+ .) We observe that if we have the last case, when we eliminate either of the weakly essential b -arcs using an exchange move and the

surgery in Figure 25, the sum $\sum_{\beta=5}^{\infty} V^e(0, \beta^e)$ will remain constant. (Referring to the labeling in Figure 25, any strongly essential b -arcs that were adjacent to the vertex w_1 will be adjacent to the vertex w_2 after the surgery. Thus, the essential valence of w_2 will increase by exactly the essential valence of w_1 .)

Next, we observe that the count $\sum_{\beta=5}^{\infty} V^e(0, \beta^e)$ is invariant under our change in foliation illustrated in Figure 20. This is proved by examining the changes in the H_θ -sequence under the change in foliation (which reverses the order of the associated singularities).

So we allow in \mathcal{A}_+ the conversion of type $(0, 3)$ vertices to type $(0, 2)$ vertices, and the elimination of type $(0, 2)$ vertices if both of the adjacent b -arcs are weakly essential. Such alterations to \mathcal{A}_+ keep $\sum_{\beta=5}^{\infty} V^e(0, \beta^e)$ constant. If the number of type $(0, 2)$ vertices which we cannot eliminate in this manner is arbitrarily large then, since for each such vertex there is a strongly essential (with respect to \mathbf{F}_+) b -arc, the braid index of X_+ will be unbounded. (Basically, \mathcal{A}_+ is forcing the existence of a block and strand tree of arbitrarily high index as described in Section 5.5.) Since we cannot have an arbitrarily large number of such type $(0, 2)$ vertices adjacent to strongly essential b -arcs, and they balance out $\sum_{\beta=5}^{\infty} V^e(0, \beta^e)$, this sum must be bounded.

A similar argument applies to $\sum_{\beta=3}^{\infty} V^e(1, \beta^e)$. Also, we can interchange the role X_+ with X_- , using \mathcal{A}_- instead of \mathcal{A}_+ . (A subtle point is that we will have to change the orientation of \mathcal{A}_- to match that of X_- .) This completes the proof of Lemma 6.3.2. \square

Remark 6.3.2 It is interesting to note the similarities between (6–4) and [10, Equation (7)]. Given any Seifert surface, \mathbf{F} with Euler characteristic $\chi(\mathbf{F})$, assume that \mathbf{F} is tiled by $aa-$, $ab-$ and $bb-$ tiles. Using the notation introduced earlier for $V(\alpha, \beta)$, we have

$$(6-10) \quad V(1, 1) + 2V(0, 2) + V(0, 3) = 4\chi(\mathbf{F}) + V(2, 1) + 2V(3, 0) + \sum_{v=4}^{\infty} \sum_{\alpha=0}^v (v + \alpha - 4)V(\alpha, v - \alpha).$$

In our proof of statements (7) and (8) of Lemma 6.3.2 we used the fact that \mathcal{A}_+ can be extended to a Seifert surface \mathbf{F}_+ bounded by X_+ . In Equation (6–4) we noticed that when dealing with a bounded braid index, growth in the values $V(0, 2)$, $V(0, 3)$ is balanced by growth in the values $V(\alpha, v - \alpha)$ for $\alpha = 0, 1$ and $4 \leq v$. So in Equation (6–10), if there is any additional growth in the values $V((\alpha, v - \alpha)$ for $0 \leq \alpha$ and $4 \leq v$ it must be balanced out by (positive) growth in the value $-\chi(\mathbf{F}_+)$. (If there are any vertices in $\mathbf{F}_+ \setminus \mathcal{A}_+$ that contribute to the count of $V(0, 2)$ or $V(0, 3)$ then by our argument in [10] they would have been eliminated through changes in foliation and

exchange moves.) In other words, one can think of \mathcal{A}_+ as being the largest subannuli in \mathbf{F}_+ such that when Equation (6–10) is specialized to the surface $\mathbf{F}_+ \setminus \mathcal{A}_+$, growth in the sum $V(2, 1) + 2V(3, 0) + \sum_{v=4}^{\infty} \sum_{\alpha=0}^v (v + \alpha - 4)V(\alpha, v - \alpha)$ is balanced precisely by growth in the genus of \mathbf{F}_+ . As the number of tiles of \mathbf{F}_+ grow the only way $\mathcal{A}_+ \subset \mathbf{F}_+$ can intersect these additional tiles is by s -bands going through aa - or ab -tiles parallel to the X_+ boundary. The braiding of s -bands of \mathcal{A}_+ that comes from them running through aa - or ab -tiles of \mathbf{F}_+ can be seen as accounting for some of the braiding that occurs in the fixed blocks of a block-strand diagram. \diamond

