

COARSE HYPERBOLICITY AND CLOSED ORBITS FOR QUASIGEODESIC FLOWS

STEVEN FRANKEL

ABSTRACT. We prove a conjecture of Calegari's, that every quasigeodesic flow on a closed hyperbolic 3-manifold contains a closed orbit.

1. INTRODUCTION

In 1950, Seifert conjectured that every nonsingular flow on the 3-sphere must contain a closed orbit [27]. The first counterexamples appeared in 1974, when Schweitzer showed that every nonsingular flow on a 3-manifold can be deformed to a C^1 flow with no closed orbits [26]. These examples have been generalized considerably, and it is known that the flow can be taken to be smooth [19] or volume-preserving [18].

On the other hand, there are certain geometric constraints that ensure the existence of closed orbits. Taubes' 2007 proof of the 3-dimensional Weinstein conjecture shows that every Reeb flow on a closed 3-manifold has a closed orbit [28]. Reeb flows are *geodesible*, i.e. there is a Riemannian metric in which the flowlines are geodesics. In 2010, Rechtman showed that any real analytic geodesible flow on a closed 3-manifold has a closed orbit, unless the manifold is a torus bundle with reducible monodromy [25].

Geodesibility is a global geometric condition. In contrast, a flow is *quasigeodesic* if the path taken by each point is coarsely comparable to a geodesic, a local condition. In this paper, we show that every quasigeodesic flow on a closed hyperbolic 3-manifold contains a closed orbit, as conjectured by Calegari.

Our argument works by studying the transverse behavior of a quasigeodesic flow, which displays a kind of "coarse hyperbolicity" whenever the ambient manifold is hyperbolic. In fact, our argument can be applied to a larger class of flows with coarsely hyperbolic transverse structures.

In addition, one can think of this coarse transverse hyperbolicity as a topological analogue of the transverse hyperbolicity exhibited by a pseudo-Anosov flow. This analogy is taken further in a proof, announced by the author, of Calegari's Flow Conjecture, which asserts that any quasigeodesic flow on a closed hyperbolic 3-manifold can be deformed to a flow that is both quasigeodesic and pseudo-Anosov.

1.1. Quasigeodesic flows. A *flow* on a manifold M is a continuous map

$$\Phi_{(\cdot)}(\cdot) : \mathbb{R} \times M \rightarrow M$$

such that

- $\Phi_0(x) = x$ for all $x \in M$, and
- $\Phi_s(\Phi_t(x)) = \Phi_{t+s}(x)$ for all $x \in M$ and $t, s \in \mathbb{R}$.

Date: November 13, 2017.

For each $t \in \mathbb{R}$, the *time- t map* $\Phi_t(\cdot) : M \rightarrow M$ is a homeomorphism, so we can think of a flow as a continuous action $M \curvearrowright \mathbb{R}$, writing $x \cdot t = \Phi_t(x)$. A flow is *nonsingular* if it has no global fixed points, i.e. for each $x \in M$ there is a time $t \in \mathbb{R}$ such that $\Phi_t(x) \neq x$.

A flow Φ on a manifold M is *quasigeodesic* if each orbit lifts to a quasigeodesic in the universal cover \widetilde{M} . That is, each lifted orbit admits a parametrization $\gamma : \mathbb{R} \rightarrow \widetilde{M}$ satisfying

$$\frac{1}{k} \cdot d(\gamma(s), \gamma(t)) - \epsilon \leq |s - t| \leq k \cdot d(\gamma(x), \gamma(y)) + \epsilon$$

for constants $k > 0, \epsilon \geq 0$ that may depend on the orbit. It is *uniformly quasigeodesic* if the constants k, ϵ can be chosen independent of the orbit.

If M is a closed hyperbolic 3-manifold, then the universal cover \widetilde{M} is identified with \mathbb{H}^3 , so it has a natural compactification to a closed 3-ball $\widetilde{M} \sqcup S_\infty^2$. Here, the *sphere at infinity* S_∞^2 is identified with the boundary of hyperbolic space in the unit ball model. The action of $\pi_1(M)$ on the universal cover is isometric, so it extends to the sphere at infinity.

A flow Φ on M lifts to a flow $\widetilde{\Phi}$ on \widetilde{M} , but the orbits of the lifted flow need not behave well with respect to the sphere at infinity. In particular, they may remain in bounded subsets of \widetilde{M} , or accumulate on arbitrary closed subsets of S_∞^2 . When Φ is quasigeodesic, however, the following so-called *Morse Lemma* implies that each lifted orbit has well-defined and distinct endpoints in S_∞^2 . See [14], [16, Corollary 3.44], or [3, §III.H].

Morse Lemma. *Every quasigeodesic in \mathbb{H}^3 lies at a bounded distance from a unique geodesic. Furthermore, there are constants $C(k, \epsilon)$ such that every (k, ϵ) -quasigeodesic in \mathbb{H}^3 lies in the $C(k, \epsilon)$ -neighborhood of its associated geodesic.*

In addition, the endpoints of lifted orbits vary continuously, and this behavior characterizes the quasigeodesic flows on a closed hyperbolic 3-manifold. They are exactly the flows that can be studied “from infinity” in the following sense.

Proposition 1.1 ([9, Theorem B] & [4, Lemma 4.3]). *Let Φ be a flow on a closed hyperbolic 3-manifold M , and let $\widetilde{\Phi}$ be the lifted flow on the universal cover \widetilde{M} . Then Φ is quasigeodesic if and only if*

- (1) *each orbit of $\widetilde{\Phi}$ has well-defined and distinct endpoints in S_∞^2 , and*
- (2) *the positive and negative endpoints of $x \cdot \mathbb{R}$ vary continuously with $x \in \widetilde{M}$.*

The simplest examples of quasigeodesic flows come from fibrations.

Example 1.2. Zeghib showed that any flow on a closed 3-manifold M (not necessarily hyperbolic) that is transverse to a fibration is quasigeodesic [31]. The idea is to lift such a flow to the infinite cyclic cover dual to a fiber $\Sigma \subset M$, which may be identified with $\Sigma \times \mathbb{R}$ in such a way that the lifts of Σ are of the form $\Sigma \times \{i\}$ for $i \in \mathbb{Z}$. Quasigeodesicity follows from the observation that there are upper and lower bounds on the distance between adjacent lifts, as well as the time it takes for the flow to move points from one lift to the next.

On the other hand, there are many quasigeodesic flows that are not transverse to fibrations, even virtually (i.e. after passing to a finite cover).

Example 1.3. Gabai showed that any nontrivial second cohomology class on a closed 3-manifold represents the depth-zero leaf of a taut, finite-depth foliation [13]. Fenley and Mosher showed that there is a quasigeodesic flow that is transverse or “almost transverse” to such a foliation [9]. If one takes a cohomology class that is not virtually represented by a fiber, then the associated quasigeodesic flow is not virtually transverse to a fibration.

1.2. Transverse hyperbolicity. For motivation, we recall the Anosov Closing Lemma, which produces closed orbits for Anosov and pseudo-Anosov flows.

A smooth flow Φ on a 3-manifold M is *Anosov* if it preserves a splitting of the tangent bundle TM into three one-dimensional sub-bundles

$$TM = E^s \oplus T\Phi \oplus E^u$$

where $T\Phi$ is tangent to the flow, and the flow exponentially contracts the *stable bundle* E^s and exponentially expands the *unstable bundle* E^u . Here and throughout, we use the convention that “expansion” means contraction in backwards time. This gives rise to a transverse pair of two-dimensional foliations, the *weak stable* and *weak unstable foliations*, obtained by integrating the 2-dimensional sub-bundles $T\Phi \oplus E^s$ and $T\Phi \oplus E^u$.

A flow on a 3-manifold is *pseudo-Anosov* if it is Anosov everywhere except near a collection of isolated closed orbits, and the weak stable and unstable foliations on the complement of these orbits extend to *singular* foliations on the entire manifold. Figure 1 illustrates the local picture near a singularity of order 4.




FIGURE 1. Singularities of stable and unstable foliations.

Let $\phi : \Sigma \rightarrow \Sigma$ be a homeomorphism of a closed surface, and consider the surface bundle M_ϕ with monodromy ϕ , i.e. the closed 3-manifold

$$M_\phi = \Sigma \times [0, 1] / (x, 1) \sim (\phi(x), 0).$$

The semi-flow $(x, y) \cdot t = (x, y + t)$ on $M \times [0, 1]$ glues up to produce a flow Φ on M_ϕ called the *suspension flow* associated to ϕ .

The suspension flow associated to an arbitrary homeomorphism is quasigeodesic since it is transverse to the fibration of M_ϕ by the images of the surfaces $\Sigma \times \{y\}$. In addition, the suspension flow associated to a pseudo-Anosov homeomorphism

is a pseudo-Anosov flow. Its 2-dimensional (singular) weak stable and unstable foliations can be seen explicitly by flowing the 1-dimensional (singular) stable and unstable foliations of the associated homeomorphism, thought of as living on a fiber.

An *almost-cycle* in a flow is a closed loop obtained by concatenating a long flow segment with a short arc. More concretely, an (ϵ, T) -cycle consists of flow segment of the form $x \cdot [0, t]$ for $t \geq T$, concatenated with arc from x to $x \cdot t$ with length at most ϵ . The Anosov Closing Lemma leverages the contracting/expanding, or “transversely hyperbolic” behavior of an Anosov flow to find closed orbits near almost-cycles.

