Canad. J. Math. Vol. **68** (6), 2016 pp. 1201–1226
http://dx.doi.org/10.4153/CJM-2016-002-1

Monodromy Action on Unknotting Tunnels in Fiber Surfaces

Jessica Banks and Matt Rathbun

Abstract. In a 2012 paper, the second author showed that a tunnel of a tunnel number one, fibered link in S^3 can be isotoped to lie as a properly embedded arc in the fiber surface of the link. In this paper we observe that this is true for fibered links in any 3-manifold, we analyze how the arc behaves under the monodromy action, and we show that the tunnel arc is nearly clean, with the possible exception of twisting around the boundary of the fiber.

1 Introduction

The Berge Conjecture is a long-standing conjecture that attempts to classify all knots in S^3 that admit Dehn surgeries resulting in a lens space. Such a classification is foundational to understanding Dehn surgery on 3-manifolds and has been a motivating topic of research in low dimensional topology for decades. The so-called Berge knots are conjectured to be all knots admitting such surgeries and are known to be both *tunnel number one* and *fibered*. Y. Ni [26] also proved that any knot admitting such a surgery must be fibered. In light of this, we aim to understand tunnel number one, fibered knots and links.

In Section 2, we will define three well-understood operations on fibered links: Stallings twisting, Hopf plumbing, and its inverse Hopf de-plumbing. All three of these operations can be characterized by arcs that are *clean*, *i.e.,* disjoint from their images under the monodromy map (except at their endpoints).

Our goal in this paper is to understand how the monodromy acts on tunnels sitting as arcs in the fiber. We show that such tunnels sit *weakly* cleanly in the fiber. We prove the following theorem.

Theorem 1.1 *Let F be a compact, connected, orientable surface with one or more boundary components and let $h\colon F \to F$ be an orientation-preserving homeomorphism. Let $M = (F \times I)/h$, and denote by F the surface $F \times \{0\}$ in M. Let τ be an arc properly embedded in F such that $M \smallsetminus N(\tau)$ is a (genus two) handlebody, where $N(\tau)$ is a regular neighborhood of τ in M. Then there is an arc that is freely ambient isotopic in F to $h(\tau)$ and is disjoint from τ.*

Received by the editors June 25, 2014; revised September 16, 2015.
Published electronically June 23, 2016.
AMS subject classification: **57M25**.
Keywords: fibered, monodromy, tunnel, clean.

We then obtain the following theorem about link exteriors in 3-manifolds as a corollary.

Theorem 1.2 *Suppose K is a tunnel number one, fibered link in a 3-manifold M, with fiber F, monodromy h, and a properly embedded arc τ in F that is an unknotting tunnel for K. Then there exists a properly embedded arc $\beta \subset F$, freely ambient isotopic in F to $h(\tau)$, so that $\tau \cap \beta = \varnothing$. In particular, up to isotopy* rel ∂F, *there exists a regular neighborhood of ∂F outside of which τ and $h(\tau)$ do not intersect.*

Johnson [20] investigated closed surface bundles with genus two Heegaard splittings. Johnson's work gives a description of the monodromy of a fibered, tunnel number one knot, but it does not tell us about the case of a two-component link. Kobayashi and Johannson independently proved that for once-punctured torus bundles, an unknotting tunnel could be isotoped into a fiber so that the arc is disjoint from its image under the monodromy of the bundle (see [31]). According to a survey article by Sakuma [32], Kobayashi and Johannson also independently proved the same result for arbitrary punctured surface bundles. However, both references are talks, and it is unclear what the relevant restrictions or equivalence classes on the monodromy map and/or the arcs are meant to be. This paper is meant to help clarify some of the various technical distinctions, particularly between surface bundles and link exteriors, and provide a written proof of the proper result. In Section 4 we will discuss examples of tunnel number one, fibered links in S^3 with tunnels α that are properly embedded in a fiber F, but are not disjoint from their images under the monodromy unless we allow the free-isotopy mentioned in Theorem 1.2.

This paper is organized as follows. Section 2 details definitions, background, and motivation for the statement and proof of the main theorem, found in Section 3. Section 4 discusses limitations of the theorem owing to difficulties associated with (fractional) Dehn twists around the boundary of the fiber surface. And finally, Section 5 provides an application to bounding the cusp area for hyperbolic, fibered knots.

2 Definitions and Background

Definition 2.1 A manifold M with boundary is said to have *tunnel number one* if there exists an arc τ (an *unknotting tunnel*) properly embedded in the manifold so that $M \smallsetminus n(\tau)$ is a handlebody. We say that a link K is *tunnel number one* if the link exterior has such an unknotting tunnel.

A tunnel number one link can therefore have at most two link components, and in this case, the tunnel must have one endpoint on each component. Tunnel number one knots and links have been studied in great depth (see, for example, [17,19,24,33]). Cho and McCullough [6] have given a bijective correspondence between tunnel number one knots (with their unknotting tunnels) and a subset of vertices of a certain tree related to a subcomplex of non-separating disks in a genus two handlebody. They are further able to parameterize all tunnel number one knots by a sequence of "cabling" operations (see [5,7]). While the cabling operation is a very natural way of describing and modifying knots, it is generally not clear how properties of the exterior change.

Definition 2.2 Let K be a link in a 3-manifold M (with an orientation for each component). A *Seifert surface* for K is a compact, orientable surface F, with no closed components, embedded in M such that $\partial F = K$ (and with a boundary orientation that agrees with the orientation of K).

Definition 2.3 If F is a compact, orientable surface (possibly with boundary), I is the unit interval $[0,1]$, and $h: F \to F$ is an orientation-preserving homeomorphism from F to itself, then a *surface bundle* is the 3-manifold obtained from the Cartesian product of the surface and the interval, $F \times I$, by identifying $F \times \{0\}$ with $F \times \{1\}$ via the homeomorphism h. That is, the surface bundle is homeomorphic to the quotient $(F \times I)/\sim$, where $(x, 0) \sim (h(x), 1)$ for all $x \in F$. We may also denote this by $(F \times I)/h$. The map h is called the *monodromy* of the bundle, and the image of each $F \times \{t\}$ is called a *fiber*.

Note 2.4 The monodromy h of a surface bundle is well defined up to free isotopy of the homeomorphism (preserving the boundary, set-wise), and also up to conjugation by elements of the mapping class group of F.

Definition 2.5 A link $K \subset M$ is said to be *fibered* if $M \setminus n(K)$ is a surface bundle where each fiber is a Seifert surface for the link K.

If we drill out a neighborhood of a fibered link from the manifold M, then there is a natural marking on each boundary component by a meridian that encodes the original manifold M. If we then remove a neighborhood of a Seifert surface, the result is homeomorphic to $F \times I$, but we still retain a marking on the boundary $(\partial F) \times I \subset \partial(F \times I)$, which still encodes the original manifold M. Forgetting this marking or losing track of it by twisting along boundary components, however, fails to encode the manifold M. This motivates a slightly more restrictive definition when we are interested in preserving this meridian information.

Definition 2.6 The *monodromy* of a fibered link K in M is a choice of homeomorphism $h: F \to F$ so that h is the identity on the boundary of F, $h|_{\partial F} = \mathrm{Id}$, the exterior of K is homeomorphic to the surface bundle determined by h, $M \setminus n(K) \cong (F \times I)/h$, and further, filling each toral boundary component with a solid torus so that the curve arising from the quotient of $\{\mathrm{pt}\} \times I$ bounds a disk in the solid torus results in the manifold M.

Note 2.7 If $\widetilde{h}$ differs from h by a product of Dehn twists about curves each parallel to a component of ∂F, then $(F \times I)/h \cong (F \times I)/\widetilde{h}$. Now, however, Dehn filling the toral boundary along the curve(s) defined by $\{x\} \times I$ where $x \in \partial F$ in each case may result in different closed 3-manifolds, related to the original by $\pm(1/n)$-surgeries. So the requirement that the loops from "vertical" slopes give rise to meridians for the link K in M restricts the free isotopy class of h. The monodromy h of a fibered link is still only well defined up to conjugation by an element of the mapping class group of F.

Fibered knots and links have also been studied in great depth (see, for example, [1,2,18,25]). Stallings [35] described a pair of operations on fibered links that result in new fibered links, which are now called the *Murasugi sum* and *Stallings twists*. Harer [18] then showed that twists and a certain type of Murasugi sum called *Hopf plumbing* (and its inverse, *Hopf de-plumbing*) were sufficient to transform any fibered link in S^3 into any other fibered link in S^3. (In fact, recent work of Giroux and Goodman [15] showed that Stallings twists are not necessary.)

These constructions are intimately connected to arcs in a fiber surface with certain properties of disjointness from their images under monodromy maps. However, we will take care to distinguish monodromy maps in surface bundles versus link complements, so we must be cautious about the setting in which we are discussing these arcs.

Definition 2.8 We will say that an arc α properly embedded in a fiber F of a surface bundle with monodromy h is *weakly clean* if there is a representative of h, say $\widetilde{h}$, so that $\alpha \cap \widetilde{h}(\alpha) = \varnothing$.

Definition 2.9 An arc α properly embedded in a fiber F of a link complement in a manifold M with monodromy h is said to be *clean* if there is a representative of h, say $\widetilde{h}$, so that $\alpha \cap \widetilde{h}(\alpha) = \partial\alpha = \partial\widetilde{h}(\alpha)$.

In this language, Theorem 1.1 says that unknotting tunnels in the fibers of surface bundles are weakly clean, while Theorem 1.2 and the discussion in Section 4 show that unknotting tunnels in the fibers of fibered link exteriors may not be clean, owing only to boundary-twisting.

There is good reason to inquire about the cleanliness of arcs in the fiber of a fibered link. Suppose α is a clean arc in a fiber of a fibered link with monodromy h. There are two distinct behaviors of $h(\alpha)$ near the boundary of α, each of which have implications for the topology of the fiber surface.

Definition 2.10 Let α be a clean arc in a fiber F of a fibered link, and let $\alpha \times [0,1]$ be a small product neighborhood of the arc α in F. We say that α is *alternating* if the image of α under the monodromy must intersect both $\alpha \times \{0\}$ and $\alpha \times \{1\}$ in a neighborhood of the endpoints. Otherwise, say that α is *non-alternating*.

Clean, alternating arcs are related to Hopf plumbings.