Proposition 6.3.1 *Choose any positive integer m . Then for each fixed positive integer $n \leq m$ the set of templates in $\mathcal{T}(m, n)$ is finite. Moreover, while the finitely many templates in $\mathcal{T}(m, n)$ depend on the braid indices m and n of $\partial\mathcal{A}$, they do not depend in any other way on the choice of \mathcal{A} .*

Proof The idea behind the assertion that $\mathcal{T}(m)$ is finite is that the parts of the foliated clasp annulus \mathcal{CA} which can grow without bound when we fix the braid index of the boundary are all inside the blocks. In this regard observe that a block of braid index $k < m$ can contain an unbounded number of distinct k -braids, and of course in any one example the k -braid assignment to the block contributes to the foliation of \mathcal{CA} . The hard part of the proof is to show that in *all* cases where aspects of the foliation of \mathcal{CA} grow without bound, the growth in a template $\mathcal{T} = (D_+, D_-)$ can be understood as occurring inside the blocks of D_+ (which are also the blocks of D_-).

We begin by defining a subset of \mathcal{CA} which contains precisely the information that we need to construct a template in $\mathcal{T}(m)$. In this regard we remark that one of the beautiful features of block-strand diagrams is that most of the detailed information about the links that they support is concealed in the blocks, however we do not need to know details of what is in the blocks to construct the templates. Therefore we really need a rather limited amount of information from the foliation of \mathcal{CA} to construct the templates in $\mathcal{T}(m)$.

A b -arc in \mathcal{CA} is said to be *near* X_ϵ , $\epsilon = \pm$, if it has a vertex endpoint that meets an a_ϵ -arc. The subset \mathcal{BS} of \mathcal{CA} which is of interest to us now is the union of all s -arcs and a_\pm arcs in \mathcal{CA} , together with all b -arcs which are near X_+ or X_- , enlarged to a closed neighborhood in \mathcal{CA} which is chosen so that its boundary (which include both X_+ and X_-) is a union of simple closed curves which are transverse to the foliation of \mathcal{CA} . We call it the *boundary support* of \mathcal{PA} , because it is the subset of \mathcal{PA} which determines the embeddings of D_+ and D_- in 3-space, by Proposition 3.1.1 and the construction in Section 5.4. In the example in Figure 56 the boundary support includes everything except the b -arcs which join vertices 5 and 11, and also vertices 2 and 14. In generic examples \mathcal{BS} will be a very small subset of \mathcal{PA} .

By definition \mathcal{BS} includes the thin annuli, \mathcal{S}_+ (resp \mathcal{S}_-) which are associated to X_+ (resp. X_-), and all b -arcs in \mathcal{CA} that are near $X_+ \sqcup X_-$. A clasp arc pair (γ_+, γ_-) in \mathcal{PA} induces a clasp arc pair (γ_+, γ_-) in \mathcal{BS} . It follows from the definition of \mathcal{BS} that clasp arcs in \mathcal{CA} are in \mathcal{BS} if and only if they are doubly long. For, if it happens that γ_ϵ^i is long but $\gamma_{-\epsilon}^i$ is not, then \mathcal{BS} will be embedded near X_ϵ , even though \mathcal{CA} is not embedded near γ^i , and we are studying the part of \mathcal{PA} which determines the embedding of the template in 3-space. Notice further that an s -arc is in \mathcal{CA} if and only if it is in \mathcal{BS} . From now on we will drop the parity subscripts for a -arcs when we talk about the induced foliation on \mathcal{BS} .

As before, $\mathcal{N} \subset \mathcal{PA}$ be the union of all normal neighborhoods of the clasp arcs and let \mathcal{N}' be its complement in \mathcal{PA} . Initially, we let \mathcal{N}_{dl} be the union of all normal neighborhoods of all doubly long clasp arcs in \mathcal{PA} . These are the clasp arcs with the property that we can push X_+ across it all the way to X_- by a sequence of microflypes that amalgamate to a flype. (Other clasp arcs are not doubly long initially, but may become doubly long after G-exchange moves.) We enlarge \mathcal{N}_{dl} by adjoining to it any ab -tile that has only weakly essential b -arcs intersecting the boundary of our initial \mathcal{N}_{dl} . We allow continual enlargement of \mathcal{N}_{dl} in this fashion until any ab -tile that intersects \mathcal{N}_{dl} and has only weakly essential b -arcs is also in \mathcal{N}_{dl} . Next, let \mathcal{N}'_{dl} be the complement of \mathcal{N}_{dl} in \mathcal{N} . Let \mathcal{S} be the union of all bands of s -arcs. It will be convenient to divide the foliation of \mathcal{BS} into parts:

- $\mathcal{BS}_1 = \mathcal{S}$, the union of all bands of s -arcs.
- $\mathcal{BS}_2 = \mathcal{BS} \cap \mathcal{N}_{\text{dl}}$.
- $\mathcal{BS}_3 = \mathcal{BS} \cap \mathcal{N}'_{\text{dl}}$.
- $\mathcal{BS}_4 = \mathcal{BS} \cap \mathcal{N}$, the intersection of \mathcal{BS} with the union of all normal neighborhoods of all clasp arcs.
- $\mathcal{BS}_5 = \mathcal{BS} \setminus (\mathcal{BS}_1 \cup \mathcal{BS}_4)$, ie the part of \mathcal{BS} that is non-trivially tiled and outside all of the normal neighborhoods.

If A is a subset of \mathcal{PA} , let $|A|$ denote the number of connected components in A .

Claim $|\mathcal{BS}_i|$ is bounded for $i = 1, 2, 3, 4, 5$.

Proof of claim $|\mathcal{BS}_1| = |\mathcal{S}|$ is bounded, because every band of s -arcs has 2 s -edges, but by (5) of [Lemma 6.3.2](#) we know that the number of s -edges in the foliation of \mathcal{CA} is bounded. Suppose next that $|\mathcal{BS}_2| = |\mathcal{BS} \cap \mathcal{N}_{\text{dl}}|$ is unbounded. Since the clasp arcs in \mathcal{N}_{dl} are doubly long, there must be some connected component of \mathcal{PA} split along the bounded set \mathcal{S} which has the property that as one travels along X_\pm in this component one passes from \mathcal{N}_{dl} to \mathcal{N}'_{dl} an unbounded number of times. However, studying the

regions in Figure 43 we see that this would violate (4), (5) and (6) of Lemma 6.3.2, so this cannot happen. The identical argument shows that $|\mathcal{BS}_3| = |\mathcal{BS} \cap \mathcal{N}'_{\text{dl}}|$ is also bounded. (Remark: this does not say that there is a bound on the number of clasp arcs, in fact no such bound exists.) Since $\mathcal{BS}_4 = \mathcal{BS}_2 + \mathcal{BS}_3$ it follows that $|\mathcal{BS}_4|$ is also bounded. Since \mathcal{BS}_5 is the part of \mathcal{BS} that is non-trivially tiled and outside the union of all normal neighborhoods, it follows that as we travel along a component of X_{\pm} we will intersect components of $\mathcal{BS}_1, \mathcal{BS}_5$ and \mathcal{BS}_4 . We never pass from a component of \mathcal{BS}_5 to another component of \mathcal{BS}_5 without passing through a component of \mathcal{BS}_1 or \mathcal{BS}_4 . Since $|\mathcal{BS}_1|$ and $|\mathcal{BS}_4|$ are both bounded, it follows that $|\mathcal{BS}_5|$ must be bounded too. This proves the claim. \square

Our next task is examine the contributions of the components of \mathcal{BS}_i to D_+ and D_- . For this we need to investigate in detail (in a more general setting) the construction in Section 5.4. Recall that to construct a template (D_+, D_-) we needed to understand four aspects of its structure: the moving blocks, the moving strands, the fixed blocks and the fixed strands. We analyze each separately.

Moving blocks An amalgamating block \mathbf{B} will be moved to an amalgamating block \mathbf{B} if and only if \mathbf{B} and \mathbf{B} are related amalgamating blocks as defined in Section 5.6. This assumes that they are associated to clasp arcs which are doubly long in \mathcal{PA} . Thus the moving blocks will be associated to \mathcal{N}_{dl} . Note that there may be some choices involved when we select the amalgamating blocks. We make those choices in such a way that the set of all moving blocks has minimal cardinality.

