



Geometry & Topology

Volume 28 (2024)

Taut foliations, left orders, and pseudo-Anosov mapping tori

JONATHAN ZUNG

Taut foliations, left orders, and pseudo-Anosov mapping tori

JONATHAN ZUNG

For a large class of 3–manifolds with taut foliations, we construct an action of $\pi_1(M)$ on \mathbb{R} by orientation-preserving homeomorphisms which captures the transverse geometry of the leaves. This action is complementary to Thurston’s universal circle. Applications include the left-orderability of the fundamental groups of every nontrivial surgery on the figure-eight knot. Our techniques also apply to at least 2598 manifolds representing 44.7% of the non-L-space rational homology spheres in the Hodgson–Weeks census of small closed hyperbolic 3–manifolds.

20F60, 37C85, 57R30

1. Introduction	4191
2. Notation and conventions	4194
3. Foliations on surgeries on pseudo-Anosov mapping tori	4197
4. Structure of branching in L	4198
5. Gluing branches	4204
6. The flat connection $\hat{\mathcal{J}}$	4209
7. Computations	4228
8. Remarks and questions	4229
References	4230

1 Introduction

A taut foliation on a 3–manifold M is a valuable structure. Taut foliations may be used to certify nontriviality of a transverse loops or to certify surfaces of minimum genus in their homology classes, as shown in classic work of Novikov [22], Roussarie [25], Thurston [26] and Gabai [17]. The existence of a taut foliation puts constraints on $\pi_1(M)$. For example, Thurston showed that if M is an atoroidal, irreducible 3–manifold with a taut foliation, then $\pi_1(M)$ admits a faithful action on S^1 by homeomorphisms; see Thurston [27] and Calegari and Dunfield [9]. In another direction, Kronheimer and Mrowka [20; 21] made a connection to Floer theory via contact and symplectic geometry, showing that a taut foliation

gives rise to a nontrivial class in monopole Floer homology. Ozsváth and Szabó [23] established parallel results in the setting of Heegaard Floer homology.

The L-space conjecture, formulated in parts by Boyer, Gordon and Watson [4] and Ozsváth and Szabó [23], is a proposed sharpening of the connections outlined above. It posits that the following are equivalent for an orientable, irreducible rational homology sphere M :

- (i) M has a co-orientable taut foliation.
- (ii) $\pi_1(M)$ is left-orderable, ie $\pi_1(M)$ acts faithfully on \mathbb{R} by orientation-preserving homeomorphisms.
- (iii) M is not an L-space, ie its Heegaard Floer homology $\widehat{\text{HF}}(M)$ satisfies the strict inequality $\text{rank}(\widehat{\text{HF}}(M)) > |H^2(M; \mathbb{Z})|$.

Technology for deciding conditions (i) and (iii) is well developed; for example, Dunfield verified the equivalence of (i) and (iii) for 99.8% of the manifolds in his census of $\sim 300,000$ small hyperbolic rational homology spheres [13]. Techniques for deciding (ii) are harder to come by. For every non-left-orderable group, there is a finite length certificate proving its non-left-orderability. On the other hand, left-orderability is not known to be decidable for three-manifold groups. See Calegari and Dunfield [9] for a discussion. Here are a few practical methods for proving the left-orderability of a 3-manifold group:

- Try to lift $\text{PSL}(2, \mathbb{R})$ representations of $\pi_1(M)$ to $\widetilde{\text{PSL}}(2, \mathbb{R})$, which acts on the universal cover of the circle at infinity in \mathbb{H}^2 . See Eisenbud, Hirsch and Neumann [14] or Culler and Dunfield [11] for examples of this technique.
- If M has a taut foliation, try to lift Thurston's action of $\pi_1(M)$ on S^1 to an action on \mathbb{R} . This works whenever the Euler class of the plane field tangent to the foliation vanishes. See Calegari and Dunfield [9, Section 7] or Boyer and Hu [5, Section 5] for examples of this technique.
- We say that a foliation \mathcal{F} is \mathbb{R} -covered if the leaf space of the lift of \mathcal{F} to the universal cover of M is homeomorphic to \mathbb{R} (eg in the case of the foliation of a fibered 3-manifold by fiber surfaces). Since $\pi_1(M)$ always acts on the leaf space, we get an action on \mathbb{R} .

The third technique is appealing since it directly uses the transverse geometry of the foliation, but is limited in generality since most taut foliations are not \mathbb{R} -covered. In this paper, we demonstrate a method for improving the third technique to work for more general taut foliations. We study a family of 3-manifolds with taut foliations which are not \mathbb{R} -covered, but whose leaf spaces admit a map to \mathbb{R} such that the action of π_1 descends to an action on \mathbb{R} . The question of the existence of taut foliations and left-orderings compatible in this sense was first raised by Thurston [27]; see also Calegari [7, Section 8.1].

We prove the following:

Theorem 1.1 *Let Σ be an orientable closed surface and $\varphi: \Sigma \rightarrow \Sigma$ a pseudo-Anosov map with orientable invariant foliations. Suppose further that φ preserves the orientation of these foliations. Let M_φ be the mapping torus of φ . Let $M_\varphi(\mathbf{p}; \mathbf{q})$ be the result of nonzero surgery along any collection of closed orbits of φ . If the surgery slopes all have the same sign, then $M_\varphi(\mathbf{p}; \mathbf{q})$ has left-orderable fundamental group.*

Here we take the zero slope, also known as the *degeneracy slope*, to be the one which crosses no prongs. See Convention 2.4 for a full explanation of our slope conventions.

Theorem 1.2 *With the assumptions of Theorem 1.1, there is a taut foliation \mathcal{F} on $M_\varphi(\mathbf{p}; \mathbf{q})$. Let L be the leaf space of $\tilde{\mathcal{F}}$. Then there is a continuous, monotone map $f: L \rightarrow \mathbb{R}$ such that the action of $\pi_1(M_\varphi(\mathbf{p}; \mathbf{q}))$ on L descends to a nontrivial action on \mathbb{R} .*

Here monotone means that f respects the natural partial order on L , ie if leaf λ_1 may be connected to leaf λ_2 by an oriented arc positively transverse to $\tilde{\mathcal{F}}$, then $f(\lambda_1) \leq f(\lambda_2)$.

Corollary 1.3 *For any $n \geq 1$, any nontrivial surgery on the n -fold cyclic branched cover of the figure-eight knot has left-orderable fundamental group.*

Previously, orderability for surgeries on the figure-eight knot was known only for slopes in $[-4, 4] \cup \mathbb{Z}$. The range $(-4, 4)$ was treated by the representation-theoretic approach of Boyer, Gordon and Watson [4], while the toroidal exceptional surgeries $\{4, -4\}$ were resolved by Clay, Lidman and Watson [10] using a gluing theorem for amalgamations of left-orderable groups. Fenley's work [15] on \mathbb{R} -covered Anosov flows applies to show that integer surgeries on the figure-eight knot are left-orderable. Hu [19] recently gave another approach to the case of integer slopes by showing that certain taut foliations on these manifolds have vanishing Euler class. Hu also proves a negative result: for slopes outside $[-2, 2] \cup \mathbb{Z}$, there does not exist a co-orientable taut foliation which both has trivial Euler class and is transverse to the Dehn surgery core. Thus, the action of the universal circle on such foliations does not lift to an action on \mathbb{R} . We discuss this further in Remark 3.4.

In Section 3, we set up notation and construct taut foliations on the manifolds of Theorem 1.1. In Section 4, we analyze the branching in the leaf spaces of these foliations.

Our approach to defining the map $f: L \rightarrow \mathbb{R}$ is to glue together certain branches of the leaf space L . This point of view is outlined in Section 5. More formally, in Section 6 we define an \mathbb{R} -bundle with structure group $\text{Homeo}^+(\mathbb{R})$ and a flat *partial connection* $\mathcal{J}^{\text{partial}}$. We complete this bundle by adding a point at infinity to each fiber. The resulting S^1 bundle has an honest flat connection \mathcal{J} . Moreover, this S^1 bundle has vanishing Euler class and hence lifts to a flat \mathbb{R} bundle.

Finally, in Section 7 we report on computations showing that manifolds satisfying the hypotheses of Theorem 1.1 are abundant in the Hodgson–Weeks census.

Acknowledgements The author would like to thank David Gabai and Sergio Fenley for several helpful discussions about this work. Nathan Dunfield and Mark Bell generously shared data from their census of monodromies of small hyperbolic manifolds. A special thanks is owed to the referee, whose exceptionally detailed and patient reading resulted in many corrections and improvements to the article. The author is indebted to Peter Ozsváth for his consistent encouragement and guidance during this project.

2 Notation and conventions

In this section, we set up notation that will be used throughout the article. Let Σ be a closed orientable surface. Let $\varphi: \Sigma \rightarrow \Sigma$ be a pseudo-Anosov map such that its stable and unstable foliations, denoted by $\mathcal{F}^s|_\Sigma$ and $\mathcal{F}^u|_\Sigma$, are orientable. This implies that each singularity of φ has an even number of prongs. Suppose further that φ preserves the orientations of $\mathcal{F}^s|_\Sigma$ and $\mathcal{F}^u|_\Sigma$.

Let M_φ denote the mapping torus of φ , and let $K_1, \dots, K_n \subset M_\varphi$ be the suspensions of any n periodic orbits of φ . For ease of exposition, we will always assume that the suspensions of the singularities of φ are included in K_1, \dots, K_n .

Let $M_\varphi^\circ = M_\varphi \setminus \{K_i\}$. We will use $M_\varphi(\mathbf{p}; \mathbf{q})$ to denote the manifold obtained by slope $(p_i; q_i)$ surgery along K_i . Let $K_i(\mathbf{p}; \mathbf{q})$ denote the core of the Dehn filling of K_i in $M_\varphi(\mathbf{p}; \mathbf{q})$. Our slope conventions are explained below in Convention 2.4.

Let $\Lambda|_\Sigma$ be the stable invariant lamination of φ produced by splitting open $\mathcal{F}^s|_\Sigma$ along the prongs at each singularity. Let \mathcal{F}^s , \mathcal{F}^u and Λ be the suspensions of $\mathcal{F}^s|_\Sigma$, $\mathcal{F}^u|_\Sigma$ and $\Lambda|_\Sigma$ in M_φ . (To be precise, we should take the suspension of $\Lambda|_\Sigma$ using an appropriate blowup of the flow φ .) The orientability constraints are equivalent to the orientability of \mathcal{F}^s .

Example 2.1 The figure-eight knot complement fibers over the circle with a genus 1 fiber and pseudo-Anosov monodromy. We can choose coordinates on the fiber $T^2 \setminus (0, 0)$ so that the monodromy is $\begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix}$. Since this matrix has distinct positive real eigenvalues, the monodromy preserves the orientation of the invariant foliations as desired.

Proof of Corollary 1.3 Consider the n^{th} cyclic branched cover of the figure-eight knot. The lift of the figure-eight knot to this cyclic branched cover is a fibered knot with monodromy $\begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix}^n$. With our slope conventions, surgery along the zero slope (ie the degeneracy slope) yields the n^{th} cyclic branched cover. By Theorem 1.1, surgery along any nonzero slope yields a manifold with left-orderable fundamental group. \square

Example 2.2 We can generate examples with a given fiber genus and prescribed singularities by enumerating periodic splitting sequences of train tracks, as in [24]. One of the lowest volume examples appearing in the genus 2, 1–singularity enumeration is the 1–cusped hyperbolic manifold m038.

Example 2.3 The $(-2, 3, 7)$ pretzel knot is fibered with a genus 5 fiber and pseudo-Anosov monodromy. In this case, the monodromy φ has a single 18–prong singularity at the boundary of the fiber, so the invariant foliations are orientable. However, φ reverses this orientation. The branched double cover of the $(-2, 3, 7)$ pretzel knot does preserve the orientation of the invariant foliations, and so satisfies the given condition.

Convention 2.4 (slope conventions) For each i , let m_i be the period of the orbit K_i . The m_i singularities in $K_i \cap \Sigma$ have the same number, denoted by k_i , of prongs. We assume throughout that k_i is even. Define ω_i to be the integer in $[0, k_i)$ such that $2\pi\omega_i$ is the counterclockwise angle by which φ^{m_i} rotates one of these singularities. We are thinking of each singularity as a cone point with angle $2\pi k_i$.

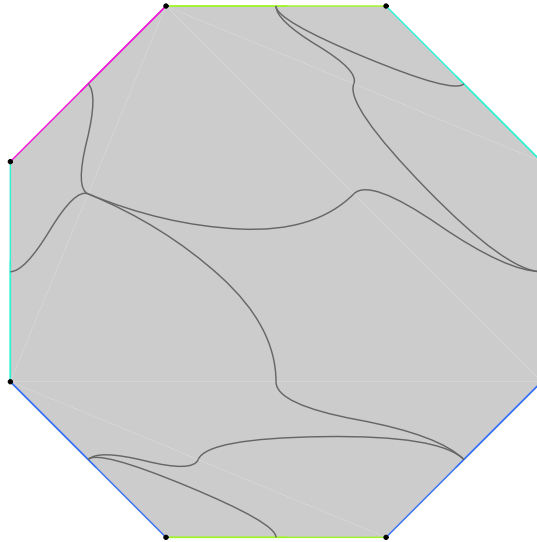


Figure 1: This figure shows a train track carrying an invariant lamination for the monodromy of the fibered, 1-cusped 3-manifold $m038$. The fiber surface has genus 2 with one puncture at the vertices of the octagon. The complement of the train track is a punctured ideal hexagon. The figure was generated using Mark Bell's program `flipper` [2].

We describe slopes in a slightly nonstandard way. A slope of $(p_i; q_i)$ corresponds to a curve in $\partial N(K_i)$ defined as follows. Choose a point near $K_i \cap \Sigma$ and flow it along the suspension flow of φ for time $m_i q_i$. (Here we set the speed of the flow that the time-1 flow of the suspension flow intersects each fiber once and returns to the starting fiber.) Typically, the resulting curve will not close up since φ^{m_i} may rotate singularities. We close it by appending a path in Σ which walks around the relevant singularity by a clockwise angle of $2\pi p_i$. Not all such pairs represent slopes; a pair $(p_i; q_i)$ corresponds to a closed curve if and only if $p_i = \omega_i q_i \pmod{k_i}$.

We always assume $q_i \geq 0$. We say that a slope is positive (resp. negative) when p_i is positive (resp. negative). We say that the slope is ∞ when $q_i = 0$. The ∞ slope corresponds to the fiber slope and is declared to be both positive and negative. The zero slope or the degeneracy slope is $(p_i; q_i) = (0; k_i / \gcd(\omega_i, k_i))$. In an abuse of notation, the zero slope and the infinity slope may intersect more than once.

\mathcal{F} will denote the taut foliation to be constructed in Construction 3.2. Let \mathcal{G} denote the guts of Λ , ie the compact subspace of $M \setminus \Lambda$ obtained by chopping off the ends of the ideal polygon bundles. Let \mathcal{I} denote the interstitial region, ie the part of $M \setminus \Lambda$ we just cut off. Topologically, \mathcal{I} is a disjoint union of I-bundles over half-infinite cylinders. We have $M_\varphi(\mathbf{p}; \mathbf{q}) = \Lambda \cup \mathcal{G} \cup \mathcal{I}$. Note that the distinct components of \mathcal{G} have disjoint closures.

It will also be convenient to work with another decomposition $M_\varphi(\mathbf{p}; \mathbf{q}) = \Lambda' \cup \mathcal{G}'$. We obtain Λ' by blowing down \mathcal{I} and define \mathcal{G}' to be the closure of $M_\varphi(\mathbf{p}; \mathbf{q}) \setminus \Lambda'$. Some of the leaves of Λ' are branched




Figure 2: Blowing down \mathcal{I} . The shaded region on the left is \mathcal{G} , and the shaded region on the right is \mathcal{G}' .

surfaces instead of surfaces, so Λ' is not a lamination, but a branched lamination. The advantage of Λ' over Λ is that there is a Solv metric supported on Λ' .

The following lemma will be useful later:

Lemma 2.5 *Let S be a fiber surface in the fibered manifold $M_\varphi(\mathbf{p}; \mathbf{q}) \setminus \mathcal{G}$. Suppose γ is a path in $S \cap (M_\varphi(\mathbf{p}; \mathbf{q}) \setminus \mathcal{G})$ with endpoints on two walls of \mathcal{G} . Suppose that γ makes N_{tan} tangencies with \mathcal{F} . Then there is an arc γ' in $S \cap (M_\varphi(\mathbf{p}; \mathbf{q}) \setminus \mathcal{G}')$ with endpoints on the corresponding two walls of \mathcal{G}' and in the same homotopy class rel endpoints as γ such that γ' makes at most N_{tan} tangencies with Λ' . (Here the homotopy happens in $S \cap (M_\varphi(\mathbf{p}; \mathbf{q}) \setminus \mathcal{G})$ and we are identifying relative homotopy classes of paths in $(S \cap (M_\varphi(\mathbf{p}; \mathbf{q}) \setminus \mathcal{G}), S \cap \partial\mathcal{G})$ and $(S \cap (M_\varphi(\mathbf{p}; \mathbf{q}) \setminus \mathcal{G}'), S \cap \partial\mathcal{G}')$ in the obvious way.) The converse also holds.*

Proof First, homotope γ so that all of its tangencies with \mathcal{F} and all of its self-intersections occur on the interior of \mathcal{I} . This may be done without introducing any new tangencies. Let \mathcal{I}_ε be the subset of \mathcal{I} of thickness less than ε . We begin by blowing down \mathcal{I}_ε for some $\varepsilon > 0$. Let \mathcal{F}_ε and Λ_ε be the images of \mathcal{F} and Λ under this blowdown.

