ISSN 1364-0380

Geometry & Topology
Volume 4 (2000) 457–515
Published: 14 December 2000

The Geometry of $\mathbb{R}$–covered foliations

Danny Calegari

Department of Mathematics, Harvard University
Cambridge, MA 02138, USA

Email: `dannyc@math.harvard.edu`

Abstract

We study $\mathbb{R}$–covered foliations of 3–manifolds from the point of view of their *transverse geometry.* For an $\mathbb{R}$–covered foliation in an atoroidal 3–manifold M, we show that $\widetilde{M}$ can be partially compactified by a canonical cylinder $S^1_{\text{univ}} \times \mathbb{R}$ on which $\pi_1(M)$ acts by elements of $Homeo(S^1) \times Homeo(\mathbb{R})$, where the S^1 factor is canonically identified with the circle at infinity of each leaf of $\widetilde{\mathcal{F}}$. We construct a pair of *very full genuine laminations* $\Lambda^\pm$ transverse to each other and to $\mathcal{F}$, which bind every leaf of $\mathcal{F}$. This pair of laminations can be blown down to give a transverse regulating pseudo-Anosov flow for $\mathcal{F}$, analogous to Thurston's structure theorem for surface bundles over a circle with pseudo-Anosov monodromy.

A corollary of the existence of this structure is that the underlying manifold M is *homotopy rigid* in the sense that a self-homeomorphism homotopic to the identity is isotopic to the identity. Furthermore, the product structures at infinity are rigid under deformations of the foliation $\mathcal{F}$ through $\mathbb{R}$–covered foliations, in the sense that the representations of $\pi_1(M)$ in $Homeo((S^1_{\text{univ}})_t)$ are all conjugate for a family parameterized by t. Another corollary is that the ambient manifold has word-hyperbolic fundamental group.

Finally we speculate on connections between these results and a program to prove the geometrization conjecture for tautly foliated 3–manifolds.

AMS Classification numbers Primary: 57M50, 57R30

Secondary: 53C12

Keywords: Taut foliation, $\mathbb{R}$–covered, genuine lamination, regulating flow, pseudo-Anosov, geometrization

Proposed: David Gabai
Seconded: Dieter Kotschick, Walter Neumann

Received: 18 September 1999
Revised: 23 October 2000

1 Introduction

The success of the work of Barbot and Fenley [13] in classifying $\mathbb{R}$–covered Anosov flows on 3–manifolds, and the development by Thurston of a strategy to show that 3–manifolds admitting *uniform* $\mathbb{R}$–covered foliations are geometric suggests that the idea of studying foliations via their transverse geometry is a fruitful one. The tangential geometry of foliations can be controlled by powerful theorems of Cantwell and Conlon [1] and Candel [7] which establish that an atoroidal irreducible 3–manifold with a codimension one taut foliation can be given a metric in which the induced metrics on the leaves make every leaf locally isometric to hyperbolic space.

A foliation of a 3–manifold is $\mathbb{R}$*–covered* if the pullback foliation of the universal cover is the standard foliation of $\mathbb{R}^3$ by horizontal $\mathbb{R}^2$'s. This topological condition has geometric consequences for leaves of $\mathcal{F}$; in particular, leaves are *uniformly properly embedded* in the universal cover. This leads us to the notion of a *confined leaf*. A leaf λ in the pullback foliation of the universal cover $\widetilde{M}$ is *confined* when some δ–neighborhood of λ entirely contains other leaves.

The basic fact we prove about confined leaves is that the confinement condition is *symmetric for* $\mathbb{R}$*–covered foliations*. Using this symmetry condition, we can show that an $\mathbb{R}$–covered foliation can be blown down to a foliation which either slithers over S^1 or has no confined leaves. This leads to the following corollary:

Corollary 2.4.3 *If $\mathcal{F}$ is a nonuniform $\mathbb{R}$–covered foliation then after blowing down some regions we get an $\mathbb{R}$–covered foliation $\mathcal{F}'$ such that for any two intervals $I, J \subset L$, the leaf space of $\widetilde{\mathcal{F}}'$, there is an $\alpha \in \pi_1(M)$ with $\alpha(I) \subset J$.*

A more refined notion for leaves which are not confined is that of a *confined direction*, specifically a point at infinity on a leaf such that the holonomy of some transversal is bounded along every path limiting to that point.

A further refinement is a *weakly confined direction*, which is a point at infinity on a leaf such that the holonomy of some transversal is bounded along a quasi-geodesic path approaching that point. Thurston shows in [33] that the existence of nontrivial harmonic transverse measures imply that with probability one, a random walk on a leaf will have bounded holonomy for *some* transversal. For general $\mathbb{R}$–covered foliations, we show that these weakly confined directions allow one to construct a natural *cylinder at infinity* C_∞ foliated by the circles at infinity of each leaf, and prove the following structure theorem for this cylinder.

Theorem 4.6.4 *For* any $\mathbb{R}$*–covered foliation with hyperbolic leaves, not necessarily containing confined points at infinity, there are two natural maps*

$$\phi_v\colon\, C_\infty \to L, \quad \phi_h\colon\, C_\infty \to S^1_{\mathbf{univ}}$$

such that:

- ϕ_v *is the projection to the leaf space.*
- ϕ_h *is a homeomorphism for every circle at infinity.*
- *These functions give co-ordinates for* C_∞ *making it homeomorphic to a cylinder with a pair of complementary foliations in such a way that* $\pi_1(M)$ *acts by homeomorphisms on this cylinder preserving both foliations.*

In the course of the proof of this theorem, we need to treat in detail the case that there is an *invariant spine* in C_∞ — that is, a bi-infinite curve intersecting every circle at infinity exactly once, which is invariant under the action of $\pi_1(M)$. In this case, our results can be made to actually characterize the foliation $\mathcal{F}$ and the ambient manifold M, at least up to isotopy:

Theorem 4.7.2 *If* C_∞ *contains a spine* Ψ *and* $\mathcal{F}$ *is* $\mathbb{R}$*–covered but not uniform, then* M *is a Solvmanifold and* $\mathcal{F}$ *is the suspension foliation of the stable or unstable foliation of an Anosov automorphism of a torus.*

In particular, we are able to give quite a detailed picture of the asymptotic geometry of leaves:

Theorem 4.7.3 *Let* $\mathcal{F}$ *be an* $\mathbb{R}$*–covered taut foliation of a closed* 3*–manifold* M *with hyperbolic leaves. Then after possibly blowing down confined regions,* $\mathcal{F}$ *falls into exactly one of the following four possibilities:*

- $\mathcal{F}$ *is uniform.*
- $\mathcal{F}$ *is (isotopic to) the suspension foliation of the stable or unstable foliation of an Anosov automorphism of* T^2*, and* M *is a Solvmanifold.*
- $\mathcal{F}$ *contains no confined leaves, but contains strictly semi-confined directions.*
- $\mathcal{F}$ *contains no confined directions.*

In the last two cases we say $\mathcal{F}$ is *ruffled.*

Following an outline of Thurston in [35] we study the action of $\pi_1(M)$ on this universal circle and for M atoroidal we construct a pair of genuine laminations transverse to the foliation which describes its lack of uniform quasi-symmetry.

Say that a vector field transverse to an $\mathbb{R}$–covered foliation is *regulating* if every integral leaf of the lifted vector field in the universal cover intersects every leaf of the lifted foliation. A torus transverse to $\mathcal{F}$ is regulating if it lifts to a plane in the universal cover which intersects every leaf of the lifted foliation. With this terminology, we show:

Theorem 5.3.13 *Let $\mathcal{F}$ be an $\mathbb{R}$–covered foliation of an atoroidal manifold M. Then there are a pair $\Lambda^{\pm}$ of essential laminations in M with the following properties:*

- *The complementary regions to $\Lambda^{\pm}$ are ideal polygon bundles over S^1.*
- *Each $\Lambda^{\pm}$ is transverse to $\mathcal{F}$ and intersects $\mathcal{F}$ in geodesics.*
- *Λ^{+} and Λ^{-} are transverse to each other, and bind each leaf of $\mathcal{F}$, in the sense that in the universal cover, they decompose each leaf into a union of compact finite-sided polygons.*

If M is not *atoroidal but $\mathcal{F}$ has hyperbolic leaves, there is a regulating essential torus transverse to $\mathcal{F}$.*

Finally we show that the construction of the pair of essential laminations $\Lambda^{\pm}$ above is *rigid* in the sense that for a family of $\mathbb{R}$–covered foliations parameterized by t, the representations of $\pi_1(M)$ in $Homeo((S^1_{\text{univ}})_t)$ are all conjugate. This follows from the general fact that for an $\mathbb{R}$–covered foliation which is not uniform, any embedded $\pi_1(M)$–invariant collection of transversals at infinity is contained in the fibers of the projection $C_\infty \to S^1_{\text{univ}}$. It actually follows that the laminations $\Lambda^{\pm}$ do not depend (up to isotopy) on the underlying $\mathbb{R}$–covered foliation by means of which they were constructed, but reflect somehow some more meaningful underlying geometry of M.

Corollary 5.3.22 *Let $\mathcal{F}_t$ be a family of $\mathbb{R}$–covered foliations of an atoroidal M. Then the action of $\pi_1(M)$ on $(S^1_{\text{univ}})_t$ is independent of t, up to conjugacy. Moreover, the laminations $\Lambda^{\pm}_t$ do not depend on the parameter t, up to isotopy.*

This paper is foundational in nature, and can be seen as part of Thurston's general program to extend the geometrization theorem for Haken manifolds to all 3–manifolds admitting taut foliations, or more generally, essential laminations. The structures defined in this paper allow one to set up a dynamical system, analogous to the dynamical system used in Thurston's proof of geometrization for surface bundles over S^1, which we hope to use in a future paper to show that 3–manifolds admitting $\mathbb{R}$–covered foliations are geometric. Some of this picture is speculative at the time of this writing and it remains to be seen whether key results from the theory of quasi-Fuchsian surface groups — eg, Thurston's double limit theorem — can be generalized to our context. However, the rigidity result for actions on S^1_{univ} is evidence for this general conjecture. For, one expects by analogy with the geometrization theorem for surface bundles over a circle, that the sphere at infinity $S^2_\infty(\widetilde{M})$ of the universal cover $\widetilde{M}$ is obtained from the universal circle S^1_{univ} as a quotient. Since the action on this sphere at infinity is independent of the foliation, we expect the action on S^1_{univ} to be rigid too, and this is indeed the case.

It is worth mentioning that we can obtain similar results for taut foliations with one-sided branching in the universal cover in [4] and weaker but related results for arbitrary taut foliations in [5] and [6]. The best result we obtain in [6] is that for an arbitrary minimal taut foliation $\mathcal{F}$ of an atoroidal 3–manifold M, there are a pair $\Lambda^\pm$ of genuine laminations of M transverse to each other and to $\mathcal{F}$. Finally, the main results of this paper are summarized in [3].

Acknowledgements I would like to thank Andrew Casson, Sérgio Fenley and Bill Thurston for their invaluable comments, criticisms and inspiration. A cursory glance at the list of references will indicate my indebtedness to Bill for both general and specific guidance throughout this project. I would also like to thank John Stallings and Benson Farb for helping me out with some remedial group theory. In addition, I am extremely grateful to the referee for providing numerous valuable comments and suggestions, which have tremendously improved the clarity and the rigour of this paper.

I would also like to point out that I had some very useful conversations with Sérgio after part of this work was completed. Working independently, he went on to find proofs of many of the results in the last section of this paper, by somewhat different methods. In particular, he found a construction of the laminations $\Lambda^\pm$ by using the theory of earthquakes as developed by Thurston.

1.1 Notation

Throughout this paper, M will always denote a closed orientable 3–manifold, $\widetilde{M}$ its universal cover, $\mathcal{F}$ a codimension 1 co–orientable $\mathbb{R}$–covered foliation and $\widetilde{\mathcal{F}}$ its pullback foliation to the universal cover. M will be atoroidal unless we explicitly say otherwise. L will always denote the leaf space of $\widetilde{\mathcal{F}}$, which is homeomorphic to $\mathbb{R}$. We will frequently confuse $\pi_1(M)$ with its image in $Homeo(L) = Homeo(\mathbb{R})$ under the holonomy representation. We denote by $\phi_v\colon\ \widetilde{M} \to L$ the canonical projection to the leaf space of $\widetilde{\mathcal{F}}$.

2 Confined leaves

2.1 Uniform foliations and slitherings

The basic objects of study throughout this paper will be *taut $\mathbb{R}$–covered foliations of 3–manifolds.*

Definition 2.1.1 A *taut* foliation $\mathcal{F}$ of a 3–manifold is a foliation by surfaces with the property that there is a circle in the 3–manifold, transverse to $\mathcal{F}$, which intersects every leaf of $\mathcal{F}$. On an atoroidal 3–manifold, taut is equivalent to the condition of having no torus leaves.

Definition 2.1.2 Let $\mathcal{F}$ be a taut foliation of a 3–manifold M. Let $\widetilde{\mathcal{F}}$ denote the foliation of the universal cover $\widetilde{M}$ induced by pullback. $\mathcal{F}$ is $\mathbb{R}$*–covered* iff $\widetilde{\mathcal{F}}$ is the standard foliation of $\mathbb{R}^3$ by horizontal $\mathbb{R}^2$'s.

In what follows, we assume that all foliations are oriented and co-oriented. Note that this is not a significant restriction, since we can always achieve this condition by passing to a double cover. Moreover, the results that we prove are all preserved under finite covers. This co-orientation induces an invariant orientation and hence a total ordering on L. For λ, μ leaves of L, we denote this ordering by $\lambda > \mu$.

The following theorem is found in [7]:

Theorem 2.1.3 (Candel) *Let Λ be a lamination of a compact space M with 2–dimensional Riemann surface leaves. Suppose that every invariant transverse measure supported on Λ has negative Euler characteristic. Then there is a metric on M such that the inherited path metric makes the leaves of Λ into Riemann surfaces of constant curvature -1.*

Remark 2.1.4 The necessary smoothness assumption to apply Candel's theorem is that our foliations be *leafwise smooth* — ie, that the individual leaves have a smooth structure, and that this smooth structure vary continuously in the transverse direction. One expects that any co-dimension one foliation of a 3–manifold can be made to satisfy this condition, and we will assume that our foliations satisfy this condition without comment throughout the sequel.

By analogy with the usual Gauss–Bonnet formula, the Euler characteristic of an invariant transverse measure can be defined as follows: for a foliation of

M by Riemann surfaces, there is a leafwise 2-form which is just the curvature form. The product of this with a transverse measure can be integrated over M to give a real number — the Euler characteristic (see [7] and [9] for details).

For M an aspherical and atoroidal 3–manifold, every invariant transverse measure on a taut foliation $\mathcal{F}$ has negative Euler characteristic.

Consequently we may assume in the sequel that we have chosen a metric on M for which every leaf of $\mathcal{F}$ has constant curvature -1.

The following definitions are from [32].

Definition 2.1.5 A taut foliation $\mathcal{F}$ of M is *uniform* if any two leaves λ, μ of $\widetilde{\mathcal{F}}$ are contained in bounded neighborhoods of each other.

Definition 2.1.6 A manifold M *slithers over* S^1 if there is a fibration $\phi\colon \widetilde{M} \to S^1$ such that $\pi_1(M)$ acts on this fibration by bundle maps.

A slithering induces a foliation of $\widetilde{M}$ by the connected components of preimages of points in S^1 under the slithering map, and when $\widetilde{M} = \mathbb{R}^3$ and the leaves of the components of these preimages are planes, this foliation descends to an $\mathbb{R}$–covered foliation of M.

By compactness of M and S^1, it is clear that the leaves of $\widetilde{F}$ stay within bounded neighborhoods of each other for a foliation obtained from a slithering. That is, such a foliation is uniform. Thurston proves the following theorem in [32]:

Theorem 2.1.7 *Let $\mathcal{F}$ be a uniform foliation. Then after possibly blowing down some pockets of leaves, $\mathcal{F}$ comes from a slithering of M over S^1, and the holonomy representation in $Homeo(L)$ is conjugate to a subgroup of $\widetilde{Homeo(S^1)}$, the universal central extension of $Homeo(S^1)$.*

In [32], Thurston actually conjectured that for atoroidal M, every $\mathbb{R}$–covered foliation should be uniform. However, this conjecture is false and in [2] we construct many examples of $\mathbb{R}$–covered foliations of hyperbolic 3–manifolds which are not uniform.

2.2 Symmetry of the confinement condition

We make the following definition:

Definition 2.2.1 Say that a leaf λ of $\widetilde{\mathcal{F}}$ is *confined* if there exists an open neighborhood $U \subset L$, where L denotes the leaf space of $\widetilde{\mathcal{F}}$, such that

$$\bigcup_{\mu \in U} \mu \subset N_\delta(\lambda)$$

for some $\delta > 0$, where $N_\delta(\lambda)$ denotes the δ–neighborhood of λ in $\widetilde{M}$.

Say a leaf λ is *semi-confined* if there is a half-open interval $O \subset L$ with closed endpoint λ such that

$$\bigcup_{\mu \in O} \mu \subset N_\delta(\lambda)$$

for some $\delta > 0$.

Clearly, this definition is independent of the choice of metric on M with respect to which these neighborhoods are defined.

Observe that we can make the definition of a confined leaf for any taut foliation, not just for $\mathbb{R}$–covered foliations. However, in the presence of branching, the neighborhood U of a leaf $\lambda \in L$ will often not be homeomorphic to an interval.

Lemma 2.2.2 *Leaves of $\widetilde{\mathcal{F}}$ are* uniformly proper*; that is, there is a function $f\colon (0,\infty) \to (0,\infty)$ where $f(t) \to \infty$ as $t \to \infty$ such that for each leaf λ of L, any two points p, q which are a distance t apart in $\widetilde{M}$ are at most a distance $f(t)$ apart in λ.*

Proof Suppose to the contrary that we have a sequence of points p_i, q_i at distance t apart in $\widetilde{M}$ which are contained in leaves λ_i where the leafwise distances between p_i and q_i goes to ∞. After translating by some elements α_i of $\pi_1(M)$, we can assume that some subsequence of p_i, q_i converge to p, q in $\widetilde{M}$ which are distance t apart. Since the leaf space L is $\mathbb{R}$, and in particular is Hausdorff, p and q must lie on the same leaf λ, and their leafwise distance is $t < \infty$. It follows that the limit of the leafwise distances between p_i and q_i is t, and therefore they are bounded, contrary to assumption. □

Lemma 2.2.3 *If $\mathcal{F}$ is $\mathbb{R}$–covered then leaves of $\widetilde{\mathcal{F}}$ are quasi-isometrically embedded in their δ–neighborhoods in $\widetilde{M}$, for a constant depending on δ, where $N_\delta(\lambda)$ has the path metric inherits as a subspace of $\widetilde{M}$.*

Proof Let $r\colon N_\delta(\lambda) \to \lambda$ be a (non-continuous) retraction which moves each point to one of the points in λ closest to it. Then if $p, q \in N_\delta(\lambda)$ are distance 1 apart, $r(p)$ and $r(q)$ are distance at most $2\delta+1$ apart in $N_\delta(\lambda)$, and therefore there is a t such that they are at most distance t apart in λ, by lemma 2.2.2. Since $N_\delta(\lambda)$ is a path metric space, any two points p,q can be joined by a sequence of arcs of length 1 whose union has length which differs from $d(p,q)$ by some uniformly bounded amount. It follows that the distance in λ between $r(p)$ and $r(q)$ is at most $td(p,q)+\text{constant}$. □

Theorem 2.2.4 *For μ,λ leaves in $\widetilde{\mathcal{F}}$ there exists a δ such that $\mu\subset N_\delta(\lambda)$ iff there exists a δ' such that $\lambda\subset N_{\delta'}(\mu)$.*

Proof Let $d(p,q)$ denote the distance in $\widetilde{M}$ between points p,q.