Anosov Closing Lemma. *Let Φ be an Anosov flow on a closed 3-manifold M . Then for each $D > 0$ there are constants $T, \epsilon > 0$ such that any (ϵ, T) -cycle contains a closed orbit in its D -neighborhood.*

An analogous result holds for pseudo-Anosov flows, but one must be careful with almost-cycles that lie near the singular orbits [22].

The idea behind the Anosov Closing Lemma is illustrated in Figure 2. The left side of the figure depicts the local structure near the ends of a long flow segment $x \cdot [0, t]$, while the right side depicts the local structure near a point $x' = x \cdot \frac{t}{2}$ in the middle. Since x is close to $x \cdot t$, the local stable/unstable leaf through x intersects the local unstable/stable leaf through $x \cdot t$.




FIGURE 2. The Anosov closing lemma.

Take a point y where the stable leaf through x intersects the unstable leaf through $x \cdot t$. Flowing forward, we arrive at a point $y \cdot t'_+$ which lies very close to x' along its stable leaf; flowing backwards, we arrive at a point $y \cdot t'_-$ which lies very close to x' along its unstable leaf. This produces a flow segment $y \cdot [t'_-, t'_+]$ whose length is comparable to $x \cdot [0, t]$, but whose ends are much closer together, which we can close up to obtain an almost-cycle. Repeating, one obtains a sequence of better and better almost-cycles, which limit to a closed orbit.

1.3. Coarse transverse hyperbolicity. At first glance, the tangent condition that defines a quasigeodesic flow seems unrelated to the hyperbolic transverse structure that defines a pseudo-Anosov flow. In the presence of ambient hyperbolicity,

however, we will see that a quasigeodesic flow admits an analogous transverse structure that is *coarsely* hyperbolic.

In a pseudo-Anosov flow, a weak stable or unstable leaf is a maximal connected set of orbits that are forwards or backwards asymptotic to each other. In a quasigeodesic flow, one defines a *weak positive* or *weak negative leaf* to be a maximal connected set of orbits that share their positive or negative endpoints in the universal cover. Any pair of orbits in a weak positive or negative leaf are *coarsely* forwards or backwards asymptotic, in the sense that they contain forwards or backwards rays that lie a uniformly bounded distance apart in the universal cover.

This “coarse transverse hyperbolicity” is a far cry from the strict form of hyperbolicity needed for the Anosov closing lemma. Besides the coarseness of the contraction/expansion, there is little control over the topology of the weak leaves. They may have nontrivial interior, and may not be path-connected or even locally connected. In particular, there is no natural notion of transversality between weak positive and negative leaves. Nevertheless, we will prove a “Homotopy Closing Lemma,” which finds closed orbits in the free homotopy classes of certain almost-cycles. This can be seen as a coarse analogue of the Anosov Closing Lemma, which finds closed orbits geometrically near certain almost-cycles.

1.4. Outline and results. Fix a quasigeodesic flow Φ on a closed hyperbolic 3-manifold M , which lifts to a flow $\tilde{\Phi}$ on the universal cover $\tilde{M} \simeq \mathbb{H}^3$. To avoid confusion, the orbits of $\tilde{\Phi}$ will be called *flowlines*. In \tilde{M} , the collections of all weak positive and weak negative leaves form a pair of $\pi_1(M)$ -equivariant decompositions, \mathcal{D}_w^+ and \mathcal{D}_w^- .

Because of the topological and geometric difficulties in working with weak leaves, our argument will take a very different form than that of the Anosov Closing Lemma. Instead of working directly in the manifold, we will reduce the 3-dimensional problem of finding closed orbits to a 2-dimensional problem in the *flowspace*. This is a topological plane P , the orbit space of the lifted flow, which comes equipped with a natural action of $\pi_1(M)$. See §3.

Each point in P corresponds to a flowline in \tilde{M} , which projects to an orbit in M , and the periodicity or recurrence of this orbit can be seen in terms of the flowspace action. In particular, a point in P corresponds to a closed orbit in M if and only if it is fixed by some nontrivial element of $\pi_1(M)$. Thus, to show the existence of a closed orbit it suffices to show that the flowspace action is not free.

The flowspace is also useful for understanding the topology of the weak leaves. The decompositions \mathcal{D}_w^\pm of \tilde{M} project to decompositions \mathcal{D}^\pm of P , whose elements are simply called *positive leaves* and *negative leaves*. These *positive and negative decompositions* are monotone, unbounded, and intersect compactly. That is, each positive or negative leaf is a closed, connected, unbounded planar set, and the intersection between a positive leaf and a negative leaf is compact. Note that the unboundedness of the leaves represents a “transverse unboundedness” of the weak leaves, which would be difficult to express without using the flowspace.

We will use these properties to treat the positive and negative decompositions as a generalization of a pair of singular foliations. This analogy is realized by ignoring the local topology of leaves and focusing on their separation properties, which can be understood using Calegari’s *universal circle* [4], a topological circle S_u^1 , equipped with a faithful action of $\pi_1(M)$. Although the original construction of this circle is

abstract, coming from a cyclic ordering of the Freudenthal ends of leaves, we showed in [10] that it can be thought of as the boundary of the flowspace. That is, the disjoint union $\mathbf{P} = P \sqcup S_u^1$ has a natural topology with which it is homeomorphic to a closed disc, and the flowspace and universal circle actions combine to form an action $\pi_1(M) \curvearrowright \mathbf{P}$.

In §4, we will see that decompositions \mathcal{D}^\pm of P extend to decompositions \mathcal{D}^\pm of \mathbf{P} whose elements are called *positive sprigs* and *negative sprigs*. Each sprig has a natural set of *ends*, the points at which it intersects the universal circle, and the separation properties of a sprig are reflected in those of its ends. Although sprigs contain leaves, and hence display the same kinds of topological pathologies, they are still easier to deal with since they are compact and have nice convergence properties.

In §5, we study the relationship between the positive and negative sprig decompositions, and show that they form a *spidery pair* (cf. Definitions 4.11 and 5.7), whose topological properties are studied in §6–7. Although there is no natural notion of a transverse intersection point between a positive and negative sprig, there is a generalization of this idea called a *linked point*. A point $p \in P$ is said to be linked if the ends of its positive and negative sprigs contain 0-spheres that are linked in the 1-sphere S_u^1 . We show that the set $P_l \subset P$ of linked points has the following properties.

Recurrent Links Lemma. *The linked region P_l is closed, nontrivial, $\pi_1(M)$ -invariant, and contains an ω -recurrent point.*

Here, an ω -recurrent point is one that corresponds to an ω -recurrent orbit in M . If $p \in P$ is ω -recurrent, then a sequence of almost-cycles (γ_i) that approximate the corresponding forward orbit are represented by a sequence of elements (g_i) in the fundamental group, called an ω -sequence, with the property that $\lim g_i(p) = p$.

In §8, we show that the coarse contraction of positive sprigs is reflected in the action of an ω -sequence on the flowspace. In particular, if $p \in P$ is an ω -recurrent point, then a corresponding ω -sequence (g_i) takes any point $p' \in P$ in the same positive sprig towards an *a priori* determined compact region. We use this in §9 to prove our closing lemma.

Homotopy Closing Lemma. *Let $p \in P_l$ be an ω -recurrent point, and let $(g_i)_{i=1}^\infty$ be a corresponding ω -sequence. Then for all i sufficiently large, g_i fixes a point in P .*

Since a point in P that is fixed by a nontrivial element of $\pi_1(M)$ corresponds to a closed orbit, our main theorem follows immediately.

Closed Orbits Theorem. *Every quasigeodesic flow on a closed hyperbolic 3-manifold contains a closed orbit.*

In addition, we will see in §9.2 that certain closed orbits can be seen purely in terms of the universal circle.

1.5. Acknowledgements. I would like to thank the referee, whose comments and suggestions contributed greatly to the readability of this paper. Thanks to Thierry Barbot, Thomas Barthelme, Danny Calegari, Sergio Fenley, Dave Gabai, Yair Minsky, Lee Mosher, and Juliana Xavier for many productive conversations.

This material is based partially upon work supported by the National Science Foundation under Grant No. DMS-1611768. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the National Science Foundation.

2. TOPOLOGICAL BACKGROUND

In this section we will review some topological background that will be used throughout the sequel. See [30], [20], [21], and [15] for more details.

2.1. Limits. Let $(A_i)_{i=1}^{\infty}$ be a sequence of subsets of a metric space X . The *limit inferior*

$$\underline{\lim} A_i$$

is the set of all $x \in X$ such that each neighborhood of x intersects all but finitely many of the A_i . The *limit superior*

$$\overline{\lim} A_i$$

is the set of all $x \in X$ such that every neighborhood of x intersects infinitely many of the A_i .

In other words, $x \in \underline{\lim} A_i$ if and only if there is a sequence of points $x_i \in A_i$ that converge to x , and $y \in \overline{\lim} A_i$ if and only if there is a sequence of points $y_i \in A_i$ that accumulate on y . The limits inferior and superior are always closed, and

$$\underline{\lim} A_i \subset \overline{\lim} A_i.$$

If the limits inferior and superior agree, then (A_i) is said to be (*Kuratowski*) *convergent*, and we write

$$\lim A_i := \underline{\lim} A_i = \overline{\lim} A_i.$$

When X is a compact metric space, this is equivalent to Hausdorff convergence, and we can avail ourselves of the following useful properties of Hausdorff limits [15, §2].