Definition 2.11 Let F be a Seifert surface for a link L. Let α be an arc properly embedded in F. *Hopf plumbing along* α is a change in the surface F within a neighborhood of the arc α, as shown in Figure 1. That is, a disk is attached to F along two sub-arcs of its boundary. The positioning of the disk is defined by α, and the disk contains a full twist relative to F. This disk is referred to as a *Hopf band*. Given F and α, there are two ways to perform Hopf plumbing, distinguished by the handedness of this twisting. The result is a new surface F' and a new link $K' = \partial F'$.

Suppose F is a Seifert surface for the link ∂F, and Hopf plumbing results in a Seifert surface F' for the link $\partial F'$. Then F is a fiber surface if and only if F' is a fiber surface,

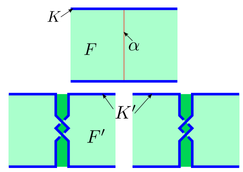

Figure 1: Hopf plumbing is a change in a surface F in the neighborhood of an arc α.

and, moreover, the monodromy of F' differs from the monodromy of F exactly by composition with a Dehn twist along the core of the Hopf band (see [11]).

De-plumbing a Hopf band corresponds exactly to cutting the fiber surface along an arc that is clean and alternating with respect to the monodromy. This is implicit in work of Gabai [12], and attributed to Sakuma [30]. For a proof, see Coward–Lackenby [8].

Clean, non-alternating arcs are related to Stallings twists.

Definition 2.12 Let c be a simple closed curve, embedded and essential in a fiber surface F of a fibered link in a manifold M so that c bounds a disk in M. Let c' be a push-off of c to one side of F, and let l be the linking number of c and c'. If $l \in \{0, 2, -2\}$, and there exist $\delta_1, \delta_2 \in \{\pm 1\}$ satisfying $l + \delta_1 = \delta_2$, then δ_2-surgery along c is called a *Stallings twist*.

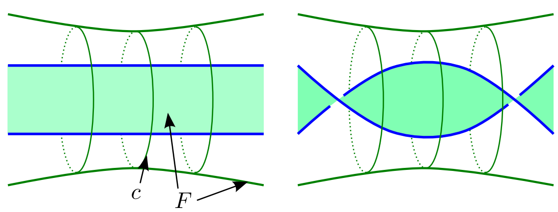

Figure 2: A Stallings twist (of type $(0,1)$) results from a ± 1-Dehn surgery on an unknotted curve in the fiber surface. (Here we show the effect of a -1-surgery on the surface.)

In the case where $l = 0$, Stallings [35] proved that the image of a fibered link under such a twist is another fibered link, with fiber surface homeomorphic to the original fiber surface. Harer [18] then extended this to the definition above . Moreover, the monodromy of the new fibration differs from the original exactly by composition with a δ_1-Dehn twist around the curve c.

Yamamoto [38] proved that the existence of a Stallings twist of a certain type (type $(0,1)$, see Figure 2) corresponds exactly to the existence of an arc that is clean and non-alternating with respect to the monodromy, and moreover that the interior of the disk bounded by c intersects the fiber surface exactly in such an arc.

Not only do these operations have close relationships to the behavior of arcs in a fiber surface, but when the arcs are unknotting tunnels, these operations also respect the nature of these tunnels. In [29], the second author showed that in a tunnel number one, fibered link exterior in S^3, an unknotting tunnel can be isotoped to lie in a fiber surface. In fact, the argument in [29] shows that this statement holds for fibered links in arbitrary 3-manifolds.

Proposition 2.13 *Suppose K is an oriented fibered link in a 3-manifold M, and τ is an unknotting tunnel for K. Then τ may be slid and isotoped until it lies in a fiber of K.*

Proof The proof in [29] for two-component fibered links depends in no way on the ambient manifold being S^3. For tunnel number one, fibered knots, the proof relies on S^3 only because of the use of [34, Proposition 4.2]. However, the argument in the final paragraph of [29, Theorem 2.10] applies equally well in the case where K is a knot and the ambient manifold is not S^3 to show that an unknotting tunnel can be made disjoint from a minimal genus Seifert surface when the knot is fibered. The rest of the argument, then, goes through in the more general case. ■

Thus, for any tunnel number one, fibered link, since an unknotting tunnel can be isotoped to lie as an arc in a fiber surface, and important operations on fiber surfaces are related to arcs properly embedded in the surface, it is a natural question to investigate what happens if we perform these operations along arcs that happen to be unknotting tunnels.

If a Hopf plumbing is performed along an unknotting tunnel lying in the fiber surface, then the resulting link is fibered and is again tunnel number one. Moreover, there is a naturally induced unknotting tunnel for the resulting link, namely the arc that runs across the Hopf band. Conversely, if de-plumbing a Hopf band corresponds to cutting along an arc that is also an unknotting tunnel lying in the fiber surface, then the resulting link is fibered and is again tunnel number one. Moreover, there is a naturally induced unknotting tunnel for the resulting link, namely the arc that spans the gap left by the cut, which is then isotopic into the new fiber surface as the arc that determined the position of the old Hopf band. One result of this is that we can start with, say, the unknot in S^3, which is fibered with fiber a disk, and progressively plumb Hopf bands along unknotting tunnels to generate tunnel number one, fibered links with fiber surfaces of increasing genus. Another is that if an unknotting tunnel, having been pushed into a fiber surface for a fibered link, is a clean, alternating arc, then the fiber surface has a de-plumbing resulting in a new tunnel number one, fibered link. We might then ask whether the next unknotting tunnel, having been pushed in the fiber surface, might also be clean and alternating, and how far this process might be continued. If, for instance, a tunnel were always clean and alternating when pushed into a fiber surface, then any tunnel number one, fibered link would come equipped with a set of instructions indicating a sequence of tunnel number one, fibered links, each obtained from the last by de-plumbing along an unknotting tunnel, resulting in a fibered link with fiber a disk.

Similarly, if a Stallings twist is performed along a curve bounding a disk that intersects a fiber surface in an arc that is also an unknotting tunnel, then the resulting link is fibered and again tunnel number one. In this case, the very same arc will persist as an unknotting tunnel.

In light of these close connections between unknotting tunnels that sit as clean arcs in a fiber surface and operations that are known to be sufficient to generate all fibered links in S^3, it would be quite interesting, although optimistic, to suspect that unknotting tunnels sitting as arcs in fiber surfaces would always be clean. However, our main result shows that the obstruction to cleanliness comes only from twisting around boundary components of the fiber. Even more surprising, in S^3, the only examples known with this obstruction appear to be 2-component links where one component is unknotted, as we will discuss in Section 4.

3 Analyzing a Tunnel in a Fiber

The main aim of this section is to prove Theorem 1.1, and Theorem 1.2 will follow quickly. We restate the first theorem here.

Theorem 1.1 *Let F be a compact, connected, orientable surface with one or more boundary components and let $h\colon F \to F$ be an orientation-preserving homeomorphism. Let $M = (F \times I)/h$, and denote by F the surface $F \times \{0\}$ in M. Let τ be an arc properly embedded in F such that $M \setminus N(\tau)$ is a (genus two) handlebody, where $N(\tau)$ is a regular neighborhood of τ in M. Then there is an arc that is freely ambient isotopic in F to $h(\tau)$ and is disjoint from τ.*

Note 3.1 Observe that since τ is an unknotting tunnel for the manifold M, M can have at most two (toral) boundary components. Thus, the boundary components of F must be permuted in one or two orbits.

Proof If F is either a disk or an annulus, the mapping class group of F is quite limited, and the result is immediate. Thus, we can assume that F is neither a disk nor an annulus. Note that since the only fibered handlebody is $S^1 \times D^2$, this implies that M is not a handlebody.

As M is not a handlebody but $M \setminus n(\tau)$ is, the arc τ must be essential in F. Set $F' = F \setminus n(\tau)$. Let τ_1 be an arc in $F \times \{1\} \subset F \times I$ and let τ_0 be an arc in $F \times \{0\} \subset F \times I$ so that, in the quotient $(F \times I)/h$, arcs τ_0 and τ_1 are both identified as the arc τ. Observe that $h(\pi(\tau_0)) = \pi(\tau_1)$, where $\pi\colon F \times I \to F$ is projection. Then for $i \in \{0,1\}$, we will refer to $(F \times \{i\}) \setminus n(\tau_i)$, contained in $(F \times I) = \overline{M \setminus F}$, as F'_i. By free isotopy of h (which corresponds to isotopy of τ_1 in $F \times \{1\}$), we can assume that $\pi(\tau_0)$ and $\pi(\tau_1)$ intersect minimally and transversely. Recall that $F \times I$ is irreducible and $F \times \{0,1\}$ is incompressible in $F \times I$.

Let A be the annulus $\overline{\partial n(\tau) \setminus \partial M}$. Then A is divided into two rectangles by F. Let A_1 be the rectangle incident to $F \times \{1\}$, and let A_0 be the rectangle incident to $F \times \{0\}$ in $\overline{M \setminus F}$. By a slight abuse of notation, we can think of A_1 as a neighborhood of τ_1 contained in $F \times \{1\}$, and similarly for $A_0 \subset F \times \{0\}$, so that $F'_i = (F \times \{i\}) \setminus A_i$ for each $i = 0, 1$.

The proof of Theorem 1.1 works by controlling certain disks within $M \smallsetminus n(\tau)$, in particular how they relate to the annulus A. We now build up some language to describe these disks.

3.1 Special Arcs

Let D be a disk properly embedded in $F \times I$ such that ∂D is transverse to $\partial F \times \{0,1\}$ and to τ_0 and τ_1.

Lemma 3.2 *No essential disk in $F \times I$ can be disjoint from $F \times \{i\}$ for $i \in \{0,1\}$.*

Proof Without loss of generality, suppose that D is an essential disk in $F \times I$ that is disjoint from $F \times \{1\}$. Then every arc in $\partial D \cap (\partial F \times I)$ is inessential in $\partial F \times I$. On the other hand, any simple closed curve in $\partial F \times I$ is either trivial or parallel to a component of $\partial F \times \{0\}$. We can therefore isotope ∂D into $F \times \{0\}$. This contradicts that $F \times \{0,1\}$ is incompressible in $F \times I$. Thus, no such disk exists. ∎

Definition 3.3 If $\partial D \cap (\partial F \times \{0,1\}) \neq \varnothing$, then the points of $\partial D \cap (\partial F \times \{0,1\})$ divide ∂D into a finite set of sub-arcs of the following six possible types.

