Groups Geom. Dyn. (Online first)
DOI 10.4171/GGD/964

Depth-one foliations, pseudo-Anosov flows and universal circles

Junzhi Huang

Abstract. Given a taut depth-one foliation $\mathcal{F}$ in a closed atoroidal 3-manifold transverse to a pseudo-Anosov flow ϕ without perfect fits, we show that the universal circle coming from leftmost sections associated to $\mathcal{F}$, constructed by Thurston and Calegari–Dunfield, is isomorphic to Fenley's ideal boundary of the flow space associated to ϕ with natural structure maps. As a corollary, we use a theorem of Barthelmé–Frankel–Mann to show that there is at most one pseudo-Anosov flow without perfect fits transverse to $\mathcal{F}$ up to orbit equivalence.

1. Introduction

There has been an important theme in 3-manifold topology to study the interaction between flows and codimension-one foliations in 3-manifolds. The simplest examples of codimension-one foliations of a 3-manifold M are fibrations, which are exactly the foliations with all leaves compact. The theory of Thurston norm organizes different ways of fibration of M into a finite number of fibered faces, and there is a one-to-one correspondence between fibered faces and isotopy classes of suspension pseudo-Anosov flows [22, 31]. The aim of this paper is to study one of the next simplest classes of foliations, namely depth-one foliations, and their interaction with transverse pseudo-Anosov flows, by comparing the π_1-actions on S^1 that arise in both settings.

A foliation in a closed 3-manifold is called a *depth-one foliation* if the restriction to the complement of compact leaves is a fibration over the circle. More precisely, there are a finite number of compact leaves, called depth-zero leaves, and the rest of the leaves (namely the depth-one leaves) are infinite-type surfaces spiraling into the depth-zero leaves. One way to construct depth-one foliations is to "spin" a fibration around an embedded surface (see [6, Example 4.8]).

Given a taut depth-one foliation $\mathcal{F}$ in a closed 3-manifold M, a result of Candel [8] shows that there exists a Riemannian metric on M such that the restrictions to the leaves of $\mathcal{F}$ are hyperbolic, giving every $\mathcal{F}$-leaf a standard hyperbolic structure in the sense of [10] (see also Section 2.3). In particular, there is a natural circle at infinity associated to the universal circle of any leaf of $\mathcal{F}$. An unpublished construction of Thurston [32], which

Mathematics Subject Classification 2020: 57K30.
Keywords: pseudo-Anosov flows, foliations, universal circles, 3-manifolds.

was later written down by Calegari–Dunfield [7], produces a circle $\mathfrak{S}_{\text{left}}$ associated to $\mathcal{F}$. We will call this circle a *universal circle from leftmost sections*. The circle $\mathfrak{S}_{\text{left}}$ is acted on faithfully by $\pi_1(M)$ and is equipped with a $\pi_1(M)$-equivariant collection of monotone structure maps $\{U_\lambda\}_{\lambda \in \widetilde{\mathcal{F}}}$ to the circles at infinity of all $\widetilde{\mathcal{F}}$-leaves, where $\widetilde{\mathcal{F}}$ is the lift of $\mathcal{F}$ to the universal cover $\widetilde{M}$ of M.

In general, there is an axiomatized notion of a universal circle associated to a taut foliation (Definition 4.1). The universal circle $\mathfrak{S}_{\text{left}}$ is a universal circle of $\mathcal{F}$ in this general sense but not a canonical one. However, when $\mathcal{F}$ is a taut depth-one foliation transverse to a pseudo-Anosov flow without perfect fits ϕ (see Section 2.1 for discussions on pseudo-Anosov flows), we will see that it is possible to relate the $\mathfrak{S}_{\text{left}}$ to a more natural object, which is the ideal boundary of the flow space of ϕ.

For a pseudo-Anosov flow ϕ, the *flow space* $\mathcal{O}$ associated to ϕ is the space of orbits of the lifted flow $\widetilde{\phi}$ in $\widetilde{M}$. It is homeomorphic to $\mathbb{R}^2$ by [2, 14, 16], and there is a compactification $\overline{\mathcal{O}} = \mathcal{O} \cup \partial\mathcal{O}$ given by Fenley [13]. The ideal boundary $\partial\mathcal{O}$ is homeomorphic to a circle, and the $\pi_1(M)$-action on $\mathcal{O}$ extends continuously to $\partial\mathcal{O}$. If ϕ has no perfect fits, we will see that the shadow of any leaf λ of $\widetilde{\mathcal{F}}$ provides a natural structure map I_λ from $\partial\mathcal{O}$ to the circle at infinity of λ (see Section 3). For these structure maps, we prove the following theorem.

Theorem 1.1. *Let M be a closed atoroidal 3-manifold with a pseudo-Anosov flow ϕ without perfect fits, and let $\mathcal{F}$ be a taut depth-one foliation in M transverse to ϕ. Then the circle $\partial\mathcal{O}$, together with the structure maps $\{I_\lambda\}_{\lambda \in \widetilde{\mathcal{F}}}$, is a universal circle for $\mathcal{F}$.*

While writing this paper, the author learned that Landry, Minsky and Taylor show a much more general version of Theorem 1.1 (recently appeared in [24]). More precisely, they prove that given a taut foliation $\mathcal{F}$ almost transverse to a pseudo-Anosov flow ϕ in a closed hyperbolic manifold M, the boundary of the flow space naturally has the structure of a universal circle for $\mathcal{F}$.

Nevertheless, we show in our setting that the universal circles $\partial\mathcal{O}$ and $\mathfrak{S}_{\text{left}}$ are isomorphic, as made precise by the following theorem.

Theorem 1.2. *Let M be a closed atoroidal 3-manifold with a pseudo-Anosov flow ϕ without perfect fits, and let $\mathcal{F}$ be a taut depth-one foliation in M transverse to ϕ. Then the $\pi_1(M)$-actions on $\partial\mathcal{O}$ and on $\mathfrak{S}_{\text{left}}$ are conjugated by a homeomorphism $T : \mathfrak{S}_{\text{left}} \to \partial\mathcal{O}$. Moreover, for any leaf λ of $\widetilde{\mathcal{F}}$, we have $I_\lambda \circ T = U_\lambda$.*

Corollary 1.3. *Let M be a closed atoroidal 3-manifold, and let $\mathcal{F}$ be a taut depth-one foliation in M. Then there is at most one pseudo-Anosov flow without perfect fits transverse to $\mathcal{F}$ up to orbit equivalence.*

Proof. Suppose there are two pseudo-Anosov flows without perfect fits ϕ and φ that are transverse to $\mathcal{F}$. Since M is atoroidal, both ϕ and φ are transitive (we say a flow is transitive if it has an orbit that is dense in both positive and negative time) [27]. Since the construction of $\mathfrak{S}_{\text{left}}$ makes no use of the transverse flows, the actions of $\pi_1(M)$ on the ideal boundaries of the orbit spaces of ϕ and φ are conjugate by Theorem 1.2. By [3, Theorem 1.5], ϕ and φ are orbit equivalent. ∎

Remark 1.4. In personal conversations, Michael Landry told the author that using the construction by Gabai–Mosher for almost transverse pseudo-Anosov flows to finite depth foliations, one can construct different pseudo-Anosov flows transverse to the same depth-one foliation, but these flows have perfect fits. While there is currently no complete proof of Gabai–Mosher's construction in the literature, the monograph [28] by Mosher contains an outline and the main ideas of the theory. See also the paper [25] by Landry–Tsang and their upcoming work [26], which are aimed at revisiting the theory using veering triangulations.

Remark 1.5. In [30], Anna Parlak constructs examples of closed and cusped hyperbolic 3-manifolds with a non-fibered face dynamically represented by two topologically inequivalent pseudo-Anosov flows using mutations of veering triangulations. However, the properties of the resulting flows, for example, whether they have perfect fits or which foliations they are transverse to, are not so clear.

A conjectural picture of "pseudo-Anosov packages" is developed in [5] by Calegari in the hope that the different structures from taut foliations, laminations, universal circles and pseudo-Anosov flows are organized and compatible in the most natural way, and it is asked to what extent the picture is true.

In particular, given a universal circle $\mathfrak{S}$ for $\mathcal{F}$, he constructs a pair of $\pi_1(M)$-invariant laminations $\Xi^{\pm}$ on $\mathfrak{S}$. In our case, one can apply the construction to the universal circle $\mathfrak{S}_{\text{left}} \cong \partial\mathcal{O}$ and get a pair of laminations on $\partial\mathcal{O}$. On the other hand, the endpoints of the singular foliations $\mathcal{F}^{u/s}_{\mathcal{O}}$ also induce a pair of laminations $\mathcal{L}^{u/s}_{\mathcal{O}}$ on $\partial\mathcal{O}$ by taking the pairs of endpoints of regular leaves and faces of singular leaves. We partially verify Calegari's picture by showing that $\Xi^{\pm}$ and $\mathcal{L}^{u/s}_{\mathcal{O}}$ coincide.

Theorem 1.6. *In the setting of Theorem 1.2, the invariant lamination Ξ^+ (resp. Ξ^-) on the universal circle $\mathfrak{S}_{\text{left}}$ equals the induced stable lamination $\mathcal{L}^s_{\mathcal{O}}$ (resp. $\mathcal{L}^u_{\mathcal{O}}$) under the isomorphism $T : \mathfrak{S}_{\text{left}} \to \partial\mathcal{O}$.*

The organization of the paper is as follows. In Section 2, we briefly recall some knowledge about pseudo-Anosov flows, depth-one foliations and circle laminations. In Section 3, we summarize the structure of the shadows of leaves of $\mathcal{F}$ in $\mathcal{O}$ developed by [12, 19] and carefully study the infinity of shadows. From there, we introduce the restriction maps I_λ and a relative version of restriction maps. We start Section 4 with a brief review of the construction of the universal circle from leftmost sections $\mathfrak{S}_{\text{left}}$ following [7], and we relate $\mathfrak{S}_{\text{left}}$ to the universal circle structure of $\partial\mathcal{O}$ we developed in Section 3. We prove Theorem 1.2 in Section 5 by explicitly constructing the homeomorphism T and proving the desired properties. We conclude with a discussion of invariant laminations and the proof of Theorem 1.6 in Section 6.

2. Preliminaries

Convention 2.1. We consider a closed Riemannian atoroidal 3-manifold M with the Riemannian metric to be determined later. For a partition, Θ (e.g., a flow or a foliation) of a space X and a point $x \in X$, we let $\Theta(x)$ be the atom of Θ containing x. More generally, if A is a subset of C, we use $\Theta(A)$ to denote the saturation of A by Θ-atoms.

2.1. Pseudo-Anosov flows

We refer to [1, 14, 27] and the recent monograph [4] for detailed discussions of pseudo-Anosov flows in 3-manifolds. The following definition follows [13].

A flow $\phi : M \times \mathbb{R} \to M$ in M is a *pseudo-Anosov flow* if it has the following properties:

- each flowline is C^1 and not a single point;
- the tangent line bundle $T\phi$ is continuous;
- there are a finite number of singular closed orbits, and a pair of 2-dimensional singular foliations $\mathcal{F}^u$ and $\mathcal{F}^s$ in M so that
  - each leaf of $\mathcal{F}^u$ or $\mathcal{F}^s$ is a union of ϕ-orbits;
  - outside of the singular orbits $\mathcal{F}^u$ and $\mathcal{F}^s$ are regular foliations whose leaves intersect transversely along ϕ-orbits;
  - for each singular orbit ω, the leaf of $\mathcal{F}^u$ or $\mathcal{F}^s$ containing ω is homeomorphic to $P_n \times [0, 1]/f$ where
$$P_n = \{re^{2ki\pi/n} \mid r \geq 0, 0 \leq k \leq n-1\} \subset \mathbb{C}$$
    is an n-prong and f is a homeomorphism from P_n to P_n. The orbit ω is the image of $\{0\} \times [0, 1]$ and n is always greater than 2 in our case;
  - orbits in the same $\mathcal{F}^s$-leaf are forward asymptotic, and orbits in the same $\mathcal{F}^u$-leaf are backward asymptotic.

The singular foliations $\mathcal{F}^s$ and $\mathcal{F}^u$ are called the *stable foliation* and the *unstable foliation* of ϕ, respectively. When the set of singular orbits is empty, ϕ is an Anosov flow.

Fix a universal cover $\widetilde{M}$ of M, and let $\widetilde{\phi}$, $\widetilde{\mathcal{F}}^u$, $\widetilde{\mathcal{F}}^s$ be the lifts of ϕ, $\mathcal{F}^u$, $\mathcal{F}^s$ to $\widetilde{M}$, respectively. The quotient of $\widetilde{M}$ by $\widetilde{\phi}$ is the *flow space* $\mathcal{O}$, which is homeomorphic to $\mathbb{R}^2$ [2,14,16]. We orient $\mathcal{O}$ so that the coorientation coincides with the flow direction, and the pictures in this paper are drawn in a way that the flow is flowing toward the reader. The deck transformation on $\widetilde{M}$ descends to an orientation-preserving $\pi_1(M)$-action on $\mathcal{O}$, and the singular foliations $\widetilde{\mathcal{F}}^u$ and $\widetilde{\mathcal{F}}^s$ descend to a pair of $\pi_1(M)$-invariant transverse singular foliations on $\mathcal{O}$, denoted by $\mathcal{F}^u_{\mathcal{O}}$ and $\mathcal{F}^s_{\mathcal{O}}$. The singular leaves of $\mathcal{F}^s_{\mathcal{O}}$ and $\mathcal{F}^u_{\mathcal{O}}$ are n-pronged with $n \geq 3$. The union of two adjacent prongs in a singular leaf of $\mathcal{F}^s_{\mathcal{O}}$ or $\mathcal{F}^u_{\mathcal{O}}$ is called a *face*.

A $\widetilde{\phi}$-orbit ω is *periodic* if there exists a non-trivial deck transformation g of $\widetilde{M}$ with $g(\omega) = \omega$. In this case, g acts on ω by translation, and ω covers a closed orbit of ϕ in M.

Similarly, a leaf of $\widetilde{\mathcal{F}}^s$ or $\widetilde{\mathcal{F}}^u$ is called *periodic* if it is fixed by some non-trivial deck transformation of $\widetilde{M}$. In particular, singular leaves are periodic. Finally, we call a point or a leaf of $\mathcal{F}^u_{\mathcal{O}}$ or $\mathcal{F}^s_{\mathcal{O}}$ in $\mathcal{O}$ periodic if the corresponding $\widetilde{\phi}$-orbit or $\widetilde{\mathcal{F}}^{s/u}$ leaf in $\widetilde{M}$ is periodic. The following fact is well known and can be found in [4, Proposition 1.4.3].

Lemma 2.2. *If ℓ is a leaf of $\mathcal{F}^s_{\mathcal{O}}$ or $\mathcal{F}^u_{\mathcal{O}}$ fixed by a non-trivial g in $\pi_1(M)$, then the g-action on ℓ has a unique fixed point. In other words, any periodic leaf of $\widetilde{\mathcal{F}}^s$ or $\widetilde{\mathcal{F}}^u$ contains a unique periodic orbit.*

A *ray* of $\mathcal{F}^u_{\mathcal{O}}$ or $\mathcal{F}^s_{\mathcal{O}}$ is an embedded closed half-line contained in a leaf with the interior disjoint from singularities. Two rays $l \in \mathcal{F}^s_{\mathcal{O}}$ and $l' \in \mathcal{F}^u_{\mathcal{O}}$ are said to form a *perfect fit* if there is an (possibly orientation-reversing) embedding

$$\iota : [0,1] \times [0,1] - (1,1) \to \mathcal{O}$$

mapping horizontal lines to $\mathcal{F}^s_{\mathcal{O}}$ leaves, vertical lines to $\mathcal{F}^u_{\mathcal{O}}$ leaves, $[0,1) \times \{1\}$ to l' and $\{1\} \times [0,1)$ to l. We say a pseudo-Anosov flow is *without perfect fits* if no two rays in $\mathcal{F}^s_{\mathcal{O}}$ and $\mathcal{F}^u_{\mathcal{O}}$ form a perfect fit. The notion of perfect fits is introduced and studied by Fenley in [17,18]. In particular, we have the following lemma which is an immediate consequence of [18, Theorem 4.8].

Lemma 2.3. *If ϕ has no perfect fits, then any non-trivial element of $\pi_1(M)$ has at most one fixed point in $\mathcal{O}$.*

Fenley introduces a compactification of $\mathcal{O}$ in [13] by building an ideal boundary $\partial\mathcal{O}$ homeomorphic to S^1, and the resulting compactified space $\overline{\mathcal{O}} = \mathcal{O} \cup \partial\mathcal{O}$ is homeomorphic to a closed 2-disk. We orient $\partial\mathcal{O}$ as the boundary of $\mathcal{O}$, and the action of $\pi_1(M)$ on $\mathcal{O}$ extends continuously to an orientation-preserving action on $\overline{\mathcal{O}}$. Each ray in $\mathcal{F}^s_{\mathcal{O}}$ or $\mathcal{F}^u_{\mathcal{O}}$ has a well-defined endpoint, and the endpoints of every leaf are distinct. When the flow has no perfect fits, a ray in $\mathcal{F}^s_{\mathcal{O}}$ and a ray in $\mathcal{F}^u_{\mathcal{O}}$ always have distinct endpoints. If we moreover assume that ϕ is not conjugate to an Anosov suspension flow (which is automatic when M is atoroidal), the action of $\pi_1(M)$ on $\partial\mathcal{O}$ is minimal [13, Main Theorem].

Convention 2.4. We assume that ϕ is a pseudo-Anosov flow without perfect fits in M.

2.2. End-periodic automorphisms

We briefly recall the basics of end-periodic automorphisms of infinite-type surfaces, which arise naturally in the study of depth-one foliations. Readers are referred to [11] for a more complete treatment of the theory.