Theorem 2.1. *Let X be compact metrizable space X . Then:*

- (1) *Every sequence of subsets has a convergent subsequence.*
- (2) *If (A_i) is a sequence of connected subsets, and $\underline{\lim} A_i$ is not empty, then $\overline{\lim} A_i$ is connected.*
- (2') *If (A_i) is a convergent sequence of connected subsets then $\lim A_i$ is connected.*

2.2. Decompositions. Let X be a topological space. A *partition* of X is a collection of pairwise disjoint subsets that cover X . A *decomposition* of X is a partition whose elements are closed. A partition or decomposition is *nontrivial* if it contains more than one element, and *monotone* if its elements are connected.

Let \mathcal{D} be a decomposition of a space X . A subset $A \subset X$ is said to be \mathcal{D} -*saturated* if each decomposition element $K \in \mathcal{D}$ that intersects A is contained in A ; equivalently, if A is a union of decomposition elements. The \mathcal{D} -*saturation* of a subset $B \subset X$ is the smallest \mathcal{D} -saturated set that contains B , which we denote by $\mathcal{D}(B)$; equivalently, $\mathcal{D}(B)$ is the union of all decomposition elements $K \in \mathcal{D}$ that intersect B . When the decomposition \mathcal{D} is implicit, we will speak simply of *saturated* sets and *saturations*.

Given a decomposition \mathcal{D} of a space X , the corresponding *decomposition space* is the identification space X/\mathcal{D} , equipped with the quotient topology. Alternatively, one can think of the decomposition space as the set \mathcal{D} itself, with the topology defined by declaring that $U \subset \mathcal{D}$ is open whenever $\bigcup_{K \in U} K$ is open in X .

Definition 2.2 ([20, §I.19]). A decomposition \mathcal{D} of a space X is *upper semicontinuous* if it satisfies the following equivalent conditions:

- (1) The quotient map $X \rightarrow X/\mathcal{D}$ is closed (i.e. the image of every closed set is closed).
- (2) $U \subset X$ is open \Rightarrow the union of all decomposition elements that are contained in U is open.
- (3) $A \subset X$ is closed \Rightarrow the union of all decomposition elements that intersect A is closed. Equivalently, the saturation of every closed set is closed.

In a compact metric space, the upper semicontinuity of a decomposition can be understood in terms of the convergence properties of its elements.

Theorem 2.3 ([21, Theorem IV.43.2]). *Let \mathcal{D} be a decomposition of a compact metrizable space X . The following are equivalent:*

- (1) \mathcal{D} is upper semicontinuous.
- (2) If $K_1, K_2, \dots \in \mathcal{D}$ is a sequence of decomposition elements, and $\varinjlim K_i$ intersects a decomposition element $K \in \mathcal{D}$, then $\varinjlim K_i \subset K$.
- (3) If $K_1, K_2, \dots \in \mathcal{D}$ is a convergent sequence of decomposition elements, then $\varinjlim K_i$ is contained in some decomposition element $K \in \mathcal{D}$.

Lemma 2.4 ([21, Theorem IV.43.1]). *Let \mathcal{D} be an upper semicontinuous decomposition of a compact Hausdorff space X . Then the decomposition space X/\mathcal{D} is compact Hausdorff.*

The following easy lemmas are useful for constructing upper semicontinuous decompositions.

Lemma 2.5 ([15, Theorem 3-31]). *If $f : X \rightarrow Y$ is a continuous map between compact metrizable spaces, then the decomposition $\{f^{-1}(y) \mid y \in Y\}$ by point-preimages is upper semicontinuous.*

The *monotonization* of a decomposition \mathcal{D} is the decomposition whose elements are the connected components of elements of \mathcal{D} .

Lemma 2.6 ([15, Theorem 3-39]). *In a compact metrizable space, the monotoneization of an upper semicontinuous decomposition is upper semicontinuous.*

If \mathcal{D} is a nontrivial monotone upper semicontinuous decomposition of a closed interval I , then it's easy to see that the associated decomposition space I/\mathcal{D} is homeomorphic to a closed interval. The following theorem generalizes this to dimension 2.

Moore's Theorem ([24]). *Let \mathcal{D} be a monotone upper semicontinuous decomposition of a closed 2-disc \mathbf{D} , and suppose that each decomposition element $K \in \mathcal{D}$ is nonseparating in \mathbf{D} . Then the decomposition space \mathbf{D}/\mathcal{D} is homeomorphic to a closed 2-disc.*

This fails in higher dimensions.

3. LEAVES

Fix, once and for all, a quasigeodesic flow Φ on a closed hyperbolic 3-manifold M . We assume throughout that M is orientable; this does not result in a loss of generality since passing to a double cover does not affect the existence of closed orbits.

We will work mostly with the lifted flow $\tilde{\Phi}$ on the universal cover $\tilde{M} \simeq \mathbb{H}^3$, whose orbits we call *flowlines*. Since M is compact, each nontrivial element $g \in \pi_1(M)$ acts as a loxodromic isometry on \tilde{M} , and has two fixed points in S_∞^2 in an attracting-repelling pair. Since M is closed, the action on S_∞^2 is minimal, in the sense that every orbit is dense. See [29].

The Morse Lemma implies that each flowline has well-defined and distinct endpoints in S_∞^2 . These endpoints vary continuously [4, Lemma 4.3], so we have a pair of $\pi_1(M)$ -equivariant maps

$$E^\pm : \tilde{M} \rightarrow S_\infty^2 \\ x \mapsto \lim_{t \rightarrow \pm\infty} x \cdot t.$$

Since quasigeodesics have distinct endpoints, we have

$$E^+(x) \neq E^-(x) \text{ for all } x \in \tilde{M}.$$

Fix a point $z \in S_\infty^2$. Then $(E^+)^{-1}(z)$ is the union of all flowlines with positive endpoint at z , and each connected component of this set is called a *weak positive leaf rooted at z* . Similarly, $(E^-)^{-1}(z)$ is the union of all flowlines with negative endpoint at z , each component of which is called a *weak negative leaf rooted at z* .

The collections of all weak positive and weak negative leaves form a pair of monotone decompositions of \tilde{M} ,

$$\mathcal{D}_w^+ := \{\text{components of } (E^+)^{-1}(z) \mid z \in S_\infty^2\} \\ \text{and } \mathcal{D}_w^- := \{\text{components of } (E^-)^{-1}(z) \mid z \in S_\infty^2\}.$$

The action of $\pi_1(M)$ preserves these decompositions, so they descend under the covering map $\pi : \tilde{M} \rightarrow M$ to a pair of monotone partitions of M . These are only partitions, not decompositions, since the image of a weak leaf need not be closed.

A quasigeodesic flow on a closed hyperbolic manifold is always uniformly quasigeodesic [4, Lemma 3.10], so there is a uniform constant $C := C(k, \epsilon)$ such that the flowline through each $x \in \tilde{M}$ lies in the C -neighborhood of the geodesic from $E^-(x)$ to $E^+(x)$. We will use this in §8 to see that the flowlines in each weak positive/negative leaf are coarsely asymptotic in the forwards/backwards direction.

3.1. The flowspace. Let P be the orbit space of the lifted flow $\tilde{\Phi}$, i.e. the set of flowlines in \tilde{M} together with quotient topology induced by the map

$$\nu : \tilde{M} \rightarrow P.$$

The deck group preserves the decomposition of \tilde{M} into flowlines, so there is an induced action $\pi_1(M) \curvearrowright P$. The space P , together with this action, is called the *flowspace* of Φ .

Using uniform quasigeodesicity, Calegari showed that P is Hausdorff, and therefore homeomorphic to the plane [4, Theorem 3.12]. The deck transformations preserve the orientation on the flowlines, as well as the orientation on the universal cover, so the flowspace action is orientation-preserving.

Although the flowspace is constructed as a quotient of the universal cover, we can use the following theorem to think of it as a transversal to the lifted flow.

Theorem 3.1 (Montgomery-Zippin [23]). *Let Ψ be a flow on \mathbb{R}^3 whose orbit space P is Hausdorff. Then the quotient map $\nu : \mathbb{R}^3 \rightarrow P$ admits a continuous section $\sigma : P \rightarrow \mathbb{R}^3$.*

Given such a section

$$\sigma : P \rightarrow \widetilde{M},$$

the map

$$\begin{aligned} \Sigma : P \times \mathbb{R} &\rightarrow M \\ (p, h) &\mapsto \sigma(p) \cdot h \end{aligned}$$

is a homeomorphism that conjugates the “vertical flow” on $P \times \mathbb{R}$, defined by $(p, h) \cdot t = (p, h + t)$, to the lifted flow $\widetilde{\Phi}$ on \widetilde{M} .

We will write

$$\langle p \rangle := \nu^{-1}(p)$$

for the flowline corresponding to a point $p \in P$, and

$$\langle A \rangle := \nu^{-1}(A)$$

for the union of flowlines corresponding to a subset $A \subset P$. From the homeomorphism Σ one sees that $\langle A \rangle$ is homeomorphic to $A \times \mathbb{R}$ for any $A \subset P$.

3.2. Leaves. Let

$$e^\pm : P \rightarrow S_\infty^2$$

be the maps that take each point $p \in P$ to the positive and negative endpoints of the corresponding flowline $\langle p \rangle$. These are just the factorizations of E^\pm through the quotient map ν , so they are continuous, $\pi_1(M)$ -equivariant, and satisfy

$$e^+(p) \neq e^-(p) \text{ for all } p \in P.$$

For each $z \in S_\infty^2$ the connected components of $(e^+)^{-1}(z)$ are called *positive leaves rooted at z* , while the connected components of $(e^-)^{-1}(z)$ are called *negative leaves rooted at z* . Equivalently, a positive or negative leaf rooted at z is the image under ν of a weak positive or weak negative leaf rooted at z . See Figure 3.