If we choose ε small enough, then γ intersects $\mathcal{F} \cap \mathcal{I}_\varepsilon$ transversely. Moreover, if we choose ε small enough, then the blowdown changes the tangent plane field of \mathcal{F} by a uniformly small amount. Thus, we can arrange that γ has N_{tan} tangencies with \mathcal{F}_ε .

Now, homotope γ so that all of its tangencies with \mathcal{F}_ε and its self intersections occur inside Λ_ε . At this stage, $\gamma \setminus \Lambda_\varepsilon$ is a collection \mathcal{C} of disjoint segments transverse to \mathcal{F}_ε . We now blow down the rest of \mathcal{I} so that each of the segments in \mathcal{C} blows down to a point. The curve γ then blows down to a curve γ' with the desired properties.

The converse is easier: Λ is obtained by splitting open the prongs in Λ' , and this may be done without introducing new tangencies with γ' . \square




Figure 3: On the left is the standard foliation of $D^2 \times S^1$, where S^1 is cut open. In the middle, we alternately comb the edges of the disks to expose their positive and negative sides. Grey shows the positive sides of the leaves, while white shows the negative sides. On the right, we show the limiting configuration which has 4 annular leaves at the boundary, which we call *walls*, and infinitely tall saddle-like leaves (homeomorphic to planes) on the interior. The interior saddle-like leaves accumulate on the walls. We can vary the number of legs of the saddle’s rider or the gluing of the top and bottom of the picture to get sutures of any desired nonmeridional slope.

3 Foliations on surgeries on pseudo-Anosov mapping tori

In this section, we describe the construction of taut foliations on our class of 3–manifolds. One might colloquially describe the construction as “stuffing the guts of the suspension of the φ –invariant lamination with monkey saddles”.

Construction 3.1 Given a sutured solid torus D such that the sutures are parallel with nonmeridional slope, there is a foliation of D by planes compatible with the sutures. Recall that the sutures on a sutured manifold divide its boundary into (possibly disconnected) positive and negative subsurfaces. We can construct the desired foliation beginning with the obvious product foliation of the solid torus by disks, and then combing the edges of the disks to expose their positive sides in the positive regions of ∂D and their negative sides in the negative regions of ∂D . This is called a foliation by a stack of monkey saddles. See Figure 3.

Construction 3.2 Let $(p_i; q_i)$ be any choice of slopes with $p_i \neq 0$. Then the manifold $M_\varphi(\mathbf{p}; \mathbf{q})$ carries a taut foliation constructed as follows. Recall that $M_\varphi^\circ = M_\varphi \setminus \{K_i\}$. Let \mathcal{F}^s be the codimension-1 weak stable foliation of the suspension flow of φ on M_φ° . Split open \mathcal{F}^s along its prongs to obtain a lamination Λ on $M_\varphi^\circ \subset M_\varphi(\mathbf{p}; \mathbf{q})$. The complement $M_\varphi(\mathbf{p}; \mathbf{q}) \setminus \Lambda$ is a collection of ideal polygon bundles over S^1 . With our conventions (see Convention 2.4), these polygons are $2p_i$ –gons since a closed curve of slope $(p_i; q_i)$ decomposes into an arc of slope $(0; q_i)$, which intersects no prongs, and an arc of slope $(p_i; 0)$, which wraps around an angle of $2\pi p_i$ and therefore intersects $2p_i$ prongs. Fill in these bundles with the foliations constructed in Construction 3.1. The resulting foliation is taut because it contains no compact leaves; in fact, all leaves are either cylinders or planes.

Example 3.3 When $p_i = 1$ for all i , the complementary regions can be blown down without inserting any leaves at all. The resulting 3–manifold carries an Anosov flow. Fried [16] shows that every transitive Anosov flow with orientable invariant foliations is obtained by this construction. The transitivity condition

is not severe; every Anosov flow on a hyperbolic 3-manifold is transitive. His construction might require the slopes to have different signs, so our results will not hold for all of these manifolds.

Remark 3.4 The Euler class of a foliation constructed via Construction 3.2 is typically nonzero, so the action of π_1 on Thurston's universal circle will not in general lift to an action on \mathbb{R} . One can see this by applying a result of Hu:

Let X be a \mathbb{Q} -homology solid torus, $k \geq 1$ the order of the longitude of X in $H_2(X, \partial X)$. Let μ be any meridian. Let S be the set of filling slopes for X yielding a 3-manifold with a co-orientable taut foliation which is transverse to the Dehn surgery core and has zero Euler class. Let x be the Thurston norm of a generator of $H_2(X, \partial X)$. If $x \neq 0$, then [19, Theorem 5.2] says that

- (i) outside the interval $(-x/k - 1, x/k + 1)$, S may contain only μ and the integer slopes;
- (ii) S is nowhere dense in $\mathbb{R} \cup \{1/0\}$. In particular, it is nowhere dense in $(-x/k - 1, x/k + 1)$.

To apply Hu's theorem, fill all but one of the boundary components of M_φ° . Since M_φ° is a pseudo-Anosov suspension, it is a cusped hyperbolic manifold. Therefore, for a generic choice of slopes, the result of filling is a 1-cusped hyperbolic manifold which we call X . Suppose further that X is a \mathbb{Q} -homology solid torus; otherwise, any filling of X has $b_1 > 0$ and we can find a left-ordering by other means. Let x be the Thurston norm of a generator of $H_2(X, \partial X)$. By hyperbolicity, $x \neq 0$. Therefore, X satisfies the hypothesis of Hu's theorem. The foliations of Construction 3.2 are transverse to the Dehn surgery core because the core is transverse to the monkey saddles. Thus, in a generic choice of filling slopes for X , the foliations of Construction 3.2 have nonzero Euler class.

As an example, take the figure-eight knot complement. In this case, $x = 1$. By Hu's result, our foliations have nonzero Euler class for filling slopes outside $(-2, 2) \cup \mathbb{Z}$.

4 Structure of branching in L

Let L be the leaf space of the lift $\tilde{\mathcal{F}}$ of \mathcal{F} to $\tilde{M}_\varphi(\mathbf{p}; \mathbf{q})$. In this section, we roughly prove that "all branching in L happens in the saddle regions". These results are not logically required for the proof of the main theorem (and indeed do not hold in full generality), but are interesting in their own right and provide motivation for subsequent constructions.

In this section, we make the following assumption:

Assumption 4.1 *All the monkey saddles used in the construction of \mathcal{F} have at least four sides.*

In the case where all the monkey saddles have two sides, the resulting manifolds carry Anosov flows. The branching in their invariant foliations was analyzed in [15].

Recall that for a taut foliation, the leaf space of the universal cover is a simply connected (possibly non-Hausdorff) 1-manifold. Each leaf of the universal cover is homeomorphic to a plane, and its stabilizer under the action of π_1 by deck transformations is the fundamental group of the projected leaf in $M_\varphi(\mathbf{p}; \mathbf{q})$.

Two points in L are said to be *nonseparated* if they coincide in the Hausdorffification of L . A *branch locus* is maximal set of at least two points in L that are pairwise nonseparable; see [8, Chapter 4].

Proposition 4.2 *Assume Assumption 4.1. For each orbit K_i there are two branch loci called B_i^+ and B_i^- (unique up to covering transformations) each of which is finite and has p_i points corresponding to coherent lifts of the positively or negatively oriented walls of the filling saddle region. Moreover, these are all of the branch loci of $\tilde{\mathcal{F}}$ up to covering transformations.*

Proof The core of a saddle region is a curve transverse to \mathcal{F} and is therefore noncontractible. Therefore, the S^1 worth of saddle-like leaves in a saddle region lifts to an \mathbb{R} worth of saddle-like leaves (plus translates thereof) in $\tilde{\mathcal{F}}$. This \mathbb{R} worth of leaves limits to lifts of the positive (resp. negative) walls in the $+\infty$ (resp. $-\infty$) direction, so the positively (resp. negatively) oriented walls of a saddle region form a branch locus. We choose one of the translates of this branch locus and call it B_i^+ (resp. B_i^-).

To show that there is no branching elsewhere, it suffices to show that any curve can be “pulled tight” relative to \mathcal{F} so that it is transverse to \mathcal{F} except for controlled intervals in \mathcal{G} . Roughly speaking, our strategy will be to pull γ tight relative to the natural Solv metric on Λ' .

We say that an arc in \mathcal{G} with endpoints on the interior of the walls in $\partial\mathcal{G}$ is *inessential* if it can be compressed into a wall of $\partial\mathcal{G}$, and *essential* otherwise. An essential arc with endpoints in oppositely oriented walls can be homotoped in \mathcal{G} rel boundary to be transverse to \mathcal{F} , while an essential arc with endpoints in similarly oriented walls can be homotoped in \mathcal{G} rel boundary to have a single tangency with \mathcal{F} . An essential arc of the latter type lifts to a short curve in L that connects two points in B_i^\pm . It suffices to show that any curve γ with endpoints on leaves λ_1, λ_2 can be homotoped in $M_\varphi(\mathbf{p}; \mathbf{q})$ relative to its endpoints so that either

- (i) γ is transverse to \mathcal{F} except for finitely many tangencies, each of which is contained in an essential subarc of γ in \mathcal{G} , or
- (ii) $\lambda_1 = \lambda_2$ and γ lies inside λ_1 .

We call such a curve *efficient*. Moreover, we can without loss of generality assume that λ_1 and λ_2 are walls of \mathcal{G} . For, suppose that $\tilde{\lambda}_1$ and $\tilde{\lambda}_2$ are nonseparated leaves of $\tilde{\mathcal{F}}$ in a branch locus not equal to one of the B_i^\pm . We may choose $\tilde{\lambda}_1$ and $\tilde{\lambda}_2$ to be adjacent, meaning that there is a family of leaves $\{\tilde{\lambda}_t\}_{t>0}$ limiting to both $\tilde{\lambda}_1$ and $\tilde{\lambda}_2$. Adjacency is a stronger condition than nonseparatedness, but a pair of nonseparated leaves can be connected by a sequence of adjacent leaves; we direct the reader to [8, Example 4.44, page 170] for further discussion. Without loss of generality, assume that $\tilde{\lambda}_t$ approaches $\tilde{\lambda}_1$ and $\tilde{\lambda}_2$ from above in L as $t \rightarrow 0$. Any path $\tilde{\gamma} \subset \tilde{M}_\varphi(\mathbf{p}; \mathbf{q})$ from $\tilde{\lambda}_1$ to $\tilde{\lambda}_2$ descends to a path γ in $M_\varphi(\mathbf{p}; \mathbf{q})$ that cannot be homotoped in $M_\varphi(\mathbf{p}; \mathbf{q})$ rel endpoints to become efficient. Since \mathcal{F} is taut, we can augment the beginning of γ with a descending transversal from λ_1 to a wall of \mathcal{G} . Similarly, we can augment the end of γ with a descending transversal from λ_2 to a wall of \mathcal{G} . This resulting path is also not homotopic in $M_\varphi(\mathbf{p}; \mathbf{q})$ to an efficient one, but has endpoints on walls of \mathcal{G} .

Given any arc γ with endpoints in leaves λ_1, λ_2 which are walls of \mathcal{G} , we may homotope it in $M_\varphi(\mathbf{p}; \mathbf{q})$ so that it intersects $\partial\mathcal{G}$ transversely on the interior of the walls. Define

- N_{tan} as the number of tangencies between γ and \mathcal{F} in $M_\varphi(\mathbf{p}; \mathbf{q}) \setminus \mathcal{G}$,
- N_{ess} as the number of components of $\gamma \cap \mathcal{G}$ that are essential,
- N_{iness} as the number of components of $\gamma \cap \mathcal{G}$ that are inessential.

Define the complexity function

$$I(\gamma) = N_{\text{ess}} + 1.01N_{\text{tan}} + 1.02N_{\text{iness}}.$$

Now choose γ in its relative homotopy class to minimize $I(\gamma)$. With this choice, we claim that $N_{\text{tan}} = N_{\text{iness}} = 0$.

Suppose $N_{\text{iness}} > 0$. Then we can compress an inessential arc in \mathcal{G} . This decreases N_{iness} by 1 at the cost of increasing N_{tan} by one, violating the minimality assumption.

Now suppose that $N_{\text{tan}} > 0$. Then there is a component γ_0 of $\gamma \cap (M_\varphi(\mathbf{p}; \mathbf{q}) \setminus \mathcal{G})$ containing a tangency with \mathcal{F} . The suspension flow of φ blows up to a flow X on $M_\varphi(\mathbf{p}; \mathbf{q}) \setminus \mathcal{G}$ which

- (i) is transverse to the fibering $M_\varphi(\mathbf{p}; \mathbf{q}) \setminus \mathcal{G} \rightarrow S^1$,
- (ii) preserves Λ , and
- (iii) exits through the interstitial annuli (ie the annuli in $\partial\mathcal{G} \cap \partial\mathcal{I}$) in forward time and exists for all backwards time.

We can arrange that γ_0 lies in a single fiber surface S of the fibering. We can do this by pushing γ_0 backwards along the flowlines of X into some choice of fiber surface S . This possibly slides $\gamma_0(0)$ and $\gamma_0(1)$ along their respective walls of \mathcal{G} , but does not change the complexity function. Using Lemma 2.5, we may replace γ_0 with an arc γ'_0 having the properties that

- (i) γ'_0 is in the support of Λ' ,
- (ii) γ'_0 is homotopic to γ_0 rel endpoints, and
- (iii) the number of tangencies of γ_0 with \mathcal{F} is equal to the number of tangencies of γ'_0 with Λ' .

The pseudo-Anosov structure of φ gives rise to a locally Euclidean Riemannian metric on the surface with boundary $\Lambda' \cap S$. Pull γ'_0 tight relative to this metric. Note that this tightening happens completely inside the surface S . Since the leaves of $\Lambda' \cap S$ are geodesics (ie locally straight lines) with respect to this metric, the number of tangencies between γ'_0 and Λ' does not increase during tightening. We have the following cases:

- γ'_0 tightens to a geodesic and is not contained in a single leaf of Λ' . Since the leaves of $\Lambda' \cap \Sigma$ are geodesic, γ'_0 cannot have any tangencies with leaves of Λ' . Thus, using Lemma 2.5 we can also homotope γ_0 to be transverse to \mathcal{F} and decrease N_{tan} .

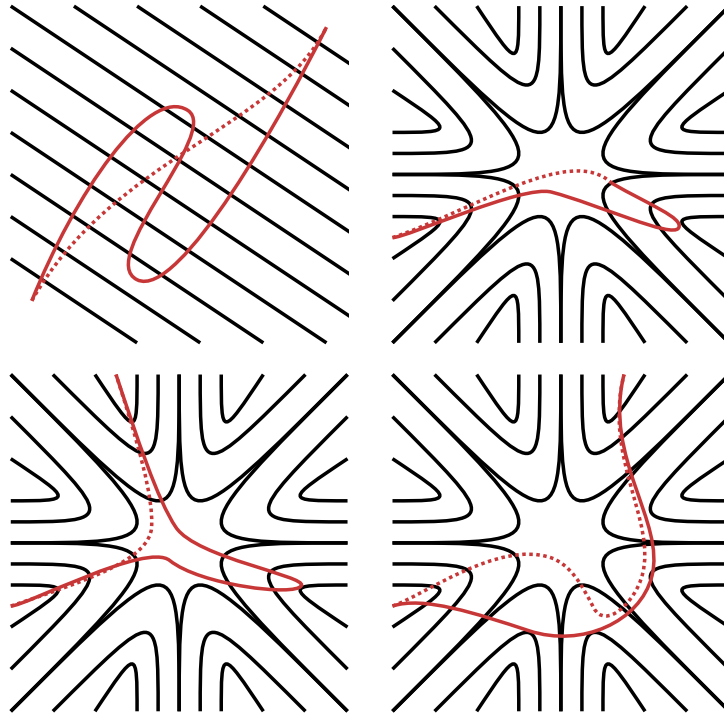


Figure 4: Various tightening moves.

- γ'_0 tightens to an arc contained in a wall of $\partial\mathcal{G}'$. Note that distinct components of the guts have distinct boundary leaves, so this arc is really contained in a single wall of some component of \mathcal{G}' . If $\gamma_0 = \gamma$, then by the correspondence in Lemma 2.5 we have successfully compressed γ into a leaf and we're done. Otherwise, using the correspondence in Lemma 2.5, we can compress γ_0 into \mathcal{G} which reduces N_{tan} by 1. If γ_0 shares an endpoint with γ , then $N_{\text{ess}} + N_{\text{iness}}$ remains constant. See Figure 4, top right. Otherwise, $N_{\text{ess}} + N_{\text{iness}}$ decreases by one (although each could individually increase). Therefore, the complexity function decreases. See Figure 4, bottom left, for an example in which N_{tan} decreases from 1 to 0 and $N_{\text{ess}} + N_{\text{iness}}$ decreases from 2 to 1.
- γ'_0 approaches \mathcal{G}' during tightening. Then there is a subarc of γ'_0 which wraps around $\partial\mathcal{G}'$, making contact along a subarc of Euclidean angle $> \pi$. Walking around the singularity, one meets a prong every π radians. Thus, the arc will cross at least two prongs of the singularity. This arc must have a tangency to Λ' . Push π of the corresponding subarc of γ_0 into \mathcal{G} so as to decrease N_{tan} by exactly one. See Figure 4, bottom right. We claim that the new arc created is essential so that N_{ess} increases by one. This fact crucially uses the assumption that the saddle region is an ideal n -gon bundle for some $n \geq 4$. A pair of walls of such a saddle region which are separated by two prongs must have distinct lifts in the universal cover of the solid torus saddle region. Therefore, the endpoints of the newly created arc in \mathcal{G} lie in different walls in this \mathbb{Z} -cover and so the new arc is essential. In total, the complexity function has decreased.

