

A random tunnel number one 3–manifold does not fiber over the circle

NATHAN M DUNFIELD

DYLAN P THURSTON

We address the question: how common is it for a 3–manifold to fiber over the circle? One motivation for considering this is to give insight into the fairly inscrutable Virtual Fibration Conjecture. For the special class of 3–manifolds with tunnel number one, we provide compelling theoretical and experimental evidence that fibering is a very rare property. Indeed, in various precise senses it happens with probability 0. Our main theorem is that this is true for a measured lamination model of random tunnel number one 3–manifolds.

The first ingredient is an algorithm of K Brown which can decide if a given tunnel number one 3–manifold fibers over the circle. Following the lead of Agol, Hass and W Thurston, we implement Brown’s algorithm very efficiently by working in the context of train tracks/interval exchanges. To analyze the resulting algorithm, we generalize work of Kerckhoff to understand the dynamics of splitting sequences of complete genus 2 interval exchanges. Combining all of this with a “magic splitting sequence” and work of Mirzakhani proves the main theorem.

The 3–manifold situation contrasts markedly with random 2–generator 1–relator groups; in particular, we show that such groups “fiber” with probability strictly between 0 and 1.

[57R22](#); [57N10](#), [20F05](#)

We dedicate this paper to the memory of Raoul Bott (1923–2005), a wise teacher and warm friend, always searching for the simplicity at the heart of mathematics.

1 Introduction

In this paper we are interested in compact orientable 3–manifolds whose boundary, if any, is a union of tori. A nice class of such manifolds are those that *fiber over the circle*, that is, are fiber bundles over the circle with fiber a surface F :

$$F \rightarrow M \rightarrow S^1$$

Equivalently, M can be constructed by taking $F \times [0, 1]$ and gluing $F \times \{0\}$ to $F \times \{1\}$ by a homeomorphism ψ of F . Manifolds which fiber over the circle are usually easier to understand than 3-manifolds in general, because many questions can be reduced to purely 2-dimensional questions about the gluing map ψ .

When M fibers over the circle, the group $H^1(M; \mathbb{Z}) \cong H_2(M, \partial M; \mathbb{Z})$ is nonzero; a nontrivial element is the fibering map to S^1 , or dually, the fiber surface F . Our main question is:

1.1 Question *If we suppose $H^1(M; \mathbb{Z}) \neq 0$, how common is it for M to fiber over the circle?*

We will give several reasons for entertaining this question below, but for now one motivation (beyond its inherent interest) is to try to estimate how much harder the Virtual Fibration Conjecture is than other variants of the Virtual Haken Conjecture. In this paper we provide evidence, both theoretical and experimental, that the answer to [Question 1.1](#) is: not very common at all. In fact, for the limited category of 3-manifolds that we study here, the probability of fibering is 0.

The type of 3-manifold we focus on here are those with tunnel number one, which we now define. Let H be an orientable handlebody of genus 2, and pick an essential simple closed curve γ on ∂H . Now build a 3-manifold M by gluing a 2-handle to ∂H along γ ; that is, $M = H \cup (D^2 \times I)$ where $\partial D^2 \times I$ is glued to ∂H along a regular neighborhood of γ . Such a manifold is said to have *tunnel number one*. There are two kinds of these manifolds, depending on whether the curve γ is separating or not. For concreteness, let us focus on those where γ is non-separating. In this case, M has one boundary component, which is a torus. A simple example of a 3-manifold with tunnel number one is the exterior of a 2-bridge knot in S^3 .

The boundary of a tunnel number one manifold M forces $H^1(M; \mathbb{Z}) \neq 0$, and so it makes sense to consider [Question 1.1](#) for all manifolds in this class. To make this question more precise, we will need a notion of a “random” tunnel number one manifold, so that we can talk about probabilities. In fact, there are several reasonable notions for this; here, we focus on two which involve selecting the attaching curve $\gamma \subset \partial H$ from the point of view of either measured laminations or the mapping class group.

For measured laminations, the setup is roughly this. We fix Dehn–Thurston coordinates on the set of multicurves (equivalently integral measured laminations) on the surface ∂H . Let $\mathcal{T}(r)$ consist of the tunnel number one 3-manifolds whose attaching curve $\gamma \subset \partial H$ has all coordinates of size less than r . As $\mathcal{T}(r)$ is finite, it makes sense to formulate a precise version of [Question 1.1](#) as: what is the proportion of $M \in \mathcal{T}(r)$ which fiber over the circle when r is large? The main theorem of this paper is:

Theorem 2.4 *The probability that $M \in \mathcal{T}(r)$ fibers over the circle goes to 0 as $r \rightarrow \infty$.*

Thus with this notion of a random tunnel number one 3–manifold, being fibered is very rare indeed. There is one technical caveat here: the set $\mathcal{T}(r)$ we consider does not cover all multicurves on ∂H , although we can always change coordinates, preserving H , to put any curve in $\mathcal{T}(r)$. See the discussion in [Section 2.2](#).

Another natural model for random tunnel number one 3–manifolds is to create them using the mapping class group of the surface ∂H . More precisely, fix a finite generating set $S = \{\psi_1, \psi_2, \dots, \psi_n\}$ of the mapping class group $\mathcal{MCG}(\partial H)$; for instance, take S to be the standard five Dehn twists. Fix also a non-separating simple closed curve γ_0 on ∂H . Now given r , create a sequence $\phi_1, \phi_2, \dots, \phi_r$ by picking each ϕ_i at random from among the elements of S and their inverses. Then set

$$\gamma = \phi_r \circ \phi_{r-1} \circ \dots \circ \phi_1(\gamma_0),$$

and consider the corresponding tunnel number one manifold M . That is, we start with γ_0 and successively mess it up r times by randomly chosen generators. Equivalently, we go for a random walk in the Caley graph of $\mathcal{MCG}(\partial H)$, and then apply the endpoint of that walk to γ_0 to get γ . [Question 1.1](#) now becomes: what is the probability that M fibers over the circle if r is large? A priori, the answer could be different from the one given in [Theorem 2.4](#). Because the number of such manifolds is countably infinite, there is no canonical probability measure on this set, so our choice of model for a random manifold is important. One might hope that all “reasonable” models give the same answer, but it should be emphasized that in some ways our two notions are fundamentally different. In any event, we will provide compelling experimental evidence for the following conjecture.

1.2 Conjecture *Let M be a tunnel number one 3–manifold created by a random walk in $\mathcal{MCG}(\partial H)$ of length r . Then the probability that M fibers over the circle goes to 0 as $r \rightarrow \infty$.*

Thus from this alternate point of view as well, it seems that nearly all tunnel number one 3–manifolds do not fiber over the circle.

1.3 Random groups

One of the fundamental tasks of 3–dimensional topology is to understand the special properties of their fundamental groups, as compared to finitely presented groups in general. From the point of view of this paper there is a surprising contrast between

these two classes of groups. While the question of whether a 3–manifold M fibers over the circle might seem fundamentally geometric, Stallings showed that it can be reduced to an algebraic question about $\pi_1(M)$ (see [Section 4](#)). For a group G , let us say that G *fibers* if there is an automorphism ρ of a free group B so that G is the algebraic mapping torus:

$$\langle t, B \mid tbt^{-1} = \rho(b) \text{ for all } b \in B \rangle$$

If M has tunnel number one, then it fibers over the circle if and only if $\pi_1(M)$ fibers in this sense ([Corollary 4.3](#)). When M has tunnel number one, its fundamental group is constructed from the free group $\pi_1(H)$ by killing the attaching curve γ of the 2–handle. Thus the fundamental group is just

$$\pi_1(M) = \langle \pi_1(H) \mid \gamma = 1 \rangle = \langle a, b \mid R = 1 \rangle,$$

that is, a 2–generator, 1–relator group.

In the spirit suggested above, we would like to compare [Theorem 2.4](#) with the situation for a random group G of the form $\langle a, b \mid R = 1 \rangle$. A natural meaning for the latter concept would be to consider the set $\mathcal{G}(r)$ of all such groups where the length of the relator R is less than r . This notion is in fact almost precisely analogous to the setup of $\mathcal{T}(r)$ for manifolds; in particular if $M \in \mathcal{T}(r)$, then the natural presentation of $\pi_1(M)$ is in $\mathcal{G}(r)$. Yet the remarkable thing is that the probability that $G \in \mathcal{G}(r)$ fibers experimentally tends to about 0.94 as $r \rightarrow \infty$. While we can't prove this, we can at least show:

Theorem 6.1 *Let p_r be the probability that $G \in \mathcal{G}(r)$ fibers. Then for all large r one has*

$$0.0006 \leq p_r \leq 0.975.$$

In particular, p_r does not limit to 0 as $r \rightarrow \infty$, in marked contrast to [Theorem 2.4](#).

As we will explain later, whether or not $G = \langle a, b \mid R = 1 \rangle$ fibers depends on the combinatorics of the relator R in a certain geometric sense. The different behavior for 3–manifold groups comes down to the fact that the curve γ is an *embedded* curve on the genus 2 surface ∂H , and this gives the relator R a recursive structure where certain “syllables” appear repeatedly at varying scales. Compare [Figure 1](#) with [Figure 2](#).

1.4 Algorithms and experiment

The original motivation for [Theorem 2.4](#), as well as the basis of [Conjecture 1.2](#), was the results of computer experiments. While there is an algorithm which decides if a




Figure 1: Two random words in the free group $F = \langle a, b \rangle$. Here a word is plotted as a walk in the plane, where a corresponds to a unit step in the positive x -direction, and b a unit step in the positive y -direction. Thicker lines indicate points transversed multiple times. The relevance of these pictures will be made clear in [Section 5](#).

general 3-manifold fibers using normal surface theory (see Schleimer [32, Section 6], Tollefson and Wang [35] or Jaco and Tollefson [16]), this is not practical for all but the smallest examples. However, special features allow one to rapidly decide if a tunnel number one 3-manifold fibers over the circle. In particular, we will show that it is possible to decide if $M \in \mathcal{T}(r)$ fibers in time which is polynomial in $\log(r)$. Our algorithm is important not just for the experimental side of this paper, but also the theoretical; it forms the basis for [Theorem 2.4](#). Indeed, the basic approach of the proof is to analyze the algorithm and show that it reports “does not fiber” with probability tending to 1 as $r \rightarrow \infty$.

Because the fundamental group of a tunnel number one manifold is so simple, one can use a criterion of Ken Brown to determine if $\pi_1(M) = \langle a, b \mid R = 1 \rangle$ fibers in the above algebraic sense. Brown’s criterion is remarkably elegant and simple to use, and is in terms of the combinatorics of the relator R . If R is not in the commutator subgroup and has length r , it takes time $O(r \log r \log \log r)$. While that may seem quite fast, some of the manifolds we examined had $r = 10^{3000}$. It’s not even possible to store the relator R in this case; after all, the number of elementary particles in the observable universe is well less than 10^{100} !

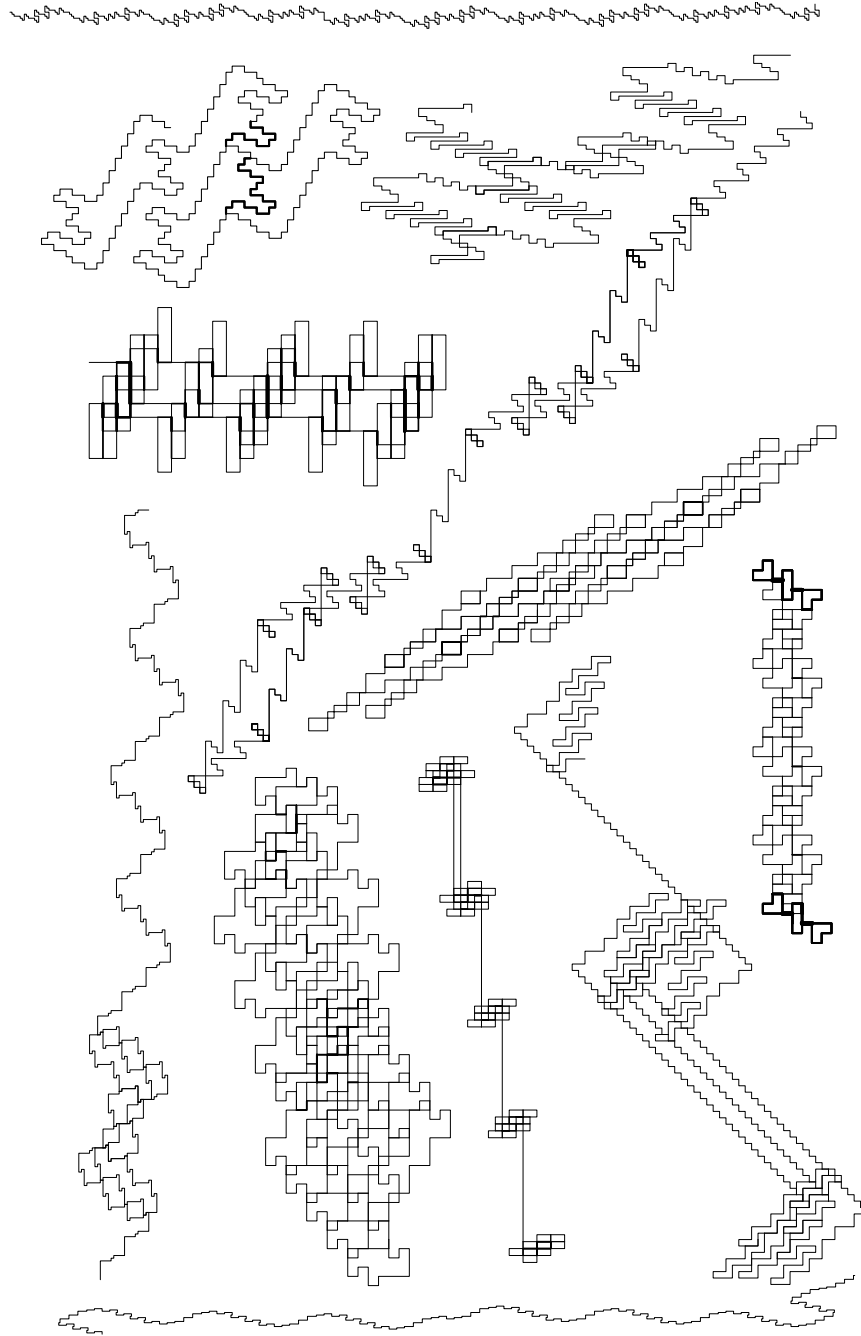


Figure 2: Relators of tunnel number one 3-manifolds, plotted in the style of [Figure 1](#). To conserve space, they are not all drawn to the same scale.

However, in the 3-manifold situation it is possible to specify R by giving the attaching curve γ , and γ can be described with only $\log(r)$ bits using either Dehn–Thurston coordinates or weights on a train track. Agol, Hass and W Thurston described [1] how to use splittings of train tracks to compute certain things about γ rapidly, eg, checking whether γ is connected, or computing its homology class. Motivated by their work, we were able to adapt Brown’s algorithm to work in this setting. The resulting algorithm uses train tracks which are labeled by “boxes” that remember a small amount of information about a segment of R . As mentioned, it can decide if $M \in \mathcal{T}(r)$ fibers in time polynomial in $\log(r)$.

1.5 The tyranny of small examples

Detailed results of our experiments are given in Section 3, and we will just highlight one aspect here. With both notions of a random tunnel number one manifold, it appears that the probability of fibering goes to 0 as the complexity increases; however, the rate at which it converges to 0 is actually quite slow, excruciatingly so in the \mathcal{MCG} context. In particular, most “small” manifolds fiber for pretty generous definitions of “small”.

Starting with the measured lamination notion, the first r for which $M \in \mathcal{T}(r)$ is *less* likely to fiber than not is about $r = 100,000$ (recall here that r is essentially the length of the relator R in the presentation of $\pi_1(M)$). The probability of fibering does not drop below 10% until about $r = 10^{14}$.

For the mapping class group version, we used the standard 5 Dehn twists as generators for $\mathcal{MCG}(\partial H)$ (Birman [4, Theorem 4.8]). Recall that the notion here is that given N , we apply a random sequence of N of these Dehn twists to a fixed base curve γ_0 to get the attaching curve γ . To get the probability of fibering to be less than 50%, you need to do $N \approx 10,000$ Dehn twists; to get it below 10% you need to take $N \approx 40,000$. It’s important to emphasize here how the \mathcal{MCG} notion relates to the measured lamination one, as it’s the later that is related to the size of the presentation for $\pi_1(M)$. For the \mathcal{MCG} notion, the length r of the relator experimentally increases *exponentially* in N . In particular, doing $N = 10,000$ Dehn twists gives manifolds in $\mathcal{T}(10^{500})$, and $N = 40,000$ gives manifold in $\mathcal{T}(10^{1750})$!

The moral here is that typically the manifolds that one can work with computationally (eg with SnapPea [37]) are so small that it is not possible to discern the generic behavior from experiments on that scale alone. For instance, about 90% of the cusped manifolds in the census of Callahan, Hildebrand and Weeks [8] are fibered (Button [7]), and most of these manifolds have tunnel number one. Without the naive version of Brown’s criterion, one would not be able to examine enough manifolds to suggest Theorem 2.4; without our improved train track version, we would not have come to the correct

version of [Conjecture 1.2](#). Indeed, initially we did experiments in the MCG case using just the naive version of Brown’s algorithm, and it was clear that the probability of fibering was converging to 1, not 0; this provoked much consternation as to why the “answer” differed from the measured lamination case. Thus one must always keep an open mind as to the possible generic behavior when examining the data at hand.

1.6 More general 3–manifolds

An obvious question that all of this presents is: What about [Question 1.1](#) for 3–manifolds which do not have tunnel number one? While there are certainly analogous notions of random manifolds for larger Heegaard genus, closed manifolds, etc., we don’t see the way to any results in that direction. Unfortunately, the method we use here is based fundamentally on Brown’s criterion, which is very specific to this case. Without this tool, it seems daunting to even try to gather enough experimental evidence to overcome the skepticism bar set by the discussion in [Section 1.5](#). However, our intuition is that for any Heegaard splitting based notion of random, the answer would remain unchanged: 3–manifolds should fiber with probability 0. For other types of models, such as random triangulations, the situation is murkier.

However, there is one generalization of [Theorem 2.4](#) that we can do. Recall that we chose to discuss tunnel number one 3–manifolds where the attaching curve γ is non-separating. If instead we look at those where γ is separating, we get manifolds M with two torus boundary components. In this case $H^1(M; \mathbb{Z}) = \mathbb{Z}^2$, and this gives us infinitely many homotopically distinct maps $M \rightarrow S^1$, any one of which could be a fibration. Thus it is perhaps surprising that the behavior here is no different than the other case:

Theorem 2.5 *Let T^s be the set of tunnel number one manifolds with two boundary components. Then the probability that $M \in T^s(r)$ fibers over the circle goes to 0 as $r \rightarrow \infty$.*

If we again compare this result to random 1–relator groups, the behavior is likely even more divergent than in the non-separating case. In particular, we conjecture that for groups $\langle a, b \mid R = 1 \rangle$ where R is in the commutator subgroup, the probability of algebraically fibering is 1.

1.7 The Virtual Fibration Conjecture

As we said at the beginning, one motivation for [Question 1.1](#) is to provide insight into:

1.8 Virtual Fibration Conjecture (W Thurston) *Let M be an irreducible, atoroidal 3-manifold with infinite fundamental group. Then M has a finite cover which fibers over the circle.*

Unlike the basic Virtual Haken Conjecture, which just posits the existence of a cover containing an incompressible surface, there is much less evidence for this conjecture. It has proven quite difficult to find interesting examples of non-fibered manifolds which can be shown to virtually fiber, though some infinite classes of tunnel number one 3-manifolds are known to have this property (Leininger [20], Walsh [36]). Our work here certainly suggests that the Virtual Fibration Conjecture is likely to be much more difficult than the Virtual Haken Conjecture. One pattern observed above suggests that the following approach is worth pondering. As discussed in [Section 1.5](#), “small” examples are still quite likely to fiber despite [Theorem 2.4](#). Presuming that this pattern persists for higher Heegaard genus, one strategy would be to try to find covers which were “smaller” than the initial manifold in some sense. For example, one measure of smallness might be the maximum length of a relator in a minimal genus Heegaard splitting.

1.9 Dynamical ingredients to the proof [Theorem 2.4](#)

In this introduction, we will not say much about the proof of [Theorem 2.4](#). However, let us at least mention the two other main ingredients besides our adaptation of Brown’s algorithm in the context of splittings of train tracks. The first is a theorem of Mirzakhani [23] which says, in particular, that non-separating simple closed curves have positive density among all multicurves (see [Theorem 11.1](#) below). This lets us sample simple closed curves by sampling multicurves. The other concerns splitting sequences of genus 2 interval exchanges. For technical reasons, we actually work with interval exchanges rather than train tracks, as we indicated above. Given a measured lamination carried by an interval exchange τ , we can split it to get a sequence of exchanges carrying the same lamination. We prove that genus 2 interval exchanges are normal, that is, any splitting sequence that can occur does occur for almost all choices of initial measured lamination ([Theorem 10.4](#)). Our proof is a direct application of a normality criterion of Kerckhoff [19].