The strands of X_+ which are incorporated into an amalgamating block lie in the black boundary of \mathcal{N}_{dl} , ie in the subset \mathcal{BS}_2 of \mathcal{BS} . We have already proved that \mathcal{BS}_2 is bounded. We know that if a subarc of $X_+ \cap \mathcal{N}_{\text{dl}}$ is related to a corresponding subarc of $X_- \cap \mathcal{N}_{\text{dl}}$, then their angular lengths coincide. So let $\{N_1, \dots, N_r\} = \mathcal{N}_{\text{dl}}$ be a listing of all of the components. For each component N_i we define its *angular support* $\angle N_i$ to be the interval $[\theta_i^0, \theta_i^1] \subset [0, 2\pi)$ for which $\theta \in [\theta_i^0, \theta_i^1]$ iff $H_{\theta} \cap [(X_+ \cup X_-) \cap N_i] \neq \emptyset$. Notice that if \mathbf{B} and \mathbf{B} are related amalgamating blocks then the angular support of every N_i component that intersects \mathbf{B} (or \mathbf{B}) must overlap, ie if each component of the subcollection $\{N_{j_1}, \dots, N_{j_r}\} \subset \{N_1, \dots, N_r\}$ intersects \mathbf{B} then $\angle N_{j_1} \cap \dots \cap \angle N_{j_r} \neq \emptyset$. Since \mathcal{N}_{dl} is a finite set there are a finite number of angular support intervals, and for those intervals there are only a finite number of possible intersection subsets. Thus, there are only a finite number of possible moving blocks in any template of $\mathcal{T}(m, n)$.

Moving Strands Every subarc of X_+ which is away from the bands of s -arcs is potentially a moving strand, however some of these potential moving strands have been amalgamated into moving blocks. We separate the surviving moving strands into two types:

- (a) Moving strands that are the subarcs of X_+ which are in $\partial\mathcal{N}_{\text{dl}}$ but were not amalgamated into moving blocks. The finiteness of this set of strands follows from the argument we used to prove that the number of moving blocks is finite.
- (b) Moving strands that are the subarcs of X_+ which intersect \mathcal{BS}_5 . Since $|\mathcal{BS}_5|$ is bounded, we can restrict ourselves to the strands which intersect a single component of \mathcal{BS}_5 .

By Proposition 3.1.1, the embedding of these strands is determined by the ordering and signs of the vertices and singularities which belong to a tile in \mathcal{BS}_5 which intersects X_+ . If we can show that the number of such vertices is bounded, it will follow that the number of singularities is also bounded. Since the number of distinct ways to assign orders and signs to a finite set of vertices and singularities is finite, it will then follow that the number of possible arrangements of the moving strands in set (b) is bounded.

The vertices in question contribute to the count of $V(1, \beta)$, and unfortunately could be unbounded. For example, since we can have arbitrarily many clasp arcs, the number of vertices contributing to the count of $V(1, 2)$ can be arbitrarily high. But, these vertices are adjacent to weakly essential b -arcs which do not add information to the positioning of our type (a) or (b) strands. Thus, the only vertices that we need to be concerned with are the ones that contribute to the count of $V^e(1, \beta^e)$. By statement (7) of Lemma 6.3.2, we know that they are bounded. Thus, there are only a finite number of possibilities for the positioning of moving strands in any template in $\mathcal{T}(m, n)$.

Fixed blocks The argument here is more subtle than the one for the moving blocks, because the fixed blocks are associated to bands of s -arcs, and so there is no tiling to work with. Nevertheless, we can relate the phenomenon of block amalgamation to the tiling, in the following way. Recall (see the construction in Section 5.4) how the fixed blocks are formed. Let \mathcal{CA} be the clasp annulus and let $\mathcal{CA}_{\text{tiled}}$ be the part that is not foliated by s -arcs. Let $\{\theta_1, \dots, \theta_r\}$ be a listing of all of the singularities in $\mathcal{BS} \cap \mathcal{CA}_{\text{tiled}}$. For each θ in one of the intervals $[\theta_i, \theta_{i+1}]$ we know there are no singularities. For each such θ we also know that H_θ contains a -arcs and b -arcs and s -arcs. So $(H_\theta \setminus [H_\theta \cap \mathcal{CA}_{\text{tiled}}])$ is a collection of discs, each containing only s -arcs (with some discs possibly containing no s -arcs). Each component of the union over all $\theta \in [\theta_i, \theta_{i+1}]$ of the set $\{(H_\theta \setminus H_\theta \cap \mathcal{CA}_{\text{tiled}})\}$ has a $D^2 \times [\theta_i, \theta_{i+1}]$ structure. Let C be one such component. If C contains bands of s -arcs, amalgamate them into a single block $\mathbf{B}(C)$. If C contained no bands or just a single band then $\mathbf{B}(C)$ is either vacuous or a single fixed strand.