For a point $p\in\widetilde{M}$ let λ_p denote the leaf in $\widetilde{\mathcal{F}}$ passing through p. We assume that δ as in the theorem has been already fixed. Let $B(p)$ denote the ball of radius δ around p in λ_p. For each leaf λ', let $C_{\lambda'}(p)$ denote the convex hull in λ' of the set of points at distance $\leq\delta$ in $\widetilde{M}$ from some $q\in B(p)$. Let

$$d(p)=\sup_{q\in C_{\lambda'}(p)} d(q,p)$$

as λ' ranges over all leaves in L such that $C_{\lambda'}(p)$ is nonempty. Let

$$s(p)=\sup_{C_{\lambda'}(p)} \operatorname{diam}(C_{\lambda'}(p)).$$

Then $d(p)$ and $s(p)$ are well-defined and finite for every p. For, if m_i,n_i are a pair of points on a leaf λ_i at distance δ_i from p converging to m,n at distance δ from p, then the hypothesis that our foliation is $\mathbb{R}$–covered implies that m,n are on the same leaf, and the leafwise distances between m_i and n_i converge to the leafwise distance between m and n.

More explicitly, we can take a homeomorphism from $B\subset\widetilde{M}$ to some region of $\mathbb{R}^3$ and consider for each leaf in the image, the convex hull of its intersection with B. Since B is contained in a compact region of $\mathbb{R}^3$, there is a continuous family of isometries of the leaves in question to $\mathbb{H}^2$ such that the intersections with B form a compact family of compact subsets of $\mathbb{H}^2$. It follows that their convex hulls form a compact family of compact subsets of $\mathbb{H}^2$ and hence their diameters are uniformly bounded.

It is clear from the construction that $d(p)$ and $s(p)$ are upper semi-continuous. Moreover, their values depend only on $\pi(p)\in M$ where $\pi\colon\widetilde{M}\to M$ is the covering projection. Hence they are uniformly bounded by two numbers which we denote d and s.

In particular, the set $C \subset \lambda$ defined by

$$C = \bigcup_{p \in \mu} C_\lambda(p)$$

is contained in $N_d(\mu)$. The hypothesis that $\mu \subset N_\delta(\lambda)$ implies that $C(p)$ is nonempty for any p. In fact, for some collection p_i of points in μ,

$$\bigcap_i B(p_i) \neq \emptyset \implies \bigcap_i C_\lambda(p_i) \neq \emptyset.$$

Moreover, the boundedness of s implies that for p, q sufficiently far apart in μ, $C_\lambda(p) \cap C_\lambda(q) = \emptyset$. For, the condition that $C_\lambda(p) \cap C_\lambda(q) \neq \emptyset$ implies that $d(p,q) \leq 2s + 2d$ in $\widetilde{M}$. By lemma 2.2.2, there is a uniform bound on the distance between p and q in μ.

Hence there is a map from the nerve of a locally finite covering by $B(p_i)$ of μ for some collection of points p_i to the nerve of a locally finite covering of some subset of C by $C_\lambda(p_i)$. We claim that this subset, and hence C, is a net in λ.

Observe that the map taking p to the center of $C_\lambda(p)$ is a coarse quasi-isometry from μ to C with its path metric. For, since the diameter of $C_\lambda(p)$ is uniformly bounded independently of p, and since a connected chain of small disks in μ corresponds to a connected chain of small disks in C, the map cannot expand distances too much. Conversely, since C is contained in the ϵ–neighborhood of μ, paths in C can be approximated by paths in μ of the same length, up to a bounded factor.

It follows by a theorem of Farb and Schwartz in [11] that the map from μ to λ sending p to the center of $C_\lambda(p)$ is coarsely onto, as promised.

But now every point in λ is within a uniformly bounded distance from C, and therefore from μ, so that there exists a δ' with $\lambda \subset N_{\delta'}(\mu)$. □

Remark 2.2.5 Notice that this theorem depends vitally upon lemma 2.2.2. In particular, taut foliations which are not $\mathbb{R}$–covered *do not* lift to foliations with uniformly properly embedded leaves. For, one knows by a theorem of Palmeira (see [28]) that a taut foliation fails to be $\mathbb{R}$–covered exactly when the space of leaves of $\widetilde{\mathcal{F}}$ is not Hausdorff. In this case there are a sequence of leaves λ_i of $\widetilde{\mathcal{F}}$ limiting to a pair of distinct leaves λ, μ. One can thus find a pair of points $p \in \lambda, q \in \mu$ and a sequence of pairs of points $p_i, q_i \in \lambda_i$ with $p_i \to p$ and $q_i \to q$ so that the leafwise distance between p_i and q_i goes to infinity, whereas the distance between them in $\widetilde{M}$ is uniformly bounded; ie, leaves are not uniformly properly embedded.

Theorem 2.2.6 *If every leaf λ of $\widetilde{\mathcal{F}}$ is confined, then $\mathcal{F}$ is uniform.*

Proof Since any two points in the leaf space are joined by a finite chain of open intervals of confinement, the previous lemma shows that the corresponding leaves are both contained in bounded neighborhoods of each other. This establishes the theorem. $\square$

2.3 Action on the leaf space

Lemma 2.3.1 *For $\mathcal{F}$ an $\mathbb{R}$–covered foliation of M, and $L \cong \mathbb{R}$ the leaf space of $\widetilde{\mathcal{F}}$, for any leaf $\lambda \in L$ the image of λ under $\pi_1(M)$ goes off to infinity in either direction.*

Proof Recall that we assume that $\mathcal{F}$ is co-oriented, so that, every element of $\pi_1(M)$ acts by an orientation-preserving homeomorphism of the leaf space L.

Suppose there is some λ whose images under $\pi_1(M)$ are bounded in some direction, say without loss of generality, the "positive" direction. Then the least upper bound λ' of the leaves $\alpha(\lambda)$ is fixed by every element of $\pi_1(M)$. Since $\mathcal{F}$ is taut, $\lambda' = \mathbb{R}^2$ and therefore $\lambda'/\pi_1(M)$ is a $K(\pi_1(M), 1)$ and is therefore homotopy equivalent to M. This is absurd since M is 3–dimensional. $\square$

We remark that for foliations which are not taut, but for which the leaf space of $\widetilde{\mathcal{F}}$ is homeomorphic to $\mathbb{R}$, this lemma need not hold. For example, the foliation of $\mathbb{R}^3 - \{0\}$ by horizontal planes descends to a foliation on $S^2 \times S^1$ by the quotient $q \to 2q$. In fact, *no* leaf goes off to infinity in both directions under the action of $\pi_1(M) = \mathbb{Z}$ on the leaf space $\mathbb{R}$, since the single annulus leaf in $\widetilde{\mathcal{F}}$ is invariant under the whole group.

Lemma 2.3.2 *For all $r > 0$ there is an $s > 0$ such that every $N_s(p) - \lambda_p$ contains a ball of radius r on either side of the leaf, for λ_p the leaf in $\widetilde{M}$ through p.*

Proof Suppose for some r that the side of $\widetilde{M}$ above λ_p contains no ball of radius r. Then every leaf above λ_p, and therefore every leaf, is confined. It follows that $\mathcal{F}$ is uniform. But in a uniform foliation, there are pairs of leaves in L which never come closer than t to each other, for any t. This gives a contradiction.

Once we know that every leaf has some ball centered at any point, the compactness of M implies that we can find an s which works for balls centered at any point. $\square$

Theorem 2.3.3 *For any leaf μ in $\widetilde{\mathcal{F}}$ and any side of μ (which may as well be the positive side), one of the following mutually exclusive conditions is true:*

(1) *μ is semi-confined on the positive side.*

(2) *For any $\lambda > \mu$ and any leaf $\mu' > \mu$, there is an $\alpha \in \pi_1(M)$ such that $\alpha([\mu, \lambda]) \subset (\mu, \mu')$.*

Remark 2.3.4 To see that the two conditions are mutually exclusive, observe that if they both hold then every leaf on one side of μ can be mapped into the semi-confined interval in L, and therefore every leaf on that side of μ is confined. Since translates of μ go off to infinity in either direction, every leaf is confined and the foliation is uniform. Since such foliations slither over S^1 (after possibly being blown down), the leaf space cannot be arbitrarily compressed by the action of $\pi_1(M)$. In particular, leaves in the same fiber of the slithering over S^1 and differing by n periods, say, cannot be translated by any α to lie between leaves in the same fiber which differ by m periods for $m < n$.

Proof If λ is in the δ–neighborhood of μ, μ is semi-confined and we are done. So suppose λ is not in the δ–neighborhood of μ for any δ.

By hypothesis therefore, μ' is not in the δ–neighborhood of μ, and conversely μ is not in the δ–neighborhood of μ', for any δ.

Let $p \in \mu, q \in \lambda$ be two points. Then $d(p,q) = t$. For $r = t + \mathrm{diam}(M)$ we know that there is a s such that any ball of radius s about a point p contains a ball of radius r on either side of λ_p. Pick a point $p' \in \mu$ which is distance at least s from μ'. Then there is a ball B of radius r between μ and μ' in the ball of radius s about p. It follows that there is an α such that $\alpha(p)$ and $\alpha(q)$ are both in B. This α has the properties we want. □

2.4 Blowing down leaves

Definition 2.4.1 For λ a confined leaf, the *umbra* of λ, denoted $\mathcal{U}(\lambda)$, is the subset of L consisting of leaves μ such that μ is contained in a bounded neighborhood of λ.

Notice that if $\mu \in \mathcal{U}(\lambda)$ then $\mathcal{U}(\mu) = \mathcal{U}(\lambda)$. Moreover, $\mathcal{U}(\lambda)$ is closed for any λ. To see this, let μ be a hypothetical leaf in $\overline{\mathcal{U}(\lambda)} - \mathcal{U}(\lambda)$. If μ is semi-confined on the side containing λ, then $\mathcal{U}(\mu) \cap \mathcal{U}(\lambda)$ is nonempty, and therefore $\mathcal{U}(\mu) = \mathcal{U}(\lambda)$ so that certainly $\mu \in \mathcal{U}(\lambda)$. Otherwise, μ is not semi-confined on

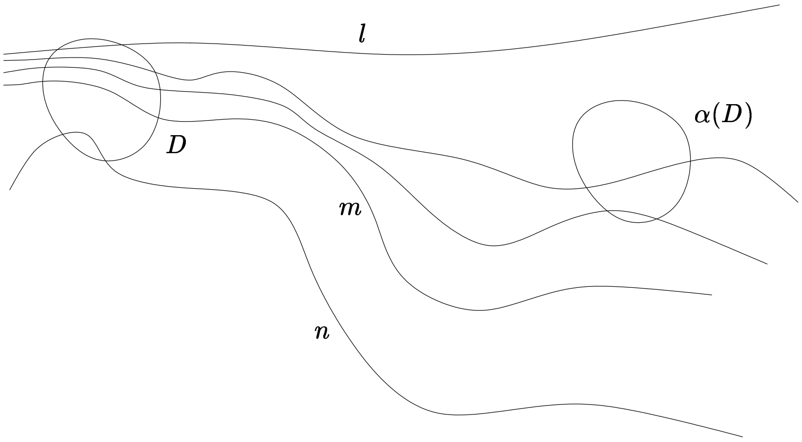

Figure 1: If l is not semi-confined, for any nearby leaf m and any other leaf n, there is an element $\alpha \in \pi_1(M)$ such that $\alpha(l)$ and $\alpha(n)$ are between l and m.

that side and theorem 2.3.3 implies that there is an α taking $[\lambda, \mu]$ inside $\mathcal{U}(\lambda)$. But then $\mathcal{U}(\lambda) = \mathcal{U}(\alpha(\lambda))$, so that $\mathcal{U}(\lambda) = \alpha^{-1}(\mathcal{U}(\lambda))$ and $\mu \in \mathcal{U}(\lambda)$ after all.

In fact, if $\alpha(\mathcal{U}(\lambda)) \cap \mathcal{U}(\lambda) \neq \emptyset$ for some $\alpha \in \pi_1(M)$ then $\alpha(\mathcal{U}(\lambda)) = \mathcal{U}(\lambda)$, and in particular, α must fix every leaf in $\partial\mathcal{U}(\lambda)$. Hence the set of elements in $\pi_1(M)$ which do not translate $\mathcal{U}(\lambda)$ off itself is a group.

We show in the following theorem that for an $\mathbb{R}$–covered foliation which is not uniform, the confined leaves do not carry any of the essential topology of the foliation.

Theorem 2.4.2 *Suppose M has an $\mathbb{R}$–covered but not uniform foliation $\mathcal{F}$. Then M admits another $\mathbb{R}$–covered foliation $\mathcal{F}'$ with no confined leaves such that $\mathcal{F}$ is obtained from $\mathcal{F}'$ by blowing up some leaves and then possibly perturbing the blown up regions.*

Proof Fix some confined leaf λ, and let G_λ denote the subgroup of $\pi_1(M)$ which fixes $\mathcal{U}(\lambda)$. The assumption that $\mathcal{F}$ is not uniform implies that some leaves are not confined, and therefore $\mathcal{U}(\lambda)$ is a compact interval. Then G_λ acts properly discontinuously on the topological space $\mathbb{R}^2 \times I$, and we claim that this action is conjugate to an action which preserves each horizontal $\mathbb{R}^2$.

This will be obvious if we can show that the action of G_λ on the top and bottom leaves λ^u and λ^l are conjugate.

Observe that λ^u and λ^l are contained in bounded neighborhoods of each other, and therefore by lemma 2.2.3 any choice of nearest point map between λ^u and λ^l is a coarse quasi-isometry. Moreover, such a map can be chosen to be G_λ–equivariant. This map gives an exact conjugacy between the actions of G_λ on their ideal boundaries $S^1_\infty(\lambda^u)$ and $S^1_\infty(\lambda^l)$. Since each of λ^u, λ^l is isometric to $\mathbb{H}^2$ and the actions are by isometries, it follows that G_λ is a torsion-free Fuchsian group.

Since every $\mu \in \mathcal{U}(\lambda)$ in isometric to $\mathbb{H}^2$, and since every choice of closest-point map from μ to λ^u is a quasi-isometry, we can identify each $S^1_\infty(\mu)$ canonically and G_λ–equivariantly with $S^1_\infty(\lambda^u)$.

Let $F = \lambda^u/G_\lambda$ be the quotient surface. Then we can find an ideal triangulation of the convex hull of F and for each boundary component of the convex hull, triangulate the complementary cylinder with ideal triangles in some fixed way. This triangulation lifts to an ideal triangulation of λ^u. Identifying $S^1_\infty(\lambda^u)$ canonically with $S^1_\infty(\mu)$ for each μ, we can transport this ideal triangulation to an ideal triangulation of each μ. The edges of the triangulation sweep out infinite strips $I \times \mathbb{R}$ transverse to $\widetilde{\mathcal{F}}$ and decompose the slab of leaves corresponding to $\mathcal{U}(\lambda)$ into a union of ideal triangle $\times\, I$. Since G_λ acts on these blocks by permutation, we can replace the foliation $\widetilde{\mathcal{F}}$ of the slab with a foliation on which G_λ acts trivially.

We can transport this action on the total space of $\mathcal{U}(\lambda)$ to actions on the total space of $\mathcal{U}(\alpha(\lambda))$ wherever it is different. Range over all equivalence classes under $\pi_1(M)$ of all such $\mathcal{U}(\lambda)$, modifying the action as described.

Now the construction implies that $\pi_1(M)$ acts on $L/\sim$ where $\mu \sim \lambda$ if $\mu \in \mathcal{U}(\lambda)$. It is straightforward to check that $L/\sim\,\cong \mathbb{R}$. Moreover, the total space of each $\mathcal{U}(\lambda)$ can be collapsed by collapsing each ideal triangle $\times\, I$ to an ideal triangle. The quotient gives a new $\mathbb{R}^3$ foliated by horizontal $\mathbb{R}^2$'s on which $\pi_1(M)$ still acts properly discontinuously. The quotient $\hat{M} = (\mathbb{R}^3/\sim)/\pi_1(M)$ is actually homeomorphic to M by the following construction: consider a covering of $\hat{M}$ by convex open balls, and lift this to an equivariant covering of $\mathbb{R}^3/\sim$. This pulls back under the quotient map to an equivariant covering of $\mathbb{R}^3$ by convex balls, which project to give a covering of M by convex balls. By construction, the coverings are combinatorially equivalent, so M is homeomorphic to $\hat{M}$.

By construction, every leaf is a limit under $\pi_1(M)$ of every other leaf, so by theorem 2.3.3, no leaf is confined with respect to any metric on M. The induced foliation on M is $\mathcal{F}'$, and the construction shows that $\mathcal{F}$ can be obtained from $\mathcal{F}'$ as required in the statement of the theorem. □

Corollary 2.4.3 *If $\mathcal{F}$ is a nonuniform $\mathbb{R}$–covered foliation then after blowing down some regions we get an $\mathbb{R}$–covered foliation $\mathcal{F}'$ such that for any two intervals $I, J \subset L$, the leaf space of $\widetilde{\mathcal{F}}'$, there is an $\alpha \in \pi_1(M)$ with $\alpha(I) \subset J$.*

In the sequel we will assume that all our $\mathbb{R}$–covered foliations have no confined leaves; ie, they satisfy the hypothesis of the preceding corollary.

3 The cylinder at infinity

3.1 Constructing a topology at infinity

Each leaf λ of $\widetilde{\mathcal{F}}$ is isometric to $\mathbb{H}^2$, and therefore has an ideal boundary $S^1_\infty(\lambda)$. We define a natural topology on $\bigcup_{\lambda \in L} S^1_\infty(\lambda)$ with respect to which it is homeomorphic to a cylinder. Once we have defined this topology and verified that it makes this union into a cylinder, we will refer to this cylinder as *the cylinder at infinity of* $\widetilde{\mathcal{F}}$ and denote it by C_∞.

Let $UT\widetilde{\mathcal{F}}$ denote the unit tangent bundle to $\widetilde{\mathcal{F}}$. This is a circle bundle over $\widetilde{M}$ which lifts the circle bundle $UT\mathcal{F}$ over M. Let τ be a small transversal to $\widetilde{\mathcal{F}}$ and consider the cylinder C which is the restriction $UT\widetilde{\mathcal{F}}|_\tau$. There is a canonical map

$$\pi_\tau \colon C \to \bigcup_{\lambda \in L} S^1_\infty(\lambda)$$

defined as follows. For $v \in UT_x\mathcal{F}$ where $x \in \lambda$, there is a unique infinite geodesic ray γ_v in λ starting at x and pointing in the direction v. This ray determines a unique point $\pi_\tau(v) \in S^1_\infty(\lambda)$. The restriction of π_τ to $UT_x\mathcal{F}$ for any $x \in \tau$ is obviously a homeomorphism. We define the topology on $\bigcup_{\lambda \in L} S^1_\infty(\lambda)$ by requiring that π_τ be a homeomorphism, for each τ.

Lemma 3.1.1 *The topology on $\bigcup_{\lambda \in L} S^1_\infty$ defined by the maps π_τ is well-defined. With respect to this topology, this union of circles is homeomorphic to a cylinder C_∞.*

Proof All that needs to be checked is that for two transversals τ, σ with $\phi_v(\tau) = \phi_v(\sigma)$, the map $\pi_\sigma^{-1}\pi_\tau \colon UT\mathcal{F}|_\tau \to UT\mathcal{F}|_\sigma$ is a homeomorphism. For ease of notation, we refer to the two circle bundles as C_τ and C_σ and $\pi_\sigma^{-1}\pi_\tau$ as f. Then each of C_τ and C_σ is foliated by circles, and furthermore f is a homeomorphism when restricted to any of these circles. For a given leaf λ

intersecting τ and σ at t and s respectively, f takes a geodesic ray through t to the unique geodesic ray through s asymptotic to it.

