


Geometry & Topology
 Volume 5 (2001) 651–682
 Published: 15 August 2001



On iterated torus knots and transversal knots

WILLIAM W MENASCO

University at Buffalo, Buffalo, New York 14214, USA

Email: menasco@tait.math.buffalo.edu

URL: <http://www.math.buffalo.edu/~menasco>

Abstract

A knot type is *exchange reducible* if an arbitrary closed n -braid representative K of \mathcal{K} can be changed to a closed braid of minimum braid index $n_{\min}(\mathcal{K})$ by a finite sequence of braid isotopies, exchange moves and \pm -destabilizations. (See Figure 1). In the manuscript [6] a transversal knot in the standard contact structure for S^3 is defined to be *transversally simple* if it is characterized up to transversal isotopy by its topological knot type and its self-linking number. Theorem 2 of [6] establishes that exchange reducibility implies transversally simplicity. Theorem 1.1, the main result in this note, establishes that iterated torus knots are exchange reducible. It then follows as a Corollary that iterated torus knots are transversally simple.

AMS Classification numbers Primary: 57M27, 57N16, 57R17

Secondary: 37F20

Keywords: Contact structures, braids, torus knots, cabling, exchange reducibility

Proposed: Yasha Eliashberg
 Seconded: Robion Kirby, Joan Birman

Received: 27 March 2001
 Revised: 17 July 2001

1 Introduction

Let $\mathcal{C} \subset S^3$ be an oriented knot, let $\mathcal{V}_{\mathcal{C}}$ be a solid torus neighborhood of \mathcal{C} and let $\partial\mathcal{V}_{\mathcal{C}} = \mathcal{T}_{\mathcal{C}} \subset S^3$ be a peripheral torus for \mathcal{C} . The oriented simple closed curve on $\mathcal{T}_{\mathcal{C}}$ that represents the homotopy class of $pm + ql$, where m is the meridian homotopy class, l is the preferred longitude homotopy class and $p, q \in \mathbb{Z}$, is called the (p, q) cable of \mathcal{C} . We will use the notion $\mathbf{C}(\mathcal{C}, (p, q))$ to indicate the resulting oriented knot of this *cabling operation*. If \mathcal{C} is the oriented unknot and (p, q) is a co-prime pair then the cabling operation produces a (p, q) -torus knot. (Our discussion and results can be adapted to the situation where (p, q) is not co-prime, but since we will be concerned with the iteration of the cabling operation, to avoid ambiguity we require that the cabling produce a knot.)

We can, of course, iterate the cabling operation. Starting with an initial knot \mathcal{C}_0 and a sequence of co-prime 2-tuples of integers $\{(p_1, q_1), (p_2, q_2), \dots, (p_h, q_h)\}$, with $p_1 < q_1$, we can construct the oriented knot

$$\mathbf{C}(\mathbf{C}(\dots \mathbf{C}(\mathbf{C}(\mathcal{C}_0, (p_1, q_1)), (p_2, q_2)) \dots, (p_{h-1}, q_{h-1})), (p_h, q_h)).$$

If \mathcal{C}_0 is the oriented unknot then any iteration of the cabling operation produces an *iterated torus knot*. Letting $(P, Q) = \{(p_1, q_1), (p_2, q_2), \dots, (p_h, q_h)\}$, the final iteration produces an oriented knot, $K_{(P, Q)}$, which is on the peripheral torus of the next to last knot in the iteration; mainly,

$$\mathcal{T}_{\mathbf{C}(\mathbf{C}(\dots \mathbf{C}(\mathcal{C}_0, (p_1, q_1)), (p_2, q_2)) \dots, (p_{h-1}, q_{h-1}))}.$$

In Section 2 of [6] three moves are discussed which take closed braids to closed braids, preserving knot type: braid isotopy, exchange moves and destabilization. Braid isotopy means an isotopy in the complement of the braid axis which preserves the braid structure. The exchange move is a special type of Reidemeister II move illustrated in Figure 1(a). Destabilization means reducing the braid index by eliminating a (positive or negative) trivial loop, as shown in Figure 1(b). Notice that braid isotopy and exchange moves preserve both algebraic crossing number and braid index, whereas destabilization changes both. For a more extensive treatment of these isotopies, see [1, 3, 4].)

As defined in [6], a knot type \mathcal{K} is *exchange reducible* if an arbitrary closed n -braid representative K of \mathcal{K} can be changed to a closed braid representative of minimum braid index, $n_{min}(\mathcal{K})$, by a finite sequence of braid isotopies, exchange moves and \pm -destabilizations. The main result of [2] established the exchange reducibility of the unknot. The main theorem in this paper is an analogous result for iterated torus knots.




Figure 1

Theorem 1.1 *Iterated torus knots are exchange reducible.*

It should be noted that in [5] it was shown that there are knots that can be represented by two non-conjugate closed 3-braids where the differing conjugacy classes are not related to each other by an exchange move. (In fact, they are related to each other by a braid preserving flype.)

The proof of Theorem 1.1 involves adapting the braid-foliation machinery developed in [4] to the situation where there is a torus in S^3 which is being foliated and a knot is embedded on this torus. It employs the result from [12] that an oriented iterated torus knot $K_{(P,Q)}$ has an unique braid representative of minimal braid index $\prod_1^h p_i$.

Theorem 1.1 has an immediate application to transversal knots. Let ξ be the standard contact structure in oriented S^3 . The structure ξ can be thought of as a plane field that is totally non-integrable. A knot K is transversal if and only if K intersects each plane in the plane field ξ transversally. A transversal isotopy of K is an isotopy of K in S^3 through transversal knots. (See [7].) If K and K' are two transversal knots that are transversally isotopic, then they are representatives of the same *transversal knot type*, \mathcal{TK} .

A classical invariant of transversal knot types is a *self-linking number*, the *Bennequin number*, $\beta(\mathcal{TK})$. The self-linking is defined by pushing the transversal knot off itself in a direction which is in the contact plane. A well-defined direction exists because S^3 is parallelizable. See [6] for a precise description. A transversal knot type \mathcal{TK} is *transversally simple* if it is determined by its topological knot type \mathcal{K} and its self-linking number. In [7] it was first shown that the unknot is transversally simple. In [8] it was established that positive transversal torus knots are transversally simple. In [6] a more general framework for understanding transversally simple knots was established.

Theorem 1.2 (See [6]) *If \mathcal{TK} is a transversal knot type with associated topological knot type \mathcal{K} , where \mathcal{K} is exchange reducible, then \mathcal{TK} is transversally simple.*

Combining Theorems 1.1 and 1.2, we have the following immediate corollary.

Corollary 1.3 *Let $\mathcal{TK}_{(P,Q)}$ be a transversal knot type with associated topological knot type that of the iterated torus knot $K_{(P,Q)}$. Then $\mathcal{TK}_{(P,Q)}$ is transversally simple.*

The outline for this note is as follows. In Section 2 we review and adapt the braid-foliation machinery for the torus that was initially introduced in [4]. We will be concerned with the situation where we are given a torus which contains a knot $K_{(P,Q)}$ and bounds a solid torus. However, we do not have a natural way of identifying the core curve of the solid torus. Hence, we will use \mathcal{T}_C as notation for the given torus containing $K_{(P,Q)}$. The foliation machinery on \mathcal{T}_C will involve understanding the manipulation of three different types of foliations—circular, mixed and tiled foliations. (These foliations will be defined in Section 2.) In Section 3 we will prove Theorem 1.1 in the special case where \mathcal{T}_C has a circular foliation. The overriding strategy of the remaining sections is to reduce the mixed and tiled foliations to circular foliations. In Section 4 we show how destabilizations and exchange moves allow one to replace a mixed foliation with a circular foliation. Similarly, in Section 5 we show how destabilizations and exchange moves allow one to replace a tiled foliation with a circular foliation.

Acknowledgments This work was partially supported by NSF grant DMS-9626884. The author wishes to thank Joan Birman and Nancy Wrinkle for encouraging him to think about a proof of Theorem 1.1 during his brief sabbatical stay at Columbia University. That stay was partially supported by NSF grant DMS-9705019 .

2 The braid foliation machinery for the torus

In this section we adapt the combinatorics of [4] to the pair (K, \mathcal{T}_C) , where $K = K_{(P,Q)}$ is an iterated torus knot which lies on the torus \mathcal{T}_C . Since the exposition in [1] supplies us with a centralized source for most of the previously developed machinery, we will use it almost exclusively as our primary reference.

Although the arguments in this note rely heavily on results in the existing literature, a reader need only consult [1] and [4] in almost all cases.

Let $K \subset S^3$ be an oriented closed n -braid with axis \mathbf{A} . Then we can choose a 2-disc fibration of the open solid torus $S^3 - \mathbf{A}$. We will refer to $\mathbf{H} = \{H_\theta \mid \theta \in [0, 2\pi]\}$ as this 2-disc fibration of $S^3 - \mathbf{A}$. Each H_θ is a disc with boundary \mathbf{A} . We consider the intersection of the H_θ 's with \mathcal{T}_C —the induced singular foliation on \mathcal{T}_C by \mathbf{H} . We have a sequence of lemmas that begin to standardize this foliation. These lemmas imitate the similar set of lemmas in Section 1 of [4] which dealt with an essential torus in the complement of a closed braid. Since our present case is slightly different (the closed braid is actually a homotopically non-trivial curve on the torus), we will only supply the additional details needed to adapt the proofs of [4] to this case.

Lemma 2.1 *We may assume that:*

- (i) *The intersections of \mathbf{A} with \mathcal{T}_C are finite in number and transverse. Also, if $p \in \mathbf{A} \cap \mathcal{T}_C$ then p has a neighborhood on \mathcal{T}_C which is radially foliated by its arcs of intersection with fibers of \mathbf{H} .*
- (ii) *All but finitely many fibers $H_\theta \in \mathbf{H}$ meet \mathcal{T}_C transversally, and those which do not (the singular fibers) are each tangent to \mathcal{T}_C at exactly one point in the interior of both \mathcal{T}_C and H_θ . Moreover, the tangencies (which are contained in singular leaves) are either local maxima, or minima, or saddle points.*
- (iii) *A leaf that does not contain a singular point (a non-singular leaf) is either an arc having endpoints on \mathbf{A} or a simple closed curve.*

Proof We use exactly the same general position argument as in [4]. □

We refer to the leaves of the foliation of \mathcal{T}_C as **b**-arcs and **c**-circles. Each **b**-arc and each **c**-circle lies in both \mathcal{T}_C and in some fiber $H_\theta \in \mathbf{H}$. Since $K \subset \mathcal{T}_C$, for all generic $H_\theta \in \mathbf{H}$, each point of $K \cap H_\theta$ is contained in a **b**-arc or **c**-circle. Finally, since K intersects each disc fiber of \mathbf{H} coherently, K must intersect each non-singular leaf coherently.