1.10 Outline of contents

[Section 2](#) contains a detailed discussion of the various notions of random tunnel number one 3-manifolds, and gives the precise setup for [Theorems 2.4](#) and [2.5](#). [Section 3](#) gives the experimental data, which in particular justifies [Conjecture 1.2](#). [Section 4](#)

just covers Stallings' theorem, which turns fibering into an algebraic question. Then [Section 5](#) discusses Brown's algorithm in its original form. [Section 6](#) is about random 2-generator 1-relator groups as mentioned in [Section 1.3](#) above. [Section 7](#) is about train tracks, interval exchanges, and our efficient version of Brown's algorithm in that setting. The rest of the paper is devoted to the proof of [Theorems 2.4](#) and [2.5](#). It starts with an outline of the main idea in [Section 8](#), where a proof is given for a indicative toy problem. [Section 9](#) is devoted to a certain "magic splitting sequence", which is one of the key tools needed. We then prove normality for genus 2 interval exchanges in [Section 10](#). Finally, [Section 11](#) completes the proof by a straightforward assembly of the various elements.

1.11 Acknowledgments

Dunfield was partially supported by the US National Science Foundation, both by grant #DMS-0405491 and as a Postdoctoral Fellow. He was also supported by a Sloan Fellowship, and some of the work was done while he was at Harvard University. Thurston was partially supported by the US National Science Foundation as a Postdoctoral Fellow. Most of the work was done while he was at Harvard University. The authors also thank Steve Kerckhoff for helpful conversations and correspondence, as well as the referee for their very careful reading of this paper and resulting detailed comments.

2 Random tunnel number one 3-manifolds

2.1 Random 3-manifolds

What is a "random 3-manifold"? Since the set of homeomorphism classes of compact 3-manifolds is countably infinite, it has no uniform, countably-additive, probability measure. However, suppose we filter the set of 3-manifolds by some notion of complexity where manifolds of bounded complexity are finite in number. Then we can consider limiting probabilities as the complexity goes to infinity. For instance, we could look at all 3-manifolds which are triangulated with less than n tetrahedra, and consider the proportion p_n which are hyperbolic. If the limit of p_n exists as $n \rightarrow \infty$, then it is a reasonable thing to call the limit the probability that a 3-manifold is hyperbolic. Of course, unless the property in question is true for only finitely many, or all but finitely many, 3-manifolds, the answer depends on the complexity that we choose. In other words, it depends on the *model* of random 3-manifolds. Nonetheless, if we just pick one of several natural models to look at, it seems worthwhile to consider these types of questions to get a better global picture of the topology of 3-manifolds. For more on different possible models, and random 3-manifolds in general, see the work of the first

author and W Thurston in [10]. Here, we focus on the special class of tunnel number one 3–manifolds because it is easy to determine whether they fiber over the circle. In the next subsection, we discuss this class of manifolds, and then give some natural notions of probability on it.

2.2 Tunnel number one 3–manifolds

Look at an orientable handlebody H of genus 2. Consider an essential simple closed curve γ on ∂H . Now one can build a 3–manifold M consisting of H and a 2–handle attached along γ ; that is, $M = H \cup (D^2 \times I)$ where $\partial D^2 \times I$ is glued to ∂H along a regular neighborhood of γ . A 3–manifold which can be constructed in this way is said to have *tunnel number one*. There are two kinds of tunnel number one 3–manifolds, depending on whether the attaching curve γ separates the surface ∂H . If γ is non-separating, then ∂M is a single torus; if it is separating, then ∂M is the union of two tori. When we want to emphasize the dependence of M on γ , we will denote it by M_γ .

There is a dual description of being tunnel number one, which makes the origin of the name clear. Consider a compact orientable 3–manifold M whose boundary is a union of tori. The manifold M has tunnel number one if and only if there exists an arc α embedded in M , with endpoints on ∂M , such that the complement of an open regular neighborhood of α is a handlebody. While there are clearly many 3–manifolds with tunnel number one, it’s worth mentioning one class with which the reader may already be familiar: the exterior of a 2–bridge knot or link in S^3 . In this case, the arc in question joins the top of the two bridges. In general, 3–manifolds with tunnel number one are a very tractable class to deal with, and much is known about them.

2.3 Measured laminations

Next, we describe our precise parameterization of the tunnel number one 3–manifolds from the measured laminations point of view. As above, let H be a genus 2 handlebody. Fix a pair of pants decomposition of ∂H combinatorially equivalent to the curves (α, δ, β) shown in Figure 3, so that each of the curves defining this decomposition bound discs in H .

We will use Dehn–Thurston coordinates to parameterize the possible attaching curves γ for the 2–handle. A *multicurve* is a disjoint collection of simple closed curves. Up to isotopy, any multicurve γ on ∂H is given by weights $w_\alpha, w_\delta, w_\beta \in \mathbb{Z}_{\geq 0}$ and twist parameters $\theta_\alpha, \theta_\delta, \theta_\beta \in \mathbb{Z}$. Here, the weights record the (minimal) number of intersections of γ with the curves (α, δ, β) , and the twists describe how the strands




Figure 3: The curve γ has weights $(1, 2, 2)$ and twists $(0, 1, -1)$ with respect to the Dehn–Thurston coordinates given by the curves (α, δ, β) .

of γ meet up across the across these curves (with respect to a certain dual marking). These coordinates are analogous to Fenchel–Nielsen coordinates on Teichmüller space; see Figure 3 for an example, and Penner and Harer [26, Section 1] or Luo and Stong [21, Section 2] for details. When one of the weights is 0, say w_α , then the twist θ_α is the number of parallel copies of α in γ ; thus in this case $\theta_\alpha \geq 0$. With this convention, Dehn–Thurston coordinates bijectively parameterize all multicurves, up to isotopy.

Now we let \mathcal{T} be the set of tunnel number one presentations defined by curves γ with the following restrictions with respect to our choice of Dehn–Thurston coordinates:

- (1) γ is a non-separating simple closed curve.
- (2) The weights $w_\alpha, w_\delta, w_\beta$ are > 0 .
- (3) Each twist θ_i satisfies $0 \leq \theta_i < w_i$.
- (4) $w_\delta \leq \min(2w_\alpha, 2w_\beta)$.

We now explain why we’re making these requirements. The second restriction removes some special cases, which are all unknots in lens spaces and $S^2 \times S^1$; these could be left in without changing the final theorem as they have asymptotic probability 0. The third restriction simply accounts for the fact that Dehn twists along (α, δ, β) extend over H ; thus any γ is equivalent to one satisfying (3).

The final restriction serves the following purpose. If we use the basis of $\pi_1(H)$ dual to the discs α, β , then condition (4) ensures that the word γ represents in $\pi_1(H)$ is cyclically reduced. This is because (4) is the same as saying that each time γ crosses δ it then intersects either α or β before intersecting δ again. It is not immediate that any γ is equivalent, under a homeomorphism of H , to one satisfying (4); this is the content of Lemma 2.11 below. The reason that we need to require (4) is to make the

core machinery of the proof work correctly; as such it is admittedly a tad artificial, and we strongly expect it is not actually needed. See [Conjecture 2.7](#) below.

Now, let $\mathcal{T}(r)$ be all elements of \mathcal{T} where $w_\alpha + w_\beta < r$; equivalently, the ones whose corresponding word in $\pi_1(H)$ has length $< r$. As there are only finitely many elements of $\mathcal{T}(r)$, it makes sense to talk about the probability that they fiber over the circle. One form of our main result is:

2.4 Theorem *Let \mathcal{T} be the set of tunnel number one manifolds described above. Then the probability that $M \in \mathcal{T}(r)$ fibers over the circle goes to 0 as $r \rightarrow \infty$.*

We can also consider the case of tunnel number one 3-manifolds with two torus boundary components; these correspond to choosing $\gamma \subset \partial H$ to be a *separating* simple closed curve, replacing condition (1). We will denote the corresponding set of manifolds \mathcal{T}^s . Each $M \in \mathcal{T}^s$ has $H^1(M; \mathbb{R}) = \mathbb{R}^2$, so they have many chances to fiber over the circle. Perhaps surprisingly, the behavior here is no different than the other case:

2.5 Theorem *Let \mathcal{T}^s be the set of tunnel number one manifolds with two boundary components. Then the probability that $M \in \mathcal{T}^s(r)$ fibers over the circle goes to 0 as $r \rightarrow \infty$.*

2.6 Mapping class group

The \mathcal{MCG} model of a random tunnel number one 3-manifold was completely defined in the introduction. In this subsection, we discuss how it differs from the measured lamination version, and what one would need to leverage [Theorem 2.4](#) into a proof of [Conjecture 1.2](#). You can skip this section at first reading, as it's a little technical, and may not make much sense if you haven't read the proof of [Theorem 2.4](#). The two models for the choice of the attaching curve γ can basically be thought of as different choices of measure on $\mathcal{PML}(\partial H)$, the space of projectivized measured laminations. In the measured lamination model, this measure is just Lebesgue measure on the sphere $\mathcal{PML}(\partial H)$, whereas for the \mathcal{MCG} model it is a certain harmonic measure that we describe below.

In proving [Theorem 2.4](#), we show that there is a certain open set $U \subset \mathcal{PML}(\partial H)$ so that $\mathcal{T} \cap U$ consists solely of γ so that M_γ is not fibered. The proof then hinges on showing that U has full Lebesgue measure in the part of $\mathcal{PML}(\partial H)$ defined by requirement (4) above. The first thing that one needs to generalize to the \mathcal{MCG} model is to prove a version of [Theorem 2.4](#) where we drop (4). In particular, for this it suffices to show:

2.7 Conjecture *There exists an open set $V \subset \mathcal{PML}(\partial H)$ of full Lebesgue measure, such that for every non-separating curve $\gamma \in V$ the manifold M_γ is not fibered.*

We are very confident of this conjecture; the proof should be quite similar to [Theorem 2.4](#) provided certain technical issues can be overcome.

Assuming this conjecture is true, what one needs to do to prove [Conjecture 1.2](#) is show that V has measure 1 with respect to the following measure. Fix generators for $MCG(\partial H)$ and a base curve γ_0 . Let W_n be the set of all words in these generators of length n . Let μ_n be the probability measure on $\mathcal{PML}(\partial H)$ which is the average of the point masses supported at $\phi(\gamma_0)$ for $\phi \in W_n$. Then we are interested in the weak limit μ of these measures, which is called the *harmonic measure* (Masur and Kaimanovich [18]). The question, then, is whether $\mu(V) \neq 1$. The relationship between μ and Lebesgue measure is not well-understood, and we don't know how to show that [Conjecture 1.2](#) follows from [Conjecture 2.7](#). In fact, we suspect that μ is mutually discontinuous with Lebesgue measure based in part on [Section 3](#), even though they do agree on V .

2.8 Other models

There are other ways we could choose γ than just the two detailed above. For instance, we could start with a one vertex triangulation of ∂H , and then flip edges in a quadrilateral to obtain a sequence of such triangulations. At the end of such a sequence of moves, select a edge in the final triangulation which is a non-separating loop and take that to be γ . Another approach would be to start with a pair of pants decomposition of ∂H , and then move along a sequence of edges in the pants complex. Then we would take one of the curves in the final decomposition as γ . We did some haphazard experiments for both these notations as well, enough to convince us that they also result in a probability of fibering of 0 and behave generally like the mapping class group experiments reported above. For moves in the pair of pants decomposition, there is a choice of how many Dehn twists to perform on average before changing the decomposition; as you increase the number of Dehn twists, the rate at which the manifolds fiber tends to increase.

Strictly speaking, we do not choose our manifolds at random from among all such manifolds with a given bound on complexity, but rather we chose from the collection of *descriptions* of bounded complexity. These are different as a manifold can have more than one such description. Focusing on the measured lamination point of view, there are two separate issues: first, a manifold can have more than one unknotting tunnel; second, having fixed an unknotting tunnel, there may be more than one $\gamma \in \mathcal{T}(r)$ describing it,

due to the action of the mapping class group of the handlebody. While we do not prove this here, we believe that, in the measured lamination case, choosing from descriptions is essentially equivalent to choosing from among manifolds, as follows. For the first issue, we strongly believe that a generic manifold in $\mathcal{T}(r)$ has a unique unknotting tunnel; in particular, we expect that the distance of the Heegaard splitting should be very large as $r \rightarrow \infty$. (Another reason why the number of unknotting tunnels is not a big concern is that this would only affect our answer if fibered manifolds had many fewer unknotting tunnels than non-fibered ones.) About the second issue, namely multiple descriptions of the same unknotting tunnel, we could further restrict the conditions (1–4) above on elements of $\mathcal{T}(r)$ to generically eliminate such multiple descriptions. As described by Berge [2], there are simple inequalities in the weights and twists which ensure that $w_\alpha + w_\beta$ is minimal among all curves equivalent under the action of the mapping class group of the handlebody. This minimal form is typically unique (up to obvious symmetries, the number of which is independent of the particular curve at hand). The exception is when there are what [2] calls “level T-transformations”; because the presence of such transformations is determined by a family of *equalities*, these occur only in an asymptotically negligible portion of $\mathcal{T}(r)$. Thus by supplementing (1–4) we could precisely parameterize pairs $(M, \text{unknotting tunnel})$. This change would make no difference in the proof of [Theorem 2.4](#).

In the case of the mapping class group setup, there is a third issue which is that there are many random walks in \mathcal{MCG} that end at the same element. One could instead work by choosing the elements in \mathcal{MCG} from larger and larger balls in the Cayley graph. This has two disadvantages. The first is that in the context of non-amenable groups such as this one, the study of random walks is probably more natural than the study of balls; eg, consider the rich and well-developed theory of the Poisson boundary (Kaimanovich [17]). The second is that it is no longer possible to generate large elements with this alternate distribution, making experiment impossible (particularly important since experiment is all we have in this case). Of course, different elements of \mathcal{MCG} may also result in the same manifold for the two reasons discussed in the measured lamination case. We expect that multiple representatives of the same curve should be quite rare since the subgroup of \mathcal{MCG} which extends over the handlebody is very small; in particular, it is of infinite index. Indeed, it is easy to show that the probability that a random walk lies in this subgroup goes to 0 as the length of the random walk goes to infinity.

2.9 Curve normal form

In this subsection, we justify the claim made in [Section 2.3](#) that given a simple closed curve $\gamma \subset \partial H$, there is a homeomorphism ϕ of the whole handlebody H so that $\phi(\gamma)$

satisfies condition (4) of Section 2.3. Equivalently, we want to find curves (α, δ, β) , arranged as in Figure 3, which satisfy

(2.10) Any subarc of γ with endpoints in $\delta \cap \gamma$ intersects $\alpha \cup \beta$.

The rest of this section is devoted to:

2.11 Lemma *Let γ be a simple closed curve on the boundary of a genus 2 handlebody H . Then we can choose $(\alpha, \delta, \beta) \subset \partial H$ bounding discs in H as above so that (2.10) holds.*

This lemma is due to Masur [22], and was also described in a much more general form by Berge [2]. The proof of the lemma is used in the algorithm for the \mathcal{MCG} case, so as the lemma is not explicitly set out in [22], and [2] is unpublished, we include a proof for completeness. You can certainly skip it at first reading.

Proof We focus on choosing α and β to make the picture as standard as possible; the right choice for δ will then be obvious. First, choose α and β to be essential non-separating, non-parallel curves that minimize the size of $\gamma \cap (\alpha \cup \beta)$. Split H open along the discs bounded by α and β to get a planar diagram as shown in Figure 4(a). Here the labeled circles, called *vertices*, correspond to the discs we cut along, and the arcs are the pieces of γ . Note vertices with the same label are the endpoints of an equal number of arcs, since these endpoints match up when we reglue to get H .

We will show that the picture can be made very similar to the one shown in Figure 4(a); then the δ shown in Figure 4(b) works to complete the proof. In particular, it is enough to show:

- (1) No arc joins a vertex to itself.
- (2) All arcs joining a pair of vertices are isotopic to each other in the complement of the other vertices.




Figure 4





Figure 5

First, suppose we do not have (1), with V_0 being the vertex with the bad arc γ_0 . Consider $V_0 \cup \gamma_0$, which separates S^2 into two regions. Both of these regions must contain a vertex, or we could isotope γ to remove an intersection with $\alpha \cup \beta$. Focus on the component which contains only one vertex V_1 . If V_0 and V_1 have the same label as shown in Figure 5(a), we have a contradiction as the V_i must be the endpoints of the same number of arcs. So we have the situation shown in Figure 5(b). Replacing α with the non-separating curve ϵ indicated reduces $\gamma \cap (\alpha \cup \beta)$, contradicting our initial choice of α and β .

For (2), there are two basic configurations, depending on whether the non-parallel arcs join vertices with the same or opposite labels:



Here, parallel arcs have been drawn as one arc; the label on that arc refers to the number of parallel copies (which may be 0). In the case at left, the gluing requirement forces

$$u + w + x + z = x + y \quad \text{and} \quad v + w + y + z = u + v,$$

which easily leads to a contradiction.

In the case at right, we must have $x = y$ and $u = v$ or else we can replace α or β by a handle slide in the spirit of Figure 5(b) to reduce $\gamma \cap (\alpha \cup \beta)$. Now reglue the α discs to get a solid torus. Looking at one of the β vertices, it is joined to α by two families of parallel arcs as shown in Figure 6(a).

Thinking of this vertex as a bead, slide it along the set of parallel γ strands past the curve α . Keep sliding past α in the same direction if possible. Either:




Figure 6

- This eventually results in an arc joining the pair of β vertices. In this case, there will also be an arc joining the pair of α vertices left over from the final bead slide. Because of these two arcs, we can't have non-parallel γ arcs joining vertices of the same type, which ensures (2).
- The bead returns to where it started, so we have something like Figure 6(b). The other β vertex must be in the same situation, running along a parallel curve on the solid torus. As γ is connected, there are no arcs not involved in the β vertex tracks. Thus after further sliding, we can make the picture completely standard, with the two β vertices next to each other. This situation satisfies (2) as well.

Since we have shown α and β can be chosen so that (1) and (2) hold, we are done. \square

3 Experimental results

In this section, we give the results of our computer experiments using the algorithm of Section 7. We begin with the measured lamination notion of random. For each fixed r , we sampled about 100,000 manifolds $M \in \mathcal{T}(r)$, and used the algorithm to decide if each one fibers. Below in Figure 7 are the results for various $r \leq 10^{20}$.

While these results are superseded by Theorem 2.4, there are still interesting things to notice about the plots. For instance, look at the *rate* at which the probability of fibering approaches 0; as we already discussed in Section 1.5, it is quite leisurely. Moreover, the convergence has a very specific form — as the log-log plot in the bottommost part of Figure 7 makes clear, it converges to 0 like $c_1 e^{-c_2 r}$ for some positive constants c_i . In the proof of Theorem 2.4, we will see why this should be the case.

Before moving on to the \mathcal{MCG} case, let us make one quick comment on why we can easily sample $M \in \mathcal{T}(r)$ uniformly at random. While it is easy to pick a random multicurve with $w_\alpha + w_\beta \approx r$, a priori there is no way to ensure that we sample only connected non-separating curves. Fortunately, Mirzakhani has shown that there is a

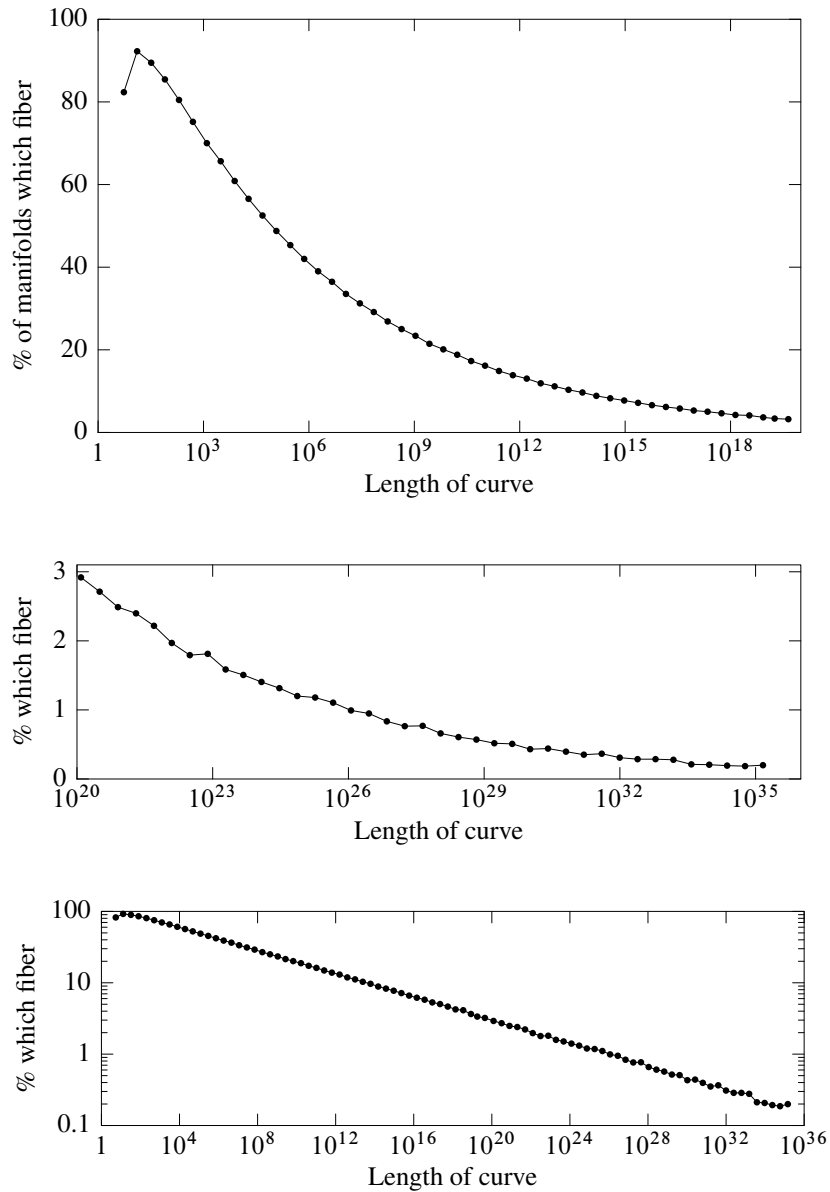


Figure 7: Data for the probability of fibering from the measured lamination point of view. The horizontal axis is the size r of the curve γ in Dehn–Thurston coordinates, or equivalently the length of the relator in the resulting presentation of π_1 . Each point represents a sample of about 100,000 manifolds.