It suffices to show that if v_i, w_i are two sequences in C_τ, C_σ with $v_i \to v$ and $w_i \to w$ with $w_i = f(v_i)$ that $w = f(v)$. The Riemannian metrics on leaves of $\widetilde{\mathcal{F}}$ vary continuously as one moves from leaf to leaf, with respect to some local product structure. It follows that the γ_{v_i} converge geometrically on compact subsets of $\widetilde{M}$ to γ_v. Furthermore, the γ_{w_i} are asymptotic to the γ_{v_i} so that they converge geometrically to a ray asymptotic to γ_v. This limiting ray is a limit of geodesics and must therefore be geodesic and hence equal to γ_w. □

The group $\pi_1(M)$ obviously acts on C_∞ by homeomorphisms. It carries a canonical foliation by circles which we refer to as the *horizontal foliation.*

3.2 Weakly confined directions

Definition 3.2.1 A point $p \in S^1_\infty(\lambda)$ for some λ is *weakly confined* if there is an interval $[\lambda^-, \lambda^+] \subset L$ containing λ in its interior and a map

$$H\colon\ [\lambda^-, \lambda^+] \times \mathbb{R}^+ \to \widetilde{M}$$

such that:

- For each $\mu \in [\lambda^-, \lambda^+]$, H maps $\mu \times \mathbb{R}^+$ to a parameterized quasigeodesic in μ.
- The quasigeodesic $H(\lambda \times \mathbb{R}^+)$ limits to $p \in S^1_\infty(\lambda)$.
- The transverse arcs $[\lambda^-, \lambda^+] \times t$ have length bounded by some constant C independent of t.

It follows from the definition that if p is weakly confined, the quasigeodesic rays $H(\mu \times \mathbb{R}^+)$ limit to unique points $p_\mu \in S^1_\infty(\mu)$ which are themselves weakly confined, and the map $\mu \to p_\mu$ is a continuous map from $[\lambda^-, \lambda^+]$ to C_∞ which is transverse to the horizontal foliation. If p is a weakly confined direction, let $\tau_p \subset C_\infty$ be a maximal transversal through p constructed by this method. Then we call τ_p a *weakly confined transversal*, and we denote the collection of all such weakly confined transversals by $\mathcal{T}$. Such transversals need not be either open or closed, and may project to an unbounded subset of L.

Lemma 3.2.2 *There exists some weakly confined transversal running between any two horizontal leaves in C_∞. Moreover, the set $\mathcal{T}$ consists of a $\pi_1(M)$–equivariant collection of embedded, mutually non-intersecting arcs.*

Proof If $\mathcal{F}$ is uniform, any two leaves of $\widetilde{\mathcal{F}}$ are a bounded distance apart, so there are uniform quasi-isometries between any two leaves which move points a bounded distance. In this case, *every* point at infinity is weakly confined.

If $\mathcal{F}$ is not uniform and is minimal, for any λ, λ' leaves of $\widetilde{\mathcal{F}}$ choose some transversal τ between λ and λ'. Then there is an $\alpha \in \pi_1(M)$ such that $\phi_v(\tau)$ is properly contained in $\alpha(\phi_v(\tau))$. It follows that we can find a square $S\colon I \times I \to \widetilde{M}$ such that $S(I,0) = \tau$, $S(I,1) \subset \alpha(\tau)$ and each $S(t,I)$ is contained in some leaf. The union of squares $S \cup \alpha(S) \cup \alpha^2(S) \cup \dots$ contains the image of an infinite strip $I \times \mathbb{R}^+$ where the $I \times t$ factors have a uniformly bounded diameter.

The square S descends to an immersed, foliated mapping torus in M which is topologically a cylinder. Let γ be the core of the cylinder. Then γ is homotopically essential, so it lifts to a quasigeodesic in $\widetilde{M}$. Since the strip $I \times \mathbb{R}^+$ stays near the lift of this core, it is quasigeodesically embedded in $\widetilde{M}$, and therefore its intersections with leaves of $\widetilde{\mathcal{F}}$ are quasigeodesically embedded in those leaves. It limits therefore to a weakly confined transversal in C_∞.

To see that weakly confined transversals do not intersect, suppose α, β are two weakly confined transversals that intersect at $p \in S^1_\infty(\lambda)$. We restrict attention to a small interval I in L which is in the intersection of their ranges. If this intersection consists of a single point p, then actually $\alpha \cup \beta$ is a subset of a single weakly confined transversal.

Corresponding to $I \subset L$ there are two infinite quasigeodesic strips $A\colon I \times \mathbb{R}^+ \to \widetilde{M}$ and $B\colon I \times \mathbb{R}^+ \to \widetilde{M}$ guaranteed by the definition of a weakly confined transversal. Let $\mu \in I$ be such that $A(\mu \times \mathbb{R}^+)$ does not limit to the same point in $S^1_\infty(\mu)$ as $B(\mu \times \mathbb{R}^+)$. By hypothesis, $A(\lambda \times \mathbb{R}^+)$ is asymptotic to $B(\lambda \times \mathbb{R}^+)$. But the uniform thickness of the strips implies that $A(\mu \times \mathbb{R}^+)$ is a bounded distance in $\widetilde{M}$ from $A(\lambda \times \mathbb{R}^+)$ and therefore from $B(\lambda \times \mathbb{R}^+)$ and consequently $B(\mu \times \mathbb{R}^+)$. But then by lemma 2.2.2 the two rays in μ limit to the same point in $S^1_\infty(\mu)$, contrary to assumption. It follows that weakly confined transversals do not intersect. $\square$

In [33] Thurston proves the following theorem:

Theorem 3.2.3 (Thurston) *For a general taut foliation $\mathcal{F}$, a random walk γ on a leaf λ of $\widetilde{\mathcal{F}}$ converging to some $p \in C_\infty$ stays a bounded distance from some nearby leaves $\lambda^\pm$ in $\widetilde{\mathcal{F}}$, with probability 1, and moreover, also with probability 1, there is an exhaustion of γ by compact sets such that outside these sets, the distance between γ and $\lambda^\pm$ converges to 0.*

It is possible but technically more difficult to develop the theory of weakly confined directions using random walks instead of quasigeodesics as suggested in [31], and this was our inspiration.

3.3 Harmonic measures

Following [21] we define a harmonic measure for a foliation.

Definition 3.3.1 A probability measure m on a manifold M foliated by $\mathcal{F}$ is *harmonic* if for every bounded measurable function f on M which is smooth in the leaf direction,

$$\int_M \Delta_{\mathcal{F}} f \, dm = 0$$

where $\Delta_{\mathcal{F}}$ denotes the leafwise Laplacian.

Theorem 3.3.2 (Garnett) *A compact foliated Riemannian manifold $M, \mathcal{F}$ always has a nontrivial harmonic measure.*

This theorem is conceptually easy to prove: observe that the probability measures on a compact space are a convex set. The leafwise diffusion operator gives a map from this convex set to itself, which map must therefore have a fixed point. There is some analysis involved in making this more rigorous.

Using the existence of harmonic measures for foliations, we can analyze the $\pi_1(M)$–invariant subsets of C_∞.

Theorem 3.3.3 *Let U be an open $\pi_1(M)$–invariant subset of C_∞. Then either U is empty, or it is dense and omits at most one point at infinity in a set of leaves of measure 1.*

Proof Let λ be a leaf of $\widetilde{\mathcal{F}}$ such that $S^1_\infty(\lambda)$ intersects U, and consequently intersects it in some open set. Then all leaves μ sufficiently close to λ have $S^1_\infty(\mu)$ intersect U, and therefore since leaves of $\mathcal{F}$ are dense, U intersects every circle at infinity in an open set.

For a point $p \in \lambda$, define a function $\theta(p)$ to be the maximum of the visual angles at p of intervals in $S^1_\infty(\lambda) \cap U$. This function is continuous as p varies in λ, and lower semi-continuous as p varies through $\widetilde{M}$. Moreover, it only depends on the projection of p to M. It therefore attains a minimum θ_0 somewhere, which must be > 0. This implies that $U \cap S^1_\infty(\lambda)$ has *full measure* in $S^1_\infty(\lambda)$, since otherwise by taking a sequence of points $p_i \in \lambda$ approaching a point of density in the complement, we could make $\theta(p_i) \to 0$.

Similarly, the supremum of θ is 2π, since if we pick a sequence p_i converging to a point p in $U \cap S^1_\infty$, the interval containing p will take up more and more of the visual angle.

Let θ_i be the time i leafwise diffusion of θ. Then each θ_i is C^∞ on each leaf, and is measurable since θ is, by a result in [21]. Define

$$\hat{\theta} = \sum_{i=1}^{\infty} 2^{-i}\theta_i$$

Then $\hat{\theta}$ satisfies the following properties:

- $\hat{\theta}$ is a bounded measurable function on M which is C^∞ in every leaf.
- $\Delta_{\mathcal{F}}\hat{\theta} \geq 0$ for every point in every leaf, with equality holding at some point in a leaf iff $\theta = 2\pi$ identically in that leaf.

To see the second property, observe that $\Delta_{\mathcal{F}}\theta = 0$ everywhere except at points where there at least two subintervals of U of largest size. For, elsewhere θ agrees with the harmonic extension to $\mathbb{H}^2 = \lambda$ of a function whose value is 1 on a subinterval of the boundary and 0 elsewhere. In particular, elsewhere θ is harmonic. Moreover, at points where there are many largest subintervals of U, $\Delta_{\mathcal{F}}\theta$ is a positive distributional function — that is, the "subharmonicity" of θ is concentrated at these points. In particular, $\Delta_{\mathcal{F}}\hat{\theta} \geq 0$ and it is $= 0$ iff there are no points in λ where there are more than one largest visual subinterval of U. But this occurs only when U omits at most 1 point from $S^1_\infty(\lambda)$.

Now theorem 3.3.2 implies that $\Delta_{\mathcal{F}}\hat{\theta} = 0$ for the support of any harmonic measure m, and therefore that $\theta = 2\pi$ for every point in any leaf which intersects the support of m.

Garnett actually shows in [21] that any harmonic measure disintegrates locally into the product of some harmonic multiple of leafwise Riemannian measure with a transverse invariant measure on the local leaf space. When every leaf is dense, as in our situation, the transverse measure is in the Lebesgue measure class. Hence in fact we can conclude that $\theta = 2\pi$ for a.e. leaf in the Lebesgue sense. □

Note that there was no assumption in this theorem that $\mathcal{F}$ contain no confined leaves, and therefore it applies equally well to uniform foliations with every leaf dense. In fact, for some uniform foliations, there are open invariant sets at infinity which omit exactly one point from each circle at infinity.

4 Confined directions

4.1 Suspension foliations

Let $\psi\colon T^2 \to T^2$ be an Anosov automorphism. ie, in terms of a basis for $H_1(T^2)$ the map ψ is given by an element of $SL(2,\mathbb{Z})$ with trace > 2. Then ψ leaves invariant a pair of foliations of T^2 by those lines parallel to the eigenspaces of the action of ψ on $\mathbb{R}^2$. These foliations suspend to two transverse foliations of the mapping torus

$$M = T^2 \times I/(x,0) \sim (\psi(x),1)$$

which we call the *stable* and *unstable* foliation $\mathcal{F}_s$ and $\mathcal{F}_u$ of M. There is a flow of M given by the vector field tangent to the I direction in the description above, and with respect to the metric on M making it a Solv-manifold, this is an Anosov flow, and $\mathcal{F}_s$ and $\mathcal{F}_u$ are the stable and unstable foliations of this flow respectively. In particular, the leaves of the foliation $\mathcal{F}_u$ converge in the direction of the flow, and the leaves of the foliation $\mathcal{F}_s$ diverge in the direction of the flow.

Both foliations are $\mathbb{R}$–covered, being the suspension of $\mathbb{R}$–covered foliations of T^2. Moreover, no leaf of either foliation is confined. To see this, observe that integral curves of the stable and unstable directions are horocycles with respect to the hyperbolic metric on each leaf. Since each leaf is quasigeodesically (in fact, geodesically) embedded in M, it can be seen that the leaves themselves, and not just the integral curves between them, diverge in the appropriate direction.

With respect to the Solv geometric structure on M, every leaf is intrinsically isometric to $\mathbb{H}^2$. One can see that every geodesic on a leaf of $\mathcal{F}_s$ which is not an integral curve of the Anosov flow will eventually curve away from that flow to point asymptotically in the direction exactly opposite to the flow. That is to say, leaves of $\mathcal{F}_s$ converge at infinity in every direction except for the direction of the flow; similarly, leaves of $\mathcal{F}_u$ converge at infinity in every direction except for the direction opposite to the flow. These are the prototypical examples of $\mathbb{R}$–covered foliations which have no confined leaves, but which have many *confined directions* (to be defined below).

4.2 Confined directions

Recapitulating notation: throughout this section we fix a 3–manifold M, an $\mathbb{R}$–covered foliation $\mathcal{F}$ with no confined leaves, and a metric on M with respect to which each leaf of $\widetilde{\mathcal{F}}$ is isometric to $\mathbb{H}^2$. We fix $L \cong \mathbb{R}$ the leaf space of $\widetilde{\mathcal{F}}$

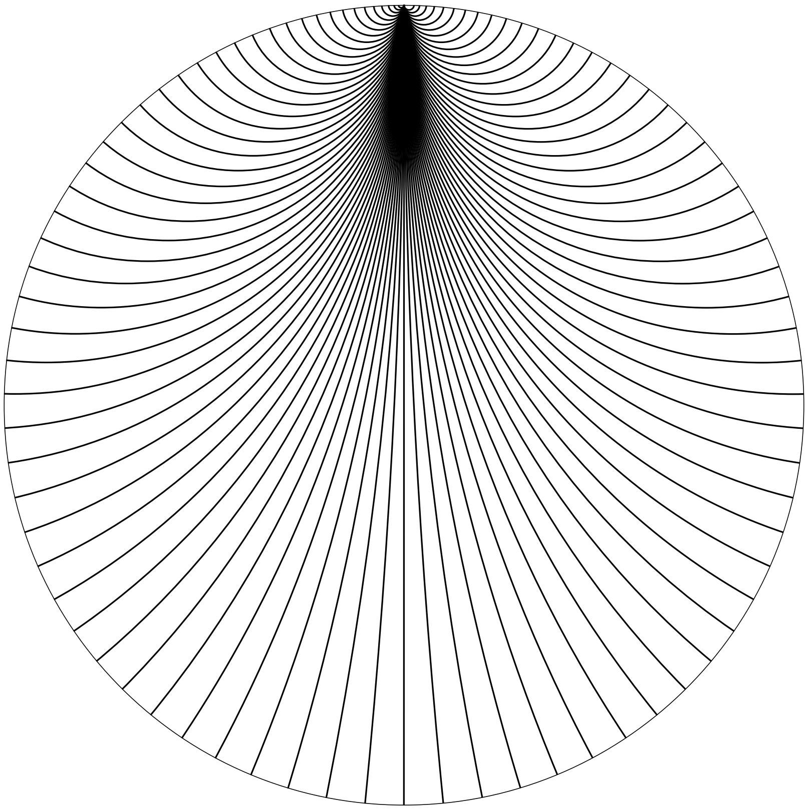

Figure 2: Each $\mathbb{H}^2$ is foliated by flow lines

and the projection $\phi_v\colon \widetilde{M} \to L$. Each leaf of $\widetilde{\mathcal{F}}$ can be compactified by the usual circle at infinity of hyperbolic space; we denote the circle at infinity of a leaf λ by $S^1_\infty(\lambda)$. We let $UT\mathcal{F}$ denote the unit tangent bundle to the foliation, and $UT\lambda$ the unit tangent bundle of each leaf λ.

Definition 4.2.1 For λ a leaf of $\widetilde{\mathcal{F}}$, we say a p a point in $S^1_\infty(\lambda)$ is a *confined point* if for *every* sequence $p_i \in \lambda$ limiting only to p, there is an interval $I \subset L$ containing λ in its interior and a sequence of transversals τ_i projecting homeomorphically to I under ϕ whose lengths are *uniformly* bounded. That is, there is some uniform t such that $\|\tau_i\| \leq t$. Equivalently, there is a neighborhood I of λ in L with endpoints $\lambda^\pm$ such that every sequence p_i as above is contained in a bounded neighborhood of both λ^+ and λ^-. If p is not confined, we say it is *unconfined.*

Remark 4.2.2 A point may certainly be unconfined and yet weakly confined.

Definition 4.2.3 For a point $p \in S^1_\infty(\lambda)$ which is unconfined, a *certificate* for p is a sequence of points $p_i \in \lambda$ limiting only to p such that for any $I \subset L$ containing λ in its interior and a sequence of transversals τ_i projecting homeomorphically to I under ϕ, the lengths $\|\tau_i\|$ are unbounded. Equivalently, there is a sequence of leaves $\lambda_i \to \lambda$ such that for any i, the sequence p_j does not stay within a bounded distance from λ_i. By definition, every unconfined point has a certificate.

For a simply connected leaf, holonomy transport is independent of the path between endpoints. The transversals τ_i defined above are obtained from τ_1 by holonomy transport.

Theorem 4.2.4 *The following conditions are equivalent:*

- *The point $p \in S^1_\infty(\lambda)$ is confined.*
- *There is a neighborhood of p in $S^1_\infty(\lambda)$ consisting of confined points.*
- *There is a neighborhood U of p in $\lambda \cup S^1_\infty(\lambda)$ such that there exists $t > 0$ and an interval $I \subset L$ containing λ in its interior such that for any properly embedded (topological) ray $\gamma\colon \mathbb{R}^+ \to \lambda$ whose image is contained in U, there is a proper map $H\colon \mathbb{R}^+ \times I \to \widetilde{M}$ such that $\phi \circ H(x,s) = s$ for all s, $H|_{\mathbb{R}^+\times\lambda} = \gamma$ and $\|H(x,I)\| \leq t$ for all x.*

Proof It is clear that the third condition implies the first. Suppose there were a sequence of unconfined points $p_i \in S^1_\infty$ converging to p. Let $p_{i,j}$ be a certificate for p_i. Then we can find integers n_i so that p_{i,n_i} is a certificate for p. It follows that the first condition implies the second. In fact, this argument shows that p is confined iff there is a neighborhood U of p in $\lambda \cup S^1_\infty(\lambda)$ and a neighborhood I of λ in L with endpoint $\lambda^\pm$ such that U is contained in a bounded neighborhood of both λ^+ and λ^-.

Assume we have such a neighborhood U of p and I of λ, and assume that $U \subset N_\epsilon(\lambda^+) \cap N_\epsilon(\lambda^-)$. Let $\gamma\colon \mathbb{R}^+ \to U \cap \lambda$ be a properly embedded ray and let x_i be a sequence of points so that $\gamma(x_i)$ is an ϵ net for the image of γ. Then there is a sequence of transversals τ_i of length bounded by $d(\epsilon)$ with $\phi(\tau_i) = I$ passing through $\gamma(x_i)$. Since $\tau_i \cap \lambda^+$ and $\tau_{i+1} \cap \lambda^+$ are at distance less than 3ϵ from each other in $\widetilde{M}$, they are distance less than $c(\epsilon)$ from each other in λ^+. A similar statement holds for $\tau_i \cap \lambda^-$ and $\tau_{i+1} \cap \lambda^-$. Therefore we can find a sequence of arcs $\alpha_i^\pm$ in $\lambda^\pm$ between these pairs of points. The circles

$$\tau_i \cup \alpha_i^+ \cup \tau_{i+1} \cup \alpha_i^-$$

bound disks of bounded diameter which are transverse to $\widetilde{\mathcal{F}}$ and whose intersection with λ is contained in the image of γ. These disks can be glued together to produce a proper map $H\colon \mathbb{R}^+ \times I \to \widetilde{M}$ with the desired properties, such that the vertical fibers $H(x, I)$ have length uniformly bounded by some function of ϵ. That is, they are uniformly bounded independently of γ. $\square$

Theorem 4.2.5 *Suppose every point $p \in S^1_\infty(\lambda)$ is confined. Then λ is a confined leaf.*

Proof By compactness, we can cover $\lambda \cup S^1_\infty$ with a finite number of open sets U_i so that there are neighborhoods I_i in L of λ with endpoints $\lambda_i^\pm$ with the property that $U_i \subset N_{\epsilon_i}(\lambda_i^+) \cap N_{\epsilon_i}(\lambda_i^-)$. (Notice that any open set U_i whose closure in λ_i is compact satisfies this property for some I_i and some ϵ_i). But this implies λ is confined, by the symmetry of the confinement condition. $\square$

Lemma 4.2.6 *Suppose that $\widetilde{\mathcal{F}}$ has no confined leaves. Let $p \in S^1_\infty(\lambda)$ be confined. Then with notation as in the proof of theorem 4.2.4, for any sequence $p_i \to p$ there are transversals τ_i with $\phi(\tau_i) = I$ such that $\|\tau_i\| \to 0$.*

Proof Let $\lambda^\pm$ be the endpoints of I. Then $U \subset N_\epsilon(\lambda^+) \cap N_\epsilon(\lambda^-)$, and therefore, if $B_{t_i}(p_i)$ denotes the ball in λ of radius t_i about p_i, we have that $B_{t_i}(p_i) \subset N_\epsilon(\lambda^+) \cap N_\epsilon(\lambda^-)$ for $t_i \to \infty$. Let $\alpha_i \in \pi_1(M)$ be chosen so that $\alpha_i(p_i) \to q \in \widetilde{M}$. Suppose no such shrinking transversals τ_i exist. Then infinitely many leaves $\alpha_i(\lambda^+), \alpha_i(\lambda^-)$ are bounded away from q. It follows that $\limsup \alpha^i(I) = J$ has non-empty interior. But by construction, the *entire* leaf through q is contained in a bounded neighborhood of the limit leaves of J. It follows that the leaf through q is confined, contrary to assumption. $\square$

Theorem 4.2.7 *The set of confined directions is open in C_∞.*

Proof For a uniform foliation, *every* direction is confined. Since every direction on a confined leaf is confined, we can assume without loss of generality that $\widetilde{\mathcal{F}}$ has no confined leaves.