A **b**-arc, $b \subset \mathcal{T}_C \cap H_\theta$, is *essential* if either $b \cap K \neq \emptyset$, or both sides of H_θ split along b are intersected by K . A **c**-circle, $c \subset \mathcal{T}_C \cap H_\theta$, is *essential* if $c \cap K \neq \emptyset$. The definition of essential **b**-arcs and **c**-arcs is an adaptation of the definition in [4], however inessential **b**-arcs (**c**-circles) are still arcs (respectively, circles) splitting off sub-discs (respectively, bounding subdiscs) of disc fibers that are not intersected by K .

Lemma 2.2 *Assume that \mathcal{T}_C satisfies (i)–(iii) of Lemma 2.1. Then \mathcal{T}_C is isotopic to a cabling torus, \mathcal{T}'_C , such that the foliation of \mathcal{T}'_C also satisfies (i)–(iii) and in addition:*

- (1) *All **b**–arcs are essential.*
- (2) *All **c**–circles are essential.*
- (3) *Any **c**–circle in the foliation is homotopically non-trivial on \mathcal{T}'_C .*

Moreover, the restriction of the isotopy to K is the identity.

Proof The argument for eliminating inessential **b**–arcs is exactly the same as the argument used in the proof of Lemma 2 of [4]. Similarly, if c is a **c**–circle in the foliation which is homotopically trivial on \mathcal{T}_C , we must have $c \cap K = \emptyset$ since K cannot intersect a homotopically trivial circle on \mathcal{T}_C coherently. Then we can, again, employ the argument in the proof of Lemma 2 [4], which relies on the fact that c bounds a sub-disc in the disc fiber that does not intersect K . □

We now consider the different types of singularities which can occur in the foliation of \mathcal{T}_C . Having two possible non-singular leaves allows for the occurrence of three possible types of singularities: a **bb**–singularity resulting from two **b**–arcs “meeting” at a saddle point; a **bc**–singularity resulting from a **b**–arc and a **c**–circle forming a saddle point; and **cc**–singularity resulting from two **c**–circles. See Figure 2.




Figure 2

Lemma 2.3 *cc–singularities do not occur.*

Proof A surgery between two \mathbf{c} -circles will produce a homotopically trivial \mathbf{c} -circle, violating statement 3 of Lemma 2.2. \square

We can now repeat the discussion in Section 1 of [4] which is based on the observation that the existence of two different types of singularities implies the occurrence of two different tiles: \mathbf{bb} -tiles and \mathbf{bc} -tiles as illustrated in Figure 2. (The reader should note that a \mathbf{bc} -tile is in fact annular.)

Proposition 2.4 *There are three possibilities for the foliation of \mathcal{T}_C :*

- *A circular foliation: every leaf is a \mathbf{c} -circle. There are no singularities.*
- *A mixed foliation: there are both \mathbf{b} -arcs and \mathbf{c} -circles. This foliation needs the occurrence of \mathbf{bc} -tiles, but allows the occurrence of \mathbf{bb} -tiles.*
- *A tiling: that is, a foliation involving only \mathbf{b} -arcs and, thus, only \mathbf{bb} -tiles. (An example of a tiled torus is given in Figure 7.)*

At this point an interested reader can reach the punch-line in rapid fashion by reviewing the statements of Propositions 4.1 and 5.1, and then proceeding to the argument in Section 3. No understanding of the machinery in the intervening subsections is needed to understand the proof of Theorem 1.1 in Section 3. (To understand Remark 3.2 it is useful to have reviewed the discussion on foliations in Section 2.1 and Section 2.2.)

2.1 Isotopies of K in \mathcal{T}_C

In this section we discuss how the positioning of K in the foliation of \mathcal{T}_C at times allows us to perform destabilizations and exchange moves on K . The reader should note that any isotopy of K induces an isotopy of \mathcal{T}_C .

We adopt the terminology of the literature, referring to the points of $\mathbf{A} \cap \mathcal{T}_C \subset \mathcal{T}_C$ as *vertices*. A vertex v is *adjacent* to a leaf in the foliation if v is an endpoint of the leaf. The *valence* of a vertex v is the number of singular leaves which end at v .

A *meridian curve* $m \subset \mathcal{T}_C$ is a curve which intersects K coherently p_h -times and bounds a disc Δ_m in S^3 , with $\Delta_m \cap \mathcal{T}_C = m$. We may assume that $m \cap \mathbf{A} = \emptyset$ and that m is transverse to the leaves of the foliation on \mathcal{T}_C . Being transverse to leaves makes m a closed braid and we orient m so that it has positive linking number with \mathbf{A} . The orientation on m is consistent with the forward direction of the foliation on \mathcal{T}_C .

The disc Δ_m is contained in the solid torus \mathcal{V}_C where $\partial\mathcal{V}_C = \mathcal{T}_C$. We also assign an orientation to \mathcal{T}_C such that the positive side of \mathcal{T}_C “points” towards \mathcal{V}_C . This choice of orientation allows us to make parity assignments to the vertices and singularities in the foliation of \mathcal{T}_C . A vertex or singularity is *positive* if the orientation of \mathbf{A} agrees with the orientation of the positive normal vector of \mathcal{T}_C at the vertex or singularity, otherwise it is *negative*. (We will deal with meridian curves and discs in Section 5. Their introduction at this point was only necessary for understanding parity assignments to vertices and singularities.)

Lemma 2.5 *Let $\Delta \subset \mathcal{T}_C$ be a sub-disc such that $\partial\Delta = \gamma \cup \rho$ where $\rho \subset K$ and γ is an arc contained in a singular leaf of the foliation. Furthermore, suppose that $\text{int}(\Delta)$ contains exactly one vertex v and no singular points. (See Figure 3.) Then an isotopy of $\rho \subset K$ across Δ to a new position which is transverse to the leaves of the foliation corresponds to a destabilization of the braid K .*




Figure 3: The vertex v has type (a).

Proof If we split \mathcal{T}_C along K we produce an annulus having two copies of K as its boundary. The foliation of this annulus will contain a valence one vertex, v , and the foliated neighborhood of v will be a *type (a) vertex* as described in Section 2.3 of [1]. (See Figure 3.) As in [1], the isotopy of ρ across Δ corresponds to a destabilization of K . \square

Lemma 2.6 *Let $\Delta \subset \mathcal{T}_C$ be a sub-disc such that $\partial\Delta = \gamma_+ \cup \gamma_- \cup \rho$, where γ_{\pm} is included in a \pm -singular leaf of the foliation and where ρ is a subarc of the knot K . (See Figure 4.) Suppose that the arcs γ_+ and γ_- have one endpoint at a common vertex v_1 and that $\text{int}(\Delta)$ contains exactly one vertex v_0 and no singular points. Then an isotopy of ρ through Δ to ρ' corresponds to an exchange move on the braid K , as illustrated in Figure 1(a).*

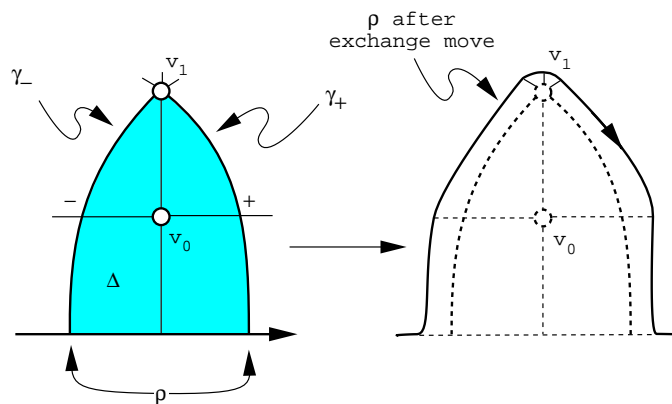


Figure 4

Proof Again, we split \mathcal{T}_C along K to produce an annulus whose boundary consists of two copies of K . The foliation of this annulus will contain a valence two vertex, v_0 . The foliated neighborhood of v_0 (see Figure 4) will correspond to a *type (a,b) vertex* as described in Section 2.4 of [1]. As in [1], the isotopy of ρ across such a neighborhood corresponds to an exchange move. \square

2.2 Manipulating the foliation of \mathcal{T}_C

Two operations, *change of foliation* and *elimination of a valence two vertex*, played an important role in establishing control over the foliation of the torus in [4]. We now discuss how these two operations are adapted to our present situation, where we must deal not just with the foliation of the torus \mathcal{T}_C , but with the pair (K, \mathcal{T}_C) .

Change of foliation Let $R \subset \mathcal{T}_C$ be the topological disc that is closure of a connected region foliated by \mathbf{b} -arcs having common endpoints at vertices $v_+, v_- \subset \mathcal{T}_C$. (The subscripts indicate the parity.) Let $s_1, s_2 \subset R \subset \mathcal{T}_C$ be the two singularities that are on the boundary of R and assume that their parity is the same. The existence of such a region R is the central assumption in the discussion in Section 2.1 of [1]. Specifically, Theorem 2.1 of [1] allows the two singular points to be either \mathbf{bc} or \mathbf{bb} singularities. Figure 5 shows how the application of this result from [1] alters the foliation of \mathcal{T}_C . To adapt the [1] change of foliation to our present situation we need only check that the presence of the knot $K \subset \mathcal{T}_C$ does not obstruct the change in foliation.

How might K intersect R ? We consider an arc $\alpha \subset K \cap R$. Since K is transverse to the foliation of \mathcal{T}_C we know that α is transverse to all of the




Figure 5: Parity of s_1 equals parity of s_2 . In case 1 there are two possible ways to change the fibration of \mathbf{H} and, thus, the foliation of \mathcal{T}_C . In case 2 there is only one way to alter the fibration.

\mathbf{b} -arcs of R and is away from the singular and vertex points of R . We can list the possible cases for a component $\alpha \subset K \cap R$ as follows:

- (1) The arc α splits the disc R into two discs, one containing the two singular points and a vertex, and the other containing only one vertex.
- (2) The arc α splits the disc R into two discs, each of which contains a vertex and a singular point.