Let L be an infinite-type surface without boundary with finitely many ends, all of which are non-planar. Given a homeomorphism $f : L \to L$, an end E of L is a contracting end of f if there is a neighborhood U_E of E and an integer $n > 0$ such that

$f^n(U_E) \subsetneq U_E$ and $\bigcap_{k\geq 0} f^{nk}(U_E)$ is empty. Such a neighborhood U_E is called a regular neighborhood of E. An end is a repelling end of f if it is a contracting end of f^{-1}, and a regular neighborhood of E for f is just a regular neighborhood for f^{-1}. A homeomorphism f is called *end periodic* if each end of L is either contracting or repelling. If f is end periodic, a multi-curve δ in L is called an *f-juncture* if δ is the boundary of a regular neighborhood of an end. If the end is contracting, δ is called a positive f-juncture. Otherwise, it is a negative f-juncture. An f-invariant choice of a positive (resp. negative) f-juncture for each contracting (resp. repelling) end is called a *system of positive (resp. negative) f-junctures*. An end-periodic homeomorphism is *atoroidal* if it does not preserve any essential multi-curve up to isotopy.

Given an end-periodic homeomorphism f, we fix a regular neighborhood U_E for every end E. Let U^+ be the union of U_E of contracting ends, and let U^- be the union of those of repelling ends. The *positive escaping set* $\mathcal{U}^+$ and the *negative escaping set* $\mathcal{U}^-$ are defined as

$$\mathcal{U}^\pm = \bigcup_{n\geq 0} f^{\mp n}(U^\pm).$$

In other words, $\mathcal{U}^+$ is the set of points whose positive iterations escape to contracting ends, and $\mathcal{U}^-$ is the set of points whose negative iterations escape to repelling ends. The mapping torus M_f is non-compact, but it is topologically tame and possesses a nice compactification which we describe below.

The mapping torus M_f is the quotient of $L \times \mathbb{R}$ by an automorphism F where F is given by

$$\begin{aligned} F : L \times \mathbb{R} &\to L \times \mathbb{R} \\ (x,t) &\mapsto (f^{-1}(x), t+1). \end{aligned}$$

We attach $\mathcal{U}^+ \times \{+\infty\}$ and $\mathcal{U}^- \times \{-\infty\}$ to $L \times \mathbb{R}$ to obtain a manifold N with boundary. The transformation F extends to an automorphism of N by setting $F(x, \pm\infty) = (f^{-1}(x), \pm\infty)$. The $\mathbb{Z}$-action on N generated by F is a covering action, and the quotient space $\overline{M_f}$ is a compact 3-manifold with interior M_f and boundary $\partial\overline{M_f} = \partial^+\overline{M_f} \cup \partial^-\overline{M_f}$, where $\partial^\pm\overline{M_f}$ is a (possibly disconnected) closed surface homeomorphic to $\mathcal{U}^\pm/f$. See [21] for a more detailed discussion of the construction. In particular, [21, Lemma 3.3] shows that $\overline{M_f}$ is atoroidal if and only if f is atoroidal.

2.3. Depth-one foliation

A foliation $\mathcal{F}$ in M is a *depth-one foliation* if $\mathcal{F}$ has finitely many compact leaves, whose union we denote by $\mathcal{F}^0$, and $\mathcal{F}$ restricted to $M - \mathcal{F}^0$ is a fibration over a circle with non-compact fibers. A connected component of $M - \mathcal{F}^0$ is called a *fibered region* of M. Any fibered region Ω is bounded by leaves in $\mathcal{F}^0$, and we denote the collection of these leaves by $\partial\Omega$. We say a leaf in $\partial\Omega$ is a *positive* (resp. *negative*) *boundary leaf* of Ω if it is on the ϕ-positive (resp. ϕ-negative) side of Ω. We denote the collection of positive/negative boundary leaves by $\partial^\pm\Omega$. Note that it is possible to have a compact leaf contained in both $\partial^+\Omega$ and $\partial^-\Omega$. Let $\overline{\Omega}$ be the union of Ω and $\partial\Omega$.

Let L be a fiber of the fibration $\mathcal{F}|_\Omega$. The leaf L limits on a compact leaf $\Sigma \subset \partial\Omega$ in the following way [9]. Let $N(\Sigma) \cong \Sigma \times [-1, 1]$ be a regular neighborhood of Σ with Σ identified with $\Sigma \times \{0\}$, and assume that L limits on Σ on the positive side. If $N(\Sigma)$ is small enough, the intersection of L and $N(\Sigma)$ is an infinite surface spiraling to Σ and covering Σ with infinite degree. More precisely, up to shrinking $N(\Sigma)$, there is a multicurve δ on Σ so that $L \cap N(\Sigma)$ is isotopic to an oriented cut-and-paste of $\delta \times (0, 1]$ and $\bigcup_{n\geq 2} \Sigma \times \{1/n\}$. A fundamental domain for the spiraling is depicted on top of Figure 1a, and a schematic picture of the spiraling neighborhood is shown in Figure 1b.

We say a foliation $\mathcal{F}$ is *taut* if for any leaf of $\mathcal{F}$, there is a transverse loop intersecting that leaf.

Convention 2.5. We fix a taut depth-one foliation $\mathcal{F}$ of M and assume that $\mathcal{F}$ is transverse to ϕ. By transversality, $\mathcal{F}$ is coorientable, and we take the orientation to be consistent with the flow direction.

By transversality of $\mathcal{F}$ to ϕ and by the compactness of M, the angle between $T\phi$ and $T\mathcal{F}$ at any point is uniformly bounded away from zero. In particular, this implies that $\phi|_\Omega$ is a suspension flow of the fibration $\mathcal{F}|_\Omega$. The flow gives us a way to identify each $\mathcal{F}$-leaf in Ω with L. Since $\mathcal{F}$ is taut, every leaf in $\mathcal{F}^0$ is homologically non-trivial and incompressible by Novikov [29]. Since M is atoroidal, every compact leaf is a closed hyperbolic surface (the possibility of being a sphere is ruled out by the Reeb stability theorem). Therefore, the fundamental domains of the spiraling of L can be chosen to be non-planar. In a spiraling neighborhood of a compact leaf, ϕ looks like the product flow. One can see that the first return map induced by ϕ is an end-periodic homeomorphism $f : L \to L$ [15]. Since M is atoroidal, so is f. The metric completion of Ω with respect to the path metric induced by the metric in M gives a compactified mapping torus of f.

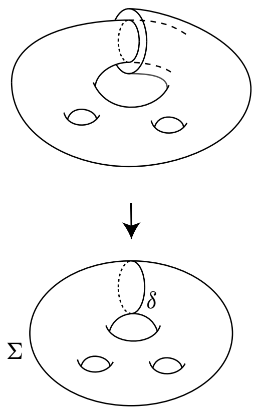

(a)

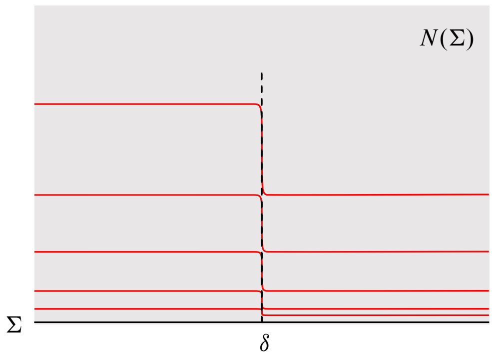

(b)

Figure 1. A schematic picture of the spiraling neighborhood of Σ.

By [8], there is a Riemannian metric on M that restricts to hyperbolic metrics on leaves of $\mathcal{F}$. We fix such a metric on M. A hyperbolic metric on a surface is *standard* if there is no embedded half-space, following [10]. The induced hyperbolic metric on any depth-one leaf of $\mathcal{F}$ has injectivity radius bounded from above. This is discussed in [23, Proof of Proposition 4.6], for instance, and the rough idea is that the junctures on any depth-one leaf have bounded length, and any point on the leaf is at a bounded distance from a juncture.

Let $\mathcal{U}^{\pm} \subset L$ be the positive or negative escaping set of f. By definition, a point $x \in L$ is in $\mathcal{U}^+$ if and only if the positive ray of $\phi(x)$ hits $\partial^+\Omega$, and it is in $\mathcal{U}^-$ if and only if the negative ray of $\phi(x)$ hits $\partial^-\Omega$. If $(\overline{\Omega}, \partial^-\Omega, \partial^+\Omega) \cong (S \times [0,1], S \times \{0\}, S \times \{1\})$ or $(S \times S^1, S \times \{1\}, S \times \{1\})$ for some closed hyperbolic surface S, we say Ω is a trivial fibered region. This is equivalent to the monodromy f being a pure translation on L and to $\mathcal{U}^- = \mathcal{U}^+ = L$ [11, Proposition 4.76].

Convention 2.6. We assume that M has no trivial product region. In particular, any compact leaf of $\mathcal{F}$ is not a fiber. All of our discussions and statements hold with trivial product regions and are readily checked, so we omit the related discussion for simplicity.

Let $\widetilde{\mathcal{F}}$ be the lift of $\mathcal{F}$ in $\widetilde{M}$. The lifted foliation $\widetilde{\mathcal{F}}$ is a foliation by planes by [29]. A connected lift of a fibered region of M is called a *product region* of $\widetilde{M}$. Any product region $\widetilde{\Omega}$ covering Ω is homeomorphic to $\widetilde{L} \times \mathbb{R}$, and we fix a homeomorphism so that $\widetilde{\Omega}$ is foliated by $\widetilde{L} \times \{t\}$ and the $\widetilde{\phi}$-orbits are $\{x\} \times \mathbb{R}$. A lift of a positive/negative boundary leaf of Ω is a positive/negative boundary leaf of $\widetilde{\Omega}$, the collection of which is denoted by $\partial^{\pm}\widetilde{\Omega}$. Define

$$\partial\widetilde{\Omega} := \partial^+\widetilde{\Omega} \cup \partial^-\widetilde{\Omega} \quad \text{and} \quad \overline{\widetilde{\Omega}} := \widetilde{\Omega} \cup \partial\widetilde{\Omega}.$$

Let $\widetilde{\mathcal{U}}^{\pm}$ be the preimage of $\mathcal{U}^{\pm}$ in $\widetilde{L}$. From the construction of the compactified mapping torus, we see that $\overline{\widetilde{\Omega}}$ is homeomorphic to

$$(\widetilde{L} \times \mathbb{R}) \cup (\widetilde{\mathcal{U}}^+ \times \{+\infty\}) \cup (\widetilde{\mathcal{U}}^- \times \{-\infty\}), \tag{2.1}$$

where $\widetilde{\mathcal{U}}^{\pm}$ is the preimage of $\mathcal{U}^{\pm}$ in $\widetilde{L}$.

For any leaf $\mu \in \partial^+\widetilde{\Omega}$, there is a component $\widetilde{\mathcal{U}}^+_{\mu}$ of $\widetilde{\mathcal{U}}^+$ such that a $\widetilde{\phi}$-orbit intersects μ if and only if it intersects $\widetilde{L}$ at a point contained in $\widetilde{\mathcal{U}}^+_{\mu}$. This gives a bijection between leaves in $\partial^+\widetilde{\Omega}$ and components of $\widetilde{\mathcal{U}}^+$. Similarly, there is a bijection between leaves in $\partial^-\widetilde{\Omega}$ and components of $\widetilde{\mathcal{U}}^-$.

A leaf of $\widetilde{\mathcal{F}}$ is called a *type-0 leaf* if it covers a leaf in $\mathcal{F}^0$. Otherwise, we call it a *type-1 leaf*. Every type-1 leaf μ is contained in a unique product region in $\widetilde{M}$, denoted by $\widetilde{\Omega}(\mu)$. Every type-0 leaf is the negative boundary of a unique production region, and the positive boundary of another different product region. A type-0 leaf λ and a type-1 leaf μ are called *adjacent* if $\lambda \subset \partial\widetilde{\Omega}(\mu)$. If moreover λ is in the positive side of μ, we say λ is *positively adjacent* to μ or $\mu \lesssim \lambda$; otherwise, we say λ is *negatively adjacent* to μ or $\mu \gtrsim \lambda$.

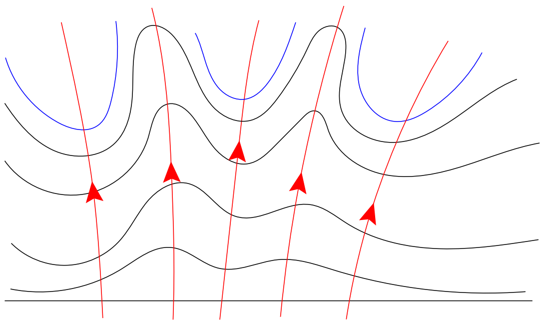

Figure 2. Type-1 leaves (black) in a product region limit to type-0 leaves (blue) in the positive boundary and stay transverse to $\widetilde{\phi}$ (red), creating non-Hausdorffness in Λ.

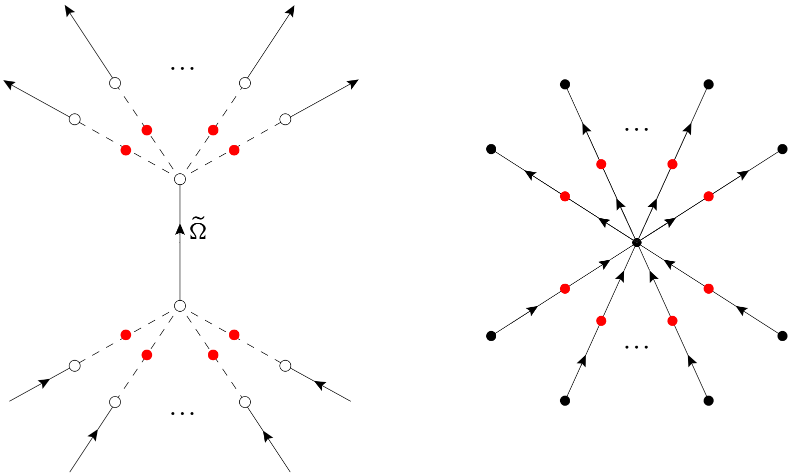

Figure 3. Left: a local picture near a product region $\widetilde{\Omega}$ in Λ. Right: the corresponding parts in Λ^*. The red vertices represent type-0 leaves in both pictures, and the arrows indicate the direction of $\widetilde{\phi}$.

Let Λ be the leaf space of $\widetilde{\mathcal{F}}$ obtained by collapsing each leaf to a point, which is a non-Hausdorff 1-manifold. Since each leaf of $\widetilde{\mathcal{F}}$ is properly embedded in $\widetilde{M}$ and hence separating, Λ is simply connected. Each product region $\widetilde{\Omega}$ projects to an oriented open interval in Λ, with the orientation induced by $\widetilde{\phi}$. Every such open interval has a countably infinite number of positive endpoints, each corresponding to a component of $\partial^+\widetilde{\Omega}$. See Figure 2 for an illustration of how leaves in the product regions in $\widetilde{M}$ limit to different type-0 leaves. The positive endpoints of the open intervals are non-separated from each other in Λ by (2.1). The same is true for negative endpoints. The closures of product regions are glued together along type-0 leaves in the boundary, so each point corresponding to a type-0 leaf is the negative endpoint of exactly one open interval associated to a product region, and the positive endpoint of exactly another different one.

It is sometimes useful to think about the dual graph Λ^*. The set of vertices is the set of product regions and type-0 leaves. For a product region or a type-0 leaf x, we denote the dual vertex in Λ^* by x^*. The edges are the pairs (Ω^*, μ^*) where Ω is a product region, μ is a type-0 leaf and $\mu \in \partial\Omega$. The dual graph Λ^* is an infinite valence tree with an orientation given by the flow (see Figure 3).

Two leaves of $\widetilde{\mathcal{F}}$ are called *comparable* if they can be connected by an oriented path in Λ. Otherwise, they are *incomparable*. If λ is comparable to μ and μ is on the positive side of λ, we write $\lambda < \mu$ or $\mu > \lambda$. Similarly, we say two vertices are comparable if there is an oriented path connecting them in Λ^* and are incomparable if otherwise. If two vertices v and w are comparable and v is on the positive side of w, we write $w < v$ or $v > w$. Two product regions are comparable/incomparable if their dual vertices are comparable/incomparable.

2.4. Laminations on S^1

We recall some definitions and constructions of abstract laminations on S^1.

Let $\mathrm{Symm}_2(S^1) := S^1 \times S^1 - \Delta/\sim$ be the space of unordered pairs of distinct points in S^1 endowed with the quotient topology, where the relation is given by $(x, y) \sim (y, x)$. Two pairs of points $\{x, y\}$ and $\{z, w\}$ on S^1 are said to be unlinked if z and w lie in the same component of $S^1 - \{x, y\}$. A lamination on S^1 is a closed pairwise unlinked subset of $\mathrm{Symm}_2(S^1)$.

By identifying S^1 with $\partial\mathbb{H}^2$, any lamination Ξ on S^1 determines a geodesic lamination Ξ_{geod} on $\mathbb{H}^2$ by taking the union of geodesics that connect the pairs in Ξ. Conversely, any geodesic lamination in $\mathbb{H}^2$ gives a lamination on S^1 consisting of endpoint pairs of leaves.

Given any subset $A \subset S^1$, the boundary of the convex hull of $\overline{A}$ is a geodesic laminations on $\mathbb{H}^2$, which can be viewed as a lamination $\partial\,\mathrm{CH}(A)$ on S^1. Note that the lamination $\partial\,\mathrm{CH}(A)$ is independent of the choice of the identification $S^1 \cong \partial\mathbb{H}^2$.

3. Infinity of shadows

Let $p : \widetilde{M} \to \mathcal{O}$ be the projection map. For any subset A of $\widetilde{M}$, the image of A under p is called the *shadow* of A. In this section, we will study the shadow of leaves of $\widetilde{\mathcal{F}}$, especially the behavior of the shadow at infinity. The main results are Lemmas 3.3 and 3.5, which hint at a universal circle structure on $\partial\mathcal{O}$.