Theorem 4.2.4 shows that the set of confined directions is open in each leaf. Moreover, it shows that if p is a confined point in $S^1_\infty(\lambda)$, then for some open neighborhood U of p in $\lambda \cup S^1_\infty(\lambda)$ and some neighborhood $I \subset L$ with limits $\lambda^\pm$, the set U is contained in $N_\epsilon(\lambda^+) \cap N_\epsilon(\lambda^-)$ for some ϵ. It is clear that for any open $V \in \lambda$ whose closure in λ is compact, we can replace U by $U \cup V$ after possibly increasing ϵ. It follows from lemma 2.2.3 that for some δ, $N_\delta(U) \cap \lambda^+$

contains an entire half-space in λ^+, and similarly for λ^-. Therefore if γ is a semi-infinite geodesic in λ emanating from v and converging to a confined point p, there is a geodesic $\gamma^+ \in \lambda^+$ which stays in a bounded neighborhood of γ.

By lemma 4.2.6 we see that the leaves $\lambda, \lambda^+, \lambda^-$ all converge near $U \cap S^1_\infty(\lambda)$. It follows that the geodesics γ and γ^+ are actually asymptotic, considered as properly embedded arcs in $\widetilde{M}$. □

Remark 4.2.8 We see from this theorem that every confined direction is weakly confined, as suggested by the terminology. The following theorem follows immediately from this observation and from theorem 3.2.2.

Theorem 4.2.9 *Let $\mathcal{C}$ denote the set of confined directions in C_∞. This set carries a $\pi_1(M)$–invariant vertical foliation transverse to the horizontal foliation, whose leaves are the maximal weakly confined transversals running through every confined point.*

Proof Immediate from theorem 3.2.2. □

4.3 Transverse vector fields

It is sometimes a technical advantage to choose a one-dimensional foliation transverse to $\mathcal{F}$ in order to unambiguously define holonomy transport of a transversal along some path in a leaf. We therefore develop some language and basic properties in this section.

Let X be a transverse vector field to $\mathcal{F}$. Then X lifts to a transverse vector field $\widetilde{X}$ to $\widetilde{\mathcal{F}}$. Following Thurston, we make the following definition.

Definition 4.3.1 A vector field X transverse to an $\mathbb{R}$–covered foliation $\mathcal{F}$ is *regulating* if every integral curve of $\widetilde{X}$ intersects every leaf of $\widetilde{\mathcal{F}}$.

Put another way, the integral curves of a regulating vector field in the universal cover map homeomorphically to L under ϕ. In fact, we will show in the sequel that *every* $\mathbb{R}$–covered foliation admits a regulating transverse vector field.

Definition 4.3.2 We say that a point $p \in S^1_\infty(\lambda)$ is *confined with respect to* X if for every sequence $p_i \to p$ there is a t and a neighborhood I of λ in L such that the integral curves σ_i of $\widetilde{X}$ passing through p_i with the property that $\phi(\sigma_i) = I$ satisfy $\|\sigma_i\| \le t$. If no integral curve of $\widetilde{X}$ passing through p_i has the property that $\phi(\sigma_i) = I$, we say that $\|\sigma_i\| = \infty$.

Theorem 4.3.3 *Let X be a regulating transverse vector field. Then a point $p \in S^1_\infty(\lambda)$ is confined iff it is confined with respect to X.*

Proof Confinement with respect to a vector field is a stronger property than mere confinement, so it suffices to show that a confined point is confined with respect to X.

Suppose we have neighborhoods U, I and a t as in Theorem 4.2.4. For a point $p \in \widetilde{M}$, let I_p be the set of leaves which intersect the ball of radius t about p. Then the integral curve σ_p of $\widetilde{X}$ passing through p with $\phi(\sigma_p) = I_p$ has length $||\sigma_p|| = f(p)$. This function is continuous in p, and depends only on the projection of p to M. Since M is compact, this function is bounded. It follows that if we have $p_i \to p$ and transversals τ_i through p_i with $||\tau_i|| < t$ that the transversals σ_i through p_i with endpoints on the same leaves as τ_i have uniformly bounded length. □

It is far from true that an arbitrary transverse vector field is regulating. However, the following is true.

Theorem 4.3.4 *Suppose $\mathcal{F}$ has no confined leaves. Let X be an* arbitrary *transverse vector field. Then a point $p \in S^1_\infty(\lambda)$ is confined iff it is confined with respect to X.*

Proof This theorem follows as above once we observe that any transverse vector field regulates the ϵ-neighborhood of every leaf for some ϵ. For, by lemma 4.2.6 we know that leaves converge at infinity near confined points. It follows that by choosing U, I suitably for a confined point p, that integral curves of $\widetilde{X}$ foliate $N_\epsilon(U) \cap \phi^{-1}(I)$ as a product, and that the length of these integral curves is uniformly bounded. Consequently, a sequence $p_i \to p$ determines a sequence σ_i of integral curves of $\widetilde{X}$ with uniformly bounded length, and p is confined with respect to X, as required. □

For uniform foliations $\mathcal{F}$, every point at infinity is confined. However, for any vector field X which is not regulating, there are points at infinity which are unconfined with respect to X. For example, the skew $\mathbb{R}$–covered foliations described in [13] and [32] have naturally defined transverse vector fields which are not regulating. Every point at infinity is confined, but there is a single point at infinity for each leaf in $\widetilde{\mathcal{F}}$ which is unconfined with respect to the non-regulating vector field. We will come back to this example in the sequel.

4.4 Fixed points in confined directions

Suppose in the remainder of this section that we have chosen some vector field X transverse to $\mathcal{F}$, which lifts to $\widetilde{X}$ transverse to $\widetilde{\mathcal{F}}$.

If K denotes the closure of the set of fixed points for the action of $\pi_1(M)$ on the cylinder C_∞, then it follows that the group $\pi_1(M)$ acts freely on the contractible manifold $\widetilde{M} \cup (C_\infty - K)$. It would be pleasant to conclude that $C_\infty - K$ is empty, since M is a $K(\pi, 1)$. However the following example shows that things are not so simple.

Example Let $\mathcal{F}$ be an $\mathbb{R}$–covered foliation with some leaf λ homeomorphic to a cylinder. Let $\hat{\mathcal{F}}$ be obtained by blowing up the leaf λ and perturbing the blown up leaves to be planes. Then this confined "pocket" of leaves gives rise to a disjoint union of cylinders at infinity, consisting entirely of confined directions, on which $\pi_1(M)$ acts without any fixed points.

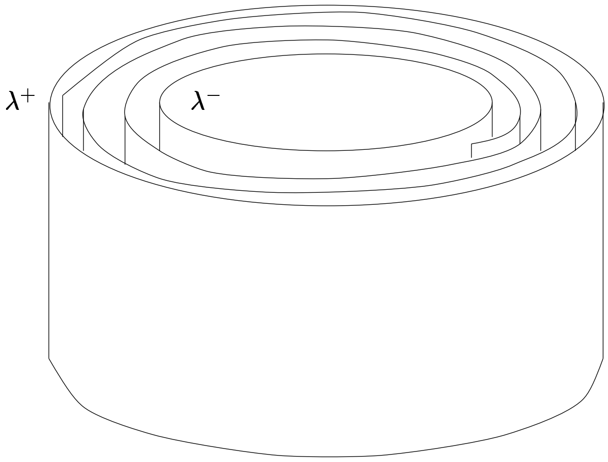

Figure 3: A cylinder is blown up to a foliated cylinder $\times I$. Then all but the boundary leaves are perturbed to planes. This pocket of leaves lifts to the universal cover to give an annulus of confined directions at infinity without any confined fixed points.

Fortunately, when every leaf is dense, we can say more about the action of $\pi_1(M)$ on C_∞. In particular, let S be any small rectangle whose boundary is contained in C_∞. We can define the (leafwise) *convex hull* $H(S)$ of S (or, generally of any subset of C_∞) to be the set of points $p \in \widetilde{M}$ such that if $p \in \lambda$, the visual angle of $\lambda_\infty \cap S$ as seen from p is $\geq \pi$. If S had the property that the translates of S under $\pi_1(M)$ were all disjoint, then the translates of the convex hull of S would also be disjoint, since there cannot be two disjoint closed arcs in a circle of angle $\geq \pi$. The following lemma quantifies the notion that every leaf of $\mathcal{F}$ is dense in M.

Lemma 4.4.1 *If $\mathcal{F}$ is a taut foliation of a manifold M such that every leaf is dense, then for every $\epsilon > 0$ there exists an R such that for any $p \in M$ and leaf λ containing p, the disk of radius R in λ with center p is an ϵ–net for M.*

Proof Observe that the such an $R(p)$ exists for every such $p \in M$. Moreover, by taking a larger $R(p)$ than necessary, we can find an $R(p)$ that works in an open neighborhood of p. Therefore by compactness of M we can find a universal R by taking the maximum of $R(p)$ over a finite open cover of M. □

In particular, for every λ, the set $\pi(\lambda \cap H(S))$ is dense in M. But now it follows that if $\tau \subset H(S)$ is any maximal integral curve of X, that there is some other maximal integral curve of X in $H(S)$, call it τ' and some $\alpha \in \pi_1(M)$ so that $\alpha(\tau') \subset \tau$. In particular, there is some $\alpha \in \pi_1(M)$ so that $\alpha(S) \cap S$ is a rectangle which is strictly bounded in the vertical direction by the upper and lower sides of S. In particular, α fixes some horizontal leaf passing through the interior of S.

More generally, we prove:

Theorem 4.4.2 *Fixed points of elements in $\pi_1(M)$ are dense in $\mathcal{C}$.*

Proof Let R be any confined rectangle. In local co-ordinates, let R be given by the set $|x| \leq 1, |y| \leq 1$ where the horizontal and vertical foliations of $\mathcal{C}$ in this chart are given by level sets of y and x respectively. Let $p \in \partial H(R)$ so that the visual angle of R is π as seen from p, and so that p is on the leaf corresponding to $y = 0$. There is some positive ϵ so that, as seen from p, there are no unconfined points within visual angle ϵ of the extreme left and right edges of R. But now we can find a q such that the visual angle of R as seen from q is at least $2\pi - \epsilon$ such that there is some $\alpha \in \pi_1(M)$ so that $\alpha(q) = p$, and so that the integral curve of $X \cap H(R)$ through q is very small compared to the integral curve of $X \cap H(R)$ through p. Moreover, the fact that the visual angle of $\alpha(R)$, as seen from p is at least $2\pi - \epsilon$, and consists entirely of confined directions, implies that the rectangles $\alpha(R)$ and R must intersect "transversely"; that is to say, $\alpha(R)$ is defined in local co-ordinates by $a < x < b, c < y < d$ where $a < -1 < 1 < b$ and $-1 < c < 0 < d < 1$. For, otherwise, the union $R \cup \alpha(R)$ would contain an entire circle at infinity, which circle could not contain any unconfined points, contrary to our assumption that no leaf is confined.

By two applications of the intermediate value theorem, it follows that α has some fixed point in R. Since R was arbitrary, it follows that confined fixed points are dense in $\mathcal{C}$. $\square$

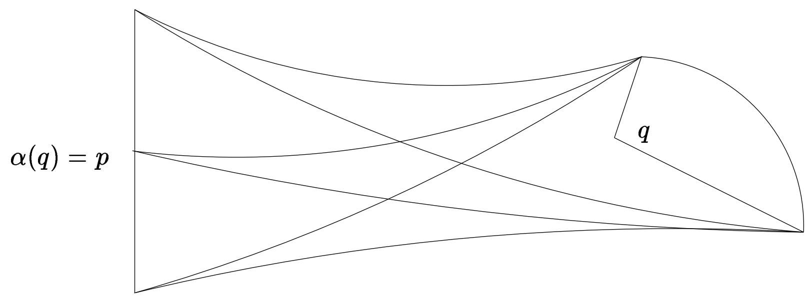

Figure 4: A sufficiently large disk about any point in any leaf is an ϵ–net for M. By going sufficiently far out towards $\mathcal{C}$ so that the vertical height of $H(R)$ is small, we can find points p, q and α as in the figure.

4.5 Semi-confined points

Given a point at infinity p and a side in C_∞ of the circle at infinity containing p, we say that p is *semi-confined* on that side if for all semi-infinite paths γ limiting to p, there is a transversal on the chosen side with one endpoint on the leaf through p which has holonomic images of bounded length along γ. If p is unconfined but still semi-confined, we say it is *strictly semi-confined.* Notice that the condition that p is unconfined implies that it can only be semi-confined on one side. It is clear from the definition that a semi-confined point can be a limit of unconfined points from only one side; that is, if p is a limit of unconfined points p_i, then the leaves containing p_i are all on the same side of p. We can actually prove the converse:

Lemma 4.5.1 *Let p be unconfined. Then on each side of p which is not semi-confined, p is a limit of unconfined points p_i.*

Proof Let R be a small rectangle in C_∞ containing p, bounded above and below by $S^1_\infty(\lambda^\pm)$ respectively. Let p lie on the leaf λ. Suppose without loss of generality that p is not semi-confined on the positive side. Then we can find a sequence of points $q_i \to p$ in λ such that the shortest transversal τ_i through q_i whose endpoints lie on λ and λ^+ has length bounded between i and $i+1$. By passing to a subsequence, we can find α_i so that $\alpha_i(q_i)$ converges to q. Let $H(R)$ denote the leafwise convex hull of R, and $\partial H(R)$ denote the leafwise boundary of this set — ie, the collection of geodesics in leaves of $\widetilde{\mathcal{F}}$ which limit to pairs of points on the vertical edges of R. Then the distance from q_i to $\partial H(R)$ gets larger and larger, so the rectangle R has visual angle $\to 2\pi$ as seen from q_i. If R contains unconfined points above p, we are done, since R was arbitrary. Otherwise the unconfined points on the leaves between λ and λ^+ are constrained to lie outside R. As seen from q_i, the visual angle of R converges to 2π, and the transversal between λ and λ^+ has length $\to \infty$. For each fixed distance $t > 0$, let $q_i(t)$ be the point on τ_i at distance t from q_i. Then the visual angle of R as seen from $q_i(t)$ also converges to 2π, since $q_i(t)$ is only a bounded distance from q_i and therefore the distance from $q_i(t)$ to $\partial H(R)$ also increases without bound. Therefore the geometric limit of $\alpha_i(R)$ is an infinite strip omitting exactly one vertical line at infinity which contains all the unconfined points. It follows that $C_\infty - \mathcal{C}$ is a single bi-infinite line containing all the unconfined points, including p. In particular, p is a limit of unconfined points from above and below. $\square$

Let p be a confined fixed point of an element $\alpha \in \pi_1(M)$. Let λ be the leaf of $\widetilde{\mathcal{F}}$ containing p. Then α acts as a hyperbolic isometry of λ, since otherwise its translation distance in $\widetilde{M}$ is 0, contradicting the fact that M is compact. Without loss of generality we can assume that p is an attracting fixed point for the action of α on λ. Let q be the other fixed point of p. Then for every point $p' \in S^1_\infty(\lambda) - q$ the sequence $\alpha^n(p') \to p$. It follows that every such p' is confined. By theorem 4.2.5 this implies that q is unconfined. Call such a q the *unconfined fixed point conjugate to* p.

Lemma 4.5.2 *Let q be the unconfined fixed point conjugate to some p in $S^1_\infty(\lambda)$. Let α fix the axis from p to q so that p is an attracting fixed point for α. Then for every sufficiently small rectangle R containing q in its interior α^{-1} takes R properly into its interior.*

Proof We can find confined transversals τ_1, τ_2 in C_∞ near q which run from λ^- to λ^+ for a pair of leaves $\lambda^\pm$ with $\lambda \in [\lambda^-, \lambda^+]$. Since α fixes a confined transversal through p, it expands this transversal, by lemma 4.2.6. It follows that α expands $[\lambda^-, \lambda^+]$ for all sufficiently close $\lambda^\pm$. Moreover, q is a repelling fixed point on $S^1_\infty(\lambda)$ for α, so the lemma follows. $\square$

4.6 Spines and product structures on C_∞

Definition 4.6.1 A π_1–invariant bi-infinite curve $\Psi \subset C_\infty$ intersecting every circle at infinity exactly once is called a *spine*.

Lemma 4.6.2 *Suppose there exists a spine Ψ. Then for any unconfined point $p \in C_\infty - \Psi$ and any pair of concentric rectangles $S \subset R$ containing p and avoiding the spine, there is some $\alpha \in \pi_1(M)$ which takes the rectangle R properly inside S.*

Proof Let I be a fixed transversal passing through the leaf λ containing p. Then there is an l such that any ball in any leaf of radius l contains a translate of some point in I. Since p is unconfined, there is a sequence $p_i \to p$ of points in λ such that the transversal with limits determined by S blows up to arbitrary length. Then we can find a p_i so that the ball of radius l in the leaf about p_i has the property that all transversals through this ball whose projection to L is equal to $\phi_v(S)$ are of length $> |I|$ on either side. For, the fact that $\mathcal{F}$ is $\mathbb{R}$–covered and M is compact implies that for any lengths l', t_1 there is a t_2 so that a transversal of length t_1 cannot blow up to length t_2 under holonomy transport of length $\leq l'$ (simply take the supremum of the lengths of holonomy transport of all transversals of length $\leq t_1$ under all paths of length $\leq l'$ and apply compactness).

But now it follows that some translate of I intersects the ball of radius l in the leaf about p_i in such a way that the translating element α maps the interval in leaf space delimited by R completely inside S. Furthermore, we can choose p_i as above so that the visual angle of S seen from any point in the ball is at least $2\pi - \epsilon$. This, together with the fact that both R and $\alpha(R)$ are the same visual angle away from the spine, as viewed from I and $\alpha(I)$ respectively, imply that $\alpha(R)$ is properly contained in S and therefore has an unconfined fixed point q in S with the desired properties. □

Theorem 4.6.3 *Let $\mathcal{F}$ be any nonuniform $\mathbb{R}$–covered foliation with dense leaves, not necessarily containing confined points at infinity. Let $\mathcal{I}$ be some nonempty π_1–invariant embedded collection of pairwise disjoint arcs transverse to the horizontal foliation of C_∞. Then at least one of the following two things happens:*

- *For any pair of leaves $\lambda < \mu$ in L, there are a collection of elements of $\mathcal{I}$ whose projection to L contains $[\lambda, \mu]$ and which intersect each of $S^1_\infty(\lambda)$ and $S^1_\infty(\mu)$ in a dense set.*

- *C_∞ contains a spine.*

In the first case, the set $\mathcal{I}$ determines a canonical identification between $S^1_\infty(\lambda)$ and $S^1_\infty(\mu)$ for any pair of leaves λ, μ.