We recall (see [1]) that there are two possible ways the foliation of R can be altered. If case 1 occurs then either of these foliation changes is permissible. (See Figure 5.) If case 2 occurs then only one of the changes in foliation is possible because only one of the changes results in K still being transverse to the foliation. Figure 5 illustrates how a case 2 arc α determines the local change in foliation in R . The proof that these changes in foliation correspond to braid isotopies of the (K, \mathcal{T}_C) pair is straight forward, but has numerous details. We will not repeat the argument in [1]. We will refer to this braid isotopy (which only alters the foliation of \mathcal{T}_C in a disc neighborhood of R) as a *change of foliation*. The following lemma describes the main features of the change in foliation that we need:

Lemma 2.7 *Let R be the closure of a region foliated by \mathbf{b} -arcs such that there are vertices $v_+, v_- \subset \partial R$ and singularities $s_1, s_2 \subset \partial R$. Then there*

exists a change of foliation such that the valence of these two vertices has been decreased.

Elimination of valence two vertices We next consider the configuration in the left sketch in Figure 6. (For the moment the reader should ignore the arcs labeled α in Figure 6. They will be referred to in the proof of Lemma 2.8.)




Figure 6: The vertex v_0 is a type (b,b) vertex. In the situation where the case 2 α arcs are absent, after an exchange move, the \mathbf{b} -arcs between v_0 and v_2 are inessential.

Let $v_0 \subset \mathcal{T}_C$ be a vertex of valence 2. Let $star(v_0)$ be the topological disc which is the closure of the union of all the \mathbf{b} -arcs having v_0 as an endpoint. Let ϵ be the parity of v_0 then $star(v_0)$ contains exactly two additional vertices, v_1 and v_2 , which will necessarily have parity $-\epsilon$. By Lemma 3.1 of [1], $star(v_0)$ must contain both $+$ and $-$ singular points. So let $s_+, s_- \subset star(v_0)$ be the singular points of positive and negative parity, respectively. In this situation we have the following lemma.

Lemma 2.8 *Let v_0 be a valence two vertex, with $star(v_0)$ topologically a disc. Suppose $star(v_0)$ contains vertices $v_1 \cup v_2$ and singularities $s_+ \cup s_-$. Then, after an isotopy of K involving only exchange moves, destabilizations and braid isotopies, we can locally eliminate vertices v_0 and v_1 (or v_0 and v_2) and singularities s_+ and s_- in the foliation of \mathcal{T}_C , as illustrated in the right sketch in Figure 6. Moreover, away from $star(v_0)$ the foliation of \mathcal{T}_C remains unchanged.*

Proof By Theorem 2.2 of [1], we know that if $K \cap star(v_0) = \emptyset$ then it is possible to perform an exchange move on K such that the foliation of \mathcal{T}_C remains unchanged (except for a change in the cyclic ordering of the vertices along \mathbf{A}). After the exchange move, the \mathbf{b} -arcs of $star(v_0)$ will be outermost,

ie, for \mathbf{b} -arc $b \subset \text{star}(v_0)$, if $b \subset H_\theta \in \mathbf{H}$ then b splits off a sub-disc $\Delta \subset H_\theta$ such that $\text{int}(\Delta) \cap K = \emptyset$. Furthermore, if $K \cap \text{star}(v_0) = \emptyset$ then \mathcal{T}_C can be isotoped along Δ , eliminating v_0 and, say, v_2 , and the two singularities, s_+ and s_- . The foliation away from $\text{star}(v_0)$ remains unchanged, but globally it has two fewer vertices and two fewer singularities.

How might we eliminate v_0 if $K \cap \text{star}(v_0) \neq \emptyset$? As with the change of foliation, we consider an arc $\alpha \subset K \cap \text{star}(v_0)$. Since K is transverse to the foliation of \mathcal{T}_C we know that α is transverse to all of the \mathbf{b} -arcs in $\text{star}(v_0)$ and does not intersect the singular points and vertices of $\text{star}(v_0)$. We list the possibilities for a component $\alpha \subset K \cap \text{star}(v_0)$:

- (1) The arc α splits the disc $\text{star}(v_0)$ into two discs, one containing the two singularities and two of the three vertices, and the other containing only one vertex. (See Figure 6.)
- (2) The arc α splits the disc $\text{star}(v_0)$ into two discs, one containing v_0 and a single singularity (say s_+), and the other containing v_1 , v_2 and s_- . (See Figure 6.)

To deal with case 1, we must in fact consider two situations: there exists \mathbf{b} -arcs of $\text{star}(v_0)$ that do not intersect K ; and K intersects every \mathbf{b} -arc of $\text{star}(v_0)$. If some \mathbf{b} -arc of $\text{star}(v_0)$ is not intersected by K then after the exchange move there will necessarily be \mathbf{b} -arcs that are inessential. The situation where K intersects every \mathbf{b} -arc of $\text{star}(v_0)$ requires a little more work.

Notice that if K intersects every \mathbf{b} -arc of $\text{star}(v_0)$, there is necessarily an arc $\gamma \subset \partial(\text{star}(v_0))$ contained in a singular leaf and having endpoints v_1, v_2 which K must intersect incoherently. The topology of \mathcal{T}_C , thus, forces the existence of a sub-arc $\gamma' \subset \gamma$, a path $\rho \subset K$, and a sub-disc $\Delta \subset \mathcal{T}_C$ such that: $\Delta \cap \gamma = \gamma'$; and $\Delta \cap K = \rho$. Since $\Delta \subset \mathcal{T}_C$, Δ inherits a foliation that basically mimics that of a Seifert disc for the unknot. Theorem 4.3 of [1] allows us to isotop ρ through Δ using exchange moves and destabilizations until K is moved off γ and we have an arc $\alpha \subset \text{star}(v_0)$ that corresponds to case 2.

Finally, to deal with case 2, notice that the sub-disc containing v_0 that α splits off has a foliation corresponding to that of a valence one vertex. (See Figure 6.) This is the configuration in Lemma 2.5. We can, thus, destabilize K to remove a case 2 intersection arc. \square

3 Proof of Theorem 1.1 in the case of a circular foliation on \mathcal{T}_C

In this section we assume that K is contained in a torus \mathcal{T}_C that does not intersect the axis \mathbf{A} , ie, \mathcal{T}_C has a circular foliation. With this assumption our argument for proving Theorem 1.1 is inductive and we need the following result.

Corollary 3.1 *Let \mathcal{K} be an exchange reducible knot type. Then $\mathbf{C}(\mathcal{K}, (p, q))$ is also exchange reducible.*

Proof Let $K_{(p,q)}$ be any closed braid representative of $\mathbf{C}(\mathcal{K}, (p, q))$ and \mathcal{T}_C be a cabling torus. By Proposition 5.1 we know that if the foliation of \mathcal{T}_C is a tiling then through a sequence of exchange moves and destabilizations we can replace this tiling with a mixed foliation. By Proposition 4.1 we can, through a sequence of exchange moves and destabilizations, replace a mixed foliation of \mathcal{T}_C with a circular foliation. The core of this circularly foliated torus is a braid representative of the cabling knot \mathcal{K} which we call K . By assumption K represents an exchange reducible knot type.

Now, referring back to Figure 1, we notice that we can alter the destabilizing move in (b) by replacing the weight of 1 with a weight of p , ie, we think of p parallel strands instead of 1 strand. Similarly, in (a) we can replace the weight of 1 on the strands involved in the exchange move isotopy with a weight of p . Thus, a destabilization of K results in p destabilizations of $K_{(p,q)}$, and an exchange move on K results in an exchange move on $K_{(p,q)}$ —the peripheral torus of K is the cabling torus circularly foliated. By a classical result in [12] (cf Satz 23.2), once K is of minimal braid index, we will have $K_{(p,q)}$ achieving its minimal braid index when the cabling torus is circularly foliated. \square

Now, let U be the unknot. By Theorem 1 of [2], U is exchange reducible. Then, by Corollary 3.1 $\mathbf{C}(U, (p_1, q_1))$ is exchange reducible.

Inductively, suppose that

$$\mathbf{C}(\cdots \mathbf{C}(\mathbf{C}(\mathcal{C}_0, (p_1, q_1)), (p_2, q_2)) \cdots, (p_{i-1}, q_{i-1}))$$

is exchange reducible. Then, again, by Theorem 1 of [2]

$$\mathbf{C}(\mathbf{C}(\cdots \mathbf{C}(\mathbf{C}(\mathcal{C}_0, (p_1, q_1)), (p_2, q_2)) \cdots, (p_{i-1}, q_{i-1})), (p_i, q_i))$$

is exchange reducible. Thus, $K_{(P,Q)}$ is exchange reducible and Theorem 1.1 is established for this special case of the foliation of \mathcal{T}_C .

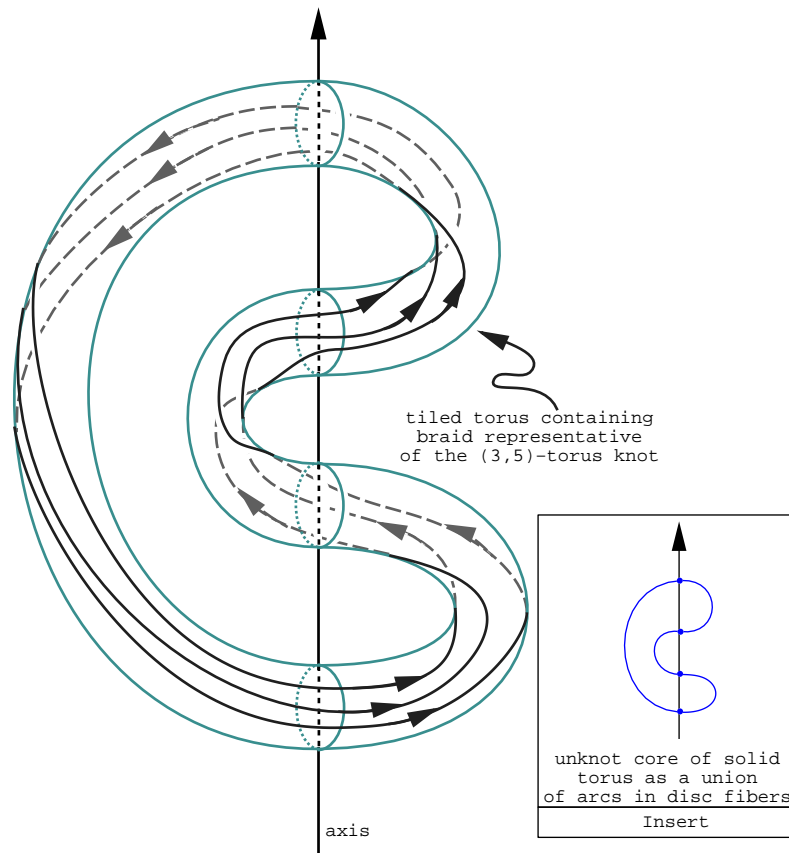


Figure 7

Remark 3.2 The ease with which we are able to prove Theorem 1.1 for this special case is due to the fact that the core of $\mathcal{V}_{\mathcal{C}}$ can be canonically chosen as a closed braid in the complement of \mathbf{A} , ie, $\mathcal{C} \subset \mathcal{V}_{\mathcal{C}}$ can be taken as the union of the “centers” of the sub-discs in the disc fibers of \mathbf{H} that are bounded by \mathbf{c} -circle. The difficulty with the remaining two cases—mixed and tiled foliations—is that there is no similar canonical choice for the core of $\mathcal{V}_{\mathcal{C}}$. In Figure 7, we give an example of a tiled $\mathcal{T}_{\mathcal{C}}$ containing the $(3,5)$ -torus knot $K_{(3,5)}$. This $\mathcal{T}_{\mathcal{C}}$ can be seen as the peripheral torus of an unknot that is represented by the union of four arcs that have their endpoints on \mathbf{A} and are contained in disc fibers of \mathbf{H} . The insert in Figure 7 depicts this representation of the unknot—a core of $\mathcal{V}_{\mathcal{C}}$. Since this core intersects \mathbf{A} at four points, $\mathcal{T}_{\mathcal{C}}$ will necessarily intersect \mathbf{A} in eight points, ie four times two. Because of Euler characteristic considerations, the tiling of $\mathcal{T}_{\mathcal{C}}$ will then have eight \mathbf{b} tiles. Notice that there are several possible