Now it may happen that there is another connected component C' with its associated block $\mathbf{B}(C')$, such that we have the *amalgamation condition*

$$(6-11) \quad \{s\text{-arcs}\} \cap C' \cap H_{\theta_{i+1}} \subset \{s\text{-arcs}\} \cap C \cap H_{\theta_{i+1}}.$$

If this happens, amalgamate $\mathbf{B}(C')$ and $\mathbf{B}(C)$ and delete the singularity at θ_{i+1} , so that we have a single block $\mathbf{B}(C \cup C')$. (If $\mathbf{B}(C')$ was vacuous or a single strand then $\mathbf{B}(C \cup C')$ is essentially still $\mathbf{B}(C)$.) Continue this amalgamation process as long as possible. Among all amalgamated blocks discovered in this way, choose one such that the set of fixed blocks has minimum cardinality. In this way we will have eliminated some number of singularities, ie the ones which separated the new-amalgamated blocks. Let $\theta^1, \dots, \theta^p$ be the angles which remain. We need to show that this listing is bounded, when $b(X_+)$ and $b(X_-)$ are fixed.

To do this we consider the effect a singularity has in the H_θ -sequence when weakly essential b -arcs are used. We refer to Figure 18, where the possibilities for the b -edges of bb - and ab -tiles are illustrated. Each could be either strongly or weakly essential.




Figure 62: The possibilities for b -edges of bb - and ab -tiles. A black undotted b -edge means strongly essential. A black dotted b -edge means weakly essential.

First, note that occurrences of aa -singularities will effect the block amalgamation condition stated previously, ie (6–11) above. So the occurrence of aa -singularities will register in our listing of remaining angles $\theta^1, \dots, \theta^p$. But, by statements (1), (2) and (3) of Lemma 6.3.2 we will have a bounded number of such singularities.

Second, notice that among the possibilities for bb -tiles, (1) and (2) will not effect our amalgamation condition, because weakly essential b -arcs split off regions that contain no s -arcs. This is also true for (1) and (2) in our possibilities for ab tiles.

Thus, the only possibilities that effect the amalgamation condition are (4), (5) and (6) for bb -tiles, and (3) for ab -tiles.

Now by statements (1), (2), (4), (6), (7) and (8) of Lemma 6.3.2 there are a bounded number of strongly essential b -arcs in a fixed fiber of \mathbf{H} . So there are a bounded number of singularities that can effect our amalgamation condition. Thus, the growth in our remaining angles $\theta^1, \dots, \theta^p$ is bounded. (It is interesting to notice that since

the tiles in the normal neighborhoods have all of their sides labeled w_i they do not effect the amalgamation condition.)

We have in fact established more than just that there are a finite number of fixed blocks. We have established that there are only a finite number of possible positions for fixed blocks. This is because there are only a finite number of positions for vertices that contribute to the count of $V^e(\alpha, \beta^e)$ and, thus, a finite number of singularities that correspond to possibilities (3)–(5) for bb tiles and possibilities (2) and (3) for ab tiles. Therefore any combinatorial information having to do with cyclic ordering of vertices on \mathbf{A} and cyclic ordering of singularities in \mathbf{H} is also finite.

Fixed strands If the number of moving blocks, moving strands and fixed blocks is finite then the number of fixed strands must be finite. Using the observation we employed to establish the finiteness of positions of fixed blocks, we can establish finiteness of the positions of fixed strands.

Thus everything is bounded, and so the number of block-strand diagram pairs is bounded. The proof of [Proposition 6.3.1](#) is complete. But then, so is the proof of [Theorem 2](#). \square

7 Open problems

(1) As was pointed out in [Section 1.1](#), Markov's Theorem is just one example of a class of theorems about 3-manifolds in which some form of stabilization and destabilization play an important role. A very different and very important example is the Kirby Calculus [21], relating two surgery representations of a 3-manifold. The stabilization move is the addition of an unknot with standard framing to the link that is to be surgered. In this setting, what should be the analogue of the MTWS? Presumably, the first entry in the complexity function should be the number of components in the surgery link. But unfortunately, in this regard, we do not know of any tools that could be put to work, to play the role that was played by the geometry of braid foliations and the group structure of the braid group in the proof of the MTWS.

(2) There are many analogies between the study of knots via their closed braid representatives and the study of 3-manifolds via their Heegaard diagrams (or equivalently via their 'Heegaard gluing maps' in the mapping class group of a closed orientable surface of genus g). In the latter setting equivalence classes of Heegaard splittings are in 1–1 correspondence with double cosets in the mapping class group \mathcal{M}_g modulo the mapping class group \mathcal{H}_g of a handlebody. We pose as an open problem to find moves which (like the moves in $\mathcal{T}(m)$) change the equivalence class of a Heegaard splitting of a 3-manifold without increasing its Heegaard genus. A strategy for finding

such moves is given by the second author in [24], however (lacking an invariant) there is no proof that this strategy actually produces inequivalent splittings. In his PhD thesis [33] Joel Zablow made a relevant contribution in his study of waves in Heegaard diagrams. Is there a tool which plays the role of braid foliations in the situation of Heegaard splittings of 3-manifolds? This seems to be a very interesting area for future investigations.