Proof Observe that there is some element τ of $\mathcal{I}$ whose projection $\phi_v(\tau)$ contains $[\lambda, \mu]$, by corollary 2.4.3. Let I_i be an exhaustion of L by compact intervals, and let τ_i be a sequence of elements of $\mathcal{I}$ such that $I_i \subset \phi_v(\tau_i)$. Then we can extract a subsequence of τ_i which converges on compact subsets to a bi-infinite $\hat{\tau}$ which is transverse to the horizontal foliation of C_∞ and which does not cross any element of $\mathcal{I}$ transversely. Call such a $\hat{\tau}$ a *long transversal.* Let U be the complement of the closure of the set of long transversals. Then U is open and π_1–invariant, and is therefore either empty or omits at most one point in a.e. circle at infinity, by theorem 3.3.3. In the second case, it is clear that there is a unique long transversal, which must be a spine. In the first case, pick a point p in the cylinder limited by $S^1_\infty(\lambda)$ and $S^1_\infty(\mu)$. There is a long transversal arbitrarily close to p, and by the definition of a long transversal, there are elements of $\mathcal{I}$ stretching arbitrarily far in either direction of L arbitrarily close to such a long transversal. It follows that there is an element of $\mathcal{I}$ whose projection to L contains $[\lambda, \mu]$ arbitrarily close to p. The elements of $\mathcal{I}$ are disjoint, and therefore they let us canonically identify a dense subset of $S^1_\infty(\lambda)$ with a dense subset of $S^1_\infty(\mu)$; this identification can be extended uniquely by continuity to the entire circles. □

Theorem 4.6.4 *For any $\mathbb{R}$–covered foliation with hyperbolic leaves, not necessarily containing confined points at infinity, there are two natural maps*

$$\phi_v \colon C_\infty \to L, \quad \phi_h \colon C_\infty \to S^1_{\mathrm{univ}}$$

such that:

- *ϕ_v is the projection to the leaf space.*
- *ϕ_h is a homeomorphism for every circle at infinity.*
- *These functions give co-ordinates for C_∞ making it homeomorphic to a cylinder with a pair of complementary foliations in such a way that $\pi_1(M)$ acts by homeomorphisms on this cylinder preserving both foliations.*

Proof If $\mathcal{F}$ is uniform, any two leaves of $\widetilde{\mathcal{F}}$ are quasi-isometrically embedded in the slab between them, which is itself quasi-isometric to $\mathbb{H}^2$. It follows that the circles at infinity of every leaf can be canonically identified with each other, producing the product structure required. Furthermore, it is obvious that the product structure can be extended over blow-ups of leaves. We therefore assume that $\mathcal{F}$ is not uniform and has no confined leaves.

Consider $\mathcal{T}$, the union of weakly confined transversals. By theorem 4.6.3, we only need to consider the case that C_∞ has a spine; for otherwise there is a canonical identification of $S^1_\infty(\lambda)$ with $S^1_\infty(\mu)$ for any $\mu, \lambda \in L$, so we can fix some $S^1_\infty(\lambda) = S^1_{\text{univ}}$ and let ϕ_h take each point in some $S^1_\infty(\mu)$ to the corresponding point in $S^1_\infty(\lambda)$. It is clear that the fibers of this identification give a foliation of C_∞ with the required properties.

It follows that we may reduce to the case that there is a spine Ψ. Let Y be the vector field on $\widetilde{\mathcal{F}}$ which points towards the spine with unit length. Observe Y descends to a vector field on $\mathcal{F}$.

Definition 4.6.5 Say a semi-infinite integral curve $\gamma \subset \lambda$ of Y pointing towards the spine is *weakly expanding* if there exists an interval $I \subset L$ with λ in its interior such that holonomy transport through integral curves of Y keeps the length of a transversal representing I uniformly bounded below. That is there is a $\delta > 0$ such that for any map $H\colon [-1,1] \times \mathbb{R}^+ \to \widetilde{M}$ with the properties

- $\phi \circ H(*, t)\colon [-1,1] \to I$ is a homeomorphism for all t
- $H(r, *)\colon \mathbb{R}^+ \to \widetilde{M}$ is an integral curve of Y
- $H(0, *)\colon \mathbb{R}^+ \to \widetilde{M}$ is equal to the image of γ

we have $\|H([-1,1], t)\| > \delta$ independent of t and H.

Suppose that a periodic weakly expanding integral curve γ of Y exists. That is, there is $\alpha \in \pi_1(M)$ with $\alpha(\gamma) \subset \gamma$. By periodicity, we can choose I as above so that $\alpha(I) \subset I$, since a transversal representing I cannot shrink too small as it flows under Y. Then we claim *every* semi-infinite integral curve γ' of Y is *uniformly* weakly expanding. That is, there is a universal ϵ such that any interval $I \subset L$ with the property that the shortest transversal τ through the initial point of γ' with $\phi(\tau) = I$ has $\|\tau\| > \epsilon$ will have the properties required for the definition of a weakly expanding transversal, for some δ independent of γ' and depending only on ϵ.

To see this, let D be a fundamental domain for M centered around the initial point p of γ. Let R be a rectangle transverse to the integral curves of Y with top and bottom sides contained in leaves of $\widetilde{\mathcal{F}}$ and $\phi_v(R) = I$ such that D projects through integral curves of Y to a proper subset of R. Then projection through integral curves of Y takes the vertical sides of R properly inside the vertical sides of $\alpha(R)$, since the flow along Y shrinks distances in leaves.

Furthermore, since $\alpha(I) \subset I$, the top and bottom lines in R flow to horizontal lines which are above and below respectively the top and bottom lines of $\alpha(R)$

Thus, holonomy transport of any vertical line in R through integral curves of Y keeps its length uniformly bounded below by some δ. For any interval $J \subset L$ with $\phi_v(R) \subset J$ therefore, an integral curve of Y beginning at a point in D is weakly expanding for the interval J and for some universal δ as above. Since D is a fundamental domain, this proves the claim.

By theorem 3.2.3 there is some point $p \in C_\infty$ not on Ψ, a pair of leaves $\lambda^\pm$ above and below the leaf λ containing p, and a sequence of points p_i in λ converging to p such that the distance from p_i to $\lambda^\pm$ converges to 0. Let D be a disk in C_∞ about p. Then the visual angle of D, as seen from p_i, converges to 2π. Moreover, there are a sequence of transversals τ_i between $\lambda^\pm$ passing through p_i whose length converges to 0. Since there is a uniform t so that any disk in a leaf of radius t intersects a translate of τ_1, we can find points p_i' in λ within a distance t of p_i so that there exists α_i with $\alpha_i(p_i') = p_1$. This α_i must satisfy $\alpha_i([\lambda^-, \lambda^+]) \subset [\lambda^-, \lambda^+]$ and furthermore it must fix Ψ, since Ψ is invariant under every transformation. If the visual angle of D seen from p_i' is at least $2\pi - \epsilon$ where D is at least ϵ away from the spine, as seen from p_1, then α_i must also fix a point in D. It follows that a semi-infinite ray contained in the axis of α_i going out towards Ψ is a periodic weakly expanding curve. This implies, as we have pointed out, that every semi-infinite integral curve of Y is uniformly weakly expanding.

We show now that the fact that every integral curve of Y is uniformly weakly expanding is incompatible with the existence of unconfined points off the spine.

For, by lemma 4.6.2 the existence of an unconfined point q implies that there are α_i fixing points at infinity near q which take a fixed disk containing q into arbitrarily small neighborhoods of q. This implies that as one goes out to infinity away from the spine along the axes of the α_i that some transversal is blown up arbitrarily large. Conversely, this implies that going along these axes in the opposite direction — towards the spine — for any t, ϵ we can find shortest transversals of length $\geq t$ which are shrunk to transversals of length $\leq \epsilon$ by flowing along Y. This contradicts the uniformly weakly expanding property of integral curves of Y. This contradiction implies that there are no unconfined points off the spine.

In either case, then we have shown that there are a dense set of vertical leaves in $\mathcal{C}$ between μ and λ. This lets us canonically identify the entire circles at infinity μ and λ. Since μ and λ were arbitrary, we can define ϕ_h to be the canonical identification of every circle at infinity with $S^1_\infty(\mu)$. $\square$

Remark 4.6.6 The identification of all the circles at infinity of every leaf with a single "universal" circle generalizes Thurston's universal circle theorem (see [31] or [5] for details of an alternative construction) to $\mathbb{R}$–covered foliations. The universal circle produced in [31] is not necessarily canonically homeomorphic to every circle at infinity; rather, one is guaranteed a monotone map from this universal circle to the circle at infinity of each leaf.

Remark 4.6.7 There is another approach to theorem 4.6.3 using "leftmost admissible trajectories". It is this approach which generalizes to the context of taut foliations with branching, and allows one to prove Thurston's universal circle theorem.

4.7 Spines and Solvmanifolds

Corollary 4.7.1 *If there exists at least one semi-confined point in C_∞ and if every semi-confined point is confined, the unconfined points lie on a spine.*

Proof Let R_1 be a closed rectangle containing some unconfined point p. We can find such an R_1 so that the left and right vertical edges of R are confined. Then if K_1 denotes the intersection of the unconfined points with R_1, $\phi_v(K_1)$ is a closed subset of an interval. Suppose it does not contain the entire interval. Then its image contains a limit point which is a limit of points from below but not from above. This pulls back to an unconfined point in R_1, which point must necessarily be semi-confined, contrary to assumption. Hence $\phi_v(K_1)$ is the entire closed interval. But R_1 was arbitrary, so by the density of vertical confined directions, we can take a sequence R_i limiting in the Hausdorff sense to a single vertical interval containing p. Since $\phi_v(K_i)$ is still the entire interval, it follows that the entire interval τ containing p is unconfined. If α_i is a sequence of elements of $\pi_1(M)$ which blow up $\phi_v(\tau)$ to all of L, then every α_i must preserve the vertical leaf containing τ, since otherwise there would be an interval of leaves containing at least two unconfined points. It follows that there is a single bi-infinite vertical leaf of unconfined directions, which must be π_1–invariant, and which contains p. But p was an arbitrary unconfined point, and therefore *every* such point is contained in the spine. □

Theorem 4.7.2 *If C_∞ contains a spine Ψ and $\mathcal{F}$ is $\mathbb{R}$–covered but not uniform, then M is a Solvmanifold and $\mathcal{F}$ is the suspension foliation of the stable or unstable foliation of an Anosov automorphism of a torus.*

Proof Since leaves of $\widetilde{\mathcal{F}}$ come close together as one goes out towards infinity in a confined direction, it follows that the map ϕ_h is compatible with the projective structures on each circle at infinity coming from their identifications with the circle at infinity of $\mathbb{H}^2$. More explicitly, a transverse vector field X to $\mathcal{F}$ regulates a uniform neighborhood of any leaf. Transport along integral curves of X determines a quasi-conformal map between the subsets of two leaves λ and μ which are sufficiently close together, and the modulus of dilatation can be bounded in terms of the length of integral curves of X between the leaves. Since this length goes to 0 as we go off to infinity anywhere except the spine, the map is more and more conformal as we go off to infinity, and in fact is a 1–quasisymmetric map at infinity, away from the spine, and is therefore symmetric (see [24] or [25]). Hence it preserves the projective structure on these circles.

It follows that $\pi_1(M)$ acts as a group of projective transformations of $(S^1,*)$, which is to say, as a group of similarities of $\mathbb{R}$. For, given $\alpha \in \pi_1(M)$ and any leaf $\lambda \in \widetilde{\mathcal{F}}$, the map $\alpha\colon \lambda \to \alpha(\lambda)$ is an isometry and therefore induces a projective map $\lambda_\infty \to (\alpha(\lambda))_\infty$; but ϕ_v is projective on every circle at infinity, by the above discussion, and so $\phi_v \circ \alpha$ is a projective map from the universal circle at infinity to itself. There is a homomorphism to $\mathbb{R}$ given by logarithm of the distortion; the image of this is actually discrete, since it is just the translation length of the element acting on a leaf of $\widetilde{\mathcal{F}}$, now identified with $\mathbb{H}^2$. Such translation lengths are certainly bounded away from 0 since M is a compact manifold and has a lower bound on its geodesic length spectrum. Hence we can take this homomorphism to $\mathbb{Z}$. But the kernel of this homomorphism is abelian, so $\pi_1(M)$ is solvable and M is a torus bundle over S^1, as required. □

It follows that we have proved the following theorem:

Theorem 4.7.3 *Let $\mathcal{F}$ be an $\mathbb{R}$–covered taut foliation of a closed 3–manifold M with hyperbolic leaves. Then after possibly blowing down confined regions, $\mathcal{F}$ falls into exactly one of the following four possibilities:*

- *$\mathcal{F}$ is uniform.*
- *$\mathcal{F}$ is (isotopic to) the suspension foliation of the stable or unstable foliation of an Anosov automorphism of T^2, and M is a Solvmanifold.*
- *$\mathcal{F}$ contains no confined leaves, but contains strictly semi-confined directions.*
- *$\mathcal{F}$ contains no confined directions.*

Remark 4.7.4 We note that in [32], Thurston advertises a forthcoming paper in which he intends to prove that uniform foliations are geometric. We expect that the case of strictly semi-confined directions cannot occur; any such example must be quite bizarre. We make the following conjecture:

Conjecture *If an $\mathbb{R}$–covered foliation has no confined leaves then it has no strictly semi-confined directions.*

Remark 4.7.5 In fact, we do not even know the answer to the following question in point set topology: suppose a finitely generated group Γ acts by homeomorphisms on $\mathbb{R}$ and on S^1. Let it act on the cylinder $\mathbb{R} \times S^1$ by the product action. Suppose $K \subset \mathbb{R} \times S^1$ is a minimal closed, invariant set for the action of Γ with the property that the projection to the $\mathbb{R}$ factor is $1-1$ on a dense set of points. Does K contain the non-constant continuous image of an interval?

Remark 4.7.6 Finally, we note that foliations with no confined directions do, in fact, exist, even in atoroidal 3–manifolds. A construction is given in [2].

5 Ruffled foliations

5.1 Laminations

In this section we study ruffled foliations, and in particular their interactions with essential laminations.

We begin with some definitions that will be important to what follows.

Definition 5.1.1 A *lamination* in a 3–manifold is a foliation of a closed subset of M by 2–dimensional leaves. The complement of this closed subset falls into connected components, called *complementary regions*. A lamination is *essential* if it contains no spherical leaf or torus leaf bounding a solid torus, and furthermore if C is the closure (with respect to the path metric) of a complementary region, then C is irreducible and ∂C is both incompressible and *end incompressible* in C. Here an end compressing disk is an embedded $(D^2 - (\text{closed arc in } \partial D^2))$ in C which is not properly isotopic rel ∂ in C to an embedding into a leaf. Finally, an essential lamination is *genuine* if it has some complementary region which is not an I–bundle.

Each complementary region falls into two pieces: the *guts*, which carry the essential topology of the complementary region, and the *interstitial regions*, which are just I bundles over non-compact surfaces, which get thinner and thinner as they go away from the guts. The interstitial regions meet the guts along annuli. Ideal polygons can be properly embedded in complementary regions, where the cusp neighborhoods of the ideal points run up the interstitial regions as $I \times \mathbb{R}^+$. An end compressing disk is just a properly embedded monogon which is not isotopic rel ∂ into a leaf. See [20] or [18] for the basic properties of essential laminations.

Definition 5.1.2 A lamination of $\mathbb{H}^2$ is an embedded collection of bi-infinite geodesics which is closed as a subset of $\mathbb{H}^2$.

Definition 5.1.3 A lamination of a circle S^1 is a closed subset of the space of unordered pairs of distinct points in S^1 such that no two pairs link each other.

If we think of S^1 as the circle at infinity of $\mathbb{H}^2$, a lamination of S^1 gives rise to a lamination of $\mathbb{H}^2$, by joining each pair of points in S^1 by the unique geodesic in $\mathbb{H}^2$ connecting them. A lamination Λ_{univ} of $\mathcal{S}^1_{\mathrm{univ}}$ invariant under the action of $\pi_1(M)$ determines a lamination in each leaf of $\widetilde{\mathcal{F}}$, and the union of these laminations sweep out a lamination $\widetilde{\Lambda}$ of $\widetilde{M}$ which, by equivariance of the construction, covers a lamination Λ in M. By examining $\widetilde{\Lambda}$ one sees that Λ is genuine.

5.2 Invariant structures are vertical

Definition 5.2.1 Let $\mathcal{F}$ be an $\mathbb{R}$–covered foliation of M with dense hyperbolic leaves. If $\mathcal{F}$ is neither uniform nor the suspension foliation of an Anosov automorphism of a torus, then say $\mathcal{F}$ is *ruffled.*

The definition of "ruffled" therefore incorporates both of the last two cases in theorem 4.7.3.

Lemma 5.2.2 *Let $\mathcal{F}$ be ruffled. Then the action of $\pi_1(M)$ on S^1_{univ} is minimal; that is, the orbit of every point is dense. In fact, for any pair I, J or intervals in S^1_{univ}, there is an $\alpha \in \pi_1(M)$ for which $\alpha(I) \subset J$.*

Proof For $p \in S^1_{\mathrm{univ}}$, let o_p be the closure of the orbit of p in S^1_{univ}, and let V_p be the union of the leaves of the vertical foliation of C_∞ corresponding to o_p. By theorem 4.6.3 the set V_p is either all of C_∞ or there is a spine; but $\mathcal{F}$ is ruffled, so there is no spine.

Now let I, J be arbitrary. There is certainly some sequence α_i so that $\alpha_i(I)$ converges to a single point p, since we can look at a rectangle $R \subset C_\infty$ with $\phi_h(R) = I$ and choose a sequence of points in $\widetilde{M}$ from which the visual angle of R is arbitrarily small, and choose a convergent subsequence. Conversely, we can find a sequence of elements β_i so that $\beta_i(J)$ converges to the complement of a single point q. Now choose some γ so that $\gamma(p) \neq q$. Then $\beta_j^{-1}\gamma\alpha_i(I) \subset J$ for sufficiently large i, j. □

Lemma 5.2.3 *Let $\mathcal{F}$ be a ruffled foliation. Then for any rectangle $R \subset C_\infty$ with vertical sides in leaves of the vertical foliation and horizontal sides in leaves of the horizontal foliation, for every $p \in S^1_{\text{univ}}$, and for every weakly confined transversal τ dividing R into two rectangles R_l, R_r, there are a sequence of elements $\alpha_i \in \pi_1(M)$ so that*

$$\phi_v(\alpha_i(R')) \to L \text{ and } \phi_h(\alpha_i(R')) \to p$$

for R' one of R_l, R_r.

Proof We have seen that weakly confined transversals are dense in C_∞. Let τ be such a transversal such that $\phi_v(R) \subset \phi_v(\tau)$, and observe that τ divides R into two rectangles R_l, R_r. There is a sequence of elements α_i in $\pi_1(M)$ which blow up τ to an arbitrarily long transversal, as seen from some fixed $p \in \widetilde{M}$ such that $\phi_v(p) \in \phi_v(\alpha_i(R))$. Let λ be a leaf in $\phi_v(R)$. Then the points in λ from which the visual angle of both R_r and R_l are bigger than ϵ, are contained in a bounded neighborhood of a geodesic ray in λ limiting to $\tau \cap S^1_\infty(\lambda)$. Since τ is a weakly confined transversal, the length of a shortest transversal σ with $\phi_v(\sigma) = \phi_v(R)$ running through such a point is uniformly bounded. It follows that for our choice of p as above, for at least one of R_l, R_r (say R_l) the visual angle of $\alpha_i(R_l)$ goes to zero, as seen from p. It follows that there is a subsequence of α_i for which $\phi_v(\alpha_i(R_l)) \to L$ and $\phi_h(\alpha_i(R_l)) = q$. If β_i is a sequence of elements for which $\beta_i(q) \to p$, then the sequence $\beta_i\alpha_{n_i}$ for n_i growing sufficiently fast will satisfy

$$\phi_v(\beta_i\alpha_{n_i}(R_l)) \to L \text{ and } \phi_h(\beta_i\alpha_{n_i}(R_l)) \to p. \qquad \square$$

The method of proof used in theorem 4.6.4 is quite general, and may be understood as showing that for a ruffled foliation, certain kinds of $\pi_1(M)$–invariant structures at infinity must come from $\pi_1(M)$–invariant structures on the universal circle S^1_{univ}. For, any group-invariant structure at infinity can be "blown up" by the action of $\pi_1(M)$ so that it varies less and less from leaf to leaf. By extracting a limit, we can find a point $p \in S^1_{\text{univ}}$ corresponding to a vertical leaf in C_∞ where the structure is constant. Either this vertical leaf is unique, in which case it is a spine and M is Solv, or the orbit of p is dense in S^1_{univ} by theorem 3.3.3 and our structure is constant along all vertical leaves in C_∞ — that is, it comes from an invariant structure on S^1_{univ}.