ways we can ϵ -push this unknotted core off \mathbf{A} so that resulting \mathcal{C} is transverse to \mathbf{H} . In particular, it is possible to produce a transverse \mathcal{C} such that any \mathbf{b} -arc $b \subset \mathcal{T}_{\mathcal{C}} \cap H_t \subset H_t$ with $b \cap K_{(3,5)} = \emptyset$ splits the set $\mathcal{C} \cap H_t$ in H_t , for $H_t \in \mathbf{H}$. That is, all of the \mathbf{b} -arcs of $\mathcal{T}_{\mathcal{C}}$ are *essential* in the complement of $\mathcal{C} \sqcup K_{(3,5)}$. Using an exchange move or destabilization isotopy of \mathcal{C} to induce an isotopy of $\mathcal{T}_{\mathcal{C}}$ then becomes unworkable. Thus, the key strategy in the remaining two cases will be to find inessential \mathbf{b} -arcs on $\mathcal{T}_{\mathcal{C}}$ in the knot complement that will enable us to simplify $\mathcal{T}_{\mathcal{C}}$.

4 Replacing mixed foliations with circular foliations

In this section we assume that we are given the braid-torus pair $(K, \mathcal{T}_{\mathcal{C}})$, and that the induced foliation on $\mathcal{T}_{\mathcal{C}}$ satisfies the conclusions of Lemmas 2.1, 2.2 and 2.3. Moreover, we assume that the foliation of $\mathcal{T}_{\mathcal{C}}$ contains \mathbf{c} -circles. Our goal is to reduce this case to the special case when the foliation contains only \mathbf{c} -circles.

Proposition 4.1 *Let $\mathcal{T}_{\mathcal{C}}$ have a mixed foliation. Then after a sequence of exchange moves and destabilizations of K and isotopies of $\mathcal{T}_{\mathcal{C}}$ the foliation of $\mathcal{T}_{\mathcal{C}}$ may be assumed to be circular.*

Proof Following the discussion in Section 3 of [4], we observe that a singularity between a \mathbf{b} -arc and a \mathbf{c} -circle foliates a \mathbf{bc} -annulus. Moreover, since a \mathbf{c} -circle must have a \mathbf{bc} -singularity both in the forward and backward direction, these \mathbf{bc} -annuli occur in pairs.

If we adjoin two \mathbf{bc} -annuli along their common \mathbf{c} -circle we will produce an annular region, W , that has each boundary curve the union of two \mathbf{b} -arcs and, thus, contains two vertices. See Figure 8.

We can cut W open along two new disjoint edges, e and e' , each having its endpoints at vertices which are on distinct components of ∂W , as in the bottom of Figure 8. We see that W is the union of two \mathbf{be} -tiles where the boundary of a \mathbf{be} -tile has two \mathbf{b} -arcs and two of the new \mathbf{e} -arcs. The number of \mathbf{be} -tiles constructed in a mixed foliation is exactly equal to the number of \mathbf{bc} -annuli. The vertices in the \mathbf{be} -tiling still correspond to the points of $\mathbf{A} \cap \mathcal{T}_{\mathcal{C}}$.

The cellular decomposition of $\mathcal{T}_{\mathcal{C}}$ coming from the \mathbf{be} -tiling yields an ‘‘Euler characteristic’’ formula (equation (1) of [4]). Namely, let $V(\beta, \epsilon)$ be the number




Figure 8

of vertices in the **be**-tiling of \mathcal{T}_C that is adjacent to β **b**-arcs and ϵ **e**-arcs. Then we have:

$$2V(2,0) + 2V(1,1) + V(2,1) + V(3,0) = \sum_{i=5}^{\infty} \sum_{\beta=\lfloor i/2 \rfloor}^i (i-4)V(\beta, i-\beta) \quad (1)$$

where both the left hand side and the right hand side are non-negative.

Referring to the discussion in [4] we recall that $V(1,1) = V(2,1) = 0$, since both a **be**-vertex (respectively, **bbe**-vertex) cannot be geometrically realized. A **be**-vertex (respectively, **bbe**-vertex) is one that is cyclically adjacent to **b**-arcs and then **e**-arcs (respectively, **b**-arcs and then, after a singularity, to **b**-arcs and then **e**-arcs). This vertex notation generalizes in the obvious manner. If $V(3,0) \neq 0$ then we have a **bbb**-vertex. Such a vertex, v , must be adjacent to singular leaves of both positive and negative parity (cf Lemma 3.1, [1]). Thus, since v is of odd valence, $star(v)$ will contain a sub-disc region satisfying the assumptions of Lemma 2.7. We can then perform the change of foliation decreasing the valence of v so that it becomes a **bb**-vertex.

If $V(2,0) \neq 0$ then there exists a vertex, v , that is a **bb**-vertex. $star(v)$ will satisfy the assumption of Lemma 2.8. After possibly some sequence of exchange moves and destabilizations, we can simplify the tiling of \mathcal{T}_C so that $V(2,0) = 0$.

Since we can now assume the left side of equation (1) is zero, we can assume that the only possible vertices in the tiling of \mathcal{T}_C are **bebe**-vertices. Figure 9 illustrates the foliation of the annular neighborhood around such valence four vertices. Notice that the **c**-circles must be intersected coherently p_h -times by K . If K does intersect all of the **b**-arcs of this local foliation then there

will exist a sub-disc satisfying the assumptions of Lemma 2.5. We can then destabilize K so that, as in Figure 9, there will be \mathbf{b} -arcs that K does not intersect. An innermost such annulus (one having consecutive vertices on \mathbf{A}) will then have an inessential \mathbf{b} -arc. We may eliminate it by an isotopy of \mathcal{T}_C . \square




Figure 9: K enters into the foliated region and intersects a singular leaf twice. This allows for a destabilization of K . The shaded region Δ corresponds to the Δ of Lemma 2.5.

5 Replacing tilings with mixed foliations

As in the previous section, we start by assuming that we are given the braid-torus pair (K, \mathcal{T}_C) such that the induced foliation on \mathcal{T}_C satisfies the conclusions of Lemmas 2.1, 2.2 and 2.3. But, in this section, we assume that the foliation of \mathcal{T}_C contains no \mathbf{c} -circles. Our goal in this section is to prove that after a sequence of exchange moves, braid isotopies and destabilizations of K , and isotopies of \mathcal{T}_C , we may assume that the foliation of \mathcal{T}_C can be assumed to be mixed.

Recall $m \subset \mathcal{T}_C$ is a meridian curve and $\Delta_m \subset \mathcal{V}_C$ is a meridian disc. Since Δ_m is a spanning surface having a closed braid boundary, the induced foliation on the disc Δ_m satisfies the conclusion of Theorem 1.1, part (i), of [1]. We review the most important features of this foliation, all of which are developed in detail in [1]. If we consider $\Delta_m \cap H_\theta \subset H_\theta$, where H_θ is a generic disc fiber, we see that this intersection contains two types of arcs: the first, called \mathbf{a} -arcs, have one endpoint on \mathbf{A} and one endpoint on m ; the second, called \mathbf{b} -arcs, have both endpoints on \mathbf{A} . The singularities are of three types: \mathbf{aa} -singularities

(saddle-point singularities formed by two **a**-arcs); **ab**-singularities (saddle-point singularities formed by an **a**-arc and a **b**-arc); and **bb**-singularities (saddle-point singularities formed by two **b**-arcs). Note that the orientation on m induces an orientation on Δ_m , and that this allows us to assign parities to the vertices (points in $\Delta_m \cap \mathbf{A}$) and singularities in the foliation in the same manner as we did in Section 2 for \mathcal{T}_C .

The argument in this section will involve “simplifying”, first, the pair (K, \mathcal{T}_C) then, second, the triple (K, \mathcal{T}_C, m) and, finally, the quadruple $(K, \mathcal{T}_C, m, \Delta_m)$. Our goal is to prove the following proposition.

Proposition 5.1 *Let (K, \mathcal{T}_C) be a knot–torus pair such that the foliation of \mathcal{T}_C is a tiling. Then, after a sequence of exchange moves and destabilizations of K , the foliation of \mathcal{T}_C may be assumed to be mixed.*

5.1 The knot–torus pair

In this subsection we define the *complexity* of (K, \mathcal{T}_C) to be $\Lambda(K, \mathcal{T}_C) = (n_1, n_2)$, where n_1 is the braid index of K and n_2 is the number of vertices in the foliation of \mathcal{T}_C . We use lexicographical ordering on the 2-tuples (n_1, n_2) to give an ordering on $\Lambda(K, \mathcal{T}_C)$.

Following the discussion in Section 2 of [4], we recall that the foliation of \mathcal{T}_C (absent of **c**-circles) yields a tiling of \mathcal{T}_C by **bb**-tiles. This tiling in turn gives us a cellular decomposition of \mathcal{T}_C and, thus, an “Euler characteristic” formula (compare with equation (1) of [4]). Let $V(\beta)$ be the number of vertices in the **bb**-tiling of \mathcal{T}_C that are adjacent to β **b**-arcs. We have:

$$2V(2) + V(3) = V(5) + 2V(6) + 3V(7) + \dots \quad (2)$$

where both the left hand side and the right hand side are non-negative.