Let λ be a type-0 leaf of $\widetilde{\mathcal{F}}$, and let Σ be the leaf of $\mathcal{F}$ it covers. Denote the covering map by π_0. Since λ separates $\widetilde{M}$ and the flow $\widetilde{\phi}$ crosses λ positively, each flowline intersects λ at most once. It follows that p restricted to λ is a continuous bijection to its image. This map $p|_\lambda : \lambda \to p(\lambda)$ also has a continuous inverse, mapping a point in $p(\lambda)$ to the intersection of the corresponding orbit with λ. Therefore, $p|_\lambda$ is a homeomorphism to $p(\lambda)$. The map $\pi := \pi_0 \circ (p|_\lambda)^{-1} : p(\lambda) \to \Sigma$ is then a covering map, so we can view $p(\lambda)$ as a universal cover of Σ. The deck transformation group action on $p(\lambda)$ is simply the action of $\pi_1(M)$ on $\mathcal{O}$ restricted to a subgroup in the conjugacy class of $\pi_1(\Sigma)$ that stabilizes $p(\lambda)$. We identify this subgroup with $\pi_1(\Sigma)$. By transversality, the intersection of $\mathcal{F}^{s/u}$ with Σ induces a singular foliation on Σ, denoted by $\mathcal{F}^{s/u}_\Sigma$. Similarly, let the intersection of $\widetilde{\mathcal{F}}^{s/u}$ with λ be denoted by $\widetilde{\mathcal{F}}^{s/u}_\lambda$. These foliations are related by the relations $p(\widetilde{\mathcal{F}}^{s/u}_\lambda) = \mathcal{F}^{s/u}_{\mathcal{O}}|_{p(\lambda)}$ and $\pi(\mathcal{F}^{s/u}_{\mathcal{O}}|_{p(\lambda)}) = \mathcal{F}^{s/u}_\Sigma$.

Shadows of type-0 leaves. Assume λ is a type-0 leaf, so Σ is an embedded closed surface in M which is not a fiber by Convention 2.6. The shape of the shadow of a non-fibered transverse closed embedded surface is carefully studied in [12] when ϕ is a pseudo-Anosov suspension flow, later generalized by [19] to general pseudo-Anosov flows. Each component of the topological boundary of $p(\lambda)$ in $\mathcal{O}$ is either a regular leaf of $\mathcal{F}^u_{\mathcal{O}}$ or $\mathcal{F}^s_{\mathcal{O}}$ or a face of a singular leaf that is regular on the side containing $p(\lambda)$ [19, Proposition 4.3]. We call a boundary component of $p(\lambda)$ in $\mathcal{O}$ a *side* of $p(\lambda)$. Consider a side e of $p(\lambda)$, which we assume to be contained in a leaf of $\mathcal{F}^s_{\mathcal{O}}$ for concreteness. We collect some useful facts about the local dynamics at e from [19] in the following proposition.

Proposition 3.1 ([19]). *The stabilizer of e in $\pi_1(M)$ is isomorphic to $\mathbb{Z}$ and contained in $\pi_1(\Sigma)$. There is a generator g_e acting on $p(\lambda)$ with the following dynamics (Figure 4):*

(1) *The element g_e acts as a contraction on e with a unique fixed point x_e.*

(2) *The element g_e fixes and expands $l_e := \mathcal{F}^u_{\mathcal{O}}(x_e) \cap p(\lambda)$, which is the interior of a ray of $\mathcal{F}^u_{\mathcal{O}}$. We have that l_e projects via π to a closed leaf α_e of $\mathcal{F}^u_\Sigma$, whose free homotopy class is represented by g_e.*

(3) *For any point x other than x_e in e, the intersection $l_x := \mathcal{F}^u_{\mathcal{O}}(x) \cap p(\lambda)$ is connected, and it projects via π to a non-compact leaf of $\mathcal{F}^u_\Sigma$ that spirals into α_e. If we orient l_x and l_e so that they point toward e and let $\pi(l_x)$ and α_e inherit the orientation, then $\pi(l_x)$ and α_e are asymptotic in the forward direction.*

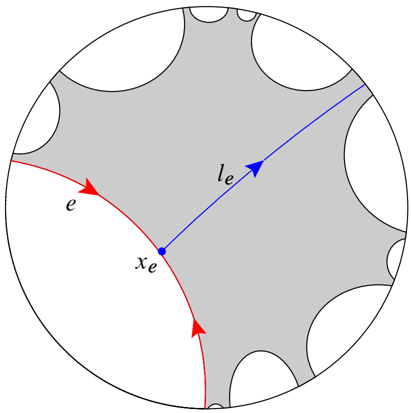

Figure 4. The dynamics near a side of the shadow of a type-0 leaf.

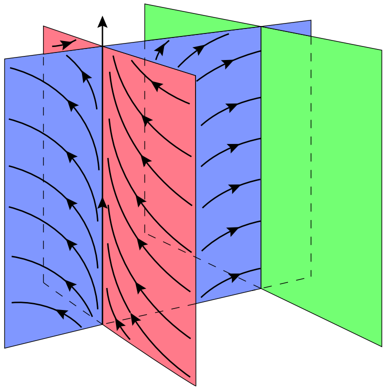

Figure 5. The vertical arrowed line represents the periodic orbit σ_e (assumed to be regular), the red plane is $\widetilde{\mathcal{F}}^s(\sigma_e)$, the blue plane is $\widetilde{\mathcal{F}}^u(\sigma_e)$ and the green plane is λ.

The 3-dimensional picture is the following (Figure 5). In $\widetilde{M}$, the orbit $\sigma_e := p^{-1}(x_e)$ is a periodic orbit disjoint from λ. The leaf $p^{-1}(e) \subset \widetilde{\mathcal{F}}^s(\sigma_e)$ does not intersect λ, and $p^{-1}(l_e) \subset \widetilde{\mathcal{F}}^u(\sigma_e)$ intersects λ transversely in a line $\widetilde{\alpha}_e$ so that all the $\widetilde{\phi}$-orbits in $p^{-1}(l_e)$ cross λ positively. The line $\widetilde{\alpha}_e$ covers the simple closed curve α_e in Σ. The element g_e is a translation that fixes σ_e and λ.

For a side e contained in a leaf of $\mathcal{F}^u_{\mathcal{O}}$, we have a similar picture. From now on, given a side e of $p(\lambda)$, we will continue to use the notations g_e, l_e, x_e and α_e for the objects described in Proposition 3.1. In particular, we always take g_e to be the generator in the stabilizer of e that contracts e.

Let $\partial p(\lambda)$ be the boundary of $p(\lambda)$ *in* $\overline{\mathcal{O}}$, consisting of sides and ideal boundary points at infinity of $p(\lambda)$.

Lemma 3.2. *The intersection $\partial p(\lambda) \cap \partial\mathcal{O}$ is nowhere dense.*

Proof. First we observe that $\partial p(\lambda) \cap \partial\mathcal{O}$ is a closed subset of $\partial\mathcal{O}$. This is because the complement of $\partial p(\lambda) \cap \partial\mathcal{O}$ on $\partial\mathcal{O}$ is a union of disjoint open intervals bounded by the endpoints of sides of $p(\lambda)$. Suppose that there is a maximal closed interval A contained in $\partial p(\lambda) \cap \partial\mathcal{O}$ with endpoints η_1 and η_2. Then there is a side e_1 of $p(\lambda)$ such that η_1 is an endpoint of e_1. Take g_{e_1}, x_{e_1} and l_{e_1} as before. Let c be the endpoint of l_{e_1} at $\partial\mathcal{O}$.

There is another side e_2 of $p(\lambda)$ such that η_2 is an endpoint of e_2. If c does not lie in A, then e_2 is between c and η_1 (Figure 6). The action of g_{e_1} fixes η_1 and c, but it cannot fix e_2. Otherwise, g_{e_1} will have two fixed points on $\mathcal{O}$ by Lemma 2.2, contradicting the assumption that ϕ has no perfect fits by Lemma 2.3. Then one of $g_{e_1}^{\pm 1}(\eta_2)$ will be in A, contradicting our choice of A.

Hence, c must lie in A. Now l_{e_1} divides $p(\lambda)$ into two connected components, one of which contains no side of $p(\lambda)$. In particular, this component contains no $\pi_1(\Sigma)$-translation of l_{e_1}. Translated to the hyperbolic plane by Q_λ, this means there is a simple

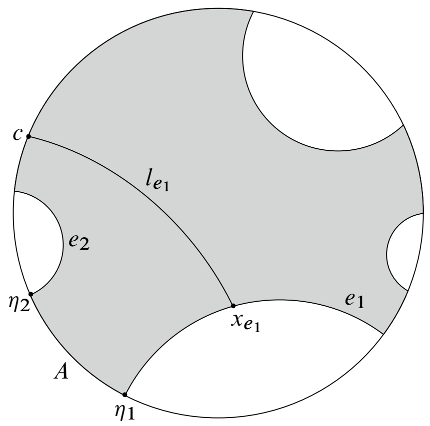

Figure 6. Proof of Lemma 3.2.

closed curve $\alpha_{e_1} \in \Sigma$ and a lift $Q_\lambda(l_{e_1})$ of α_{e_1} in λ such that there is no other lift on one side of $\widetilde{\alpha}$. But that is impossible. We conclude that such an interval A does not exist. ■

The boundary $\partial p(\lambda)$ is homeomorphic to a circle, and $\overline{p(\lambda)} := \partial p(\lambda) \cup p(\lambda)$ is the closure of $p(\lambda)$ in $\overline{\mathcal{O}}$. Note that our choice of the metric on M restricts to a hyperbolic metric on Σ, so λ is isometric to $\mathbb{H}^2$. We define $\overline{\lambda}$ to be the usual compactified hyperbolic plane with ideal boundary $\partial_\infty\lambda$. The next lemma reveals how $\partial p(\lambda)$ is related to $\partial_\infty\lambda$.

A map g from S^1 to S^1 is *monotone* if the preimage of any point on S^1 is contractible. A *gap* of g is a maximal closed interval of positive length in S^1 that is collapsed to a single point by g. The *core* of g is the complement of the union of the interiors of gaps, denoted by $\operatorname{core}(g)$. We remark that the gaps are sometimes taken to be open intervals in the literature, different from our convention.

Lemma 3.3. *Let λ be a type-0 leaf of $\widetilde{\mathcal{F}}$. Then the homeomorphism $p^{-1}|_{p(\lambda)} : p(\lambda) \to \lambda$ extends continuously to a map Q_λ from $\overline{p(\lambda)}$ to $\overline{\lambda} := \partial_\infty\lambda \cup \lambda$. The map restricted to $\partial p(\lambda)$ is a monotone map to $\partial_\infty\lambda$ with core $\overline{p(\lambda)} \cap \partial\mathcal{O}$.*

Proof. To begin with, we define $Q_\lambda = p^{-1}|_{p(\lambda)}$ in the interior of $p(\lambda)$. To extend Q_λ to $\partial p(\lambda)$, we will first define the map on the sides of $p(\lambda)$ and then extend it to the entire $\partial p(\lambda)$.

Let e be a side of $p(\lambda)$. Take $g_e \in \operatorname{Stab}(e)$, $x_e \in e$ and l_e as in Proposition 3.1. As an element of $\pi_1(\Sigma)$, g_e acts as a hyperbolic element on $\overline{\lambda}$ with a contracting fixed point $\partial^- g_e$ and a repelling fixed point $\partial^+ g_e$ at infinity. We define Q_λ on e as the constant map to $\partial^- g_e$. Different sides are sent to different points in $\partial_\infty(\lambda)$ because of Lemma 2.3.

Since $Q_\lambda(l_e)$ projects to a simple closed curve α_e in Σ, it is a quasi-geodesic in λ with well-defined endpoints in $\partial_\infty\lambda$. Moreover, α_e represents the free homotopy class of g_e up to taking the inverse, by Proposition 3.1. Therefore, the endpoints of $Q_\lambda(l_e)$ are $\partial^\pm g_e$. By the way we choose g_e (item (2) of Proposition 3.1), it contracts l_e near e. Therefore, if we

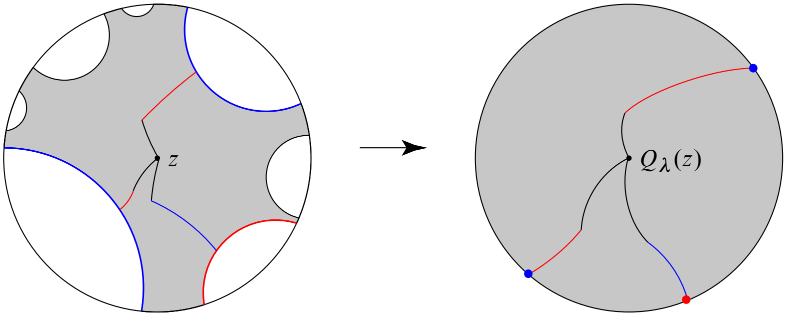

Figure 7. Q_λ preserves the cyclic order.

orient l_e to point toward e and give $Q_\lambda(l_e)$ the induced orientation, the forward endpoint of $Q_\lambda(l_e)$ is $\partial^- g_e$. This shows Q_λ is continuous at x_e when restricted to $l_e \cup x_e$. This further ensures that Q_λ preserves the cyclic order of the boundary leaves, which is a consequence of the fact that Q_λ is an orientation-preserving homeomorphism in the interior. For example, one can argue as follows (indicated in Figure 7). Consider three sides e_1, e_2 and e_3 in clockwise order, and take x_{e_i} and l_{e_i} as before. Take a point $y_{e_i} \in l_{e_i}$, and let $l'_{e_i} \subset l_{e_i}$ be the segment between x_{e_i} and y_{e_i}. We may ensure l'_{e_i} are pairwise disjoint by taking y_{e_i} close enough to x_{e_i}. Take any point $z \in p(\lambda)$ that is not in the union of l'_{e_i}, and connect y_{e_i} to z to get an embedded 3-prong P with z as the center and x_{e_i} as the endpoints. The image $Q_\lambda(P)$ is an embedded 3-prong in $\overline{\lambda}$ with well-defined endpoints $Q_\lambda(e_i)$. The cyclic order of the edges of R is preserved under Q_λ, assuring that $Q_\lambda(e_1)$, $Q_\lambda(e_2)$ and $Q_\lambda(e_3)$ are arranged clockwise.

Moreover, if γ is an element of $\pi_1(\Sigma)$, it translates e to another side γe of $p(\lambda)$. We have $g_{\gamma e} = \gamma g_e \gamma^{-1}$, so $Q_\lambda(\gamma e) = \gamma Q_\lambda(e)$. By the minimality of the $\pi_1(\Sigma)$ action on $\partial_\infty \lambda$, the Q_λ-images of all the sides of $p(\lambda)$ are dense in $\partial_\infty \lambda$. Now there is a unique way to define Q_λ on $\partial p(\lambda)$ such that it is continuous when restricted to $\partial p(\lambda)$. That is, for any point $x \in \partial p(\lambda)$, one can find a sequence of sides e_n converging to x and define $Q_\lambda(x)$ as the limit of $Q_\lambda(e_n)$ by Lemma 3.2. The limit exists and is independent of e_n because Q_λ preserves the cyclic order of sides and the image of sides is dense in $\partial_\infty \lambda$.

To complete the proof of Lemma 3.3, what is left is to check that Q_λ is continuous on $\overline{p(\lambda)}$.

For a side e (which is assumed to be contained in a leaf of $\mathcal{F}^s_{\mathcal{O}}$ for concreteness) of $p(\lambda)$ and a point $x \in e$, we take a rectangular neighborhood of x as the image of an embedding $\rho : (0,1) \times [0,1) \to \overline{p(\lambda)}$ such that

- $\rho(1/2, 0) = x$;
- $\rho((0,1) \times \{0\})$ is contained in e;
- for any $s \in [0,1)$, $\rho((0,1) \times \{s\})$ is contained in a leaf of $\mathcal{F}^s_{\mathcal{O}}$;
- for any $t \in (0,1)$, $\rho(\{t\} \times (0,1))$ is contained in a leaf of $\mathcal{F}^u_{\mathcal{O}}$.

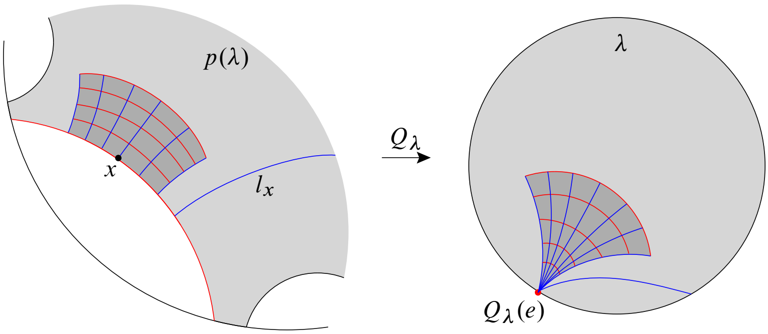

Figure 8. The Q_λ-image of $\mathcal{F}^s_{\mathcal{O}}$-leaves near an unstable boundary leaf e.

Such a neighborhood always exists because $\mathcal{F}^s_{\mathcal{O}}$ is regular on the side of e that contains $p(\lambda)$.

Fix a rectangular neighborhood $\mathrm{RN}(x)$ of x. By Proposition 3.1, we know that $Q_\lambda(\rho(\{t\} \times (0,1)))$ is asymptotic to $Q_\lambda(l_e)$, the latter being a quasi-geodesic ray in λ since it is a lift of a simple closed curve in Σ. This shows that $Q_\lambda(\mathrm{RN}(x))$ is a wedge-shaped region in λ with one ideal point $Q_\lambda(e)$ (Figure 8).

Consider a sequence of points x_n in $p(\lambda)$ converging to $x \in \partial p(\lambda)$. If x is on a side e, we can take a rectangular neighborhood of x that will eventually contain x_n. The shape of the rectangular neighborhoods under Q_λ shows that $Q_\lambda(x_n)$ converges to $Q_\lambda(x) = Q_\lambda(e)$.