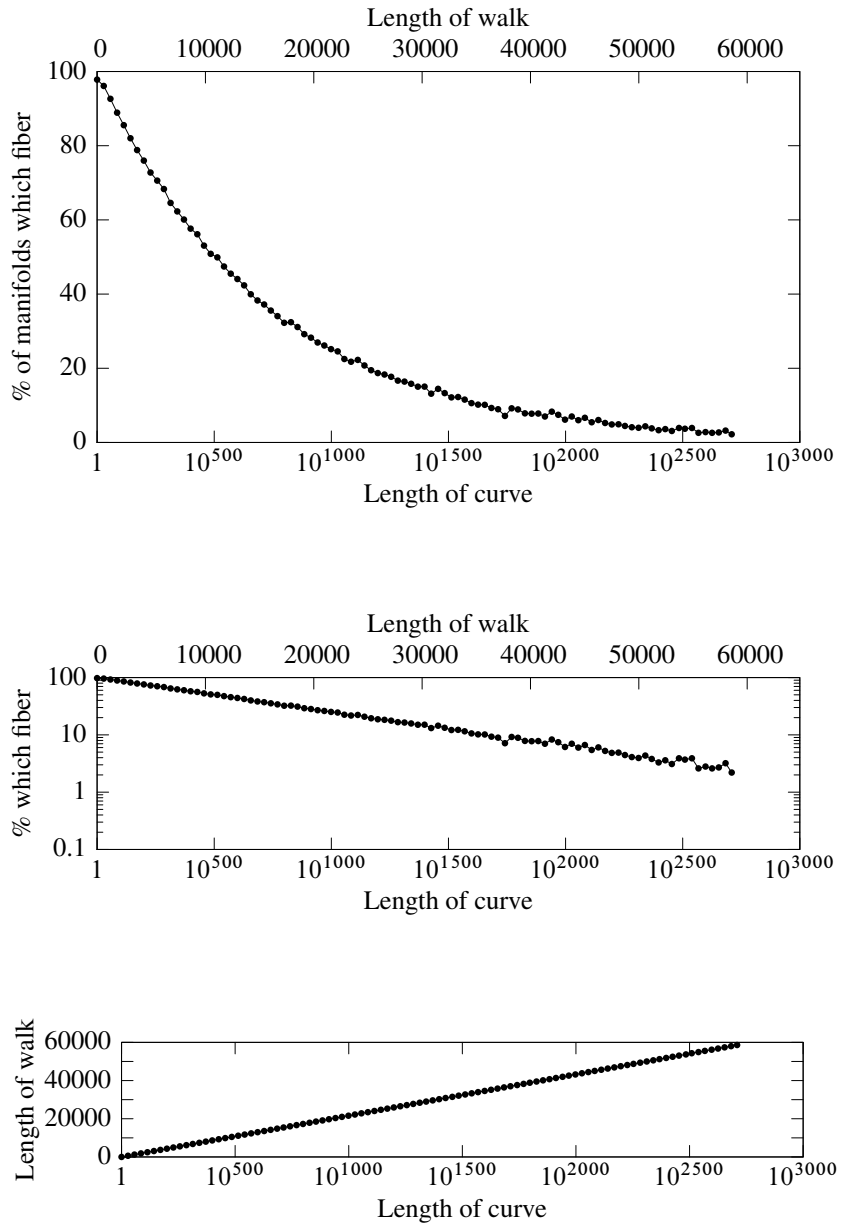


Figure 8: Data for the MCG case. All the points represent samples of at least 1000 manifolds. The first half or so of the data represent 10,000 manifolds.

definite probability, roughly 1/5, that a randomly chosen multicurve is of this form (see [Theorem 11.1](#) below). Thus one simply samples multicurves at random, ignoring all of those which are not of the desired form.

We turn now to the \mathcal{MCG} case. For this, we choose the standard five Dehn twists as our generating set for $\mathcal{MCG}(\partial H)$. The results are shown in [Figure 8](#). There are two horizontal scales on each of the upper two plots. Along the top is the number of Dehn twists done to create the manifold, ie, the length of the walk in $\mathcal{MCG}(\partial H)$. To give a scale at which to compare it to the previous figure, along the bottom is the size of the resulting attaching curve γ in terms of the standard Dehn–Thurston coordinates. As the plot at the bottom shows, the Dehn–Thurston size grows exponentially in length of the walk, which justifies the use of the two scales on the upper graphs.

One thing to notice here is just how slowly the probability goes to zero in terms of the Dehn–Thurston size; in the earlier [Figure 7](#), the probability of fibering was less than 0.3% for $r = 10^{35}$, but here the probability is still greater than 40% at $r = 10^{500}$. This reinforces the point made in [Section 2.6](#) that the γ resulting from the \mathcal{MCG} process are not generic with respect to Lebesgue measure on $\mathcal{PML}(\partial H)$.

Because of how large some of these curves are, we had to use much smaller samples than in the earlier case; this is why the graph looks so jumpy. However, if we look at the middle plot, we again see near perfect exponential decay, just as in the measure lamination case. Thus we are quite confident that [Conjecture 1.2](#) is correct.

3.1 Fibering in slices of $\mathcal{PML}(\partial H)$

The parameter space \mathcal{T} of tunnel number one 3-manifolds is a subset of $\mathcal{ML}(\partial H; \mathbb{Z})$. Let us projectivize, and so view \mathcal{T} as a subset of $\mathcal{PML}(\partial H; \mathbb{R})$, which is just the 5-sphere. If we take a two dimensional projectively linear slice of $\mathcal{PML}(\partial H; \mathbb{R})$, we can plot the fibered points of \mathcal{T} in the following sense. Fix some positive number r . Divide the slice into little boxes, and in each box pick a random $\gamma \in \mathcal{T}(r)$ and plot whether or not M_γ fibers. Of course, as $r \rightarrow \infty$ the probability of fibering goes to zero, so it is much more informative to plot *how many steps* the algorithm takes before it reports “not fibered”. [Figure 9](#) shows the results for one such slice, where we fixed $w_\alpha \approx w_\beta \approx (2/3)w_\delta \approx 2\theta_\delta$, and took $r \approx 10^{11}$. The horizontal and vertical axes are θ_α and θ_β ; since they are well defined modulo w_α and w_β , the figure should be interpreted as living on the torus.

3.2 Knots in S^3

As mentioned in the introduction, we suspect that the pattern exhibited above should persist for any Heegaard splitting based model of random manifold, regardless of genus.




Figure 9: A slice of $\mathcal{PML}(\partial H; \mathbb{R})$, with curves where M_γ is fibered plotted in white, and where M_γ is non-fibered plotted in shades of gray; darker grays indicate that the program reported “non-fibered” in fewer steps.

It is less clear what would happen for, say, random triangulations. We did a little experiment for knots in S^3 , where one filters (isomorphism classes of) prime knots by the number of crossings. Rather than address the difficult question of whether they fiber, we looked instead at whether the lead coefficient of the Alexander polynomial is ± 1 , ie, the polynomial is monic. A monic Alexander polynomial is necessary for fibering, but not sufficient. For alternating knots, however, it is sufficient (Murasugi [24]), and for non-alternating knots with few crossings there are probably not many non-fibered knots with monic Alexander polynomials. The results are shown in Figure 10. We used the program Knotscape [13] with knot data from Hoste, Thistlethwaite and Weeks [14] and Rankin, Flint and Schermann [28; 29]. In light of Section 1.5, we do not wish to draw any conclusions from this data. Really, what needs to be done is to figure out how to generate a random prime knot with, say, 100 crossings with close to the uniform distribution. It would be quite interesting to do so even for alternating knots; for this case, a place to start might be Poulalhon, Schaeffer and Zinn-Justin [27; 30].

3.3 Implementation notes

The algorithm used in the experiments is the one described in Section 7, though some corners were cut in the implementation of the Dehn twist move; so our resulting program is not completely efficient in the sense of that section. In the \mathcal{MCG} case,

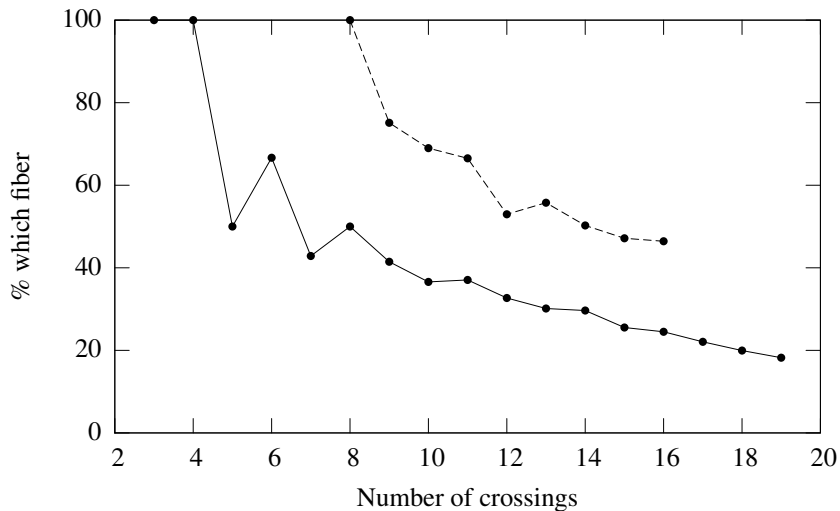


Figure 10: Proportion of prime knots whose complements fiber. The lower line plots alternating knots, and the upper one non-alternating knots. The vertical axis is really the proportion with non-monic Alexander polynomial; as such it is an *upper* bound on the proportion that fiber.

we had to deal with the fact that the resulting attaching curve γ might not be in \mathcal{T} ; that is, it fails to satisfy conditions (2–4) of Section 2.3. To rectify this, we used the method of the proof of Lemma 2.11 to get an equivalent curve in \mathcal{T} . The complete source code for our program can be downloaded from the front page for this paper: DOI:10.2140/gt.2006.10.2431

4 Algebraic criterion for fibering

In the next two sections, we describe how to determine if a tunnel number one 3-manifold fibers over the circle. There is an exponential-time algorithm from normal surface theory for deciding such questions in general [32], but that is impractical for our purposes. The criterion we give below is purely combinatorial in terms of the word in the fundamental group of the handlebody given by the attaching curve of the 2-handle. In this section, we give a theorem of Stallings which reduces this geometric question to an algebraic one about the fundamental group. In the next section, we describe an algorithm of K. Brown which then completely solves the algebraic question for the special type of groups coming from tunnel number one 3-manifolds.

4.1 Stallings' Theorem

Suppose a 3–manifold M fibers over the circle:

$$F \longrightarrow M \xrightarrow{\phi} S^1$$

The map $\phi: M \rightarrow S^1$ represents an element of $H^1(M; \mathbb{Z})$, which is Poincaré dual to the element of $H_2(M; \mathbb{Z})$ represented by the fiber F . Now, take M to be any compact 3–manifold. Continuing to think of a $[\phi] \in H^1(M; \mathbb{Z})$ as a homotopy class of maps $M \rightarrow S^1$, it makes sense to ask if $[\phi]$ can be represented by a fibration over S^1 . Associated to $[\phi] \in H^1(M; \mathbb{Z})$ is the infinite cyclic cover \tilde{M} of M whose fundamental group is the kernel of $\phi_*: \pi_1(M) \rightarrow \pi_1(S^1) = \mathbb{Z}$. If $[\phi]$ can be represented by a fibration, then \tilde{M} is just (fiber) $\times \mathbb{R}$. In particular, the kernel of ϕ_* is $\pi_1(\text{fiber})$, and hence finitely generated. The converse to this is also true:

4.2 Theorem (Stallings [33]) *Let M be a compact, orientable, irreducible 3–manifold. Consider a $[\phi] \neq 0$ in $H^1(M; \mathbb{Z})$. Then $[\phi]$ can be represented by a fibration if and only if the kernel of $\phi_*: \pi_1(M) \rightarrow \mathbb{Z}$ is finitely generated.*

The irreducibility hypothesis here is just to avoid the Poincaré Conjecture; it rules out the possibility that M is the connect sum of a fibered 3–manifold and a nontrivial homotopy sphere. When M has tunnel number one, the irreducibility hypothesis can be easily dropped without presuming the Poincaré Conjecture, as follows. Consider M as a genus 2 handlebody H with a 2–handle attached along $\gamma \subset \partial H$. By Jaco's Handle Addition Lemma (Jaco [15], Scharlemann [31]), M is irreducible if $\partial H \setminus \gamma$ is incompressible in H . If instead $\partial H \setminus \gamma$ compresses, then it is not hard to see that M is the connected sum of a lens space with the exterior of the unknot in the 3–ball. As lens spaces trivially satisfy the Poincaré Conjecture, we have:

4.3 Corollary *Let M^3 have tunnel number one. Then M fibers over the circle if and only if there exists a $[\phi] \neq 0$ in $H^1(M, \mathbb{Z})$ such that the kernel of $\phi_*: M \rightarrow \mathbb{Z}$ is finitely generated.*

In general, if G is a finitely presented group and $G \rightarrow \mathbb{Z}$ an epimorphism, deciding if the kernel is finitely generated is a very difficult question. Note that if H is a finitely presented group, then H is trivial if and only if the obvious epimorphism $H * \mathbb{Z} \rightarrow \mathbb{Z}$ has finitely generated kernel. Thus our question subsumes the problem of deciding if a given H is trivial, and hence is algorithmically undecidable. Thus, it is not at all clear that Stallings' Theorem can be leveraged to an algorithm to decide if a 3–manifold fibers. However, as we'll see in the next section, the algebraic problem is solvable in the case of a presentation with two generators and one relation, giving us a practical algorithm to decide if a tunnel number one 3–manifold fibers over the circle.

5 Brown’s Algorithm

Consider a two-generator, one-relator group $G = \langle a, b \mid R = 1 \rangle$. Given an epimorphism $\phi: G \rightarrow \mathbb{Z}$, Kenneth Brown gave an elegant algorithm which decides if the kernel of ϕ is finitely generated [6]. Brown was interested in computing the Bieri–Neumann–Strebel (BNS) invariant of G , which is closely related to this question. We will first discuss Brown’s algorithm for a fixed ϕ , and then move to the BNS context to understand what happens for all ϕ at once.

Let us explain Brown’s criterion with a geometric picture. Regard the group $G = \langle a, b \mid R = 1 \rangle$ as the quotient of the free group F on $\{a, b\}$. Think of F as the fundamental group of a graph Γ with one vertex and two loops. The cover $\tilde{\Gamma}$ of Γ corresponding to the abelianization map $F \rightarrow \mathbb{Z}^2$ can be identified with the integer grid in \mathbb{R}^2 ; the vertices of $\tilde{\Gamma}$ form the integer lattice $\mathbb{Z}^2 \subset \mathbb{R}^2$ and correspond to the abelianization of F . A homomorphism $\phi: F \rightarrow \mathbb{Z}$ can be thought of as a linear functional $\mathbb{R}^2 \rightarrow \mathbb{R}$. Now consider our relator R , which we take to be a cyclically reduced word in F . Let \tilde{R} be the lift of the word R to $\tilde{\Gamma}$, starting at the origin (see Figure 11). An epimorphism $\phi: F \rightarrow \mathbb{Z}$ descends to G if and only if $\phi(R) = 1$. Geometrically, this means that the kernel of ϕ is a line in \mathbb{R}^2 joining the terminal point of \tilde{R} to the origin. Turing this around, suppose R is not in the commutator subgroup of F so that the endpoints of \tilde{R} are distinct; in this case there is essentially only one ϕ , namely projection orthogonal to the line joining the endpoints. (To be precise, one should scale this projection so that ϕ takes values in \mathbb{Z} rather than \mathbb{R} , and is surjective.)

Now fix a ϕ which extends to G , and think of ϕ as a function on the lifted path \tilde{R} . Brown’s criterion is in terms of the number of global mins and maxes of ϕ along \tilde{R} . Roughly, $\ker(\phi) \leq G$ is finitely generated if and only if ϕ has the fewest extrema




Figure 11: The lift of the word $R = b^2 abab^{-1} ab^{-1} a^{-2}$ to the cover $\tilde{\Gamma}$.




Figure 12: At left is the lift of $R = b^2abab^{-1}ab^{-1}a^{-2}$ to $\tilde{\Gamma}$. The essentially unique homomorphism ϕ from $G = \langle a, b \mid R = 1 \rangle$ to \mathbb{Z} is indicated by the three diagonal lines which are among its level sets. In this case, ϕ has two global maxes and one global min. Hence the kernel of ϕ is *not* finitely generated. At right is another example with $R' = Ra$. While the words differ only slightly, in this case, the kernel of $\phi': \langle a, b \mid R' = 1 \rangle \rightarrow \mathbb{Z}$ is finitely generated as the global extrema of ϕ on \tilde{R}' are unique.

possible on \tilde{R} ; that is, it has only one global min and one global max. Figure 12 illustrates the two possibilities.

To be precise about Brown's criterion, one needs some additional conventions. First, extrema are counted with multiplicities: if \tilde{R} passes through the same point of $\tilde{\Gamma}$ twice and $\phi: \tilde{R} \rightarrow \mathbb{R}$ is maximal there, then this counts as two maxes. Also, the endpoints of $\tilde{\Gamma}$ can be extrema, and we include only one of them in our count. Finally, if the kernel of ϕ is horizontal or vertical, then there will be infinitely many global extrema; in this case we count unit length segments of extrema. To ensure that there is no ambiguity, we state Brown's theorem a little more combinatorially. For our relator word $R \in F$, let R_i denote the initial subword consisting of the first i letters of R . The value that ϕ takes on the i^{th} vertex of \tilde{R} is then $\phi(R_i)$.

5.1 Theorem [6, Theorem 4.3] *Let $G = \langle a, b \mid R = 1 \rangle$, where R is a nontrivial cyclically reduced word in the free group on $\{a, b\}$. Let R_1, \dots, R_n be initial subwords of R , where $R_n = R$. Consider an epimorphism $\phi: G \rightarrow \mathbb{Z}$.*

If $\phi(a)$ and $\phi(b)$ are both nonzero, then $\ker(\phi)$ is finitely generated if and only if the sequence $\phi(R_1), \dots, \phi(R_n)$ has a unique minimum and maximum. If one of $\phi(a)$ and $\phi(b)$ is zero, then the condition is that there are exactly 2 mins and maxes in the sequence and that R is not a^2 or b^2 .

The statement above is equivalent to our earlier geometric one; in the generic case, extrema of ϕ on \tilde{R} must occur at vertices.

We'll now briefly outline the proof of Brown's theorem in a way which elucidates its connections to the classical Alexander polynomial test for non-fibering of a 3-manifold. Given a two-generator, one-relator group G and an epimorphism $\phi: G \rightarrow \mathbb{Z}$, we can always change generators in the free group to express G as $\langle t, u \mid R = 1 \rangle$ where $\phi(t) = 1$ and $\phi(u) = 0$. The kernel of ϕ as a map from the free group $\langle t, u \rangle$ is (freely) generated by $u_k = t^k u t^{-k}$ for $k \in \mathbb{Z}$; this is because the cover of Γ corresponding to the kernel of ϕ is just a line with a loop added at each integer point. As $\phi(R) = 0$, we have

$$(5.2) \quad R = u_{k_1}^{\epsilon_1} u_{k_2}^{\epsilon_2} \cdots u_{k_n}^{\epsilon_n} \quad \text{where each } \epsilon_i = \pm 1.$$

The geometric condition of [Theorem 5.1](#) implies that the kernel of $\phi: G \rightarrow \mathbb{Z}$ is finitely generated if and only if the sequence k_1, k_2, \dots, k_n has exactly one max and min. The "if" part is elementary. For instance, suppose there is a unique minimum k_i , which we can take to be 0 by replacing R with $t^m R t^{-m}$. For any u_j , the relation $t^j R t^{-j}$ now implies that u_j can be expressed as a product of u_l 's with $l > j$. Similarly, a unique maxima allows us to express u_j as a product of u_l 's with strictly smaller indices. Thus the kernel of ϕ is generated by $u_{\min(k_i)}, \dots, u_{\max(k_i)}$. The "only if" direction is more subtle, and uses the fact that the relator in a one-relator group is in a certain sense unique.