We can make this precise as follows:

Theorem 5.2.4 *Let $\mathcal{F}$ be a ruffled foliation, and let $\mathcal{I}$ be a π_1–invariant collection of embedded pairwise-disjoint arcs in C_∞ transverse to the horizontal foliation by circles. Then $\mathcal{I}$ is* vertical*: that is, the arcs in $\mathcal{I}$ are contained in the vertical foliation of C_∞ by preimages of points in S^1_{univ}.*

Proof Since $\mathcal{F}$ is ruffled, C_∞ does not admit a spine. Therefore by theorem 4.6.3, we know that for any pair of leaves $\lambda < \mu$, there are a set of arcs in $\mathcal{I}$ whose projection to L includes $[\lambda, \mu]$ and intersect each of $S^1_\infty(\lambda)$ and $S^1_\infty(\mu)$ in a dense set of points. It follows that there is a product structure $C_\infty = S^1_{\mathcal{I}} \times \mathbb{R}$ so that the elements of $\mathcal{I}$ are contained in the vertical foliation $\mathcal{F}_{\mathcal{I}}$ for this product structure.

We claim that this foliation agrees with the vertical foliation by preimages of points in S^1_{univ} under ϕ_h^{-1}.

For, let τ_1, τ_2 be two segments of $\mathcal{F}_{\mathcal{I}}$ running between leaves λ, μ so that $\phi_h(\tau_1)$ and $\phi_h(\tau_2)$ are disjoint. Then we can find a rectangle R with vertical sides in the vertical foliation of C_∞ and $\phi_v(R) = \phi_v(\tau_1) = \phi_v(\tau_2)$ which is divided into rectangles R_l, R_r by a weakly confined transversal as in the hypothesis of lemma 5.2.3 so that $\tau_1 \subset R_l$ and $\tau_2 \subset R_r$. Then lemma 5.2.3 implies that for any $p \in S^1_{\text{univ}}$, there are a sequence of elements α_i so that for some j, $\phi_v(\alpha_i(\tau_j)) \to L$ and $\phi_h(\alpha_i(\tau_j)) \to p$. It follows that there is a vertical leaf of $\mathcal{F}_{\mathcal{I}}$ which agrees with $\phi_h^{-1}(p)$. Since p was arbitrary, the foliation $\mathcal{F}_{\mathcal{I}}$ agrees with the vertical foliation of C_∞; that is, $\mathcal{I}$ is vertical, as required. ☐

Theorem 5.2.5 *Let $\mathcal{F}$ be a ruffled foliation. Let Λ be any essential lamination transverse to $\mathcal{F}$ intersecting every leaf of $\mathcal{F}$ in quasi-geodesics. Then Λ is regulating. That is, the pulled-back lamination $\widetilde{\Lambda}$ of $\widetilde{M}$ comes from a $\pi_1(M)$–invariant lamination in S^1_{univ}.*

Proof Let λ be a leaf of $\widetilde{\Lambda}$. Then λ intersects leaves of $\widetilde{\mathcal{F}}$ in quasi-geodesics whose endpoints determines a pair of transverse curves in C_∞. These transverse curves are continuous for the following reason. We can straighten λ leafwise in its intersection with leaves of $\widetilde{\mathcal{F}}$ so that these intersections are all geodesic.

This "straightening" can be done continuously; for if $s, s' = \lambda \cap \nu, \nu'$ for ν, ν' leaves of $\widetilde{\mathcal{F}}$, and σ, σ' are long segments of s, s', then the straightenings of σ, σ' stay very close to the straightenings of s, s' along most of their interiors. In particular, the straightenings of σ and σ' are very close, since the leaves ν, ν' are close along σ, σ'. Thus the straightenings of s, s' will be close wherever ν, ν' are close, which is the definition of continuity. If τ is a transversal to $\widetilde{\mathcal{F}}$ contained in λ, then we can identify $UT\mathcal{F}|_\tau$ with a cylindrical subset of C_∞. The endpoints of λ can be identified with $UT\mathcal{F}|_\tau \cap T\lambda$ and therefore sweep out continous curves.

By theorem 5.2.4, these transverse curves are actually leaves of the vertical foliation of C_∞, and therefore each leaf of $\widetilde{\Lambda}$ comes from a leaf of a $\pi_1(M)$–invariant lamination of S^1_{univ}. □

If Λ is transverse to $\mathcal{F}$ but does not intersect quasigeodesically, we can nevertheless make the argument above work, except in extreme cases. For, if μ is a leaf of $\widetilde{\mathcal{F}}$ and λ is a leaf of $\widetilde{\Lambda}$ such that $\mu \cap \lambda = \alpha$, then we can look at the subsets $\alpha^\pm$ of $S^1_\infty(\mu)$ determined by the two ends of α. If these are both proper subsets, we can "straighten" α to a geodesic $\overline{\alpha}$ running between the two most anticlockwise points in $\alpha^\pm$. This straightens Λ to $\overline{\Lambda}$ which intersects $\mathcal{F}$ geodesically. Of course, we may have collapsed Λ somewhat in this process.

5.3 Constructing invariant laminations

In this section we show that for M atoroidal and $\mathcal{F}$ ruffled, there exist a pair of essential laminations $\Lambda^\pm$ with solid toroidal complementary regions which intersect each other and $\mathcal{F}$ transversely, and whose intersection with $\mathcal{F}$ is geodesic. By theorem 5.2.5 such laminations must come from a pair of transverse invariant laminations of S^1_{univ}, but this is actually the method by which we construct them.

Definition 5.3.1 A *quadrilateral* is an ordered 4–tuple of points in S^1 which bounds an embedded ideal rectangle in $\mathbb{H}^2$.

Let S_4 denote the space of ordered 4–tuples of distinct points in S^1 whose ordering agrees with the circular order on S^1. We fix an identification of S^1 with $\partial\mathbb{H}^2$. To each 4–tuple in S_4 there corresponds a point $p \in \mathbb{H}^2$ which is the center of gravity of the ideal quadrilateral whose vertices are the four points in question. Let $\overline{S_4}$ denote the space obtained from S_4 by adding limits of 4–tuples whose center of gravity converges to a definite point in $\mathbb{H}^2$. For $R \in \overline{S_4}$

let $c(R) =$ center of gravity. We say a sequence of 4–tuples *escapes to infinity* if their corresponding sequence of centers of gravity exit every compact subset of $\mathbb{H}^2$. We will sometimes use the terms 4–tuple and quadrilateral interchangeably to refer to an element of $\overline{S_4}$, where it should be understood that the geometric realization of such a quadrilateral may be degenerate. Let $S_4' = \overline{S_4} - S_4$ be the set of degenerate quadrilaterals whose center of gravity is well-defined, but the vertices of the quadrilateral have come together in pairs.

Corresponding to an ordered 4–tuple of points $\{a, b, c, d\}$ in $S^1 = \partial\mathbb{H}^2$ there is a real number known as the *modulus* or *cross-ratio*, defined as follows. Identify S^1 with $\mathbb{R} \cup \infty$ by the conformal identification of the unit disk with the upper half-plane. Let $\alpha \in PSL(2, \mathbb{R})$ be the unique element taking a, b, c to $0, 1, \infty$. Then $\mathrm{mod}(\{a, b, c, d\}) = \alpha(d)$. Note that we can extend mod to all of $\overline{S_4}$ where it might take the values 0 or ∞.

See [24] for the definition of the modulus of a quadrilateral and a discussion of its relation to quasiconformality and quasi-symmetry.

Definition 5.3.2 A group Γ of homeomorphisms of S^1 is *renormalizable* if for any bounded sequence $R_i \in \overline{S_4}$ with $|\mathrm{mod}(R_i)|$ bounded such that there exists a sequence $\alpha_i \in \Gamma$ with $|\mathrm{mod}(\alpha_i(R_i))| \to \infty$ there is another sequence β_i such that $|\mathrm{mod}(\beta_i(R_i))| \to \infty$ and $\beta_i(R_i) \to R' \in S_4'$.

Definition 5.3.3 Let $\alpha \in \hom(S^1)$. We say that α is *weakly topologically pseudo-Anosov* if there are a pair of disjoint closed intervals $I_1, I_2 \subset S^1$ which are both taken properly into their interiors by the action of α. We say that α is *topologically pseudo-Anosov* if α has $2n$ isolated fixed points, where $2 < 2n < \infty$ such that on the complementary intervals α translates points alternately clockwise and anticlockwise.

Obviously an α which is topologically pseudo-Anosov is weakly topologically pseudo-Anosov. A topologically pseudo-Anosov element has a pair of fixed points in the associated intervals I_1, I_2; such fixed points are called *weakly attracting.*

The main idea of the following theorem was communicated to the author by Thurston:

Theorem 5.3.4 (Thurston) *Let G be a renormalizable group of homeomorphisms of S^1 such that no element of G is weakly topologically pseudo-Anosov. Then either G is conjugate to a subgroup of $PSL(2, \mathbb{R})$, or there is a lamination Λ of S^1 left invariant by G.*

Proof Suppose that there is no sequence R_i of 4–tuples and $\alpha_i \in G$ such that $\mathrm{mod}(R_i) \to 0$ and $\mathrm{mod}(\alpha_i(R_i)) \to \infty$. Then the closure of G is a Lie group, and therefore either discrete, or conjugate to a Lie subgroup of $PSL(2,\mathbb{R})$, by the main result of [23]. If G is discrete it is a convergence group, and the main result of [16] or [8], building on substantial work of Tukia, Mess, Scott and others, implies G is a Fuchsian group.

Otherwise the assumption of renormalizability implies there is a sequence R_i of 4–tuples with $|\mathrm{mod}(R_i)|$ bounded and a sequence $\alpha_i \in G$ such that

$$\mathrm{mod}(\alpha_i(R_i)) \to \infty$$

and $c(R_i)$ and $c(\alpha_i(R_i))$ both converge to particular points in $\mathbb{H}^2$. A 4–tuple can be *subdivided* as follows: if a,b,c,d,e,f is a cyclically ordered collection of points in S^1 we say that the two 4–tuples $\{a,b,e,d\}$ and $\{b,c,d,e\}$ are obtained by subdividing $\{a,c,d,f\}$. If we subdivide R_i into a pair of 4–tuples R_i^1, R_i^2 with moduli approximately equal to $\frac{1}{2}\mathrm{mod}(R_i)$, then a subsequence in $\mathrm{mod}(\alpha_i(R_i^j))$ converges to infinity for some fixed $j \in \{1,2\}$. Subdividing inductively and extracting a diagonal subsequence, we can find a sequence of 4–tuples which we relabel as R_i with

$$\mathrm{mod}(R_i) \to 0 \text{ and } \mathrm{mod}_i(\alpha_i(R_i)) \to \infty$$

with $c(R_i)$ and $c(\alpha_i(R_i))$ bounded in $\mathbb{H}^2$. Extracting a further subsequence, it follows that there are a pair of geodesics γ, τ of $\mathbb{H}^2$ such that the points of R_i converge in pairs to the endpoints of γ, and the points of $\alpha_i(R_i)$ converge in pairs to the endpoints of τ, in such a way that the partition of R_i into convergent pairs is different in the two cases. Informally, a sequence of "long, thin" rectangles is converging to a core geodesic. Its images under the α_i are a sequence of "short, fat" rectangles, converging to another core geodesic. We can distinguish a "thin" rectangle from a "fat" rectangle by virtue of the fact that the R_i are *ordered* 4–tuples, and therefore we know which are the top and bottom sides, and which are the left and right sides.

We claim that no translate of γ can intersect a translate of τ. For, this would give us a new sequence of elements α_i which were manifestly weakly topologically pseudo-Anosov, contrary to assumption. It follows that the unions $G(\gamma)$ and $G(\tau)$ are *disjoint* as subsets of $\mathbb{H}^2$.

We point out that this is actually enough information to construct an invariant lamination, in fact a pair of such. For, since no geodesic in $G(\gamma)$ intersects a geodesic in $G(\tau)$, the connected components of $G(\tau)$ separate the connected components of $G(\gamma)$ — in fact, since $G(\tau)$ is a union of geodesics, it separates the *convex hulls* of the connected components of $G(\gamma)$. Let C_i be the convex hulls of the connected components of $G(\gamma)$. It is straightforward to see that there are infinitely many C_i. Each C_i has nonempty boundary consisting of a collection of geodesics ∂C_i, and the invariance of $G(\gamma)$ under G implies $\bigcup_i \partial C_i$ has closure a geodesic lamination. A similar construction obviously works for the connected components of $G(\tau)$.

But in fact we can show that *a priori* the closure of one of $G(\gamma)$ or $G(\tau)$ is a lamination. For, suppose $\alpha(\gamma)$ intersects γ transversely for some γ. Then if $R_i \to \gamma$ with $\alpha_i(R_i) \to \tau$, we must have $\alpha_i\alpha(\gamma) \to \tau$. It follows that τ is a limit of leaves of $G(\gamma)$. If now for some β we have $\beta(\tau)$ intersects τ transversely, then $\beta(\tau)$ intersects $\alpha_i\alpha(\gamma)$ transversely for sufficiently large i, and therefore some element of G is weakly topologically pseudo-Anosov. $\square$

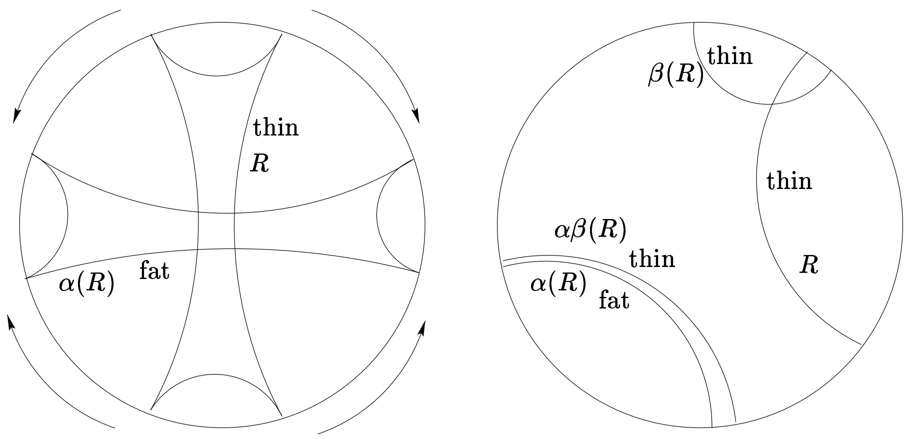

Figure 5: A fat rectangle cannot cross a thin rectangle, or some element would act on S^1 in a weakly pseudo-Anosov manner. Similarly, if a thin rectangle crosses a thin rectangle, a translate of this thin rectangle "protects" fat rectangles from being crossed by fat rectangles.

Theorem 5.3.4 is especially important in our context, in view of the following observation:

Lemma 5.3.5 *Let $\pi_1(M) \to S^1_{univ}$ be the standard action, where S^1_{univ} inherits the symmetric structure from $S^1_\infty(\lambda)$ for some leaf λ of $\widetilde{\mathcal{F}}$. Then this action is renormalizable.*

Proof Let D be a fundamental domain for M intersecting λ. Suppose we have a sequence of 4–tuples R_i in S^1_{univ} whose moduli, as measured by the identification of S^1_{univ} with $S^1_\infty(\lambda)$, goes to 0. Then this determines a sequence of rectangles in λ with moduli $\to 0$, whose centers of mass can all be translated by elements β_i of $\pi_1(M)$ to intersect D. By compactness of D, as we sweep the rectangles $\beta_i(R_i)$ through the leaf space of $\widetilde{\mathcal{F}}$ to λ, their modulus does not distort very much, and their centers of mass can be made to land in a fixed compact region of λ. If α_i is a sequence in $\pi_1(M)$ such that the moduli of $\alpha_i(R_i)$ converges to ∞ as measured in $S^1_\infty(\lambda)$, we can translate the corresponding rectangles in λ back to D by γ_i without distorting their moduli too much. This shows the action is renormalizable, as required. □

We discuss the implications of these results for the action of $\pi_1(M)$ on S^1_{univ}.

Lemma 5.3.6 *The action of $\pi_1(M)$ is one of the following three kinds:*

- *$\pi_1(M)$ is a convergence group, and therefore conjugate to a Fuchsian group.*
- *There is an invariant lamination Λ_{univ} of S^1_{univ} constructed according to theorem 5.3.4.*
- *There are two distinct pairs of points p, q and r, s in S^1_{univ} which link each other so that for each pair of closed intervals I, J in $S^1_{univ} - \{r, s\}$ with $p \in I$ and $q \in J$ the sequence α_i restricted to the intervals I, J converge to p, q uniformly as $i \to \infty$, and α_i^{-1} restricted to the intervals $S^1 - (I \cup J)$ converge to r, s uniformly as $i \to \infty$.*

Proof If $\pi_1(M)$ is not Fuchsian, by lemma 5.3.5, there are a sequence of 4–tuples R_i with moduli $\to 0$ converging to γ and a sequence $\alpha_i \in \pi_1(M)$ so that $\text{mod}(\alpha_i(R_i)) \to \infty$ and $\alpha_i(R_i) \to \tau$. Either all the translates of γ are disjoint from τ and vice versa, or we are in the situation of the third alternative.

If all the translates of γ avoid all the translates of τ, the closure of the union of translates of one of these gives an invariant lamination. □

In fact we will show that the second case cannot occur. However, the proof of this relies logically on lemma 5.3.6. It is an interesting question whether one can show the existence of a family of weakly topologically pseudo-Anosov elements of $\pi_1(M)$ directly.

We analyze the action of $\pi_1(M)$ on S^1_{univ} in the event of the third alternative provided by lemma 5.3.6.

Lemma 5.3.7 *Suppose $\pi_1(M)$ acts on S^1_{univ} in a manner described in the third alternative given by lemma 5.3.6. Let γ be the geodesic joining p to q and γ' the geodesic joining r to s. Then the closure of $\pi_1(M)(\gamma)$ is an invariant lamination Λ^+_{univ} of S^1_{univ}, and similarly the closure of $\pi_1(M)(\gamma')$ is an invariant lamination Λ^-_{univ} of S^1_{univ}.*

Proof All we need do to prove this lemma is to show that no translate of γ intersects itself. Let $\alpha(\gamma)$ intersect γ transversely. Then the endpoints of $\alpha(\gamma)$ avoid I, J for some choice of I, J containing p, q respectively. We know α_i does not fix any leaf of $\widetilde{\mathcal{F}}$, since otherwise its action on S^1_{univ} would be topologically conjugate to an element of $PSL(2,\mathbb{R})$. For sufficiently large i, depending on our choice of I, J, the dynamics of α_i imply that there are two fixed points p_i, q_i for α_i, very close to p, q; in particular, they are contained in I, J. Let γ_i be the geodesic joining p_i to q_i, and let π be the corresponding plane in $\widetilde{M}$ obtained by sweeping γ_i from leaf to leaf of $\widetilde{\mathcal{F}}$. Then α_i stabilizes π, and quotients it out to give a cylinder C which maps to M. The hypothesis on α implies that $\alpha(\gamma_i)$ intersects γ_i transversely, and therefore π intersects $\alpha(\pi)$ in a line in $\widetilde{M}$. If we comb this intersection through $\widetilde{M}$ in the direction in which α_i^{-1} translates leaves, we see that the projection of this ray of intersection to C must stay in a compact portion of C. For otherwise, the translates of $\alpha(\gamma_i)$ under α_i^n would escape to an end of γ_i, which is incompatible with the dynamics of α_i. But if this ray of intersection of C with itself stays in a compact portion of C, it follows that it is *periodic* — that is, the line $\pi \cap \alpha(\pi)$ is stabilized by some power of α_i. For, there is a compact sub-cylinder $C' \subset C$ containing the preimage of the projection of the line of intersection. C' maps properly to M, and therefore its self-intersections are compact. The image of the ray in question is therefore compact and has at most one boundary component. In particular, it must be a circle, implying periodicity in π.