It follows that the minimizer γ has $N_{\text{tan}} = N_{\text{iness}} = 0$, so every tangency between γ and \mathcal{F} occurs inside an essential arc in \mathcal{G} . Finally, we pull each component of $\gamma \cap \mathcal{G}$ tight so as to contain either 0 or 1 tangencies with \mathcal{F} . The arc γ is now efficient. \square

Remark 4.3 Although Construction 3.2 works more generally for manifolds with a pseudo-Anosov flow having orientable invariant foliations, Proposition 4.2 does not hold. For example, there are many manifolds supporting Anosov flows whose invariant foliations have branching [15]. The existence of a Solv metric on $M_\varphi \setminus \mathcal{G}$ and the presence of “negative curvature” in \mathcal{G} from Assumption 4.1 are crucial.

Proposition 4.4 *Assume Assumption 4.1. The stabilizers of B_i^+ and B_i^- in $\pi_1(M_\varphi(\mathbf{p}; \mathbf{q}))$ are both equal to the infinite cyclic group $\text{Stab } B_i^\pm$ generated by a conjugate to the Dehn filling core $K_i(\mathbf{p}; \mathbf{q})$. $\text{Stab } B_i^\pm$ preserves a circular order on the points of B_i^+ .*

Proof Let R be the relevant saddle region. Observe that $\text{Stab } R = \text{Stab } B_i^+ = \text{Stab } B_i^-$. Every leaf in R is homeomorphic to a plane, and so has trivial stabilizer. Therefore, there is at most one element of $\text{Stab } R$ mapping between lifts $\lambda_1, \lambda_2 \in \tilde{\mathcal{F}}$ of a leaf $\lambda \in R$. There is always a power of $K_i(\mathbf{p}; \mathbf{q})$ which accomplishes this transformation, so $\text{Stab } R$ consists only of such elements. The circular order on walls is the order in which they appear as sides of a fiber of the ideal polygon bundle structure on R . \square

Let us now recall a standard fact about periodic orbits of φ .

Lemma 4.5 *The nonplanar leaves of Λ are all cylindrical. Moreover, they are in correspondence with primitive periodic orbits of φ . This correspondence is k -to-one for the leaves intersecting $\partial\mathcal{G}$, where k is the number of walls of the incident component of \mathcal{G} . It is one-to-one for all other leaves.*

Proof In what follows, we regard the various prongs of a singular leaf of \mathcal{F}^s or $\mathcal{F}^s \cap \Sigma$ as separate leaves. Since φ preserves no closed curve, the intersection of any leaf in \mathcal{F}^s with Σ is either a prong (homeomorphic to $[0, \infty)$) or a copy of \mathbb{R} . Leaves of \mathcal{F}^s with nontrivial topology are suspensions of those leaves of $\mathcal{F}^s \cap \Sigma$ which are preserved by some power of φ , and hence are homeomorphic to cylinders. Since φ is pseudo-Anosov, its action contracts lengths in each leaf of $\mathcal{F}^s \cap \Sigma$. A contraction on $[0, \infty)$ or \mathbb{R} has a unique fixed point, so each cylindrical leaf of \mathcal{F}^s contains exactly one primitive orbit of the suspension flow of φ .

Conversely, the suspension of a primitive periodic orbit of φ is either contained in a unique leaf of \mathcal{F}^s or is a singular orbit contained in two or more prongs.

We obtain Λ from \mathcal{F}^s by splitting open the leaves of \mathcal{F}^s which contain orbits in $\{K_i\}$. This has the effect of doubling such leaves and then gluing any leaves of the form $[0, \infty) \times S^1$ in pairs; the conclusion of the lemma follows for the leaves intersecting $\partial\mathcal{G}$. All other leaves are preserved in the passage from \mathcal{F}^s to Λ . \square

Proposition 4.6 *Assume Assumption 4.1. No element of $\text{Stab } B_i^\pm$ stabilizes any point in the Hausdorff-ification of L aside from B_i^+ and B_i^- .*

Proof $\text{Stab } B_i^\pm$ is generated by a loop freely homotopic to the Dehn surgery core $K_i(\mathbf{p}; \mathbf{q})$. The stabilizer of a cylindrical leaf in Λ is represented by a loop freely homotopic to the suspension of a periodic point of φ . We must show that there are no nontrivial free homotopies among loops $K_i(\mathbf{p}; \mathbf{q})$ (or powers thereof) or suspensions periodic orbits of φ (or powers thereof). Call the set of these curves \mathcal{O} . For simplicity, we replace each $K_i(\mathbf{p}; \mathbf{q})$ with a power thereof that is homotopic to one living inside a wall of $\partial\mathcal{G}$.

First, we rule out a nontrivial homotopy between two curves in \mathcal{O} that stays inside $M_\varphi(\mathbf{p}; \mathbf{q}) \setminus \mathcal{G}$. Now $M_\varphi(\mathbf{p}; \mathbf{q}) \setminus \mathcal{G}$ embeds in the fibered manifold M_φ so that the elements of \mathcal{O} are suspensions of periodic orbits of φ . Since φ is pseudo-Anosov, these suspensions are never homotopic to one another.

Now we must rule out a nontrivial homotopy between curves in \mathcal{O} that might pass through the guts \mathcal{G} . The idea is that the saddle regions contains lots of negative curvature, but an annulus giving rise to the purported homotopy has Euler characteristic zero and therefore can't cut across any saddle region. To be more precise, let γ and η be two elements of \mathcal{O} . Suppose γ is homotopic to η in $M_\varphi(\mathbf{p}; \mathbf{q})$. Replacing γ and η by γ^k and η^k if necessary, we can assume that γ and η both lie in leaves of \mathcal{F} . Let $A = [0, 1] \times S^1$ and choose an immersion $\iota: A \rightarrow M_\varphi(\mathbf{p}; \mathbf{q})$ realizing a homotopy from γ to η . By Roussarie [25] and Thurston [26], we can homotope the annulus so that the induced foliation has no critical points. We can also make A transverse to $\partial\mathcal{G}$ and the corners of $\partial\mathcal{G}$. Let $C = \iota^*(\partial\mathcal{G})$; by our transversality conditions, C is a collection of disjoint closed polygonal curves in A . Choose ι to minimize number of components of C .

Case 1 C contains an innermost loop α that is inessential in A . Let D be the disk in A bounded by α . The boundary of D is a $2n$ -gon with sides alternating between arcs transverse to the foliation and arcs tangent to the foliation. Here, $2n$ is the number of corners of $\partial\mathcal{G}$ that ∂D intersects. Moreover, the angles between adjacent arcs are all convex or all concave, depending on whether $\text{int } D$ lies inside or outside \mathcal{G} . See Figure 5, left, for a picture of the convex case with $n = 4$. The foliation on D is oriented and has no critical points. An application of the Poincaré–Hopf index theorem shows that $n = 2\chi(D) = 2$ in the convex case, or $n = -2\chi(D) = -2$ in the concave case, which is impossible. Since every saddle region is an ideal polygon bundle with fiber having at least four sides, a meridian in $\partial\mathcal{G}$ crosses at least eight corners. Therefore, α is not a meridian or a multiple of a meridian. Furthermore, α cannot be homotopic to a multiple of a core of \mathcal{G} since α is zero in $\pi_1(M_\varphi(\mathbf{p}; \mathbf{q}))$. Therefore, α must be inessential in $\partial\mathcal{G}$. We can then compress D into $\partial\mathcal{G}$ (while maintaining Roussarie–Thurston general position) and eliminate the curve α from C , contradicting our minimality assumption.

Case 2 C contains only loops that are essential in A . Let α be any element of C . As in the first case, $\alpha \cap A$ alternates between arcs between transverse and tangent to \mathcal{F} , and the angles between adjacent arcs




Figure 5: The annulus A is shown along with the induced foliation in a neighborhood of α in two disallowed configurations. If α crosses the sutures in $\partial\mathcal{G}$ as shown, there is no way to extend the foliation to the rest of A without critical points, contradicting Roussarie–Thurston.

all either convex or concave depending on which side of α is on the interior of \mathcal{G} . Let n be the number of alternations between transverse and tangent arcs in α . See Figure 5, right, for the case $n = 4$. Let R be the annulus cobounded by γ and α . Since $\mathcal{F} \cap R$ has no singularities, we have $n = 2\chi(R) = 0$ and we must be in the convex case. Thus, α actually intersects no corners of $\partial\mathcal{G}$ and is either entirely transverse to \mathcal{F} (ie is contained in a suture of $\partial\mathcal{G}$) or is contained in a wall of \mathcal{G} . Moreover, α is not contractible in ∂G since α is homotopic in $M_\varphi(\mathbf{p}; \mathbf{q})$ to γ which is even noncontractible in $M_\varphi(\mathbf{p}; \mathbf{q})$. Therefore, α is a closed curve parallel to the sutures in some component of $\partial\mathcal{G}$.

The arguments above hold for any $\alpha \in C$. The curves in C cut A into a collection of annuli which we label A_1, \dots, A_M . Each A_k is either an annulus living in $M_\varphi(\mathbf{p}; \mathbf{q}) \setminus \mathcal{G}$ or an annulus in \mathcal{G} . In the first alternative, A_k represents a free homotopy in $M_\varphi(\mathbf{p}; \mathbf{q}) \setminus \mathcal{G}$; as noted at the beginning of the proof, the two boundary components of such an A_k must represent the same element of \mathcal{O} . In the latter alternative, we can trivially make the same conclusion. Therefore, γ and η represent the same element of \mathcal{O} , as desired. \square

5 Gluing branches

In this section, we give some motivation for the constructions in Section 6. We wish to define on L an equivalence relation, denoted by \sim , such that $(L/\sim) \cong \mathbb{R}$.

5.1 Hausdorff and non-Hausdorff quotients

We first describe this gluing process in a simpler, abstract case where L is replaced by an infinite binary tree T . Consider an infinite rooted binary tree T , each of whose vertices has a distinguished left and right child. We think of T not combinatorially, but as a topological space. Let v_0 be the root vertex of T . Given a vertex $v \in T$, let vL represent its left child and vR represent its right child. We use the notation $\{L, R\}^k$ to denote the set of words of length k in the two letter alphabet $\{L, R\}$. We use fractional powers of R to denote points that are on edges of the tree; for example, $vR^{\frac{1}{2}}$ is the midpoint of the path $[v, vR]$. Then any point on the tree is of the form $v_0\{L, R\}^k L^\alpha$ or $v_0\{L, R\}^k R^\alpha$ for some $k \geq 0$ and $0 \leq \alpha < 1$.




Figure 6: The red paths are of the form $[v, vLLL\dots)$ and the blue paths are of the form $[v, vRRR\dots)$. Points which are equivalent with respect to the gluing are marked with the same number.

In order to illustrate the kinds of difficulties we will encounter in the next section, we will show two superficially similar equivalence relations on T such that the quotient is $[0, \infty)$ in the first but is non-Hausdorff in the latter.

Example 1 Let us define the first equivalence relation. At each vertex v , we glue the infinite paths $[v, vLLL\dots)$ to $[v, vRRR\dots)$ together by the obvious map $[0, \infty) \rightarrow [0, \infty)$ which preserves depths of vertices. See Figure 6. This has the effect of collapsing the entire tree down to a copy of $[0, \infty)$.

Example 2 Let us construct a different equivalence relation such that the quotient is not Hausdorff. At each vertex $v \in T$ we glue the infinite descending paths $[v, vLRRR\dots)$ and $[v, vRLRRR\dots)$ by a homeomorphism which sends $[v, vL]$ to $[v, vRL]$ by a dilation by a factor of two, and sends $[vL, vLRRR\dots)$ to $[vRL, vRLRRR\dots)$ by an isometry. See Figure 7. The points v_0R^nL are identified for all $n \geq 0$. Therefore, $[v_0, v_0RRR\dots)$ is not properly embedded in the quotient space. It turns out that the quotient space is not $[0, \infty)$ but a tree with infinite valence at each vertex. For example, the points v_0LL and v_0RLL are not comparable in the partial order induced on the quotient. For points $a, b \in T$, let $[a]$ and $[b]$ denote their images in the quotient. Write $a \geq b$ if a is an ancestor of b in T , and $[a] \geq [b]$ if $[a]$ is an ancestor of $[b]$ in the quotient. The interested reader is invited to show the following, which give a precise picture of the quotient of T with respect to our second gluing.

- (i) Each equivalence class under the gluing has a canonical representative of the form $v_0\{L, R\}^k LLR^s$, v_0LR^s , or v_0R^s for some $s \geq 0$.
- (ii) Suppose a and b are canonical representatives of their equivalence classes. Write $a = rWR^s$ for some word W in L and R , and $s \in \mathbb{R}$ maximal. Then $[a] \geq [b]$ if and only if one of the following holds:
 - (a) $a \geq b$.
 - (b) W is the empty word and $v_0WR^nL \geq b$ for some integer $n \leq s$.
 - (c) $W = W'L$ and $v_0W'R^kLL \geq b$ for some integer $k \geq 0$.

5.2 Gluing branches in L

We wish to perform a similar gluing on the leaf space L . The leaf space should be thought of as a kind of tree, but possibly with a dense set of branching points. There are two new features in this case. First, there may be branching in both the upward and downward directions. Second, we must do this gluing in a π_1 -equivariant way so that the action of π_1 will descend to the quotient L/\sim . The gluings we performed in the case of the binary tree are equivariant with respect to the semigroup of isometric self-maps of the binary tree which preserve the left and right children at each node. More care will be required in L since the stabilizer of a branch point may be nontrivial.

Let b_i be a generator for $\text{Stab } B_i^\pm$. Let a_i be the smallest positive integer such that $b_i^{a_i}$ stabilizes B_i pointwise. In other words, a_i is the integer such that $b_i^{a_i}$ is freely homotopic to the essential loop in a wall of \mathcal{G} . We fix the orientation of b_i by asking that the pushoff of $b_i^{a_i}$ to the outside of \mathcal{G} has negative intersection number with Σ . With this choice, the holonomy of \mathcal{F} on the outside of \mathcal{G} along a curve freely homotopic to $b_i^{a_i}$ is always attracting. This implies that the action of $b_i^{a_i}$ on the leaf space is locally expanding.

For each $x \in B_i^+$ (ie a lift of a wall of a saddle region), let $(x, \infty) \subset L$ be the lift to L of the transverse loop obtained by pushing the essential loop in x slightly outside the saddle region corresponding to B_i^+ . Another way to think of (x, ∞) is as the path in L corresponding with the set of leaves of $\tilde{\mathcal{F}}$ intersecting a prong of the unstable invariant foliation \mathcal{F}^u . Let $[x, \infty) = \{x\} \cup (x, \infty)$. The paths $\{[x, \infty) \mid x \in B_i^+\}$ will play the role of the paths $[v, vRRR \dots)$ or $[v, vLLL \dots)$ in the case of the binary tree.

Lemma 5.1 *The set $[x, \infty)$ is a $b_i^{a_i}$ invariant path from x to ∞ in L . Furthermore, b^{a_i} acts on $[x, \infty)$ by an expanding dilation fixing x .*

Proof Since (x, ∞) is a lift of a loop homotopic to a_i times the core of the corresponding saddle region, $[x, \infty)$ is $b_i^{a_i}$ invariant. Suppose $[x, \infty)$ had a greatest lower bound y . Then either y and $b^{a_i} y$ are distinct nonseparated leaves or $y = b^{a_i} y$. By Proposition 4.6, this never happens. Therefore, $[x, \infty)$ is properly embedded in L . Since the holonomy around the essential loop in x is repelling on the outside of \mathcal{G} and $b_i^{a_i}$ has no fixed points in $[x, \infty)$ besides x , it must be that $b_i^{a_i}$ acts by a homeomorphism conjugate to a dilation with stretch factor > 1 . □

Call $[x, \infty)$ the *invariant path* at x . We are now ready to give a first attempt at defining the equivalence relation \sim on L .

Construction 5.2 Recall that B_i^+ is an arbitrarily chosen representative among its translates by deck transformations. For each point x in the branch locus B_i^+ (resp. B_i^-), construct the upward (resp. downward) oriented path $[x, \infty) \subset L$ from x to $+\infty$ (resp. $-\infty$), which is invariant under $b_i^{a_i}$ as in Lemma 5.1. Now b_i acts on $\bigcup_{x \in B_i^+} [x, \infty)$. We shall now glue together the various paths in the set $\{[x, \infty) \mid x \in B_i^+\}$. Up




Figure 8: On the left, the action of b_i on $\bigcup_{x \in O} [x, \infty)$ is shown for the case $a_i = 3$. Here, $O = \{x_1, x_2, x_3\} \subset B_i^-$. The action of b_i descends to the quotient by \sim . This action on the quotient is shown on the right.

to reparametrizing $[x, \infty)$, we can assume that $b_i^{a_i}$ acts on each $[x, \infty)$ by dilation by a factor 2^{a_i} . For each $x \in B_i^+$, each $y \in [x, \infty)$ and each $m \in \mathbb{Z}$, declare $y \sim (1/2^m)b_i^m y$. This has the effect of gluing together the paths $\{[x, \infty) \mid x \in O\}$ for $O \subset B_i^+$ an orbit of the $\text{Stab } B_i^+$ action on B_i^+ . See Figure 8. The action of b_i descends to the quotient L/\sim . In the quotient, b_i acts on $[x, \infty)/\sim$ as a dilation by a factor of 2.

We have now glued together $[x_i, \infty)$ and $[x_j, \infty)$ whenever x_i, x_j are points of B_i^+ lying in the same orbit of the $\text{Stab } B_i^+$ action. However, the $\text{Stab } B_i^+$ action on B_i^+ may not be transitive. For example, consider the case when K_i is an untwisted orbit with $2k$ prongs. In the nonsingular case $k = 1$, we have $|B_i^+| = |B_i^-| = 1$. In the singular case $k > 1$, the action of b_i on B_i^+ is not transitive because it does not permute the k points in B_i^+ . We also wish to additionally glue those paths when x_i and x_j are in different orbits. Moreover, we wish to do this in some b_i equivariant way. Observe that the induced action of b_i on the quotient is the same on $[x_i, \infty)/\sim$ and $[x_j, \infty)/\sim$; both induced actions are expanding dilations. Thus, when x_i and x_j lie in different orbits, we simply choose some gluing map from $[x_i, \infty)$ to $[x_j, \infty)$ that intertwines their respective $\text{Stab } B_i^+$ actions.