The collections of all positive and negative leaves form a pair of $\pi_1(M)$ -invariant monotone decompositions of P , the *positive* and *negative decompositions*

$$\begin{aligned} \mathcal{D}^+ &:= \{\text{components of } (e^+)^{-1}(z) \mid z \in S_\infty^2\} \\ \text{and } \mathcal{D}^- &:= \{\text{components of } (e^-)^{-1}(z) \mid z \in S_\infty^2\}. \end{aligned}$$

These have two important properties.

Proposition 3.2 ([4, Lemmas 4.8 & 5.8]).

- (1) *Each leaf $K \in \mathcal{D}^+ \cup \mathcal{D}^-$ is unbounded.*
- (2) *If $K \in \mathcal{D}^+$ is a positive leaf and $L \in \mathcal{D}^-$ is a negative leaf, then $K \cap L$ is compact (and possibly empty).*

We say that \mathcal{D}^+ and \mathcal{D}^- are *unbounded decompositions that intersect compactly*. We will use these properties to treat the positive and negative decompositions as a broad generalization of a pair of foliations. Property (2) can be generalized as follows.




FIGURE 3. The positive and negative leaves through a point $p \in P$.

Lemma 3.3. *Let A and B be disjoint compact subsets of S_∞^2 . Then $(e^-)^{-1}(A) \cap (e^+)^{-1}(B)$ is compact.*

Proof. By uniform quasigeodesicity, $(E^-)^{-1}(A) \cap (E^+)^{-1}(B)$ is contained in the C -neighborhood of the union of all geodesics from A to B . Since A and B are compact and disjoint, we can find a compact set $D \subset \widetilde{M}$ that intersects every one of these flowlines. Then $(e^-)^{-1}(A) \cap (e^+)^{-1}(B)$ compact since it is a closed subset of the compact set $\nu(D)$. \square

Remark 3.4. Let Ψ be a pseudo-Anosov flow on a closed 3-manifold N , not necessarily hyperbolic. Then the flowspace $\pi_1(N) \curvearrowright P$, defined in the same manner, is also a topological plane. The 2-dimensional singular weak stable and unstable foliations lift to the universal cover and project to 1-dimensional singular foliations of the flowspace P which we'll call the *stable and unstable foliations*. See [8].

If N is hyperbolic, and Ψ is quasigeodesic in addition to pseudo-Anosov, then the positive and negative decompositions that come from its quasigeodesic structure are exactly the stable and unstable foliations that come from its pseudo-Anosov structure.

3.3. Dynamics in the flowspace. Each point $p \in P$ corresponds to a flowline $\langle p \rangle$ in \widetilde{M} , and an orbit $\pi(\langle p \rangle)$ in M .

Lemma 3.5. *A point $p \in P$ corresponds to a closed orbit in M if and only if there is a nontrivial element $g \in \pi_1(M)$ such that $g(p) = p$. Any such g represents a multiple of the free homotopy class of the corresponding orbit.*

Proof. Let g be a nontrivial element of $\pi_1(M)$ that fixes p . Then g fixes the corresponding flowline $\langle p \rangle \subset \widetilde{M}$, so the image $\pi(\langle p \rangle)$ is a closed orbit, and g represents a multiple of its free homotopy class.

If $\pi(\langle p \rangle)$ is closed, then it is homotopically nontrivial, since it has a lift $\langle p \rangle$ that is homeomorphic to a line. Take a point $x \in \langle p \rangle$ as the basepoint for \widetilde{M} , and

$x_0 = \pi(x)$ as the basepoint for M . Then an element $g \in \pi_1(M, x_0)$ that represents its homotopy class is nontrivial, and fixes $\langle p \rangle$ and hence p . \square

Thus the Closed Orbits Theorem reduces to showing that the flowspace action is not free.

A point $x \in M$ is called ω -recurrent (α -recurrent) if there is a sequence of times $t_i \rightarrow \infty$ (resp. $t_i \rightarrow -\infty$) such that $\lim_{i \rightarrow \infty} x \cdot t_i = x$. If $x, y \in M$ are in the same orbit, then x is α - or ω -recurrent if and only if y is, so we can speak of *orbits* being α - or ω -recurrent. A *recurrent* point or orbit is one that is either α - or ω -recurrent. A point $p \in P$ is said to be *recurrent*, α -recurrent, or ω -recurrent when this holds for the corresponding orbit in M .

Lemma 3.6. *A point $p \in P$ is recurrent if and only if there is a sequence of nontrivial elements $g_1, g_2, \dots \in \pi_1(M)$ such that $\lim_{i \rightarrow \infty} g_i(p) = p$.*

Proof. Take a point $x \in \langle p \rangle$ as the basepoint for \widetilde{M} , and $x_0 = \pi(x)$ as the basepoint for M .

If $\pi(\langle p \rangle)$ is ω -recurrent, then we can find a sequence of times $t_i \rightarrow \infty$ such that $\lim_{i \rightarrow \infty} x_0 \cdot t_i = x_0$. For each i , let g_i be the element of $\pi_1(M, x_0)$ that represents the almost-cycle obtained by concatenating the flow segment $x_0 \cdot [0, t_i]$ with an arc c_i of shortest possible length. The corresponding lift is the concatenation of $x \cdot [0, t_i]$ with the lift \tilde{c}_i of c_i that starts at $x \cdot t_i$, and g_i takes the terminal endpoint of \tilde{c}_i to x . The length of \tilde{c}_i goes to 0, so we have $\lim_{i \rightarrow \infty} g_i(x \cdot t_i) = x$, and hence $\lim_{i \rightarrow \infty} g_i(p) = p$. A similar argument applies if p is α -recurrent.

On the other hand, suppose that we have a sequence $(g_i)_{i=1}^\infty$ of nontrivial elements of $\pi_1(M)$ such that $\lim_{i \rightarrow \infty} g_i(p) = p$. Then we can find a sequence of times t_i such that $\lim_{i \rightarrow \infty} g_i(x \cdot t_i) = x$, which means that $\lim_{i \rightarrow \infty} x_0 \cdot t_i = x_0$. Since the injectivity radius of M is bounded below, it follows that the t_i are unbounded, so we can take a subsequence with either $t_i \rightarrow \infty$ or $t_i \rightarrow -\infty$. Thus p is either ω - or α -recurrent. \square

Such a sequence $(g_i)_{i=1}^\infty$ is called a *recurrence sequence* for the recurrent point p . Note that p is ω -recurrent (α -recurrent) if and only if we can find a recurrence sequence with the property that $\lim_{i \rightarrow \infty} g_i(x \cdot t_i) = x$ for a point $x \in \langle p \rangle$ and a sequence of times $t_i \rightarrow \infty$ (resp. $t_i \rightarrow -\infty$). Such a sequence is called an ω -*sequence* (resp. α -*sequence*) for p .

More generally, we will denote the ω - and α -limit sets of a point $x \in M$ by $\omega(x)$ and $\alpha(x)$. That is, $y \in \omega(x)$ ($y \in \alpha(x)$) if and only if there is a sequence of times $t_i \rightarrow \infty$ (resp. $t_i \rightarrow -\infty$) such that $\lim_{i \rightarrow \infty} x \cdot t_i = y$. As with recurrence, we can treat ω - and α -limit sets on the level of orbits. Note that a point or orbit is ω - or α -recurrent if and only if it is contained in its own ω - or α -limit set. In the flowspace, we write $q \in \omega(p)$ or $q \in \alpha(p)$ whenever this holds for the corresponding orbits in M .

Lemma 3.7. *Let $p, q \in P$. Then $q \in \omega(p) \cup \alpha(p)$ if and only if there is a sequence of nontrivial elements $g_1, g_2, \dots \in \pi_1(M)$ such that $\lim_{i \rightarrow \infty} g_i(q) = p$.*

The proof is similar to that of the preceding lemma, and we can extend the notions of ω - and α -sequences in the obvious way. In particular, $q \in \omega(p)$ if and only if there is a sequence $(g_i)_{i=1}^\infty$ such that $\lim_{i \rightarrow \infty} g_i(x \cdot t_i) = y$ for points $x \in \langle p \rangle$ and $y \in \langle q \rangle$ and times $t_i \rightarrow \infty$. We call this an ω -*sequence* for $q \in \omega(p)$.

In the sequel, we will restrict attention to ω -sequences, though all of our results have corresponding versions for α -sequences.

4. SPRIGS

A *compactification* of a space X consists a compact space Y , together with an identification of X with a dense subset of Y . In [10], we showed that any finite collection $\{\mathcal{D}^i\}_{i=1}^k$ of unbounded decompositions of a plane P that intersect compactly determines a *universal compactification* to a closed disc $\mathbf{P} = P \sqcup S_u^1$ with interior P and boundary circle S_u^1 . This may be characterized by the following properties:

- (1) the closure of each $K \in \bigcup_i \mathcal{D}^i$ in \mathbf{P} intersects the boundary circle in a nontrivial totally disconnected set $\overline{K} \cap S_u^1$,
- (2)

$$\bigcup_{K \in \bigcup_i \mathcal{D}^i} \overline{K} \cap S_u^1$$

is dense in S_u^1 ,

- (3) any other compactification with these properties is a quotient of \mathbf{P} .

It follows that any group action $\Gamma \curvearrowright P$ that preserves each \mathcal{D}^i extends uniquely to an action on the corresponding universal compactification.