If $V(3) \neq 0$ then we have (as in [4]) a **bbb**-vertex, v . Since $star(v)$ must contain both positive and negative singularities (see Lemma 3.1 of [1]), it must contain a sub-disc region satisfying the assumptions of Lemma 2.7. After a change of foliation we can assume that the valence of v is two. We can then apply the procedure in Lemma 2.8 to reduce n_2 and, thus, $\Lambda(K, \mathcal{T}_C)$.

If $V(2) \neq 0$ then we can apply the procedure in Lemma 2.8 straight away to reduce $\Lambda(K, \mathcal{T}_C)$. After some number of changes of these two types we may assume that the LHS of Equation (2) is zero. But then the RHS is too. We can conclude that the only possibility is that $V(4)$ is non-zero. Moreover, if

$v \in \mathcal{T}_C$ is a **bbbb**-vertex then the four singular leaves which intersect at v are $+, -, +, -$ in that cyclic order. (Otherwise, a change of foliation could be performed to reduce the valence of v to three.) The salient features of this *standard tiling* are: all vertices of \mathcal{T}_C are valence four; the parity pattern on the tiling is a “checkerboard” pattern; and all **b**-arcs are essential. The following proposition summarizes the above discussion.

Proposition 5.2 *Let (K, \mathcal{T}_C) be a knot–torus pair where the foliation of \mathcal{T}_C is a tiling. If the tiling of \mathcal{T}_C is non-standard then, through the use of exchange moves, we may replace the knot–torus pair with (K', \mathcal{T}'_C) such that \mathcal{T}'_C has a foliation that is either a standard tiling or a mixed foliation. Moreover, $\Lambda(K', \mathcal{T}'_C) < \Lambda(K, \mathcal{T}_C)$.*

Next, we define four graphs, $G_{+,+}$, $G_{+,-}$, $G_{-,+}$ and $G_{-,-}$ in \mathcal{T}_C . The vertices of $G_{\delta,\epsilon}$ are the vertices of \mathcal{T}_C having parity δ and the edges of $G_{\delta,\epsilon}$ are sub-arcs of singular leaves which join the two δ vertices in the **bb**-tiles having a singularity of parity ϵ . Notice that the definition of these graphs forces $G_{+,+}$ ($G_{+,-}$) to be disjoint from $G_{-,-}$ (respectively, $G_{-,+}$). Moreover, the parity checkerboard pattern to the standard tiling implies that each vertex of any of the four graphs is adjacent to exactly two graph edges. Thus, each component of $G_{\delta,\epsilon}$ is a simple closed curve (scc) on \mathcal{T}_C which, by Lemma 3.8(i) of [1], is also homotopically non-trivial on \mathcal{T}_C . (For a more complete analysis of $G_{\delta,\epsilon}$, see Section 3 of [1].)

Lemma 5.3 *Let $\mathcal{C} \subset \mathcal{T}_C$ be a curve such that $\mathcal{C} \in \{K, m\}$. Then we can apply either a destabilization or an exchange move to \mathcal{C} in the following situations.*

- (a) *Suppose that as \mathcal{C} is traversed, a sub-arc $\rho \subset \mathcal{C}$ has the following sequential intersections with our four graphs: $G_{\delta,\epsilon} \rightarrow G_{-\delta,-\epsilon} \rightarrow G_{\delta,-\epsilon} \rightarrow G_{-\delta,-\epsilon} \rightarrow G_{\delta,\epsilon}$. Then we can destabilize \mathcal{C} along ρ .*
- (b) *Suppose that as \mathcal{C} is traversed, a sub-arc $\rho \subset \mathcal{C}$ has the following sequential intersections with our four graphs: $G_{\delta,\epsilon} \rightarrow G_{-\delta,-\epsilon} \rightarrow G_{\delta,-\epsilon} \rightarrow G_{-\delta,\epsilon}$. Then \mathcal{C} admits an exchange move along ρ .*

Proof Referring to Figure 10(a), we see that a sequential intersection pattern $G_{\delta,\epsilon} \rightarrow G_{-\delta,-\epsilon} \rightarrow G_{\delta,-\epsilon} \rightarrow G_{-\delta,-\epsilon} \rightarrow G_{\delta,\epsilon}$ implies the existence of a sub-disc in $\mathcal{T}_C - \mathcal{C}$ satisfying the assumptions in Lemma 2.5. Conclusion (a) follows.

Referring again to Figure 10(b), we observe that a sequential intersection pattern $G_{\delta,\epsilon} \rightarrow G_{-\delta,-\epsilon} \rightarrow G_{\delta,-\epsilon} \rightarrow G_{-\delta,\epsilon}$ implies the existence of a sub-disc in $\mathcal{T}_C - \mathcal{C}$ satisfying the assumptions in Lemma 2.6. Conclusion (b) follows. □

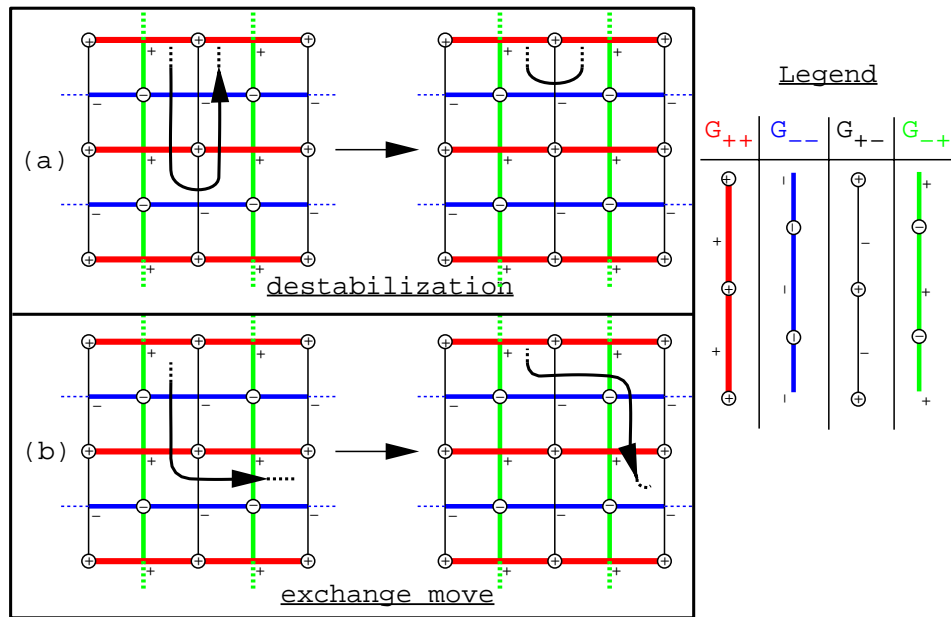


Figure 10: In (a) the intersection pattern is $G_{-,-} \rightarrow G_{+,+} \rightarrow G_{-,+} \rightarrow G_{+,+} \rightarrow G_{-,-}$. In (b), the intersection pattern is $G_{-,-} \rightarrow G_{+,+} \rightarrow G_{+,-} \rightarrow G_{-,+}$.

Lemma 5.4 *Let $\mathcal{C} \subset \mathcal{T}_C$ be a curve such that $\mathcal{C} \in \{K, m\}$. If \mathcal{C} intersects a component of $G_{\delta, \epsilon}$ incoherently then, after a sequence of exchange moves, we can reduce the braid index of \mathcal{C} using a destabilization.*

Proof Suppose that \mathcal{C} intersects $G_{+,+}$ incoherently. Then there exists sub-arcs $\gamma \subset G_{+,+}$ and $\rho \subset \mathcal{C}$ and a sub-disc $\Delta \subset \mathcal{T}_C$ such that $\partial\Delta = \gamma \cup \rho$. Without loss of generality, we can assume that $int(\Delta) \cap \mathcal{C} = \emptyset$. Now, notice that the standard tiling of \mathcal{T}_C forces the existence of a sub-arc $\rho' \subset \rho$ which satisfies the sequential intersection pattern (a) or (b) of Lemma 5.3. Moreover, it can be assumed that the existence of the type (a) or (ab) vertex from the proofs of Lemmas 2.5 and 2.6 (which the intersection sequences of Lemma 5.3 invoke) are contained in Δ .

If the sequence in Lemma 5.3(a) occurs then we immediately have the conclusion of our lemma. If the sequence in Lemma 5.3(b) occurs then, after performing the exchange move on ρ' , we will again have sub-arcs γ and ρ . But, they now bound a sub-disc Δ' that has fewer vertices in its induced foliation than Δ . Iterating this procedure we conclude that at some point ρ must have an intersection sequence with the graphs that corresponds to the sequence in Lemma 5.3(a). \square

We now focus on understanding how K is contained in the foliation of \mathcal{T}_C . Specifically, we have the following application of Lemma 5.4.

Proposition 5.5 *Let (K, \mathcal{T}_C) be a knot–torus pair where \mathcal{T}_C has a standard tiling. Then there exists a knot $K' \subset \mathcal{T}_C$ that is transverse to the foliation of \mathcal{T}_C and is obtained from K through a sequence of exchange moves and destabilizations, with $\Lambda(K', \mathcal{T}_C) \leq \Lambda(K, \mathcal{T}_C)$ at every change, such that K' coherently intersects components of the graphs $G_{\pm, \pm}$ and $G_{\pm, \mp}$.*

Proof If K does not intersect graph components coherently then, after repeated application of Lemma 5.4, we can assume that K has been replaced by K' satisfying the conclusion of the proposition. If any of the \mathbf{b} –arcs of \mathcal{T}_C are now inessential, we perform the necessary isotopy of \mathcal{T}_C to remove them. The new foliation of \mathcal{T}_C will have fewer vertices and singularities. We then repeat the applications of Lemmas 2.7 and 2.8, emulating the argument at the begin of this section, so that $V(2)$ and $V(3)$ of equation (2) are again zero and the tiling of \mathcal{T}_C is checkerboard by the singularity parity values. All operations on K and the tiling of \mathcal{T}_C are non-increasing on the complexity measure. \square

For a scc $\mathcal{C} \subset \mathcal{T}_C$ transverse to the foliation of \mathcal{T}_C , let $\mathcal{S}_C \subset \mathcal{T}_C$ be the closure of the union of all the \mathbf{b} –arcs that \mathcal{C} intersects in the foliation of \mathcal{T}_C . We call \mathcal{S}_C the \mathbf{b} –support of \mathcal{C} . The definition of \mathbf{b} –support implies a useful fact about the “width” of \mathcal{S}_K : If $\gamma \subset \mathcal{S}_K$ is properly embedded (that is, $(\text{int}(\gamma), \partial\gamma) \subset (\text{int}(\mathcal{S}_K), \partial\mathcal{S}_K)$, which is the union of \mathbf{b} –arcs) then γ is in fact a single \mathbf{b} –arc which is non-parallel to the boundary of \mathcal{S}_K . We will refer to this fact as (\star) .