In total, we have glued together $|B_i^+|$ semiinfinite paths from B_i^+ to ∞ together in a $\text{Stab } B_i^\pm$ -equivariant way. Perform this gluing procedure for B_i^- as well. Finally, we perform a gluing at each of the translates of B_i^\pm in the way dictated by π_1 -equivariance. This concludes the construction.

Note that if $x \sim y$, then $gx \sim gy$ for any $g \in \pi_1(M_\varphi(\mathbf{p}; \mathbf{q}))$. Therefore, Construction 5.2 gives a π_1 -equivariant equivalence relation \sim on leaves. Unfortunately, L/\sim is not homeomorphic to \mathbb{R} due to a phenomenon like that of Figure 7. In Section 6, we present a coarsening of this equivalence relation which finishes the job.

6 The flat connection $\hat{\mathcal{J}}$

We prove the main theorems in this section. Recall that $M_\varphi^\circ = M_\varphi \setminus \{K_i\}$. Our first goal is to construct a trivial S^1 -bundle (with structure group $\text{Homeo}^+(S^1)$) on M_φ° with a flat connection \mathcal{J} whose monodromy around the filling curves on the toroidal ends is trivial. This S^1 -bundle will lift to an \mathbb{R} -bundle with a flat connection $\hat{\mathcal{J}}$ having the same property.

For the purposes of defining \mathcal{J} and proving Theorem 1.1, we will need only the singular foliations \mathcal{F}^u and \mathcal{F}^s . The foliation \mathcal{F} will figure in the proof of Theorem 1.2 where we make use of the fact that \mathcal{F} blows down to \mathcal{F}^s .

6.1 Preliminaries on connections and partial connections

We will use the following convention throughout the rest of the paper.

Convention 6.1 (concatenation of paths) Let X be a topological space. Given paths γ_1 and γ_2 in X , we define their concatenation $\gamma_1 * \gamma_2$ to be the path which first follows γ_1 , then follows γ_2 . We use $*$ as multiplication in $\pi_1(X)$. With this convention, the action of $\pi_1(X)$ on \tilde{X} by deck transformations is determined as follows. Choose a basepoint $x_0 \in \tilde{X}$. For any $g \in \pi_1(M)$ and $q \in \tilde{X}$, let γ_1 be a path representing g and let γ_2 be a path in X that lifts to a path from x_0 to q . Then $\gamma_1 * \gamma_2$ lifts to a path from x_0 to gq .

Let B be a topological space. Given a topological bundle $F \rightarrow E \xrightarrow{\pi} B$, a *connection* \mathcal{H} is a choice of a homeomorphism $\mathcal{H}_\gamma: \pi^{-1}(\gamma(0)) \rightarrow \pi^{-1}(\gamma(1))$ for each continuous path $\gamma: [0, 1] \rightarrow B$. These homeomorphisms are required to be independent of the parametrization of γ and to satisfy functoriality conditions with respect to concatenation of paths. We refer to \mathcal{H}_γ as *parallel transport* along γ with respect to \mathcal{H} . The functoriality conditions are:

- (i) $\mathcal{H}_{\gamma * \lambda} = \mathcal{H}_\lambda \circ \mathcal{H}_\gamma$, where λ and γ are concatenable paths in B , and $\gamma * \lambda$ is their concatenation.
- (ii) $\mathcal{H}_{\gamma * \gamma^{-1}} = \text{id}$, where γ^{-1} is γ traversed backwards.
- (iii) Parallel transport along a trivial path is the identity map.

Said in a different way, let \mathcal{C} be the category whose objects are topological spaces homeomorphic to F and whose morphisms are homeomorphisms. An F -bundle over B assigns to each point $x \in B$ the fiber over x which is an object in \mathcal{C} . A connection is an extension of this assignment to a functor from the groupoid of paths in B to \mathcal{C} . This perspective will be useful because we will sometimes specify a connection by defining it on a set of generators for the groupoid of paths in B . We might also replace the groupoid of paths in B with a slightly larger but equivalent groupoid.

Remark 6.2 For the purposes of this paper, we will not need to ask that \mathcal{H}_γ is close to the identity when γ is a short path. Thus, our definition of a connection makes sense even in the absence of a local product structure on the bundle. This is useful because we will actually use a connection to define the local product structure.

A connection is called *flat* if $\mathcal{H}_\gamma = \text{id}$ whenever γ is contractible. A flat connection gives rise to a homomorphism of $\pi_1(B)$ into $\text{Homeo}(F)$ defined by $[\gamma] \mapsto \mathcal{H}_\gamma^{-1}$ for $[\gamma] \in \pi_1(B)$.

A section of π is called *flat* with respect to \mathcal{H} or \mathcal{H} -flat if $\mathcal{H}_\gamma(s(\gamma(0))) = s(\gamma(1))$ for all paths $\gamma: [0, 1] \rightarrow B$. When \mathcal{H} is understood, we will suppress mentioning it and simply say that a section is flat.

We may also refer to sections of π over a curve $\gamma: [0, 1] \rightarrow B$; this simply means a section of the pullback bundle over $[0, 1]$. We will usually express such a section as a map $s: [0, 1] \rightarrow E$. A flat section over a path is also called a *parallel section*. For any path $\gamma: [0, 1] \rightarrow B$ and $x \in \pi^{-1}(\gamma(0))$, there is a unique parallel section s_x over γ satisfying $s_x(0) = x$. It is traced out by parallel transport of x along γ . In equations, this means $s_x(t) = \mathcal{H}_{\gamma|_{[0,t]}}(x)$, where $\gamma|_{[0,t]}$ is the restriction of γ to $[0, t]$.

A *partial connection* is similar to a connection except that the homeomorphisms \mathcal{H}_γ may not be defined on all of $\pi^{-1}(\gamma(0))$. A partial connection \mathcal{H} is a choice of a homeomorphism \mathcal{H}_γ for each path $\gamma: [0, 1] \rightarrow B$ from some (topological) subspace of $\pi^{-1}(\gamma(0))$ to some subspace of $\pi^{-1}(\gamma(1))$. The subspaces may depend on the path γ . The homeomorphisms are again required to be independent of the parametrization of γ and functorial with respect to composition of paths:

- (i) $\mathcal{H}_{\gamma*\lambda}(x) = \mathcal{H}_\lambda(\mathcal{H}_\gamma(x))$ for all $x \in \pi^{-1}(\gamma(0))$ at which the right side is defined. In particular, $\mathcal{H}_{\lambda*\gamma}(x)$ is defined whenever the right side is defined.
- (ii) $\mathcal{H}_{\gamma*\gamma^{-1}}(x) = x$ for all $x \in \pi^{-1}(\gamma(0))$ at which \mathcal{H}_γ is defined. In particular, $\mathcal{H}_{\gamma*\gamma^{-1}}(x)$ is defined whenever $\mathcal{H}_\gamma(x)$ is defined.
- (iii) Parallel transport along a trivial path is defined on the entire fiber $\pi^{-1}(\gamma(0))$ and is equal to the identity map.

A partial connection is called *flat* if for each point $x \in B$ and each compact set $W \subset \pi^{-1}(x)$, there exists a neighborhood U of x such that for every contractible path $\gamma: [0, 1] \rightarrow U$, \mathcal{H}_γ is defined on all of W and agrees with the identity map.

Remark 6.3 In contrast with the case of flat connections, the monodromy of a flat partial connection around a long, contractible loop need not agree with the identity map on its domain of definition.

A section is called *flat* with respect to a partial connection \mathcal{H} if $\mathcal{H}_\gamma(s(\gamma(0))) = s(\gamma(1))$ for all paths $\gamma: [0, 1] \rightarrow B$ along which \mathcal{H}_γ is defined on $s(\gamma(0))$. Define *flat* or *parallel* sections over paths analogously to the case of connections. Given a path $\gamma: [0, 1] \rightarrow B$ and a point $x \in \pi^{-1}(\gamma(0))$, one may attempt to define a parallel section s_x over γ by $s_x(t) = \mathcal{H}_{\gamma|_{[0,t]}}(x)$. However, the right side may fail to be




Figure 9: The total space for the bundle is shown with the fiber direction vertical. Parallel sections for the partial connection \mathcal{H} are shown in red. Each parallel section is defined only for a limited time.

defined for large t . In this case, let $t_{\max}(x) = \sup\{t \mid \mathcal{H}_\gamma|_{[0,t]}(x) \text{ is defined}\}$. Then we say that the parallel section s_x *blows up* at time $t_{\max}(x)$.

A flat partial connection is roughly the same thing as a foliated bundle, though we will not use that language since the total space of our bundle may not a priori have a topology making it a manifold.

Example 6.4 A partial connection \mathcal{H} on the bundle $(0, 1) \rightarrow (0, 1) \times \mathbb{R} \rightarrow \mathbb{R}$ may be defined as follows. For any γ parametrizing a curve in \mathbb{R} that is monotonically increasing or decreasing, \mathcal{H}_γ is a homeomorphism between an open subinterval of the fiber over $\gamma(0)$ and an open subinterval of the fiber over $\gamma(1)$. We define $\mathcal{H}_\gamma(x) = x + \gamma(0) - \gamma(1)$ for any x satisfying $x \in (0, 1)$ and $x + \gamma(0) - \gamma(1) \in (0, 1)$. For γ which is not monotonically increasing or decreasing, \mathcal{H}_γ is defined by composition of monotonic paths. In this case, the range and domain of \mathcal{H}_γ will be smaller than that specified in the formula. See Figure 9.

6.2 The partial connection $\mathcal{J}^{\text{partial}}$

Let Σ° be a typical fiber of M_φ° , ie Σ punctured at its intersections with the K_i . M_φ° comes equipped with transverse stable and unstable codimension 1 foliations \mathcal{F}^s and \mathcal{F}^u . It also has a natural incomplete Solv metric g locally expressible as $dt^2 + \lambda^t dx^2 + \lambda^{-t} dy^2$, where x and y are local coordinates on Σ° , $\lambda > 1$ is the stretch factor of φ , and t is the coordinate transverse to Σ° . By the orientability constraints, we may consistently establish cardinal directions north, south, east and west. We choose our orientations on \mathcal{F}^s and \mathcal{F}^u so that $\mathcal{F}^s \cap \Sigma^\circ$ is oriented eastward and $\mathcal{F}^u \cap \Sigma^\circ$ is oriented northward. For us, a *rectangle* will always mean a rectangle in a fiber surface free of singularities on its interior, with top and bottom parallel to the east–west direction, and left and right sides parallel to the north–south direction. Recall that we constructed \mathcal{F} by splitting open \mathcal{F}^s . Let Σ_p° be the fiber surface containing a point p .

A *generalized leaf* of $\mathcal{F}^u \cap \Sigma^\circ$ is one of

- (i) a leaf of $\mathcal{F}^u \cap \Sigma^\circ$ which is not a prong, or
- (ii) the concatenation of two prongs of $\mathcal{F}^u \cap \Sigma^\circ$ which meet at the same singularity and make an angle of π , ie a limit of leaves of $\mathcal{F}^u \cap \Sigma^\circ$ approaching the singularity.




Figure 10: The unstable foliation near a singularity with $k_i = 4$. The fibers over p_{west} and p_{east} are each unions of two prongs.

We may similarly define a *generalized leaf* of \mathcal{F}^u . The concatenation of two prongs P_1 and P_2 of $\mathcal{F}^u \cap \Sigma^\circ$ incident with the same singularity q formally requires an extra dummy point since q itself is not a point in M_φ° . We usually call the dummy point q^* so that the generalized leaf is $P_1 \cup \{q^*\} \cup P_2$. If there is more than one generalized leaf in play, we will use q^{**}, q^{***} , etc to denote their dummy points.

A point on a prong of \mathcal{F}^u is contained in exactly two generalized leaves, but it would be better if each point were contained in exactly one generalized leaf. Thus, we will formally double each point p on a prong P of \mathcal{F}^u into two points p_{west} and p_{east} . The points p_{west} and p_{east} should be regarded as points in M_φ° infinitesimally perturbed to the west and east of P respectively. Let Z be the space obtained from M_φ° by doubling each point on a prong of \mathcal{F}^u . The open sets of the topology on Z are the pullbacks of open sets in M_φ° under the obvious projection. Note that Z is not Hausdorff and not even T_0 because for any point p on a prong, p_{west} and p_{east} are not topologically distinguishable. In Z , each point is contained in exactly one generalized leaf of \mathcal{F}^u . One can construct a path in Z from p_{west} to p_{east} whose image is the two point set $\{p_{\text{west}}, p_{\text{east}}\}$; such a path should be interpreted as an infinitesimal path in M_φ° crossing P from p_{west} to p_{east} . As a consequence, Z is path-connected.

For any $p \in Z$, let Σ_p° be the fiber surface containing p . (Here we are abusing notation slightly by using Σ_p° to refer to its lift to Z . Let

$$E_\pi = \{(p, y) \mid p, y \in Z \text{ and } y \text{ is on the generalized leaf of } \mathcal{F}^u \cap \Sigma_p^\circ \text{ which contains } p\}.$$

Let $\pi : E_\pi \rightarrow Z$ be the projection map $\pi : (p, y) \mapsto p$. Then π defines an \mathbb{R} bundle over Z . Intuitively, one should think of π as the \mathbb{R} -bundle over M_φ° whose fiber over a point $p \in M_\varphi^\circ$ is the generalized leaf of $\mathcal{F}^u \cap \Sigma_p^\circ$ containing p .

Define the auxiliary map $\theta : E_\pi \rightarrow Z$ by $(p, y) \mapsto y$. In practice, we visualize points in E_π via their images under θ .




Figure 11: This is a picture in M_φ° . The leaves of \mathcal{F}^s are shown in red. The point $h_\gamma(y_0)$ lies on the same stable leaf as y_0 . Moreover, there exists a path from y_0 to $h_\gamma(y_0)$ in a leaf of \mathcal{F}^s , and this path is homotopic to γ through paths connecting $\theta \circ \pi^{-1}(\gamma(0))$ to $\theta \circ \pi^{-1}(\gamma(1))$. These properties uniquely specify $h_\gamma(y_0)$. On the other hand, h_γ is not defined at y'_0 .

Note that there is no obvious local product topology on E_π , since the leaves of \mathcal{F}^u diverge around every singular point. We will eventually define a topology by using a flat connection in Section 6.4.

In the rest of this section, when there is no chance of confusion, we will conflate Z and M_φ° . This means that we will think of p_{west} and p_{east} as points in M_φ° , think of π as an \mathbb{R} -bundle over M_φ° , and think of θ as a map from E_π to M_φ° . We will also think of the path from p_{west} to p_{east} in Z as an infinitesimal path in M_φ° with distinct endpoints. Since M_φ° embeds in Z (for example by always choosing $p \mapsto p_{\text{west}}$ when p lies on a prong), a connection on Z restricts to a connection on M_φ° . Thus, we only gain generality by working in Z .

For a curve $\gamma: [0, 1] \rightarrow M_\varphi^\circ$, there is a partially defined map

$$h_\gamma: \theta \circ \pi^{-1}(\gamma(0)) \rightarrow \theta \circ \pi^{-1}(\gamma(1)),$$

which we refer to as the holonomy of \mathcal{F}^s along γ . To define h_γ , choose a lift of γ to the leaf space of $\tilde{\mathcal{F}}^s$. (Here, $\tilde{\mathcal{F}}^s$ is a foliation on \tilde{M}_φ° , obtained by lifting the restriction of \mathcal{F}^s to M_φ° .) Choose a compatible lift of $\theta \circ \pi^{-1}(\gamma(0))$ to a line embedded in the leaf space of $\tilde{\mathcal{F}}^s$. Here, compatible means that the lift of $\theta \circ \pi^{-1}(\gamma(0))$ agrees with the lift of γ at $\gamma(0)$. Similarly, choose a lift of $\theta \circ \pi^{-1}(\gamma(1))$ that agrees with the lift of γ at $\gamma(1)$. Then for $y_0 \in \theta \circ \pi^{-1}(\gamma(0))$, we define $h_\gamma(y_0)$ to be the point in $\theta \circ \pi^{-1}(\gamma(1))$ whose lift to the leaf space of $\tilde{\mathcal{F}}^s$ coincides with the lift of y_0 , if such a point exists. Note that h_γ is only defined on the subinterval of $\theta \circ \pi^{-1}(\gamma(0))$ whose lift intersects the lift of $\theta \circ \pi^{-1}(\gamma(1))$. See Figure 11 for a picture in M_φ° .

Now we shall define a partial connection $\mathcal{J}^{\text{partial}}$ on the bundle π . We initially define $\mathcal{J}^{\text{partial}}$ on three types of short paths which generate a dense set of paths in M_φ° . Whenever h_γ exists, we define $\mathcal{J}_\gamma^{\text{partial}}$ so that

$$(1) \quad \theta \circ \mathcal{J}_\gamma^{\text{partial}} = h_\gamma \circ \theta.$$

In other words, the θ -image of a parallel section tries to stay on the same leaf of \mathcal{F}^s . This is what happens in type 1 and type 2 curves below. However, $\mathcal{J}^{\text{partial}}$ will also be defined on some curves (type 3 curves below) along which holonomy does not exist.

Type 1 Suppose γ is a path in a generalized leaf μ of \mathcal{F}^u . In this case, we simply define $\mathcal{J}^{\text{partial}}$ using the holonomy of \mathcal{F}^s . In $\tilde{\mu} \cong \mathbb{R}^2$, the leaves of $\tilde{\mathcal{F}}^s \cap \tilde{\mu}$ (ie the lifts of flow lines of the suspension flow of φ) and the lifts to $\tilde{\mu}$ of the θ -images of fibers of π form two transverse codimension-1 foliations. This pair of foliations is topologically conjugate to the horizontal and vertical foliations of \mathbb{R}^2 . It follows that h_γ is fully defined on $\theta \circ \pi^{-1}(\gamma(0))$. We may thus define $\mathcal{J}_\gamma^{\text{partial}}$ as in equation (1). In other words, we are pushing θ -images of fibers of π along the suspension flow of φ .