We can now explain the promised connection to the Alexander polynomial. Let $\Delta(t)$ denote the Alexander polynomial associated to the cyclic cover corresponding to ϕ . Recall the classic test in the 3-manifold context is that if the lead coefficient of $\Delta(t)$ is not ± 1 (that is, $\Delta(t)$ is not *monic*), then ϕ cannot be represented by a fibration. Let us see why this is true for groups of the form we are looking at here. First notice that $\Delta(t)$ is just what you get via the formal substitution $u_k \mapsto t^k$ in (5.2), where multiplication is turned into addition (eg $u_2 u_1^{-2} u_2^{-1} u_0 \mapsto 1 - 2t$). Thus in the "fibered" case where the kernel of ϕ is finitely generated, we have that the lead coefficient of $\Delta(t)$ is indeed monic, as expected. Of course, $\Delta(t)$ can be monic and G still not be fibered. Essentially this is because $\Delta(t)$ is only detecting homological information; geometrically, if we look at the lift of R to the cover corresponding to the kernel of $\phi: \langle u, t \rangle \rightarrow \mathbb{Z}$, the issue is that the Alexander polynomial only sees the homology class of the lift of R , whereas Brown's criterion sees the whole lift. Thus you can regard Brown's test as a variant of the Alexander polynomial test that looks at absolute geometric information instead of homological information, and thereby gives an exact criterion for fibering instead of only a necessary one.

5.3 Remark One thing that is interesting to note about the proof sketch above is that when the kernel of ϕ is finitely generated, then in fact the group G is the mapping

torus (or HNN extension, if you prefer) of an automorphism of a free group. The free group in question here is just $u_{\min(k_i), \dots, u_{\max(k_i)-1}}$ (see [6, Section 4] for the details). When G is the fundamental group of a tunnel number one 3-manifold M , this makes sense as the fiber will be a surface with boundary, whose fundamental group is free.

5.4 BNS invariants

Let $G = \langle a, b \mid R = 1 \rangle$ be a two-generator one-relator group. To apply Stallings' [Theorem 4.2](#), we need to be able to answer this broader version of our preceding question: does there exist an epimorphism $\phi: G \rightarrow \mathbb{Z}$ with finitely generated kernel? So far, we just know how to answer this for a particular such ϕ . If the relator R is not in the commutator subgroup of the free group $F = \langle a, b \rangle$ then there is, up to sign, a unique such ϕ . So we only need to consider the case where $R \in [F, F]$; equivalently, the relator R lifts to $\tilde{\Gamma}$ as a closed loop. Now there are infinitely many ϕ to consider, as every $\phi: F \rightarrow \mathbb{Z}$ extends to G . Fortunately, the geometric nature of [Theorem 5.1](#) allows for a clean statement. It is natural to give the answer in term of Brown's original context, namely the Bieri–Neumann–Strebel (BNS) invariant of a group. This subsection is devoted to the BNS invariant and giving Brown's full criterion. The reader may want to skip ahead to [Section 5.9](#) at first reading; the current subsection will only be referred to in [Section 6.6](#) on random groups of this form. In particular, the main theorems about tunnel number one 3-manifolds are independent of it.

Let G be a finitely-generated group. Broadening our point of view to get a continuous object, consider nontrivial homomorphisms $\phi: G \rightarrow \mathbb{R}$. For reasons that will become apparent later, we will consider such ϕ up to *positive* scaling. Let $S(G)$ denote the set of all such equivalence classes; $S(G)$ is the sphere

$$S(G) = \left(H^1(G, \mathbb{R}) \setminus 0 \right) / \mathbb{R}^+.$$

The BNS invariant of G is a subset Σ of $S(G)$, which captures information about the kernels of the ϕ . Rather than start with the definition, let us give its key property (see [3; 6] for details).

5.5 Proposition *Let ϕ be an epimorphism from $G \rightarrow \mathbb{Z}$. Then the kernel of ϕ is finitely generated if and only if ϕ and $-\phi$ are both in Σ .*

To define Σ , first some notation. For $[\phi] \in S(G)$, let $G_\phi = \{g \in G \mid \phi(g) \geq 0\}$, which is a submonoid, but not subgroup, of G . Let G' denote the commutator subgroup of G , which G_ϕ acts on by conjugation. If H is a submonoid of G , we say that G' is

finitely generated over H if there is a finite set $K \subset G'$ such that $H \cdot K$ generates G' . Then the BNS invariant of G is

$$\Sigma = \{[\phi] \in S(G) \mid G' \text{ is finitely generated over some finitely generated submonoid of } G_\phi\}.$$

The BNS invariant has some remarkable properties—for instance, it is always an open subset of $S(G)$. When G is the fundamental group of a 3-manifold, Σ is symmetric about the origin and has the following natural description:

5.6 Theorem [3, Theorem E] *Let M be a compact, orientable, irreducible 3-manifold. Then Σ is exactly the projection to $S(G)$ of the interiors of the fibered faces of the Thurston norm ball in $H^1(M; \mathbb{R})$.*

In the BNS context, Brown’s Theorem 5.1 has the following reformulation:

5.7 Theorem [6, Theorem 4.3] *Let $G = \langle a, b \mid R = 1 \rangle$, where R is nontrivial and cyclically reduced. Let R_i be initial subwords of R and let $[\phi] \in S(G)$. If $\phi(a)$ and $\phi(b)$ are non-zero, then ϕ is in Σ if and only if the sequence $\phi(R_1), \dots, \phi(R_n)$ has a unique maximum. If one of $\phi(a)$ or $\phi(b) = 0$ vanishes, the condition is that there are exactly 2 maxes.*

Now consider the case when R is in the commutator subgroup so that $S(G)$ is a circle. To describe Σ , begin by letting \tilde{R} be the lift of the relator to $\tilde{\Gamma}$ thought of as a subset of $\mathbb{R}^2 = H_1(G; \mathbb{R})$. The focus will be on the convex hull C of \tilde{R} . For a vertex v of C , let F_v be the open interval in $S(G)$ consisting of ϕ so that the *unique* max of ϕ on C occurs at v . Geometrically, if we pick an inner product on $H_1(G; \mathbb{R})$ so we can identify it with its dual $H^1(G; \mathbb{R})$, then F_v is the interval of vectors lying between the external perpendiculars to the sides adjoining v . (Equivalently, we can think of the dual polytope $D \subset H^1(G; \mathbb{R})$ to C . Then F_v is projectivization into $S(G)$ of the interior of the edge of D dual to v .) We call a vertex of C *marked* if \tilde{R} passes through it more than once. Theorem 5.7 easily gives

5.8 Theorem [6, Theorem 4.4] *Let $G = \langle a, b \mid R = 1 \rangle$, where R is a nontrivial cyclically reduced word which is in the commutator subgroup. Then the BNS invariant Σ of G is*

$$\bigcup \{F_v \mid v \text{ is an unmarked vertex of } C\}$$

together with those ϕ whose kernels are horizontal or vertical if the edge of C where their maxima occur has length 1 and two unmarked vertices.




Figure 13: Here $R = b^{-1}a^2ba^{-1}b^{-1}aba^{-1}b^2a^{-1}b^{-1}a^{-1}b^{-1}aba^{-1}b^{-1}a$. At left is the convex hull C of \tilde{R} ; at right is BNS invariant Σ , shown as subset of the unit circle $S(G)$ in $H^1(G; \mathbb{R})$ with respect to the dual basis $\{a^*, b^*\}$.

A simple example is shown in Figure 13. The BNS picture can also be connected to the Alexander polynomial, in particular to the coefficients which occur at the vertices of the Newton polygon [9].

5.9 Boxes and Brown's Criterion

In this subsection, we show how to apply Brown's criterion by breaking up the relator R into several subwords, examining each subword individually, and then combining the information. This works by assigning what we call *boxes* to the subwords, together with rules for multiplying boxes. This is crucial for adapting Brown's criterion to efficiently incorporate the topological constraints when R is the relator for a tunnel number one 3-manifold. That said, the contents of this subsection apply indiscriminantly to any relator.


Let $F = \langle a, b \rangle$ be the free group on two generators. Let $R \in F$ be a cyclically reduced word, and $1 = R_0, R_1, R_2, \dots, R_n = R$ be the initial subwords. (Note that we are including $R_0 = 1$, which differs from our conventions earlier.) Suppose that $\phi: F \rightarrow \mathbb{Z}$ is an epimorphism with $\phi(R) = 0$. To apply Brown's criterion, we are interested in the sequence $\{\phi(R_i)\}$, and, in particular, in the number of (global) extreme values. We can think of $\{\phi(R_i)\}$ as a walk in \mathbb{Z} . Thus, we are lead to consider the set of finite walks \mathcal{W} on \mathbb{Z} which start at 0, where steps of any size are allowed, including pausing; that is, an element of \mathcal{W} is simply a finite sequence of integers whose first term is 0.

We now introduce *boxes* to record certain basic features of a walk $w \in \mathcal{W}$. In particular, we want to remember:

- The final position of w , which we call the *shift* and denote by s .

- The maximum value of w , called the *top* and denoted t .
- The minimum value of w , called the *bottom* and denoted b .
- The number of times the top is visited, denoted n^t . We count in a funny way: each time the top is visited counts twice, and we subtract one if the first integer is the top, and subtract one if the last integer is the top. Unless the walk is just $\{0\}$, this amounts to counting a visit at the beginning or end with a weight of 1 and all others with a weight of 2. We count in this way to make boxes well behaved under operations discussed below.
- The number of times the bottom is visited, counted in the same way, denoted n^b .

Abstractly, a *box* is simply 5 integers (s, t, b, n^t, n^b) satisfying $t \geq 0, b \leq 0, b \leq s \leq t$, and $n^t, n^b \geq 0$. Graphically, we denote boxes as




and the set of all boxes is denoted \mathcal{B} . We have a natural map $\text{Box}: \mathcal{W} \rightarrow \mathcal{B}$ implicit in our description above.

Given two walks $w_1, w_2 \in \mathcal{W}$ we can concatenate them into a walk $w_1 * w_2$ by translating all of w_2 so that its initial point matches the terminal point of w_1 , dropping the first element of the translated w_2 , and joining the two lists. For example

$$\{0, -1, 0, 1\} * \{0, 1, 0, 1\} = \{0, -1, 0, 1, 2, 1, 2\}$$

where the picture at right is in terms of the graphs of the walks (see Figure 14 if this is unclear). This operation makes \mathcal{W} into a monoid, with identity element the 1-element walk $\{0\}$. The set \mathcal{B} of boxes also has a monoid structure for which $\text{Box}: \mathcal{W} \rightarrow \mathcal{B}$ is a morphism. Pictorially, the box multiplication is given by



or algebraically by the following rule. For $i = 1, 2$, let B_i be the box $(s_i, t_i, b_i, n_i^t, n_i^b)$. Then $B_1 * B_2$ is, by definition, the box which has


- Shift $s = s_1 + s_2$.
- Top $t = \max(t_1, s_1 + t_2)$.
- Number of top visits

$$n^t = \begin{cases} n_1^t & \text{if } t_1 > s_1 + t_2, \\ n_2^t & \text{if } t_1 < s_1 + t_2, \\ n_1^t + n_2^t & \text{if } t_1 = s_1 + t_2. \end{cases}$$


and the corresponding rules for the bottom and n^b . The identity element for this monoid is the box

$$(0, 0, 0, 0, 0).$$

We can also reverse a walk w into another walk $\text{rev}(w)$ of the same length by translating by the negative of its last element and reversing the list. For example



This is an anti-automorphism of the monoid structure on \mathcal{W} . There is a corresponding anti-automorphism of the monoid structure on \mathcal{B} which is compatible with the map Box , given pictorially by



Algebraically, $\text{rev}(s, t, b, n^t, n^b)$ is the box with:

- Shift $s' = -s$.
- Top $t' = t - s$.
- Number of top visits $n^{t'} = n^t$.

and similarly for the bottom and number of bottom visits.

Let us now return to the setting of words in the free group F . Suppose w is a reduced word in F ; here w need *not* be cyclically reduced. Let $1 = w_0, w_1, \dots, w_n = w$ be

the initial subwords of w . Let ϕ be any epimorphism from F to \mathbb{Z} so that $\phi(a)$ and $\phi(b)$ are non-zero. (We will deal with the case when $\phi(a)$ or $\phi(b)$ is zero below.) We set

$$\text{Box}_\phi(w) = \text{Box}(\{\phi(w_i)\})$$

Now suppose that v is a reduced word in F so that the concatenation of w and v is also reduced. Then we have

(5.10)
$$\text{Box}_\phi(wv) = \text{Box}_\phi(w) * \text{Box}_\phi(v)$$

(5.11)
$$\text{Box}_\phi(w^{-1}) = \text{rev}(\text{Box}_\phi(w)).$$

The fact that there is no cancellation when we multiply w with v is important here; $\text{Box}_\phi: F \rightarrow \mathcal{B}$ is not a morphism. If we want to think of it as a morphism, we would need to take the domain to be the monoid of strings in $\{a^{\pm 1}, b^{\pm 1}\}$.

An alternate way to describe Box_ϕ is to give the values on the generators. If we assume that $s_1 = \phi(a) > 0$ and $s_2 = \phi(b) > 0$, then

$$\begin{aligned} \text{Box}(a) &= (s_1, s_1, 0, 1, 1) & \text{Box}(a^{-1}) &= (-s_1, 0, -s_1, 1, 1) \\ \text{Box}(b) &= (s_2, s_2, 0, 1, 1) & \text{Box}(b^{-1}) &= (-s_2, 0, -s_2, 1, 1) \end{aligned}$$

and Box is multiplicative on reduced words.

In the case when $\phi(a)$ is 0, we instead set

$$\begin{aligned} \text{Box}(a) &= (0, 0, 0, 2, 2) & \text{Box}(a^{-1}) &= (0, 0, 0, 2, 2) \\ \text{Box}(b) &= (1, 1, 0, 0, 0) & \text{Box}(b^{-1}) &= (-1, 0, -1, 0, 0) \end{aligned}$$

and extend by multiplicativity on reduced words. In this case n^t and n^b are twice the number of segments of extrema.

Now, let's restate Brown's criterion in terms of boxes. We say that the top (resp. bottom) of a $\text{Box}_\phi(w)$ is *marked* if $n^t > 2$ (resp. $n^b > 2$). Then [Theorem 5.1](#) can be restated as:

5.12 Theorem [6, Theorem 4.3] *Let $G = \langle a, b \mid R = 1 \rangle$, where R is a cyclically reduced word in the free group $\langle a, b \rangle$. Consider an epimorphism $\phi: G \rightarrow \mathbb{Z}$. Then $\ker(\phi)$ is finitely generated if and only if neither the top nor the bottom of $\text{Box}_\phi(R)$ are marked.*

Now we relate this to our original question of how to apply Brown's criterion by breaking R up into pieces. Suppose

$$R = w_1 w_2 w_3 \cdots w_k$$

where each w_i is reduced and the above product involves no cancellation to get R in reduced form. Then we have

$$(5.13) \quad \text{Box}_\phi(R) = \text{Box}_\phi(w_1) * \text{Box}_\phi(w_2) * \cdots * \text{Box}_\phi(w_k).$$

Notice that if B_1 and B_2 are two boxes with marked tops, then $B_1 * B_2$ also has a marked top. Hence, if it happens that each $\text{Box}_\phi(w_i)$ has a marked top, it follows that $\text{Box}_\phi(R)$ does as well without working out the product (5.13). This yields

5.14 Lemma *Let $G = \langle a, b \mid R = 1 \rangle$, where R is a cyclically reduced word, and consider an epimorphism $\phi: G \rightarrow \mathbb{Z}$. Suppose $R = w_1 w_2 w_3 \cdots w_k$ where each w_i is a reduced word, and the product has no cancellations. If each $\text{Box}_\phi(w_i)$ has a marked top, then the kernel of ϕ is infinitely generated.*

In light of the above lemma, it will be useful to have criteria for when a word has a marked top. The one we will need is based on the following simple observation: suppose B is a box with shift $s = 0$ and $n^t \geq 2$. Then $B * B$ has a marked top. To apply this, suppose $w \in F$ is a nontrivial cyclically reduced word with $\phi(w) = 0$; taking $B = \text{Box}_\phi(w)$, we claim that our observation implies that $\text{Box}_\phi(w^2) = B * B$ has a marked top. If neither $\phi(a)$ or $\phi(b)$ is 0, then it is easy to see that B has $n_t \geq 2$. If $\phi(a)$ vanishes, then there are words where B has $n^t = 0$, eg $w = b^{-1}ab$; however, any such word is *not* cyclically reduced. When w is cyclically reduced, n^t must be at least 2. This proves the claim that $\text{Box}_\phi(w^2) = B * B$ has a marked top. More generally

5.15 Lemma *Let F and $\phi: F \rightarrow \mathbb{Z}$ be as above. Suppose $w \in F$ is a nontrivial cyclically reduced word such that $\phi(w) = 0$. If w' is any subword of w^n of length at least twice that of w , then $\text{Box}_\phi(w')$ has a marked top.*

Proof By conjugating w , we can assume that $w' = w^2 r$ where r is an *initial* subword of w^n for some $n \geq 0$. We have $\text{Box}_\phi(w^2) = \text{Box}_\phi(w)^2$, and since $\phi(w) = 0$ and w is nontrivial, this implies that $\text{Box}_\phi(w^2)$ has a marked top. As r is a subword of w^n and $\phi(w) = 0$, the top of $\phi(w^2)$ forms part of the top of $\phi(w')$; hence $\text{Box}_\phi(w')$ has a marked top as well. \square

6 Random 1-relator groups

In this section, we consider the following natural notion of a random 2-generator 1-relator group. Let $\mathcal{G}(r)$ be the set of presentations $\langle a, b \mid R = 1 \rangle$ where the relator R

is a cyclically reduced word of length r . While properly the elements of $\mathcal{G}(r)$ are presentations, we will usually refer to them as groups. A random 2–generator 1–relator group of complexity r is then just an element of the finite set $\mathcal{G}(r)$ chosen uniformly at random. Now given any property of groups, consider the probability p_r that $G \in \mathcal{G}(r)$ has this property; we are interested in the behavior of p_r as $r \rightarrow \infty$. When p_r has a limit p , it is reasonable to say that “a random 2–generator 1–relator group has this property with probability p ”; of course, p is really an asymptotic quantity dependent on our choice of filtration of these groups, namely word length of the relator. An example theorem is that a 2–generator 1–relator group is word-hyperbolic with probability 1 (Gromov [12], Ol’shanskii [25]).

In analogy with the 3–manifold situation, we say that a group G *fibers* if it has an epimorphism to \mathbb{Z} with finitely generated kernel. As we noted in Remark 5.3, for these types of groups fibering is equivalent to being the mapping torus of an automorphism of a free group, which was the definition of fibered discussed in Section 1.3. This section is devoted to showing that for 2–generator 1–relator groups the probability of fibering is strictly between 0 and 1. In particular:

6.1 Theorem *Let p_r be the probability that $G \in \mathcal{G}(r)$ fibers. Then for all large r one has*

$$0.0006 < p_r < 0.975.$$

Experimentally, p_r seems to limit to 0.94. It seems quite remarkable to us that the probability a 2–generator 1–relator group fibers is neither 0 nor 1. In slightly different language, that most one relator groups fiber was independently discovered experimentally by Kapovich, Sapir, and Schupp [5, Section 1]; in that context, the proof of Theorem 6.1 shows that [5, Theorem 1.2] does *not* suffice to show that $G \in \mathcal{G}$ group is residually finite with probability 1.

Theorem 6.1 is strikingly different than the corresponding result (Theorem 2.4) for tunnel number one 3–manifolds; these fiber with probability 0. The setups of the two theorems are strictly analogous. Indeed, the parameter space $\mathcal{T}(r)$ of tunnel number one 3–manifolds is essentially just those $G \in \mathcal{G}(r)$ which are *geometric* presentations of the fundamental group of a tunnel number one 3–manifold. The differing results can happen because $\mathcal{T}(r)$ is a vanishingly small proportion of $\mathcal{G}(r)$ as $r \rightarrow \infty$; looking at Dehn–Thurston coordinates, it is clear that $\#\mathcal{T}(r)$ grows polynomially in r , whereas $\#\mathcal{G}(r)$ grows exponentially. Another example of differing behavior is word-hyperbolicity — because of the boundary torus, the groups in \mathcal{T} are almost never hyperbolic, whereas those in \mathcal{G} almost always are. Still, the different behavior with respect to fibering is surprising. As the proof of Theorem 2.4 will eventually make clear, the difference stems from the highly recursive nature of the relators of $G \in \mathcal{T}(r)$.