Next, we define two modifications of the \mathbf{b} –support of \mathcal{C} .

Exchange move with type-I support Let $v_\epsilon, v_{-\epsilon} \in \partial\mathcal{S}_K$ be two vertices such that if $b \subset \mathcal{S}_K$ is a \mathbf{b} –arc having v_o as an endpoint then v_i is also an endpoint of b . More descriptively, these conditions correspond to having v_o as an *outside corner* and v_i is an *inside corner* of $\partial\mathcal{S}_K$. For each point of $K \cap b$ we can find an arc neighborhood $b \subset \gamma \subset K$ such that γ has the sequential intersection pattern of $G_{-\epsilon, -\delta} \rightarrow G_{\epsilon, \delta} \rightarrow G_{\epsilon, -\delta} \rightarrow G_{-\epsilon, \delta}$ as depicted in Figure 10(b). The alteration in the \mathbf{b} –support that results from performing the exchange move in Figure 11 to each γ arc intersecting b is referred to as an *exchange move with type-I support*.

Exchange move with type-II support Let $v_\epsilon^t, v_{-\epsilon}^1, v_{-\epsilon}^2, v_{-\epsilon}^3 \in \partial\mathcal{S}_K$ be four vertices such that if $b \subset \mathcal{S}_K$ is a \mathbf{b} –arc having v_ϵ^t as an endpoint then its other b endpoint is contained in $\{v_{-\epsilon}^1, v_{-\epsilon}^2, v_{-\epsilon}^3\}$. Let $b^1, b^2, b^3 \subset \partial\mathcal{S}_K$ be three \mathbf{b} –arcs




Figure 11: The alteration in \mathcal{S}_K from left to right corresponds to an exchange move with type-I support. The vertex labeled v_o is an outside vertex and the vertex labeled v_i is an inside vertex. The shaded region is a portion of \mathcal{S}_K .

such that $v_{-\epsilon}^i$ is an endpoint of b^i , $1 \leq i \leq 3$. Assume the i superscript is in correspondence with the angular order of the \mathbf{b} -arcs around the vertex v_{ϵ}^t . Descriptively, the vertex v_{ϵ}^t is at a *tee vertex* of \mathcal{S}_K . The singular leaf that is common to vertices $v_{\epsilon}^t, v_{-\epsilon}^1, v_{-\epsilon}^2$ has its singular point on $\partial\mathcal{S}_K$. This singular point, s^t , is a *tee singularity* of \mathcal{S}_K . For any point of $K \cap b^2$ we can find an arc neighborhood $b \subset \gamma \subset K$ such that either (path 1) γ intersects in order b^1, b^2, b^3 , or (path 2) γ intersects in order b^1, b^2, b' where b' is a \mathbf{b} -arc having $v_{-\epsilon}^2$ as an endpoint but not v_{ϵ} . If only path 1 occurs then $v_{-\epsilon}^2$ is a type of corner vertex of $\partial\mathcal{S}_K$ and we have the situation described in the exchange move with type-I support. So assume that both path 1 and path 2 γ -arcs occur. Since the sequential intersection pattern for path 1 γ -arcs corresponds to that of Figure 10(b), we can perform a sequence of exchange moves to K until v_{ϵ} no longer exists as a tee of \mathcal{S}_K . Figure 12 shows the alternation to \mathcal{S}_K due to such a sequence of exchange moves. We refer to this alternation as an *exchange move with type-II support*.

A few final remarks about exchange moves with types I and II support will be useful. First, the four possible states of a fixed vertex on $\partial\mathcal{S}_K$ are: inside corner; outside corner; tee; and null, ie neither corner nor tee. Second, suppose as we transverse a boundary component $c \subset \partial\mathcal{S}_K$ we encounter sequentially vertices v^1, v^2, v^3 . Assume v^2 is an outside corner. If v^1 (or v^3) has a null state then after an exchange move with type-I support at v^2 , we will have v^1 (or v^3) being an outside corner. If v^1 (or v^3) is an inside corner then after an exchange move with type-I support at v^2 it has a null state. If v^1 (or v^3)




Figure 12: The alteration in \mathcal{S}_K from left to right corresponds to an exchange move with type-II support. The vertex labeled v_ϵ^t is a tee vertex and the singularity labeled s^t is a tee singularity. The shaded region indicates a portion of \mathcal{S}_K .

is a tee then after an exchange move with type-I support at v^2 it has a null state. Notice if K is coherent with respect to $G_{\epsilon,\delta}$ neither v^1 nor v^3 may be an outside corner. Now, assume v_2 is a tee vertex. If v^1 and v^3 states are null then after an exchange move with type-II support at v^2 , one vertex, say v^1 , will be a tee vertex and the other vertex, v^3 , will be an inside corner vertex. If v^1 is an outside corner vertex then after an exchange move with type-II support at v^2 , v^1 state will be null. If there is a tee singularity being v^1 and v^2 then after an exchange move with type-II support at v^2 , v^1 will become a tee vertex. Notice by the assumption that K is coherent with respect to $G_{\epsilon,\delta}$ we cannot have either v^1 or v^3 being an inside corner or tee vertex.

We now consider the possible subsurface types the \mathbf{b} -support might assume.

Lemma 5.6 *Let (K, \mathcal{T}_C) be a knot–torus pair where \mathcal{T}_C has a standard tiling. Assume that K intersects all components of the graphs $G_{\pm,\pm}$ and $G_{\pm,\mp}$ coherently and that the \mathbf{b} -support of K is an annulus in \mathcal{T}_C . If $\partial\mathcal{S}_k$ contains a corner then there exists a (K', \mathcal{T}'_C) that is obtained from (K, \mathcal{T}_C) by a sequence of exchange moves such that $\Lambda(K', \mathcal{T}'_C) < \Lambda(K, \mathcal{T}_C)$.*

Proof First, notice that if $v \subset \partial\mathcal{S}_K$ is an outside or inside corner and the unique \mathbf{b} -arc in \mathcal{S}_K that has v as its endpoint is an outermost essential arc, then, after an exchange move with type-I support at v , this \mathbf{b} -arc will become inessential.

Second, suppose that as we transverse a boundary component $c \subset \partial\mathcal{S}_K$ we encounter sequentially vertices v^1, v^2, \dots, v^l , where v_l is the only corner vertex in our list. Furthermore, suppose that v_1 is adjacent to an outermost essential \mathbf{b} -arc, $b \subset \mathcal{S}_K$. We can then “walk” the corner at v_l up to v_1 : perform an exchange move with type-I support which will result in v_{l-1} being a corner; and iterate this process along c until v_1 is a corner. Finally, we perform an exchange move with type-I support at v_1 to make b inessential.

Removing the inessential \mathbf{b} -arc will produce a new knot–torus pair having decreased complexity. \square

Lemma 5.7 *Let (K, \mathcal{T}_C) be a knot–torus pair where \mathcal{T}_C has a standard tiling. Assume that K intersects all components of the graphs $G_{\pm, \pm}$ and $G_{\pm, \mp}$ coherently and that $c \subset \partial\mathcal{S}_K$ is a scc which bounds a disc $\Delta_c \subset \mathcal{T}_C - \text{int}(\mathcal{S}_K)$. If $|\Delta_c|_v$ is the number of vertices in the interior of Δ_c then there exists an exchange move with special support that will reduce $|\Delta_c|_v$.*

Proof Since c is the union of edge-paths in the graphs $G_{\epsilon, \delta}$ and since the tiling of \mathcal{T}_C is standard (specifically, all vertices are valence four), we know that c must contain either an inside corner or a tee vertex. If c contains an inside corner then after an exchange move with type-I support, the number of vertices contained in the interior of Δ_c will decrease. If c contains a tee vertex then after an exchange move with type-II support, again, the number of vertices contained in the interior of Δ_c will decrease. \square

Proposition 5.8 *Let (K, \mathcal{T}_C) be a knot–torus pair where \mathcal{T}_C has a standard tiling and assume that K coherently intersects all components of the graphs $G_{\pm, \pm}$ and $G_{\pm, \mp}$. Then there exists a knot $K' \subset \mathcal{T}_C$ that is transverse to the foliation of \mathcal{T}_C and is obtained from K through a sequence of exchange moves, with $\Lambda(K', \mathcal{T}'_C) = \Lambda(K, \mathcal{T}_C)$, such that either:*

- (i) $\mathcal{S}_{K'}$ is an annulus with $\partial\mathcal{S}_{K'} = c_1 \cup c_2$ where c_1 is a component of $G_{\epsilon, \delta}$ and c_2 is a component of $G_{-\epsilon, -\delta}$;
- (ii) $\mathcal{S}_{K'}$ is a torus-minus-a-disc with $\partial\mathcal{S}_{K'} = \alpha_1 \cup \beta_1 \cup \alpha_2 \cup \beta_2$ where α_1 and α_2 are arcs in $G_{\epsilon, \delta}$ and β_1 and β_2 are arcs in $G_{\epsilon, -\delta}$. In particular, $\partial\mathcal{S}_K$ does not contain any corner vertices.

Proof Since $\mathcal{S}_K \subset \mathcal{T}_C$ and $\mathcal{T}_C - K$ is an annulus, we can distinguish the cases for \mathcal{S}_K topologically as being either an annulus, an annulus minus discs, or a torus minus discs.

- (i-a) If \mathcal{S}_K is an annulus then, by Lemma 5.6 we have the annulus in (i).
- (i-b) If \mathcal{S}_K is an annulus minus discs let $c \subset \partial\mathcal{S}_K$ be a component that bounds a disc $\Delta_c \subset \mathcal{T}_C - \text{int}(\mathcal{S}_K)$. By Lemma 5.7 we can reduce the number of vertices contained in $\text{int}(\Delta_c)$ iteratively until we eliminate a component of $\partial\mathcal{S}_K$. Thus, an annulus minus discs can be reduced to an annulus. Appealing to Lemma 5.6 again, we reduce to the annulus in (i).
- (ii) If \mathcal{S}_K is a torus minus discs, we can apply the argument above to reduce the number of components of $\partial\mathcal{S}_K$ to one component. If this single boundary component contains a corner then, as in the proof of Lemma 5.6, we can “walk” that corner past any outermost essential **b**-arc in \mathcal{S}_K . So \mathcal{S}_K will be a torus minus a disc as described in (ii). □

5.2 The knot, torus, meridian triple

We now consider the triple (K, \mathcal{T}_C, m) where $m \subset \mathcal{T}_C$ is a meridian curve that intersects K coherently p_h -times. We expand our measure of complexity to be the 3-tuple $\Lambda(K, \mathcal{T}_C, m) = (n_1, n_2, n_3)$ where n_3 is the braid index of m and, as before, n_1 is the braid index of K and n_2 is the number of vertices in the foliation of \mathcal{T}_C .