Type 2 Suppose R is a rectangle in Σ_p° for some p . Let γ be a path in R from the east to the west side of R . In this case, h_γ identifies the east and west sides of the rectangle by the obvious isometry, so we define $\mathcal{J}^{\text{partial}}$ as in equation (1). (If an endpoint of γ lies on a prong, then we make the same definition of $\mathcal{J}_\gamma^{\text{partial}}$ regardless of whether the endpoint is infinitesimally perturbed to the east or the west. We also allow the north or south side of the rectangle to contain a singularity.)

Type 3 Let p be a point on a prong P of $\mathcal{F}^u \cap \Sigma_p^\circ$, and let γ be the infinitesimal path from p_{west} to p_{east} . Let q be the singularity terminating P . Let P_{east} and P_{west} be the prongs at q adjacent to P so that

$$\theta \circ \pi^{-1}(p_{\text{west}}) = P \cup \{q^*\} \cup P_{\text{west}} \quad \text{and} \quad \theta \circ \pi^{-1}(p_{\text{east}}) = P \cup \{q^{**}\} \cup P_{\text{east}}.$$

Briefly conflating $\theta \circ \pi^{-1}(p_{\text{west}})$ with $\pi^{-1}(p_{\text{west}})$, let us define

$$\mathcal{J}_\gamma^{\text{partial}}: P \cup \{q^*\} \cup P_{\text{west}} \rightarrow P \cup \{q^{**}\} \cup P_{\text{east}}$$

separately in two subcases:

- **Type 3a** If p lies to the south of q , then declare that $\mathcal{J}_\gamma^{\text{partial}}$ sends q^* to q^{**} , acts as the identity on P , and stretches P_{west} by a factor of $\lambda^{m_i q_i / p_i}$, where $(p_i; q_i)$ is the surgery slope at the relevant singularity and m_i is the φ -period of the singularity. Here, λ is the stretch factor of φ and the relevant metric is the singular Euclidean metric induced on $\tilde{\Sigma}_p^\circ$ by the incomplete Solv metric g , as introduced at the beginning of Section 6.2. It is instructive to look at Figure 13, which illustrates how parallel sections over type 3a curves behave.
- **Type 3b** If p lies to the north of q , then make the same definition except that we use the inverse stretch factor $\lambda^{-m_i q_i / p_i}$ instead.

Note that h_γ is defined on P , but not on P_{west} . Thus, we had some freedom to choose the dilation factors in types 3a and 3b. Our choice is designed to make Proposition 6.8 work.

Now we will extend the definition of $\mathcal{J}^{\text{partial}}$ to arbitrary curves in M_φ° . By composition of parallel transport maps, we may define $\mathcal{J}^{\text{partial}}$ for concatenations of type 1–3 curves. Given an arbitrary path γ in M_φ° , we may C^0 approximate γ by a concatenation of type 1–3 curves. There may be many different

ways to do this approximation, resulting in different parallel transport maps with slightly different domains of definition; for example, a type 2 curve may be written as a concatenation of a type 2, type 3 and a type 2 curve. We define $\mathcal{J}_\gamma^{\text{partial}}$ by patching together all of these parallel transport maps. More precisely, for $x \in \pi^{-1}(\gamma(0))$, we say that $\mathcal{J}_\gamma^{\text{partial}}(x) = y$ if there is a sequence of paths γ_i from $\gamma(0)$ to $\gamma(1)$ such that

- γ_i is a concatenation of type 1–3 curves,
- the sequence of paths γ_i converges to γ in the C^0 topology as $i \rightarrow \infty$, and
- $\mathcal{J}_{\gamma_i}^{\text{partial}}(x) = y$.

Proposition 6.7 below guarantees that the various different C^0 approximations of γ yield parallel transport maps that agree on the intersection of their domains of definition. So $\mathcal{J}_\gamma^{\text{partial}}$ is well defined.

Remark 6.5 The definition of $\mathcal{J}^{\text{partial}}$ along Type 3 curves is motivated by Construction 5.2. The images in L of P_{west} and P_{east} are paths the form $[x, \infty)$ constructed in Section 5. The failure of $\mathcal{J}^{\text{partial}}$ to be a connection is analogous to the failure of L/\sim to be a line in Construction 5.2.

Definition 6.6 A commutator of type i and type j curves is a contractible loop formed by concatenating four arcs which alternate between type i and type j curves.

Proposition 6.7 Given any point $p \in M_\phi^\circ$ and a compact subset W of $\pi^{-1}(p)$, the monodromy of $\mathcal{J}^{\text{partial}}$ around sufficiently small contractible loops based at p formed by concatenating type 1, 2 and 3 curves is defined on W and equal to the identity. It follows that $\mathcal{J}^{\text{partial}}$ is flat.

Proof Without loss of generality, assume $p \in \text{int } \theta(W)$. If p does not lie on a prong, then one can find a tall, skinny rectangle R in Σ_p° containing a neighborhood of $\theta(W)$. Thicken this rectangle by ε in the direction transverse to Σ_p° . Now holonomy of transversals to \mathcal{F}^s in $R \times (-\varepsilon, \varepsilon)$ exists along all curves in a neighborhood p , so the required monodromies are trivial in this case.

Suppose instead that $p = p_{\text{west}}$ lies on a prong $P \subset \Sigma_p^\circ$. This time, choose a tall, skinny rectangle R_{west} such that its east side contains $\theta(W)$. Choose another tall, skinny rectangle R_{east} such that its west side contains $\theta(\mathcal{J}_\gamma^{\text{partial}} W)$, where γ is the type 3 curve from p_{west} to p_{east} . Let U be an ε -thickening of $R_{\text{west}} \cup R_{\text{east}}$ in the direction transverse to Σ_p° . By construction, p lies on the interior of U .

We need to show that monodromy around a contractible loop in U based at p is trivial on W . The null-homotopy of such a loop may be decomposed into disks whose boundaries are commutators of type i and type j arcs. Moreover, we can arrange that the only commutators crossing $P \times (-\varepsilon, \varepsilon)$ are commutators of type 3 and type 1 arcs. This can be achieved by subdividing every type 2 arc crossing $P \times (-\varepsilon, \varepsilon)$ into a concatenation of a type 2 arc in $R_{\text{west}} \times (-\varepsilon, \varepsilon)$, a type 3 arc crossing $P \times (-\varepsilon, \varepsilon)$ and a type 2 arc in $R_{\text{east}} \times (-\varepsilon, \varepsilon)$. This replacement only increases the domain of definition of the $\mathcal{J}^{\text{partial}}$. Any commutator contained in $R_{\text{west}} \times (-\varepsilon, \varepsilon)$ or $R_{\text{east}} \times (-\varepsilon, \varepsilon)$ has trivial monodromy on W . In particular,

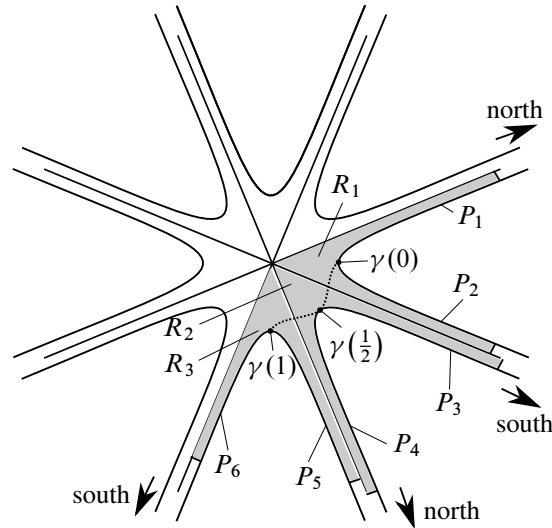


Figure 12

this includes any commutator involving a type 2 arc. So it only remains to check that the monodromy around a commutator of a type 1 arc and a type 3 arc vanishes. Parallel transport along a type 1 arc which moves a distance α in the direction transverse to Σ_p° acts as a dilation by factor λ^α (relative to the induced Euclidean metric on fibers). On the other hand, parallel transport along a type 3 arc acts as a piecewise dilation. These two maps commute, as desired. Finally, it's easy to check that parallel transport along any path $\gamma \subset U$ starting at p is defined on W . \square

Now we will prove one of the key properties that make our constructions work.

Proposition 6.8 *The monodromy of $\mathcal{J}^{\text{partial}}$ along a curve parallel to a filling slope on a toroidal end of M_ϕ° is trivial. To be precise, let p be a point on an unstable prong incident to an end of M_ϕ° . For any compact subspace $W \subset \pi^{-1}(p)$, the monodromy of $\mathcal{J}^{\text{partial}}$ around a small enough loop γ based at p and homotopic to the Dehn filling meridian is defined on W and equal to the identity.*

Proof The boundary curve in $N(K_i)$ can be written as the suspension of a small perturbation of a point in $\Sigma \cap K_i$ under the map $\phi^{m_i q_i}$ and an arc in Σ° traveling a clockwise angle of $2\pi p_i$ around the singularity. Parallel transport along the first arc stretches distances by a factor of $\lambda^{m_i q_i}$. Roughly speaking, the second arc contains p_i subarcs of type 3b and p_i subarcs which are (the inverse of) type 3a. Together, $\mathcal{J}^{\text{partial}}$ -parallel transport along these arcs produces a dilation of a factor of $(\lambda^{-m_i q_i / p_i})^{p_i}$. Thus, the composite monodromy is dilation by a factor $\lambda^{m_i q_i} (\lambda^{-m_i q_i / p_i})^{p_i} = 1$.

In more detail, let us analyze parallel transport around a short arc traveling an angle of 2π around the singularity. Figure 12 shows the foliation $\mathcal{F}^u \cap \Sigma^\circ$ near a singularity along with an arc γ traveling an angle of 2π around the singularity. The arc γ is chosen so that $\gamma(0)$, $\gamma(1/2)$ and $\gamma(1)$ lie on prongs of the

stable foliation $\mathcal{F}^s \cap \Sigma^\circ$ (not shown). The θ -images of half fibers of π are labeled P_1, \dots, P_6 , so that, for example, $\theta \circ \pi^{-1}(\gamma(0)) = P_1 \cup P_2$ and $\gamma(0) = P_1 \cap P_2$. We will abuse notation and use P_1 to refer to both a half leaf $\mathcal{F}^u \cap \Sigma^\circ$ and its θ -preimage in $\pi^{-1}(\gamma(0))$, and similarly for P_2, \dots, P_6 .

Let R_1, R_2 and R_3 be rectangles bounded on the east and west sides by the P_i 's and the prongs of $\mathcal{F}^u \cap \Sigma^\circ$ as shown shaded in Figure 12. We might not be able to make the rectangles arbitrarily long in the north–south direction due to hitting other singularities. However, if $\gamma(0), \gamma(1/2)$ and $\gamma(1)$ hug very close to the singularity, these rectangles can be chosen to be as long in the north–south direction, as desired.

Restrict attention to points in the P_i such that the θ -images of their parallel transports along γ stay inside $R_1 \cup R_2 \cup R_3$. Call the set of such points N . For example, $N \cap (P_1 \cup P_2)$ is an interval containing $\gamma(0)$.

Now $\gamma|_{[0,1/2]}$ decomposes as a type 2 curve in R_1 , the inverse of a type 3a curve crossing the prong, and a type 2 curve in R_2 . Therefore, $\mathcal{J}_{\gamma|_{[0,1/2]}}^{\text{partial}}$ maps $P_2 \cap N$ to $P_3 \cap N$ by an isometry and maps $P_1 \cap N$ to $P_4 \cap N$ by a stretch by a factor $\lambda^{-m_i q_i / p_i}$. (It is worth checking in Figure 12 that the prong is indeed crossed from east to west, so γ follows the inverse of a type 3a curve and the corresponding dilation factor has negative exponent.) Similarly, $\gamma|_{[1/2,1]}$ decomposes as a type 2 curve in R_2 , a type 3b curve, and a type 2 curve in R_3 . Thus, $\gamma|_{[1/2,1]}$ maps $P_4 \cap N$ to $P_5 \cap N$ by an isometry and $P_3 \cap N$ to $P_6 \cap N$ by a stretch by a factor $\lambda^{-m_i q_i / p_i}$. The composite $\mathcal{J}_\gamma^{\text{partial}}$ maps $(P_1 \cup P_2) \cap N$ to $(P_5 \cup P_6) \cap N$ by a dilation of $\lambda^{-q_i m_i / p_i}$. This justifies and makes precise the claim made in the first paragraph of the proof.

Finally, if we take $\gamma(0), \gamma(1)$ and $\gamma(1/2)$ to hug close to the singularity, then R_1, R_2 and R_3 may be made as long as needed, and in turn N may be made as large as desired. By the accounting of dilation factors from the first paragraph, the monodromy of γ is the identity on as large a subspace of $\pi^{-1}(p)$ as desired. □

6.3 Blowup time for parallel sections

Let $\gamma: [0, \infty) \rightarrow \Sigma^\circ$ parametrize an eastward ray in a leaf of $\Sigma^\circ \cap \mathcal{F}^s$ by arclength. (Here, we are conflating Σ° with its lift to Z , and will continue to do so without comment. So, for example, $\gamma(0)$ might be of the form p_{west} or p_{east} for some p on a prong.) Given a point $x \in \pi^{-1}(\gamma(0))$, recall from Section 6.1 that one may attempt to construct a parallel section $s_x: [0, \infty) \rightarrow E_\pi$ with $s_x(0) = x$. However, s_x typically blows up in finite time and can only be defined over $[0, t_{\max}(x))$ for some $t_{\max}(x) > 0$. In this section, we show that t_{\max} is locally a homeomorphism from $\pi^{-1}(\gamma(0))$ to \mathbb{R} .

Remark 6.9 In what follows, we will abuse notation and treat s_x as a real-valued function of t whose value is the signed north–south distance in $\tilde{\Sigma}^\circ$ from $\gamma(t)$ to $\theta \circ s_x(t)$. Similarly, we will think of x as a real number.

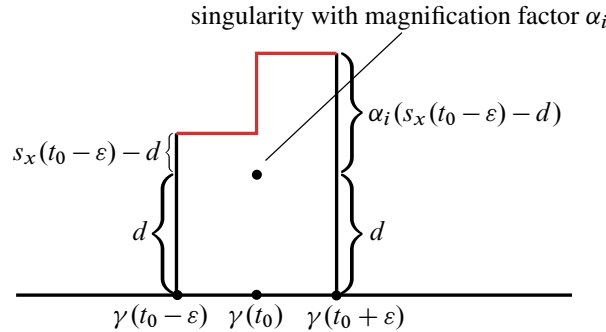


Figure 13: Constructing $s_x(t)$ by pushing a fiber of π past a singularity, as seen in the image of the developing map.

Recall that we defined λ so that $\lambda > 1$. We also assumed that p_i all have the same sign. Without loss of generality, we assume $p_i \geq 0$. Let $\alpha_i = \lambda^{m_i q_i / p_i}$. Note that $\alpha_i \geq 1$. Let $\alpha_{\max} = \max_i \lambda^{m_i q_i / p_i}$. If $\alpha_{\max} = 1$, then $M_\varphi(\mathbf{p}; \mathbf{q})$ is fibered and there is no blowup of sections. So we assume that $\alpha_{\max} > 1$. Now we can define $\alpha_{\min} = \min\{\lambda^{m_i q_i / p_i} \mid q_i > 0\}$. We call a singularity *magnifying* if the associated coefficient satisfies $\alpha_i > 1$. A *ragged rectangle* is the region in $\tilde{\Sigma}^\circ$ swept out by a interval in $\mathcal{F}^s \cap \Sigma^\circ$ under the time t northward flow along $\mathcal{F}^u \cap \Sigma^\circ$ (ie the map that sends each point t units due north). It looks like a rectangle with some vertical slits cut into the northern edge.

It is best to visualize all our constructions via the developing map $D: \tilde{\Sigma}^\circ \rightarrow \mathbb{R}^2$, which is locally an isometry. We can arrange that $D(\widetilde{\mathcal{F}^s \cap \Sigma^\circ})$ and $D(\widetilde{\mathcal{F}^u \cap \Sigma^\circ})$ are the horizontal and vertical foliations of \mathbb{R}^2 , respectively. Thus, the θ -images of fibers of π also correspond to vertical lines in this picture. Ragged rectangles project via D to honest rectangles. We can also arrange that the image of γ is the positive x axis.

Given $x \in \pi^{-1}(\gamma(0))$, the section $s_x(t)$ may be constructed geometrically. If $x > 0$, push the vertical segment $[\gamma(0), \theta(x)]$ eastward along γ . The endpoint sweeps out the θ image of the section $s_x(t)$. See Figure 13. Whenever the segment hits a singularity in orbit i at distance d from $\gamma(t)$ with $0 < d < s_x(t)$, the segment continues on the other side with length $d + \alpha_i(s_x(t) - d)$. If $x < 0$, then the same procedure works except that the new length of the segment is $d + (1/\alpha_i)(s_x(t) - d)$.

When $x < 0$, the section $s_x(t)$ stays bounded for all time because γ can be expressed as a composite of type 2 and type 3b curves, neither of which increases $|s_x(t)|$. On the other hand, when $x > 0$, $s_x(t)$ grows with t and possibly blows up. Lemma 6.10 below guarantees that we can expect many singularities to be encountered during this procedure.