For tunnel number one manifolds with two boundary components, [Theorem 2.5](#) says that the probability of fibering is still 0, despite the fact that there are now many epimorphisms to \mathbb{Z} . In contrast, let $\mathcal{G}'(r)$ be those groups in $\mathcal{G}(r)$ whose defining relation is a commutator; then it seems very likely that:

6.2 Conjecture *Let p_r be the probability that $G \in \mathcal{G}'(r)$ fibers. Then $p_r \rightarrow 1$ as $r \rightarrow \infty$.*

We will explain our motivation for this conjecture in [Section 6.6](#).

6.3 A random walk problem

Let us first reformulate the question answered by [Theorem 6.1](#) in terms of random walks. This will suggest a simplified toy problem whose solution will make it intuitively clear why [Theorem 6.1](#) is true. We take the point of view of [Section 5.9](#), which runs as follows. Start with the free group $F = \langle a, b \rangle$ and an epimorphism $\phi: F \rightarrow \mathbb{Z}$. A word $R \in F$ gives 1-dimensional random walk $w = \{\phi(R_i)\}$ on \mathbb{Z} , where the R_i are the initial subwords of R . Assuming neither $\phi(a)$ or $\phi(b)$ is 0, Brown's Criterion is then that $\langle a, b \mid R = 1 \rangle$ fibers if and only if w visits its minimum and maximum value only once.

Unfortunately, from the point of view of the 1-dimensional walk w , things are a little complicated:

- (1) The walk w has two different step sizes, namely $\phi(a)$ and $\phi(b)$. Moreover, the condition that R is reduced means, for instance, that you aren't allowed to follow a $\phi(a)$ step by a $-\phi(a)$ step.
- (2) The walk must end at 0.
- (3) Worst of all, the step sizes themselves are determined by the relator, as it is the endpoint of R in the plane that determines ϕ in the first place. Thus one can't really remove the 2-dimensional nature of the problem.

To get a more tractable setup, let us consider instead walks on \mathbb{Z} where at each step we move one unit to the left or right with equal probability. Let $W(r)$ denote the set of such walks which both start and end at 0 (thus r must be even). For simplicity, let's just focus on the maxima. Then:

6.4 Proposition [\[11\]](#) *A walk $w \in W(r)$ visits its maximum value more than once with probability $1/2$.*




Figure 14

So the toy problem at least exhibits the neither 0 or 1 behavior of [Theorem 6.1](#). While the proposition is well known, we include a proof which, for our limited purposes, is more direct than those in the literature. The argument is also similar to what we will use for [Theorem 6.1](#) itself.

Proof We focus on the graph of a walk $w \in W(r)$, which we think of as a sequence of up and down segments (see [Figure 14](#)). Let U be those walks with a unique maximum. To compute the size of U , we relate it to the set D of walks which end on a down segment. Given $w \in U$ take the down segment immediately after the unique maxima, and shift it to the end to produce an element in D . This is a bijection; the inverse $D \rightarrow U$ is to move the final down to immediately after the *leftmost* maximum. Thus $\#U = \#D = (1/2)\#W(r)$, completing the proof. \square

If the toy problem was an exact model for [Theorem 6.1](#), we would expect the much lower value of $(1/2)^2 = 1/4$ for the probability of fibering, rather than the 0.94 that was experimentally observed. Next, we consider a slightly more accurate model, where the probability of a unique maxima rises. Consider the case where $\phi(a) = \phi(b) = 1$. Then condition (1) above becomes a momentum condition — at each step there is a $2/3$ chance of continuing in the same direction and a $1/3$ chance of turning and going the other way. Intuitively, this increases the chance of a unique max since it is less likely that a repeat max is created by a simple up-down-up-down segment. In this case we have:

6.5 Proposition *Consider random walks on \mathbb{Z} with momentum as described above. As the length of the walk tends to infinity, the probability of a unique maximum limits to $2/3$.*

Proof We will just sketch the argument, ignoring certain corner cases which are why the probability $2/3$ occurs only in the limit. The set of walks of length r is still

$W(r)$, unchanged from the previous proposition. What has changed is the probability measure P on $W(r)$ — it is no longer uniform. While we still have a bijection $f: U \rightarrow D$ as above, it is no longer measure preserving. Let U_d denote the set of walks with a unique max which end with a down, and U_u those that end with a up. Then $P(U_d) = P(f(U_d))$ whereas $P(U_u) = 2P(f(U_u))$. Also $f(U_d)$ consists of walks in D which end in *two* downs; thus $f(U_d)$ contributes $2/3$ of the measure of D , whereas $f(U_u)$ contributes only $1/3$. Combining gives $P(U) = 4/3P(D) = 2/3$ as desired. \square

Unfortunately, our approach seems to fail when we allow differing step sizes as in (1), even ignoring the momentum issue. The problem is that while the maps between U and D are still defined, they are no longer bijective. We turn now to the proof of [Theorem 6.1](#) which uses similar but cruder methods which have no hope of being sharp.

Proof As above, let $\mathcal{G}(r)$ be our set of 1-relator groups, which we will always think of as the set of cyclically reduced words R in $F = \langle a, b \rangle$ of length r . As a first step, we compute $\#\mathcal{G}(r)$. Counting reduced words, as opposed to cyclically reduced words, is easy: there are 4 choices for the first letter and 3 choices for each successive one, for a total of $4 \cdot 3^{r-1}$. What we need to find $\#\mathcal{G}(r)$ is the probability that a reduced word is cyclically reduced. Thinking of a reduced word w as chosen at random, the relationship between the final letter and the initial one is governed by a Markov chain whose distribution converges rapidly to the uniform one. Thus the distribution of the final letter is (nearly) independent of the first letter, and so the odds that w is cyclically reduced is $3/4$. Thus $\#\mathcal{G}(r)$ is asymptotic to 3^r . A more detailed analysis, not needed for what we do here, shows that $\#\mathcal{G}(r) = 3^r + 1$ when r is odd, and $3^r + 3$ when r is even.

Let $\mathcal{G}_0(r)$ denote those R which are not in the commutator subgroup, and so that the unique epimorphism $\phi: \langle a, b \mid R = 1 \rangle \rightarrow \mathbb{Z}$ does not vanish on either a or b . It is not hard to see that the density of $\mathcal{G}_0(r)$ in $\mathcal{G}(r)$ goes to 1 as $r \rightarrow \infty$. Thus in the remainder of the proof, we work to estimate the probability p'_r that $\langle a, b \mid R = 1 \rangle$ fibers for $R \in \mathcal{G}_0(r)$. To bound it from above, we construct an injection from $\mathcal{G}_0(r-4)$ into the *non-fibered* subset of $\mathcal{G}_0(r)$. In particular, given $R \in \mathcal{G}_0(r-4)$, go to the first global maximum and insert a commutator as shown in [Figure 15\(a\)](#). As we inserted a commutator, ϕ is unchanged, but we now have enough maxima to see that it is non-fibered. To see that this map is injective, observe that there is an inverse process: go to the first global maximum and delete the next 4 letters. Thus

$$1 - p'_r \geq \frac{\#\mathcal{G}_0(r-4)}{\#\mathcal{G}_0(r)} \approx 3^{-4} \quad \text{and hence } p'_r < 0.988 \text{ for large } r.$$




Figure 15

To improve this, note that we can also insert a commutator at the global minimum; the images of these two injections of $\mathcal{G}_0(r - 4)$ into the non-fibered words have some overlap coming from $\mathcal{G}_0(r - 8)$. Thus, pretending for convenience that $\#\mathcal{G}_0(n) = 3^n$, we have

$$1 - p'_r \geq 2 \cdot 3^{-4} - 3^{-8} \quad \text{and hence } p'_r < 0.975 \text{ for large } r,$$

as desired.

To estimate the number of fibered words in $\mathcal{G}_0(r)$, we inject $\mathcal{G}_0(r - 8)$ into them by inserting a commutator at both the first global min and the first global max as shown in Figure 15(b). In fact, there are four distinct ways of doing this, depending on which way we orient the two commutators. Thus

$$p'_r \geq 4 \frac{\#\mathcal{G}_0(r - 8)}{\#\mathcal{G}_0(r)} \approx 4 \cdot 3^{-8} > 0.0006 \quad \text{for large } r,$$

completing the proof. □

6.6 The case of \mathcal{G}'

We end this section by giving the motivation for Conjecture 6.2. Consider $G \in \mathcal{G}'(r)$. Every $\phi: F \rightarrow \mathbb{Z}$ extends to G . If we fix ϕ , the proof of Theorem 6.1 shows that the probability that $\ker(\phi)$ is finitely generated is at least 0.0006. The intuition is that if we fix two such epimorphisms ϕ and ϕ' , then as $r \rightarrow \infty$ the event that $\ker(\phi)$ is finitely generated becomes independent from the corresponding event for ϕ' . This should extend to any finite collection of ϕ_i , and as each event has probability at least 0.006, independence means that at least one of the ϕ_i will have finitely generated kernel with very high probability. Increasing the number of ϕ_i would allow one to show that the probability is at least $1 - \epsilon$ for any ϵ , and hence the probability limits to 1. The different ϕ_i should have independent behavior for the following reason. As described in Section 5.4, whether G fibers depends on the convex hull of the relator R .




Figure 16: Train tracks on a 4-punctured S^2 and on a torus.


In particular, at the vertices of the convex hull we care about whether R passes through them multiple times. Any given ϕ picks up its global extrema from some pair of these vertices. As long as the extreme vertices associated to ϕ and ϕ' are distinct, then whether they are repeated vertices should be independent. Thus the key issue is simply that the number of vertices on the convex hull of R should grow as $r \rightarrow \infty$. In the related questions that have been studied, the number of vertices grows like $\log(r)$, and we don't expect the situation here to be any different. See Steele [34] and the references therein for details.

7 Efficient implementation of Brown's algorithm

In this section, we discuss how to efficiently decide whether a given curve in our parameter space gives a fibered tunnel number one manifold. Here, “efficiently” means in time which is polynomial in the \log of the Dehn–Thurston coordinates. The method we present is also crucial to the proof of [Theorem 2.4](#).

7.1 Train tracks

Our main combinatorial tools for studying curves on surfaces are train tracks and their generalizations. Roughly, a *train track* in a surface Σ is a 1-dimensional CW complex τ embedded in Σ , which is made up of 1-dimensional branches joined by trivalent switches. Here, each switch has one incoming branch and two outgoing branches. See [Figure 16](#) for examples, and [\[26\]](#) for details. Associated to a train track τ is a *space of weights* (or transverse measures). This consists of assignments of weights $w_e \in \mathbb{R}_{\geq 0}$ to each branch e of τ , which satisfy the switch condition: at each switch the sum of the weights on the two outgoing edges is equal to the weight on the incoming one. The space of weights is denoted $\mathcal{ML}(\tau, \mathbb{R})$, and $\mathcal{ML}(\tau, \mathbb{Z})$ denotes those where each $w_e \in \mathbb{Z}$. As shown at right, an integral measure $w \in \mathcal{ML}(\tau, \mathbb{Z})$ naturally specifies a multicurve which lies in a small neighborhood of τ , that is, is *carried by* τ . More generally, $\mathcal{ML}(\tau, \mathbb{R})$






Figure 17: At left is an interval exchange τ on a 5-punctured S^2 . At right is a regular neighborhood of τ .

parameterizes measured laminations carried by τ . For suitable train tracks, called *complete* train tracks, $\mathcal{ML}(\tau, \mathbb{R})$ gives a chart on the space of measured lamination on the underlying surface Σ .

7.2 Interval exchanges

A more generalized notion of train tracks is to allow switches where there are an arbitrary number of incoming and outgoing branches. Here, we will focus on the class where there is just one switch. These are called *interval exchanges* for reasons we will see shortly. An example on a 5-punctured S^2 is shown in [Figure 17](#). As you can see from that figure, a regular neighborhood of such an interval exchange can be decomposed into a thickened interval (shaded) whose top and bottom are partitioned into subintervals which are exchanged by means of bands (the thickened branches). A $w \in \mathcal{ML}(\tau, \mathbb{R})$ can be thought of as assigning widths to the bands so that the total length of the top and bottom intervals agree.

There are two kinds of bands. Those that go from the top to the bottom are termed *orientation preserving* since that is how they act on their subintervals. Those joining a side to itself are called *orientation reversing*. An interval exchange gives rise to a natural dynamical system which we describe in the next subsection. In that context, they have been studied extensively since the 1970s. However, generally only orientation preserving bands are allowed in that literature; we will refer to such exchanges as *classical interval exchanges*.

For the rest of this section, one could easily work with train tracks instead of interval exchanges. However, the use of interval exchanges has important technical advantages in the proof of [Theorem 2.4](#).




Figure 18: Splitting an interval exchange.

7.3 Rauzy induction and determining connectivity

Suppose τ is an interval exchange in a surface Σ , and $w \in \mathcal{ML}(\tau, \mathbb{Z})$. We will now describe how to determine the number of components of the associated multicurve. This method is also the basis for our efficient form of Brown's algorithm. The basic operation is called Rauzy induction in the context of interval exchanges, and splitting or sliding in the context of train tracks. Starting with τ and $w \in \mathcal{ML}(\tau, \mathbb{R})$, we will construct a new pair (τ', w') realizing the same measured lamination in Σ . To begin, consider the rightmost bands, t and b , on the top and bottom respectively. First suppose that $w_t > w_b$. Then we slice as shown in Figure 18 to construct (τ', w') . The new band t' has weight $w'_{t'} = w_t - w_b$, and the other modified band b' has weight $w'_{b'} = w_b$. The other weights are of course unchanged. If instead $w_b < w_t$ one does the analogous operation, flipping the picture about the horizontal axis.

If $w_t = w_b$, then one simply cuts through the middle interval and amalgamates the bands b and t together. If t and b happen to be the same band, this splits off an annular loop. Hence we enlarge our notion of an interval exchange by allowing the addition of a finite number of such loops. We will denote Rauzy induction, which we usually call *splitting*, by $(\tau, w) \searrow (\tau', w')$.

Now suppose we start with an integral measure $w \in \mathcal{ML}(\tau, \mathbb{Z})$ and want to determine the number of components of the associated multicurve. We can split repeatedly to get a sequence (τ_i, w_i) carrying the same multicurve. Here, there is no reason to remember the (increasingly complicated) embeddings of the τ_i into Σ . That is, you should think of the τ_i as *abstract* interval exchanges, not embedded in any particular surface. At each stage, either some weight of w_i is strictly reduced, or we split off a loop and reduce the number of bands in play. In the end, we are reduced to a finite collection of

loops labeled by elements of $\mathbb{Z}_{\geq 0}$; the sum of these labels is the number of components in the multicurve.

The procedure just described is not always more efficient than the most naive algorithm; in particular the number of steps can be equal to $|w| = \max\{w_b\}$. For instance, take τ to be the exchange on the torus shown in [Figure 16](#) and set the weights on the bands to be 1 and n . However, Agol, Hass, and W. Thurston have shown that if one adds a “Dehn twist” operation, then the number of steps becomes polynomial in $\log |w|$ [\[1\]](#).

7.4 Computing other information

As pointed out in [\[1\]](#), one can adapt this framework to compute additional invariants of the multicurve γ given by $w \in \mathcal{ML}(\tau, \mathbb{Z})$. The cases of interest for us are derived from the following setup. Suppose we know that γ is connected, and fix generators for $\pi_1(\Sigma)$. Let’s see how to find a word in $\pi_1(\Sigma)$ representing γ in terms of splitting. We can take the basepoint for $\pi_1(\Sigma)$ to lie in the base interval for τ , which we think of as very small. An oriented band of τ thus gives rise to an element of $\pi_1(\Sigma)$. We think of each band as being *labeled* with this element. Now suppose we do a splitting $(\tau, w) \searrow (\tau', w')$. Resuming the notation of [Figure 18](#), we presume $w_t > w_b$. Orient the bands t and b so that both orientations point vertically at the right-hand side of the base interval. The new bands b' and t' of τ' now inherit orientations as well. If we use L to denote our $\pi_1(\Sigma)$ labels, then these transform via

$$(7.5) \quad L(b') = L(b) \cdot L(t) \quad \text{and} \quad L(t') = L(t),$$

with all the other labels remaining unchanged. Since we are presuming that γ is connected, if we continue splitting in this manner we eventually arrive at a single loop with weight 1. The label on that loop is then a word representing γ in $\pi_1(\Sigma)$.

Since in the end we recover a full word representing γ , this splitting algorithm takes time at least proportional to the size of that word; this can certainly be as large as $|w|$. The real payoff is when we want to compute something derived from this word which carries much less information. For instance, suppose we want to know the class of γ in $H_1(\Sigma)$. Then we can use labels which are the images of the $\pi_1(\Sigma)$ labels under the quotient $\pi_1(\Sigma) \rightarrow H_1(\Sigma)$. In this way, we can compute the class of γ in $H_1(\Sigma)$ in time polynomial in $\log |w|$ [\[1\]](#).

7.6 The algorithm: boxes on interval exchanges

We now turn to the main question at hand. Suppose H is our genus 2 handlebody, and $\gamma \in \mathcal{T}$ a non-separating simple closed curve. We want to (efficiently) decide

if the associated tunnel number one manifold M_γ fibers. Suppose that γ is given to us in terms of weights w on a train track τ_0 . Using the technique of the last subsection, we can quickly compute the element γ represents in $H_1(H)$. Let us further suppose that this is not 0; we now have determined the essentially unique epimorphism $\phi: \pi_1(M_\gamma) \rightarrow \mathbb{Z}$. In light of [Corollary 4.3](#), to decide if M_γ fibers we just need to apply Brown's Criterion to decide if the kernel of ϕ is finitely generated. In [Section 5.9](#), we described how to implement Brown's Criterion by breaking the defining relation R up into subwords and using boxes to capture the needed information about these subwords. Roughly, we initially label τ by corresponding words of $F = \pi_1(H) = \langle a, b \rangle$, and then immediately replace each word v with $\text{Box}_\phi(v)$. Then at each split, we will combine the boxes via box multiplication following the rule [\(7.5\)](#). At the end we will be left with a single loop labeled by $\text{Box}_\phi(R)$ to which we can apply [Theorem 5.12](#). However, in order for the final box to really be $\text{Box}_\phi(R)$, we must restrict the initial train track $\tau: \text{Box}_\phi$ is not a morphism unless we take the domain to be the monoid of words in $\{a^{\pm 1}, b^{\pm 1}\}$, rather than the free group F itself. In particular, we must ensure that each time we split the interval exchange there is no cancellation in the F labels.

While one way of thinking about the final label on γ is via the splitting process, it can be also thought of less dynamically. Focus on a neighborhood of τ , and think of each band as having a vertical dividing line in the middle of its length. Fix a transverse orientation for the divider. Suppose we label each band by a word in the generators and their inverses. The label for each band should be thought of as affixed to its dividing line. A connected curve γ carried by τ has a sequence of intersections with the dividers; reading off the labels as we go around γ (inverting the label if the direction of travel does not match the transverse orientation of the divider) and taking the product gives the final word. The final word is well-defined up to the choice of starting point and choice of orientation of γ . We say that γ is *tight* if the final word is cyclically reduced.

7.7 Definition Let τ be an interval exchange with bands labeled by elements of F . We say that τ is *tightly labeled* if every γ carried by τ is tight.

An example of a tightly labeled interval exchange is given in [Figure 19](#); the point is that as we run along γ , if we cross a label which is a power of a , then the next label other than 1 that we encounter is a power of b . Concluding the above discussion, we have:

7.8 Lemma *Let τ be an interval exchange tightly labeled by F . Suppose $\phi: F \rightarrow \mathbb{Z}$ is an epimorphism, and let γ be a connected simple closed curve carried by τ . Split τ until we get a single loop labeled by R . Now start back at the beginning and replace the labels on τ by $L \mapsto \text{Box}_\phi(L)$, and again split until we get a single loop labeled with a box B . Then $B = \text{Box}_\phi(R)$.*




Figure 19: A tightly labeled interval exchange.