If $x \in \partial p(\lambda) \cap \partial\mathcal{O}$, we can trap $Q_\lambda(x_n)$ using segments of l_e, forcing them to converge to the right points. The argument is similar to the one we use to prove that Q_λ preserves the cyclic order of the sides. See Figure 9 for an illustration. More precisely, take sequences of sides $\{e_m^+\}$ and $\{e_m^-\}$ of $p(\lambda)$ that approximate x from two sides respectively (using Lemma 3.2). Fixing m, take a short segment $l'_{e_m^\pm} \subset l_{e_m^\pm}$ with one endpoint

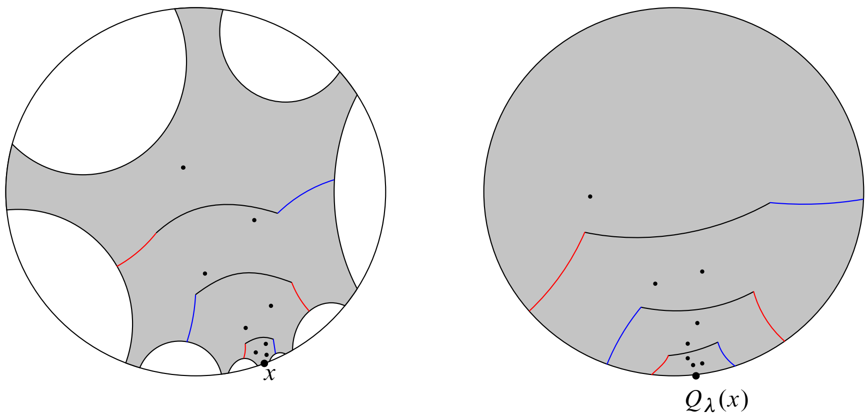

Figure 9. Continuity of Q_λ at $\partial p(\lambda) \cap \partial\mathcal{O}$.

at $x_{e_m^\pm}$ so that $l'_{e_m^+}$ and $l'_{e_m^-}$ are disjoint. Connect the other endpoints of the two segments by a path in $p(\lambda)$, and denote the resulting path by α_m. The image $Q_\lambda(\alpha_m)$ is an embedded line in $\widetilde{\lambda}$ with disjoint endpoints, separating $\widetilde{\lambda}$ into two half-planes. We let the half-plane containing $Q_\lambda(x)$ be $H_m(x)$. The fact that $Q_\lambda|_{p(\lambda)}$ is an orientation-preserving homeomorphism guarantees that $Q_\lambda(x_n)$ eventually enters $H_m(x)$ for all m. We can arrange $H_m(x)$ to be nested so that $\bigcap_m H_m(x)$ has exactly one ideal point $Q_\lambda(x)$. Since $\{x_n\}$, hence $\{Q_\lambda(x_n)\}$, escapes every compact set, we know that $\{Q_\lambda(x_n)\}$ must limit to $Q_\lambda(x)$.

We have proved that Q_λ is continuous on $\overline{p(\lambda)}$, finishing the proof of Lemma 3.3. ■

Shadows of type-1 leaves. In this subsection, we would like to study the shadow of type-1 leaves in $\widetilde{\mathcal{F}}$. Note that by [20, Proposition 4.1], the shadow of a leaf of a transverse foliation is always bounded by regular leaves or faces of singular leaves of the stable/unstable foliations. These boundary leaves are in general not periodic for leaves in an arbitrary transverse foliation. However, we will show in the next lemma that they are in fact periodic for type-1 leaves in a depth-one foliation.

We now assume λ is a type-1 leaf of $\widetilde{F}$ contained in a product region $\widetilde{\Omega} \cong \widetilde{L} \times \mathbb{R}$ and λ covers a non-compact leaf identified with L. Similar to the case of type-0 leaves, we will use $\partial p(\lambda)$ to denote the boundary of $p(\lambda)$ in $\overline{\mathcal{O}}$.

Lemma 3.4. *The shadow $p(\lambda)$ is a proper open subset of $\mathcal{O}$, bounded by periodic regular leaves of $\mathcal{F}^s_{\mathcal{O}}$ or $\mathcal{F}^u_{\mathcal{O}}$ or faces of singular leaves which are regular on the side containing $p(\lambda)$. We call a component of the boundary of $p(\lambda)$ (as a subset of $\mathcal{O}$) a side of $p(\lambda)$.*

Moreover, if a side e of $p(\lambda)$ is contained in a leaf of $\mathcal{F}^s_{\mathcal{O}}$, then there is a type-0 leaf μ negatively adjacent to λ (see Section 2.3 for the definition) such that e is also a side of $p(\mu)$. If e is contained in a leaf of $\mathcal{F}^u_{\mathcal{O}}$, then there is a type-0 leaf μ positively adjacent to λ such that e is also a side of $p(\mu)$.

Proof. Suppose $z \in \mathcal{O}$ is a boundary point of $p(\lambda)$. Let $\sigma_z = p^{-1}(z)$ be the $\widetilde{\phi}$-orbit that projects to z. The orbit σ_z cannot intersect any product region that is incomparable to $\widetilde{\Omega}(\lambda)$, which is the product region containing λ. This is because if σ_z intersects such a product region $\widetilde{\Omega}$, z will have a neighborhood U such that every orbit that projects to U also intersects $\widetilde{\Omega}$. This forces them to be disjoint from $\widetilde{\Omega}(\lambda)$. Then U is not in $p(\lambda)$, and z is not in the boundary of $p(\lambda)$. Similarly, it cannot intersect any type-0 leaf incomparable to $\widetilde{\Omega}(\lambda)$ either.

Since $p(\lambda)$ is open, z is not contained in $p(\lambda)$. The orbit σ_z induces an oriented path γ_z in the dual graph Λ^* consisting of all the type-0 leaves and product regions that σ_z travels through (Section 2). Every vertex in γ_z is comparable to the dual vertex $\widetilde{\Omega}(\lambda)^*$ by the previous paragraph, but $\widetilde{\Omega}(\lambda)^*$ is not in γ_z. Take any vertex v in γ_z, and suppose $v > \widetilde{\Omega}(\lambda)^*$. Since Λ^* is a tree, any vertex w satisfying $v > w > \widetilde{\Omega}(\lambda)^*$ lies in the unique oriented interval from v to $\widetilde{\Omega}(\lambda)^*$. So the path γ_z has to stay in the interval in the negative direction. The case where $v < \widetilde{\Omega}(\lambda)^*$ is similar, and we conclude that the orbit σ_z will eventually stay in one product region in either the positive or the negative direction.

If γ_z is on the negative side of $\widetilde{\Omega}(\lambda)^*$, then σ_z eventually stays in a product region, denoted by $\widetilde{\Omega}_0$, in the positive direction. Take the unique shortest oriented path γ from $\widetilde{\Omega}_0^*$ to $\widetilde{\Omega}(\lambda)^*$ in Λ^*, and let $\widetilde{\Omega}_1^*$ be the last vertex in γ before $\widetilde{\Omega}(\lambda)^*$. There is a unique type-0 leaf μ that is positively adjacent to $\widetilde{\Omega}_1$ and negatively adjacent to $\widetilde{\Omega}(\lambda)$. Any $\widetilde{\phi}$-orbit intersecting both $\widetilde{\Omega}_0$ and λ has to enter $\widetilde{\Omega}(\lambda)$ via μ. Conversely, every $\widetilde{\phi}$-orbit intersecting μ also meets λ. This shows that if an orbit σ intersects $\widetilde{\Omega}_0$, then $p(\sigma) \in p(\mu)$ if and only if $p(\sigma) \in p(\lambda)$. Since $p(\widetilde{\Omega}_0)$ is open, we have also shown that if x is in $p(\widetilde{\Omega}_0)$, then $x \in \partial p(\lambda)$ if and only if $x \in \partial p(\mu)$.

Since z is in the interior of $p(\widetilde{\Omega}_0)$, the above reasoning shows that z is in the boundary of $p(\mu)$. Let e_z be the side of $p(\mu)$ containing z. Note that all the orbits in the same $\widetilde{\mathcal{F}}^s$-leaf as σ_z are positively asymptotic to σ_z. In particular, they intersect $\widetilde{\Omega}_0$ and are disjoint from μ. We see that e_z is contained in a leaf of $\mathcal{F}^s_{\mathcal{O}}$, and using a similar argument to the last paragraph, we have $e_z \subset \partial p(\lambda)$.

To summarize, we have shown the following: for a point $z \in \partial p(\lambda)$, if γ_z is on the negative side of $\widetilde{\Omega}(\lambda)^*$, then there is type-0 leaf μ negatively adjacent to λ and a side $e_z \in \partial p(\mu)$ contained in a leaf of $\mathcal{F}^s_{\mathcal{O}}$ such that $z \in e_z$ and $e_z \subset \partial p(\lambda)$. One can apply the same argument to the case where γ_z is on the positive side of $\widetilde{\Omega}(\lambda)^*$, finishing the proof. ∎

Similar to the case of type-0 leaves, we define $\overline{p(\lambda)}$ to be $p(\lambda) \cup \partial p(\lambda)$. The following lemma is the type-1 version of Lemma 3.3.

Lemma 3.5. *Let λ be a type-1 leaf of $\widetilde{\mathcal{F}}$. Then the homeomorphism $(p_\lambda)^{-1} : p(\lambda) \to \lambda$ extends continuously to a map Q_λ from $\overline{p(\lambda)}$ to $\partial_\infty \lambda$. The map restricted to $\partial p(\lambda)$ is a monotone map to $\partial_\infty \lambda$ with core $\overline{p(\lambda)} \cap \partial \mathcal{O}$.*

Proof. We define $Q_\lambda = (p|_\lambda)^{-1}$ in $p(\lambda)$ and extend Q_λ to the boundary following the same strategy as in the proof of Lemma 3.3.

Let e be any side of $p(\lambda)$, and we assume it to be a leaf of $\mathcal{F}^s_{\mathcal{O}}$ for concreteness. By Lemma 3.4, e is also a side of the shadow of a type-0 leaf μ negatively adjacent to λ. Take g_e, $x_e \in e$ and l_e as in Proposition 3.1. We orient l_e to be pointing toward e. By [20, Theorem C], each ray of $\widetilde{\mathcal{F}}^s \cap \lambda$ or $\widetilde{\mathcal{F}}^u \cap \lambda$ has a well-defined endpoint at $\partial_\infty \lambda$. In particular, $Q_\lambda(l_e)$ has a well-defined forward endpoint at infinity. We define Q_λ on each side e of $p(\lambda)$ to be the constant map to the forward endpoint of $Q_\lambda(l_e)$.

Claim. Different sides are mapped to different points by Q_λ.

Proof of the claim. Recall that $\widetilde{\Omega}(\lambda)$ is the product region containing λ, and it covers a fibered region $\Omega(\lambda)$ in M with fiber L and the monodromy $h : L \to L$. Since M is atoroidal, h is an atoroidal end-periodic map. The fundamental group of $\Omega(\lambda)$ is isomorphic to the semidirect product $\mathbb{Z} \ltimes \pi_1(L)$ with $\mathbb{Z}$ acting on $\pi_1(L)$ by h_*. The group $\pi_1(\Omega(\lambda))$ stabilizes $p(\lambda)$, and we can define an action of $\pi_1(\Omega(\lambda))$ on λ by

$g(x) := Q_\lambda \circ g \circ p(x)$ for $g \in \pi_1(\Omega(\lambda))$ and $x \in \lambda$. If the element g has a trivial $\mathbb{Z}$-factor, then the action of g is a covering transformation. Otherwise, g acts as a lift of some power of the monodromy h.

Let e_1 and e_2 be different sides of $p(\lambda)$, and take g_{e_i} and x_{e_i} as before for $i = 1, 2$. By Proposition 3.4, there is a type-0 leaf μ_i adjacent to λ so that e_i is a side of $p(\mu_1)$. The type-0 leaf μ_i covers a compact leaf $\Sigma_i \subset \partial\Omega(\lambda)$. The element g_{e_i}, being a deck transformation for the covering map $\mu_i \to \Sigma_i$, can be viewed as an element of $\pi_1(\Omega(\lambda))$. Therefore, g_{e_i} stabilizes $p(\lambda)$ and acts on λ either as a hyperbolic isometry or as a lift of some power of h. In both cases, the action extends continuously to an automorphism of $\partial_\infty\lambda$ by [10]. Since g_{e_i} stabilizes l_{e_i} in $p(\lambda)$, we know g_{e_i} stabilizes $Q_\lambda(l_{e_i})$ in λ, hence also the point $Q_\lambda(e_i)$ by the definition of $Q_\lambda(e_i)$.

Suppose for contradiction that $Q_\lambda(e_1) = Q_\lambda(e_2) =: q$. Then g_{e_1} and g_{e_2} fix the same point q on $\partial\lambda$. Note that g_{e_1} and g_{e_2} represent different elements in $\pi_1(\Omega(\lambda))$ and do not share a non-trivial power by Lemma 2.3. Also note that for any lift $\widetilde{h} : \lambda \to \lambda$ of some power of the monodromy h, the action of $\widetilde{h}$ on $\partial\lambda$ will not fix any fixed point of an element in $\pi_1(L)$ that acts as a hyperbolic isometry on λ. Otherwise, h will fix a closed geodesic α on L up to isotopy. A regular neighborhood $U(\alpha)$ of α will also be fixed by h up to isotopy, so a power of h fixes components of $U(\alpha)$, contradicting the assumption that h is atoroidal. Therefore, if one of the g_{e_i} is a hyperbolic isometry, then g_{e_1} and g_{e_2} have no common fixed point on $\partial_\infty\lambda$. If both of them are lifts of some power of h, up to taking powers we may assume that they are lifts of the same power of h. Then there is an element $\gamma \in \pi_1(\lambda)$ such that $g_{e_1} = \gamma g_{e_2}$. The common fixed point q will also be fixed by γ, a contradiction. We have proved that g_{e_1} and g_{e_2} do not have common fixed points, which indicates that $Q_\lambda(e_1) \neq Q_\lambda(e_2)$. ■

We continue our proof of Lemma 3.5. The map Q_λ preserves the cyclic order of the sides by the same reason as in the proof of Lemma 3.3, and the image is dense by the minimality of the $\pi_1(L)$-action on $\partial_\infty\lambda$. Indeed, the hyperbolic structure on L has bounded injectivity radius, so the limit set of $\pi_1(L)$ is the entire circle at infinity.

We also have the following lemma, analogous to Lemma 3.2.

Lemma 3.6. *The intersection $\partial p(\lambda) \cap \partial\mathcal{O}$ is nowhere dense in $\partial\mathcal{O}$.*

Proof. Let e be a side of $p(\lambda)$. By Lemma 3.4, e is a side of a shadow of a type-0 leaf. So we can take g_e as in Proposition 3.1. By the discussion above, g_e stabilizes $p(\lambda)$, and the action near e has the desired expanding-contracting dynamics because of Proposition 3.1. The rest of the proof is the same as that of Lemma 3.2. ■

To continue the proof of Lemma 3.5, note that by Lemma 3.6, the map Q_λ can be extended to $\partial p(\lambda)$ by approaching any point in $\partial p(\lambda) \cap \partial\mathcal{O}$ by sides of $\partial p(\lambda)$, as in the proof of Lemma 3.3. The extension is well defined and is continuous on $\partial p(\lambda)$. To show that the extension is continuous on $p(\lambda) \cup \partial p(\lambda)$, we again need the nice asymptotic property of rectangular neighborhoods of the sides. What we will show next is basically that Figure 8 is also a correct picture when λ is a type-1 leaf.

Again, let e be any side of $p(\lambda)$, assumed to be a leaf of $\mathcal{F}^s_{\mathcal{O}}$. Let μ be a type-0 leaf negatively adjacent to λ such that e is also a side of $p(\mu)$, using Lemma 3.4. Finally, let g_e, $x_e \in e$ and l_e be as in Proposition 3.1, and orient l_e to be pointing at x_e. Let $x' \in e$ be a point different from x_e and set l' to be a small segment of $\mathcal{F}^u_{\mathcal{O}}(x') \cap p(\mu)$ with an endpoint at x'. We orient l' to be pointing toward e as well. By Proposition 3.1, $Q_\mu(l')$ is forward asymptotic to $Q_\mu(l_e)$ under the orientation induced from l' and l_e. Since $\mathcal{F}^s_{\mathcal{O}}$ is regular near e on the side of $p(\lambda)$, there is a pair of points $p_e \in l_e$ and $p' \in l'$ so that $p' \in \mathcal{F}^s_{\mathcal{O}}(p_e)$. Let σ_e and σ' be the $\widetilde{\phi}$-orbits corresponding to p_e and p', respectively. Since they lie in the same leaf of $\widetilde{\mathcal{F}}^s$, σ_e and σ' are positively asymptotic. It follows that the distance between $\sigma_e \cap \lambda$ and $\sigma' \cap \lambda$ is bounded by the distance between $\sigma_e \cap \mu$ and $\sigma' \cap \mu$ up to a uniform multiplicative constant. In other words, $d_\lambda(Q_\lambda(p_e), Q_\lambda(p'))$ is coarsely bounded by $d_\mu(Q_\mu(p_e), Q_\mu(p'))$. Moreover, the pair $\{p_e, p'\}$ can be chosen arbitrarily close to e. Hence, $Q_\lambda(l')$ and $Q_\lambda(l_e)$ are forward asymptotic. If $\mathrm{RN}(x')$ is a rectangular neighborhood of x' in $p(\mu)$, then we have shown that $Q_\lambda(\mathrm{RN}(x'))$ is a wedge-shaped domain in λ meeting $\partial_\infty \lambda$ at exactly one point.

Since we have a similar description of the Q_λ-image of a rectangular neighborhood, we can carry out the same argument and conclude that the extended Q_λ is continuous on $p(\lambda) \cup \partial p(\lambda)$, finishing the proof of Lemma 3.5. ■

For any leaf λ of $\widetilde{\mathcal{F}}$, either type-0 or type-1, we define a map $I_\lambda : \partial\mathcal{O} \to \partial_\infty \lambda$ as follows. For any $\zeta \in \partial\mathcal{O}$, if ζ is in $\overline{p(\lambda)}$, set $I_\lambda(\zeta) = Q_\lambda(\zeta)$. If ζ is contained in an open interval V_ζ in $\partial\mathcal{O} \setminus \overline{p(\lambda)}$, then there is a boundary leaf e of $p(\lambda)$ with the same endpoints

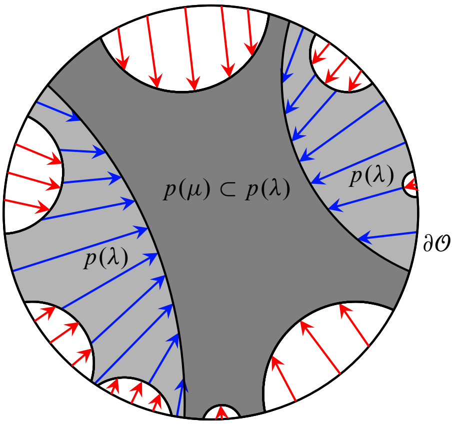

Figure 10. The union of the shaded area is $p(\lambda)$, and the heavily shaded area is $p(\mu)$, which is a subset of $p(\lambda)$. The red arrows represent the map Q_λ, and the blue ones represent $Q_{\lambda\mu}$. One should think of each side of $p(\lambda)$ or $p(\mu)$ as a single point at infinity.

as V_ζ. In this case, we define $I_\lambda(\zeta)$ to be $Q_\lambda(e)$. It is immediate from the definition that the maps I_λ are monotone surjections, therefore continuous, and $\pi_1(M)$-equivariant, that is, for any $g \in \pi_1(M)$, $\lambda \in \widetilde{\mathcal{F}}$ and $x \in \partial\mathcal{O}$, we have $gI_\lambda(x) = I_{g\lambda}(gx)$, where the action $g : \partial_\infty\lambda \to \partial_\infty(g\lambda)$ is induced by the isometry $g : \lambda \to g\lambda$.