Proposition 5.9 *Let (K, \mathcal{T}_C, m) be a knot-torus-meridian triple where the pair (K, \mathcal{T}_C) satisfies the conclusion of Proposition 5.5. Then there exists a meridian curve $m' \subset \mathcal{T}_C$ that is transverse to the foliation of \mathcal{T}_C and is obtained from m through a sequence of exchange moves and destabilization such that m' intersects all components of the graphs $G_{\pm, \pm}$ and $G_{\pm, \mp}$ coherently. Moreover, $\Lambda(K, \mathcal{T}_C, m') \leq \Lambda(K, \mathcal{T}_C, m)$.*

Proof We model our proof on the proof of Proposition 5.5. If m does not intersect graph components coherently then, after repeated application of Lemma 5.4, we can assume that m has been replaced by m' satisfying the conclusion of the proposition. Notice that if m is destabilized using the isotopy in Figure 10, this isotopy will occur away from K . Moreover, if m is isotoped through a Figure 10 exchange move then m will still intersect K p_h -times. All operations on m are non-increasing on the complexity measure. □

Let v_+ and v_- be two vertices in the foliation of \mathcal{T}_C which are endpoints of a common **b**-arc. Consider the rectangular region in \mathcal{T}_C that is the closure of the union of all the **b**-arcs that have v_+ and v_- as their endpoints. We shall call such a region a **b**-rectangle.

We introduce the notation of a *parallel push-off* of a meridian curve m . Specifically, we can choose a scc $m' \subset \mathcal{T}_C$ such that: there exists an annulus $\mathcal{A} \subset \mathcal{T}_C$ with $\partial\mathcal{A} = m \cup m'$; the induced foliation on \mathcal{A} has only **aa**-singularities; and each vertex in the foliation of \mathcal{A} is adjacent to two singular leaves. See Figure 13. The oriented arc m is isotopic on \mathcal{T}_C to the oppositely oriented arc m' . We use the notation \mathcal{S}_m for the **b**-support of the meridian curve m .

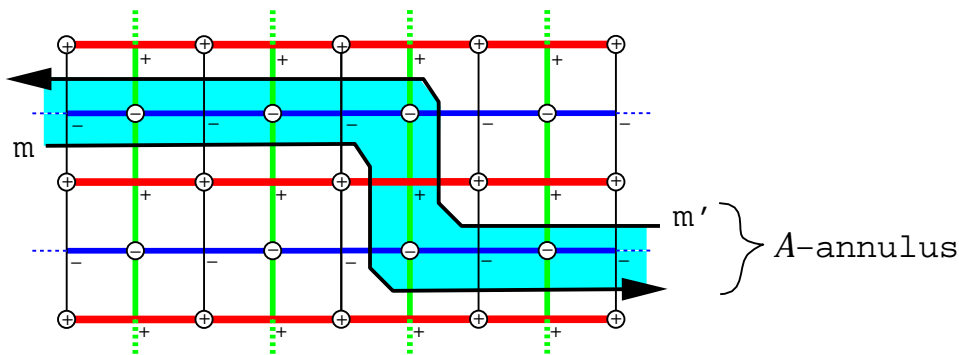


Figure 13: The shaded region between m and m' is \mathcal{A} . Notice that the induced foliation of \mathcal{A} has only **aa**-singularities.

Proposition 5.10 *Let (K, \mathcal{T}_C, m) be a knot-torus-meridian triple satisfying the conclusions of Propositions 5.5, 5.8 and 5.9. Then we can assume that $\mathcal{S}_K \cap \mathcal{S}_m$ is a union of disjoint **b**-rectangles.*

Proof By Lemma 5.6, we know that \mathcal{S}_K is either an annulus or a torus minus a disc. Moreover, we also know by the same lemma that $\partial\mathcal{S}_K$ has no corners.

Using the same argument that was used in the proof of Lemma 5.6, we can assume that each boundary component of \mathcal{S}_m has at most one outside (inside) corner.

If \mathcal{S}_K is an annulus of the type-(ii) in Proposition 5.8 and $\partial\mathcal{S}_m$ has no corners then we immediately have the conclusion of the proposition. If $\partial\mathcal{S}_m$ has any corners then we can “walk” these corners (as in the proof of Lemma 5.6) so that the corners of $\partial\mathcal{S}_m$ are away from $\mathcal{S}_K \cap \mathcal{S}_m$. After this repositioning, $\mathcal{S}_K \cap \mathcal{S}_m$ will be disjoint **b**-rectangles.

It is easily seen that if $\mathcal{S}_K \cup \mathcal{S}_m$ is not a disjoint union of **b**-rectangles then $\mathcal{S}_K \cup \mathcal{S}_{m'}$ will be a disjoint union of **b**-rectangles where m' is a parallel push-off of m . □

5.3 The knot, torus, meridian and meridian-disc

We now consider the quadruple $(K, \mathcal{T}_C, m, \Delta_m)$ where Δ_m is a meridian disc which m bounds inside the solid torus \mathcal{V}_C which \mathcal{T}_C bounds. Recall that the fibration \mathbf{H} induces a foliation on Δ_m . Again, we expand our measure of complexity to be the 4-tuple $\Lambda(K, \mathcal{T}_C, m, \Delta_m) = (n_1, n_2, n_3, n_4)$ where $n_4 = |\mathbf{A} \cap \Delta_m|$ and n_1, n_2 and n_3 are as before. If $n_4 = 0$ then the foliation of \mathcal{T}_C contains \mathbf{c} -circles and we can appeal to the arguments in Section 4.

Before we proceed further it will be useful to review the induced foliation on Δ_m . Again, a comprehensive reference is [1]. When $n_4 = 1$, Δ_m is *radially* foliated by \mathbf{a} -arcs adjacent to the unique vertex. When $n_4 > 1$, the singular foliation of Δ contains the vertices $\Delta_m \cap \mathbf{A}$ and saddle singularities of three possible types— \mathbf{aa} -, \mathbf{ab} - and \mathbf{bb} -singularities. Thus, Δ_m is tiled by \mathbf{aa} , \mathbf{ab} and \mathbf{bb} tiles. Since m induces on Δ_m an orientation, we can assign a parity to each vertex and singularity using the same assignment scheme employed in the foliation of \mathcal{T}_C . Having a parity assignment on vertices and singularities allows us to define graphs $G_{\pm, \pm}$ and $G_{\pm, \mp}$ as before with the proviso that we treat m as if it were a negative vertex. (This implies that $G_{-, -}$ and $G_{-, +}$ will contain edges having endpoints on m .)

Now consider how the foliations of Δ_m and $\Delta_{m'}$ are related to one another, when m' is a parallel push-off of m . Let $\mathcal{A} \subset \mathcal{T}_C$ be the annulus in Figure 13 that has m and m' as its boundary curves. Notice that the core circle in \mathcal{A} that is contained in a union of all of the singular leaves in the induced foliation of \mathcal{A} is a union of sub-arcs of the graphs $G_{\epsilon, \delta} \subset \mathcal{T}_C$. Similar in flavor to the arguments in Proposition 5.8, we can produce a sequence of exchange moves by applying Lemma 5.3 so that either $\partial\mathcal{S}_m$ contains no corners, or each component of $\partial\mathcal{S}_m$ contains exactly one inside corner and one outside corner. If $\partial\mathcal{S}_m$ contains no corners then the foliation of \mathcal{A} will have singularities all of the same parity. If each component of $\partial\mathcal{S}_m$ has exactly one inside and outside corner then the core circle in \mathcal{A} is the union of an arc in $G_{\epsilon, +}$ and an arc in $G_{\epsilon, -}$.

Now the foliation of $\Delta_{m'}$ is simply the foliation of $\Delta_m \cup_m \mathcal{A}$. Specifically, for each singularity in $s \in \mathcal{A} \subset \mathcal{T}_C$ there is a singularity $t \in G_{+, \pm} \subset \Delta_{m'}$ (coming from the “ \mathcal{A} portion” of Δ_m) that is in a regular neighborhood of s on the negative side of the oriented \mathcal{T}_C . It is easily checked by appealing to the H_t -sequence of $\mathcal{T}_C \cup \Delta_m \subset \mathbf{H}$ that if s is a positive (negative) singularity then t is the singularity that occurs immediately before (after) s . (See Figure 14.) If $\partial\mathcal{S}_m$ contains no corners then the $G_{\epsilon, \delta}$ of $\Delta_{m'}$ contains a loop—the core circle coming from \mathcal{A} . But by Lemma 3.8 of [1] such a loop cannot exist. So we

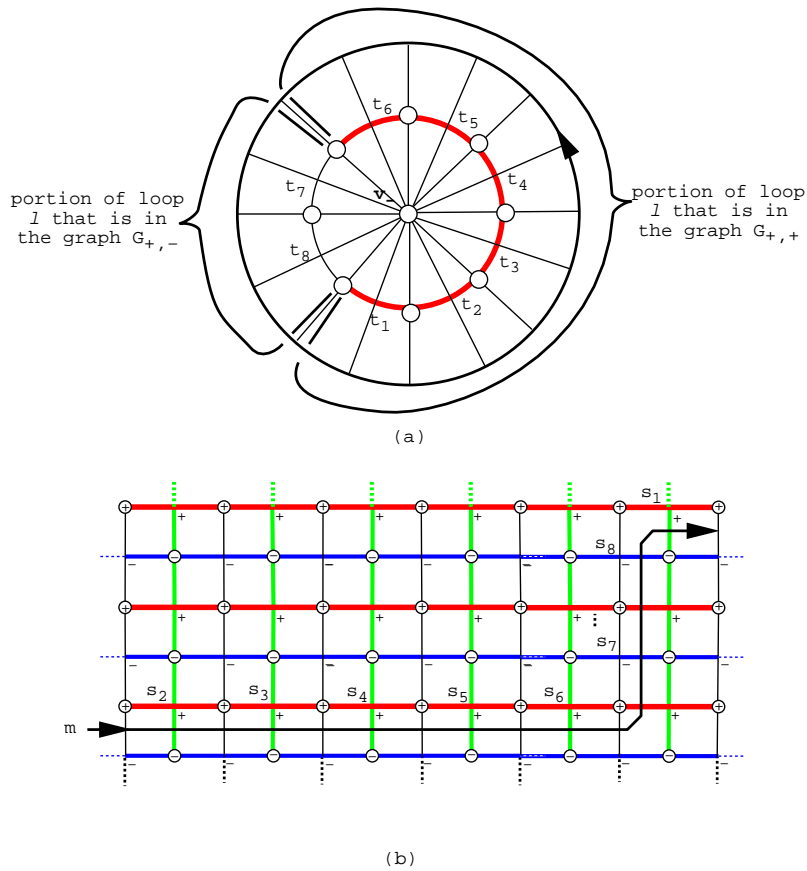


Figure 14: In (a) we have the foliation of Δ_m illustrated in the special case where l contains two negative and six positive singularities, and there is a single negative vertex. In (b) we have the corresponding path m in a standard tiling of \mathcal{T}_C . The arrowheads are identified to form a scc. The labels $\{t_1, \dots, t_8\}$ of the singularities in (a) show the sequential correspondence to the singularities in (b) with the labels $\{s_1, \dots, s_8\}$.

are left having $\Delta_{m'}$ containing a loop (which is the core circle of $\mathcal{A} \subset \Delta_{m'}$) that is the union of an arc in the graph $G_{+,+} \subset \Delta_{m'}$ and an arc in the graph $G_{-,-} \subset \Delta_{m'}$. This loop contains all of the positive vertices of $\Delta_{m'}$. As noted above, a singularity $s \in \partial\mathcal{S}_m$ will in \mathbf{H} sequentially correspond to a singularity $t \in l$. An application of the argument in Lemma 3.8 (ii) allows us to conclude this discussion of the foliation of $\Delta_{m'}$ with the following summarizing result.