As at the beginning of this subsection, suppose we are given $\gamma \in \mathcal{T}$, our parameter space of tunnel number one 3-manifolds. We assume that γ is given to us in terms of Dehn–Thurston coordinates as in Section 2.3. As we next describe, the constraints (2–4) in Section 2.3 on the Dehn–Thurston coordinates of $\gamma \in \mathcal{T}$ allow us to put γ on a tight interval exchange closely related to the one given in Figure 19. In terms of Figure 3, consider the punctured torus T bounded by δ containing α . The intersection of γ with T consists of at most 3 parallel families of arcs. Thinking homologically, it is easy to see that we can orient things so that the labels on these families are a^i , a^j and a^{i+j} where $i, j \in \mathbb{Z}$. Condition (4) of Section 2.3 means that none of $\{i, j, i + j\}$ are zero. (If there are fewer than 3 families of arcs, we add in empty families to increase the number to 3, making the discussion uniform. This can be done fairly arbitrarily, and we can thus ensure that $\{i, j, i + j\}$ are all nonzero.) Each family of arcs will contribute one band to our final τ . The same picture is true for the other punctured torus. We can now make an interval exchange τ by taking these bands in the punctured tori and adding one additional band to allow us to effect the twist around δ . The result is shown




Figure 20: A standard starting interval exchange. Here the elements a^i , a^j , a^{i+j} , b^k , b^l , and b^{k+l} of F are not the identity.

in Figure 20. (If you are worried here about whether the final band is consistent with the orientation of the twisting about δ , note that since a full Dehn twist about δ extends

over H , we can make this twisting have any sign we like without changing M_γ .) We will call interval exchanges of this form *standard starting interval exchanges*. The same reasoning used above shows that τ is tightly labeled. For future reference we record:

7.9 Lemma *Every $\gamma \in \mathcal{T}$ is carried by one of countably many standard starting interval exchanges, each of which is tightly labeled by F .*

To summarize, here is the procedure to efficiently decide if M_γ fibers over S^1 , provided that γ is nonzero in $H_1(H)$. (If γ is zero in $H_1(H)$, there is not a unique ϕ to test. While Brown's algorithm adapts to work very elegantly to this situation (see [Section 5.4](#)), it is unclear if it can be implemented efficiently in this case. You may have to remember too much in the appropriate labels.) First, it is straightforward in terms of the Dehn–Thurston coordinates to put γ on a standard initial exchange τ_0 . Then run the splitting once using $H_1(H)$ labels to determine ϕ . Once ϕ is known, we go back the beginning and relabel τ_0 with Box_ϕ labels. Now run the splitting to the bottom again. In light of [Lemma 7.8](#) and [Theorem 5.12](#), the label on the final loop determines if the kernel of ϕ is finitely generated, and hence if M_γ fibers over the circle. Since there isn't much to a box, really just 5 numbers bounded by the square of the initial weights, the running time will still be polynomial in the log of initial weights, or equivalently in the size of the Dehn–Thurston coordinates.

8 The idea of the proof of the main theorems

In this section, we explain the outline of the common proof of the main results of this paper, [Theorems 2.4](#) and [2.5](#). For concreteness, let us focus attention on [Theorem 2.4](#). The basic idea is to analyze the algorithm given in the last section, and prove that it will report “non-fibered” with probability tending to 1 as the input curve γ becomes more and more complicated. Recall the setup is that we are given a connected curve γ on ∂H carried by a tightly labeled copy of τ_0 . Let $G = \pi_1(M_\gamma) = \langle a, b \mid R = 1 \rangle$, and let us assume we are in the generic case where there is a unique epimorphism $\phi: G \rightarrow \mathbb{Z}$. In the algorithm of [Section 7.6](#), we start with Box_ϕ labels on τ_0 , and then split (τ_0, γ) repeatedly, at each stage replacing one of the box labels by its product with another box. After many splittings, we are left with a single loop. What we want to show is that the top of the remaining box is very likely to be marked, and thus $\ker(\phi)$ is infinitely generated by [Theorem 5.12](#).

Suppose at some point in the splitting sequence we get to a state where *every* box label has a marked top. As [Lemma 5.14](#) shows, this property will persist as we continue to split, and so this state alone implies that $\ker(\phi)$ is infinitely generated. Thus we can

stop at this stage, even though we may still have a million splits to go before we get down to a loop. This is in fact what happens when you implement the algorithm of [Section 7.6](#) — typically, you end up with marked boxes on all the bands long before you have completed splitting. The strategy for the proof is therefore to show that having all boxes marked becomes increasingly likely as you do more splits.

We begin by discussing a much simpler problem, whose solution follows the same strategy, and then explain how the approach must be modified to account for the constraints in the actual topological situation.

8.1 A toy problem

Our toy problem is the following. Suppose we start with a finite set of boxes; for concreteness, let us say there are 7 of them to match the number of bands in one of the standard starting interval exchanges shown in [Figure 20](#). At each “split” we pick two distinct boxes A and B at random, and replace A with a random selection from

$$A * B, \quad A * B^{-1}, \quad B * A, \quad \text{or} \quad B^{-1} * A,$$

where here $B^{-1} = \text{rev}(B)$ is the reverse or “inverse” box described in [Section 5.9](#). The simplification here is that any pair of boxes can interact at any stage, whereas for interval exchanges the pair is fixed by the topology. However, for interval exchanges one expects that over time all pairs of labels will interact, and so it is reasonable to hope that the behavior of the toy problem will tell us something about the real case of interest.

For any choice of the initial 7 boxes where at least one is non-trivial, we’ll show:

8.2 Theorem *The probability that all boxes have marked tops after n splits goes to 1 as $n \rightarrow \infty$.*

This result is an easy consequence of the following lemma.

8.3 Lemma *Let B_1, B_2, \dots, B_7 be boxes where at least one is nontrivial. Then there is a sequence of 14 splits so that the resulting boxes all have marked tops.*

Assuming the lemma, here’s the proof of the theorem. Start with our initial boxes, split 14 times, then 14 more times, etc. While we don’t know anything about the state of the boxes at the start of each chunk of 14 splits, the lemma tells us that there is at least a 28^{-14} chance that the next 14 splittings result in all boxes marked. Since the splittings in distinct chunks are chosen independently, the probability that some box is

not marked after $14k$ splits is at most $(1 - 28^{-14})^k$. Hence the probability that all of the boxes are marked converges to 1 as the number of splits goes to infinity. To see that we are on the right track, notice that this exponential decay of the probability of fibering is consistent with the experimental data in [Section 3](#).

We now head toward the proof of [Lemma 8.3](#); before continuing, the reader may want to review the notation of [Section 5.9](#). First, if X is a nontrivial box with shift 0, and then X^2 has a marked top. Moreover, we can create a box with shift 0 from any pair of boxes A, B by taking the commutator $X = A * B * A^{-1} * B^{-1}$. Now, given our 7 boxes, it is not easy to create a commutator by the splitting moves; however, one thing we can do is that if A, B, C are three of the boxes, then we can replace C with

$$(8.4) \quad C' = C * (A * B * A^{-1} * B^{-1})^2$$

by doing 8 splits. Roughly, if the box C is shorter than $X = A * B * A^{-1} * B^{-1}$ then the top of C' should come from the X^2 term, and hence C' will have a marked top as well. We turn now to the details.

Proof of [Lemma 8.3](#) Let B_1, B_2, \dots, B_7 be our initial boxes. Given the way splittings work, we can replace B_i with B_i^{-1} without really changing anything, so let's normalize things so that the shifts satisfy $s_i \geq 0$. We will denote the top of B_i by t_i . Let A be the nontrivial B_i with largest top, which we denote t_a . Pick two of the remaining B_i , and denote them B and C . The splitting sequence we will use is: first do the sequence of 8 splits replacing C by C' in [\(8.4\)](#); then for each B_i which is not C' , do a split to replace B_i with $B_i * C'$.

To see that all boxes are marked at the end of this, first consider $X = A * B * A^{-1} * B^{-1}$. By choice of A , the top of X satisfies $t_x \geq t_a \geq t_i$ for all i . Then C' above has a marked top coming from X^2 at height $t_x + s_c$, where s_c is the shift of C . Thus for any B_i we have that $B_i * C'$ has a marked top at height $s_i + s_c + t_x$ since $t_x \geq t_i$. Thus we can make all the boxes marked in only 14 splits. \square

8.5 Outline of the proof of the main theorems

With [Theorem 8.2](#) in hand, we now explain how the approach generalizes to [Theorem 2.4](#). As mentioned above, the difficulty we need to incorporate is that with interval exchanges, we have much less freedom in how the boxes are changed at each step. Despite this, the analog of [Lemma 8.3](#) is still true. In [Section 9](#), we prove that there is a single “magic” splitting sequence which always gives marked boxes. Interestingly, unlike [Lemma 8.3](#), some assumptions must be made on the starting boxes for this to be true; however, these always hold for those boxes arising in the algorithm of

[Section 7.6](#) (see [Remark 9.8](#) for more). In the toy problem, going from [Lemma 8.3](#) to [Theorem 8.2](#) was essentially immediate. The key features were that each block of 14 splits is independent of the others, and the desired splitting sequence always has a definite probability of occurring. For interval exchanges, these things are more subtle; essentially, what we need is that splitting complete genus 2 interval exchanges is “normal”. This is shown in [Section 10](#), relying on work of Kerckhoff [19]. Finally, in [Section 11](#) we assemble the pieces just discussed with work of Mirzakhani [23] to complete the proof of [Theorem 2.4](#).

9 The magic splitting sequence

This section is devoted to a lemma which is one of the central ingredients in the proof of the main theorems. Suppose γ is a connected curve on the boundary of our genus 2 handlebody H . Suppose γ is carried by some interval exchange τ_0 . Roughly, we show that if the splitting sequence of (τ_0, γ) has a certain topological form, then the tunnel number one 3–manifold M_γ does not fiber over the circle. Before stating the lemma, we discuss its precise context.

9.1 Complete interval exchanges

We will work with interval exchanges τ in a surface Σ for which $\mathcal{ML}(\tau, \mathbb{R})$ gives a chart for $\mathcal{ML}(\Sigma)$; in particular, we work with complete interval exchanges, which we now define. Let Σ be a closed surface of genus at least 2. An interval exchange τ in Σ is called *recurrent* if there is a $w \in \mathcal{ML}(\tau, \mathbb{R})$ where every band has positive weight. For an interval exchange, the switch condition is that the sum of the weights of the orientation reversing bands on the top is equal to the corresponding quantity for the bottom. Thus, τ is recurrent if and only if there are orientation reversing bands on both sides, or no such bands at all. The exchange τ is *complete* if it is recurrent, and every complementary region is an ideal triangle. When τ is complete, the natural map $\mathcal{ML}(\tau, \mathbb{R}) \rightarrow \mathcal{ML}(\Sigma)$ is a homeomorphism onto its image; if we restrict the domain to w which are nowhere zero, then we get a homeomorphism onto an open subset of $\mathcal{ML}(\Sigma)$. (When working with train tracks, one also requires transverse recurrence in the definition of completeness. However, an easy application of [26, Corollary 1.3.5] shows that any interval exchange is transversely recurrent.)

In what follows, the embedding of τ into Σ is not really relevant. Thus, we will tend to think of interval exchanges abstractly, that is, as not embedded in any particular surface. If one presumes that the complementary regions of τ are ideal polygons, then one can reconstruct Σ from the combinatorics of τ alone. Thus, it makes sense to speak of an abstract interval exchange as being a complete interval exchange on a genus 2 surface.

9.2 Statement of the lemma

Now let's give the setup for the main result of this section. We are interested in splitting sequences of interval exchanges. Suppose τ is an interval exchange. Given $w \in \mathcal{ML}(\tau, \mathbb{R})$, as described in Section 7.3 we can split (τ, w) to (τ', w') , which is denoted by $(\tau, w) \searrow (\tau', w')$. Independent of the choice for w , there are (at most) 3 distinct possibilities for τ' ; in the notation of Figure 18, the 3 possibilities correspond to $w_t > w_b$, $w_t < w_b$, and $w_t = w_b$. We will use the notation $\tau \searrow \tau'$ to indicate that (τ, w) splits to (τ', w') for some $w \in \mathcal{ML}(\tau, \mathbb{R})$.

From now on we will look at complete interval exchanges on a genus 2 surface. Of special importance is the exchange shown in Figure 20, which we will call τ_0 . Suppose we have a splitting sequence

$$S: \tau_0 \searrow \sigma_1 \searrow \sigma_2 \searrow \cdots \searrow \sigma_n$$

where the σ_i are also complete genus 2 interval exchanges. For a multicurve $\gamma \in \mathcal{ML}(\tau_0, \mathbb{Z})$, we say that γ exhibits S if the initial part of the splitting sequence of (τ_0, γ) is

$$\tau_0 \searrow \tau_1 \searrow \tau_3 \searrow \cdots \searrow \tau_m$$

where the tail $\tau_{m-n} \searrow \tau_{m-n} \searrow \cdots \searrow \tau_m$ is abstractly isomorphic to S . The point of this section is to prove:

9.3 Lemma There exists a splitting sequence of complete genus 2 interval exchanges

$$S: \tau_0 \searrow \tau_1 \searrow \tau_2 \searrow \cdots \searrow \tau_n$$

such that the following holds. Suppose γ is a connected simple closed curve on ∂H carried by a tightly labeled copy of τ_0 . If the splitting sequence for (τ_0, γ) exhibits S , then the manifold M_γ does not fiber over the circle.

Note that in the lemma γ is allowed to be either separating or non-separating. The splitting sequence S will be referred to as the *magic splitting sequence*. It is quite complicated, and so even to describe it, we must first give another point of view on splitting interval exchanges.

9.4 Flexible splitting of interval exchanges

Let τ be an interval exchange with some initial measured lamination $w \in \mathcal{ML}(\tau, \mathbb{R})$. Let $I = [0, L]$ be the base interval of the exchange. In the notation of Section 7.3, during each splitting we reduce the length of I by $\min(w_b, w_t)$. We now describe a way of seeing the result of several splittings at once. Consider a subinterval $J = [0, L']$

for some $L' < L$. Now take a knife to (τ, w) and begin to slice it starting at some notch between bands, following the lamination as you go. Here J should be viewed as indestructible, and when the knife collides with J you stop. Repeat for each of the other notches until no more progress can be made. Thus we have created a new pair (τ', w') which describes the same lamination. Moreover (τ', w') is a stage of the splitting sequence for (τ, w) , in particular the one right before the base interval shrinks to a proper subinterval of J . In describing this cutting process, there is no need to cut each notch down to J in one go — we can start somewhere, cut for a bit, and work somewhere else before coming back to finish the job.

9.5 The magic sequence

We now describe the first part of the magic sequence, using the setup just given. In [Figure 21](#), we start in the upper left with τ_0 drawn as a train track; in this picture of τ_0 the top and the bottom of the vertical segment are identified and this convention persists throughout the figure. Also marked on τ_0 is the base interval I and the smaller initial subinterval J . [Figure 21](#) describes a splitting sequence S_1 where we split τ_0 down to J . The final train track is again the interval exchange τ_0 . A choice of $w \in \mathcal{ML}(\tau_0, \mathbb{R})$ which induces this splitting is not indicated; it can be determined a posteriori by choosing nowhere zero weights w' on the final copy of τ_0 , and working backwards up the sequence to determine w . In proving [Lemma 9.3](#), it will be very important to know what the labels (in the sense of [Section 7.6](#)) are on the final copy of τ_0 . The initial τ_0 is labeled by $\{x, a, b, c, d, e, f\}$ in $\pi_1(H)$. When following through what happens in [Figure 21](#), you should view each label as sitting precisely at the indicated arrow. The final labels are as follows, where $A = a^{-1}$, etc. and the vertical bars should simply be ignored for now.

$$\begin{aligned}
 x' &= XCFBEADcfbead \\
 a' &= DAEBFCdaebf|bead \\
 b' &= DAEBFCdae|ad \\
 c' &= DAEBFCd| \\
 d' &= XCFBEAD|cx \\
 e' &= XCFBEADcf|bfcx \\
 f' &= XCFBEADcfbe|aebfcx
 \end{aligned}$$

(9.6)

The second part of the magic splitting sequence is much easier to describe. Suppose $w \in \mathcal{ML}(\tau_0, \mathbb{R})$ has larger weight on the x band than the c band. Then splitting one step gives us τ_0 again. We refer to this as the *stable splitting* of τ_0 . Let S_2 consist

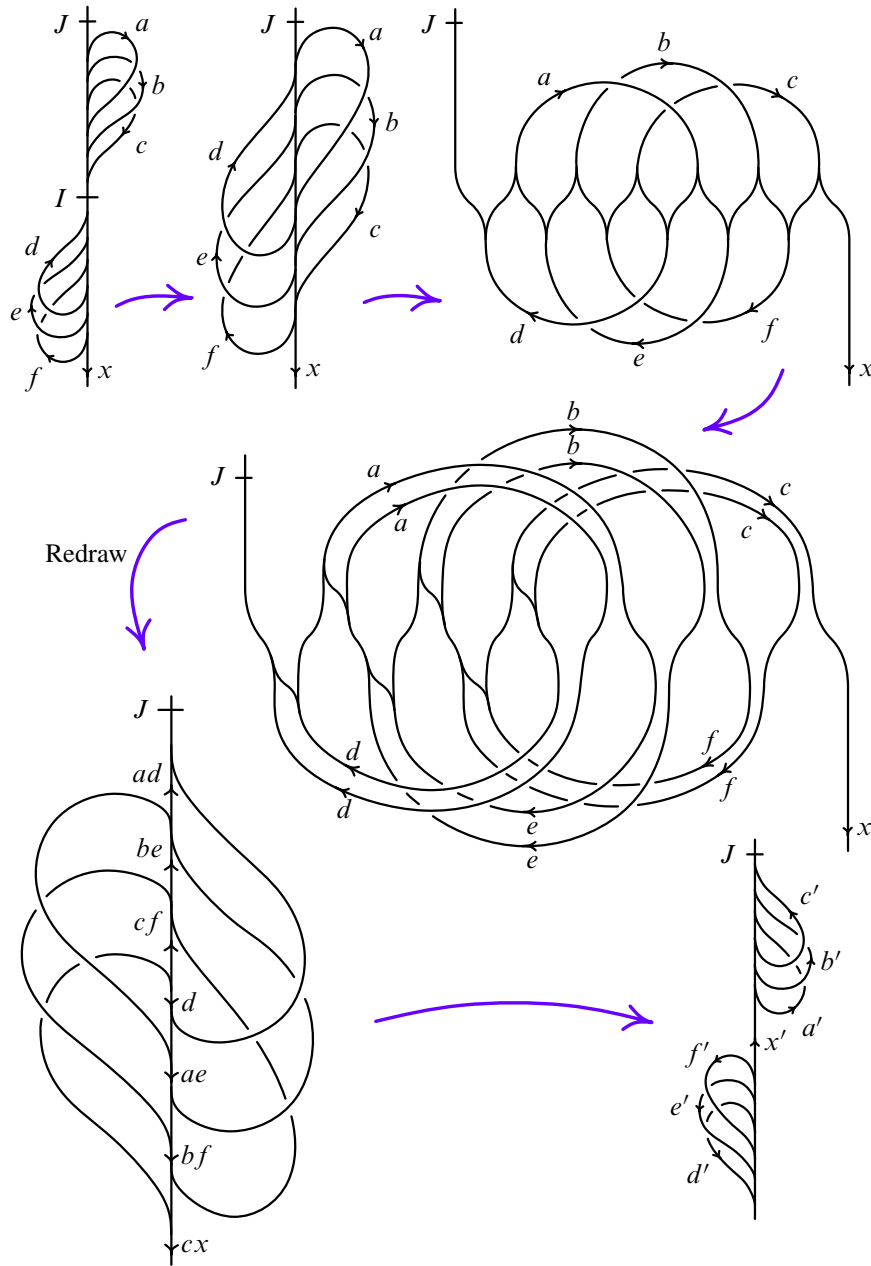


Figure 21: The first part S_1 of the magic splitting sequence, which starts in the top right with τ_0 , and ends in the bottom left with another copy of τ_0 .

of repeating the stable splitting 6 times. The key properties of S_2 that we will use are as follows. The sequence S_2 affects the labels by replacing a with Xax and the same for b and c . Suppose the initial splitting sequence of some $w \in \mathcal{ML}(\tau_0, \mathbb{R})$ is S_2 , what does that tell us about w ? Let w_x be the width of the x -band, and $w_r = 2(w_a + w_b + w_c)$ be the difference between the length of the base interval of τ_0 and the x -band. It is not hard to see that the splitting sequence for (τ_0, w) starts off with S_2 if and only if $w_x \geq w_r$. More generally, the splitting sequence starts off with n copies of S_2 if and only if $w_x \geq nw_r$.

Finally, the magic sequence itself is

$$S = S_1 \searrow S_2 \searrow S_2 \searrow S_2 \searrow S_2 \searrow S_2 \searrow S_2 \searrow S_2.$$

9.7 Proof of the lemma

We first outline the approach to proving [Lemma 9.3](#). Recall the setup is that we are given a connected curve γ on ∂H carried by a tightly labeled copy of τ_0 . Supposing the splitting sequence for (τ_0, γ) exhibits S , then we need to show that M_γ does not fiber over the circle. If $G = \pi_1(M_\gamma) = \langle a, b \mid R = 1 \rangle$, for each epimorphism $\phi: G \rightarrow \mathbb{Z}$ we need to show that $\ker(\phi)$ is infinitely generated. From now on, we view ϕ as fixed. Roughly, the strategy of the proof is we start with Box_ϕ labels on τ_0 as in [Section 7.6](#), and then split until the sequence S occurs. At that point, all the box labels will have their tops *marked* in the sense of [Section 5.9](#). Thus the relation R for G coming from γ is a product where each factor has Box_ϕ with a marked top. By [Lemma 5.14](#), the subgroup $\ker(\phi)$ is then infinitely generated. For technical reasons, the proof of the lemma deviates slightly from the above sketch, though we suspect it could be made to work on the nose by (further) complicating S . We turn now to the details.