(i) Sub-arcs in $F \times \{0\}$ parallel in F to τ_0; call these *τ_0-arcs*.
(ii) Sub-arcs in $F \times \{1\}$ parallel in F to τ_1; call these *τ_1-arcs*.
(iii) Sub-arcs in $\partial F \times I$; call these *boundary arcs*.
(iv) Sub-arcs in $F \times \{0\}$ or $F \times \{1\}$ that are trivial in F; call these *extra arcs*.
(v) Sub-arcs in $F \times \{i\}$ for $i \in \{0,1\}$ that are essential in F, are not τ_i-arcs, and can be isotoped (fixing endpoints) to be disjoint from τ_i; call these *special arcs*.
(vi) Sub-arcs in $F \times \{i\}$ for $i \in \{0,1\}$ that are essential in F, are not τ_i-arcs, and necessarily intersect τ_i; call these *bad arcs*.

For $i \in \{0,1\}$, label each sub-arc of ∂D with i if it is contained in $F \times \{i\}$.

We will show in Lemma 3.8 that the disks of interest to us do not contain bad arcs.

Definition 3.4 An extra arc that is outermost in $F \times \{i\}$ can be isotoped off $F \times \{i\}$, along the subdisk it cuts off from $F \times \{i\}$, joining two sub-arcs on $\partial F \times I$ into a single boundary arc. Call this a *tightening-move*. Notice that this does not affect the isotopy type of any essential arc in $F \times \{0,1\}$, and has the effect of deleting an i-label from the labeling of ∂D.

If τ is incident to two boundary components of F, the following definition gives two isotopy classes of arcs in F (hence four isotopy classes in $\partial(F \times I)$) that will be of special interest to us. These arcs are boundary-parallel in F' and have both endpoints on the same component of ∂F. The two isotopy classes are distinguished by which component of ∂F contains the endpoints of the arc.

Definition 3.5 Call a special arc a *τ_2-arc* if it is parallel in F to the union of the two arcs in $\partial A_i \smallsetminus \partial F$ and one of the two components of $\partial F \smallsetminus A_i$. See Figure 3. Roughly speaking, it runs parallel to τ_i, around ∂F while avoiding τ_i, and then back parallel to τ_i.

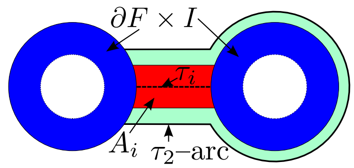

Figure 3: A τ_2-arc runs parallel to τ_i, around ∂F, and back parallel to τ_i.

It is interesting to note that if a τ_2-arc α exists in ∂D, pushing part of α across the disk of F' cut off by α into the component of $\partial F \times I$ that does not contain the endpoints of α would change the special arc α into two τ_i-arcs and one boundary arc.

The significance of τ_2-arcs is their appearance in following lemma.

Lemma 3.6 *If ∂D contains exactly one special arc and no bad arcs, then either D is essential in $F \times I$ or τ is incident to two boundary components of F and the special arc is a τ_2-arc.*

Proof Without loss of generality, assume the special arc α is labeled 1. Perform as many tightening-moves as possible to remove all extra arcs. This neither creates any new special or bad arcs, nor alters α. Then $\partial D \cap (F \times \{1\})$ consists of α together with some number of τ_1-arcs.

Suppose D is boundary parallel in $F \times I$, and let D' be the disk in $\partial(F \times I)$ to which D is parallel. Consider the two components of $\overline{(F \times \{1\}) \setminus \partial D}$ adjacent to α, one of which is a subsurface of D'. Call this D''. Note that D'' is planar, and all but one of the components of $\partial D''$ are contained in $\operatorname{int}(D')$ and are therefore components of ∂F. Any such components of ∂F must also bound disks in $\partial(F \times I)$. Since there are no such components of ∂F, we see that D'' is a disk. There can be at most two τ_1-arcs in $\partial D''$, and exactly one copy of α.

If $\partial D''$ contained no τ_1-arcs, then D'' would provide an isotopy of α into $\partial(F \times \{1\})$, which is not possible. If $\partial D''$ contained exactly one τ_1-arc, then D'' would provide an isotopy of α onto τ_1, which is also impossible. Therefore, $\partial D''$ contains two τ_1-arcs. Notice that $\partial D'' \setminus \alpha$ is contained in $\partial F_1'$. Suppose τ (and therefore τ_1) is incident to a single component of $\partial(F \times \{1\})$. Then $\partial F_1'$ has two components, with one copy of τ in each. Therefore, no such disk D'' could exist. Hence, τ is incident to two components of ∂F. The disk D'' demonstrates that α is a τ_2-arc. ■

3.2 Special Disks

Lemma 3.6 shows that disks whose boundaries contain no bad arcs and only one special arc are important. This motivates the following definition.

Definition 3.7 Given a disk D properly embedded in $F \times I$ such that ∂D is transverse to $\partial F \times \{0,1\}$, say that D is *special* if it is essential in $F \times I$, and there are no bad arcs and at most one special arc in ∂D. We will call D a 0-*special* or 1-*special* disk depending on the label and location of the special arc if one exists. If there is no special arc, then Lemma 3.2 says that there must be both τ_1- and τ_0-arcs, so for convenience we will distinguish one τ_1-arc as a special arc and say the disk is 1-special.

Lemma 3.8 *There exist special disks in* $\overline{(M \smallsetminus n(\tau)) \smallsetminus F'}$.

Proof As $M \smallsetminus n(\tau)$ is a genus two handlebody, we know that $\partial(M \smallsetminus n(\tau))$ is compressible in $M \smallsetminus n(\tau)$. Let D' be a compression disk such that $\partial D' \cap A$ consists of straight arcs, each essential in A and running from one component of $A_i \cap \partial M$ to the other, and such that $|D' \cap F'|$ is minimal among such disks. Since ∂M is incompressible in $M \smallsetminus n(\tau)$, we know that $\partial D'$ runs across A at least once.

If $D' \cap F' = \varnothing$, then D' is a disk in $F \times I$ and $\partial D'$ contains no special or bad arcs. Note that D' is essential in $F \times I$, since it is essential in $M \smallsetminus n(\tau)$ and F' is not a disk. Therefore, D' is a special disk.

If $D' \cap F' \neq \varnothing$, then notice that $D' \cap F'$ consists only of arcs, since circles of intersection innermost in D' and essential in F would give rise to compressions for F, and inessential ones could be removed to reduce $|D' \cap F'|$. Moreover, as τ is essential in F, the minimality of $|D' \cap F'|$ implies that every arc of $D' \cap F'$ is essential in F. Knowing this, the minimality of $|D' \cap F'|$ further implies that no arc of $D' \cap F'$ is isotopic to τ in F.

Consider an arc α of $D' \cap F'$ that is outermost in D', cutting off a subdisk D from D'. Now view D as a disk in $F \times I$. Without loss of generality, assume α is labeled 1. Note also that α is a special arc. Because ∂D contains exactly one special arc and no bad arcs, by Lemma 3.6 either the disk D is essential in $F \times I$ as required, or α is a τ_2-arc. In this case, α would cut off a disk from F'. This disk might contain other arcs of $D' \cap F'$. Boundary compressing D' along this disk would reduce $|D' \cap F'|$, creating at least two disks, at least one of which would contradict the minimality condition in the choice of D'. Therefore, D is essential, and so is a special disk. ∎

Lemma 3.9 *If D is an i-special disk in $F \times I$ for some $i \in \{0,1\}$, then ∂D contains at least one τ_{1-i}-arc.*

Proof Without loss of generality, assume that D is 1-special. Perform as many tightening-moves on D as possible. This does not change that D is 1-special and does not alter any τ_0-arcs in ∂D. Having done this, we see that $\partial D \cap (F \times \{0\})$ consists only of τ_0-arcs. As D is essential in $F \times I$, Lemma 3.2 implies that there must be at least one arc of $\partial D \cap (F \times \{0\})$ remaining, which is therefore a τ_0-arc. ∎

Now, consider the vertical product disks $E_0 = \tau_0 \times I$, and $E_1 = \tau_1 \times I$. We would like to find an i-special disk D, for some $i \in \{0,1\}$, such that ∂D and ∂E_i do not intersect on $F \times \{1-i\}$. Since $\pi(E_i) = \pi(\partial E_i \cap (F \times \{1-i\})) = \pi(\tau_i)$, and a τ_{1-i}-arc in ∂D projects under π to $\pi(\tau_{1-i})$, this would show that $\pi(\tau_0)$ and $\pi(\tau_1)$ are disjoint.

Recall that $h(\pi(\tau_0)) = \pi(\tau_1)$, and that we have assumed that the monodromy has been isotoped (including along the boundary) so as to minimize $|\pi(\tau_0) \cap \pi(\tau_1)|$.

Definition 3.10 The *size* of a special disk D is the triple $(|\partial D \cap (F \times \{0,1\})|, |D \cap E_j|, |\partial D \cap \partial E_j \cap (F \times \{0,1\})|)$, where D is a j-special disk. We will compare the size of two special disks using the lexicographical order.

That is, we order disks first by the total number of τ_0-, τ_1-, extra, and special arcs, second by the number of arcs and simple closed curves of intersection with the product disk E_j, and finally by the number of endpoints of these intersection arcs that lie on F.

It is worth noting that this situation looks similar to that found in [16, Lemma 2.3]. It appears that one could conclude immediately that a special disk was *boundary compressible towards* $F \times \{1\}$, and repeat such compressions until one arrived at a product disk. This is the idea of our proof, but we need to show some additional care as we want the arcs of $\partial D \cap (F \times \{0,1\})$ to stay parallel to τ_0 and τ_1 so that we can conclude something about the tunnel arc.

Since we know that special disks exist, we take a special disk with minimal size and call it D. Recall that if ∂D contains no special arc, then we have agreed to pick a τ_1-arc and call it special so that D is 1-special. On the other hand, if it does contain a special arc, then we can assume without loss of generality (by flipping $[0,1]$) that D is 1-special. In either case, call the special arc α.

Lemma 3.11 *There are no extra arcs in* ∂D.