In particular, this construction can be applied to the decompositions \mathcal{D}^\pm that come from a quasigeodesic flow (in which case the boundary circle, together with the restricted action, is identified with the universal circle constructed by Calegari in [4]). In [11], we showed that the endpoint maps $e^\pm : P \rightarrow S_\infty^2$ extend uniquely to $\pi_1(M)$ -equivariant maps $\mathbf{e}^\pm : \mathbf{P} \rightarrow S_\infty^2$. These agree on the boundary circle, in the sense that $\mathbf{e}^+(s) = \mathbf{e}^-(s)$ for all $s \in S_u^1$.

4.1. The compactified flow space. For convenience, we will work with the variant of this compactification provided by the following theorem.

Theorem 4.1. *There is a compactification of P to a closed disc $\mathbf{P} = P \sqcup S_u^1$ with the following properties:*

- (1) *the flow space action $\pi_1(M) \curvearrowright P$ extends uniquely to an action $\pi_1(M) \curvearrowright \mathbf{P}$;*
- (2) *the endpoint maps $e^\pm : P \rightarrow S_\infty^2$ extend uniquely to $\pi_1(M)$ -equivariant maps*

$$\mathbf{e}^\pm : \mathbf{P} \rightarrow S_\infty^2;$$

- (3) *the extended endpoint maps agree on the universal circle, i.e.*

$$\mathbf{e}^+(s) = \mathbf{e}^-(s) \text{ for all } s \in S_u^1;$$

and

- (4) *for each $z \in S_\infty^2$,*

$$(\mathbf{e}^+)^{-1}(z) \cap S_u^1 = (\mathbf{e}^-)^{-1}(z) \cap S_u^1$$

is totally disconnected.

Proof. In [10, Theorem 7.9 & Construction 5.7] and [11, Theorem 2.10 & Proposition 2.11], we constructed a compactification $\mathbf{P} = P \sqcup S_u^1$ of P with extended endpoint maps $\mathbf{e}^\pm : \mathbf{P} \rightarrow S_\infty^2$ that satisfies properties (1)–(3). To obtain (4) we will pass to a quotient of this compactification.

Let

$$\mathcal{C} := \{\text{components of } (\mathbf{e}^+)^{-1}(z) \cap S_u^1 \mid z \in S_\infty^2\} \cup \{\{p\} \mid p \in P\}.$$

This is a monotone upper semicontinuous of \mathbf{P} by nonseparating subsets, so Moore's Theorem says that the decomposition space \mathbf{P}/\mathcal{C} is a closed disc.

The decomposition \mathcal{C} is trivial on the interior of \mathbf{P} , so the interior of the quotient is still identified with P . It is preserved by the action of $\pi_1(M)$, so there is an induced action on the quotient. Each element of \mathcal{C} maps to a single point under \mathbf{e}^+ and \mathbf{e}^- , so these descend to $\pi_1(M)$ -equivariant maps on the quotient. Thus the theorem is satisfied after replacing \mathbf{P} by \mathbf{P}/\mathcal{C} . \square

The space \mathbf{P} , together with the action of $\pi_1(M)$, is called the *compactified flow-space*. The boundary circle S_u^1 , together with the restricted action, is called the *universal circle*.

Since the extended endpoint maps agree on the universal circle, we will denote their mutual restriction by

$$\mathbf{e} : S_u^1 \rightarrow S_\infty^2.$$

This is $\pi_1(M)$ -equivariant, and since the action on S_∞^2 is minimal, its image must be the entire sphere at infinity. This generalizes the Cannon-Thurston Theorem, which produces such equivariant sphere-filling curves for fibered hyperbolic 3-manifolds [6].

4.2. Ends of leaves. Given a subset $A \subset \mathbf{P}$, we define $\partial A := A \cap S_u^1$. Given a positive or negative leaf $K \in \mathcal{D}^\pm$, the points in $\partial \bar{K}$ are called *ends* of K .¹

Lemma 4.2. *Each leaf has a nontrivial and totally disconnected set of ends.*

Proof. Nontriviality follows from that fact that the leaves are unbounded subsets of P . The ends of a leaf rooted at z are contained in the totally disconnected set $\mathbf{e}^{-1}(z)$, and hence totally disconnected. \square

Although distinct positive leaves $K, L \in \mathcal{D}^+$ are disjoint in the flow-space, their closures may intersect in the universal circle, so that they “share an end” $s \in \partial \bar{K} \cap \partial \bar{L}$. Negative leaves may also share ends, and a positive leaf may share an end with a negative leaf.

Lemma 4.3. *Any two leaves (both positive, both negative, or one of each) that share an end are rooted at the same point.*

Proof. If s is an end of a leaf $K \in \mathcal{D}^\pm$, then K is rooted at $\mathbf{e}(s)$. Indeed, if K is a positive leaf then its root is $e^+(K) = \mathbf{e}^+(\bar{K}) = \mathbf{e}^+(s) = \mathbf{e}(s)$, and if K is a negative leaf then its root is $e^-(K) = \mathbf{e}^-(\bar{K}) = \mathbf{e}^-(s) = \mathbf{e}(s)$.

Therefore, two leaves K and L that share an end $s \in \partial \bar{K} \cap \partial \bar{L}$ are both rooted at $\mathbf{e}(s)$. \square

On the other hand, positive and negative cannot simultaneously intersect and share ends.

Lemma 4.4. *If $K \in \mathcal{D}^+$ and $L \in \mathcal{D}^-$ share an end, then $K \cap L = \emptyset$.*

¹Our usage of the word “end” differs slightly from that of [10] and [11], where it refers to a *Freudenthal end*. Each Freudenthal end of a leaf K maps to a point in the universal circle, and $\partial \bar{K}$ is the closure of the image of K 's Freudenthal ends. We will not need this, but the reader may refer to [10, Lemma 7.8] for more details.

Proof. If K and L share an end $s \in \partial\bar{K} \cap \partial\bar{L}$, then we have $e^+(K) = \mathbf{e}(s) = e^-(L)$ by the preceding lemma. Then K and L must be disjoint, since any point $p \in K \cap L$ would have $e^+(p) = e^-(p)$, which contradicts the fact that quasigeodesics have distinct endpoints. \square

4.3. Sprigs. Since distinct positive/negative leaves can share ends, the collection of positive/negative leaf-closures does not form a decomposition of the compactified flowspace. In addition, there may be points in S_u^1 that are not in the closure of a positive or negative leaf. However, there are natural decompositions of \mathbf{P} obtained using the extended endpoint maps.

For each $z \in S_\infty^2$, the connected components of $(\mathbf{e}^+)^{-1}(z)$ and $(\mathbf{e}^-)^{-1}(z)$ are called, respectively, *positive sprigs rooted at z* and *negative sprigs rooted at z* . The collections of all positive and negative sprigs form a pair of $\pi_1(M)$ -invariant monotone decompositions of \mathbf{P} , the positive and negative *sprig decompositions*

$$\begin{aligned} \mathcal{D}^+ &:= \{\text{components of } (\mathbf{e}^+)^{-1}(z) \mid z \in S_\infty^2\} \\ \mathcal{D}^- &:= \{\text{components of } (\mathbf{e}^-)^{-1}(z) \mid z \in S_\infty^2\}. \end{aligned}$$

Let $\mathbf{K} \in \mathcal{D}^\pm$ be a positive or negative sprig. Then the points in $\partial\mathbf{K} = \mathbf{K} \cap S_u^1$ are called *ends* of \mathbf{K} , while $\mathring{\mathbf{K}} := \mathbf{K} \cap P$ is called its *bounded part*. Note that

$$\mathbf{K} = \mathring{\mathbf{K}} \sqcup \partial\mathbf{K}.$$

Lemma 4.5. *Each sprig has a nontrivial and totally disconnected set of ends.*

The proof is the same as in the case of leaves. A sprig $\mathbf{K} \in \mathcal{D}^\pm$ is said to be *trivial* if $\mathring{\mathbf{K}} = \emptyset$, and *nontrivial* otherwise.

Lemma 4.6. *Let $\mathbf{K} \in \mathcal{D}^\pm$ be a positive/negative sprig rooted at z .*

- (1) *If \mathbf{K} is trivial, then it consists a single point in S_u^1 .*
- (2) *If \mathbf{K} is nontrivial, then each component of $\mathring{\mathbf{K}}$ is a positive/negative leaf rooted at z .*

Proof. Suppose without loss of generality that \mathbf{K} is a positive sprig. If \mathbf{K} is trivial, then $\mathbf{K} = \partial\mathbf{K}$ it is a connected subset of the totally disconnected set $\mathbf{e}^{-1}(z)$, and hence a single point. If \mathbf{K} is nontrivial, then each component of $\mathring{\mathbf{K}}$ is a connected component of $(\mathbf{e}^+)^{-1}(z) \cap S_u^1 = (\mathbf{e}^+)^{-1}(z)$, which is a positive sprig rooted at z . \square

On the other hand:

Lemma 4.7. *Any two positive leaves that share an end are contained in the same positive sprig, and any two negative leaves that share an end are contained in the same negative sprig.*

Proof. If K and L are positive leaves that share an end $s \in \partial\bar{K} \cap \partial\bar{L}$, then $\bar{K} \cup \bar{L}$ is connected. By Lemma 4.3, $e^+(K \cup L)$ is a single point, so $\mathbf{e}^+(\bar{K} \cup \bar{L})$ is a single point. The same argument holds for negative sprigs using e^- and \mathbf{e}^- . \square

Distinct positive sprigs are disjoint by definition, as are distinct negative sprigs. However, a positive sprig might share an end with a negative sprig, and we have the following analogue of Lemmas 4.3 and 4.4.