Proof of [Lemma 9.3](#) Continuing with the notation above, we can assume that the *initial* splitting sequence of (τ_0, γ) is S . Beyond the tightness restraint, the only thing we need to show about the initial labels on τ_0 is that $\phi(x) = 0$, as follows. Consider the curve ϵ carried by τ_0 where the only nonzero weight is $w_x = 1$. From the way ϵ divides up τ_0 , we see that ϵ is a *separating* curve in ∂H . Thus the word $x \in \pi_1(H)$ associated to ϵ lies in the commutator subgroup, and so $\phi(x) = 0$.

Now split (τ_0, γ) along S_1 , getting back to τ_0 with the labels as in [\(9.6\)](#). Next do the S_2 splitting twice, so that the we have new labels

$$a'' = (X')^2 a'(x')^2 \quad b'' = (X')^2 b'(x')^2 \quad c'' = (X')^2 c'(x')^2$$

with the others unchanged. Changing tacks, rather than implement the 4 remaining S_2 splits, we just use the fact noted above that this means that the x' band is much wider than all the other bands put together. In particular, we now split starting from the *left* side of the base interval rather than the right, and do two analogs of the S_2 splittings there. This has the effect of changing the labels by

$$c'' = (x')^2 c'(X')^2 \quad e'' = (x')^2 e'(X')^2 \quad f'' = (x')^2 f'(X')^2$$

We claim that the boxes of $a'', b'', c'', d'', e'', f''$ all have marked tops. Consider for instance

$$a'' = (DAEBFCdaebfcx)^2 \cdot DAEBFCdaebf|bead \cdot (XCFBEADcfbead)^2.$$

Now we have $\phi(x') = \phi(X') = 0$, by the same argument that shows $\phi(x) = 0$. If we look at the part of a'' lying to the left of the vertical line, we see a subword of $(X')^3$ that is long enough so [Lemma 5.15](#) implies that its box has a marked top. Similarly, the right half of a'' also has a box with a marked top, and thus $\text{Box}_\phi(a'')$ has a marked top. The same argument works for all the other labels except x' , where the division of the word into two parts is indicated in [\(9.6\)](#).

To conclude the proof, it would be enough to know that $\text{Box}_\phi(x')$ has a marked top. Rather than show this, first note that $\text{Box}_\phi((x')^n)$ has a marked top for any $n \geq 2$. We still have two S_2 splits “left” that we haven’t used. This means the weight on the x' band is large enough that the places where the x' band is attached on the top and bottom have considerable overlap. As a result, anytime the curve γ enters the x' band it goes around it at least twice before moving to a different band. Thus the defining relation R for G is a (non-canceling) product of the words

$$a'', b'', c'', d'', e'', f'', (x')^n \quad \text{for } n \geq 2$$

and their inverses. Since all of these have marked boxes, the subgroup $\ker(\phi)$ is not finitely generated by [Lemma 5.14](#), completing the proof. \square

9.8 Remark The subtle issue in the proof is not that marked boxes tend to persist, but why any marked boxes must be created in the first place. For the latter, the fact that $\phi(x) = 0$ is crucial, as otherwise the box of a power like $(x')^2$ need not have a marked top. We found situations where you start with a train track and put random box labels on it so that no sequence of splitting operations can produce marked boxes on all the bands.

10 Ubiquity of splitting sequences

In this section, we study splitting sequences of complete genus 2 interval exchanges. We show that certain finite splitting sequences (including the magic one) are ubiquitous — in a suitable sense, they occur in the splitting sequence of almost every multicurve. Before giving the precise statement, it is worth considering the analogous fact for continued fractions. Let s_1, s_2, \dots, s_n be a sequence of positive integers. Given a rational number $p/q \in [0, 1]$, we can look at the partial quotients in its continued fraction expansion. We can ask if the sequence $\{s_i\}$ occurs as successive terms in these partial quotients. In fact, the probability that this occurs goes to 1 as $\max(p, q) \rightarrow \infty$. The point of this section is to prove the corresponding result, [Theorem 10.1](#), for genus 2 interval exchanges.

Among complete interval exchanges on a genus 2 surface, we focus on the exchange τ_0 shown in [Figure 20](#). Suppose we have a splitting sequence

$$S: \tau_0 \searrow \sigma_1 \searrow \sigma_2 \searrow \cdots \searrow \sigma_n$$

where the σ_i are also complete genus 2 interval exchanges. As in the last section, for a multicurve $\gamma \in \mathcal{ML}(\tau_0, \mathbb{Z})$, we say that γ *exhibits* S if its splitting sequence contains a copy of S . We will show that, asymptotically, the set of γ in $\mathcal{ML}(\tau_0, \mathbb{Z})$ which exhibit S has density 1. More precisely:

10.1 Theorem *Let τ_0 be the complete genus 2 interval exchange specified above, and S a splitting sequence of τ_0 consisting of complete interval exchanges. Set*

$$\mathcal{S} = \{\gamma \in \mathcal{ML}(\tau_0, \mathbb{Z}) \mid \gamma \text{ exhibits } S\}.$$

If U is a bounded open set in $\mathcal{ML}(\tau_0, \mathbb{R})$, then

$$\frac{\#(\mathcal{S} \cap tU)}{\#(\mathcal{ML}(\tau_0, \mathbb{Z}) \cap tU)} \rightarrow 1 \quad \text{as } t \rightarrow \infty.$$

This theorem may strike the reader as excessively narrow: what is so special about genus 2 and this particular choice of τ_0 ? The choice of τ_0 can be broadened, but not to the point where every complete genus 2 exchange can play its role. The restriction to genus 2 is necessary for the proof as given; it is unclear to us if it is actually needed. Both of these issues are discussed at length below. The proof of [Theorem 10.1](#) is based on applying a criterion of Steve Kerckhoff [\[19\]](#), which we discuss in the next two subsections.

10.2 Splitting interval exchanges as a dynamical system

The bulk of the proof of [Theorem 10.1](#) will be carried out in a slightly different setting. This subsection is devoted to describing this setting, and stating the analog therein of [Theorem 10.1](#). Let \mathcal{I} be the set of complete genus 2 interval exchanges up to (abstract) isomorphism. It is easy to see that any $\tau \in \mathcal{I}$ has 7 bands, so \mathcal{I} is finite. Consider the set

$$X = \coprod_{\tau \in \mathcal{I}} \mathcal{ML}(\tau, \mathbb{R}).$$

with the measure coming from Lebesgue measure on each $\mathcal{ML}(\tau, \mathbb{R})$. We can construct a transformation of X by sending (τ, w) to its splitting (τ', w') . This is well-defined provided τ' is also complete; equivalently, we are not in the corner case where the weights on the rightmost bands are equal, ie, $w_t = w_b$ in the notation of [Figure 18](#). As this only concerns a measure 0 subset of X , we simply delete all (τ, w) in X which pass through the corner case somewhere in their splitting sequence. (Although we have just removed all the integral points which are the focus of [Theorem 10.1](#), this will not inconvenience us greatly.) Thus given a (τ, w) in X , we get an infinite sequence $\tau = \tau_1 \searrow \tau_2 \searrow \cdots$ of elements of \mathcal{I} . We are interested in the behavior of this sequence for generic choices of the initial (τ, w) , where here generic means except for a measure 0 subset of X . The goal is to show that X with this transformation is *normal*; that is, a finite splitting sequence

$$S: \sigma_1 \searrow \sigma_2 \searrow \cdots \searrow \sigma_n$$

which can happen must happen infinitely often for almost every (τ, w) . Here, the phrase “can happen” must be interpreted correctly, as we next discuss.

Make \mathcal{I} into a directed graph by putting an edge from τ to τ' if $\tau \searrow \tau'$. The splitting sequence of some (τ, w) now corresponds to an infinite directed path in this graph. A complication is that \mathcal{I} is not directedly connected; that is, there are τ and τ' so that no directed path starting at τ ends at τ' . In terms of our dynamical system X , this gives rise to transient states which generically appear only finitely many times in a splitting sequence. If τ can be joined to τ' by a directed path starting at τ , we write $\tau \searrow\!\!\!\searrow \tau'$. Let

$$\mathcal{I}_S = \{\tau \in \mathcal{I} \mid \sigma \searrow\!\!\!\searrow \tau \text{ for every } \sigma \in \mathcal{I}\},$$

which we refer to as the *sink* of \mathcal{I} . A priori, \mathcal{I}_S could be empty; for instance, this is the case if \mathcal{I} is not connected.

We now restrict attention to the subset of X sitting over the sink

$$X_S = \coprod_{\tau \in \mathcal{I}_S} \mathcal{ML}(\tau, \mathbb{R})$$

which is closed under the splitting transformation. Our precise definition of normality is:

10.3 Definition Suppose the sink \mathcal{I}_S is non-empty. Then the transformation on X_S is *normal* if every finite splitting sequence

$$\sigma_1 \searrow \sigma_2 \searrow \cdots \searrow \sigma_n \quad \text{which is contained in } \mathcal{I}_S$$

occurs infinitely often in the splitting sequence of almost all $(\tau, w) \in X_S$.

We will show the following, which will easily imply [Theorem 10.1](#):

10.4 Theorem *As above, let X_S be the set of weights on complete genus 2 interval exchanges lying in the sink \mathcal{I}_S . Then the splitting transformation on X_S is normal. Moreover, τ_0 lies in the sink \mathcal{I}_S .*

Our proof of [Theorem 10.4](#) is a direct application of a normality criterion of Steve Kerckhoff [\[19\]](#), which he used to prove the analogous result for *classical* interval exchanges, namely those without orientation reversing bands. Readers familiar with [\[19\]](#) may wonder why we are working with non-classical interval exchanges rather than just using train tracks, as [\[19\]](#) states the analog of [Theorem 10.1](#) for complete train tracks in any genus. There are two reasons for this. The first is that using interval exchanges makes it much easier to understand \mathcal{I} explicitly. The other reason is that the proof given in [\[19\]](#) for the train track case is incomplete. The proof there involves two steps, the first is to establish that a certain combinatorial criterion implies normality, and the second to check that this criterion holds for train tracks. The criterion, which is the one we use here, is certainly strong enough to ensure normality; the problem occurs in the second step, as the criterion is violated for certain explicit train tracks. Kerckhoff informs us that he noticed this problem as well, and that there should be a weaker combinatorial condition which still ensures normality but also holds in the train track setting. Kerckhoff is planning on publishing a correction along these lines.

We turn now to the combinatorics of \mathcal{I} and its sink. It turns out that there are 201 genus 2 interval exchanges in \mathcal{I} , and 190 in the sink \mathcal{I}_S . Also, as mentioned above, $\tau_0 \in \mathcal{I}_S$. We checked this by brute force computer enumeration, and will not further justify these facts here. This is not a difficult calculation, requiring only 100 lines of code and a few minutes of computer time. The part that's actually used in [Theorem 10.1](#), namely that τ_0 is in the sink of its connected component of \mathcal{I} , is particularly straightforward: First, start with τ_0 and keeps splitting until you don't generate any new exchanges. Then for each of the 190 exchanges so generated, check that you can split them all back to τ_0 .

Enumerating all of \mathcal{I} , in particular to see that it is connected, requires a little more thought to set up. One approach is to think of a complete genus 2 interval exchange τ as constructed by starting with a disc and adding 7 bands. The boundary of the disc is divided into 28 segments, alternating between a place where we attach a band, and a gap between bands. Further, two of the gaps are distinguished because they form the vertical sides of the thickened interval that is the base of the exchange; we call these special gaps *smoothed*. If we ignore the smoothings, the complementary regions of τ are ideal polygons of valence either $\{3, 3, 3, 5\}$ or $\{3, 3, 4, 4\}$. Conversely, provided the complementary regions are of this form, choosing smoothings of gaps in the larger complementary regions gives a genus 2 interval exchange. As there are only 135,135 possible gluings for the bands, one can simply try them all and thus calculate \mathcal{I} .

For train tracks, the structure of \mathcal{I} and its sink can also be quite complicated. For a 4-punctured sphere, for instance, the vast majority of the complete train tracks do not lie in the sink. It would be quite interesting to answer the following question.

10.5 Question *Find a type of train track like object where the elements of the sink can be characterized topologically, and for which the corresponding dynamical system is normal.*

10.6 Normality after Kerckhoff

This subsection is devoted to the proof of [Theorem 10.4](#). We begin by describing Kerckhoff's normality criterion. First, we will need to work with *labeled* interval exchanges, that is, interval exchanges where the bands are labeled by integers from 1 to the number of bands. If τ is a labeled exchange, and we split τ to τ' , the convention for the labels on τ' is as follows. The bands of τ' that come unchanged from τ retain their labels; in the notation of [Figure 18](#), the two modified bands t' and b' get the labels of t and b respectively.

The set of labeled complete genus 2 exchanges is of course finite, and we focus on the subset \mathcal{I}' where the underlying unlabeled exchange lies in the sink \mathcal{I}_S . Again, \mathcal{I}' has structure of a directed graph with edges given by splittings. The forgetful map $\mathcal{I}' \rightarrow \mathcal{I}_S$ is a covering map. Fix a connected component \mathcal{I}'_0 of \mathcal{I}' ; the covering map $\mathcal{I}'_0 \rightarrow \mathcal{I}_S$ is also surjective. (It appears that \mathcal{I}' consists of two components.) We claim that \mathcal{I}'_0 is its own sink. Call a directed graph *strongly connected* if for all vertices v_1 and v_2 , there is a directed path from v_1 to v_2 . What we need is equivalent to:

10.7 Lemma *Let G be a finite strongly connected directed graph, and $H \rightarrow G$ a finite covering map. If H is connected then it is strongly connected.*

Proof A cycle in a directed graph is a union of edges which form a directed closed loop, where no vertex is visited more than once. Observe that a connected directed graph is strongly connected if and only if every edge is part of a cycle — the point here is that going almost all the way around a cycle effectively allows us go backwards along a directed edge. The result now follows by noting that the preimage of a cycle in G is a disjoint union of cycles in H . \square

From now on, we work to show that

$$Z = \coprod_{\tau \in \mathcal{I}'_0} \mathcal{ML}(\tau, \mathbb{R})$$

with its splitting transformation is normal; this suffices to prove [Theorem 10.4](#). We now set up the terminology needed to state Kerckhoff’s normality criterion. Let τ be a labeled interval exchange. The rightmost places where the bands are glued on the top and bottom are called the *critical positions*. The bands in those positions are called the *critical bands*, and are denoted t and b respectively. During a splitting move, the band whose width is reduced is said to be split by the other. For example, in [Figure 18](#) where $w_t > w_b$, we say that t is split by b , or equivalently b splits t . A *block* is a cycle in \mathcal{I}'_0 , that is, a splitting sequence

$$\tau_1 \searrow \tau_2 \searrow \cdots \searrow \tau_n \searrow \tau_1$$

starting and ending at the same point. The key definition is as follows:

10.8 Definition A block is said to be *isolating* if we can partition the labels into non-empty subsets $V = \{v_i\}$ and $W = \{w_j\}$ such that:

- (1) Every v_i splits some v_j and is split by some v_k .
- (2) No w splits a v .

Then we have:

10.9 Theorem (Kerckhoff) *If there are no isolating blocks, then the splitting transformation on Z is normal.*

The way we have stated things differs slightly from [\[19\]](#), so we now say how to directly connect our presentation to his work there (readers unfamiliar with [\[19\]](#) will want to skip ahead to the proof of [Theorem 10.4](#)). First, the notion of an isolated block is defined at the bottom of page 262 of [\[19\]](#). It is given there in terms of a partition of vertices of a simplex Σ . We have that $\mathcal{ML}(\tau, \mathbb{R})$ is a convex cone which is the intersection of $\mathbb{R}_{\geq 0}^7$ with a hyperplane. Here the coordinates of $\mathbb{R}_{\geq 0}^7$ correspond to the

labeled bands. Kerckhoff's simplex Σ is just the convex hull of the positive unit vectors along each of the coordinate axes. Splitting the interval exchange corresponds to adding vertices of Σ , as detailed in the proof of Prop. 1.4 of [19]. [Theorem 10.9](#) above is essentially just Theorem 2.1 of [19], with a slight modification because $\mathcal{ML}(\tau, \mathbb{R})$ is not all of $\mathbb{R}_{\geq 0}^7$. This modification is justified in the 2nd paragraph of page 268 of [19]. (The problem mentioned above with the proof of normality for train tracks occurs later, namely in the proof of Proposition 2.2 of [19].) Finally, the notion of normality is not precisely defined in [19], as it is a standard concept in dynamical systems. In particular, there is no mention there of the graph structure of \mathcal{I} or the need to focus on splitting sequences lying in the sink. However, these notions are implicit in [19], see in particular the second sentence of the proof of Corollary 1.9. We return now to the matter at hand.

Proof of [Theorem 10.4](#) The proof of this theorem is a little involved, but is purely combinatorial, and essentially self-contained. By [Theorem 10.9](#), we just need to show that the splitting transformation on Z has no isolating blocks. Suppose to the contrary we have a block

$$\tau_1 \searrow \tau_2 \searrow \cdots \searrow \tau_n \searrow \tau_1$$

isolating band subsets $V = \{v_i\}$ and $W = \{w_j\}$ as above. Throughout, we will think of a τ_i as being specified by two lists of band labels, one each for the top and bottom interval of the exchange. For instance, the standard exchange τ_0 is given by

$$\tau_0 = \frac{1234234}{5675671}.$$

A statement like “ v_i lies to the right of w_j ” means that one occurrence of v_i lies to the right of w_j in its list. For τ_0 , both the statements “2 lies to the right of 3” and “3 lies to the right of 2” are correct; they simply refer to different occurrences of the labels. We begin with:

10.10 Lemma *No w ever enters the critical position. Moreover, on both the top and the bottom interval, every v lies to the right of every w .*

Proof First, we argue as in Proposition 1.4 of [19] that no w ever enters one of the critical positions. By axiom (1) of isolation, at some stage along the block both of the critical positions are occupied by V bands. Reindexing, we can assume that this is the case for τ_1 . Consider the largest k so that τ_k has a band w_j in the critical position. As a single splitting only changes one of the labels in the critical positions, the other critical band for τ_k is some v_i . As the next stage τ_{k+1} has only V bands in the critical positions, for $\tau_k \searrow \tau_{k+1}$ we must have w_j splitting v_i , violating axiom (2). So no w_j ever enters a critical position.

The rest of the proof of the lemma is based on considering the following quantity:

$$C = \#(v_i \text{ to the right of all } w_j \text{ on top}) + \#(v_i \text{ to the right of all } w_j \text{ on bottom}) .$$

How does C change as we split τ_k to τ_{k+1} ? For notation, suppose v_i splits v_j . Then the critical end of v_i is removed, decreasing C by 1. On the other hand, the non-critical end of v_j is divided into two, and so overall either:

- (a) The non-critical end of v_j lies to the right of all w_k , in which case C is unchanged.
- (b) The non-critical end of v_j lies to left of some w_k , and C decreases by 1.

Axiom (2) of isolation forces each v_i to be split, and hence case (b) does occur somewhere along the block. But as C is non-increasing, this gives a contradiction as the block starts and ends at the same interval exchange, and so C must be unchanged after running all the way through the block. \square

The above lemma says that τ_k splits up, in a certain sense, into two distinct interval exchanges which have been stuck next to each other. Here one exchange consists of the W bands, and the other of the V_k bands; we denote them by W and V_k respectively (W is unchanging and so needs no subscript). For example, we might have:

(10.11)
$$\frac{w_1 w_2 w_3 w_1 | v_1 v_2 v_3 v_1}{w_3 w_2 | v_4 v_3 v_2 v_4} .$$

Note that W and V_k may be quite degenerate as interval exchanges, in particular they need not be recurrent. Moreover, this decomposition is purely at the combinatorial level; a measure $\mu \in \mathcal{ML}(\tau_k, \mathbb{R})$ is typically not the result of taking $\mu_w \in \mathcal{ML}(W, \mathbb{R})$ and $\mu_v \in \mathcal{ML}(V_k, \mathbb{R})$ and amalgamating them. We now work to acquire more information about W and the V_k , eventually deriving a contradiction.

10.12 Lemma *Both W and V_k have orientation reversing bands. Moreover, V_k has such bands on both top and bottom.*

Proof Suppose W has only orientation preserving bands. Then for all measure laminations $\mu \in \mathcal{ML}(\tau_k, \mathbb{R})$ the gaps on the top and bottom between the W bands and the V bands line up exactly. Thus we can slice through τ_k at that point to get a (generalized) train track carrying μ which is the disjoint union of W and V . Since the complementary regions to τ_k are ideal triangles, it follows that one of the complementary regions to μ is *not* an ideal triangle. But as this is true for *every* μ , the exchange τ_k cannot be complete as laminations with triangle complementary regions are dense in $\mathcal{ML}(\Sigma)$. So W has an orientation reversing band.