Proof If there is an extra arc in ∂D, we can perform a tightening-move. This will reduce the number of extra arcs without changing the number of τ_0-, τ_1-, or special arcs. This, therefore, reduces the size of D, a contradiction. ∎

Lemma 3.11 implies that $\partial D \cap (F \times \{0\})$ consists only of τ_0-arcs. Let $E' = E_1$. Although it is not necessary, for notational convenience we will continue to assume that, for $i \in \{0,1\}$, all τ_i-arcs are contained within the rectangle A_i and run straight from one component of $A_i \cap \partial F$ to the other.

Lemma 3.12 *Every arc of* ∂D *on* $F \times \{1\}$ *is disjoint from* $\partial E'$.

Proof Choose $\varepsilon > 0$ such that $(\partial F \times [1-\varepsilon, 1)) \cap (\partial D \cup \partial E')$ consists of disjoint embedded arcs that are essential in the half-open annulus $\partial F \times [1-\varepsilon, 1)$. Let $F^+ = (F \times \{1\}) \cup (\partial F \times [1-\varepsilon, 1))$. Since $\partial D \cap (F \times \{1\})$ contains only τ_1- and special arcs, there is an isotopy of $\partial D \cap F^+$, fixed on ∂F^+, that makes ∂D disjoint from $\partial E'$ on $F \times \{1\}$. See Figure 4. Because $\partial E' \cap F^+$ is a single arc, this isotopy can be chosen so that it does not increase $|\partial D \cap \partial E' \cap F^+|$ at any point. Note that such an isotopy does not change the type of any arc of $\partial D \cap (F \times \{1\})$. Therefore, this means that the isotopy can be extended to an isotopy of D that does not increase $|D \cap E'|$. If $\partial D \cap \partial E' \cap (F \times \{1\}) \neq \varnothing$ before the isotopy then the isotopy strictly reduces the size of D, which is a contradiction. Thus, no such isotopy is required and $\partial D \cap \partial E' \cap (F \times \{1\}) = \varnothing$. ∎

Lemma 3.13 *We can assume that the endpoints of the arc* $\partial E' \cap (F \times \{0\})$ *are disjoint from* A_0, *and that every arc of* $\partial E' \cap A_0$ *connects opposite sides of* A_0 *and intersects each* τ_0*-arc of* ∂D *exactly once.*

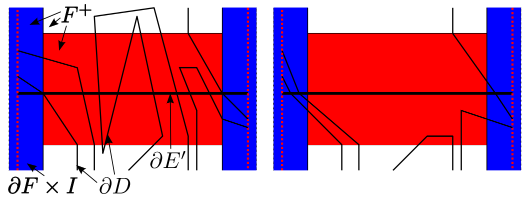

Figure 4: Arcs of $\partial D \cap (F \times \{1\})$ can be made disjoint from the arc of $\partial E' \cap (F \times \{1\})$ without increasing $|D \cap E'|$.

Proof By Lemma 3.11, $\partial D \cap (F \times \{0\})$ consists only of τ_0-arcs. We have assumed that each of these lies in A_0, connecting the two components of $A_0 \cap \partial(F \times \{0\})$. By isotoping $A_0 \cap F_0'$ in $(F \times \{0\}) \smallsetminus \partial D$, we can assume that $\partial E'$ is transverse to ∂A_0.

Consider the arcs of $\partial E' \cap A_0$. Each of the two sides of A_0 on $\partial(F \times \{0\})$ contains at most one endpoint of these arcs. All other endpoints must lie on the two components of $A_0 \cap F_0'$. Choose $\varepsilon > 0$ such that $((A_0 \cap \partial F) \times (0, \varepsilon]) \cap (\partial D \cup \partial E')$ consists of disjoint embedded arcs each having one endpoint on $(A_0 \cap \partial F) \times \{0\}$ and one endpoint on $(A_0 \cap \partial F) \times \{\varepsilon\}$. Let $A_0^+ = A_0 \cup ((A_0 \cap \partial F) \times (0, \varepsilon])$. As in the proof of Lemma 3.12, there is an isotopy of ∂D within A_0^+, fixed on ∂A_0^+, to minimize $|\partial D \cap \partial E' \cap A_0|$, and, moreover, this isotopy can be chosen so that it extends to an isotopy of D that does not increase the size of D (see Figure 5). Again, if this isotopy strictly reduced $|\partial D \cap \partial E' \cap A_0|$, then it would strictly reduce the size of D, contradicting that D was chosen to have minimal size. Therefore, no such isotopy is needed, and the arcs of $\partial E' \cap A_0^+$ have minimal intersection in A_0 with the arcs of $\partial D \cap A_0^+$.

Let γ be an arc of $\partial E' \cap A_0^+$. If the endpoints of γ lie on distinct components of $A_0 \cap F_0'$, then we see that γ intersects each arc of $\partial D \cap A_0$ exactly once, and because $|\partial D \cap (F \times \{0\})|$ has not increased, we know that this intersection occurs within A_0. If the endpoints of γ lie on the same component of $A_0 \cap F_0'$, then we find that γ is disjoint from ∂D. In this case we may isotope $A_0 \cap F_0'$ to remove γ from $\partial E' \cap A_0$ without affecting $\partial D \cap A_0$ (again, see Figure 5). If γ has one endpoint on $A_0 \cap \partial(F \times \{0\})$ and the other on $A_0 \cap F_0'$, then $\gamma \cap A_0$ is disjoint from $\partial D \cap A_0$, and again we can isotope ∂A_0 to remove γ from $\partial E' \cap A_0$. Finally, suppose that γ has both endpoints on components of $A_0 \cap \partial(F \times \{0\})$. Then γ is a τ_1-arc. Since $\pi(\gamma) = \pi(\tau_0)$ this shows that τ and $h(\tau)$ are isotopic in F, and in this case the proof of Theorem 1.1 is complete. ∎

Lemma 3.14 *Let γ be an arc of $\partial E' \cap F_0'$. If γ has both endpoints on $A_0 \cap F_0'$, then γ does not co-bound a disk in F_0' with $A_0 \cap F_0'$. If γ has one endpoint on $A_0 \cap F_0'$ and one on $\partial(F \times \{0\}) \smallsetminus A_0$, then γ does not cut off from F_0' a disk whose boundary consists of γ, a single sub-arc of $A_0 \cap F_0'$ and a single sub-arc of $\partial(F \times \{0\}) \smallsetminus A_0$.*

Proof Given Lemma 3.13, this follows immediately from the minimality of $|\pi(\tau_0) \cap \pi(\tau_1)|$ (see Figure 6). ∎

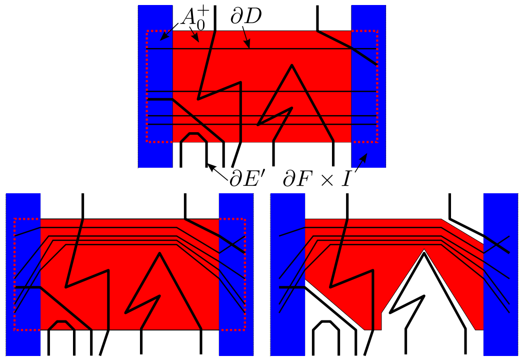

Figure 5: $|\partial D \cap \partial E' \cap A_0|$ and $|A_0 \cap F' \cap \partial E'|$ can be minimized without increasing $|D \cap E'|$.

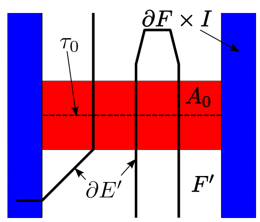

Figure 6: Arcs of $\partial E' \cap F'$ do not cut off certain types of disk.

Now consider $D \cap E'$. By innermost disk arguments, any simple closed curves of intersection could be removed, since $F \times I$ is irreducible. Thus, since D has minimal size, the intersection consists of arcs. From Lemma 3.12 we know that none of these intersection arcs have endpoints on $F \times \{1\}$. We will show that there are also no arcs of intersection with an endpoint on $F \times \{0\}$. There are three types of arcs that we will be concerned with: type 0 will be arcs with both endpoints on the same component of $\partial E' \cap (\partial F \times I)$; type I will be those with one endpoint on $F \times \{0\}$, and the other on $\partial F \times I$; type II will be arcs with both endpoints incident to $F \times \{0\}$ (see Figure 7). Showing that none of these arcs exist, and hence $\partial D \cap \partial E' \cap (F \times \{0\}) = \varnothing$, will complete the proof of Theorem 1.1. (Note that there may be arcs of intersection of $D \cap E'$ which have endpoints on different components of $\partial E' \cap (\partial F \times I)$, but these have no impact on the removal of arcs with an endpoint on $F \times \{0\}$.)

https://doi.org/10.4153/CJM-2016-002-1 Published online by Cambridge University Press

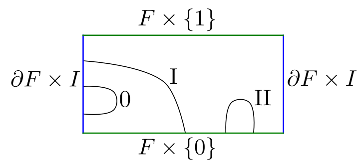

Figure 7: Arcs of $D \cap E'$ of type 0, type I, and type II in E'.

3.3 Arcs of Type 0

Suppose there is an arc of $D \cap E'$ with both endpoints on the same component of $E' \cap (\partial F \times I)$. Choose such an arc that is outermost in E', and let E be the subdisk of E' it cuts off. Compress D along E, reducing $|D \cap E'|$ without altering the arcs of $\partial D \cap (F \times \{0,1\})$. This gives two disks, D^* and D^{**}. Take D^* to be the one containing α in its boundary. At least one of D^* and D^{**} is essential, and neither has more than one special arc or any bad arcs in its boundary. In addition,

$$|D^* \cap (F \times \{0,1\})| \le |D \cap (F \times \{0,1\})| \quad \text{and} \quad |D^* \cap E'| < |D \cap E'|,$$

while $|D^{**} \cap (F \times \{0,1\})| < |D \cap (F \times \{0,1\})|$. Therefore, at least one of D^* and D^{**} is special and has smaller size than D, which is a contradiction. Hence, no arcs of type 0 exist.

3.4 Arcs of Type II

If there is an arc of type II, then there is an arc of type II that is outermost in E'. Call this arc δ, and call the subdisk of E' that it cuts off E. Let $\gamma = \partial E \setminus \delta$. Boundary compressing D along E reduces $|D \cap E'|$ and gives two disks, D^* and D^{**}, at least one of which is essential. Take D^* to be the resulting disk containing α in its boundary. The endpoints of γ must both be on τ_0-arcs.