This implies that

$$\alpha\alpha_i^m\alpha^{-1} = \alpha_i^n$$

for some n, m. The co-orientability of $\mathcal{F}$ implies that n, m can both be chosen to be positive. It follows that α permutes the fixed points of α_i. But this is true *for all sufficiently large i*. The definition of the collection $\{\alpha_i\}$ implies that the only fixed points of α_i are in arbitrarily small neighborhoods of p, q, r, s, for sufficiently large i. It follows that α permutes p, q, r, s and that these are the only fixed points of any α_i. Since $\alpha(\gamma)$ intersects γ transversely, it follows that α permutes $\{p, q\}$ and $\{r, s\}$. But this means that it permutes an attracting point of α^n with a repelling point of α^m, which is absurd.

Observe that the roles of p,q and r,s are interchanged by replacing the α_i by α_i^{-1}, so no translate of γ' intersects γ' either, and the closure of its translates is an invariant lamination too. □

Corollary 5.3.8 *Let M be a 3–manifold with an $\mathbb{R}$–covered foliation $\mathcal{F}$. Then either M is Seifert fibered or solv, or there is a genuine lamination Λ of M transverse to $\mathcal{F}$.*

Proof If the action of $\pi_1(M)$ on S^1_{univ} is Fuchsian, then M is either solv or Seifert fibered by a standard argument (see eg [29]). Otherwise lemma 5.3.6 and lemma 5.3.7 produce Λ. □

Corollary 5.3.9 *If M is atoroidal and admits an $\mathbb{R}$–covered foliation, then $\pi_1(M)$ is δ–hyperbolic in the sense of Gromov.*

Proof This follows from the existence of a genuine lamination in M, by the main result of [19]. □

We analyze now how the hypothesis of atoroidality of M constrains the topology of the lamination Λ transverse to $\mathcal{F}$.

Lee Mosher makes the following definition in [27]:

Definition 5.3.10 A genuine lamination of a 3–manifold is *very full* if the complementary regions are all finite-sided ideal polygon bundles over S^1. Put another way, the gut regions are all sutured solid tori with the sutures a finite family of parallel curves nontrivially intersecting the meridian.

Lemma 5.3.11 *If M is atoroidal, the lamination Λ is very full, and the complementary regions to Λ_{univ} are all finite sided ideal polygons. Otherwise, there exist reducing tori transverse to $\mathcal{F}$ which are* regulating*. M can be split along such tori to produce simpler manifolds with boundary tori, inheriting taut foliations which are also $\mathbb{R}$–covered.*

Proof Let G be a gut region complementary to Λ, and let A_i be the collection of interstitial annuli, which are subsets of the boundary of G. Let $\widetilde{G}$ be a lift of G to $\widetilde{M}$ and $\widetilde{A}_i$ a collection of lifts of the A_i compatible with $\widetilde{G}$. Let α_i be the element of $\pi_1(M)$ stabilizing $\widetilde{A}_i$, so that $\widetilde{A}_i/\alpha_i = A_i$.

The first observation is that the interstitial annuli A_i can be straightened to be transverse to $\mathcal{F}$. Firstly, we can find a core curve $a_i \subset A_i$ and straighten A_i

leafwise so that $A_i = a_i \times I$ where each I is contained in a leaf of $\mathcal{F}$. Then, we can successively push the critical points of a_i into leaves of $\mathcal{F}$. One might think that there is a danger that the kinks of a_i might get "caught" on something as we try to push them into a leaf; but this is not possible for an $\mathbb{R}$–covered foliation, since obviously there is no obstruction in $\widetilde{M}$ to doing so, and since the lamination Λ is transverse to $\mathcal{F}$, we can "slide" the kinks along leaves of Λ whenever they run into them. The only danger is that the curves $\widetilde{a_i}$ might be "knotted", and therefore that we might change crossings when we straighten kinks. But a_i is isotopic into each of the boundary curves of A_i, and these lift to embedded lines in leaves of $\widetilde{\Lambda}$ which are properly embedded planes. It follows that the $\widetilde{a_i}$ are not knotted, and kinks can be eliminated.

Now, the boundary of a gut region is a compact surface transverse to $\mathcal{F}$. It follows that it has Euler characteristic 0, and is therefore either a torus or Klein bottle. By our orientability/co-orientability assumption, the boundary of a gut region is a torus. If M is atoroidal, this torus must be inessential and bounds a solid torus in M (because the longitude of this torus is non-trivial in $\pi_1(M)$). One quickly sees that this solid torus is exactly G, and therefore Λ is very full.

One observes that a pair of leaves λ, μ of Λ which have an interstitial annulus running between them must correspond to geodesics in Λ_{univ} which run into a "cusp" in S^1_{univ} — ie, they have the same endpoint in S^1_{univ}. For, by the definition of an interstitial region, the leaves λ, μ stay very close away from the guts, whereas if the corresponding leaves of Λ_{univ} do not have the same endpoint, they eventually diverge in any leaf, and one can find points in the interstitial regions arbitrarily far from either λ or μ, which is absurd. It follows that the annuli A_i are *regulating*, and each lift of a gut region of Λ corresponds to a finite sided ideal polygon in S^1_{univ}.

Conversely, if the boundary of some gut region is an *essential* torus, it can be pieced together from regulating annuli and regulating strips of leaves, showing that this torus is itself regulating. It follows that we can decompose M along such regulating tori to produce a taut foliation of a (possibly disconnected) manifold with torus boundary which is also $\mathbb{R}$–covered. □

Corollary 5.3.12 *If M admits an $\mathbb{R}$–covered foliation $\mathcal{F}$ then any homeomorphism $h\colon M \to M$ homotopic to the identity is isotopic to the identity.*

Proof This follows from the existence of a very full genuine lamination in M, by the main result of [18]. □

Theorem 5.3.13 *Let $\mathcal{F}$ be an $\mathbb{R}$–covered foliation of an atoroidal manifold M. Then there are a pair $\Lambda^{\pm}$ of essential laminations in M with the following properties:*

- *The complementary regions to $\Lambda^{\pm}$ are ideal polygon bundles over S^1.*
- *Each $\Lambda^{\pm}$ is transverse to $\mathcal{F}$ and intersects $\mathcal{F}$ in geodesics.*
- *Λ^+ and Λ^- are transverse to each other, and bind each leaf of $\mathcal{F}$, in the sense that in the universal cover, they decompose each leaf into a union of compact finite-sided polygons.*

If M is not *atoroidal but $\mathcal{F}$ has hyperbolic leaves, there is a regulating essential torus transverse to $\mathcal{F}$.*

Proof We have already shown the existence of at least one lamination Λ^+_{univ} giving rise to a very full lamination Λ^+ of M with the requisite properties, and we know that it is defined as the closure of the translates of some geodesic γ, which is the limit of a sequence of 4–tuples R_i with modulus $\to 0$ for which there are α_i so that $\text{mod}(\alpha_i(R_i)) \to \infty$ and $\alpha_i(R_i) \to \tau$. In fact, by passing to a minimal sublamination, we may assume that γ is a boundary leaf of Λ_{univ}, so that there are a sequence γ_i of leaves of Λ_{univ} converging to γ.

Fix a leaf λ of $\widetilde{\mathcal{F}}$ and an identification of $S^1_\infty(\lambda)$ with S^1_{univ}. Now, an element $\alpha_i \in \pi_1(M)$ acts on a 4–tuple R_i in S^1_{univ} in the following manner; let $Q_i \subset \lambda$ be the ideal quadrilateral with vertices corresponding to R_i. Then there is a unique ideal quadrilateral $Q_i' \subset \alpha_i^{-1}(\lambda)$ whose vertices project to the elements of R_i in S^1_{univ}. The element α_i translates Q_i' isometrically into λ, where its vertices are a 4–tuple of points in $S^1_\infty(\lambda)$ which determines $\alpha_i(R_i)$ in S_4. By definition, the moduli of the Q_i converge to 0, and the moduli of the Q_i' converge to ∞. The possibilities for the moduli of $\beta(R_i)$ as β ranges over $\pi_1(M)$ are constrained to be a subset of the moduli of the ideal quadrilaterals Q_i' obtained by sweeping Q_i through $\widetilde{M}$.

Let P be an ideal polygon which is a complementary region to Λ^+_{univ}, corresponding to a lift of a gut region G of Λ^+. $\widetilde{G}$ is foliated by ideal polygons in leaves of $\widetilde{\mathcal{F}}$. As we sweep through this family of ideal polygons in $\widetilde{G}$, the moduli of the polygons P_λ in each leaf λ corresponding to P stay bounded, since they cover a compact family of such polygons in M. Let α be an element of $\pi_1(M)$ stabilizing $\widetilde{G}$. Then after possibly replacing α with some finite power, α acts on S^1_{univ} by fixing P pointwise, and corresponds to the action on $S^1_\infty(\lambda)$ defined by sweeping through the circles at infinity from λ to $\alpha(\lambda)$ and then translating back by α^{-1}. Without loss of generality, γ is an edge of P. We label the endpoints of γ in S^1_{univ} as p, q. Note that p, q are fixed points of α.

A careful analysis of the combinatorics of the action of α and the α_i on S^1_{univ} will reveal the required structure.

We have quadrilaterals $Q_i \subset \lambda$ corresponding to the sequence R_i, and the vertices of these quadrilaterals converge in pairs to the geodesic γ_λ in λ corresponding to γ. Suppose there are fixed points m, n, r, s of α so that p, m, n, q, r, s are cyclically ordered. Then the moduli of all quadrilaterals Q_i' obtained by sweeping Q_i through $\widetilde{M}$, for i sufficiently large, are uniformly bounded. For, there is an ideal hexagon bundle in M corresponding to p, m, n, q, r, s and the moduli of these hexagons are bounded, by compactness. The pattern of separation of the vertices of this hexagon with R_i implies the bound on the moduli of the Q_i. It follows that there is at most one fixed point of α between p, q on some side. See figure 6a.

If there is *no* fixed point of α between p and q on one side, then α acts as a translation on the interval between p and q on that side. Obviously, the side of γ containing no fixed points of α must lie outside P, since the other vertices of P are fixed by α. It follows that the γ_i are on the side on which α acts as a translation. But this implies that for sufficiently large i, $\alpha(\gamma_i)$ crosses γ_i, which is absurd since the γ_i are leaves of an invariant lamination. Hence there is exactly one fixed point of α on one side of γ, and this point must be attracting for either α or α^{-1}. See figure 6b.

It follows that we have shown *in each complementary interval of the vertices of P, there is exactly one fixed point of α which is attracting for either α or α^{-1}.*

The same argument actually implies that Λ_{univ} was already minimal, since otherwise for γ' a leaf of Λ_{univ} which is a diagonal of P, the modulus of any sequence of 4–tuples converging to γ' is bounded under the image of powers of α, and therefore under the image of all elements of $\pi_1(M)$. This would contradict the definition of Λ_{univ}. Likewise, τ cannot be a diagonal of Λ_{univ}, since again the dynamics of α would imply that for any sequence of 4–tuples $R_i \to \tau$, the modulus of translates of R_i by any element of $\pi_1(M)$ would be bounded. It follows that if no translate of τ crosses any translate of γ, then the closure of the union of translates of τ is exactly equal to Λ_{univ}.

To summarize, we have established the following facts:

- Λ_{univ} is minimal.
- Either τ may be chosen transverse to γ, so that we are in the third alternative of lemma 5.3.6 and lemma 5.3.7 applies, or else the closure of the union of the translates of τ is equal to Λ_{univ}.

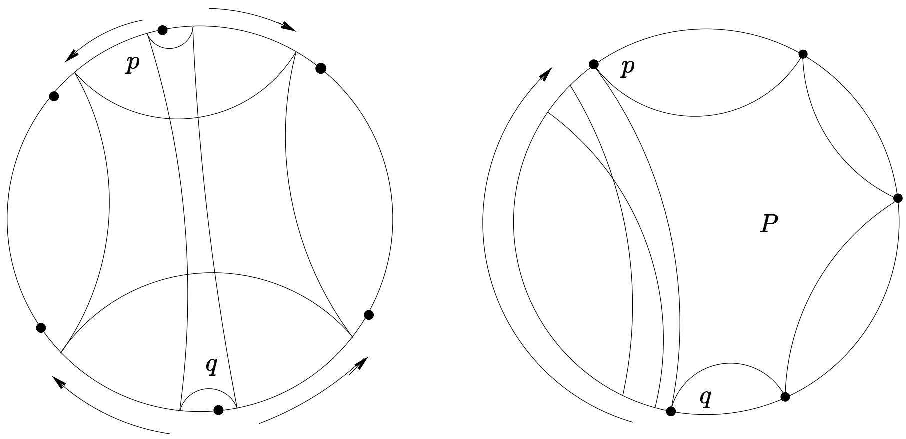

Figure 6: p, q are the vertices of a boundary leaf γ of P. If α has at least two fixed points on either side of p, q, the moduli of rectangles nested between these fixed points are bounded *above* by the action of $\pi_1(M)$. If α has *no* fixed points on some side of p, q, the fact that γ is not isolated on one side says that some nearby geodesic γ_i intersects its translate under α. The solid dots in the figure are fixed points of α. The arrows indicate the dynamics of α.

In fact, we will see that the fixed points of α in the complementary intervals to the vertices of P are *all* attracting points for α or for α^{-1}. For, suppose otherwise, so that there are consecutive vertices p, q, r of P and between them, points s, t which are repelling and attracting fixed points of α respectively, so that p, s, q, t, r are circularly ordered. Let γ' be the geodesic from s to q. Choose $s_i \to s$ from the side between s and q, and $q_i \to q$ from the side between q and t. Then $R_i' = \{s, s_i, q, q_i\}$ is a sequence of 4–tuples with $\text{mod}(R_i') \to 0$ and $R_i' \to \gamma'$ so that there are n_i with $\text{mod}(\alpha^{n_i}(R_i')) \to \infty$. It follows that there is a minimal lamination Λ'_{univ} constructed in exactly the same manner as Λ_{univ} which contains γ' as a leaf. Observe that α acts as a translation on the interval of S^1_{univ} from s to q, so that γ' must be the boundary of some complementary region P' of Λ'_{univ}. But then the core α' of the gut region in the complement of Λ' corresponding to P' is isotopic into the cylinder obtained by suspending γ', as is α, so in fact α and α' are freely isotopic, and correspond to the same element of $\pi_1(M)$ in our lift. It follows that α can have only one fixed point on the other side of γ', contradicting the fact that it fixes p and r there.

The end result of this fixed-point chase is that the fixed points of α in the complementary intervals to the vertices of P are all attracting fixed points for α (say) and therefore the vertices of P are all repelling fixed points of α.

Since Λ_{univ} is minimal, we can find β_i taking γ very close to γ_i. For γ_i

sufficiently close to γ, there is not much room for the image of P under β_i; on the other hand, the modulus of $\beta(P)$ cannot be distorted too much, since it varies in a compact family. Hence all the vertices of P but one are carried very close to one endpoint of γ_i. We can find a 4–tuple R_i'' with modulus close to 0 and vertices close to the endpoints of γ_i so that $\mathrm{mod}(\alpha^n\beta_i^{-1}(R_i'')) \to \infty$ and this sequence of rectangles converges to one of the geodesics joining s to an adjacent fixed point of α, which fixed point depending on which vertices of P are taken close to each other. It follows that for a sequence n_i growing sufficiently quickly, the sequence of rectangles R_i'' and the sequence $\alpha^{n_i}\beta_i^{-1}(R_i'')$ have moduli going to 0 and to ∞ respectively, and converge to a pair of transverse geodesics.

This establishes that we are in the third alternative of lemma 5.3.6, and therefore lemma 5.3.7 applies. That is, there are *two* laminations $\Lambda^{\pm}_{\mathrm{univ}}$ which are minimal, and transverse to each other, and these two laminations are exactly the closure of the union of the translates of γ and of τ respectively. Every complementary region to either lamination is finite sided, and therefore every complementary region to the union of these laminations is finite sided. To show that these laminations bind every leaf (ie, these finite sided regions are *compact*), it suffices to show that for p a vertex of a complementary region to $\Lambda^{+}_{\mathrm{univ}}$, say, there is a sequence of leaves in $\Lambda^{-}_{\mathrm{univ}}$ which nest down around p. This is actually an easy consequence of minimality of $\Lambda^{\pm}$, the fact that they are transverse, and the fact that M is compact. For completeness, and because it is useful in the sequel, we prove this statement as lemma 5.3.14. $\square$

Lemma 5.3.14 *Let $p \in S^1_{univ}$ be arbitrary. Then there is a sequence λ_i of leaves in either Λ^{+}_{univ} or Λ^{-}_{univ} which nest down around p.*

Proof Since both $\Lambda^{\pm}$ are minimal and transverse, it follows that there is a uniform t such that any leafwise geodesic γ contained in $\mu \cap \widetilde{\Lambda}^{+}$, for some leaf μ in $\widetilde{\mathcal{F}}$, must intersect a leaf of $\mu \cap \widetilde{\Lambda}^{-}$ with a definite angle within every subinterval of length t. It follows that these intersections determine leaves of $\Lambda^{-}_{\mathrm{univ}}$ which nest down to the point in S^1_{univ} corresponding to the endpoint of γ. It follows that endpoints of leaves of $\Lambda^{\pm}_{u}$ enjoy the property required by the lemma.

Now, there is a uniform t so that if $\gamma \subset \mu$ is an arbitrary geodesic, it intersects some leaf of $\widetilde{\Lambda}^{\pm} \cap \mu$ within every subinterval of length t, by the fact that $\widetilde{\Lambda}^{\pm} \cap \mu$ bind μ, and the compactness of M. There is a T and an ϵ such that every subinterval of length T must contain an intersection with angle bounded below by ϵ. For, if γ intersects $\widetilde{\Lambda}^{+}$ with a very small angle, it must stay close to a leaf of $\widetilde{\Lambda}^{+} \cap \mu$ for a long time, and therefore within a bounded time must intersect a leaf of $\widetilde{\Lambda}^- \cap \mu$ with a definite angle. It follows that some subsequence contained in either Λ^+_{univ} or Λ^-_{univ} must nest down to the point in S^1_{univ} corresponding to the endpoint of γ. Since γ was arbitrary, we are done. $\square$

Theorem 5.3.15 *Every $\alpha \in \pi_1(M)$ acts on S^1_{univ} in a manner either conjugate to an element of $PSL(2,\mathbb{R})$, or it is topologically pseudo-Anosov, or it has no fixed points and a finite power is topologically pseudo-Anosov.*

Proof Suppose α has non-isolated fixed points. Then either α is the identity, or it has a fixed point p which is a limit of fixed points on the left but not on the right. Let λ_i be a sequence of leaves of Λ^+_{univ} say nesting down to p. Then for some integer i, $\alpha^i(\lambda_j)$ intersects λ_j transversely, which is absurd. It follows that the fixed points of α are isolated. Again, the existence of a nesting sequence λ_i for every p implies that α must move all sufficiently close points on one side of p clockwise and on the other side, anticlockwise.

If α has no fixed points at all, either it is conjugate to a rotation, or some finite power has a fixed point and we can apply the analysis above. $\square$

Notice that for any topologically pseudo-Anosov α, the fixed points of α are alternately the vertices of a finite-sided complementary region to $\Lambda^+_{\text{univ}}, \Lambda^-_{\text{univ}}$ respectively.

In fact, we showed in theorem 5.3.13 that for α corresponding to the core of a lift $\widetilde{G}$ of a gut region G of Λ^+, the attracting fixed points of α are exactly the ideal vertices of the corresponding ideal polygon in S^1_{univ}, and the repelling fixed points are exactly the ideal vertices of a "dual" ideal polygon, corresponding to a lift of a gut region of Λ^-.