Figure 22: Splitting apart W and V_k results in amalgamating two vertices of the triangle T into a punctured monogon.

The same argument shows that V_k must have an orientation reversing band. Suppose it has such a band only on one side, say the top. As it is not possible to create a reversing band without a reversing band on the same side, it follows that none of the V_k have a reversing band on the bottom. Since there are no reversing bands on the bottom, we cannot ever decrease the number of such bands on the top. As the block starts and ends at the same exchange, it follows that the number of reversing bands is constant throughout. However, by axiom (1), at some point one of the reversing bands is split, necessarily by an orientation preserving band; as this increases the number of reversing bands, we have a contradiction. So each V_k must have orientation reversing bands on both sides. \square

What are the complementary regions of W and V_k , thought of as abstract interval exchange? If we slice across to separate τ_k into W and V_k , we amalgamate two ideal triangles at a pair of vertices. As this cutting separates τ_k into two pieces, in fact we must be amalgamating two ideal vertices of the *same* ideal triangle, as shown in Figure 22. Thus the complementary regions to W or V_k are all ideal triangles with one exceptional region, which we call the *outside* region. One of the outside regions is an ideal monogon, and the other is smooth. Next, we use this to show:

10.13 Lemma *The exchange W has at least 3 bands, and V_k has at least 4 bands.*

Proof Let's begin with W , which we know has an orientation reversing band. If W has only that one band, then the outside complementary region is smooth, but the interior complementary region is an ideal monogon. So we must add a second band with one end glued inside the monogon to "break" it. There are two possibilities depending on whether the second band reverses or preserves orientation, but these have outside regions a triangle and digon respectively. So W must have at least three bands.

Turning to V_K , we know it has orientation reversing bands on both sides. The only way to get rid of the associated interior monogons by adding a single additional band is

$$\frac{v_1 v_2 v_1}{v_3 v_2 v_3}$$

which has a digon complementary region. So V_k needs at least 4 bands. □

Now, a complete genus 2 exchange has 7 bands, so it follows from the lemma that V_k consists of exactly 4 bands. We still have work to do, as the example of (10.11) is consistent with what we know so far. To conclude the proof, we consider the possible choices for the number r of reversing bands of V_k .

$r = 4$: To avoid interior monogons, we must have two bands on each side, interlocking to form

$$\frac{v_1 v_2 v_1 v_2}{v_3 v_4 v_3 v_4}.$$

But then the only complementary region is a hexagon. So no V_k has $r = 4$.

$r = 3$: Since we have reversing bands on both sides, we can assume we have two such bands on the top and one on the bottom. The top reversing bands must interlock since we have only a single additional orientation preserving band to break any interior monogons and digons. Thus we must have

$$\frac{v_1 v_2 v_1 v_2}{v_3 v_4 v_3}$$

where the top end of v_4 has not yet been attached. There are 5 possibilities for the placement of v_4 , and an easy check shows that all of them have the wrong complementary regions. Thus no V_k has $r = 3$.

$r = 2$: As the example of (10.11) shows, this is possible as a stand-alone exchange; we argue instead that such an exchange cannot lie in an isolating block — that is, to return to where we start in a block we must involve some of the W bands. By Lemma 10.12 and the above cases, we must have that every V_k has exactly two reversing bands, one on each side. But by axiom (1) at some point a reversing band is split, which either creates or destroys an orientation reversing band, a contradiction.

Thus we have ruled out all possibilities for an isolating block, proving the theorem. □

10.14 Remark For surfaces of higher genus, it seems very likely that there are isolating blocks. Indeed, in genus 3 take

$$W = \frac{w_1 w_2 w_3 w_1}{w_4 w_3 w_2 w_4}$$

and let V_1 be a complete exchange for a once punctured genus 2 surface. If the monogon complementary region is in the right place, then $\tau_1 = W \cup V_1$ is a complete genus 3 exchange. If V_1 is in some closed strongly connected subgraph of such exchanges, then completeness will allow us to construct a splitting sequence of V_1 back to itself which satisfies axiom (1). What is not entirely clear is whether this isolating block is in the sink for genus 3 exchanges.

10.15 Proof of Theorem 10.1

We end this section by deriving Theorem 10.1 from Theorem 10.4.

Proof of Theorem 10.1 As in the statement of the theorem, let τ_0 be our particular complete genus 2 exchange, and S some splitting sequence of τ_0 consisting of complete exchanges. Set

$$S = \{\gamma \in \mathcal{ML}(\tau_0, \mathbb{Z}) \mid \gamma \text{ exhibits } S\}.$$

Let U be a bounded open set in $\mathcal{ML}(\tau_0, \mathbb{R})$. We need to show

$$\frac{\#(S \cap tU)}{\#(\mathcal{ML}(\tau_0, \mathbb{Z}) \cap tU)} \rightarrow 1 \quad \text{as } t \rightarrow \infty.$$

Let V_m be the set of $\mu \in \mathcal{ML}(\tau_0, \mathbb{R})$ such that the splitting sequence of μ passes through S ending at the m^{th} stage. The set V_m is defined by a sequence of strict inequalities in the weights of the bands in the successive critical positions; in particular, it is open. If we set $V = \bigcup V_m$, then by Theorem 10.4, the complement of V has measure 0 in $\mathcal{ML}(\tau_0, \mathbb{R})$. Moreover, as $S = \mathcal{ML}(\tau_0, \mathbb{Z}) \cap V$ and V is invariant under positive scaling

$$S \cap tU = \mathcal{ML}(\tau_0, \mathbb{Z}) \cap V \cap tU = \mathcal{ML}(\tau_0, \mathbb{Z}) \cap t(V \cap U).$$

$$\begin{aligned} \text{Therefore } \frac{\#(S \cap tU)}{\#(\mathcal{ML}(\tau_0, \mathbb{Z}) \cap tU)} &= \frac{\#((V \cap U) \cap t^{-1}\mathcal{ML}(\tau_0, \mathbb{Z}))}{\#(U \cap t^{-1}\mathcal{ML}(\tau_0, \mathbb{Z}))} \\ &= \frac{t^{-6} \cdot \#((V \cap U) \cap t^{-1}\mathcal{ML}(\tau_0, \mathbb{Z}))}{t^{-6} \cdot \#(U \cap t^{-1}\mathcal{ML}(\tau_0, \mathbb{Z}))} \end{aligned}$$

As $t \rightarrow \infty$, the top and bottom of the right-hand fraction converge to the Lebesgue measures of $V \cap U$ and U respectively, since these sets are open. As the complement of V has measure 0, the sets $V \cap U$ and U have the same measure. Hence the fraction limits to 1 as $t \rightarrow \infty$, completing the proof. \square

11 Proof of the main theorems

Here, we complete the proofs of Theorems 2.4 and 2.5. In addition to the results of Sections 9 and 10, we need one additional ingredient, which says that specific types of multicurves, eg, connected non-separating curves, have positive density in the set of all multicurves:

11.1 Theorem (Mirzakhani) *Let Σ be a closed surface of genus $g \geq 2$. Fix a multicurve $\gamma \in \mathcal{ML}(\Sigma, \mathbb{Z})$. Consider the set $\mathcal{C} = \text{MCG}(\Sigma) \cdot \gamma$ of all multicurves in Σ of the same topological type. Then for every bounded open subset U of $\mathcal{ML}(\Sigma, \mathbb{R})$ we have:*

$$\frac{\#(\mathcal{C} \cap tU)}{\#(\mathcal{ML}(\Sigma, \mathbb{Z}) \cap tU)} \rightarrow d_\gamma \in \frac{1}{\pi^{6g-6}} \mathbb{Q},$$

where d_γ is positive and independent of U .

For us, the exact value of d_γ will not be important, merely the fact that it is positive; however, Mirzakhani does provide a recursive procedure for computing it. The above theorem is a slight restatement of Theorem 6.4 of [23] where we have set $d_\gamma = c_\gamma/b_{g,0}$ in the notation there, and rewritten the measure $\mu_{t,\gamma}$ in terms of $\#(\mathcal{ML}(\Sigma, \mathbb{Z}) \cap tU)$ rather than t^{6g-6} via the proof of Theorem 3.1 of [23].

Let us now recall the statement of Theorem 2.4 and the notation of Section 2.3. Let H be our genus 2 handlebody, and fix Dehn–Thurston coordinates

$$(w_\alpha, w_\beta, w_\delta, \theta_\alpha, \theta_\beta, \theta_\delta)$$

for ∂H compatible with H as discussed in Section 2.3 and shown in Figure 3. We are interested in the set \mathcal{T} of attaching curves $\gamma \subset \partial H$ parameterizing tunnel number one 3-manifolds with one boundary component, where γ satisfies the restrictions:

- (1) γ is a non-separating simple closed curve.
- (2) The weights $w_\alpha, w_\delta, w_\beta$ are > 0 .
- (3) Each twist satisfies $0 \leq \theta < w$.
- (4) $w_\delta \leq \min(2w_\alpha, 2w_\beta)$.

We then consider the finite set $\mathcal{T}(r)$ of $\gamma \in \mathcal{T}$ where $w_\alpha + w_\beta < r$. Theorem 2.4 is that the probability that $M_\gamma \in \mathcal{T}(r)$ fibers over S^1 goes to 0 as $r \rightarrow \infty$.

To prove this, we first reinterpret the setup in the context of $\mathcal{ML}(\partial H, \mathbb{R})$. The Dehn–Thurston coordinates on $\mathcal{ML}(\partial H, \mathbb{Z})$ have natural extensions to coordinates on all of $\mathcal{ML}(\partial H, \mathbb{R})$ [26, Theorem 3.11], which we denote in the same way. Let $W \subset$

$\mathcal{ML}(\partial H, \mathbb{R})$ consist of measured laminations satisfying conditions (2–4) above. The subset W is a polyhedral cone in $\mathcal{ML}(\partial H, \mathbb{R}) \cong \mathbb{R}^6$ with some faces removed. Set

$$\mathcal{G} = \{\gamma \in \mathcal{ML}(\Sigma, \mathbb{Z}) \mid \gamma \text{ is connected and non-separating}\}.$$

We have $\mathcal{T} = W \cap \mathcal{G}$, and if we let U be the open subset of $\mathcal{ML}(\partial H, \mathbb{R})$ defined by $w_\alpha + w_\beta < r$ then

$$\mathcal{T}(r) = \mathcal{T} \cap rU = W \cap \mathcal{G} \cap rU.$$

A slightly more general result immediately implying [Theorem 2.4](#) is:

11.2 Theorem *Let $U \subset \mathcal{ML}(\partial H, \mathbb{R})$ be an open set such that $W \cap U$ is bounded. Then the probability that M_γ fibers over the circle for $\gamma \in \mathcal{T} \cap rU$ goes to 0 as $r \rightarrow \infty$.*

The rest of this section is devoted to the proof of [Theorem 11.2](#). The proof of [Theorem 2.5](#) which concerns tunnel number one 3–manifolds with two boundary components is identical if one simply replaces all occurrences of “non-separating” with “separating”.

Proof of Theorem 11.2 Let \mathcal{S} be the magic splitting sequence of [Lemma 9.3](#). Let \mathcal{S} consist of those laminations $\mu \in \mathcal{ML}(\partial H, \mathbb{R})$ which can be carried by some tightly labeled interval exchange τ where the splitting sequence of (τ, μ) exhibits \mathcal{S} . By [Lemma 9.3](#), for $\gamma \in \mathcal{S} \cap \mathcal{T}$ the manifold M_γ does not fiber. So if \mathcal{S}^c denotes the complement of \mathcal{S} , it suffices to show

$$(11.3) \quad \frac{\#(\mathcal{S}^c \cap \mathcal{T} \cap rU)}{\#(\mathcal{T} \cap rU)} \rightarrow 0 \quad \text{as } r \rightarrow \infty.$$

We next show that [Theorem 11.1](#) allows us to replace \mathcal{T} with $\mathcal{ML}(\partial H, \mathbb{Z}) \cap V$ in the above limit, where $V = W \cap U$. In particular

$$\begin{aligned} \frac{\#(\mathcal{S}^c \cap \mathcal{T} \cap rU)}{\#(\mathcal{T} \cap rU)} &\leq \frac{\#(\mathcal{S}^c \cap \mathcal{ML}(\Sigma, \mathbb{Z}) \cap rV)}{\#(\mathcal{G} \cap rV)} \\ &= \frac{\#(\mathcal{S}^c \cap \mathcal{ML}(\Sigma, \mathbb{Z}) \cap rV)}{\#(\mathcal{ML}(\Sigma, \mathbb{Z}) \cap rV)} \cdot \frac{\#(\mathcal{ML}(\Sigma, \mathbb{Z}) \cap rV)}{\#(\mathcal{G} \cap rV)} \end{aligned}$$

By [Theorem 11.1](#), the second factor in the final expression converges to a positive number as $r \rightarrow \infty$. Thus to show [\(11.3\)](#) it suffices to prove

$$(11.4) \quad \frac{\#(\mathcal{S}^c \cap \mathcal{ML}(\partial H, \mathbb{Z}) \cap rV)}{\#(\mathcal{ML}(\partial H, \mathbb{Z}) \cap rV)} \rightarrow 0 \quad \text{as } r \rightarrow \infty.$$

Now by [Lemma 7.9](#), W is covered by a countable collection of charts $\mathcal{ML}(\tau, \mathbb{R})$ where τ is one of the standard tightly labeled interval exchanges. In each chart,

the proof of [Theorem 10.1](#) shows that [Theorem 10.4](#) gives a *homogeneous open set* $Y_\tau \subset \mathcal{ML}(\tau, \mathbb{R}) \cap \mathcal{S}$ whose complement in $\mathcal{ML}(\tau, \mathbb{R})$ has measure 0. Then

$$Y = \bigcup_{\tau} \text{int}(Y_\tau) \subset \mathcal{S}$$

is an open subset of W whose complement has measure 0. As in the proof of [Theorem 10.1](#), it is easy to show that

$$(11.5) \quad \frac{\#(Y^c \cap \mathcal{ML}(\partial H, \mathbb{Z}) \cap rV)}{\#(\mathcal{ML}(\partial H, \mathbb{Z}) \cap rV)} \rightarrow 0 \quad \text{as } r \rightarrow \infty,$$

since the top and bottom converge to the Lebesgue measure of $Y^c \cap V$ and V respectively. This implies [\(11.4\)](#) and hence the theorem. \square

References

- [1] **I Agol, J Hass, W Thurston**, *The computational complexity of knot genus and spanning area*, Trans. Amer. Math. Soc. 358 (2006) 3821–3850 [MR2219001](#)
- [2] **J Berge**, *Documentation for the program Heegaard*, preprint (c1990) Available at <http://www.math.uic.edu/~t3m/>
- [3] **R Bieri, W D Neumann, R Strebel**, *A geometric invariant of discrete groups*, Invent. Math. 90 (1987) 451–477 [MR914846](#)
- [4] **J S Birman**, *Braids, links, and mapping class groups*, Annals of Mathematics Studies 82, Princeton University Press (1974) [MR0375281](#)
- [5] **A Borisov, M Sapir**, *Polynomial maps over finite fields and residual finiteness of mapping tori of group endomorphisms* [arXiv:math.GR/0309121](#)
- [6] **K S Brown**, *Trees, valuations, and the Bieri-Neumann-Strebel invariant*, Invent. Math. 90 (1987) 479–504 [MR914847](#)
- [7] **J O Button**, *Fibred and virtually fibred hyperbolic 3-manifolds in the censuses*, Experiment. Math. 14 (2005) 231–255 [MR2169525](#)
- [8] **P J Callahan, M V Hildebrand, J R Weeks**, *A census of cusped hyperbolic 3-manifolds, with microfiche supplement*, Math. Comp. 68 (1999) 321–332 [MR1620219](#)
- [9] **N M Dunfield**, *Alexander and Thurston norms of fibered 3-manifolds*, Pacific J. Math. 200 (2001) 43–58 [MR1863406](#)
- [10] **N M Dunfield, W P Thurston**, *Finite covers of random 3-manifolds*, Invent. Math. 166 (2006) 457–521
- [11] **M Dwass**, *Simple random walk and rank order statistics*, Ann. Math. Statist. 38 (1967) 1042–1053 [MR0215463](#)

- [12] **M Gromov**, *Hyperbolic groups*, from: “Essays in group theory”, Math. Sci. Res. Inst. Publ. 8, Springer (1987) 75–263 [MR919829](#)
- [13] **J Hoste**, **M Thistlethwaite**, *Knotscape* (1999) Available at <http://www.math.utk.edu/~morwen>
- [14] **J Hoste**, **M Thistlethwaite**, **J Weeks**, *The first 1,701,936 knots*, Math. Intelligencer 20 (1998) 33–48 [MR1646740](#)
- [15] **W Jaco**, *Adding a 2-handle to a 3-manifold: an application to property R*, Proc. Amer. Math. Soc. 92 (1984) 288–292 [MR754723](#)
- [16] **W Jaco**, **JL Tollefson**, *Algorithms for the complete decomposition of a closed 3-manifold*, Illinois J. Math. 39 (1995) 358–406 [MR1339832](#)
- [17] **V A Kaimanovich**, *Poisson boundary of discrete groups*, preprint (2001)
- [18] **V A Kaimanovich**, **H Masur**, *The Poisson boundary of Teichmüller space*, J. Funct. Anal. 156 (1998) 301–332 [MR1636940](#)
- [19] **S P Kerckhoff**, *Simplicial systems for interval exchange maps and measured foliations*, Ergodic Theory Dynam. Systems 5 (1985) 257–271 [MR796753](#)
- [20] **C J Leininger**, *Surgeries on one component of the Whitehead link are virtually fibered*, Topology 41 (2002) 307–320 [MR1876892](#)
- [21] **F Luo**, **R Stong**, *Dehn-Thurston coordinates for curves on surfaces*, Comm. Anal. Geom. 12 (2004) 1–41 [MR2074869](#)
- [22] **H Masur**, *Measured foliations and handlebodies*, Ergodic Theory Dynam. Systems 6 (1986) 99–116 [MR837978](#)
- [23] **M Mirzakhani**, *Growth of the number of simple closed geodesics on hyperbolic surfaces*, preprint (2004) to appear in Ann. of Math.
- [24] **K Murasugi**, *On a certain subgroup of the group of an alternating link*, Amer. J. Math. 85 (1963) 544–550 [MR0157375](#)
- [25] **A Y Ol’shanskii**, *Almost every group is hyperbolic*, Internat. J. Algebra Comput. 2 (1992) 1–17 [MR1167524](#)
- [26] **R C Penner**, **JL Harer**, *Combinatorics of train tracks*, Annals of Mathematics Studies 125, Princeton University Press (1992) [MR1144770](#)
- [27] **D Poulalhon**, **G Schaeffer**, *Optimal coding and sampling of triangulations*, from: “Automata, languages and programming”, Lecture Notes in Comput. Sci. 2719, Springer, Berlin (2003) 1080–1094 [MR2080766](#)
- [28] **S Rankin**, **O Flint**, **J Schermann**, *Enumerating the prime alternating knots. I*, J. Knot Theory Ramifications 13 (2004) 57–100 [MR2039118](#)
- [29] **S Rankin**, **O Flint**, **J Schermann**, *Enumerating the prime alternating knots. II*, J. Knot Theory Ramifications 13 (2004) 101–149 [MR2039119](#)

- [30] **G Schaeffer, P Zinn-Justin**, *On the asymptotic number of plane curves and alternating knots*, Experiment. Math. 13 (2004) 483–493 [MR2118273](#)
- [31] **M Scharlemann**, *Outermost forks and a theorem of Jaco*, from: “Combinatorial methods in topology and algebraic geometry (Rochester, N.Y. 1982)”, Contemp. Math. 44, Amer. Math. Soc., Providence, RI (1985) 189–193 [MR813113](#)
- [32] **S Schleimer**, *Almost normal Heegaard splittings*, PhD thesis, Berkeley (2001)
- [33] **J Stallings**, *On fibering certain 3-manifolds*, from: “Topology of 3-manifolds and related topics (Proc. The Univ. of Georgia Institute, 1961)”, Prentice-Hall (1962) 95–100 [MR0158375](#)
- [34] **J M Steele**, *The Bohnenblust-Spitzer algorithm and its applications*, Probabilistic methods in combinatorics and combinatorial optimization, J. Comput. Appl. Math. 142 (2002) 235–249 [MR1910531](#)
- [35] **J L Tollefson, N Wang**, *Taut normal surfaces*, Topology 35 (1996) 55–75 [MR1367275](#)
- [36] **G S Walsh**, *Great circle links and virtually fibered knots*, Topology 44 (2005) 947–958 [MR2153978](#)
- [37] **J Weeks**, *SnapPea* Available at <http://www.geometrygames.org>

*Mathematics 253-37, California Institute of Technology
Pasadena, CA 91125, USA*

*Mathematics, Barnard College, Columbia University MC 4436
New York, NY 10027, USA*

dunfield@caltech.edu, dthurston@barnard.edu

Proposed: Cameron Gordon
Seconded: Rob Kirby, Joan Birman

Received: 8 April 2006
Accepted: 13 November 2006