First suppose that $\gamma \subset A_0$. Then by Lemma 3.13 we know that the endpoints of γ lie on distinct τ_0-arcs of ∂D. Let β^* and β^{**} be the sub-arcs of $\partial D^* \cap (F \times \{0\})$ and $\partial D^{**} \cap (F \times \{0\})$, respectively, that contain copies of γ. Then β^* and β^{**} are both extra arcs (see Figure 8), so neither D^* nor D^{**} has any bad arcs or more than one special arc in its boundary. Moreover, it is again the case that $|D^* \cap (F \times \{0,1\})| \le |D \cap (F \times \{0,1\})|$ and $|D^* \cap E'| < |D \cap E'|$, while $|D^{**} \cap (F \times \{0,1\})| < |D \cap (F \times \{0,1\})|$. This tells us that at least one of D^* and D^{**} is special and has smaller size than D, a contradiction.

Now assume instead that $\gamma \not\subset A_0$. Then it runs between two τ_0-arcs that are outermost in A_0. That is, γ runs from a sub-arc of ∂D, across one of the sides of ∂A_0 incident to F_0', through F_0', then across a side of ∂A_0 and to another sub-arc of ∂D. There are, then, two possibilities. Either γ returns to the same side of ∂A_0 (see Figure 9), or it returns to the other side of ∂A_0 (see Figure 10).

If γ returns to the same side of ∂A_0, then both endpoints must be incident to the same component of $\partial D \cap A_0$ (see Figure 9), and ∂D^{**} is a simple closed curve in

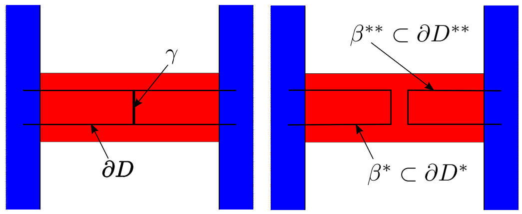

Figure 8: If $\gamma \subset A_0$, then β^* and β^{**} are extra arcs.

$F \times \{0\}$. Lemma 3.14 shows that ∂D^{**} does not bound a disk in F, so this means that D^{**} is a compression disk for F, contradicting that $F \times \{0\}$ is incompressible in $F \times I$.

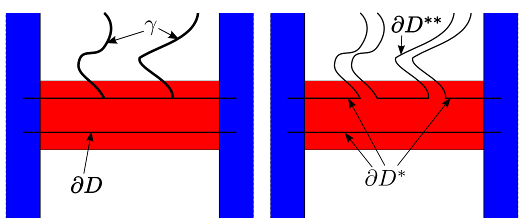

Figure 9: If $\gamma \not\subset A_0$, and returns to A_0 on the same side, then D^{**} is a compression disk for F.

If γ returns to the other side of ∂A_0, then the orientation on D implies that there are at least two τ_0-arcs in ∂D. Let β^* and β^{**} be the sub-arcs of ∂D^* and ∂D^{**}, respectively, that contain copies of γ (see Figure 10).

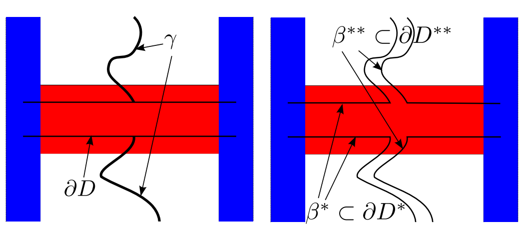

Figure 10: If $\gamma \not\subset A_0$, then $|D^* \cap (F \times \{0,1\})| + |D^{**} \cap (F \times \{0,1\})| = |D \cap (F \times \{0,1\})|$.

https://doi.org/10.4153/CJM-2016-002-1 Published online by Cambridge University Press

There are no bad arcs in either ∂D^* or ∂D^{**}, and there is at most one special arc in ∂D^{**}. As before, $|D^* \cap (F \times \{0,1\})| \leq |D \cap (F \times \{0,1\})|$ and $|D^* \cap E'| < |D \cap E'|$, while $|D^{**} \cap (F \times \{0,1\})| < |D \cap (F \times \{0,1\})|$.

If D^{**} is essential, then it is a special disk with smaller size than D, which is a contradiction. Suppose otherwise. Then D^* is essential. Additionally, by Lemma 3.6, β^{**} is either an extra arc, a τ_0-arc, or a τ_2-arc.

If β^{**} is a τ_0-arc, then τ is incident to only one boundary component of F, since the endpoints of β^{**} lie on the same component of ∂F. However, there is an arc parallel to τ_0 in A_0 that is disjoint from β^{**} and whose endpoints interleave on $\partial(F \times \{0\})$ with those of β^{**}. It is therefore impossible that these two arcs together bound a disk in $F \times \{0\}$. This shows that β^{**} is not a τ_0-arc.

If β^{**} is a τ_2-arc, then τ is incident to two boundary components of F and β^* is an extra arc. Thus, D^* is a special disk with smaller size than D, a contradiction.

If β^{**} is an extra arc, then τ is incident to two boundary components of F and β^* is a τ_2-arc. Let F^* be the subdisk of F_0' that β^* cuts off. Now, $(\partial E' \cap (F \times \{0\})) \smallsetminus \gamma$ consists of two arcs; call these γ' and γ''. From their endpoints that meet γ, both γ' and γ'' run to the opposite side of A_0, by Lemma 3.13. At this point, therefore, one of γ' and γ'' lies closer than the other in A_0 to the component of $A_0 \cap \partial(F \times \{0\})$ containing the endpoints of β^*. Take this to be γ'. See Figure 11.

Consider the path of γ'' from the endpoint that meets γ. When it first leaves A_0, γ'' enters the disk F^*. As we continue to follow its path, it can either end on the component of $\partial F \times \{0\}$ that contains the endpoints of β^{**} or else return to $A_0 \cap F_0'$ (necessarily on the other side, by Lemma 3.14). We see, therefore, that γ'' spirals around one boundary component of $F \times \{0\}$ some number of times before ending on this component of $\partial(F \times \{0\})$. Consider the final section of γ'', from where it last leaves A_0 to where it reaches $\partial(F \times \{0\})$. This cuts off a disk from F_0', the remainder of whose boundary consists of a single sub-arc of $A_0 \cap F_0'$ and a single sub-arc of $\partial(F \times \{0\}) \smallsetminus A_0$. This contradicts Lemma 3.14. It is therefore not possible that β^{**} is an extra arc.

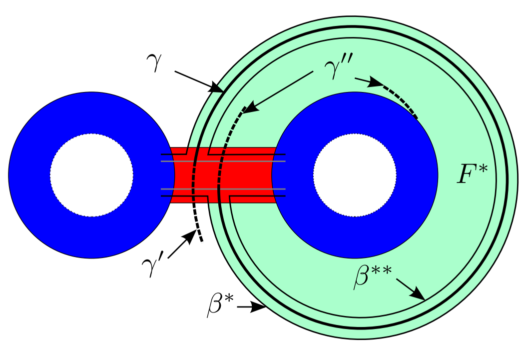

Figure 11: If β^{**} is an extra arc, then $|\partial F| = 2$ and γ'' spirals around one component of ∂F.

Thus, we conclude that there are no arcs of type II in $D \cap E'$.

3.5 Arcs of Type I

Since we now know there are no arcs of types 0 or II, if there are arcs of type I, then one of them is outermost in E'. Again, call one of these arcs δ, and call the subdisk of E' that it cuts off E. Let $\gamma = \partial E \smallsetminus \delta$. Then γ consists of two sub-arcs. Let $\gamma_0 = \gamma \cap (F \times \{0\})$ and $\gamma_\partial = \gamma \cap (\partial F \times I)$. Observe that γ_0 has one endpoint on a τ_0-arc of ∂D and the other end on $\partial F \smallsetminus A_0$, given Lemma 3.13. Note that, since γ_0 is disjoint on its interior from ∂D, Lemma 3.13 also tells us that $\gamma_0 \cap A_0$ is a single sub-arc of γ_0.

As before, boundary compressing D along E results in two disks, D^* and D^{**}, at least one of which is essential. Again let D^* be the one that contains α in its boundary. Let β^* and β^{**} be the sub-arcs of ∂D^* and ∂D^{**}, respectively, that contain copies of γ_0. As before, $|D^* \cap E'| < |D \cap E'|$, neither ∂D^* nor ∂D^{**} contains any bad arcs, and ∂D^{**} contains at most one special arc. Now

$$|\partial D^* \cap (F \times \{0,1\})| + |\partial D^{**} \cap (F \times \{0,1\})| = |\partial D \cap (F \times \{0,1\})| + 1.$$

In addition, $|\partial D^* \cap (F \times \{0,1\})| \geq 2$, while $|\partial D^{**} \cap (F \times \{0,1\})| \geq 1$. Therefore, $|\partial D^* \cap (F \times \{0,1\})| \leq |\partial D \cap (F \times \{0,1\})|$ and $|\partial D^{**} \cap (F \times \{0,1\})| < |\partial D \cap (F \times \{0,1\})|$.

From Lemma 3.14, we know that neither β^* nor β^{**} is an extra arc. If D^{**} is essential, then it is a special disk with smaller size than D, which is a contradiction. Suppose otherwise. Then D^* is essential. Additionally, by Lemma 3.6, β^{**} is either a τ_0-arc or a τ_2-arc.

If β^{**} is a τ_2-arc, then β^* is a τ_0-arc. Thus, D^* is a special disk that is smaller than D, a contradiction.

If β^{**} is a τ_0-arc, then, as β^{**} is disjoint from a copy of τ_0 in A_0, together these arcs bound a disk in F. If τ is incident to a single boundary component of F, the presence of this disk tells us that the endpoints of β^{**} do not interleave on ∂F with those of τ_0. Therefore the disk contains β^* and β^* is an extra arc, a contradiction.