Lemma 6.10 *There exists a positive constant κ such that for every $\varepsilon > 0$, there exists a large enough A_ε so that every ragged rectangle of area $A > A_\varepsilon$ has number of magnifying singularities in the range $[(\kappa - \varepsilon)A, (\kappa + \varepsilon)A]$.*

Proof First, we may use the action of φ to turn any ragged rectangle of area A into a new ragged rectangle of width $O(1)$ and height $O(A)$ having the same number of magnifying singularities. Now the lemma follows from the ergodicity of the translation flow on flat surfaces. Masur’s criterion states that the translation flow on a flat surface is ergodic whenever the corresponding flow on Teichmüller space given by multiplying the metric by

$$\begin{pmatrix} e^t & 0 \\ 0 & e^{-t} \end{pmatrix}$$

stays in some compact set; see [28, Section 3.7]. A pseudo-Anosov map corresponds to a closed orbit under this flow. Therefore, Masur’s criterion is fulfilled. \square

We now define A_* to be a constant large enough that every ragged rectangle of area A_* contains a magnifying singularity. We also remind the reader that in what follows, we are using the identifications introduced in Remark 6.9.

Lemma 6.11 *There exists a constant C independent of γ such that whenever $x > S > 0$, $s_x(t)$ blows up in time $t_{\max}(x) < C/S$.*

Proof By Lemma 6.10, a prong from a magnifying singularity intersects the image of γ in Σ° at north–south distance $< S/2$ in time $t < A_* \cdot (2/S)$. Therefore,

$$(2) \quad s_x\left(\frac{2A_*}{S}\right) > \frac{S}{2} + \alpha_{\min} \frac{S}{2} = \frac{1 + \alpha_{\min}}{2} S.$$

Repeating the argument to find more nearby singularities, we find that

$$(3) \quad s_x\left(\frac{2A_*}{S} \sum_{i=0}^{n-1} \left(\frac{1 + \alpha_{\min}}{2}\right)^{-i}\right) > \left(\frac{1 + \alpha_{\min}}{2}\right)^n S$$

as long as the left side is defined. Taking $n \rightarrow \infty$, we find that

$$(4) \quad t_{\max}(x) < \frac{2A_*}{S} \left(1 - \frac{2}{1 + \alpha_{\min}}\right)^{-1}.$$

The right side is of the form C/S , as desired. \square

Lemma 6.12 *There exists a constant c independent of γ such that whenever $0 \leq x < S$, the section $s_x(t)$ exists for time c/S .*

Proof By Lemma 6.10, we can choose a constant A_κ so that every ragged rectangle of area $A > A_\kappa$ has fewer than $2\kappa A$ singularities. Choose $c = A_\kappa / ((\alpha_{\max})^{2\kappa A_\kappa})$. Choose any $S > 0$ and any $0 \leq x < S$.

Consider a ragged rectangle R whose southwest corner lies at $\gamma(0)$ and which has an area A_κ , with height $(\alpha_{\max})^{2\kappa A_\kappa} S$ and width $c/S = A_\kappa / ((\alpha_{\max})^{2\kappa A_\kappa} S)$. This ragged rectangle has fewer than $2\kappa A_\kappa$ singularities. Since $|s_x(t)|$ grows by at most α_{\max} each time it encounters a singularity, $\theta \circ s_x(t)$ does not grow enough times to exit through the top of R . Therefore, $s_x(t)$ exists for time at least the width of the ragged rectangle, c/S . \square

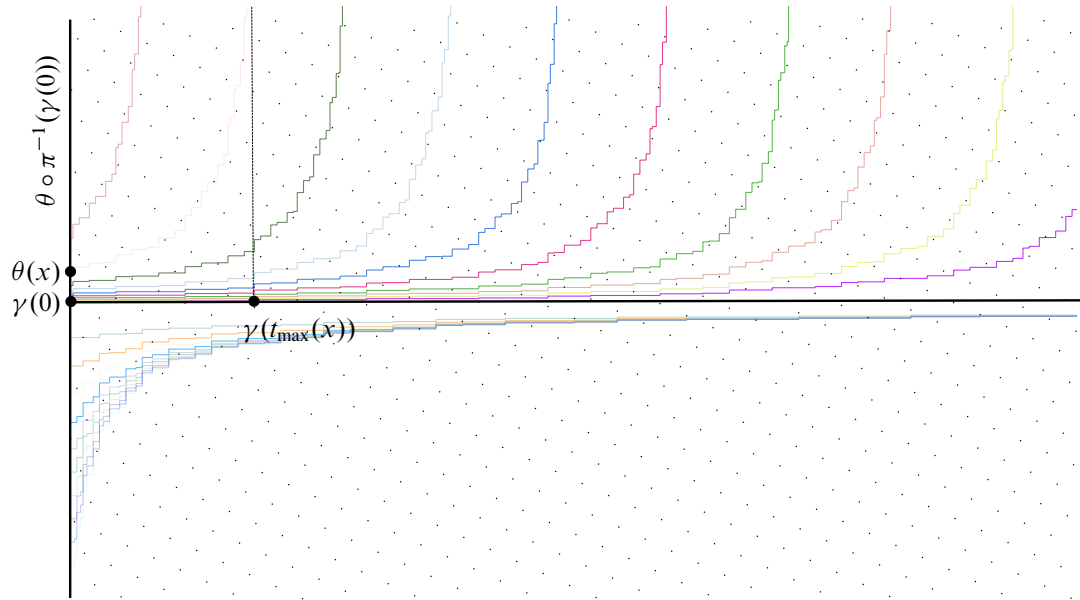


Figure 14: In this example, $\Sigma^\circ = T^2$, φ is the monodromy of the figure-eight knot, and the surgery coefficient is $(p_1; q_1) = (5, 1)$. This picture shows $D(\tilde{\Sigma}^\circ) \cong \mathbb{R}^2$ with the 2-pronged singularities of φ drawn as dots. Objects in Σ have been lifted to $\tilde{\Sigma}^\circ$, and then projected to \mathbb{R}^2 via D ; for convenience in labeling the picture, we identify objects with their images under this procedure. Sections $s_x(t)$ are shown in colors for various choices of x . These sections have been projected to Σ° by θ , then lifted to $\tilde{\Sigma}^\circ$, and finally projected to $D(\tilde{\Sigma}^\circ)$. Whenever $x \in \pi^{-1}(\gamma(0))$ satisfies $x < 0$, the section $s_x(t)$ exists for all time. When $x > 0$, $s_x(t)$ blows up in finite time.

Proposition 6.13 *The map $t_{\max} : \{x \in \pi^{-1}(\gamma(0)) \mid x > 0\} \rightarrow (0, \infty)$ is surjective.*

Proof We first show that the image of t_{\max} is dense. For any t_0 and any $\varepsilon > 0$, Lemma 6.11 guarantees that if $x \in \pi^{-1}(\gamma(t_0))$ is large enough, the parallel section over γ passing through x will blow up before $t = t_0 + \varepsilon$. A section can always be extended backwards in time because as one travels west, the section is always decreasing and positive. See Figure 14. Thus, we have found a section defined at $t = 0$ and which blows up between time t_0 and time $t_0 + \varepsilon$. Our choices of t_0 and ε were arbitrary, so t_{\max} has dense image.

Since t_{\max} is nondecreasing and has dense image, it is surjective. □

Proposition 6.14 *The map $t_{\max} : \{x \in \pi^{-1}(\gamma(0)) \mid x > 0\} \rightarrow (0, \infty)$ is injective.*

Proof Choose $x, y \in \pi^{-1}(\gamma(0))$. Assume $0 < x < y$. We want to show that s_y explodes strictly before s_x .

Case 2 ($s_y(t)/s_x(t)$ is unbounded) Choose t so that $s_y(t)/s_x(t) > C/c$, where C and c are the constants from Lemmas 6.11 and 6.12. By Lemma 6.11 with $S = s_y(t)$, s_y blows up before time $t + C/s_y(t)$. By Lemma 6.12 with $S = s_y(t)c/C$, s_x exists at time $t + C/s_y(t)$. Therefore, s_y explodes strictly before s_x .

Case 2 ($s_y(t)/s_x(t)$ stays bounded) Let $\{t_j\}$ be the sequence of times at which s_x grows, and let b_j be the height of the singularity encountered at time t_j . Note that $0 < b_j < s_x(t_j)$. Let $t_j - \varepsilon$ be a time just before crossing the singularity and $t_j + \varepsilon$ a time just after. This helps us avoid talking about $s_x(t_j)$ which is problematic since s_x is discontinuous at t_j . We have

$$\begin{aligned} \frac{s_y(t_j + \varepsilon) - s_x(t_j + \varepsilon)}{s_x(t_j + \varepsilon)} &= \frac{\alpha_{t_j} \cdot (s_y(t_j - \varepsilon) - s_x(t_j - \varepsilon))}{\alpha_{t_j} \cdot (s_x(t_j - \varepsilon) - b_j) + b_j} = \frac{s_y(t_j - \varepsilon) - s_x(t_j - \varepsilon)}{s_x(t_j - \varepsilon) - b_j(1 - 1/\alpha_{t_j})} \\ &= \frac{1}{1 - (1 - 1/\alpha_{t_j}) \cdot b_j/s_x(t_j - \varepsilon)} \cdot \frac{s_y(t_j - \varepsilon) - s_x(t_j - \varepsilon)}{s_x(t_j - \varepsilon)}, \end{aligned}$$

where α_{t_j} is the magnification factor at the singularity encountered at time t_j . Since $\alpha_{t_j} > 1$, we have $0 < 1 - 1/\alpha_{t_j} < 1$. Since α_{t_j} takes on at most finitely many values, $1 - 1/\alpha_{t_j}$ is bounded away from 0. If $b_j/s_x(t_j - \varepsilon)$ is also bounded away from 0, then $(s_y(t) - s_x(t))/s_x(t)$ grows by a factor bounded away from 1. Since we assumed that $s_y(t)/s_x(t)$ remains bounded, we must have $\lim_{j \rightarrow \infty} b_j/s_x(t_j - \varepsilon) = 0$.

Let $N \in \mathbb{N}$ and $\varepsilon > 0$ be constants to be determined. Choose n large enough that $b_i/s_x(t_i) < \varepsilon$ for all $i > n$. By Lemma 6.11, there exists $m > n$ such that

$$(5) \quad (\alpha_{\max})^{N-1} s_x(t_n) < s_x(t_m) < (\alpha_{\max})^{N+1} s_x(t_n),$$

$$(6) \quad t_m - t_n < \frac{C}{s_x(t_n)}.$$

Now consider the ragged rectangle R with base $\gamma([t_n, t_m])$ and height $\varepsilon s_x(t_m)$. We have that

$$(7) \quad \text{Area}(R) = \varepsilon s_x(t_m)(t_m - t_n) < \varepsilon s_x(t_n) \frac{C}{s_x(t_n)} (\alpha_{\max})^{N+1} = C \varepsilon (\alpha_{\max})^{N+1}.$$

Since $b_j/s_x(t_j) < \varepsilon$ for every $j \in [n, m]$, the ragged rectangle contains every singularity which the section s_x encounters between time t_n and t_m . Taking into account equation (5) and the fact that s_x can grow by a factor of at most α_{\max} each time it encounters a singularity, there must be at least N singularities in R . Taking N large enough and then ε small enough, we find a ragged rectangle of arbitrarily small area containing arbitrarily many singularities. Such a rectangle may be extended northward to a ragged rectangle of any desired larger area with at least as many singularities. This contradicts Lemma 6.10. \square

6.4 The full connection \mathcal{J}

The uniqueness statements for blowup proven in the previous section suggest that we add a point at infinity to each fiber of π , obtaining a new bundle $S^1 \rightarrow E_\Pi \xrightarrow{\Pi} M_\varphi^\circ$. Via the embedding of fibers of π into fibers of Π , $\mathcal{J}^{\text{partial}}$ is naturally a partial connection on Π as well. We now define an honest connection \mathcal{J} on Π which extends $\mathcal{J}^{\text{partial}}$, ie for all $\gamma \subset M_\varphi^\circ$, we have $\mathcal{J}_\gamma = \mathcal{J}_\gamma^{\text{partial}}$ on the domain of $\mathcal{J}_\gamma^{\text{partial}}$ in $\Pi^{-1}(\gamma(0))$.

On a type 1 curve γ , the domain of $\mathcal{J}_\gamma^{\text{partial}}$ is all of $\pi^{-1}(\gamma(0))$. We define \mathcal{J} in the only sensible way which extends $\mathcal{J}^{\text{partial}}$, asking that \mathcal{J}_γ sends the point at infinity in $\Pi^{-1}(\gamma(0))$ to the point at infinity in $\Pi^{-1}(\gamma(1))$.




Figure 15: We continue the example from Figure 14. As before, we are drawing all objects lifted to $\tilde{\Sigma}^\circ$ and then projected via D to \mathbb{R}^2 . The figure shows the θ -image of a parallel section for \mathcal{J} over a curve γ oriented east–west. The curve γ is drawn as a horizontal black line in the figure. For any t , $\theta \circ \pi^{-1}(\gamma(t))$ would be a vertical line in the picture. The section is piecewise constant. It starts out equal to s_0 , wraps around ∞ at $t = t_{\max}(x)$, and continues as s_1 .

It remains to define \mathcal{J} on curves oriented east–west, ie those in leaves of $\Sigma^\circ \cap \mathcal{F}^s$ which were considered in Section 6.3. Such a curve may be represented as a (possibly infinite) concatenation of type 2 and type 3 curves. Given $\gamma: [0, t] \rightarrow M_\varphi^\circ$ parametrizing an eastward interval in a leaf of $\Sigma^\circ \cap \mathcal{F}^s$ by arclength and a point $x \in \Pi^{-1}(\gamma(0))$, let s_0 be the parallel section over γ traced out by parallel transport of x along γ using $\mathcal{J}^{\text{partial}}$. Suppose s_0 blows up towards $+\infty$ at time $t_{\max}(x)$. (When $x = \infty$, we say that $t_{\max}(x) = 0$, and when $x < 0$ we have $t_{\max}(x) = \infty$.) By Propositions 6.13 and 6.14, there exists a unique parallel section s_1 defined on $(t_{\max}(x), \infty)$ that blows up towards $-\infty$ as t approaches $t_{\max}(x)$ from the right. We now define \mathcal{J} in cases:

$$(8) \quad \mathcal{J}_\gamma(x) = \begin{cases} s_0(t) & \text{if } t < t_{\max}(x), \\ s_1(t) & \text{if } t > t_{\max}(x), \\ \infty & \text{if } t = t_{\max}(x). \end{cases}$$

Proposition 6.15 \mathcal{J} is flat.

Proof Consider three types of short arcs in M_φ° : eastward arcs, northward arcs and transverse arcs, ie arcs contained in flowlines of the suspension flow of φ . Every curve in M_φ° can be C^0 approximated by

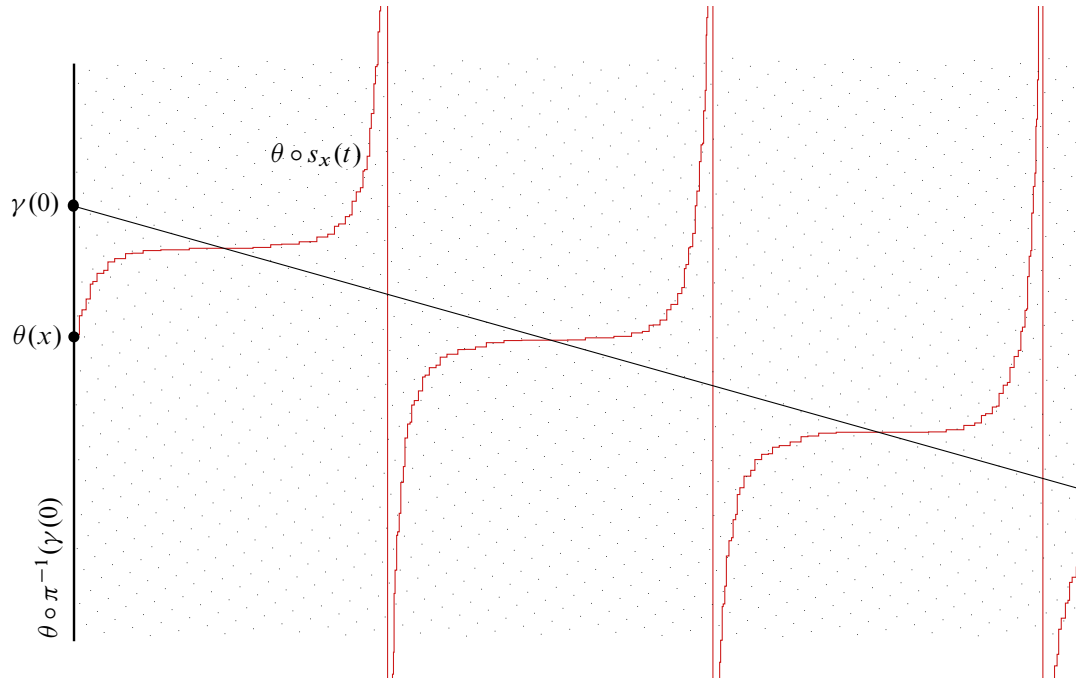


Figure 16: The θ -image of a parallel section for \mathcal{J} over a curve γ of negative slope relative to the stable foliation on Σ° . This picture is drawn with the same conventions as in Figure 15. The section s_x intersects the ∞ -section infinitely many times. To see this, first homotope γ in Σ° to a curve which is piecewise southward or eastward with each straight segment having length at least ε for some $\varepsilon > 0$. This does not change $\theta \circ s_x$, since \mathcal{J} is flat. Since γ had negative slope and $\theta \circ s_x$ is nondecreasing, the north-south distance between $\theta \circ s_x(t)$ and $\gamma(t)$ increases to C/ε in finite time, where C is the constant from Lemma 6.11. By that lemma, s_x blows up and crosses the ∞ -section within the next eastward segment of γ . This process repeats indefinitely and produces infinitely many intersections between s_x and the ∞ -section.

such curves. As in Proposition 6.7, it is enough to show that the monodromy of \mathcal{J} around all commutators is trivial.