Lemma 4.8. *Let $\mathbf{K} \in \mathcal{D}^+$ and $\mathbf{L} \in \mathcal{D}^-$ be positive and negative sprigs with $\partial\mathbf{K} \cap \partial\mathbf{L} \neq \emptyset$. Then $\mathbf{e}^+(\mathbf{K}) = \mathbf{e}^-(\mathbf{L})$, and $\mathring{\mathbf{K}} \cap \mathring{\mathbf{L}} = \emptyset$*

Conversely, sprigs whose bounded parts intersect must have disjoint ends, which we express as follows.

Corollary 4.9. *For each $p \in P$, we have $\partial\mathcal{D}^+(p) \cap \partial\mathcal{D}^-(p) = \emptyset$.*

Here, we are using the notation from §2.2: $\mathcal{D}^+(p)$ and $\mathcal{D}^-(p)$ are the \mathcal{D}^+ - and \mathcal{D}^- -saturations of the set $\{p\}$, which are just the positive and negative sprigs through p .

As illustrated in Figure 4, one can have a sequence of positive leaves that limits on more than one positive leaf. Note, however, that the two limit leaves share an end, and are therefore contained in a single positive sprig. This is a result of the following lemma, which is an immediate consequence of Lemmas 2.5 and 2.6.

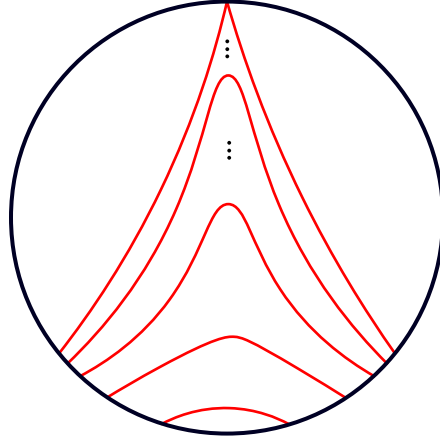


FIGURE 4. Leaves that split in the limit.

Lemma 4.10. *The sprig decompositions \mathcal{D}^+ and \mathcal{D}^- are upper semicontinuous.*

In particular, the sprig decompositions have the following properties from Definition 2.2 and Theorem 2.3.

- (1) Let $\mathbf{K}_1, \mathbf{K}_2, \dots$ be a sequence of positive (resp. negative) sprigs for which $\underline{\lim} \mathbf{K}_i \neq \emptyset$. Then $\overline{\lim} \mathbf{K}_i$ is contained in a single positive (negative) sprig.
- (2) Let p_1, p_2, \dots be a convergent sequence of points in \mathbf{P} . Then $\overline{\lim} \mathcal{D}^\pm(p_i) \subset \mathcal{D}^\pm(\lim p_i)$.
- (3) If $A \subset \mathbf{P}$ is closed, then the saturations $\mathcal{D}^+(A)$ and $\mathcal{D}^-(A)$ are closed.

We can abstract the properties of the sprig decompositions as follows.

Definition 4.11. A decomposition \mathcal{D} of the closed disc $\mathbf{D} = D \sqcup S^1$ is *spidery* if it is upper semicontinuous, and each decomposition element $\mathbf{K} \in \mathcal{D}$ intersects the boundary circle in a nontrivial totally disconnected set $\mathbf{K} \cap S^1$.

We do not yet have a satisfactory notion of transversality between the spidery decompositions \mathcal{D}^\pm . This is done in §5.5.

5. MASTER SPRIGS

Using the compactified flowspace, we will construct a compactification of \widetilde{M} , called the *flow ideal compactification*, that is especially adapted to the lifted flow. The boundary of this compactification is a 2-sphere S_u^2 , called the *universal sphere*, which we will use to understand the relationship between positive and negative sprigs.

Remark 5.1. Cannon-Thurston constructed flow ideal compactifications for pseudo-Anosov suspension flows [6] and Fenley generalized this to all pseudo-Anosov flows [8].

5.1. Master sets. For each $z \in S_\infty^2$, the set

$$Z = (\mathbf{e}^+)^{-1}(z) \cup (\mathbf{e}^-)^{-1}(z) \subset \mathbf{P},$$

which consists of all positive and negative sprigs rooted at z , is called the *master set rooted at z* .

Lemma 5.2. *Each master set is connected.*

Proof. Fix a master set Z rooted at a point $z \in S_\infty^2$, and let $D_1 \supset D_2 \supset \dots$ be a nested sequence of open discs in S_∞^2 centered at z , with $\bigcap_{i=1}^\infty D_i = z$. For each i , let $H_i \subset \widetilde{M}$ be the union of all geodesics with both endpoints in D_i .

Each H_i projects to a connected subset $X_i := \nu(H_i)$ of the flowspace, where $p \in X_i$ if and only if $\langle p \rangle$ intersects H_i . The closures \overline{X}_i are compact connected subsets of the compactified flowspace, so $\bigcap_i \overline{X}_i$ is compact and connected. To complete the proof, we will show that $\bigcap_i \overline{X}_i = Z$.

Let us show that $Z \subset \bigcap_i \overline{X}_i$. Let $p \in Z$. If $p \in Z \cap P$, then either $e^+(p) = z$ or $e^-(p) = z$, so we have $p \in X_i \subset \overline{X}_i$ for every i , and hence $p \in \bigcap_i \overline{X}_i$. If $p \in Z \cap S_u^1$, then $\mathbf{e}^+(p) = z = \mathbf{e}^-(p)$. Take a sequence of points $p_j \in P$ that converge to p , and note that $\lim e^+(p_j) = \mathbf{e}^+(p) = z$. Then after taking a subsequence of these points, we can assume that $e^+(p_j) \subset D_i$ for all $j \geq i$, which means that $p_j \in X_i$ for all $j \geq i$. Then $p = \lim_{j \rightarrow \infty} p_j \in \overline{X}_i$ for all i , and hence $p \in \bigcap_i \overline{X}_i$.

Now we will show that $\mathbf{P} \setminus Z$ is disjoint from $\bigcap_i \overline{X}_i$. Observe that if $A \subset S_\infty^2$ is compact, and $z \notin A$, then union of all flowlines with both ends in A is eventually disjoint from H_i , since it is contained in the C -neighborhood of union of all geodesics with both endpoints in A . Alternatively, if $B \subset P$ has the property that $e^+(B) \cup e^-(B)$ is bounded away from z , then we eventually have $B \cap X_i = \emptyset$.

Let p be a point in $\mathbf{P} \setminus Z$, which means that $\mathbf{e}^+(p) \neq z \neq \mathbf{e}^-(p)$. If $p \in P$, then this is the same as $e^+(p) \neq z \neq e^-(p)$, so we can find an open neighborhood $U \subset P$ of p such that $e^+(U) \cup e^-(U)$ is bounded away from z , and hence U is eventually disjoint from X_i . Then U is eventually disjoint from \overline{X}_i , and hence $p \notin \bigcap_i \overline{X}_i$. If $p \in S_u^1$, then we can find an open neighborhood $U \subset \mathbf{P}$ of p such that $\mathbf{e}^+(U) \cup \mathbf{e}^-(U)$ is bounded away from z . Then $e^+(U \cap P) \cup e^-(U \cap P)$ is bounded away from z , so $U \cap P$ is eventually disjoint from X_i . Any sequence of points in P that approaches a point in $U \cap S_u^1$ is eventually contained in $U \cap P$, so U is eventually disjoint from \overline{X}_i , and hence $p \notin \overline{X}_i$. \square

Let Z be a master set rooted at a point z . As with sprigs, we write $Z = \overset{\circ}{Z} \sqcup \partial Z$, where the points in $\partial Z = Z \cap S_u^1$ are called *ends* of Z and $\overset{\circ}{Z} := Z \cap P$ is called its *bounded part*. Since $\partial Z = \mathbf{e}^{-1}(z)$, each master sprig has a nontrivial, closed, and

totally disconnected set of ends. We say that Z is *trivial* if $\mathring{Z} = \emptyset$, in which case $Z = \partial Z$ consists of a single point, being a connected subset of a totally disconnected set.

Note, however, that the master sets do not form a decomposition of \mathbf{P} , since each point $p \in P$ is contained in two master sets, those rooted at $e^+(p)$ and $e^-(p)$.

5.2. The flow ideal compactification. Recall from §3.1 that we can identify \widetilde{M} with the open cylinder

$$C := P \times \mathbb{R}$$

in such a way that the flowlines of $\widetilde{\Phi}$ correspond to vertical lines $\{p\} \times \mathbb{R}$ in C . We can compactify C by thinking of it as the interior of the closed cylinder

$$\mathbf{C} := \mathbf{P} \times \overline{\mathbb{R}},$$

where $\overline{\mathbb{R}} = [-\infty, \infty]$ is the usual two-point compactification of $\mathbb{R} = (-\infty, \infty)$. The action of $\pi_1(M)$ on $C \simeq \widetilde{M}$ can be extended to the upper and lower horizontal faces

$$\partial_{\pm} \mathbf{C} := \mathbf{P} \times \{\pm\infty\}$$

of the closed cylinder, since they are identified with the compactified flowspace, but not to the vertical face

$$\partial_v \mathbf{C} := S_u^1 \times \overline{\mathbb{R}}.$$

However, it does extend to the “closed lens”

$$\mathbf{L} := \mathbf{C} / \{\{s\} \times \overline{\mathbb{R}} \mid s \in S_u^1\}$$

obtained by collapsing the vertical lines in the vertical face. This space, together with the action $\pi_1(M) \curvearrowright \mathbf{L}$, is called the *flow ideal compactification*.