It remains only to consider the case where τ has its endpoints on distinct components of ∂F, as does β^{**}. Again, if the disk between β^{**} and τ_0 contains β^*, then β^* is an extra arc, a contradiction. Accordingly, the disk does not contain β^*, and β^* is a τ_2-arc, cutting off from F_0' a disk F^*. Let $\gamma_0' = (\partial E' \cap (F \times \{0\})) \smallsetminus \gamma_0$. This is an arc with one endpoint on a τ_0-arc of ∂D, where it meets γ_0, and the other endpoint on $\partial(F \times \{0\}) \smallsetminus A_0$. Given the definition of E', this endpoint lies on the opposite component of $\partial(F \times \{0\})$ to the other endpoint of γ_0. That is, γ_0' does not meet the same component of $\partial(F \times \{0\})$ as β^* does. See Figure 12. Consider the path of γ_0' from where it meets γ_0. It first runs through A_0, and passes through $A_0 \cap F_0'$ into the disk F^*. As we continue to follow its path, it can either end on $\partial(F \times \{0\})$ or else return to $A_0 \cap F_0'$ (necessarily on the other side, by Lemma 3.14). We see that, like the arc γ'' above, γ_0' spirals around one boundary component of $F \times \{0\}$ some number of times before ending on the same component of $\partial(F \times \{0\})$. Consider the final section of γ_0', from where it last leaves A_0 to where it reaches $\partial(F \times \{0\})$. This cuts off a disk from F_0', the remainder of whose boundary consists of a single sub-arc of $A_0 \cap F_0'$ and a single sub-arc of $\partial(F \times \{0\}) \smallsetminus A_0$. This contradicts Lemma 3.14.

Thus, there are no arcs of type I. This completes the proof of Theorem 1.1. ■

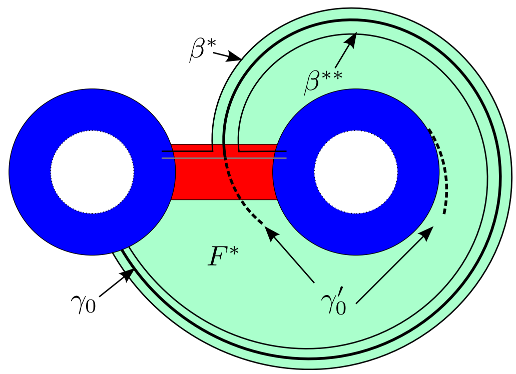

Figure 12: If β^{**} is a τ_0-arc and $|\partial F| = 2$, then γ_0' spirals around one component of ∂F.

By Proposition 2.13, we know that an unknotting tunnel τ for a fibered, tunnel number one link K in a manifold M can be isotoped to lie in a fiber F. We can now prove the following theorem.

Theorem 1.2 *Suppose K is a tunnel number one, fibered link in a 3-manifold M, with fiber F, monodromy h, and a properly embedded arc τ in F that is an unknotting tunnel for K. Then there exists a properly embedded arc $\beta \subset F$, freely ambient isotopic in F to $h(\tau)$, so that $\tau \cap \beta = \varnothing$. In particular, up to isotopy* rel ∂F, *there exists a regular neighborhood of ∂F outside of which τ and $h(\tau)$ do not intersect.*

Proof Recall that $(M \setminus n(K))$ is a surface bundle, that h is a particular monodromy for the bundle, and that by definition of unknotting tunnel, $(M \setminus n(K)) \setminus n(\tau)$ is a genus two handlebody. So, the hypotheses of Theorem 1.1 apply, and the statement follows. ∎

4 Boundary Twisting and Fractional Dehn Twists

In this section, we will discuss why the free isotopy mentioned in Theorem 1.2 is necessary, why a stronger claim about unknotting tunnels being clean cannot be made in general, and some remaining open questions.

4.1 Full Twisting

We first consider full twists around boundary components of the fiber surface.

Example 4.1 First, consider a surface bundle $M = (F \times I)/h$ as in Theorem 1.1, and suppose M is tunnel number one (*i.e.,* that there is an arc $\tau \subset F$ such that $M \setminus n(\tau)$ is a genus two handlebody). Let T_∂ be a Dehn twist along a curve in F that is parallel to a component of ∂F. Then for all $n \in \mathbb{Z}$, the maps h and $T_\partial^n \circ h$ are freely isotopic, so that $(F \times I)/h \cong (F \times I)/(T_\partial^n \circ h)$. In fact, $\tau \subset F$ is still an unknotting tunnel for $(F \times I)/(T_\partial^n \circ h)$. However, even if τ is clean with respect to h, there will be intersections between τ and $(T_\partial^n \circ h)(\tau)$ in a neighborhood of ∂F for all sufficiently high values of $|n|$. These intersections can be removed by freely isotoping $(T_\partial^n \circ h)(\tau)$ independently of τ, but then the arc does not correspond to the image of τ under the map $(T_\partial^n \circ h)$. If we consider the surface bundle as the exterior of a link in some 3-manifold, then these twists can be thought to affect the meridian(s) of the link, and can be viewed as changing the ambient 3-manifold in which the fibered link sits. So generically, weak cleanliness of unknotting tunnels is the best that can be hoped for.

One might hope that this type of indeterminacy would improve if we restrict our attention to knots and links in S^3, as this would specify the representative monodromy map by determining the meridian(s). We next, therefore, consider an example in S^3, suggested to the authors by Ken Baker.

Example 4.2 Suppose τ is the upper (or lower) tunnel for a fibered 2-bridge knot K in S^3 (see [23]), sitting in a fiber surface F as a clean arc such that $h(\tau) \neq \tau$. Now, perform a Hopf plumbing along an arc that is parallel into ∂F, but has endpoints interleaved on K with those of τ. The result is $K\#L \subset S^3$, where L is a Hopf link and has a monodromy map h' that is a composition of h with a Dehn twist around the core curve of the Hopf band. The choice of sign for the Hopf band determines the orientation on the link as well as the sign of the Dehn twist. Either way, τ is an unknotting tunnel of $K\#L$, since τ together with the unknotted component of the link is actually equivalent to one of the dual upper (or lower) tunnels for K (see [23]). Although one choice results in a monodromy under which τ is still clean, the other results in a monodromy under which it is not, since the extra twist forces an intersection between τ and $h'(\tau)$ in a neighborhood of the boundary of the fiber. See Figure 13.

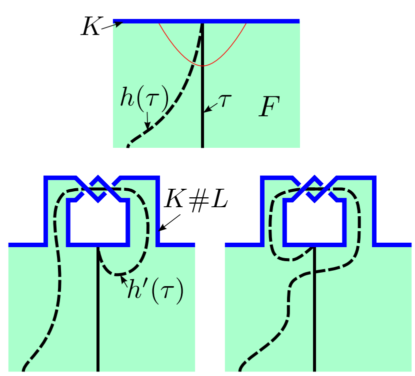

Figure 13: One choice of Hopf plumbing gives a clean tunnel while the other does not.

In fact, it is not only in the case of a connected sum with a Hopf link that this complication with boundary twisting arises. Kai Ishihara pointed out to the authors that if L is a tunnel number one, fibered, two-component link in S^3 with one trivial component, K, and linking number ± 2 or 0, then modifying the monodromy by n Dehn twists ($n = \mp 1$ or n arbitrary, respectively) along a curve in the fiber parallel to K corresponds to performing Stallings twists, and produces tunnel number one, fibered links in S^3. Each of these links has an unknotting tunnel that intersects its image (several times) in a neighborhood of the boundary of the fiber, precisely because the twisting was performed around a curve parallel (in the fiber) to a boundary component of the fiber.

Example 4.3 One such example is the Whitehead link, which has linking number zero and is also hyperbolic. Figure 14 (left) shows the link resulting from twisting $n = 3$ times around one of the components of the Whitehead link, along with an unknotting tunnel, τ, for this link. One can check that the surface illustrated is a fiber (since it is genus one, *i.e.,* minimal genus), and that τ is an unknotting tunnel for the link. The arc τ is not clean, as the image of τ under the monodromy is indicated. Alternatively, one can see that τ cannot be clean, because cutting the fiber surface along the tunnel arc produces a surface whose boundary is the 5_2 knot. If τ were clean and alternating, then it would correspond to a plumbed Hopf band, the de-plumbing of which would result in a genus one fiber surface with a connected boundary, so the boundary would be a trefoil or figure-eight knot. On the other hand, if τ were clean and non-alternating, then cutting along τ would result in a pre-fiber surface (see [22]), which itself would be a (genus one) compressible surface, implying that the boundary was the unknot.

Twisting the same component of the Whitehead link an arbitrary n times also results in a new tunnel number one, fibered link. In Figure 14 (right), the light gray arc still indicates an unknotting tunnel, and the black train track with weights determines the arc that is the image of this tunnel under the monodromy for this surface.

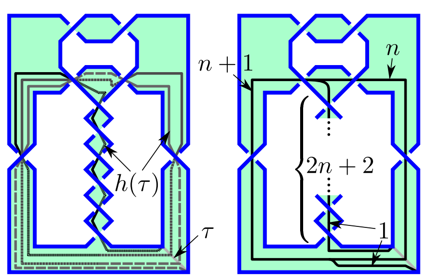

Figure 14: A hyperbolic, tunnel number one, fibered link in S^3 with an unclean tunnel obtained by twisting the Whitehead link around an unknotted component $n = 3$ (left) or $n \geq 1$ (right) times.

https://doi.org/10.4153/CJM-2016-002-1 Published online by Cambridge University Press

In spite of the examples discussed above, it remains possible that the following question has an affirmative answer.

Question 1 If a tunnel number one, fibered link of two components in S^3 has an unclean unknotting tunnel, then must one of the components be unknotted?

We will see that this kind of full twisting around the boundary cannot occur for (nontrivial) tunnel number one, fibered knots in S^3. However, fractional twisting remains possible.

4.2 Fractional Dehn Twists

We next consider partial twisting around boundary components of the fiber surface.

Thurston classified automorphisms of a (hyperbolic) surface. Every automorphism $f: F \to F$ is freely isotopic to one, $\widetilde{f}$, that is either (1) *reducible*, (2) *periodic*, or (3) *pseudo-Anosov* (see [4, 37]). In all cases, $\widetilde{f}$ is called the *Thurston representative* of f. (We follow the convention of referring to a map as reducible only if it is not periodic.)

By Thurston's Hyperbolization Theorem a surface bundle over S^1 is hyperbolic if and only if the (Thurston representative of the) monodromy map is pseudo-Anosov (see [27, 28, 36]). Since the Whitehead link is hyperbolic, the Thurston representative of its monodromy is correspondingly pseudo-Anosov. Observe that this means the family of examples given in Figure 14 are all hyperbolic, since all of their respective monodromies are freely isotopic to the monodromy of the Whitehead link.