Lemma 5.11 *Let $(K, \mathcal{T}_C, m, \Delta_m)$ be a quadruple where the triple (K, \mathcal{T}_C, m)*

satisfies the conclusions of Propositions 5.5, 5.8, 5.9 and 5.10. We can assume that the initial foliation of Δ_m satisfies the following conditions.

- (a) There are only **ab**-singularities.
- (b) There exists a closed loop l that is the union of an arc in $G_{+,+} \subset \Delta_m$ and an arc in $G_{+,-} \subset \Delta_m$.
- (c) When $n_4 = 3$, there is a single negative vertex v_- . When $n_4 > 3$, there is a negative vertex v_- adjacent to two singularities of common parity that are consecutive on l .

Moreover, the singularities contained in each component of $\partial\mathcal{S}_m$ are in one-to-one correspondence with, and in \mathbf{H} sequentially correspond to, the singularities contained in l . (See Figure 14.)

We now use this initial tiling of Δ_m to reduce the complexity of the quadruple $(K, \mathcal{T}_C, m, \Delta_m)$. Notice that the number of singularities in the loop $l \subset \Delta_m$ is naturally two or greater. The next lemma shows us that when l contains exactly two singularities the foliation of a tiled \mathcal{T}_C must have an inessential **b**-arc.

Lemma 5.12 *Let $(K, \mathcal{T}_C, m, \Delta_m)$ be a quadruple, where the triple (K, \mathcal{T}_C, m) satisfies the conclusions of Propositions 5.5, 5.8, 5.9, 5.10 and 5.11. If $n_4 = 3$ then the foliation of \mathcal{T}_C has an inessential **b**-arc.*

Proof If $n_4 = 3$ then l must have contained one positive and one negative singularity. Thus, referring to the terminology of [4], \mathcal{T}_C must be a type **k** torus. A type **k** torus embedding has the property that a meridian curve on \mathcal{T}_C can be represented by the union of two **b**-arcs, each being outermost in a disc fiber of \mathbf{H} . Since K can only intersect one of these two **b**-arcs, \mathcal{T}_C must contain an inessential **b**-arc. (Refer back to Remark 3.2 and Figure 7 for an example of a type **k** torus.) □

We next deal with the case where l contains more than two singularities.

Lemma 5.13 *Let $(K, \mathcal{T}_C, m, \Delta_m)$ be a quadruple where the triple (K, \mathcal{T}_C, m) satisfies the conclusions of Propositions 5.5, 5.8, 5.9, 5.10 and 5.11. If $n_4 > 3$ then there exists an isotopy of \mathcal{T}_C such that n_2 is reduced.*

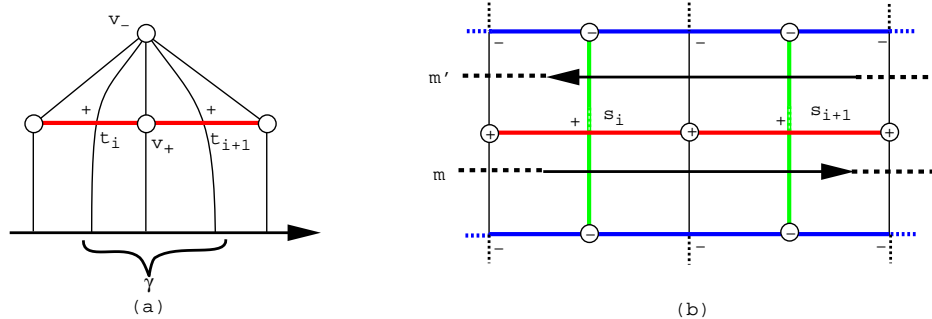


Figure 15

Proof If $n_4 > 3$ then we employ the negative vertex v_- in condition (c) of Lemma 5.11. Figure 15(a) illustrates this situation when the parity of these two singularities is positive. In Figure 15(a) these singularities are labeled t_i and t_{i+1} . The common positive vertex that is adjacent to both of their associated singular leaves is v_+ . A regular neighborhood of the portion of m that contains endpoints of \mathbf{a} -arcs adjacent to v_+ is labeled γ . If we look at the image of γ in the foliation of \mathcal{T}_C we see that γ intersects a positive singular leaf containing the singular point s_i associated to t_i ; then a negative singular leaf; then another positive singular leaf containing the singular point s_{i+1} associated to t_{i+1} . In the H_t -sequence of $\Delta_m \cup \mathcal{T}_C$ we know by the discussion in the proof of Lemma 5.11 that t_i occurs immediately before s_i and t_{i+1} occurs immediately before s_{i+1} .

By Proposition 5.10, there are now three possibilities for the arc γ : (1) $\gamma \cap K = \emptyset$; (2) γ intersects the singular leaf belonging to s_i first, then the negative singular leaf and then K ; or (3) γ intersects the singular leaf belonging to s_i first, then intersects K some number of times, and then intersects the negative singular leaf.

If (3) occurs we replace m with m' , altering $\Delta_{m'}$ so that its foliation again satisfies the initial condition foliation of Lemma 5.11. This produces a new γ which corresponds to possibility (2). Now notice in the situation of (2) we can isotopy K in \mathcal{T}_C so that K intersects m between the time m crosses the t_{i+1} singular leaf and the time it crosses the s_{i+1} singular leaf. Thus, we can assume that $\gamma \cap K = \emptyset$.

Since t_i and t_{i+1} are of the same parity and adjacent to the unique negative vertex v_- in Δ_m , we can perform a change of foliation. (This is an application to the foliation of Δ_m of the alteration in Figure 5.) Figure 16(a) show how this change of foliation results in a type (a) vertex, $v_+ \subset \Delta_m$, which locally

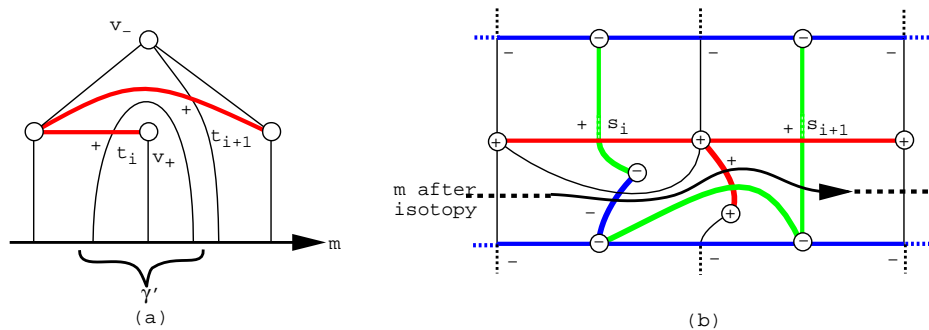


Figure 16

has the same foliation as that in Figure 3. The boundary of this valence 1 vertex contains $\gamma' \subset m$. (See Figure 16(a).) We now destabilize Δ_m through this valence 1 vertex and, since $\gamma' \cap K = \emptyset$, we can drag \mathcal{T}_C along through this destabilization of m without altering the embedding of K . Figure 16(b) illustrates how the foliation of \mathcal{T}_C is altered by this destabilization of m . Notice that two new vertices are introduced into the foliation of \mathcal{T}_C , one positive and one negative, and that both of these new vertices are valence 2. We can now appeal to Figure 6 to eliminate four vertices of \mathcal{T}_C . This reduces n_2 and, thus, the complexity of our quadruple.

Finally, it should be noted that we have made a choice of the parity of $\{t_i, t_{i+1}, s_i, s_{i+1}\}$ for reasons having to do with the clarity of the expository. This does not reduce the generality of the argument. \square

We have proved Proposition 5.1 and, thus, Theorem 1.1.

References

- [1] **J Birman, E Finkelstein**, *Studying surfaces via closed braids*, Journal of Knot Theory and its Ramifications, 7 (1998) 267–334
- [2] **J Birman, W Menasco**, *Studying Links via Closed Braids V: Closed braid representations of the unlink*, Transactions of the AMS, 329 (1992) 585–606
- [3] **J Birman, W Menasco**, *Studying Links via Closed Braids IV: Split links and composite links*, Inventiones Mathematicae, 102 (1990) 115–139
- [4] **J Birman, W Menasco**, *Special positions for essential tori in link complements*, Topology, 33 (1994) 525–556
- [5] **J Birman, W Menasco**, *Studying Links via Closed Braids IV: Closed links which are 3–braids*, Pacific Journal of Mathematics, 161 (1993) 25–113

- [6] **J Birman, N C Wrinkle**, *On transversally simple knots*, preprint (1999)
- [7] **Y Eliashberg**, *Legendrian and transversal knots in tight contact 3-manifolds*, *Topological Methods in Modern Mathematics*, (1991) 171–193
- [8] **J Etnyre**, *Transversal torus knots*, *Geometry and Topology*, 3 (1999) 253–268
- [9] **W Jaco**, *Lectures on Three-Manifold Topology*, AMS Regional conference series, No. 43
- [10] **J Los**, *Knots, braid index and dynamical type*, *Topology*, 33 (1994) 257–270
- [11] **D Rolfsen**, *Knots and Links*, Mathematics Lecture Series 7, Publish or Perish, Inc. (1976)
- [12] **H Schubert**, *Knoten und Vollringe*, *Acta Math.* 90 (1953) 131–226