If λ is a type-1 leaf and μ is a type-0 leaf adjacent to λ, then $p(\lambda)$ contains $p(\mu)$, and the monotone quotient maps I_λ and I_μ satisfy the property that for any $\xi_1, \xi_2 \in \partial\mathcal{O}$, $I_\lambda(\xi_1) = I_\lambda(\xi_2)$ implies $I_\mu(\xi_1) = I_\mu(\xi_2)$. It follows that there is a continuous monotone surjection $I_{\lambda\mu} : \partial_\infty\lambda \to \partial_\infty\mu$ such that $I_\mu = I_{\lambda\mu} \circ I_\lambda$. More precisely, for any point $\xi \in \partial_\infty\lambda$, $I_{\lambda\mu}(\xi)$ is defined to be $Q_\mu(Q_\lambda^{-1}(\xi))$. The maps I_λ and $I_{\lambda\mu}$ can be visualized as in Figure 10.

4. Markers and universal circles

The outline of this section is the following. We first recall the definition of a universal circle (Definition 4.1) and prove Theorem 1.1. Then we review in Section 4.1 the construction from [7] of a particular universal circle $\mathfrak{S}_{\text{left}}$, which we call the universal circle from leftmost sections, for any taut foliation. The circle $\mathfrak{S}_{\text{left}}$ arises from a collection of special sections, called the leftmost sections, of a circle bundle E_∞ over Λ whose fibers are the circle at infinity of the leaves. The construction will then be examined carefully for our depth-one foliation $\mathcal{F}$ in Section 4.2, where we study what a leftmost section looks like inside a product region, and in Section 4.3, where we analyze the behavior of a leftmost section at adjacent type-0 and type-1 leaves. The punchlines of this section are Lemmas 4.14 and 4.15, where we show that the leftmost sections can be determined by the structure of $\partial\mathcal{O}$ developed in Section 3.

The following axiomatic definition of a universal circle for $\mathcal{F}$ first appears in [7]. It is worth remarking that although condition (2) seems not at all natural at first glance, it provides the universal circle with more interesting structures. In particular, it is necessary for the construction of invariant laminations in [5] (cf. Theorem 1.6).

Definition 4.1 (Universal circle). A *universal circle* for $\mathcal{F}$ is a circle $\mathfrak{S}$ with a faithful $\pi_1(M)$-action and a monotone map $U_\lambda : \mathfrak{S} \to \partial_\infty\lambda$, called a *structure map*, for any leaf λ of $\widetilde{\mathcal{F}}$ such that

(1) for any leaf λ and any $\gamma \in \pi_1(M)$, the following diagram commutes:

$$\begin{array}{ccc} \mathfrak{S} & \xrightarrow{U_\lambda} & \partial_\infty(\lambda) \\ \gamma\downarrow & & \downarrow\gamma \\ \mathfrak{S} & \xrightarrow{U_{\gamma\lambda}} & \partial_\infty(\gamma\lambda) \end{array}$$

(2) if λ and μ are incomparable leaves, then the core of U_{λ_1} is contained in a single gap of U_{λ_2}, and vice versa.

Two universal circles $\{\mathfrak{S}, U_\lambda\}$ and $\{\mathfrak{S}', U'_\lambda\}$ are isomorphic if there is a $\pi_1(M)$-equivariant homeomorphism $h : \mathfrak{S} \to \mathfrak{S}'$ such that $U'_\lambda \circ h = U_\lambda$ for all λ.

Proof of Theorem 1.1. All the conditions of a universal circle in Definition 4.1 are obvious by properties of $\partial\mathcal{O}$ and the way we define I_λ except for condition (2). Suppose λ and μ are incomparable, then their shadows are disjoint. Otherwise, there is an orbit of the flow $\widetilde{\phi}$ intersecting both leaves, contradicting their incomparability. Condition (2) is easily seen to be satisfied. ∎

4.1. Calegari–Dunfield's construction

In [7], Calegari–Dunfield describe an explicit construction of a universal circle $\mathfrak{S}_{\text{left}}$ for any taut foliation. We briefly review their construction below. For simplicity, we will stick to our $\mathcal{F}$ instead of more general taut foliations.

The bundle E_∞ is a circle bundle over Λ whose fiber at any leaf λ is the circle at infinity $\partial_\infty\lambda$. The topology of E_∞ is defined as follows. For any transversal τ of $\widetilde{\mathcal{F}}$, τ embeds into Λ, and we identify τ with its embedding image in Λ. The unit tangent bundle of $\widetilde{\mathcal{F}}$ restricted to τ is the circle bundle $\mathrm{UT}\,\widetilde{\mathcal{F}}|_\tau$, and there is a natural map $\mathrm{UT}\,\widetilde{\mathcal{F}}|_\tau \to E_\infty|_\tau$ sending a tangent vector of a leaf to the ideal point it points toward. We require the map to be a homeomorphism. It is shown in [7] that this topology is well defined, that is, independent of the choice of τ.

Since $\mathcal{F}$ is a taut foliation, there is an $\varepsilon_1 > 0$ such that every leaf of $\widetilde{\mathcal{F}}$ is quasi-isometrically embedded in its ε_1-neighborhood by [7, Lemma 2.4]. By the structure of depth-one foliations, there is a constant $\varepsilon_2 > 0$ so that the ε_2-neighborhood of $\mathcal{F}^0$ is contained in a spiraling neighborhood. Fix ε_0 to be $\min\{\varepsilon_1/3, \varepsilon_2\}$.

Definition 4.2. A *marker* for $\widetilde{\mathcal{F}}$ is an embedding

$$m : [0,1] \times \mathbb{R}_{\geq 0} \to \widetilde{M}$$

such that

- for any $x \in [0,1]$, $m(\{x\} \times \mathbb{R}_{\geq 0})$ is a geodesic ray in a leaf of $\widetilde{\mathcal{F}}$;
- for any $y \in \mathbb{R}_+$, $m([0,1] \times \{y\})$ is a transversal with length bounded by ε_0.

Any marker m gives a section s of $E_\infty|_\tau$, where τ is the image of $m([0,1] \times \{y\})$ in Λ, such that for any leaf $\lambda \in \tau$, $s(\lambda)$ is the ideal endpoint of $\mathrm{Image}(m) \cap \lambda$. The image of τ under s is called the *end* of m.

Note that our choice of ε_0 is different from but no larger than the constant chosen in [7]. Shrinking the constant will not affect the main results in their paper. In general, different ε might give rise to different universal circles, but in our case, it can be seen that for ε_0 small enough (i.e., smaller than ε_2 above), the leftmost universal circles are all isomorphic.

Lemma 4.3 ([7]). *Given two marker ends, either they are disjoint or their union is an embedded closed interval in E_∞ transverse to fibers.*

A point $\xi \in \partial\lambda$ is called a *marker endpoint* if there is a marker m so that the end of m intersects $\partial_\infty\lambda$ at ξ. The following theorem was originally announced by Thurston in an unpublished manuscript [32], and the proof is carefully written down in [7]. Heuristically, it says that the leaves of $\widetilde{\mathcal{F}}$ stay close in many directions.

Theorem 4.4 (Thurston's leaf pocket theorem). *For any leaf λ of $\widetilde{\mathcal{F}}$, the set of marker endpoints in $\partial_\infty\lambda$ is dense.*

For any leaf λ of $\widetilde{\mathcal{F}}$ and any point $\xi \in \partial_\infty\lambda$, there is a special section s_ξ of E_∞, called the *leftmost section starting from* ξ, built as follows.

In Λ, there is a neighborhood τ of λ homeomorphic to a closed interval, and $E_\infty|\tau$ is a cylinder. We adopt the convention that the flow $\widetilde{\phi}$ is flowing upward, and we are facing the cylinder $E_\infty|\tau$ from the outside. Take a finite collection C of marker ends in $E_\infty|\tau$ so that each fiber intersects at least one element in C. This is possible by Theorem 4.4. We build a path α_C by starting from ξ, heading left horizontally in a fiber until we hit the first marker end in C and following the marker end to move upward. After we reach the top of the marker end, we turn left again, staying in a fiber until we hit the next marker end in C, and follow the same rules to move on until we reach the top of $E_\infty|_\tau$. We call this the leftmost-up rule, following [7]. We can also move downward from ξ, but in the rightmost-down way. This procedure gives us a staircase path α_C in $E_\infty|_\tau$, which is an approximation to s_ξ (Figure 11).

To go from the staircase approximations to the leftmost section s_ξ, we define $s_\xi|_\tau$ to be the (rightmost) supremum above ξ and the (leftmost) infimum below ξ among all possible α_C. To be precise, we view $E_\infty|_\tau$ as $\tau \times (\mathbb{R}/\mathbb{Z})$. For a leaf $\mu \in \tau$ above λ, we define

$$s_\xi(\mu) = \sup_C(\min(\alpha_C \cap \partial_\infty\mu)).$$

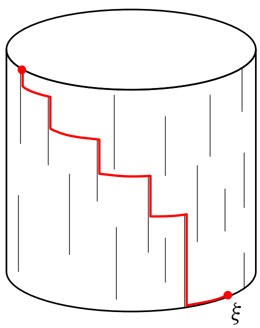

Figure 11. Approximating the leftmost section on $E_\infty|_\tau$ by starting from ξ and going leftmost up.

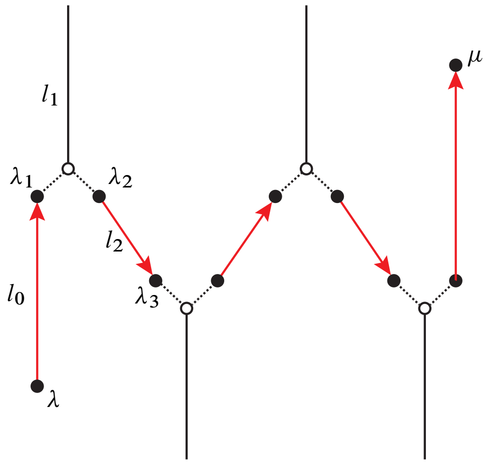

Figure 12. Extending a leftmost section to incomparable leaves.

For $\mu \in \tau$ below λ, define

$$s_\xi(\mu) = \inf_C(\max(\alpha_C \cap \partial_\infty \mu)).$$

It was proved in [7] that the supremum and the infimum exist, and s_ξ is indeed a continuous section of E_∞ over τ. We can define s_ξ for all leaves comparable to λ following this procedure.

Finally, we can branch out in Λ by turning around to reach incomparable leaves where s_ξ is not yet defined. More precisely, suppose μ is a leaf incomparable to λ. There is a sequence of leaves

$$\lambda_0 = \lambda, \lambda_1, \ldots, \lambda_n = \mu$$

so that λ_{2i} and λ_{2i+1} are comparable and λ_{2i+1} and λ_{2i+2} are non-separated. To illustrate the idea, we assume λ_1 is above λ (see Figure 12 for the case when $n = 9$). In this case, there is a product region $\widetilde{\Omega}_1$ so that $\lambda_1, \lambda_2 \in \partial^- \widetilde{\Omega}_1$. Let l_0 be the segment $[\lambda, \lambda_1]$ in Λ, and let l_1 be the image of $\widetilde{\Omega}_1$ in Λ. By the above construction, s_ξ is already defined over l_0, λ_1 and l_1. It is a consequence of Lemma 4.3 and Theorem 4.4 that $s_\xi|_{l_1}$ has a well-defined endpoint at $\partial_\infty \lambda_2$ (see [7, Lemma 6.18]). We extend s_ξ to λ_2 continuously, and follow the rightmost-down rule to define it over the segment $l_2 := [\lambda_2, \lambda_3]$. We continue along the sequence λ_i until we have defined $s_\xi(\mu)$. Since the dual graph Λ^* is a tree, there is a unique way to reach any incomparable μ from λ through such a sequence of λ_n. In the end, we have a section s_ξ that is well defined on the whole Λ.

The process of extending s_ξ is a process of branching out from λ and sweeping Λ. The values at leaves that are closer to λ are defined first, and the values at leaves farther away from λ are determined by the closer values. At each point of Λ, there is a direction of extension of s_ξ that points toward the direction away from λ, along which s_ξ is defined.

There is a unique leftmost section starting from any point in E_∞. The set of leftmost sections is denoted by LS. The images of two different leftmost sections might coalesce but can never cross each other. If ℓ is a line in Λ, the bundle $E_\infty|_\ell$ is homeomorphic to a cylinder, and the leftmost sections restricted to ℓ give embedded lines on $E_\infty|_\ell$ transverse to the fiber. For any three different leftmost sections, there is an embedded line in Λ so that the restrictions of the sections to this line have a well-defined cyclic order, and the cyclic order is independent of the choice of the line [7, Lemma 6.25]. The completion of LS with respect to the cyclic order is homeomorphic to a circle, denoted by $\mathfrak{S}_{\text{left}}$. The fundamental group $\pi_1(M)$ acts naturally on LS, and the action extends to an action on $\mathfrak{S}_{\text{left}}$. For any $\lambda \in \Lambda$, there is a valuation map $U_\lambda : \text{LS} \to \partial_\infty \lambda$ given by

$$U_\lambda(s) = s(\lambda).$$

The map U_λ can be extended to a monotone map $U_\lambda : \mathfrak{S}_{\text{left}} \to \partial_\infty \lambda$.

Theorem 4.5 ([7]). *The circle $\mathfrak{S}_{\text{left}}$ together with the $\pi_1(M)$-action and the set of structure maps $\{U_\lambda\}_{\lambda\in\Lambda}$ is a universal circle for $\mathcal{F}$.*

In order to prove Theorem 1.2, we need to analyze more carefully what marker ends and the leftmost sections look like on E_∞. We will continue to use the terminologies from Section 3.

4.2. Markers contained in a product region

We first consider the ends of markers that are contained in a product region. The identification of $\widetilde{\Omega}$ with $\widetilde{L} \times \mathbb{R}$ gives a canonical identification of $E_\infty|_{\widetilde{\Omega}}$ as $\partial_\infty \widetilde{L} \times \mathbb{R}$. Here we implicitly use that for any homeomorphism between two infinite-type surfaces with standard hyperbolic structures, any lift to their universal covers extends continuously to a homeomorphism between their boundaries at infinity, and the extension is unique [10]. Denote the leaf $\widetilde{L} \times \{t\}$ by λ_t. Again, each λ_t is identified with $\widetilde{L}$.

Lemma 4.6. *For any $\xi \in \partial_\infty \lambda_t$, there is an $\varepsilon > 0$ so that $\{\xi\} \times [t, t+\varepsilon]$ is the end of a marker.*

Proof. We will use a tightening method described in [7, Section 5.3]. Take any point $x \in \lambda_t$ and consider the geodesic ray γ from x to ξ. Since depth-one leaves in the same fibered region have asymptotic ends, for every small $\delta > 0$, there is an ε such that any flowline of $\widetilde{\phi}$ between λ_t and $\lambda_{t+\varepsilon}$ has length $< \delta$. Moreover, the map from λ_t to any λ_s with $s \in (t, t+\varepsilon]$ induced by flowing λ_t to λ_s is K-bi-Lipschitz for some uniform K. Therefore, the flow image γ_s of γ in each such λ_s is a family of uniform quasi-geodesics with ideal endpoints (ξ, s). We can then tighten γ_s to geodesics γ_s^* on λ_s. We claim that the union of the λ_s^* for $s \in [t, t+\varepsilon]$ is a continuous one-ended band with bounded width. That is because any pair of γ_{s_1} and γ_{s_2} are bounded Hausdorff distance from each other, and so are their geodesic tightenings $\gamma_{s_1}^*$ and $\gamma_{s_2}^*$. By the continuity of the leafwise hyperbolic metric, this only happens when γ_s^* is a continuous family of geodesic rays. The union of γ_s^* has bounded width because the tightening process only shifts the rays by a bounded amount. Finally, we can take ε even smaller to obtain a genuine marker with width $< \varepsilon_0$ and with the end $\{\xi\} \times [t, t+\varepsilon]$. ■

The next corollary follows immediately from the construction of leftmost sections on comparable leaves.

Corollary 4.7 (Leftmost sections on a product region are vertical). *Suppose s is any leftmost section of E_∞. If $s(\lambda_{t_0}) = (\xi, t_0)$ for some $t_0 \in \mathbb{R}$, then $s(\lambda_t) = (\xi, t)$ for all t.*

Proof. By the construction of leftmost sections, one can see that any leftmost section s has the following property: when extending s upward, if s meets the end e of a marker at a point $x \in e$, then s has to contain the part of e above x. The same is true if we are extending s downward: if s meets e at a point x, it must also contain everything in e below x. Since through any point in $E_\infty|_{\widetilde{\Omega}}$ there is a vertical marker end in both directions, the leftmost section s restricted to $\widetilde{\Omega}$ is forced to be vertical. ■

4.3. Markers intersecting a type-0 leaf

We now consider the ends of markers intersecting a type-0 leaf. Let μ be a type-0 leaf covering a compact leaf Σ, and assume for the rest of this section that μ is in the positive boundary of $\widetilde{\Omega}$. The case where μ is in the negative boundary of $\widetilde{\Omega}$ is similar.

As in the previous discussion, we identify every depth-one leaf in Ω with L, and every type-1 leaf in $\widetilde{\Omega}$ with $\widetilde{L}$. This gives us a homeomorphism between $E_\infty|_{\widetilde{\Omega}}$ and $\partial_\infty \widetilde{L} \times \mathbb{R}$. Recall that in Section 3, we define a continuous map

$$I_{\lambda_t \mu} : \partial_\infty \lambda_t \to \partial_\infty \mu$$

for any $\lambda_t \in \widetilde{\Omega}$. Under the homeomorphism $\partial_\infty \lambda_t \cong \partial_\infty \widetilde{L}$, the map $I_{\lambda_t \mu}$ is the same map for any t when viewed as a map from $\partial_\infty \widetilde{L}$ to $\partial_\infty \mu$. We denote this map by $I_{\widetilde{\Omega}\mu}$.