(3) We pass to open questions about [Theorem 2](#) of this paper. In principle the templates in $\mathcal{T}(m)$ can be enumerated, but the actual enumeration is non-routine. We pose this as an open problem for $m = 4, 5$ and any other cases which prove to be computable. A very interesting special case are to classify the templates in $\mathcal{T}(m, m)$. In particular, these relate any two closed braid representatives when both have minimum braid index. There should be applications. For example, knowing that the only templates that we need for braid index 3 are the flype and destabilization templates, it is a simple matter to classify links which are closed 3-braids (replacing the complicated argument we used in [11]), and it is to be expected that if $\mathcal{T}(4)$ is computed, then one would learn more about the classification of links of braid index 4. One could expect many applications, if such an advance in knowledge could be achieved.

(4) Although actual enumeration of $\mathcal{T}(m)$ may not be routine, if we restrict our attention to a single type of isotopy such enumeration or characterization may be reasonably doable. Specifically, referring back to [Section 5.7](#), the characterization of all knot complements that admit a cyclic move predicated on the existence of an essentially embedded standard annulus in its complement would be of interest. Moreover, such knot complements could be divided into two classes: the first class would have the components of G_{++} and G_{--} being homeomorphic to $[0, 1]$; and, the second class would have the components of G_{++} and G_{--} being homeomorphic to S^1 . The first class would use a positive stabilization and destabilization to begin and end the cyclic move. The second class would use a negative stabilization and destabilization for the cyclic move. The question is how to determine when a knot is not in both classes, since the first class corresponds to a transversal isotopy and the second class does not. (See our article [6].)

(5) In the manuscript [7] the authors proved that an arbitrary closed braid representative of a composite knot or link may be modified by the use of exchange moves to a prime summand of the same braid index. We do not know whether the work in that paper can be incorporated into the proof of the MTWS. The difficulty is that we do not know whether the splitting 2-sphere which realizes the connect-sum operation can be modified to one which intersects the clasp annulus \mathcal{CA} in a ‘nice’ way.

(6) Some knots or links, for example the unlink [9] and most iterated torus links [25], have unique closed braid representatives of minimum braid index. On the other hand,

there are links of braid index 3 which have more than one conjugacy class of 3–braid representatives [11], and this pathology clearly persists as braid index is increased. We pose the open problem: find general conditions which suffice for a knot or link type to have a *unique* conjugacy class of closed braid representative of minimum braid index.

(7) Referring back to Remark 6.3.2, it would be of interest to understand exactly how the structure of a minimal genus Seifert surface can restrict the ability of a given link type having minimal braid index n to be carried by any template of $\mathcal{T}(m, n)$. In particular, if the Seifert surface has a foliation composed of aa tiles at minimal braid index is $\mathcal{T}(m, n)$ necessarily empty for all values of m .

(8) There are special knots and links, for example the unlink [9] and most iterated torus links [25], for which the MTWS is very simple: the moves are simply braid isotopy and exchange moves. We say that such links are *exchange-reducible*. Are there other examples of exchange-reducible links? Does the conclusion hold under weaker hypotheses?

We remark that by the main result in Birman–Wrinkle [14], if a knot type \mathcal{X} is exchange-reducible, then every transversal knot type $\mathcal{T}\mathcal{X}$ associated to \mathcal{X} is transversally simple, ie determined up to transversal isotopy by \mathcal{X} and the Thurston–Bennequin invariant. Since G–exchange moves and positive flypes are realized by transversal isotopy, it would be equally interesting if the condition ‘exchange-reducible’ was weakened to ‘exchange and positive flype-reducible’.