First, let's check that the monodromy around a commutator of eastward and northward arcs vanishes (ie the monodromy around the boundary of a rectangle is trivial). Let γ_1 and γ_2 be the north and south sides of the rectangle respectively. The θ -image of any $\mathcal{J}^{\text{partial}}$ -parallel section over γ_1 is also the θ -image of a $\mathcal{J}^{\text{partial}}$ -parallel section over γ_2 ; this follows from flatness of $\mathcal{J}^{\text{partial}}$ combined with the fact that the θ -image of a $\mathcal{J}^{\text{partial}}$ -parallel section along any northward arc is a single point. A \mathcal{J} -parallel section is formed by patching together $\mathcal{J}^{\text{partial}}$ parallel sections at their asymptotes according to equation (8). Therefore, the θ -images of \mathcal{J} -parallel sections over γ_1 are also θ -images of \mathcal{J} -parallel sections over γ_2 . It follows that the desired monodromy is the identity.

Second, the commutator of northward and transverse arcs vanishes because both are type 1 curves.

Finally, consider the commutator of a transverse arc of length ε and an eastward arc γ . Observe that the entire construction of \mathcal{J}_γ for an eastward arc γ is equivariant with respect to dilations in the north–south or east–west directions, and in particular, a dilation by λ^ε in the north–south direction and $\lambda^{-\varepsilon}$ in the east–west direction where ε . This is equivalent to the vanishing of the desired monodromy. \square

We can now rectify the lack of a good topology on E_Π . A flat connection gives a local product structure on E_Π , so we endow E_Π with the corresponding local product topology.

Π has a special section, called the ∞ –section and denoted by r_∞ , whose value at any point is the point at infinity in the corresponding fiber. There is another section, called the zero–section and denoted by r_0 , defined by the property

$$\theta(r_0(p)) = p \quad \text{for all } p \in M_\varphi^\circ.$$

Note that neither of these sections is flat with respect to \mathcal{J} . Nevertheless, these sections are continuous.

Lemma 6.16 *The ∞ –section and the zero–section are both continuous sections of Π .*

Proof Lemma 6.12 gives a quantitative bound on how fast flat sections can explode to ∞ in the east and west directions. In the north–south and transverse directions, flat sections do not blow up. It follows that the graph of the ∞ –section is closed and r_∞ is continuous.

Now let’s check continuity of the zero–section near a point $p \in M_\varphi^\circ$. Take a neighborhood U of a point $p \in M_\varphi^\circ$ which \mathcal{F}^s and \mathcal{F}^u foliate as products. The set

$$U' = \{y \in E_\pi \mid \theta(y) \in U, \pi(y) \in U\}$$

is a neighborhood of $r_0(p)$. Since holonomy of \mathcal{F}^s exists for all transverse arcs in U , the topology on U' is just the standard product topology. The graph of the zero–section is clearly closed in this topology. \square

Lemma 6.17 *The zero–section is flat over any leaf of \mathcal{F}^s .*

Proof This follows immediately from the way we defined \mathcal{J} to agree with the holonomy of \mathcal{F}^s in equation (1). If λ is a leaf of \mathcal{F}^s and γ is a path in λ , then we have

$$\begin{aligned} h_\gamma(\gamma(0)) &= \gamma(1), \\ \mathcal{J}_\gamma(r_0(\gamma(0))) &= \theta^{-1} \circ h_\gamma \circ \theta(r_0(\gamma(0))) = \theta^{-1} \circ h_\gamma(\gamma(0)) = r_0(\gamma(1)). \end{aligned} \quad \square$$

Proposition 6.18 *The monodromy of \mathcal{J} around a curve parallel to the Dehn filling slope on a toroidal end of M_φ° is trivial. Moreover, a parallel section over such a curve has zero intersection number with the ∞ –section.*

Proof Since \mathcal{J} is a flat connection, the conjugacy class of its monodromy around a closed curve is invariant under free homotopies of the closed curve. Thus, it suffices to show that the desired monodromy around some freely homotopic curve is trivial. Fix a point p on an unstable prong incident to the

toroidal end. Consider a loop based at p freely homotopic to our curve. Proposition 6.8 tells us that the monodromy of $\mathcal{J}^{\text{partial}}$ around this loop is the identity, and moreover that the domain of $\mathcal{J}^{\text{partial}}$ can be made as large as we like by choosing the loop to hug very close to the toroidal end. Since \mathcal{J} agrees with $\mathcal{J}^{\text{partial}}$ when they are both defined, we conclude that the monodromy is the identity for all points in the fiber not equal to ∞ . Since \mathcal{J} is a bijection from $\Pi^{-1}(p)$ to itself, it must be equal to the identity at ∞ as well.

For the second statement, it suffices to check that just one parallel section of \mathcal{J} has zero intersection number with the ∞ -section. Simply take any parallel section of $\mathcal{J}^{\text{partial}}$ along the meridian; as noted in the previous paragraph, such a section has no intersections with the ∞ -section. \square

Proposition 6.18 guarantees that (Π, \mathcal{J}) extends to a bundle over $M_\varphi(\mathbf{p}; \mathbf{q})$ with a flat connection. Moreover, the ∞ -section extends to a section over $M_\varphi(\mathbf{p}; \mathbf{q})$. Thus, the Euler class of Π vanishes and the S^1 -bundle with flat connection (Π, \mathcal{J}) unrolls to an \mathbb{R} -bundle with flat connection which we call $(\widehat{\Pi}, \widehat{\mathcal{J}})$. We choose this unrolling so that the ∞ -section lifts to a section of $\widehat{\Pi}$. More detail on this construction is provided in Remark 6.19.

Remark 6.19 Suppose X is a space with an S^1 bundle $S^1 \rightarrow E \xrightarrow{\Pi} X$ and a flat connection \mathcal{H} . A fiberwise cover of Π is an bundle $\mathbb{R} \rightarrow \widehat{E} \xrightarrow{\widehat{\Pi}} X$ along with projection map β making the following diagram commute:

$$\begin{array}{ccccc}
 \mathbb{R} & \longrightarrow & \widehat{E} & & \\
 \text{mod } 1 \downarrow & & \downarrow \beta & \searrow \widehat{\Pi} & \\
 S^1 & \longrightarrow & E & \xrightarrow{\Pi} & X
 \end{array}$$

Fiberwise covers of Π are classified by elements of $H^1(E, \mathbb{Z})$ which evaluate to 1 on each fiber. Such an element exists if and only if the Euler class of Π vanishes. By Poincaré duality, such elements are in one-to-one correspondence with homotopy classes of sections of Π . Given such a class $\varphi \in H^1(E, \mathbb{Z})$, we may construct \widehat{E} as the corresponding Abelian cover of E . Let $\beta: \widehat{E} \rightarrow E$ be the corresponding covering map and let $\widehat{\Pi} = \Pi \circ \beta$. Since φ evaluates to 1 on each fiber of Π , the $\widehat{\Pi}$ preimage of a point in X is a copy of \mathbb{R} as desired. Observe that under this construction, any section of E Poincaré dual to φ lifts to a section of \widehat{E} . In the context of this article, we are using the fiberwise lift corresponding to the ∞ -section.

Suppose now that Π is equipped with a flat connection \mathcal{H} . Then on any fiberwise cover $\widehat{\Pi}$, there is an induced flat connection $\widehat{\mathcal{H}}$. One may define $\widehat{\mathcal{H}}$ as follows. A flat connection on the bundle Π is the same thing as a foliation of E by leaves transverse to the S^1 fibers; the flat connection may be recovered as holonomy of this foliation. The lift of this foliation to \widehat{E} is transverse to the fibers of $\widehat{\Pi}$. Moreover, the holonomy of the lifted foliation exists along any path γ in the base. This holonomy map is what we call $\widehat{\mathcal{H}}_\gamma$.

Proof of Theorem 1.1 The flat connection $\widehat{\mathcal{J}}$ provides a homomorphism from $\pi_1(M_\varphi(\mathbf{p}; \mathbf{q}))$ into $\text{Homeo}^+(\mathbb{R})$. This map is nontrivial since the monodromy around any closed orbit of the restriction of

the suspension flow of φ to M_φ° is nontrivial; indeed, the monodromy of $\mathcal{J}^{\text{partial}}$ around such a curve is a dilation. Theorem 1.1 of [6] states that for fundamental groups of irreducible 3-manifolds, the existence of any nontrivial map to a left-orderable group is equivalent to the existence of a left-ordering. \square

6.5 The fiber of $\widehat{\Pi}$ and the leaf space

In this section we will prove Theorem 1.2. Our strategy is to construct a monotone map from the leaf space of $\widetilde{\mathcal{F}}$ to a chosen fiber of $\widehat{\Pi}$. Before proceeding with the proof, we briefly summarize the relevant constructions from the previous section. Π is an S^1 bundle over $M_\varphi(\mathbf{p}; \mathbf{q})$ with trivial Euler class. \mathcal{J} is a flat connection on Π . We defined a section r_0 of Π called the zero-section. The zero-section is defined by the property that $\theta(r_0(p)) = p$. The \mathbb{R} -bundle $\widehat{\Pi}$ is a fiberwise cover of Π , and is equipped with a flat connection $\widehat{\mathcal{J}}$.

Proof of Theorem 1.2 Let $\overline{M}_\varphi^\circ$ be the lift of M_φ° to $\widetilde{M}_\varphi(\mathbf{p}; \mathbf{q})$. We use bars to denote the lifts of objects in M_φ° (for example \mathcal{F}^s) to objects in $\overline{M}_\varphi^\circ$ (for example $\overline{\mathcal{F}}^s$). Recall that L is the leaf space of $\widetilde{\mathcal{F}}$. Let K be the leaf space of $\overline{\mathcal{F}}^s$, where prongs incident with the same singularity are considered the same leaf. This is the same as saying that K is the leaf space of the induced stable foliation in $\widetilde{M}_\varphi(\mathbf{p}; \mathbf{q})$. Let $f_1: L \rightarrow K$ be the monotone, $\pi_1(M_\varphi(\mathbf{p}; \mathbf{q}))$ -equivariant map which crushes each interval in L corresponding with a lift of a saddle region down to a point.

By construction, the ∞ -section of the S^1 -bundle Π lifts to a section (not just a multivalued section) of the \mathbb{R} -bundle $\widehat{\Pi}$. Refer to Remark 6.19 for more discussion of this point. The zero-section is homotopic to the ∞ section, so it also lifts to a section of $\widehat{\Pi}$. Choose such a lift and call it \widehat{r}_0 .

Select a basepoint $p \in M_\varphi^\circ$ and a lift $\bar{p} \in \overline{M}_\varphi^\circ$. We define a monotone, $\pi_1(M_\varphi(\mathbf{p}; \mathbf{q}))$ -equivariant map

$$f_2: K \rightarrow \widehat{\Pi}^{-1}(p)$$

as follows. For any leaf $\bar{\lambda}$ in $\overline{\mathcal{F}}^s$, choose a point \bar{q} on $\bar{\lambda}$ and a path $\bar{\gamma}: [0, 1] \rightarrow \overline{M}_\varphi^\circ$ from \bar{q} to \bar{p} . Let q, λ and γ be the projections of $\bar{q}, \bar{\lambda}$ and $\bar{\gamma}$ to M_φ° , and similarly let γ be the projection of $\bar{\gamma}$ to M_φ° . We now define

$$f_2(\bar{\lambda}) = \widehat{\mathcal{J}}_\gamma(\widehat{r}_0(q)).$$

Since $\widehat{\mathcal{J}}$ extends to a flat bundle over $M_\varphi(\mathbf{p}; \mathbf{q})$, this definition is independent of the choice of γ in its relative homotopy class in $M_\varphi(\mathbf{p}; \mathbf{q})$. As discussed in Lemma 6.17, the zero-section is flat on λ . Therefore, our definition is independent of the choice of q on λ . Finally, it's easy to check that two prongs of $\overline{\mathcal{F}}^s$ corresponding with the same point in K have the same value of f_2 . One way to see this is to connect two adjacent prongs by a short path γ_p , note that $h_{\gamma_p}(\gamma_p(0)) = \gamma_p(1)$, and then argue as in Lemma 6.17 that f_2 has the same value on these two prongs. Then one could iterate to show that all prongs incident to the same singularity have the same value of f_2 .

Let's check that f_2 is π_1 -equivariant, ie that it intertwines the action on K by deck transformations and the action on $\hat{\Pi}^{-1}(p)$ induced by $\hat{\mathcal{J}}$. Given $g \in \pi_1(M_\varphi(\mathbf{p}; \mathbf{q}))$, represent g by a loop γ_2 in M_φ° based at p . We remind the reader that g acts on $\hat{\Pi}^{-1}(p)$ as $\hat{\mathcal{J}}_{\gamma_2^{-1}}$, as explained in Section 6.1. Now we have

$$\begin{aligned} f_2(g\bar{\lambda}) &= \hat{\mathcal{J}}_{\gamma * \gamma_2^{-1}}(\hat{r}_0(q)) && \text{since } \gamma * \gamma_2^{-1} \text{ lifts to a path from } g\bar{\lambda} \text{ to } \bar{p}, \\ &= \hat{\mathcal{J}}_{\gamma_2^{-1}} \circ \hat{\mathcal{J}}_\gamma(\hat{r}_0(q)) \\ &= \hat{\mathcal{J}}_{\gamma_2^{-1}} f_2(\bar{\lambda}) \\ &= g f_2(\bar{\lambda}), && \text{since } g \text{ acts on } \hat{\Pi}^{-1}(p) \text{ as } \hat{\mathcal{J}}_{\gamma_2^{-1}}. \end{aligned}$$

Next, we will check that for any leaf λ of $\overline{\mathcal{F}^s}$, there is an interval in K around λ which f_2 maps homeomorphically onto an interval in $\hat{\Pi}^{-1}(p)$. (Though K is not a 1-manifold, the notion of interval still makes sense: we mean the lift to K of an interval in $\overline{M}_\varphi^\circ$ which is transverse to $\overline{\mathcal{F}^s}$.) Roughly speaking, the local leaf space of $\overline{\mathcal{F}^s}$ is homeomorphic to the local leaf space of \mathcal{F}^s , and θ gives rise to a homeomorphism between the local leaf space of \mathcal{F}^s and an interval in a fiber of $\hat{\Pi}$. Finally, $\hat{\mathcal{J}}$ gives a homeomorphism from any fiber of $\hat{\Pi}$ to $\hat{\Pi}^{-1}(p)$. The function f_2 can be expressed as a composition of these maps, and hence is locally a homeomorphism.

In more detail, choose a short curve $\bar{\gamma}_3 : [0, 1] \rightarrow \overline{M}_\varphi^\circ$ that is transverse to $\overline{\mathcal{F}^s}$, lies in the strong unstable foliation, and passes through the leaf λ . Then $\bar{\gamma}_3$ parametrizes an interval in K containing λ . We will check that f_2 is a homeomorphism on this interval. Let $\bar{q} = \bar{\gamma}_3(1)$. Let $\bar{\gamma}$ be a path from \bar{q} to \bar{p} . Let γ , γ_3 and q be the projections of $\bar{\gamma}$, $\bar{\gamma}_3$ and \bar{q} to M_φ° . Since γ_3 lies in the strong unstable foliation, γ_3 is the $\hat{\theta}_q$ -image of some curve γ_4 in $\hat{\Pi}^{-1}(q)$. Here, $\hat{\theta}_q : \hat{\Pi}^{-1}(q) \rightarrow M_\varphi^\circ$ is the composition of θ with the covering map $\hat{\Pi}^{-1}(q) \rightarrow \Pi^{-1}(q)$. A priori, there are many choices for γ_4 because the covering map $\hat{\Pi}^{-1}(q) \rightarrow \Pi^{-1}(q)$ is \mathbb{Z} to 1. We eliminate this ambiguity by asking that

$$(9) \quad \gamma_4(1) = \hat{r}_0(\gamma_3(1)).$$

Using the definition of \mathcal{J} on the type 1 curve γ_3 , we have

$$\theta \circ \mathcal{J}_{\gamma_3|_{[t,1]}}(r_0(\gamma_3(t))) = \theta \circ r_0(\gamma_3(t)) = \gamma_3(t).$$

Therefore,

$$(10) \quad \hat{\theta}_q \circ \hat{\mathcal{J}}_{\gamma_3|_{[t,1]}}(\hat{r}_0(\gamma_3(t))) = \gamma_3(t),$$

$$(11) \quad \hat{\mathcal{J}}_{\gamma_3|_{[t,1]}}(\hat{r}_0(\gamma_3(t))) = \gamma_4(t).$$

A priori, equation (11) holds only up to deck transformations of the covering $\hat{\Pi}^{-1}(q) \rightarrow \Pi^{-1}(q)$. However, the normalization in equation (9) is designed so that equation (11) holds exactly.