The *fractional Dehn twist coefficient* of a surface automorphism h at a boundary component of the surface measures the amount of twisting around that boundary component necessary to freely isotope h to its Thurston representative. While the details differ slightly between the cases of the different Thurston types, the fractional Dehn twist coefficient is a rational number p/q, (with p and q relatively prime), which corresponds to a $2p\pi/q$ rotation around a boundary component of the surface.

The relevance of fractional Dehn twist coefficients for the question of clean arcs in fibers of fibered links is as follows. Suppose h is the monodromy of a fibered link complement with fiber F, α is an arc properly embedded in a fiber surface with both endpoints on the same component of ∂F, and $\widetilde{h}$ is the Thurston representative of h (with respect to a fixed hyperbolic structure on the fiber). Take $\overline{\alpha}$ to be the geodesic arc freely isotopic to α, and $\overline{h(\alpha)}$ to be the geodesic arc freely isotopic to $h(\alpha)$, which will also be freely isotopic to $\widetilde{h}(\alpha)$. As these are geodesic arcs, they intersect minimally among their free isotopy representatives. Let A be a small annular neighborhood of the boundary component to which α is incident. Since h has the property that $h|_{\partial F} = \mathrm{Id}$, the arc $h(\alpha)$ can be realized by replacing $\overline{h(\alpha)} \cap A$ with arcs that monotonically spiral around A with rotation $2p\pi/q$, where h has fractional Dehn twist coefficient of p/q at the relevant boundary component. This spiralling may necessarily result in intersections between α and $h(\alpha)$ in their interiors, as indicated in the statement of Theorem 1.2.

When the surface bundle is a *knot* complement in S^3, works of Gabai [13] and Kazez and Roberts [21] have shown that the fractional Dehn twist coefficient is either 0 or $1/n$ for some integer n, $|n| \geq 2$. In particular, this means that if we orient an arc

α in a fiber surface for a knot in S^3, then an initial sub-arc of α and an initial sub-arc of $h(\alpha)$ will not have any intersections in a neighborhood of the boundary of the fiber owing to fractional Dehn twisting. Thus, the only intersections that could be introduced by fractional Dehn twisting will be between an initial sub-arc of α and a terminal sub-arc of $h(\alpha)$, or vice versa. In fact, if the fractional Dehn twist coefficient is $1/n$ with $|n| > 2$, then only one of these two can occur. For an unknotting tunnel sitting as an arc in the fiber, these are the only intersections that occur at all, so we get the following slight refinement of Theorem 1.2.

Theorem 4.4 *Suppose K is a tunnel number one, fibered knot in S^3, with fiber F, monodromy h, and a properly embedded arc τ in F that is an unknotting tunnel for K. Then τ and $h(\tau)$ can be ambient isotoped in F rel ∂F so that $|\operatorname{int}(\tau) \cap \operatorname{int}(h(\tau))| \le 2$, and any such intersections occur in a regular neighborhood of ∂F; moreover, if $|\operatorname{int}(\tau) \cap \operatorname{int}(h(\tau))| = 2$, then h has a fractional Dehn twist coefficient of $\pm 1/2$.*

We ask the following optimistic question.

Question 2 Is the unknotting tunnel of a tunnel number one, fibered *knot* in S^3 always clean?

While Theorem 4.4 does not rule out a negative answer, the authors know of no examples demonstrating one.

Should the answers to both Questions 1 and 2 turn out to be yes, might all tunnel number one, fibered links in S^3 be obtainable by a sequence of operations like twisting around unknotted boundary components, plus Hopf plumbing, de-plumbing, and Stallings twisting restricted to locations determined by unknotting tunnels?

5 An Application to Hyperbolic Cusps

In [10], Futer and Schleimer study the hyperbolic structure on a hyperbolic surface bundle M. Each boundary component of M is a cusp in the hyperbolic structure. If we pick one boundary component, expanding a regular neighborhood of the corresponding cusp until it "bumps into itself" gives a well-defined "maximal cusp". The geometric properties of the bounding torus of this neighborhood are invariants of the manifold M. Futer and Schleimer relate this geometry to the action of the (pseudo-Anosov) monodromy on the arc complex of the fiber surface.

Given a compact, connected surface F with boundary, the *arc complex* $\mathcal{A}(F)$ is a simplicial complex. The vertices of the complex are free isotopy classes of essential arcs properly embedded in F. Distinct vertices span a simplex exactly when the free isotopy classes of arcs can be simultaneously realized disjointly in F. (Note that sometimes $\mathcal{A}(F)$ is also used to denote the complex whose vertices are essential arcs up to isotopy rel ∂, though this is not the usage here.) A homeomorphism h of F induces a homeomorphism h_* of $\mathcal{A}(F)$. The *translation distance* $\mathrm{d}_{\mathcal{A}}(h)$ of h is

$$\mathrm{d}_{\mathcal{A}}(h) = \min_{v \in \mathcal{A}^{(0)}(F)} \mathrm{d}(v, h_*(v)).$$

https://doi.org/10.4153/CJM-2016-002-1 Published online by Cambridge University Press

Here the distance d is measured in the 1-skeleton $\mathcal{A}^{(1)}(F)$, where each edge has length 1. The *stable translation distance* $\overline{\mathrm{d}}_{\mathcal{A}}(h)$ is given by

$$\overline{\mathrm{d}}_{\mathcal{A}}(h) = \lim_{n\to\infty} \frac{\mathrm{d}(v, h_*^n(v))}{n},$$

where v is any vertex of $\mathcal{A}(F)$. The triangle inequality implies that $\overline{\mathrm{d}}_{\mathcal{A}}(h) \le \mathrm{d}_{\mathcal{A}}(h)$.

We claim that a pseudo-Anosov homeomorphism cannot fix an essential arc in the surface. Assume that F is not a disk or an annulus, as there are no pseudo-Anosov homeomorphisms on disks or annuli. Also, assume that F is not a pair of pants, as a homeomorphism fixing an essential arc in a pair of pants must be isotopic either to the identity or a rotation of order 2, neither of which are pseudo-Anosov homeomorphisms. Let γ be an essential arc in F. Suppose that $h'\colon F \to F$ is a map isotopic to h with $h'(\gamma) = \gamma$.

First, suppose that γ has its endpoints on the same component of ∂F. The endpoints of γ divide the boundary component of F into two arcs. Let γ_1, γ_2 be the simple closed curves given by combining each of these two arcs with a copy of γ. At least one of γ_1 or γ_2 must be an essential curve in F, for otherwise they are both isotopic to boundary components of F, and F would be a pair of pants. Then, since at least one of γ_1 or γ_2 is essential, either h' fixes γ_1 and γ_2, and we have an essential curve fixed by h', or h' exchanges them, in which case $\gamma_1 \cup \gamma_2$ is an essential multi-curve fixed by h'.

On the other hand, suppose γ has its endpoints on distinct components of ∂F. Let γ' be a simple closed curve that runs parallel to γ, around one boundary component of ∂F on which γ has an endpoint, back parallel to γ and around the other boundary component. Then, up to isotopy, $h'(\gamma') = \gamma'$. The curve γ' must be essential, else F would be a pair of pants, and h', again, fixes an essential curve.

Thus, since a pseudo-Anosov homeomorphism cannot fix an essential multi-curve, it cannot fix an essential arc.

Written in this language, Theorem 1.1 says the following.

Corollary 5.1 *If the surface bundle $(F\times I)/h$ has tunnel number one, then $\mathrm{d}_{\mathcal{A}}(h) \le 1$. If h is pseudo-Anosov, then $\mathrm{d}_{\mathcal{A}}(h) = 1$.*

Given this, [10, Theorem 1.5] yields the following result.

Theorem 5.2 *If the surface bundle $(F \times I)/h$ has tunnel number one, $|\partial F| = 1$, and h is pseudo-Anosov, then the area of the maximal cusp is bounded above by $9\chi(F)^2$, and the height of the cusp is strictly less than $-3\chi(F)$.*

Here the *height* of the cusp torus is its area divided by the length of the longitude.

We remark that [10, Theorem 1.5] also gives lower bounds on these quantities in terms of $\overline{\mathrm{d}}_{\mathcal{A}}(h)$. In [14], Gadre and Tsai studied the analogous distance in the curve complex, giving an explicit lower bound. It seems plausible that such a bound could likewise be obtained for the arc complex.

David Futer pointed out the following corollary of Corollary 5.1 to the authors.

Corollary 5.3 *There exists a family of fibered knots K_n, each having monodromy with translation distance* 1*, such that the cusp area grows linearly with the knot genus.*

Proof For $n \geq 1$, let K_n be the $(6n+1)$-crossing knot with diagram D_n formed from the blocks in Figure 15, taking one of each of the outer two blocks and n of the inner one. In addition, let R_n be the Seifert surface for K_n constructed by combining the

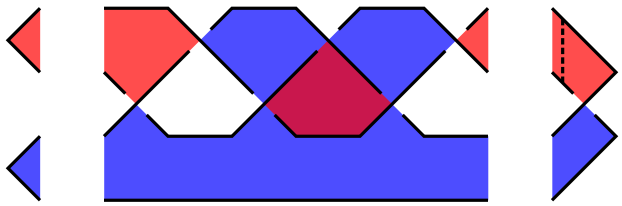

Figure 15: We build the knot K_n by combining n copies of the middle block with one copy of each of the outer blocks.

pieces of surface shown in Figure 15. As D_n is alternating, this surface has minimal genus. Note that $\chi(R_n) = 1 - 4n$, so K_n has genus $2n$.

For $m \in \mathbb{N}$, let f_m denote the m-th term of the Fibonacci sequence (so $f_1 = f_2 = 1$, $f_3 = 2$, $f_4 = 3$, $f_5 = 5$, etc.). Then K_n is the rational knot corresponding to the fraction f_{6n+1}/f_{6n+2}. A rational knot with fraction $1/q$ for some q is a torus knot, and all other rational knots are hyperbolic (see, for example, [3]). Two fractions p_1/q_1 and p_2/q_2 (with p_i coprime to q_i) correspond to the same rational knot if and only if $p_1 = p_2$ and either $q_1 \cong q_2 \mod p_1$ or $q_1 q_2 \cong 1 \mod p_1$. Since $f_{6n+1} \neq 1$ for $n \geq 1$, this shows that K_n is hyperbolic for each n.