In Λ, $\widetilde{\Omega} \cup \mu$ is a half-open interval. Let $\mathcal{M}_{\mu,\widetilde{\Omega}}$ be the set of markers m in $\mu \cup \widetilde{\Omega}$ with one side lying in μ. By Corollary 4.7 and Lemma 4.3, the end of such an m intersects $\partial_\infty \mu$ at a single point $\xi_\mu(m)$ and intersects $E_\infty|_{\widetilde{\Omega}}$ in a vertical segment

$$\{(\xi_{\widetilde{\Omega}}(m), t) \mid t > T\},$$

where T is a constant depending on m, and $\xi_{\widetilde{\Omega}}(m)$ is a point in $\partial_\infty \widetilde{L}$ depending on m and independent of t (Figure 13).

Define the set $\xi_{\widetilde{\Omega}}(\mathcal{M}_{\mu,\widetilde{\Omega}}) := \{\xi_{\widetilde{\Omega}}(m) \mid m \in \mathcal{M}_{\mu,\widetilde{\Omega}}\} \subset \partial_\infty \widetilde{L}$. Intuitively, these are the directions on $\widetilde{L}$ in which μ does not diverge from λ_t.

Recall that the positive escaping set $\mathcal{U}^+ \subset L$ is the set of points whose ϕ-orbit escapes Ω in positive time. Let $\widetilde{\mathcal{U}}^+$ be the preimage of $\mathcal{U}^+$ in $\widetilde{L}$. There is a component $\widetilde{\mathcal{U}}^+_\mu$ of $\widetilde{\mathcal{U}}^+$ so that $x \in \lambda_t$ is a point of $\widetilde{\mathcal{U}}^+_\mu$ if and only if $\widetilde{\phi}(x)$ hits μ. The hitting map $H_\mu : \widetilde{\mathcal{U}}^+_\mu \to \mu$ defined by

$$H_\mu(x) = \widetilde{\phi}(x) \cap \mu$$

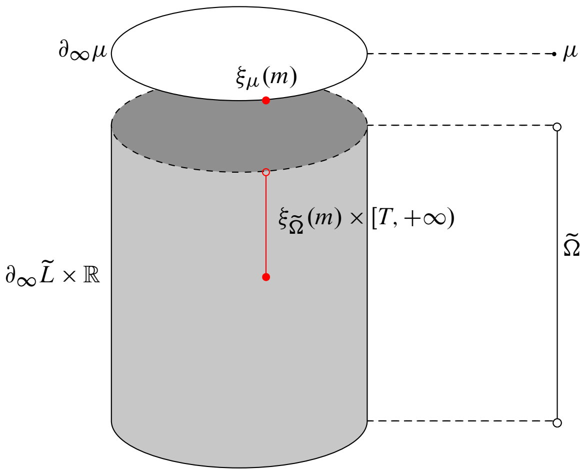

Figure 13. The end of a marker m (in red) in $\mathcal{M}_{\mu,\widetilde{\Omega}}$.

is a homeomorphism since it is an open bijection. It is tautological that

$$p|_\mu \circ H_\mu = p|_{\widetilde{\mathcal{U}}_\mu^+}, \tag{4.1}$$

where p is the projection to $\mathcal{O}$ as in Section 3.

Lemma 4.8. *Let α^* be an oriented simple closed geodesic in Σ with trivial $\mathcal{F}$-holonomy on the side of Ω, and let $\widetilde{\alpha}^*$ be a lift of α^* to μ. Denote the forward endpoint at infinity of $\widetilde{\alpha}^*$ by $\partial^+\widetilde{\alpha}^*$. Then there is a marker $m \in \mathcal{M}_{\mu,\widetilde{\Omega}}$ so that $\xi_\mu(m) = \partial^+\widetilde{\alpha}^*$ and $\xi_{\widetilde{\Omega}}(m) \in I^{-1}_{\widetilde{\Omega}\mu}(\partial^+\widetilde{\alpha}^*)$.*

Proof. Perform a homotopy of α^* along ϕ into Ω for a short distance so that the final image is a simple closed curve on a depth-one leaf in Ω. This is possible because α^* has trivial $\mathcal{F}$-holonomy on the side of Ω. The full homotopy image is an annulus, denoted by A. We lift A to $\widetilde{M}$ to get a two-ended infinite band $\widetilde{A}$ in $\widetilde{\Omega} \cup \mu$ with one side being $\widetilde{\alpha}^*$ and the other side on λ_T for some T. Note that A has bounded width because it is a lift of an annulus.

Fix a base point $x \in \widetilde{\alpha}^*$, and let $\widetilde{\alpha}^*(x)$ be the oriented ray in $\widetilde{\alpha}^*$ starting from x toward $\partial^+\widetilde{\alpha}^*$. We restrict the lifted homotopy to the ray $\widetilde{\alpha}^*(x)$, and the restricted homotopy image is a one-ended infinite band $B \subset \widetilde{A}$ (Figure 14). The band B has finite width between μ and λ_T, and we can make the width arbitrarily small by cutting the homotopy at $\lambda_{T'}$ for large enough T'. When $t > T'$, the ray $B \cap \lambda_t$ has a well-defined endpoint at infinity because it is contained in a lift of a simple closed curve on a hyperbolic depth-one leaf. Similar to the proof of Lemma 4.6 and the tightening method in [7, Section 5.3], by taking even larger T' and pulling tight the intersection of B and λ_t, we get a marker $m \in \mathcal{M}_{\mu,\widetilde{\Omega}}$. Note that different from Lemma 4.6, here we do not have uniform bi-Lipschitz maps

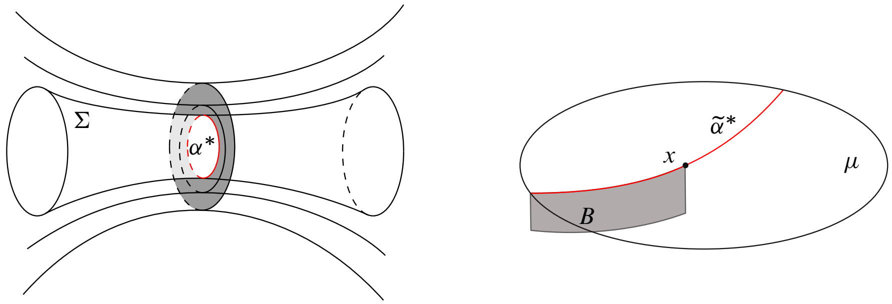

Figure 14. An annulus without holonomy (shaded on the left) lifts to a band B (shaded on the right) that gives a marker.

from μ to type-1 leaves in $\widetilde{\Omega}$. However, we still have that the rays $B \cap \lambda_t$ are a family of uniform quasi-geodesics by the continuity of the leafwise metric, so the same tightening argument as in [7, Section 5.3] still works. It is clear from the construction that $\xi_\mu(m) = \partial^+\widetilde{\alpha}^*$. The point $\xi_{\widetilde{\Omega}}(m)$ is the endpoint of the ray

$$B \cap \lambda_t = H_\mu^{-1}(\widetilde{\alpha}^*(x)).$$

By (4.1), the projections of $\widetilde{\alpha}^*(x)$ and $H_\mu^{-1}(\widetilde{\alpha}^*(x))$ to $\mathcal{O}$ are identical, and both rays escape to infinity in the leaf. This implies $\partial^+\widetilde{\alpha}^* \in \partial p(\mu) \cap \partial p(\lambda_t)$ and

$$\xi_{\widetilde{\Omega}}(m) \in I_{\widetilde{\Omega}\mu}^{-1}(\partial^+\widetilde{\alpha}^*). \qquad \blacksquare$$

To state the next lemma, we need one more definition. A simple closed curve $\beta \subset L$ is called a Σ-*juncture* if β is a connected positive f-juncture and β covers a simple closed curve in Σ under the covering map $\mathcal{U}^+ \to \partial^+\Omega$ (note that Σ is only a component of $\partial^+\Omega$).

Corollary 4.9. *Let $\beta \subset L$ be a Σ-juncture, and let $\widetilde{\beta}$ be a lift of β in $\widetilde{L}$ that lies in $\widetilde{\mathcal{U}}_\mu^+$. Then both endpoints of $\widetilde{\beta}$ are in $\xi_{\widetilde{\Omega}}(\mathcal{M}_{\mu,\widetilde{\Omega}})$.*

Proof. Suppose β covers a simple closed curve α on Σ. Fix an orientation of β, which induces an orientation of α. The curve α has trivial $\mathcal{F}$-holonomy on the side of Ω because β is an f-juncture. For any T, we can view $\widetilde{\beta}$ as an embedded line in λ_T. There is a lift $\widetilde{\alpha} \subset \mu$ of α given by $\widetilde{\alpha} = H_\mu(\widetilde{\beta})$. Let $\partial^+\widetilde{\alpha}$ be the forward endpoint of $\widetilde{\alpha}$. Applying Lemma 4.8 to the geodesic tightening α^* of α, we obtain a marker $m \in \mathcal{M}_{\mu,\widetilde{\Omega}}$. But since $\widetilde{\alpha}$ is obtained by flowing $\widetilde{\beta}$ to μ, the proof of Lemma 4.8 implies that $\xi_{\widetilde{\Omega}}(m)$ is exactly the forward endpoint of $\widetilde{\beta}$. The same proof applies for the backward endpoint of $\widetilde{\beta}$, proving the corollary. $\blacksquare$

We define the limit set $\mathcal{L}(\widetilde{\mathcal{U}}_\mu^+)$ of $\widetilde{\mathcal{U}}_\mu^+$ as the intersection $\mathrm{cl}(\widetilde{\mathcal{U}}_\mu^+) \cap \partial_\infty\widetilde{L}$, where $\mathrm{cl}(\widetilde{\mathcal{U}}_\mu^+)$ is defined to be the closure taken in $\widetilde{L} \cup \partial_\infty\widetilde{L} \cong \overline{\mathbb{H}^3}$.

Lemma 4.10. *The closure in $\partial_\infty\widetilde{L}$ of $\xi_{\widetilde{\Omega}}(\mathcal{M}_{\mu,\widetilde{\Omega}})$ is $\mathcal{L}(\widetilde{\mathcal{U}}_\mu^+)$.*

Proof. Suppose m is a marker in $\mathcal{M}_{\mu,\widetilde{\Omega}}$. The intersection of the image of m with any λ_t is a ray r_t. After identifying λ_t with L, we claim that r_t is contained in $\widetilde{\mathcal{U}}^+_\mu$. This is because the image of m is ε_0-close to μ. By our choice of ε_0, the ε_0-neighborhood of Σ in M is contained in an $\mathcal{F}$-spiraling neighborhood $N(\Sigma)$ of Σ, and every ϕ-orbit intersecting the spiraling neighborhood will hit Σ. So r_t is contained in $\widetilde{\mathcal{U}}^+_\mu$, and the ideal endpoint of r_t, which is exactly $\xi_{\widetilde{\Omega}}(m)$ by definition, will then be contained in $\mathcal{L}(\widetilde{\mathcal{U}}^+_\mu)$. This shows $\xi_{\widetilde{\Omega}}(\mathcal{M}_{\mu,\widetilde{\Omega}}) \subset \mathcal{L}(\widetilde{\mathcal{U}}^+_\mu)$.

By Corollary 4.9, it now suffices to show that the endpoints of lifts of any Σ-juncture in $\widetilde{\mathcal{U}}^+_\mu$ are dense inside $\mathcal{L}(\widetilde{\mathcal{U}}^+_\mu)$. We first recall some known facts and constructions. By [11], the geodesic tightenings of a system of positive f-junctures limit to the negative Handel–Miller geodesic lamination Λ^-_{HM} under negative iterations of f, and Λ^-_{HM} is independent of the choice of junctures. On the other hand, the intersection of $\mathcal{F}^s$ with L induces a singular foliation $\mathcal{F}^s_L$ on L. We define W^- as the restriction of $\mathcal{F}^s_L$ to the complement of $\mathcal{U}^+$. The complement of $\mathcal{U}^+$ is saturated by leaves of $\mathcal{F}^s_L$ and W^- is a singular sublamination of $\mathcal{F}^s_L$ by [23]. Let $\widetilde{W}^-$ be the lift of W^- to $\widetilde{L}$. The singular lamination $\widetilde{W}^-$ determines an abstract lamination on $\partial_\infty \widetilde{L}$ by a standard construction (Section 2). It follows from [23, Theorem 8.4] that $\widetilde{W}^- = \Lambda^-_{\mathrm{HM}}$ as abstract laminations (the paper proves it for circular pseudo-Anosov flows, but the same method applies to any pseudo-Anosov flow without perfect fits).

Now fix a Σ-juncture β. Suppose that β is the boundary of a contracting neighborhood of a contracting end $\mathcal{E}$, and n is the smallest positive integer so that $f^n(\mathcal{E}) = \mathcal{E}$. The above facts imply that the endpoints of $\partial\widetilde{\mathcal{U}}^+_\mu$ can be approximated by endpoints of lifts of $\{f^{-in}(\beta)\}_{i\geq 0}$. The limit set $\mathcal{L}(\widetilde{\mathcal{U}}^+_\mu)$ is nowhere dense by the next lemma (Lemma 4.11), so points in $\mathcal{L}(\widetilde{\mathcal{U}}^+_\mu)$ are approximated by endpoints of $\partial\mathcal{L}(\widetilde{\mathcal{U}}^+_\mu)$. Also note that if β is a Σ-juncture, so is $f^{-in}(\beta)$. Therefore, the endpoints of lifts of any Σ-juncture in $\widetilde{\mathcal{U}}^+_\mu$ are dense inside $\mathcal{L}(\widetilde{\mathcal{U}}^+_\mu)$, which completes the proof. ∎

We give a proof of the following fact used in the proof of Lemma 4.10.

Lemma 4.11. *The limit set $\mathcal{L}(\widetilde{\mathcal{U}}^+_\mu)$ is nowhere dense in $\partial_\infty \widetilde{L}$.*

Proof. The subsurface $\mathcal{U}^+ \subset L$ is bounded by leaves of $\mathcal{F}^s_L$ [23], each of which has distinct well-defined endpoints at infinity. In particular, $\mathcal{L}(\widetilde{\mathcal{U}}^+_\mu)$ is a proper subset of $\partial_\infty \widetilde{L}$. If $\mathcal{L}(\widetilde{\mathcal{U}}^+_\mu)$ contains a non-trivial interval I, we can find a hyperbolic element $g \in \pi_1(L)$ with a fixed point in I. Here we are using the fact that L has bounded injectivity radius. But $g(\widetilde{\mathcal{U}}^+_\mu)$ is either disjoint from or equal to $\widetilde{\mathcal{U}}^+_\mu$ since $\mathcal{U}^+$ is embedded. So $g^n(I)$ is contained in $\mathcal{L}(\widetilde{\mathcal{U}}^+_\mu)$ for any integer n, a contradiction. ∎

Lemma 4.12. *There is a subset $\mathcal{M}'_{\mu,\widetilde{\Omega}}$ of $\mathcal{M}_{\mu,\widetilde{\Omega}}$ such that*

(1) *the set $\xi_\mu(\mathcal{M}'_{\mu,\widetilde{\Omega}})$ is dense in $\partial_\infty \mu$;*

(2) *the set $\xi_{\widetilde{\Omega}}(\mathcal{M}'_{\mu,\widetilde{\Omega}})$ is dense in $\mathcal{L}(\widetilde{\mathcal{U}}^+_\mu)$;*

(3) *for any* $m \in \mathcal{M}'_{\mu,\widetilde{\Omega}}$*, we have*

$$\xi_\mu(m) = I_{\widetilde{\Omega}\mu}(\xi_{\widetilde{\Omega}}(m)).$$

Proof. The subset $\mathcal{M}'_{\mu,\widetilde{\Omega}}$ can be taken to be the set of markers in $\mathcal{M}_{\mu,\widetilde{\Omega}}$ that arise from Lemma 4.8. Item (3) is exactly the content of Lemma 4.8. Item (1) is true because the endpoints of the lifts of a simple closed curve are dense in $\partial_\infty \mu$. The second part of the proof of Lemma 4.10 only used markers in $\mathcal{M}'_{\mu,\widetilde{\Omega}}$, so we have already proved item (2). ■

Lemma 4.13. *If* η *is a point in* $\partial_\infty \mu$ *such that* $I^{-1}_{\widetilde{\Omega}\mu}(\eta)$ *is a closed interval with positive length, then* $\eta \notin \xi_\mu(\mathcal{M}_{\mu,\widetilde{\Omega}})$*.*

Proof. Suppose there is a marker $m \in \mathcal{M}_{\mu,\widetilde{\Omega}}$ such that $\xi_\mu(m) = \eta$. Let r_μ be the geodesic ray $m \cap \mu$, and let r_t be the geodesic ray $m \cap \lambda_t$. Let e be the side of $p(\mu)$ corresponding to η. Then the interior of e is contained in $p(\lambda_t)$ for all t. Take a quasi-geodesic ray r' in μ such that $p(r')$ has an endpoint in the interior of e. For example, using the notations in Proposition 3.1, for any $x \in e$, we can take r' to be a ray in l_x with one end at x. Since r' and r_μ both have endpoint μ, they have bounded Hausdorff distance from each other on μ. By the choice of ε_0 in Definition 4.2, each r_t in m flows forward in bounded time to hit μ, and the image on μ is exactly $H_\mu(r_t)$. Therefore, r' has bounded Hausdorff distance from r_t and $H_\mu(r_t)$ as well. Since ε_0 is smaller than a separation constant (Definition 4.2), r' and $H_\mu(r_t)$ are also a bounded Hausdorff distance from each other in μ.