(9) As noted in problem (4) above, the unlink is exchange-reducible. This fact proves that *there exists* a monotonic and very rapid (perhaps even a quadratic) algorithm for recognizing the unlink, through the use of exchange moves. Unfortunately, however, the complexity function that would translate this existence theorem into a working algorithm needs new techniques, as the complexity function is concealed in the invisible family of discs which the unlink bounds. (One of these days the first author will write a short note to show that it is also concealed in the auxiliary ‘extended braid word’ of Birman and Hirsch [5].) We note that the unknot recognition algorithms in [5] and the finite unknot recognition algorithm in the very new paper by Ivan Dynnikov [17], which is based upon related foliation techniques, are exponential. A vague (but we feel realistic) problem is to find an ‘energy functional’ (AKA complexity function) which uses the monotonic reduction process that is guaranteed to exist because of exchange-reducibility. A wild guess is that it is encoded in notions based upon Ricci curvature.

(10) The MTWS begins with a choice of a closed braid representative $X_- \in \mathcal{X}$ which has braid index $b(\mathcal{X})$, however at this writing we do not know how to compute $b(\mathcal{X})$. The most useful tool that we know is the ‘Morton–Franks–Williams inequality’,

however the MFW inequality is doomed to fail in certain situations (see Kawamuro [20]). An important open problem is to develop new techniques for computing the braid index of a knot or link. As can be seen from [20], templates can be a very useful tool.

(11) Conjecture:

- Every block-strand diagram in a template in $\mathcal{T}(m)$ has at least one block.

We remark that our attempts to find a counterexample have been unsuccessful, but we lack a proof that it cannot happen. Note that the number of counterexamples is necessarily finite.

(12) In a standard annulus, the graph $G_{-, \delta}$ (resp. $G_{+, \delta}$) is topologically equivalent to a circle which cobounds with a component of X_+ (resp. X_-) an embedded annulus. That annulus is foliated without singularities. From this it follows that, if we regard $G_{-, \delta}$ and $G_{+, \delta}$ as defining knots in \mathbb{R}^3 , then they will have the same knot type as the component in question of \mathcal{X} . But in fact more can be said. The graph $G_{-, \delta}$ (resp. $G_{+, \delta}$) is a union of arcs, each a branch in a singular leaf, which join up a string of negative (resp. positive) vertices in a cycle. Each arc lies in a fiber of \mathbf{H} and has its endpoints on \mathbf{A} , and so this representation of the component of \mathcal{X} has an ‘arc presentation’, in the sense defined by Cromwell [16] and Dynnikov [17]. Indeed, Ivan Dynnikov has been engaged in a project which begins with the introduction of braid foliations, and goes on to study the foliations of the associated Seifert surface bounded by $G_{-, \delta}$ and $G_{+, \delta}$, adapting the braid foliation machinery in [7] and [9] to arc presentations. One expects that there will be similar adaptations of the work in this paper to arc presentations, although the adaptation is almost certainly non-trivial.

One reason why arc presentations are of interest is because they give a filtration of all knots and links, using the number of arcs as a measure of complexity, and with that filtration there are always finitely many arc presentations which represent a given knot type and have complexity at most the complexity of a given example. This fact is important if one wishes to use the braid foliation machinery to construct algorithmic solutions to the knot and link problem.

References

- [1] **J W Alexander**, *A lemma on systems of knotted curves*, Proc. Nat. Acad. Sci. USA 9 (1923) 93–95
- [2] **D Bennequin**, *Entrelacements et équations de Pfaff*, from: “Third Schnepfenried geometry conference, Vol. 1 (Schnepfenried, 1982)”, Astérisque 107, Soc. Math. France, Paris (1983) 87–161 [MR753131](#)