Let S_u^2 denote the boundary sphere of \mathbf{L} , which, together with the action of $\pi_1(M)$, is called the *universal sphere*. This consists of two copies of the compactified flowspace glued along their universal circles. We will denote these by \mathbf{P}^+ and \mathbf{P}^- , and think of them as the upper and lower hemispheres of S_u^2 .

5.3. Master sprigs. We can think of the extended endpoint maps as maps

$$\mathbf{e}^{\pm} : \mathbf{P}^{\pm} \rightarrow S_{\infty}^2$$

supported on the two hemispheres of the universal sphere. These agree on the equator, so they combine to form a $\pi_1(M)$ -equivariant map

$$m : S_u^2 \rightarrow S_{\infty}^2.$$

Alternatively, the identifications $\mathbf{P}^{\pm} \simeq \mathbf{P}$ define a *flattening map*

$$f : S_u^2 \rightarrow \mathbf{P},$$

and we can define m by

$$m(p) = \begin{cases} \mathbf{e}^+ \circ f(p), & \text{for } p \in \mathbf{P}^+ \\ \mathbf{e}^- \circ f(p), & \text{for } p \in \mathbf{P}^-. \end{cases}$$

For each $z \in S_{\infty}^2$, the set $m^{-1}(z) \subset S_u^2$ consists of all positive sprigs rooted at z lying in the upper hemisphere, together with all negative sprigs rooted at z lying in the lower hemisphere. We call this the *master sprig rooted at z* . Note that the flattening map f takes each master sprig to the master set rooted at the same point. See Figure 5.




FIGURE 5. The universal sphere.

Although the master *sets* do not form a decomposition of \mathbf{P} , the master *sprigs* form a decomposition of S_u^2 , the *master decomposition*

$$\mathcal{M} := \{m^{-1}(z) \mid z \in S_\infty^2\}.$$

This is upper semicontinuous by Lemma 2.5, and monotone by the following lemma.

Lemma 5.3. *Each master sprig is connected.*

Proof. Suppose that some master sprig $Z = m^{-1}(z)$ is disconnected. Then we can write it as a disjoint union

$$Z = A \sqcup B$$

of nontrivial compact sets. Since f restricts to a homeomorphism from the equator to S_u^1 , and takes only points in the equator to S_u^1 , it follows that $f(A) \cap S_u^1$ is disjoint from $f(B) \cap S_u^1$. Furthermore, $f(A) \cap P$ is disjoint from $f(B) \cap P$, since $p \in f(A) \cap f(B) \cap P$ would imply that $e^+(p) = e^-(p)$, which contradicts the fact that quasigeodesics have distinct endpoints. Thus we can write the corresponding master set as a disjoint union of compact sets

$$f(Z) = f(A) \sqcup f(B),$$

which contradicts Lemma 5.2. \square

5.4. Recovering the sphere at infinity. The map $m : S_u^2 \rightarrow S_\infty^2$ takes each master sprig to a point, so it factors through a $\pi_1(M)$ -equivariant map

$$\hat{m} : S_u^2 / \mathcal{M} \rightarrow S_\infty^2$$

defined on the decomposition space of \mathcal{M} .

Lemma 5.4. \widehat{m} is a homeomorphism.

Proof. Since \mathcal{M} is upper semicontinuous, the decomposition space S_u^2/\mathcal{M} is compact Hausdorff by Lemma 2.4. Since \mathcal{M} is monotone, and m is surjective, the map \widehat{m} is bijective, and a continuous bijection between compact Hausdorff spaces is a homeomorphism. \square

Thus we recover the sphere at infinity as a quotient of the universal sphere. As a consequence, we have the following observation.

Lemma 5.5. Each master sprig is nonseparating in S_u^2 .

Proof. Suppose that some master sprig $Z = m^{-1}(z)$ separates S_u^2 . Choose some complementary component U of Z , and let V be the union of all other complementary components. Then U and V are open unions of master sprigs, so they map to disjoint open sets $\widehat{m}(U), \widehat{m}(V) \subset S_\infty^2$, and $S_\infty^2 = \widehat{m}(U) \sqcup z \sqcup \widehat{m}(V)$. But this means that z is a cutpoint of S_∞^2 , which is impossible. \square

5.5. Spidery pairs. Let $\mathbf{K} \in \mathcal{D}^+$ and $\mathbf{L} \in \mathcal{D}^-$ be positive and negative sprigs. If \mathbf{K} and \mathbf{L} share more than one end, then we can think of them as forming a kind of “ideal bigon” in \mathbf{P} . The following lemma says that this cannot happen. One also show that there are no “ideal polygons.” That is, one cannot have an alternating sequence of positive and negative sprigs $\mathbf{K}_0, \mathbf{K}_1, \mathbf{K}_2, \dots, \mathbf{K}_{n-1}$ where each \mathbf{K}_i shares an end with $\mathbf{K}_{i+1} \pmod{n}$. We leave this as an exercise since it will not be used directly.

Lemma 5.6 (No bigons). Let $\mathbf{K} \in \mathcal{D}^+$ and $\mathbf{L} \in \mathcal{D}^-$ be positive and negative sprigs. Then $\partial\mathbf{K}$ intersects $\partial\mathbf{L}$ in at most one point.

Proof. We will use following fact from classical analysis situs [30, Theorem II.5.28a]: If A and B are compact connected subsets of S^2 , and $A \cap B$ is disconnected, then $A \cup B$ separates S^2 .

Think of \mathbf{K} as a subset of \mathbf{P}^+ and \mathbf{L} as a subset of \mathbf{P}^- . If they intersect in S_u^1 then they are contained in a single master sprig. If they intersect at more than one point in S_u^1 , then this master sprig is separating, contradicting the preceding lemma. \square

It follows that the sprig decompositions satisfy the following definition.

Definition 5.7. A *spidery pair* consists of two spidery decompositions \mathcal{D}^\pm of the closed disc $\mathbf{D} = D \sqcup S^1$ with the following property: for each $\mathbf{K} \in \mathcal{D}^+$ and $\mathbf{L} \in \mathcal{D}^-$, the intersection $\mathbf{K} \cap \mathbf{L}$ is either empty, a compact subset of D , or a single point in S^1 .

5.6. Fixed sprigs. As with sprigs and master sets, a master sprig is said to be *trivial* if it is contained entirely in S_u^1 , thought of as the equator of S_u^2 . Equivalently, a master sprig $m^{-1}(z)$ is trivial if and only if the corresponding master set $f(m^{-1}(z))$ is trivial. We will use the following lemma to simplify the process of finding closed orbits.

Lemma 5.8. Let g be a nontrivial element of $\pi_1(M)$ that fixes a leaf, nontrivial sprig, nontrivial master set, or nontrivial master sprig. Then g fixes a point in P .

Proof. We will use the Brouwer Plane Translation Theorem (cf. [12]), which says that an orientation-preserving homeomorphism of the plane with a bounded forward orbit must have a fixed point.

If g fixes a leaf, nontrivial sprig, or nontrivial master sprig, then it fixes the corresponding master set, which is nontrivial. Thus we can assume that g fixes a nontrivial master set Z .

In the sphere at infinity S_∞^2 , g has exactly two fixed points α and ρ in an attracting-repelling pair, so Z must be rooted at one of these. After possibly replacing g by its inverse, we can assume that Z is rooted at the repelling fixed point ρ .

Since Z is nontrivial, we can choose a point $p \in Z \cap P$. If p is contained in a positive subleaf of Z , then $e^+(p) = \rho$ and $e^-(p) \neq \rho$, so we have $\lim_{i \rightarrow \infty} e^+(g^i(p)) = \rho$ and $\lim_{i \rightarrow \infty} e^-(g^i(p)) = \alpha$. Similarly, if p is contained in a negative subleaf, then $\lim_{i \rightarrow \infty} e^-(g^i(p)) = \rho$ and $\lim_{i \rightarrow \infty} e^+(g^i(p)) = \alpha$.

Either way, this means that $\lim_{i \rightarrow \infty} e^+(g^i(p)) \neq \lim_{i \rightarrow \infty} e^-(g^i(p))$. Since e^+ agrees with e^- on S_u^1 , this means that $g^i(p)$ stays in a bounded subset of P for all $i \geq 0$. Then g fixes some point in P by the Brouwer Plane Translation theorem, as does g^{-1} . \square

6. DECOMPOSITIONS I: SEPARATION

In this section we study the structure of the individual sprig decompositions, and relate the separation properties of sprigs with those of their ends. In particular, we will see that any two positive or negative sprigs are separated from each other by an interval's worth of positive or negative sprigs. The relationship between the positive and negative sprig decompositions is covered in the following section.

Throughout this section, take \mathcal{D} to be either \mathcal{D}^+ or \mathcal{D}^- , and “sprig” to mean a positive or negative sprig accordingly. In fact, the results in this section apply to any spidery decomposition of the closed disc (Definition 4.11).