Now fix a λ_t intersecting m, and for any $y \in \mu$, define $T(y)$ to be the time it takes for y to flow backward to λ_t. If we write $\widetilde{\phi}^T$ for flowing along $\widetilde{\phi}$ for time T, then $T(y)$ satisfies $\phi^{-T(y)}(y) \in \lambda_t$. This function T is bounded on $H_\mu(r_t)$ by the above. On the other hand, since we pick r' to have an endpoint in the interior of $p(\lambda_t)$, $T(y)$ must go to $+\infty$ as y travels along r' to infinity. We will show that this is a contradiction. To see this, take sequences $y_n \in H_\mu(r_t)$ and $z_n \in r'$ with $d_\mu(y_n, z_n) \leq C$, where C is a constant, and with x_n (and hence y_n) going to infinity. Then $T(y_n)$ is bounded, while $T(z_n)$ goes to $+\infty$. By acting with deck transformations in $\pi_1(\mu)$, we can find translates (y'_n, z'_n) of (y_n, z_n) that stay in a compact subset of μ. Since the transformations preserve $\widetilde{\mathcal{F}}$, $\widetilde{\phi}^{-T(y_n)}(y'_n)$ and $\widetilde{\phi}^{-T(z_n)}(z'_n)$ are in the same type-1 leaf in $\widetilde{\Omega}$. Assume y'_n converges to y and z'_n converges to z, and up to taking a subsequence, assume that $T(y_n)$ converges to a finite positive number T_0. Then the flowline $\widetilde{\phi}(z)$ will never intersect the type-1 leaf containing $\widetilde{\phi}^{-T_0}(y)$, which is a contradiction because every flowline intersecting $\widetilde{\Omega}$ will intersect every leaf in $\widetilde{\Omega}$. This proves that such a marker m does not exist. ■

The following lemmas, Lemmas 4.14 and 4.15, together with Corollary 4.7, will give us the complete rules to build the leftmost section starting from a given point. Lemma 4.14 tells us how to extend the leftmost section from a product region to an adjacent type-0 leaf, and Lemma 4.15 tells us how to go from a type-0 leaf to an adjacent product region.

For any $\xi \in \partial_\infty \widetilde{L}$, we use c_ξ to denote the vertical section of $E_\infty|_{\widetilde{\Omega}}$ given by $c_\xi(\lambda_t) = (\xi, t)$.

Lemma 4.14. *For any $\xi \in \partial_\infty(\widetilde{L})$, the vertical section c_ξ extends continuously to μ by setting $c_\xi(\mu) = I_{\widetilde{\Omega}\mu}(\xi)$.*

Proof. First of all, we remark that as noted in [7, Proof of Lemma 6.18], the closure of c_ξ in E_∞ is a closed interval transverse to the circle fibers, intersecting $\partial_\infty\mu$ in exactly one point. This is a consequence of the density of markers, and it implies the section c_ξ has a unique continuous extension to μ.

To determine the value of c_ξ at μ, we first consider the case where ξ is in a gap G_ξ of $I_{\widetilde{\Omega}\mu}$. Let $\partial_r G_\xi$ and $\partial_l G_\xi$ be the leftmost and the rightmost endpoints of G_ξ, respectively. By Lemma 4.12, we have sequences $\{m_n^+\}$ and $\{m_n^-\}$ in $\mathcal{M}'_{\mu,\widetilde{\Omega}} \subset \mathcal{M}_{\mu,\widetilde{\Omega}}$ with the following properties:

- the points $\xi_{\widetilde{\Omega}}(m_n^-)$ limit to $\partial_l G_\xi$ from the left and the points $\xi_{\widetilde{\Omega}}(m_n^+)$ limit to $\partial_r G_\xi$ from the right;
- we have $\xi_\mu(m_n^\pm) = I_{\widetilde{\Omega}\mu}(\xi_{\widetilde{\Omega}}(m_n^\pm))$.

These properties imply that the points $\xi_\mu(m_n^-)$ limit to $I_{\widetilde{\Omega}\mu}(\xi)$ from the left and the points $\xi_\mu(m_n^+)$ limit to $I_{\widetilde{\Omega}\mu}(\xi)$ from the right by item (3) of Lemma 4.12 and Lemma 4.3. The two sequences of marker ends pin down the endpoint of c_ξ to be $I_{\widetilde{\Omega}\mu}(\xi)$ (see the left-hand side of Figure 15).

When ξ is not in any gap of $I_{\widetilde{\Omega}\mu}$, the proof can be done similarly to above by replacing G_ξ, $\partial_r G_\xi$ and $\partial_l G_\xi$ all by the point ξ. ■

Lemma 4.15. *Let η be a point in $\partial_\infty\mu$. Suppose ν is any leaf of $\widetilde{\mathcal{F}}$, ξ is any point in $\partial_\infty\nu$ and s_ξ is the leftmost section starting from ξ. If the direction of extension of s_ξ at μ*

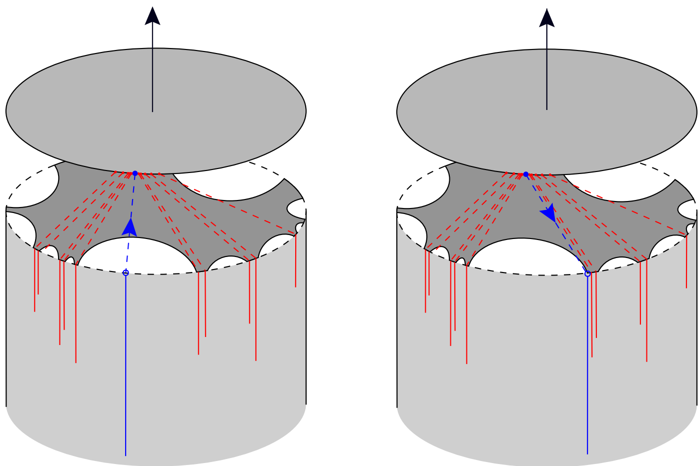

Figure 15. The subset $\mathcal{M}'_{\mu,\widetilde{\Omega}}$ (red) pins down the way to extend the leftmost sections (blue), with the blue arrows indicating the direction of extension. The left-hand side corresponds to Lemma 4.14 and the right-hand side corresponds to Lemma 4.15.

points from μ to $\widetilde{\Omega}$, that is, μ is closer to ν in Λ than $\widetilde{\Omega}$, and $s_\xi(\mu) = \eta$, then $s_\xi(\lambda_t)$ is the rightmost endpoint of $I_{\widetilde{\Omega}\mu}^{-1}(\eta)$. Here we view a singleton as a closed interval of length zero.

Remark 4.16. We remind the readers that here we are assuming μ to be in the positive boundary of $\widetilde{\Omega}$, as stated in the first paragraph of Section 4.3, so the direction of extension of s_ξ is the backward flow direction. If μ is in the negative boundary of $\widetilde{\Omega}$ and the direction of extension of s_ξ is the forward flow direction pointing from μ to $\widetilde{\Omega}$, the lemma remains true if we replace "the rightmost endpoint" by "the leftmost endpoint" with the same proof.

Proof. We refer the readers to Figure 15 for an illustration of the situation. The value of s_ξ on $\widetilde{\Omega}$ is determined by going down from η and following the rightmost-down rule, described in Section 4.1. Since, by Corollary 4.7, s_ξ can only go vertically down in $\widetilde{\Omega}$, its value in $\widetilde{\Omega}$ is the leftmost infimum of $\xi_{\widetilde{\Omega}}(m)$ over all $m \in \mathcal{M}_{\mu,\widetilde{\Omega}}$ such that $\xi_\mu(m)$ is in a small neighborhood U of η and $\xi_\mu(m)$ is not on the left of η (i.e., after identifying U with an interval in $\mathbb{R}$, $\eta \le \xi_\mu(m)$). By Lemma 4.12, there is a sequence of markers $m_n \in \mathcal{M}'_{\mu,\widetilde{\Omega}} \subset \mathcal{M}_{\mu,\widetilde{\Omega}}$ so that $\xi_\mu(m_n)$ limits to η from the right. By the monotonicity of $I_{\widetilde{\Omega}\mu}$ and item (3) of Lemma 4.12, we have that $\xi_{\widetilde{\Omega}}(m_n)$ limits to the rightmost endpoint of $I_{\widetilde{\Omega}\mu}^{-1}(\eta)$ from the right. If $I_{\widetilde{\Omega}\mu}^{-1}(\eta)$ is a single point, this implies that $s_\xi(\lambda_t) = I_{\widetilde{\Omega}\mu}^{-1}(\eta)$. If $I_{\widetilde{\Omega}\mu}^{-1}(\eta)$ is a closed interval, the existence of m_n and Lemma 4.13 imply $s_\xi(\lambda_t)$ is the rightmost endpoint of $I_{\widetilde{\Omega}\mu}^{-1}(\eta)$. ∎

5. Building the homeomorphism

Suppose λ is a leaf of $\mathcal{F}$ and ξ is a point in $\partial_\infty\lambda$. Let V_ξ denote $I_\lambda^{-1}(\xi) \subset \partial\mathcal{O}$, which is either a closed interval or a singleton. If V_ξ is a closed interval, let $\partial_l V_\xi$ and $\partial_r V_\xi$ be the leftmost and the rightmost endpoints of V_ξ, respectively. Note that by our convention, left means clockwise and right means counterclockwise. We say V_ξ is a stable (resp. unstable) gap of λ if $Q_\lambda^{-1}(\xi)$ is a boundary $\mathcal{F}^s_{\mathcal{O}}$-leaf (resp. $\mathcal{F}^u_{\mathcal{O}}$-leaf) of $p(\lambda)$. Note that a closed interval of $\partial\mathcal{O}$ cannot be a stable gap of a leaf while being an unstable gap of another. Indeed, a leaf of $\mathcal{F}^s_{\mathcal{O}}$ and a leaf of $\mathcal{F}^u_{\mathcal{O}}$ cannot bound an ideal bigon. In fact, for ϕ without perfect fits, a stronger statement is true: any two rays in $\mathcal{F}^s_{\mathcal{O}} \cup \mathcal{F}^u_{\mathcal{O}}$ have distinct endpoints at infinity [13, Lemma 3.20]. We say a closed interval of $\partial\mathcal{O}$ is a stable (resp. an unstable) gap if it is a stable (resp. an unstable) gap of a leaf of $\widetilde{\mathcal{F}}$.

Since $\mathfrak{S}_{\text{left}}$ is the completion of the set LS of leftmost sections, to define the homeomorphism $T : \mathfrak{S}_{\text{left}} \to \partial\mathcal{O}$, it suffices to define T on LS and show that it admits an extension. For any point ξ in E_∞, let s_ξ be the leftmost section starting from ξ. The set LS* of *pointed leftmost sections* is the set

$$\{(s, \xi) \mid s \in \text{LS},\ \xi \in E_\infty,\ s = s_\xi\}.$$

We point out that this definition is not redundant, for it is possible to have $s_\xi = s'_\xi$ with $\xi \neq \xi'$. There is a natural forgetful map $\pi : \mathrm{LS}^* \to \mathrm{LS}$ given by forgetting the starting point. We define a map $T^* : \mathrm{LS}^* \to \partial\mathcal{O}$ as follows. For a pointed leftmost section (s, ξ), consider V_ξ. If V_ξ is a single point, define $T^*(s, \xi) = V_\xi$. If V_ξ is an unstable gap, define $T^*(s, \xi) = \partial_l V_\xi$. If V_ξ is a stable gap, define $T^*(s, \xi) = \partial_r V_\xi$.

Theorem 1.2 follows from the following theorem.

Theorem 5.1. *The map T^* descends to a map $T : \mathrm{LS} \to \partial\mathcal{O}$ that extends continuously to a map from $\mathfrak{S}_{\mathrm{left}}$ to $\partial\mathcal{O}$, which we will again denote by T. The extension T is injective and preserves the cyclic order; hence it is a homeomorphism. Moreover, T is $\pi_1(M)$-equivariant, and for any $\lambda \in \Lambda$ and any $s \in \mathfrak{S}_{\mathrm{left}}$, we have $U_\lambda(s) = I_\lambda(T(s))$. In other words, T is an isomorphism of universal circles between $\mathfrak{S}_{\mathrm{left}}$ and $\partial\mathcal{O}$.*

The rest of this section will be dedicated to proving Theorem 5.1.

Lemma 5.2. *Suppose (s, ξ_1) is an element in LS^*. Let μ be any leaf of $\widetilde{\mathcal{F}}$ and suppose $s(\mu) = \xi$. Then $T^*(s, \xi_1) \in V_\xi$.*

Proof. By definition, we have $s = s_{\xi_1}$ is the leftmost section starting from ξ_1, where $\xi_1 \in \partial_\infty\lambda_1$ for some $\lambda_1 \in \Lambda$. We assume that λ_1 is a type-1 leaf for simplicity. The case when λ_1 is type-0 is basically the same.

We first consider the case when λ_1 and μ are comparable and $\lambda_1 < \mu$. Take a sequence of $\widetilde{\mathcal{F}}$-leaves

$$\lambda_1 \to \lambda_2 \to \lambda_3 \to \cdots \to \lambda_n = \mu,$$

where λ_{2i} is a type-0 leaf, λ_{2i+1} is a type-1 leaf (note that here we set up the notations so that λ_k is of type k mod 2) and $\lambda_i \lesssim \lambda_{i+1}$. The sequence represents the shortest path from $\widetilde{\Omega}(\lambda_1)^*$ to μ^* in Λ^*, after identifying a type-0 leaf with the dual vertex and a type-1 leaf with the vertex dual to the product region containing it. We record the value of s along this sequence by $s_i = s(\lambda_i)$, and let $V_i = V_{s(\lambda_i)}$. We have a sequence of closed intervals (possibly with length zero) $V_i \subset \partial\mathcal{O}$. The goal is to show that for all $i = 1, \ldots, n$, we have $T^*(s, \xi_0) \in V_i$. In particular, this implies $T^*(s, \xi_0) \in V_n = V_\xi$. We will show this by tracking how the intervals V_i vary along the sequence λ_i. By Lemma 4.14, we have $V_{2i-1} \subset V_{2i}$; by Lemma 4.15, we have $V_{2i+1} \subset V_{2i}$.

Lemma 5.3. If $V_{2i-1} \subsetneq V_{2i}$, then V_{2i} is a stable gap. If $V_{2i+1} \subsetneq V_{2i}$, then V_{2i} is an unstable gap and $V_{2i+1} = \partial_l V_{2i}$.

Proof. First, suppose $V_{2i-1} \subsetneq V_{2i}$. The leaf $e := Q^{-1}_{\lambda_{2i}}(I_{\lambda_{2i}}(V_{2i}))$ is a side of $p(\lambda_{2i})$ containing in the interior of $p(\lambda_{2i-1})$. The fixed point x_e in e under $\mathrm{Stab}(e)$ corresponds to a periodic $\widetilde{\phi}$-orbit γ in $\widetilde{M}$ intersecting λ_{2i-1} but not intersecting λ_{2i}. Since every $\widetilde{\phi}$-orbit in $\widetilde{\mathcal{F}}^s(\phi)$ is forward asymptotic to γ, we see that $\mathcal{F}^s_{\mathcal{O}}(x)$ is not contained in $p(\lambda_{2i})$. This means $e = \mathcal{F}^s_{\mathcal{O}}(x)$ and V_{2i} is a stable gap.

Now suppose $V_{2i+1} \subsetneq V_{2i}$. A similar argument to the above shows V_{2i} is an unstable gap. By Lemma 4.15 and Remark 4.16, if V_{2i+1} is a closed interval of positive length, there will be a boundary leaf e of $p(\lambda_{2i})$ and a boundary leaf e' of $p(\lambda_{2i+1})$ different from e and sharing the leftmost endpoint with e. But this cannot happen because when ϕ has no perfect fits, no pair of leaves in $\mathcal{F}^s_{\mathcal{O}} \cup \mathcal{F}^u_{\mathcal{O}}$ can share an endpoint. Therefore, V_{2i+1} is a single point, and it is the leftmost endpoint of V_{2i} by Lemma 4.15. ∎

We continue the proof that $T^*(s, \xi_1) \in V_i$ for all $1 \leq i \leq n$. It is obvious that $T^*(s, \xi_1) \in V_1$ by the definition of T^*. If V_1 is a single point or a stable gap of λ_1, by Lemma 5.3 we have $V_i \subset V_{i+1}$ for all i. So we have $T^*(s, \xi_1) \in V_i$.

If V_1 is an unstable gap, then there are two cases. If for all $1 \leq i \leq n$, we have $V_i = V_1$, then there is nothing to prove. If this is not the case, let N be the first positive integer so that $V_N \neq V_1$. If $V_1 \subsetneq V_N$, then V_N is a stable gap by Lemma 5.3, and we again have $V_i \subset V_{i+1}$ for all $i \geq N-1$. If $V_N \subsetneq V_1$, then $V_N = T^*(s, \xi_1) = \partial_l V_1$ by definition and Lemma 5.3. We use Lemma 5.3 again to see that $\{V_i\}_{i \geq N}$ is a monotone increasing sequence of closed intervals. In any case, we have $T^*(s, \xi_1) \in V_i$ for all i.

The case when $\mu < \lambda_1$ can be proved using a similar argument. Thus the lemma is proved for μ comparable to λ_1.

Now suppose μ is not comparable to λ_1. We again consider the shortest path from λ_1 to μ in Λ^* similar to above and track how V_i changes along the path. To illustrate the idea, we consider the following example. Suppose the shortest path from λ to μ in Λ^* is of length five:

$$\lambda_1 \to \lambda_2 \to \cdots \to \lambda_5 = \mu,$$

where $\lambda_1, \lambda_3, \lambda_5$ are type-1 leaves, λ_2, λ_4 are type-0 leaves, and they satisfy $\lambda_1 \lesssim \lambda_2 \lesssim \lambda_3$ and $\lambda_5 \lesssim \lambda_4 \lesssim \lambda_3$. Let $V_i = V_{s(\lambda_i)}$. We made a turn at λ_3 from the positive flow direction to the negative direction. Our previous discussion shows that $T^*(s, \xi_1) \in V_i$ for $i = 1, 2, 3$. Since λ_2 and λ_4 are incomparable, the core of I_{λ_2} is contained in a single unstable gap G of I_{λ_4}. The interval V_4 must be the gap G containing V_3 by Lemmas 4.14 and 4.15 and the definition of T. In particular, we have $T^*(s, \xi_1) \in V_4$. Since λ_4 is negatively adjacent to λ_3 and $V_3 \subsetneq V_4$, a similar argument to the proof of Lemma 5.3 shows that V_4 is an unstable gap, and an unstable gap will only become larger as we track V_i backward. Therefore, we have $T^*(s, \xi_1) \in V_4 \subset V_5$.