Now let us compute f_2 on the desired interval. For any $t \in [0, 1]$,

$$(12) \quad f_2(\bar{\gamma}_3(t)) = \hat{\mathcal{J}}_\gamma \circ \hat{\mathcal{J}}_{\gamma_3|_{[t,1]}}(\hat{r}_0(\gamma_3(t)))$$

$$(13) \quad = \hat{\mathcal{J}}_\gamma(\gamma_4(t))$$

$$(14) \quad = \hat{\mathcal{J}}_\gamma(\hat{\theta}_q^{-1}(\gamma_3(t))).$$

In equation (12) we used that $\gamma_3|_{[t,1]} * \gamma$ lifts to a path from $\bar{\gamma}_3(t)$ to \bar{p} . In (13), we applied (11). In (14), we have used the fact that $\hat{\theta}_q$ restricts to a homeomorphism between the images of γ_4 and γ_3 . Equation (14) expresses f_2 locally as a composition of the homeomorphisms $\hat{\mathcal{J}}$ and $\hat{\theta}^{-1}$, so f_2 is a homeomorphism of $\bar{\gamma}_3$ onto its image as desired. The chosen leaf λ was arbitrary, so it follows that f_2 is continuous and monotone. We conclude that the composition $f = f_2 \circ f_1$ is a continuous, monotone, $\pi_1(M_\varphi(\mathbf{p}; \mathbf{q}))$ -equivariant map. The $\pi_1(M_\varphi(\mathbf{p}; \mathbf{q}))$ action on the image is nontrivial, as noted in the proof of Theorem 1.1. \square

Remark 6.20 The map f_2 can be visualized quite cleanly in the setting of Figure 16. Extend γ linearly in both directions and choose $p = \gamma(0)$. The vertical red lines (plus their completions at ∞) are θ -images of fibers of Π . Lift these fibers to half-open subintervals of fibers of the \mathbb{R} -bundle $\hat{\Pi}$. Choose the lifts so that they intersect the zero-section in $\hat{\Pi}$. Now parallel transport these subintervals to $\hat{\Pi}^{-1}(\gamma(0))$ along γ^{-1} . The transported intervals are disjoint and cover $\hat{\Pi}^{-1}(\gamma(0))$. To be more precise, let $\dots, x_{-1}, x_0, x_1, x_2, \dots$ be the lifts of the point $x \in \Pi^{-1}(\gamma(0))$ to $\hat{\Pi}^{-1}(\gamma(0))$. Parallel transport along γ^{-1} maps the first vertical red line (viewed as a half-open interval in a fiber of $\hat{\Pi}$) to $[x_{-1}, x_0)$. The second vertical red line is transported to $[x_{-2}, x_{-1})$, etc. During this parallel transport, the points at $\pm\infty$ of the vertical red lines follow the section s_x drawn in the figure.

Therefore, the vertical lines (plus their completions at ∞) contain a representative from each point preimage of f_2 . The quotient of the leaf space we have constructed (ie $\hat{\Pi}^{-1}(\gamma(0))$) can now be seen as the concatenation of all of the vertical red lines in Figure 16. This picture is analogous to that of the step map in the setting of skew-Anosov flows: the vertical red lines are unstable leaves in the orbit space, while the curved red arcs are analogous to stable leaves which make perfect fits with the unstable leaves. See [1, Proposition 3.22] for a discussion of the skew-Anosov picture.

7 Computations

Building on work of Dunfield and Bell, we were able to find 2598 manifolds in the Hodgson–Weeks census which can be constructed by a surgery satisfying the hypotheses of Theorem 1.1. This represents about 44.7% of the 5801 non-L-spaces in the Hodgson–Weeks census; see Hodgson and Weeks [18] and Dunfield [13]. Dunfield and Bell [3] found monodromies for many of the fibered, orientable 1-cusped manifolds that can be triangulated with at most 9 tetrahedra. Using Bell’s program flipper [2], they were able to find invariant laminations for about 25,700 of them. About 800 of these have orientable invariant laminations and monodromy preserving these orientations. The first few such examples with genus ≥ 2 are listed in Table 1, left.

We used flipper to drill these manifolds along their pseudo-Anosov singularities and produce ideal triangulations of the resulting many-cusped manifolds. Using SnapPy, we performed surgeries with small coefficients on these manifolds satisfying the constraints of Theorem 1.1, and identified the resulting

name	genus	name	underlying fibered manifold	volume
m038	2	m003(-2,3)	m004	0.981
m120	3	m004(6,1)	m004	1.284
s090	4	m004(1,2)	m004	1.398
v0224	5	m003(-3,4)	v0650	1.414
m221	3	m009(4,1)	m023	1.414
t00448	6	m004(3,2)	m004	1.440
o9_00896	7	m004(7,1)	m004	1.463
s173	4	m004(5,2)	m004	1.529
v0248	6	m015(5,1)	t03310	1.757
m289	2	m009(5,1)	m009	1.831
t00682	4	m009(-5,1)	m009	1.831
m305	2	m011(1,3)	v1577	1.831
s296	2	m009(1,2)	m009	1.843
m310	3	m007(-5,1)	o9_31045	1.843
t00707	3	m006(-5,1)	m009	1.941

Table 1: Left: the first few 1-cusped fibered manifolds with genus ≥ 2 and monodromy satisfying the conditions of Theorem 1.1. Right: the first few (closed) manifolds in the Hodgson–Weeks census to which Theorem 1.1 applies.

manifolds in the Hodgson–Weeks census [12]. We found 2598 manifolds in the census, the first few of which are shown in Table 1, right.

Dunfield [13] obtained orderability results for many manifolds in the Hodgson–Weeks census either by constructing a taut foliation with vanishing Euler class or by constructing a $\text{PSL}(2, \mathbb{R})$ -representation that lifts to a $\widetilde{\text{PSL}}(2, \mathbb{R})$ representation. Table 2 shows the overlap in applicability between these methods and ours.

8 Remarks and questions

- For which taut foliations does there exist a π_1 -equivariant, order-preserving map from the leaf space of the universal cover to \mathbb{R} ? As a first step, we suggest the following conjecture:

Conjecture 8.1 *Theorem 1.2 holds without the condition that the surgery slopes have the same sign.*

This would greatly expand the scope of the results in this paper; for example, by [16] it would include every 3-manifold carrying a transitive pseudo-Anosov flow with orientable invariant foliations. The difficulty is that the some of dilation factors $\lambda^{m_i p_i/q_i}$ may be greater than 1 in absolute value while others may be smaller than 1. Parallel sections are no longer monotonic as shown in Figure 14, and they might blow up as $t \rightarrow +\infty$ or as $t \rightarrow -\infty$. A more detailed study of their dynamics in this case is necessary.

		Theorem 1.1 applies?	
		Yes	No
Taut foliation or $\mathrm{PSL}(2, \mathbb{R})$ representation with Euler class 0?	Yes	1795	1730
	No	803	1473

Table 2: The overlap in applicability between our method and previously used methods for proving orderability on the non-L-space rational homology spheres in the Hodgson–Weeks census.

- Can the map $f: L \rightarrow \mathbb{R}$ be upgraded to a strictly monotone map? We expect that f factors as

$$(15) \quad \begin{array}{ccccc} & & f & & \\ & & \curvearrowright & & \\ L & \xrightarrow{\mu} & R & \xrightarrow{\eta} & \widehat{\Pi}^{-1}(p) \\ & & \parallel & & \parallel \\ & & \mathbb{R} & & \mathbb{R} \end{array}$$

where μ is locally a homeomorphism onto its image, η is monotonic, and R is yet to be defined. The composition $f = \eta \circ \mu$ collapses the leaves of each saddle region in L to a point. Therefore, R should be obtained from $\widehat{\Pi}^{-1}(p)$ by blowing up the f -image of a leaf in a saddle region to a closed interval. The difficulty is that different saddle regions could conceivably map to the same point under f .

- What is the best possible analytic quality of the representations we have constructed?
- A generic pseudo-Anosov map will violate the orientability constraints of Theorem 1.1. For surgery coefficients satisfying an appropriate parity condition, we expect that the methods of this paper may be extended to give an action of π_1 on \mathbb{R} by possibly orientation-reversing homeomorphisms.
- What can be said about fillings along the degeneracy slope? Experiments suggest that these are L-spaces, and hence do not carry taut foliations.

References

- [1] **T Barthelmé**, *Anosov flows in dimension 3: preliminary version*, preprint (2017) Available at http://www.crm.umontreal.ca/sms/2017/pdf/diapos/Anosov_flows_in_3_manifolds.pdf
- [2] **M Bell**, *Flipper*, computational software (2013) Available at <https://markcbell.github.io/software.html#flipper-software>
- [3] **M C Bell, N Dunfield**, *Censuses*, data accompanying the Flipper program (2024) Available at <https://flipper.readthedocs.io/en/latest/user/samples.html>
- [4] **S Boyer, C M Gordon, L Watson**, *On L-spaces and left-orderable fundamental groups*, Math. Ann. 356 (2013) 1213–1245 MR Zbl
- [5] **S Boyer, Y Hu**, *Taut foliations in branched cyclic covers and left-orderable groups*, Trans. Amer. Math. Soc. 372 (2019) 7921–7957 MR Zbl

- [6] **S Boyer, D Rolfsen, B Wiest**, *Orderable 3-manifold groups*, Ann. Inst. Fourier (Grenoble) 55 (2005) 243–288 MR Zbl
- [7] **D Calegari**, *Problems in foliations and laminations of 3-manifolds*, from “Topology and geometry of manifolds” (G Matić, C McCrory, editors), Proc. Sympos. Pure Math. 71, Amer. Math. Soc., Providence, RI (2003) 297–335 MR Zbl
- [8] **D Calegari**, *Foliations and the geometry of 3-manifolds*, Oxford Univ. Press (2007) MR Zbl
- [9] **D Calegari, N M Dunfield**, *Laminations and groups of homeomorphisms of the circle*, Invent. Math. 152 (2003) 149–204 MR Zbl
- [10] **A Clay, T Lidman, L Watson**, *Graph manifolds, left-orderability and amalgamation*, Algebr. Geom. Topol. 13 (2013) 2347–2368 MR Zbl
- [11] **M Culler, N M Dunfield**, *Orderability and Dehn filling*, Geom. Topol. 22 (2018) 1405–1457 MR Zbl
- [12] **M Culler, N M Dunfield, M Goerner, J R Weeks**, *SnapPy, a computer program for studying the geometry and topology of 3-manifolds* Available at <http://snappy.computop.org>
- [13] **N M Dunfield**, *Floer homology, group orderability, and taut foliations of hyperbolic 3-manifolds*, Geom. Topol. 24 (2020) 2075–2125 MR Zbl
- [14] **D Eisenbud, U Hirsch, W Neumann**, *Transverse foliations of Seifert bundles and self-homeomorphism of the circle*, Comment. Math. Helv. 56 (1981) 638–660 MR Zbl
- [15] **S R Fenley**, *The structure of branching in Anosov flows of 3-manifolds*, Comment. Math. Helv. 73 (1998) 259–297 MR Zbl
- [16] **D Fried**, *Transitive Anosov flows and pseudo-Anosov maps*, Topology 22 (1983) 299–303 MR Zbl
- [17] **D Gabai**, *Foliations and the topology of 3-manifolds*, J. Differential Geom. 18 (1983) 445–503 MR Zbl
- [18] **C D Hodgson, J R Weeks**, *Symmetries, isometries and length spectra of closed hyperbolic three-manifolds*, Exp. Math. 3 (1994) 261–274 MR Zbl
- [19] **Y Hu**, *Euler class of taut foliations and Dehn filling* (2019) arXiv 1912.01645 To appear in Comm. Anal. Geom.
- [20] **P B Kronheimer, T S Mrowka**, *Scalar curvature and the Thurston norm*, Math. Res. Lett. 4 (1997) 931–937 MR Zbl
- [21] **P Kronheimer, T Mrowka**, *Monopoles and three-manifolds*, New Math. Monogr. 10, Cambridge Univ. Press (2007) MR Zbl
- [22] **S P Novikov**, *The topology of foliations*, Tr. Mosk. Mat. Obs. 14 (1965) 248–278 MR Zbl In Russian.
- [23] **P Ozsváth, Z Szabó**, *Holomorphic disks and genus bounds*, Geom. Topol. 8 (2004) 311–334 MR Zbl
- [24] **R C Penner, J L Harer**, *Combinatorics of train tracks*, Ann. of Math. Stud. 125, Princeton Univ. Press (1992) MR Zbl
- [25] **R Roussarie**, *Plongements dans les variétés feuilletées et classification de feuilletages sans holonomie*, Inst. Hautes Études Sci. Publ. Math. 43 (1974) 101–141 MR Zbl
- [26] **W P Thurston**, *A norm for the homology of 3-manifolds*, from “Two papers”, Mem. Amer. Math. Soc. 339, Amer. Math. Soc., Providence, RI (1986) 99–130 MR Zbl
- [27] **W P Thurston**, *Three-manifolds, foliations and circles, I*, preprint (1997) arXiv math/9712268

- [28] **A Zorich**, *Flat surfaces*, from “Frontiers in number theory, physics, and geometry, I” (P Cartier, B Julia, P Moussa, P Vanhove, editors), Springer (2006) 437–583 MR Zbl

*Department of Mathematics, Princeton University
Princeton, NJ, United States*

*Current address: Department of Mathematics, Massachusetts Institute of Technology
Cambridge, MA, United States*

jzung@mit.edu

Proposed: David Gabai

Seconded: Ciprian Manolescu, Cameron Gordon

Received: 6 July 2020

Revised: 27 June 2023

GEOMETRY & TOPOLOGY

msp.org/gt

MANAGING EDITOR

András I Stipsicz Alfréd Rényi Institute of Mathematics
stipsicz@renyi.hu

BOARD OF EDITORS

Mohammed Abouzaid	Stanford University abouzaid@stanford.edu	Mark Gross	University of Cambridge mgross@dpmms.cam.ac.uk
Dan Abramovich	Brown University dan_abramovich@brown.edu	Rob Kirby	University of California, Berkeley kirby@math.berkeley.edu
Ian Agol	University of California, Berkeley ianagol@math.berkeley.edu	Bruce Kleiner	NYU, Courant Institute bkleiner@cims.nyu.edu
Arend Bayer	University of Edinburgh arend.bayer@ed.ac.uk	Sándor Kovács	University of Washington skovacs@uw.edu
Mark Behrens	University of Notre Dame mbehren1@nd.edu	Urs Lang	ETH Zürich urs.lang@math.ethz.ch
Mladen Bestvina	University of Utah bestvina@math.utah.edu	Marc Levine	Universität Duisburg-Essen marc.levine@uni-due.de
Martin R Bridson	University of Oxford bridson@maths.ox.ac.uk	Ciprian Manolescu	University of California, Los Angeles cm@math.ucla.edu
Jim Bryan	University of British Columbia jbryan@math.ubc.ca	Haynes Miller	Massachusetts Institute of Technology hrm@math.mit.edu
Dmitri Burago	Pennsylvania State University burago@math.psu.edu	Tomasz Mrowka	Massachusetts Institute of Technology mrowka@math.mit.edu
Tobias H Colding	Massachusetts Institute of Technology colding@math.mit.edu	Aaron Naber	Northwestern University anaber@math.northwestern.edu
Simon Donaldson	Imperial College, London s.donaldson@ic.ac.uk	Peter Ozsváth	Princeton University petero@math.princeton.edu
Yasha Eliashberg	Stanford University eliash-gt@math.stanford.edu	Leonid Polterovich	Tel Aviv University polterov@post.tau.ac.il
Benson Farb	University of Chicago farb@math.uchicago.edu	Colin Rourke	University of Warwick gt@maths.warwick.ac.uk
David M Fisher	Rice University davidfisher@rice.edu	Roman Sauer	Karlsruhe Institute of Technology roman.sauer@kit.edu
Mike Freedman	Microsoft Research michaelf@microsoft.com	Stefan Schwede	Universität Bonn schwede@math.uni-bonn.de
David Gabai	Princeton University gabai@princeton.edu	Natasa Sesum	Rutgers University natasas@math.rutgers.edu
Stavros Garoufalidis	Southern U. of Sci. and Tech., China stavros@mpim-bonn.mpg.de	Gang Tian	Massachusetts Institute of Technology tian@math.mit.edu
Cameron Gordon	University of Texas gordon@math.utexas.edu	Ulrike Tillmann	Oxford University tillmann@maths.ox.ac.uk
Jesper Grodal	University of Copenhagen jg@math.ku.dk	Nathalie Wahl	University of Copenhagen wahl@math.ku.dk
Misha Gromov	IHÉS and NYU, Courant Institute gromov@ihes.fr	Anna Wienhard	Universität Heidelberg wienhard@mathi.uni-heidelberg.de

See inside back cover or msp.org/gt for submission instructions.

The subscription price for 2024 is US \$805/year for the electronic version, and \$1135/year (+\$70, if shipping outside the US) for print and electronic. Subscriptions, requests for back issues and changes of subscriber address should be sent to MSP. Geometry & Topology is indexed by Mathematical Reviews, Zentralblatt MATH, Current Mathematical Publications and the Science Citation Index.

Geometry & Topology (ISSN 1465-3060 printed, 1364-0380 electronic) is published 9 times per year and continuously online, by Mathematical Sciences Publishers, c/o Department of Mathematics, University of California, 798 Evans Hall #3840, Berkeley, CA 94720-3840. Periodical rate postage paid at Oakland, CA 94615-9651, and additional mailing offices. POSTMASTER: send address changes to Mathematical Sciences Publishers, c/o Department of Mathematics, University of California, 798 Evans Hall #3840, Berkeley, CA 94720-3840.

GT peer review and production are managed by EditFLOW[®] from MSP.

PUBLISHED BY

 **mathematical sciences publishers**
nonprofit scientific publishing
<http://msp.org/>

© 2024 Mathematical Sciences Publishers

GEOMETRY & TOPOLOGY

Volume 28 Issue 9 (pages 3973–4381) 2024

Teichmüller curves in genus two: square-tiled surfaces and modular curves	3973
EDUARD DURYEV	
Equivariant aspects of singular instanton Floer homology	4057
ALIAKBAR DAEMI and CHRISTOPHER SCADUTO	
Taut foliations, left orders, and pseudo-Anosov mapping tori	4191
JONATHAN ZUNG	
Small-energy isotopies of loose Legendrian submanifolds	4233
LUKAS NAKAMURA	
On the high-dimensional geography problem	4257
ROBERT BURKLUND and ANDREW SENGER	
Dual structures on Coxeter and Artin groups of rank three	4295
EMANUELE DELUCCHI, GIOVANNI PAOLINI and MARIO SALVETTI	
Fixed-point-free pseudo-Anosov homeomorphisms, knot Floer homology and the cinquefoil	4337
ETHAN FARBER, BRAEDEN REINOSO and LUYA WANG	