- [3] **J S Birman**, *Braids, links, and mapping class groups*, Princeton University Press, Princeton, N.J. (1974) [MR0375281](#)
- [4] **J S Birman, E Finkelstein**, *Studying surfaces via closed braids*, *J. Knot Theory Ramifications* 7 (1998) 267–334 [MR1625362](#)
- [5] **J S Birman, M D Hirsch**, *A new algorithm for recognizing the unknot*, *Geom. Topol.* 2 (1998) 175–220 [MR1658024](#)
- [6] **J S Birman, W W Menasco**, *Stabilization in the braid groups II: Applications to transverse knots* [arXiv:math.GT/0310280](#)
- [7] **J S Birman, W W Menasco**, *Studying links via closed braids. IV. Composite links and split links*, *Invent. Math.* 102 (1990) 115–139 [MR1069243](#)
- [8] **J S Birman, W W Menasco**, *Studying links via closed braids. I. A finiteness theorem*, *Pacific J. Math.* 154 (1992) 17–36 [MR1154731](#)
- [9] **J S Birman, W W Menasco**, *Studying links via closed braids. V. The unlink*, *Trans. Amer. Math. Soc.* 329 (1992) 585–606 [MR1030509](#)
- [10] **J S Birman, W W Menasco**, *Studying links via closed braids. VI. A nonfiniteness theorem*, *Pacific J. Math.* 156 (1992) 265–285 [MR1186805](#)
- [11] **J S Birman, W W Menasco**, *Studying links via closed braids. III. Classifying links which are closed 3-braids*, *Pacific J. Math.* 161 (1993) 25–113 [MR1237139](#)
- [12] **J S Birman, W W Menasco**, *On Markov's theorem*, *J. Knot Theory Ramifications* 11 (2002) 295–310 [MR1905686](#)
- [13] **J S Birman, M Rampichini, P Boldi, S Vigna**, *Towards an implementation of the B-H algorithm for recognizing the unknot*, *J. Knot Theory Ramifications* 11 (2002) 601–645 [MR1915499](#)
- [14] **J S Birman, N C Wrinkle**, *On transversally simple knots*, *J. Differential Geom.* 55 (2000) 325–354 [MR1847313](#)
- [15] **G Burde, H Zieschang**, *Knots*, de Gruyter Studies in Mathematics 5, Walter de Gruyter & Co., Berlin (1985) [MR808776](#)
- [16] **P R Cromwell**, *Embedding knots and links in an open book. I. Basic properties*, *Topology Appl.* 64 (1995) 37–58 [MR1339757](#)
- [17] **I Dynnikov**, *Arc presentations of links: monotonic simplification* [arXiv:math.GT/0208153](#)
- [18] **T Fiedler**, *A small state sum for knots*, *Topology* 32 (1993) 281–294 [MR1217069](#)
- [19] **J Hempel**, *3-Manifolds*, Princeton University Press, Princeton, N. J. (1976) [MR0415619](#)
- [20] **K Kawamuro**, *Failure of the Morton–Franks–Williams inequality* [arXiv:math.GT/0509169](#)

- [21] **R Kirby**, *A calculus for framed links in S^3* , Invent. Math. 45 (1978) 35–56 [MR0467753](#)
- [22] **S Lambropoulou, C P Rourke**, *Markov's theorem in 3-manifolds*, Topology Appl. 78 (1997) 95–122 [MR1465027](#)
- [23] **A A Markov**, *Über die freie Äquivalenz geschlossener Zöpfe*, Recueil Mathematique Moscou I (1935) 73–78
- [24] **W W Menasco**, *Closed braids and Heegaard splittings*, from: “Knots, braids, and mapping class groups—papers dedicated to Joan S. Birman (New York, 1998)”, AMS/IP Stud. Adv. Math. 24, Amer. Math. Soc., Providence, RI (2001) 131–141 [MR1873114](#)
- [25] **W W Menasco**, *On iterated torus knots and transversal knots*, Geom. Topol. 5 (2001) 651–682 [MR1857523](#)
- [26] **H R Morton**, *Infinitely many fibred knots having the same Alexander polynomial*, Topology 17 (1978) 101–104 [MR486796](#)
- [27] **H R Morton**, *An irreducible 4-string braid with unknotted closure*, Math. Proc. Cambridge Philos. Soc. 93 (1983) 259–261 [MR691995](#)
- [28] **H R Morton**, *Threading knot diagrams*, Math. Proc. Cambridge Philos. Soc. 99 (1986) 247–260 [MR817666](#)
- [29] **M Scharlemann, A Thompson**, *Thin position and Heegaard splittings of the 3-sphere*, J. Differential Geom. 39 (1994) 343–357 [MR1267894](#)
- [30] **J Singer**, *Three-dimensional manifolds and their Heegaard diagrams*, Trans. Amer. Math. Soc. 35 (1933) 88–111 [MR1501673](#)
- [31] **P Traczyk**, *A new proof of Markov's braid theorem*, from: “Knot theory (Warsaw, 1995)”, Banach Center Publ. 42, Polish Acad. Sci., Warsaw (1998) 409–419 [MR1634469](#)
- [32] **F Waldhausen**, *Heegaard-Zerlegungen der 3-Sphäre*, Topology 7 (1968) 195–203 [MR0227992](#)
- [33] **J Zablow**, *Loops, waves and an algebra for Heegaard splittings*, PhD thesis, City University of New York (1999)

*Department of Mathematics, Barnard College, Columbia University
2990 Broadway, New York NY 10027, USA*

Department of Mathematics, University at Buffalo, Buffalo NY 14260, USA

jb@math.columbia.edu, menasco@math.buffalo.edu

Proposed: Robion Kirby
Seconded: Benson Farb, David Gabai

Received: 23 June 2005
Accepted: 25 January 2006