In general, the path from λ to μ has a finite number of turns. If we track the interval V_i along the path, at a turn from the positive direction to the negative direction, V_i will become a larger unstable gap and can only grow even larger until the next turn happens. Similarly, if we turn from the negative direction to the positive direction, V_i will become a larger stable gap and can only grow even larger until the next turn happens. Hence, V_i will be non-decreasing after we make the first turn. But we have shown that $T^*(s, \xi_1)$ is in V_i before we made any turn in the first part of the proof. So the proof of Lemma 5.2 is completed. ∎

Corollary 5.4. *Let $(s,\xi_1) \in \mathrm{LS}^*$ be a pointed leftmost section. Then we have $T^*(s,\xi_1) = \bigcap_{\lambda\in\Lambda} V_{s(\lambda)}$.*

Proof. Lemma 5.2 already shows that $T^*(s,\xi_1) \in \bigcap_{\lambda\in\widetilde{\mathcal{F}}} V_{s(\lambda)}$, so it suffices to prove that there is some λ with $V_{s(\lambda)}$ being a singleton. Suppose ξ_1 is at the infinity of the leaf λ_1. If V_{ξ_1} is a single point, it is trivial. If V_{ξ_1} is a non-trivial closed interval, let e be the side of $p(\lambda_1)$ facing V_{ξ_1}, and let x_e be the periodic point in e as in Proposition 3.1. The periodic orbit $p^{-1}(x_e)$ intersects some leaf λ comparable to λ_1. If we take a path in Λ^* from λ_1 to λ and record the intervals V_i as in the proof of Lemma 5.2, there must be some i so that $V_{i+1} \subsetneq V_i$. Otherwise, we have $V_1 \subset V_2 \subset \cdots$ and so I_λ has a gap containing V_{ξ_1}, but that contradicts $x_e \in p(\lambda)$. By Lemma 5.3, V_{i+1} is a single point, so is the intersection $\bigcap_{\lambda\in\widetilde{\mathcal{F}}} V_{s(\lambda)}$. The lemma is proved. ∎

Corollary 5.5. *There is a map $T : \mathrm{LS} \to \partial\mathcal{O}$ so that the following diagram commutes:*

$$\begin{array}{ccc} \mathrm{LS}^* & \xrightarrow{\pi} & \mathrm{LS} \\ & {\scriptstyle T^*}\searrow & \downarrow {\scriptstyle T} \\ & & \partial\mathcal{O} \end{array}$$

Proof. For any $s \in \mathrm{LS}$, pick a starting point ξ for s and define $T(s) = T^*(s,\xi)$. By Corollary 5.4, we have $T^*(s,\xi) = T^*(s,\xi')$ for any $(s,\xi), (s,\xi') \in \mathrm{LS}^*$. Therefore, the map T is well defined. ∎

Corollary 5.6. *The map T preserves the cyclic order of elements in* LS. *In particular, T is injective.*

Proof. The cyclic order of leftmost sections is determined by the cyclic order of their values on embedded lines in the leaf space ([7, Lemma 6.25]; see also Section 4). Their images under T must follow the same cyclic order by Lemma 5.2. ∎

Lemma 5.7. *The image of* LS *under T is dense.*

Proof. It is clear that T is $\pi_1(M)$-equivariant, so $T(\mathrm{LS})$ is a $\pi_1(M)$-invariant subset of $\partial\mathcal{O}$. The lemma follows from the minimality of the $\pi_1(M)$-action on $\partial\mathcal{O}$ [20]. ∎

Lemma 5.8. *The map $T : \mathrm{LS} \to \partial\mathcal{O}$ extends continuously to a homeomorphism $T : \mathfrak{S}_{\mathrm{left}} \to \partial\mathcal{O}$.*

Proof. By Corollary 5.6, Lemma 5.7 and the fact that $\mathfrak{S}_{\mathrm{left}}$ is the completion of $\mathcal{Q}$, there is a unique continuous extension of T to $\mathfrak{S}_{\mathrm{left}}$, and the extended T is a homeomorphism between $\mathfrak{S}_{\mathrm{left}}$ and $\partial\mathcal{O}$. ∎

Proof of Theorem 5.1. It suffices to show the "moreover" part about the map T defined above. The $\pi_1(M)$-equivariance is automatic from the way we define T. The structure maps are intertwined by T because of Lemma 5.2. ∎

It is also possible to define the universal circle from rightmost sections $\mathfrak{S}_{\text{right}}$ by considering the completion of rightmost sections (i.e., the sections of E_∞ that go rightmost up and leftmost down). In general, there is no reason to expect $\mathfrak{S}_{\text{left}} = \mathfrak{S}_{\text{right}}$. However, we have the following corollary of Theorem 1.2.

Corollary 5.9. *Under the assumption of Theorem 1.2, the universal circles $\mathfrak{S}_{\text{left}}$ and $\mathfrak{S}_{\text{right}}$ are isomorphic.*

Proof. Using the same proof of Theorem 1.2, it can be shown that $\mathfrak{S}_{\text{right}}$ is isomorphic to $\partial\mathcal{O}$, hence isomorphic to $\mathfrak{S}_{\text{left}}$. ∎

Corollary 5.9 suggests that one can view the sets of leftmost and rightmost sections as different dense subsets of the same circle. Indeed, the homeomorphisms from $\mathfrak{S}_{\text{left}}$ and $\mathfrak{S}_{\text{right}}$ to $\partial\mathcal{O}$ give embeddings of the sets of leftmost sections and rightmost sections into the $\partial\mathcal{O}$, both with dense image. Corollary 5.4 is true for both embeddings, so leftmost sections and rightmost sections never cross.

6. Invariant laminations

We conclude with a discussion of the invariant laminations on $\partial\mathcal{O} \cong \mathfrak{S}_{\text{left}}$. See Section 2 for a discussion about laminations on a circle.

Any $\lambda \in \Lambda$ separates Λ into two components, the one $\Lambda^+(\lambda)$ containing the flow positive side of λ and the one $\Lambda^-(\lambda)$ containing the flow negative side of λ. The leaf λ also separates $\widetilde{M}$ into two parts, denoted by $\widetilde{M}^+(\lambda)$ and $\widetilde{M}^-(\lambda)$ with the same sign convention. Define a subset $\Xi^\pm(\lambda)$ of $\text{Symm}_2(\partial\mathcal{O})$ by

$$\Xi^\pm(\lambda) = \partial\text{CH}\Big(\bigcup_{\mu\in\Lambda^\pm(\lambda)} \text{core}(I_\mu)\Big).$$

The set $\Xi^\pm$ is then defined as

$$\Xi^\pm = \overline{\bigcup_{\lambda\in\Lambda} \Xi^\pm(\lambda)}.$$

It is proved in [5] that $\Xi^\pm$ is indeed a pair of $\pi_1(M)$-invariant laminations on $\partial\mathcal{O}$.

We are now ready to prove Theorem 1.6. Recall that $\mathcal{L}^{s/u}_{\mathcal{O}}$ is the lamination on $\partial\mathcal{O}$ induced by $\mathcal{F}^{s/u}_{\mathcal{O}}$.

Proof of Theorem 1.6. Take any type-0 leaf λ and consider the shadow $p(\lambda)$. Since λ is not a fiber, [19, Proposition 4.6] shows that $\partial p(\lambda)$ has leaves in both $\mathcal{F}^u_{\mathcal{O}}$ and $\mathcal{F}^s_{\mathcal{O}}$. Suppose that e is a side of $p(\lambda)$ that is contained in a leaf of $\mathcal{F}^s_{\mathcal{O}}$ and consider the leftmost section s starting from $Q_\lambda(e)$. If μ is a leaf in $\Lambda^+(\lambda)$, we take a path in Λ^* from λ to μ and track the closed interval V_i as in the proof of Lemma 5.2. The interval V_0 is the stable gap of I_λ facing e, and the proof of Lemma 5.2 shows that V_i is monotone increasing along the path. This means the shadow of μ is on the same side of e as $p(\lambda)$. The side e, viewed as an element of $\mathrm{Symm}_2(\partial\mathcal{O})$, is therefore in $\Xi^+(\lambda)$, hence in Ξ^+. By transitivity of ϕ, every leaf of $\mathcal{F}^s$ is dense in M. This implies that the $\pi_1(M)$-image of ∂e is dense in $\mathcal{L}^s_{\mathcal{O}}$ because ϕ has no perfect fits. Since both Ξ^+ and $\mathcal{F}^s_{\mathcal{O}}$ are $\pi_1(M)$-invariant and closed, we have $\Xi^+ \supset \mathcal{L}^s_{\mathcal{O}}$.

If $\Xi^+ - \mathcal{L}^s_{\mathcal{O}}$ is not empty, the difference must be a union of diagonals of complementary regions of $\mathcal{L}^s_{\mathcal{O}}$. Note that these diagonals cannot be approximated by leaves in $\mathcal{L}^s_{\mathcal{O}} \cup \mathcal{L}^u_{\mathcal{O}}$, so there must be such a diagonal d in $\Xi^+(\lambda)$ for some λ. The corresponding complementary polygon comes from a singular leaf l of $\mathcal{F}^s_{\mathcal{O}}$, and we denote the singularity in l by s.

Suppose d has endpoints ξ and ξ'. Then there is a sequence of leaves $\lambda_n \in \Lambda^+(\lambda)$, sides e_n of $p(\lambda_n)$ and endpoints ξ_n of e_n so that ξ_n converges to ξ. Up to taking a subsequence, we can assume that all e_n are contained in leaves of $\mathcal{F}^s_{\mathcal{O}}$ or in leaves of $\mathcal{F}^u_{\mathcal{O}}$. If all e_n are contained in leaves of $\mathcal{F}^u_{\mathcal{O}}$, then d cannot be a boundary component of the convex hull of $\bigcup_{\mu\in\Lambda^\pm(\lambda)} \mathrm{core}(I_\mu)$, because e_n will eventually cross d by the absence of perfect fits. So all e_n are contained in leaves of $\mathcal{F}^s_{\mathcal{O}}$. In particular, the singularity s can be approximated by points in shadows of leaves in $\Lambda^+(\lambda)$. We will show that this is impossible, a contradiction.

If s is in $p(\lambda)$, then there are points of $\mathrm{core}(I_\lambda)$ on both sides of d, contradicting the assumption that $d \in \Xi^+(\lambda)$. Therefore, the $\widetilde{\phi}$-orbit $p^{-1}(s)$ is disjoint from λ. If $p^{-1}(s)$ is contained in $\widetilde{M}^+(\lambda)$, it contradicts our assumption that $d \in \Xi^+(\lambda)$ by a similar reason. So $p^{-1}(s)$ is contained in $\widetilde{M}^-(\lambda)$. Note that s is not in $\partial p(\lambda)$, otherwise a face of l will be a side of $p(\lambda)$, and the face will be a leaf of $\Xi^+(\lambda)$ as we showed above. This again contradicts the assumption that $d \in \Xi^+(\lambda)$. So orbits close enough to $p^{-1}(s)$ will stay in $\widetilde{M}^-(\lambda)$, giving the desired contradiction.

Therefore, Ξ^+ must be the same as $\mathcal{L}^s_{\mathcal{O}}$. For the same reason, we have $\Xi^- = \mathcal{F}^u_{\mathcal{O}}$, finishing the proof of Theorem 1.6. ■

Acknowledgments. The author is grateful to his advisor, Yair Minsky, for being inspiring and supportive throughout this project. The author would like to thank Hyungryul Baik, Ellis Buckminster, Danny Calegari, Sergio Fenley, Michael Landry, Anna Parlak and Sam Taylor for helpful comments and conversations. The author thanks the referee for carefully reading the paper and for giving numerous insightful comments and suggestions.

Funding. The author is partially supported by NSF grant DMS-2005328.

References

[1] I. Agol and C. C. Tsang, Dynamics of veering triangulations: infinitesimal components of their flow graphs and applications. *Algebr. Geom. Topol.* **24** (2024), no. 6, 3401–3453 Zbl 1552.57031 MR 4812222

[2] T. Barbot, Caractérisation des flots d'Anosov en dimension 3 par leurs feuilletages faibles. *Ergodic Theory Dynam. Systems* **15** (1995), no. 2, 247–270 Zbl 0826.58025 MR 1332403

[3] T. Barthelmé, S. Frankel, and K. Mann, Orbit equivalences of pseudo-Anosov flows. *Invent. Math.* **240** (2025), no. 3, 1119–1192 Zbl 08038253 MR 4902161

[4] T. Barthelmé and K. Mann, Pseudo-Anosov flows: a plane approach. [v1] 2025, [v2] 2026, arXiv:2509.15375v2

[5] D. Calegari, Promoting essential laminations. *Invent. Math.* **166** (2006), no. 3, 583–643 Zbl 1106.57014 MR 2257392

[6] D. Calegari, *Foliations and the geometry of 3-manifolds*. Oxford Math. Monogr., Oxford University Press, Oxford, 2007, 363 pp. Zbl 1118.57002 MR 2327361

[7] D. Calegari and N. M. Dunfield, Laminations and groups of homeomorphisms of the circle. *Invent. Math.* **152** (2003), no. 1, 149–204 Zbl 1025.57018 MR 1965363

[8] A. Candel, Uniformization of surface laminations. *Ann. Sci. Éc. Norm. Super. (4)* **26** (1993), no. 4, 489–516 Zbl 0785.57009 MR 1235439

[9] J. Cantwell and L. Conlon, Poincaré–Bendixson theory for leaves of codimension one. *Trans. Amer. Math. Soc.* **265** (1981), no. 1, 181–209 Zbl 0484.57015 MR 0607116

[10] J. Cantwell and L. Conlon, Hyperbolic geometry and homotopic homeomorphisms of surfaces. *Geom. Dedicata* **177** (2015), no. 1, 27–42 Zbl 1359.37097 MR 3370020

[11] J. Cantwell, L. Conlon, and S. R. Fenley, Endperiodic automorphisms of surfaces and foliations. *Ergodic Theory Dynam. Systems* **41** (2021), no. 1, 66–212 Zbl 1464.57022 MR 4190052

[12] D. Cooper, D. D. Long, and A. W. Reid, Bundles and finite foliations. *Invent. Math.* **118** (1994), no. 1, 255–283 Zbl 0858.57015 MR 1292113

[13] S. Fenley, Ideal boundaries of pseudo-Anosov flows and uniform convergence groups with connections and applications to large scale geometry. *Geom. Topol.* **16** (2012), no. 1, 1–110 Zbl 1279.37026 MR 2872578

[14] S. Fenley and L. Mosher, Quasigeodesic flows in hyperbolic 3-manifolds. *Topology* **40** (2001), no. 3, 503–537 Zbl 0990.53040 MR 1838993

[15] S. R. Fenley, Asymptotic properties of depth one foliations in hyperbolic 3-manifolds. *J. Differential Geom.* **36** (1992), no. 2, 269–313 Zbl 0766.53018 MR 1180384

[16] S. R. Fenley, Anosov flows in 3-manifolds. *Ann. of Math. (2)* **139** (1994), no. 1, 79–115 Zbl 0796.58039 MR 1259365

[17] S. R. Fenley, The structure of branching in Anosov flows of 3-manifolds. *Comment. Math. Helv.* **73** (1998), no. 2, 259–297 Zbl 0999.37008 MR 1611703

[18] S. R. Fenley, Foliations with good geometry. *J. Amer. Math. Soc.* **12** (1999), no. 3, 619–676 Zbl 0930.53024 MR 1674739

[19] S. R. Fenley, Surfaces transverse to pseudo-Anosov flows and virtual fibers in 3-manifolds. *Topology* **38** (1999), no. 4, 823–859 Zbl 0926.57009 MR 1679801

[20] S. R. Fenley, Geometry of foliations and flows I: almost transverse pseudo-Anosov flows and asymptotic behavior of foliations. *J. Differential Geom.* **81** (2009), no. 1, 1–89 Zbl 1160.57026 MR 2477891

[21] E. Field, H. Kim, C. Leininger, and M. Loving, End-periodic homeomorphisms and volumes of mapping tori. *J. Topol.* **16** (2023), no. 1, 57–105 Zbl 1567.57023 MR 4532490

[22] D. Fried, Fibrations over S^1 with pseudo-Anosov monodromy. In *Travaux de Thurston sur les surfaces*, pp. 251–266, Astérisque 66–67, Société Mathématique de France, Paris, 1979 Zbl 0446.57023 MR 0568308

[23] M. P. Landry, Y. N. Minsky, and S. J. Taylor, Endperiodic maps via pseudo-Anosov flows. 2023, arXiv:2304.10620v1, to appear in *Geom. Topol.*

[24] M. P. Landry, Y. N. Minsky, and S. J. Taylor, Simultaneous universal circles. *J. Topol.* **19** (2026), no. 1, article no. e70054, 28 pp. Zbl 08143171 MR 5009301

[25] M. P. Landry and C. C. Tsang, Endperiodic maps, splitting sequences, and branched surfaces. *Geom. Topol.* **29** (2025), no. 9, 4531–4663 Zbl 08152702 MR 5017754

[26] M. P. Landry and C. C. Tsang, Pseudo-Anosov flows from sutured hierarchies. In preparation

[27] L. Mosher, Dynamical systems and the homology norm of a 3-manifold, I: efficient intersection of surfaces and flows. *Duke Math. J.* **65** (1992), no. 3, 449–500 Zbl 0754.58030 MR 1154179

[28] L. Mosher, Laminations and flows transverse to finite depth foliations. Preprint, 1996

[29] S. P. Novikov, The topology of foliations (in Russian). *Tr. Moskov. Mat. Obs.* **14** (1965), 248–278. English translation. *Trans. Moscow Math. Soc.* **14** (1965), 268–304 Zbl 0247.57006 MR 0200938

[30] A. Parlak, Mutations and faces of the Thurston norm ball dynamically represented by multiple distinct flows. *Geom. Topol.* **29** (2025), no. 4, 2105–2173 Zbl 08084075 MR 4929474

[31] W. P. Thurston, A norm for the homology of 3-manifolds. *Mem. Amer. Math. Soc.* **59** (1986), no. 339, 99–130 Zbl 0585.57006 MR 0823443

[32] W. P. Thurston, Three-manifolds, foliations and circles, II. Unfinished manuscript, 1998

Received 19 December 2024.

Junzhi Huang
Department of Mathematics, Yale University, 219 Prospect St, Floors 7-9, New Haven, CT 06511, USA; junzhi.huang@yale.edu
Author IDs: zbMATH huang.junzhi MR 1498159 ORCID 0000-0001-8451-4997