That R_n is a fiber surface can be checked directly by product disk decompositions (see [12]) — $2n$ product disk decompositions can be used to remove the "trefoil pattern" in the center of each of the n middle blocks, leaving a checkerboard surface; further product decompositions can be used to reduce the surface to a disk (by removing the white bigons in the remaining diagram).

Being rational knots, each K_n has tunnel number one, with a tunnel given by the dotted arc in Figure 15. Therefore, Corollary 5.1 applies, and the monodromy of K_n has translation distance 1.

In a link diagram, a *twist region* is a maximal collection of crossings connected in a line by bigons. Each diagram D_n is twist-reduced and has $6n-1$ twist regions. Thus, [9, Theorem 4.8] gives that, for the knot K_n, the area a_n of the maximal cusp satisfies

$$\frac{1}{12}(6n-2) \leq a_n < \frac{40}{3}(6n-2). \qquad \blacksquare$$

Corollary 5.3 shows that the dependence on the Euler characteristic in the area bound in [10, Theorem 1.5] and in Theorem 5.2 is necessary.

https://doi.org/10.4153/CJM-2016-002-1 Published online by Cambridge University Press

Acknowledgment The authors would like to give special thanks to Abigail Thompson, Ken Baker, Kai Ishihara, Dave Futer, and the reviewer, who helped to improve this paper substantially.

References

[1] K. L. Baker, J. E. Johnson, and E. A. Klodginski, *Tunnel number one, genus-one fibered knots.* Comm. Anal. Geom. **17**(2009), no. 1, 1–16. http://dx.doi.org/10.4310/CAG.2009.v17.n1.a1

[2] J. S. Birman, *A representation theorem for fibered knots and their monodromy maps.* In: Topology of low-dimensional manifolds (Proc. Second Sussex Conf., Chelwood Gate, 1977), Lecture Notes in Math., 722, Springer, Berlin, 1979, pp. 1–8.

[3] M. Brittenham and Y.-Q. Wu, *The classification of exceptional Dehn surgeries on 2-bridge knots.* Comm. Anal. Geom. **9**(2001), no. 1, 97–113. http://dx.doi.org/10.4310/CAG.2001.v9.n1.a4

[4] A. J. Casson and S. A. Bleiler, *Automorphisms of surfaces after Nielsen and Thurston*. London Mathematical Society Student Texts, 9, Cambridge University Press, Cambridge, 1988. http://dx.doi.org/10.1017/CBO9780511623912

[5] S. Cho and D. McCullough, *Cabling sequences of tunnels of torus knots.* Algebr. Geom. Topol. **9**(2009), no. 1, 1–20. http://dx.doi.org/10.2140/agt.2009.9.1

[6] ———, *The tree of knot tunnels.* Geom. Topol. **13**(2009), no. 2, 769–815. http://dx.doi.org/10.2140/gt.2009.13.769

[7] ———, *Constructing knot tunnels using giant steps.* Proc. Amer. Math. Soc. **138**(2010), no. 1, 375–384. http://dx.doi.org/10.1090/S0002-9939-09-10069-2

[8] A. Coward and M. Lackenby, *Unknotting genus one knots.* Comment. Math. Helv. **86**(2011), no. 2, 383–399. http://dx.doi.org/10.4171/CMH/227

[9] D. Futer, E. Kalfagianni, and J. S. Purcell, *Cusp areas of Farey manifolds and applications to knot theory.* Int. Math. Res. Not. IMRN **2010**, no. 23, 4434–4497. http://dx.doi.org/10.1093/imrn/rnq037

[10] D. Futer and S. Schleimer, *Cusp geometry of fibered 3-manifolds.* Amer. J. Math. **136**(2014), no. 2, 309–356. http://dx.doi.org/10.1353/ajm.2014.0012

[11] D. Gabai, *The Murasugi sum is a natural geometric operation. II.* In: Combinatorial methods in topology and algebraic geometry (Rochester, N.Y., 1982), Contemp. Math., 44, American Mathematical Society, Providence, RI, 1985, pp. 93–100. http://dx.doi.org/10.1090/conm/044/813105

[12] ———, *Detecting fibred links in S^3*. Comment. Math. Helv. **61**(1986), no. 4, 519–555. http://dx.doi.org/10.1007/BF02621931

[13] ———, *Problems in foliations and laminations.* In: Geometric topology (Athens, GA, 1993), AMS/IP Stud. Adv. Math., 2.2, American Mathematical Society, Providence, RI, 1997, pp. 1–33.

[14] V. Gadre and C.-Y. Tsai, *Minimal pseudo-Anosov translation lengths on the complex of curves.* Geom. Topol. **15**(2011), no. 3, 1297–1312. http://dx.doi.org/10.2140/gt.2011.15.1297

[15] E. Giroux and N. Goodman, *On the stable equivalence of open books in three-manifolds.* Geom. Topol. **10**(2006), 97–114. http://dx.doi.org/10.2140/gt.2006.10.97

[16] H. Goda, *Heegaard splitting for sutured manifolds and Murasugi sum.* Osaka J. Math. **29**(1992), no. 1, 21–40.

[17] C. McA. Gordon and A. W. Reid, *Tangle decompositions of tunnel number one knots and links.* J. Knot Theory Ramifications **4**(1995), no. 3, 389–409. http://dx.doi.org/10.1142/S0218216595000193

[18] J. Harer, *How to construct all fibered knots and links.* Topology **21**(1982), no. 3, 263–280. http://dx.doi.org/10.1016/0040-9383(82)90009-X

[19] H. M. Hilden, D. M. Tejada, and M. M. Toro, *Tunnel number one knots have palindrome presentations.* J. Knot Theory Ramifications **11**(2002), no. 5, 815–831. http://dx.doi.org/10.1142/S0218216502001998

[20] J. Johnson, *Surface bundles with genus two Heegaard splittings.* J. Topol. **1**(2008), 671–692. http://dx.doi.org/10.1112/jtopol/jtn018

[21] W. H. Kazez and R. Roberts, *Fractional Dehn twists in knot theory and contact topology*. Algebr. Geom. Topol. **13**(2013), no. 6, 3603–3637. http://dx.doi.org/10.2140/agt.2013.13.3603

[22] T. Kobayashi, *Fibered links and unknotting operations.* Osaka J. Math. **26**(1989), no. 4, 699–742.

[23] ———, *Classification of unknotting tunnels for two bridge knots.* In: Proceedings of the Kirbyfest (Berkeley, CA, 1998), Geom. Topol. Monogr., 2, Geom. Topol. Publ., Coventry, 1999, pp. 259–290. http://dx.doi.org/10.2140/gtm.1999.2.259

[24] K. Morimoto, M. Sakuma, and Y. Yokota, *Identifying tunnel number one knots.* J. Math. Soc. Japan **48**(1996), no. 4, 667–688. http://dx.doi.org/10.2969/jmsj/04840667
[25] T. Nakamura, *On canonical genus of fibered knot.* J. Knot Theory Ramifications **11**(2002), no. 3, 341–352. http://dx.doi.org/10.1142/S0218216502001652
[26] Y. Ni, *Knot Floer homology detects fibred knots.* Invent. Math. **170**(2007), no. 3, 577–608. http://dx.doi.org/10.1007/s00222-007-0075-9
[27] J.-P. Otal, *Le théorème d'hyperbolisation pour les variétés fibrées de dimension 3.* Astérisque, 235, 1996.
[28] ______, *The hyperbolization theorem for fibered 3-manifolds.* SMF/AMS Texts and Monographs, 7, American Mathematical Society, Providence, RI, 2001.
[29] M. Rathbun, *Tunnel one, fibered links.* Pacific J. Math. **259**(2012), no. 2, 473–481. http://dx.doi.org/10.2140/pjm.2012.259.473
[30] M. Sakuma, *Murasugi decompositions of certain Seifert surfaces and unknotting operations.* In: Proceedings of the 5th Korea-Japan School of knots and links, 1997, pp. 233–248.
[31] ______, *Unknotting tunnels and canonical decompositions of punctured torus bundles over a circle. Analysis of discrete groups (Kyoto, 1995).* Sūrikaisekikenkyūsho Kōkyūroku, **967**(1996), 58–70.
[32] ______, *The topology, geometry and algebra of unknotting tunnels.* Chaos, Solitons and Fractals **9**(1998), no. 4–5, 739–748. http://dx.doi.org/10.1016/S0960-0779(97)00101-X
[33] M. Scharlemann, *Tunnel number one knots satisfy the Poenaru conjecture.* Topology Appl. **18**(1984), no. 2–3, 235–258. http://dx.doi.org/10.1016/0166-8641(84)90013-0
[34] M. Scharlemann and A. Thompson, *Unknotting tunnels and Seifert surfaces.* Proc. London Math. Soc. (3) **87**(2003), no. 2, 523–544. http://dx.doi.org/10.1112/S0024611503014242
[35] J. R. Stallings, *Constructions of fibred knots and links.* In: Algebraic and geometric topology (Proc. Sympos. Pure Math., Stanford Univ., Stanford, Calif., 1976), Part 2, Proc. Sympos. Pure Math., 32, American Mathematical Society, Providence, RI, 1978, pp. 55–60.
[36] D. Sullivan, *Travaux de Thurston sur les groupes quasi-fuchsiens et les variétés hyperboliques de dimension* 3 *fibrées sur* S^1. In: Bourbaki Seminar, Vol. 1979/80, Lecture Notes in Math., 842, Springer, Berlin-New York, 1981, pp. 196–214.
[37] W. P. Thurston, *On the geometry and dynamics of diffeomorphisms of surfaces.* Bull. Amer. Math. Soc. (N.S.) **19**(1988), no. 2, 417–431. http://dx.doi.org/10.1090/S0273-0979-1988-15685-6
[38] R. Yamamoto, *Stallings twists which can be realized by plumbing and deplumbing Hopf bands.* J. Knot Theory Ramifications **12**(2003), no. 6, 867–876. http://dx.doi.org/10.1142/S0218216503002779

University of Hull, Hull, UK, HU6 7RX
e-mail: jessica.banks@lmh.oxon.org j.banks@hull.ac.uk

California State University, Fullerton, 800 N. State College Blvd., Fullerton, CA, 92831
e-mail: mrathbun@fullerton.edu