In [27], Lee Mosher defines a *topologically pseudo-Anosov* flow Ψ on a 3–manifold as, roughly speaking, a flow with weak stable and unstable foliations, singular along a collection of pseudohyperbolic orbits, and Ψ has a Markov partition which is "expansive". For the full definition one should consult [27], but the idea is that away from the (isolated) singular orbits, the manifold decomposes locally into a product $F \times E^s \times E^u$, where F corresponds to the flow-lines and E^s and E^u to the stable and unstable foliations, so that distances along the stable foliations are exponentially expanded under the flow, and distances along the unstable foliations are exponentially contracted under the flow. Mosher conjectures that every topological pseudo-Anosov flow on a closed 3–manifold should be smoothable — that is, there should exist a smooth structure on M with respect to which Φ is a smooth pseudo-Anosov (in the usual sense) flow.

Corollary 5.3.16 *An $\mathbb{R}$–covered foliation $\mathcal{F}$ admits a regulating transverse flow. If the ambient manifold M is atoroidal, this flow can be chosen to have isolated closed orbits. It can also be chosen to be "topologically pseudo-Anosov", as defined by Mosher in [27].*

Proof The laminations $\widetilde{\Lambda}^{\pm}$ bind every leaf of $\widetilde{\mathcal{F}}$, so we can canonically identify each leaf λ with each other leaf μ complementary region by region, where any canonical parameterization of a finite-sided hyperbolic polygon will suffice. For instance, the sided can be parameterized by arclength, and then coned off to the center of mass.

Alternatively, the method of [27] can be used to "blow down" $\widetilde{M}$ and therefore M to the lines $\widetilde{\Lambda}^+ \cap \widetilde{\Lambda}^-$. The flow along these lines descends to a flow on the blown down $\widetilde{M}$ where it is manifestly topologically pseudo-Anosov. More precisely, we can collapse, leafwise, intervals and polygons of the stratification of each leaf by its intersection with $\widetilde{\Lambda}^{\pm}$ to their boundary vertices. To see that this does not affect the homeomorphism type of M, choose a fine open cover of the blown-down manifold by open balls (such that the nerve of the cover gives a triangulation of M), and observe that its preimage gives a fine open cover of M with the same combinatorics. Theorem 5.3.15 implies that the flow so constructed satisfies the properties demanded by Mosher. To get a constant rate of expansion and contraction, pick an arbitrary metric on M and look at the expansion and contraction factors of the time t flow of an arbitrary segment σ in a leaf γ of $\widetilde{\Lambda}^+ \cap \lambda$ for some leaf λ of $\widetilde{\mathcal{F}}$, say. By construction, there are a sequence of rectangles R_i with moduli converging to 0 which nest down along γ, such that under the time t flow the moduli of the rectangles $\phi_t(R_i)$ converge to ∞. On can see from this pictures that the length of σ will shrink by a definite amount under the time t flow for some fixed t. The minimality of Λ^+ implies that the same is true for an arbitrary segment. By the usual argument, the expanding dynamics implies this flow is ergodic, and therefore the rate of expansion/contraction is asymptotically constant. One can therefore fix up the metric infinitesimally in the stable and unstable directions by looking at the asymptotic behavior, to get a rate of expansion and contraction bounded away from 1. By reparameterizing the metric in the flow direction, we can make this rate of expansion/contraction constant.

If M has a torus decomposition, but $\mathcal{F}$ has hyperbolic leaves, we have seen that the tori can be chosen to be transverse and regulating, and therefore inductively split along, and the flow found on simpler pieces.

If $\mathcal{F}$ has Euclidean or spherical leaves, it admits a transverse measure; any flow transverse to a transversely measured foliation is regulating. □

Remark 5.3.17 It is not too hard to see that all the results of this section can be made to apply to 3–manifolds with torus boundary and $\mathbb{R}$–covered foliations with hyperbolic leaves which intersect this boundary transversely. The laminations $\Lambda^{\pm}$ obtained will not necessarily have solid torus guts: they will also include components which are torus$\times I$ neighborhoods of the boundary tori. The main point is that the laminations $\Lambda^{\pm}_{\mathrm{univ}}$ of S^1_{univ} will still have *cusps*, so that they can be canonically completed to laminations with finite sided complements by adding new leaves which spiral around the boundary torus.

Remark 5.3.18 In [17], Gabai poses the general problem of studying when 3–manifold group actions on order trees "come from" essential laminations in the manifold. He further suggests that an interesting case to study is the one in which the order tree in question is $\mathbb{R}$. The previous theorem, together with the structure theorems from earlier sections, provide a collection of non-trivial conditions that an action of $\pi_1(M)$ on $\mathbb{R}$ must satisfy to have come from an action on the leaf space of a foliation. We consider it a very interesting question to formulate (even conjecturally) a list of properties that a good "realization theorem" should require. We propose the following related questions as being perhaps more accessible:

Fix an $\mathbb{R}$–covered foliation of M and consider the associated action of $\pi_1(M)$ on $\mathbb{R}$, the leaf space of the foliation in the universal cover.

- Is this action conjugate to a Lipschitz action?
- Are leaves in the foliation $\widetilde{\mathcal{F}}$ at most exponentially distorted?
- Is the pseudo-Anosov flow transverse to an $\mathbb{R}$–covered foliation of an atoroidal 3–manifold quasi-geodesic? That is, are the flowlines of the lift of the transverse regulating pseudo-Anosov flow to $\widetilde{M}$ quasigeodesically embedded?

We remark that the construction in [2] allows us to embed any finitely generated subgroup of $\widetilde{Homeo}(S^1)$ in the image of $\pi_1(M)$ in $Homeo(\mathbb{R})$ for some $\mathbb{R}$–covered foliation. In fact, we can take any finite collection of irrationally related numbers $t_1, \ldots t_n$, any collection of finitely generated subgroups of $\widetilde{Homeo}(S^1_{t_i})$ — the group of homeomorphisms of $\mathbb{R}$ which are periodic with period t_i — and consider the group they all generate in $Homeo(\mathbb{R})$. Then this group can be embedded in the image of $\pi_1(M)$ in $Homeo(\mathbb{R})$ for some $\mathbb{R}$–covered foliation of M, for some M. Probably M can be chosen in each case to be hyperbolic, by the method of [2], but we have not checked all the details of this.

It seems difficult to imagine, but perhaps all $\mathbb{R}$–covered foliations of atoroidal manifolds are at worst "mildly" nonuniform, in this sense. We state this as a

Question 5.3.19 *If $\mathcal{F}$ is an $\mathbb{R}$–covered foliation of an atoroidal 3–manifold M, is there a choice of parameterization of the leaf space of $\widetilde{\mathcal{F}}$ as $\mathbb{R}$ so that $\pi_1(M)$ acts on this leaf space by* coarse 1–quasi-isometries*? That is, is there a k_α for each α such that, for every $p, q \in \mathbb{R}$, there is an inequality*

$$\alpha(p) - \alpha(q) - k_\alpha \le p - q \le \alpha(p) - \alpha(q) + k_\alpha$$

Remark 5.3.20 A regulating vector field integrates to a 1–dimensional foliation which lifts in the universal cover to the product foliation of $\mathbb{R}^3$ by vertical copies of $\mathbb{R}$. Such a foliation is called *product covered* in [10] where they are used to study the question of when an immersed surface is a virtual fiber. It is tautological from the definition of a product covered foliation that there is an associated slithering of M over $\mathbb{R}^2$. One may ask about the quality of the associated representation $\pi_1(M) \to Homeo(\mathbb{R}^2)$.

Definition 5.3.21 A *family* of $\mathbb{R}$–covered foliations on a manifold M indexed by the unit interval I is a choice of 2–plane field D_t for each $t \in I$ such that each D_t is integrable, and integrates to an $\mathbb{R}$–covered foliation $\mathcal{F}_t$, and such that $D_t(p)$ for any fixed p varies continuously with t.

Notice that the local product structure on $\mathcal{F}_t$ in a small ball varies continuously. That is, for any sufficiently small ball B there is a 1–parameter family of isotopies $i_t\colon B \to M$ such that $i_t^*(\mathcal{F}_t)|_{i_t(B)} = \mathcal{F}_0|_B$. In particular, a family of foliations on M is a special kind of foliation on $M \times I$.

Corollary 5.3.22 *Let $\mathcal{F}_t$ be a family of $\mathbb{R}$–covered foliations of an atoroidal M. Then the action of $\pi_1(M)$ on $(S^1_{univ})_t$ is independent of t, up to conjugacy. Moreover, the laminations $\Lambda_t^\pm$ do not depend on the parameter t, up to isotopy.*

Proof Let Λ_t be one of the two canonical geodesic laminations constructed from $\mathcal{F}_t$ in theorem 5.3.13. For s, t close enough, Λ_t intersects $\mathcal{F}_s$ quasigeodesically. For, in $\mathbb{H}^2$, quasigeodesity is a local property; that is, a line in $\mathbb{H}^2$ is quasigeodesic provided the subsets of the line of some fixed length are sufficiently close to being geodesic. For s sufficiently close to t, the lines of intersection $\Lambda_t \cap \mathcal{F}_s$ are very close to being geodesic, so are quasigeodesic.

The only subtlety is that we need to know that we can choose uniformizing metrics on M so that leaves of $\mathcal{F}_t$ are hyperbolic for each t in such a way that the metrics vary continuously in t. Candel's theorem in full generality says that we can do this; for, we can consider the foliation $\mathcal{F}_I$ of $M \times I$ whose leaves are

$$\mathcal{F}_I = \bigcup_{\lambda \in \mathcal{F}_t, t \in I} \lambda \times t.$$

This is a foliation of a compact manifold with Riemann surface leaves and no invariant transverse measure of non-negative Euler characteristic, so Candel's theorem 2.1.3 applies.

It follows by theorem 5.2.5 that Λ_t comes from an invariant lamination of $(S^1_{\text{univ}})_s$. This gives a canonical, equivariant identification of $(S^1_{\text{univ}})_s$ and $(S^1_{\text{univ}})_t$ as follows: for a dense set of points $p \in (S^1_{\text{univ}})_t$ and each leaf λ of $\widetilde{\mathcal{F}}_t$ there is a leaf μ of Λ_t which intersects λ in a geodesic g, one of whose endpoints projects to p under $(\phi_h)_t$. For a leaf λ' of $\widetilde{\mathcal{F}}_s$ which contains some point of $\lambda \cap \mu$, the intersection $\lambda' \cap \mu$ is a quasigeodesic which can be straightened to a geodesic g' with the same endpoints. By choosing an orientation on μ and continuously varying orientations on λ and λ', the geodesics g and g' are oriented, so we know which of the endpoints to choose in $S^1_\infty(\lambda')$. Projecting to $(S^1_{\text{univ}})_s$ by $(\phi_h)_s$, we get a point p'. Since the structure of Λ_t in $(C_\infty)_s$ and $(C_\infty)_t$ are vertical, this construction does not depend on any choices.

Extending by continuity we get a canonical, and therefore π_1–invariant identification of $(S^1_{\text{univ}})_s$ and $(S^1_{\text{univ}})_t$. Since the laminations $\Lambda^\pm_t$ are canonically constructed from the action of $\pi_1(M)$ on the universal circle of $\mathcal{F}_t$, the fact that these actions are all conjugate implies that the laminations too are invariant. $\square$

Remark 5.3.23 Thurston has a program to construct a universal circle and a pair of transverse laminations intersecting leaves geodesically for *any* taut foliation of an atoroidal M; see [33]. In [6] we produce a pair of genuine laminations $\Lambda^\pm$ transverse to an arbitrary *minimal* taut foliation of an atoroidal M.

If an $\mathbb{R}$–covered foliation is perturbed to a non–$\mathbb{R}$–covered foliation, nevertheless this lamination stays transverse for small perturbations, and therefore the action of $\pi_1(M)$ on the universal circle of the taut foliation is the same as the action on S^1_{univ} of the $\mathbb{R}$–covered foliation. This may give a criteria for an $\mathbb{R}$–covered foliation to be a limit of non–$\mathbb{R}$–covered foliations.

One wonders whether *every* taut foliation of an atoroidal manifold M is homotopic, as a 2–plane field, to an $\mathbb{R}$–covered foliation.

Remark 5.3.24 As remarked in the introduction, Sérgio Fenley has proved many of the results in this section independently, by somewhat different methods, using the canonical product structure on C_∞ constructed in theorem 4.6.4.

5.4 Are $\mathbb{R}$–covered foliations geometric?

In 1996, W. Thurston outlined an ambitious and far-reaching program to prove that 3–manifolds admitting taut foliations are geometric. Speaking very vaguely, the idea is to duplicate the proof of geometrization for Haken manifolds as outlined in [29],[30] and [34] by developing the analogue of a quasi-Fuchsian deformation theory for leaves of such a foliation, and by setting up a dynamical system on such a deformation space which would find a hyperbolic structure on the foliated manifold, or find a topological obstruction if one existed.

This paper may be seen as foundational to such a program for geometrizing $\mathbb{R}$–covered foliations. In [12] it is shown that for $\mathbb{R}$–covered foliations of Gromov-hyperbolic 3–manifolds, leaves in the universal cover limit to the entire sphere at infinity. This is evidence that $\mathbb{R}$–covered foliations behave geometrically somewhat like surface bundles over circles. This suggests the following strategy, obviously modeled after [34]:

- Pick a leaf λ in $\widetilde{\mathcal{F}}$, and an element $\alpha \in \pi_1(M)$ which acts on L without fixed points. Then the images $\alpha^n(\lambda)$ for $-\infty < n < \infty$ go off to infinity in L in either direction.
- We can glue λ to $\alpha^n(\lambda)$ along their mutual circles at infinity by the identification of either with S^1_{univ} to get a topological S^2. We would like to "uniformize" this S^2 to get $\mathbb{C}P^1$.
- Let X be a regulating transverse vector field. This determines a map ϕ_n from λ to $\alpha^n(\lambda)$ by identifying points which lie on the same integral curve of $\widetilde{X}$.
- The map ϕ_n is uniformly quasi-isometric on regions where λ and $\alpha^n(\lambda)$ are close, but cannot be guaranteed to be uniformly quasi-isometric on *all* of λ, and probably is not so. By comparing the conformal structure on λ and $\alpha^n(\lambda)$ we get a Beltrami differential $\mu_n \frac{d\bar{z}}{dz}$ which is not necessarily in $B(\mathbb{H})_1$. Nevertheless, the fact that λ and $\alpha^n(\lambda)$ are asymptotic at infinity in almost every direction encourages one to hope that one has enough geometric control to construct a uniformizing homeomorphism of S^2 to $\mathbb{C}P^1$ with prescribed Beltrami differential.
- Taking a sequence of such uniformizing maps corresponding to differentials μ_n with $n \to \infty$ one hopes to show that there is a convergence $S^1_{\text{univ}} \to S^2$ geometrically. Then the action of $\pi_1(M)$ on S^1_{univ} will extend to S^2 since the map $S^1_{\text{univ}} \to S^2$ is canonical and therefore $\pi_1(M)$–equivariant. Does this action give a representation in $PSL(2,\mathbb{C})$?

- Group-theoretically, we can use X to let $\pi_1(M)$ act on any given leaf λ of $\widetilde{\mathcal{F}}$. $\pi_1(M)$ therefore acts on $\lambda \cup \alpha^n(\lambda)$ and so on $\mathbb{C}P^1$. We can use the barycentric extension map of Douady and Earle to extend this to a map of $\mathbb{H}^3$ to itself. We hope that some of the powerful technology developed by McMullen in [26] can be used to show that this action is nearly isometric deep in the convex hull of the image of S^1_{univ}, and perhaps a genuine isometric action can be extracted in the limit.

We stress that this outline borrows heavily from Thurston's strategy to prove that manifolds admitting *uniform* foliations are geometric, as communicated to the author in several private communications. In fact, the hope that one might generalize this strategy to $\mathbb{R}$–covered foliations was our original motivation for undertaking this research, and it has obviously greatly influenced our choice of subject and approach.

References

[1] **J Cantwell**, **L Conlon**, *Leafwise hyperbolicity of proper foliations*, Comm. Math. Helv. 64 (1989) 329–337

[2] **D Calegari**, *$\mathbb{R}$–covered foliations of hyperbolic 3–manifolds*, Geometry and Topology, 3 (1999) 137–153

[3] **D Calegari**, *The geometry and topology of $\mathbb{R}$–covered foliations*, Elec. Res. Ann. of the AMS, 6 (2000) 31–39

[4] **D Calegari**, *Foliations with one-sided branching*, preprint

[5] **D Calegari**, *Foliations and the geometry of 3–manifolds*, PhD dissertation at UC Berkeley, May 2000

[6] **D Calegari**, *Promoting essential laminations*, in preparation

[7] **A Candel**, *Uniformization of surface laminations*, Ann. Sci. Ecole Norm. Sup. 26 (1993) 489–516

[8] **A Casson**, **D Jungreis**, *Convergence groups and Seifert fibered 3–manifolds*, Invent. Math. 118 (1994) 441–456

[9] **A Connes**, *Non-commutative Geometry*, Academic Press (1994)

[10] **D Cooper**, **D Long**, *Foliations of some 3–manifolds which fiber over the circle* Proc. Amer. Math. Soc. 126 (1998) 925–931.

[11] **B Farb**, **R Schwartz**, *The large-scale geometry of Hilbert modular groups*, J. Diff. Geom. 44 (1996) 435–478

[12] **S Fenley**, *Quasi-isometric foliations*, Topology, 31 (1992) 667–676

[13] **S Fenley**, *Anosov flows in 3–manifolds*, Ann. Math. 139 (1994) 79–115

[14] **S Fenley**, personal communication

[15] **D Gabai**, *Foliations and the topology of 3–manifolds*, J. Diff. Geom. 18 (1983) 445–503

[16] **D Gabai**, *Convergence groups are Fuchsian groups*, J. Diff. Geom. 36 (1992) 269–313

[17] **D Gabai**, *Problems in foliations and laminations*, from: "Geometric Topology" (ed. W Kazez) proceedings of the 1993 Georgia International Topology Conference, Vol. 2, part 2, 1–33

[18] **D Gabai**, **W Kazez**, *Homotopy, isotopy and genuine laminations of 3–manifolds*, from: "Geometric Topology" (ed. W Kazez) proceedings of the 1993 Georgia International Topology Conference, Vol. 1, 123–138

[19] **D Gabai**, **W Kazez**, *Group negative curvature for 3–manifolds with genuine laminations*, Geometry and Topology, 2 (1998) 65–77

[20] **D Gabai**, **U Oertel**, *Essential laminations in 3–manifolds*, Ann. Math. 130 (1989) 41–73

[21] **L Garnett**, *Foliations, the ergodic theorem and Brownian motion*, Journal of Functional Analysis, 51 (1983) 285–311

[22] **M Gromov**, *Hyperbolic groups*, from: "Essays in Group Theory" (ed. S Gersten) MSRI publ. 8, Springer–Verlag (1987) 75–263

[23] **A Hinkkanen**, *Abelian and nondiscrete convergence groups on the circle*, Trans. Amer. Math. Soc. 318 (1990) 87–121

[24] **Y Imayoshi**, **M Taniguchi**, *An introduction to Teichmüller spaces*, Springer–Verlag (1992)

[25] **O Lehto**, *Univalent functions and Teichmüller spaces*, Springer–Verlag (1987)

[26] **C McMullen**, *Renormalization and 3–manifolds which fiber over the circle*, Ann. Math. Studies, Princeton Uni. Press, 142 (1996)

[27] **L Mosher**, *Laminations and flows transverse to finite depth foliations*, preprint

[28] **C Palmeira**, *Open manifolds foliated by planes*, Ann. Math. 107 (1978) 109–131

[29] **W Thurston**, *Geometry and topology of 3–manifolds*, Princeton lecture notes (1979)

[30] **W Thurston**, *Hyperbolic structures on 3–manifolds I: deformations of acylindrical manifolds*, Ann. Math. 124 (1986) 203–246

[31] **W Thurston**, *Foliations and circles*, Talk at the Pacific Northwest Geometry Seminar, Spring 1997 (available on video from the MSRI)

[32] **W Thurston**, *3–manifolds, foliations and circles I*, `arxiv:math.GT/9712268`

[33] **W Thurston**, *3–manifolds, foliations and circles II*, preprint

[34] **W Thurston**, *Hyperbolic structures on 3–manifolds II: surface groups and 3–manifolds which fiber over the circle*, `arxiv:math.GT/9801045`

[35] **W Thurston**, personal communication