Lemma 6.1. *Some sprig $\mathbf{K} \in \mathcal{D}$ has at least two ends.*

Proof. Suppose that each sprig has exactly one end. Then we can define a map

$$r : \mathbf{P} \rightarrow S_u^1$$

that sends each point $p \in \mathbf{P}$ to the end $r(p)$ of its sprig. This is continuous by upper semicontinuity, and it restricts to the identity on S_u^1 . That is, it is a retraction of \mathbf{P} onto S_u^1 , which is impossible. \square

6.1. Complementary regions and intervals. Fix an orientation on S_u^1 . Then any ordered pair of points $s, t \in S_u^1$ determines an oriented open subinterval $(s, t) \subset S_u^1$ running from s to t . This makes sense even when $s = t$, where we take $(s, s) = S_u^1 \setminus \{s\}$.

Given a closed subset $A \subset \mathbf{P}$, the connected components of $\mathbf{P} \setminus A$ are called *complementary regions* of A , while the connected components of $S_u^1 \setminus \partial A$ are called *complementary intervals* of ∂A (or of A). These are open intervals whose initial and terminal points lie in ∂A .

Lemma 6.2. *Let U be a complementary region of a sprig $\mathbf{K} \in \mathcal{D}$. Then $U \cap S_u^1$ is a complementary interval of $\partial \mathbf{K}$.*

Proof. Choose an arbitrary point $p \in U$. Then the sprig $\mathcal{D}(p)$ is contained in U , so $U \cap S_u^1$ is nontrivial since it contains $\partial\mathcal{D}(p)$. Each complementary interval of $\partial\mathbf{K}$ that intersects $U \cap S_u^1$ is contained in it, so $U \cap S_u^1$ is a nontrivial union of complementary intervals.

Suppose that $U \cap S_u^1$ contains two distinct complementary intervals I and J . Since U is path-connected (it is a connected open subspace of a locally path-connected space), we can find an arc $c \subset U$ with initial point in I and terminal point in J . This separates the endpoints of I , which are contained in \mathbf{K} , so it separates \mathbf{K} . But \mathbf{K} is connected, so $U \cap S_u^1$ must be a single complementary interval. \square

Since each complementary interval is contained in a complementary region, it follows that $U \mapsto U \cap S_u^1$ defines a bijection between the complementary regions of a sprig and the complementary intervals of its ends. In particular, an n -ended sprig \mathbf{K} has exactly n complementary regions.

Corollary 6.3. *Let $\mathbf{K}, \mathbf{L} \in \mathcal{D}$ be distinct. Then $\partial\mathbf{K}$ is contained in a single complementary interval of $\partial\mathbf{L}$.*

Note that this relies on the fact that sprigs are closed. For example, [7] contains an illustration of two disjoint, connected, non-closed subsets of the square that connect opposite pairs of corners.

The following easy corollary says that the separation properties of sprigs can be seen in terms of their ends.

Corollary 6.4. *Let $\mathbf{K}, \mathbf{K}', \mathbf{K}'' \in \mathcal{D}$ be distinct. Then \mathbf{K}' separates \mathbf{K} from \mathbf{K}'' if and only if $\partial\mathbf{K}'$ separates $\partial\mathbf{K}$ from $\partial\mathbf{K}''$.*

6.2. Saturated continua. The (\mathcal{D}) -saturation of a connected set is connected, and the saturation of a closed set is closed by upper semicontinuity, cf. §2.2. A closed, connected, saturated subset of \mathbf{P} is called a *saturated continuum*.

The preceding results generalize easily as follows.

Lemma 6.5. *Let U be a complementary region of a saturated continuum $\mathbf{A} \subset \mathbf{P}$. Then $U \cap S_u^1$ is a complementary interval of $\partial\mathbf{A}$.*

Corollary 6.6. *Let $\mathbf{A}, \mathbf{B} \subset \mathbf{P}$ be disjoint saturated continua. Then $\partial\mathbf{A}$ is contained in a single complementary interval of $\partial\mathbf{B}$.*

Corollary 6.7. *Let $\mathbf{A}, \mathbf{A}', \mathbf{A}'' \subset \mathbf{P}$ be disjoint saturated continua. Then \mathbf{A}' separates \mathbf{A} from \mathbf{A}'' if and only if $\partial\mathbf{A}'$ separates $\partial\mathbf{A}$ from $\partial\mathbf{A}''$.*

Let $\mathbf{A} \subset \mathbf{P}$ be a saturated continuum. For each complementary region U of \mathbf{A} , we define the corresponding *face* of \mathbf{A} by

$$F(U) := \overline{U} \cap \mathbf{A}.$$

Lemma 6.8. *Let U be a complementary region of a saturated continuum $\mathbf{A} \subset \mathbf{P}$. Then $F(U)$ is a connected subset of a single sprig.*

Proof. The face $F(U)$ is connected because the disc has the *Brouwer property* [30, § II.4]: If $X \subset \mathbf{P}$ is closed and connected, and U is a component of $\mathbf{P} \setminus X$, then the frontier of U is closed and connected.

Let $a, b \in \partial\mathbf{A}$ be the endpoints of the corresponding complementary interval $U \cap S_u^1 = (a, b)$, and note that $a, b \in F(U)$. Let $\mathbf{K}_a = \mathcal{D}(a)$ and $\mathbf{K}_b = \mathcal{D}(b)$ be the

sprigs through these points. We will show that $F(U) \subset \mathbf{K}_a \cup \mathbf{K}_b$. Since $F(U)$ is connected, this implies that $\mathbf{K}_a = \mathbf{K}_b$, which completes the lemma.

Fix a point $p \in F(U)$, and choose a sequence of points $p_1, p_2, \dots \in U$ that converge to p . The ends of each $\mathcal{D}(p_i)$ are contained in (a, b) , so they must accumulate on either a or b because $\overline{\lim} \mathcal{D}(p_i) \subset \mathcal{D}(p) \subset A$. Thus we have either $\mathcal{D}(p) = \mathbf{K}_a$ or $\mathcal{D}(p) = \mathbf{K}_b$. This applies for every point in $F(U)$, so $F(U) \subset \mathbf{K}_a \cup \mathbf{K}_b$ as desired. \square

6.3. Separating sprigs. Given a pair of distinct sprigs $\mathbf{K}, \mathbf{L} \in \mathcal{D}$, let $U(\mathbf{K}|\mathbf{L})$ be the complementary region of \mathbf{K} that contains \mathbf{L} , and let $U(\mathbf{L}|\mathbf{K})$ be the complementary region of \mathbf{L} that contains \mathbf{K} . The intersection

$$U(\mathbf{K}, \mathbf{L}) := U(\mathbf{K}|\mathbf{L}) \cap U(\mathbf{L}|\mathbf{K})$$

is called the *region between \mathbf{K} and \mathbf{L}* .

Lemma 6.9. *Let $\mathbf{K}, \mathbf{L} \in \mathcal{D}$ be distinct. Then $U(\mathbf{K}, \mathbf{L})$ is path-connected. Moreover,*

$$U(\mathbf{K}, \mathbf{L}) \cap S_u^1 = I(\mathbf{K}, \mathbf{L}) \sqcup I(\mathbf{L}, \mathbf{K}),$$

where $I(\mathbf{K}, \mathbf{L})$ ($I(\mathbf{L}, \mathbf{K})$) is the unique complementary interval of $\partial\mathbf{K} \cup \partial\mathbf{L}$ with initial endpoint in $\partial\mathbf{K}$ (resp. $\partial\mathbf{L}$) and terminal endpoint in $\partial\mathbf{L}$ (resp. $\partial\mathbf{K}$).

Proof. Let $\mathbf{A} = \mathbf{P} \setminus U(\mathbf{K}|\mathbf{L})$ and $\mathbf{B} = \mathbf{P} \setminus U(\mathbf{L}|\mathbf{K})$, which are disjoint, nonseparating saturated continua. Then $\mathbf{A} \cup \mathbf{B}$ is nonseparating because the disc has the *Phragmen-Brouwer property* [30, § II.4]: If $\mathbf{A}, \mathbf{B} \subset \mathbf{P}$ are disjoint and nonseparating, then $\mathbf{A} \cup \mathbf{B}$ is nonseparating. Then $U(\mathbf{K}, \mathbf{L})$ is connected, and hence path-connected, since it can be written as $U(\mathbf{K}, \mathbf{L}) = \mathbf{P} \setminus (\mathbf{A} \cup \mathbf{B})$.

Note that $U(\mathbf{K}|\mathbf{L}) \cap S_u^1 = (k, k')$ for $k, k' \in \partial\mathbf{K}$, while $U(\mathbf{L}|\mathbf{K}) \cap S_u^1 = (l', l)$ for $l', l \in \partial\mathbf{L}$. See the left side of Figure 6. Thus $U(\mathbf{K}, \mathbf{L}) \cap S_u^1 = U(\mathbf{K}|\mathbf{L}) \cap U(\mathbf{L}|\mathbf{K}) \cap S_u^1 = (k, l) \sqcup (l', k')$. By Corollary 6.3, $\partial\mathbf{K}$ and $\partial\mathbf{L}$ do not separate each other, so (k, l) is the unique complementary interval of $\partial\mathbf{K} \cup \partial\mathbf{L}$ with initial point in $\partial\mathbf{K}$ and terminal point in $\partial\mathbf{L}$, and similarly for (l', k') . \square




FIGURE 6. Separating sprigs.

It follows that a sprig $\mathbf{K}' \in \mathcal{D}$ is contained in $U(\mathbf{K}, \mathbf{L})$ if and only if $\partial\mathbf{K}' \subset I(\mathbf{K}, \mathbf{L}) \sqcup I(\mathbf{L}, \mathbf{K})$. In addition, the preceding lemma, together with Corollary 6.4, implies the following.

Lemma 6.10. *Let $\mathbf{K}, \mathbf{L} \in \mathcal{D}$ be distinct. Then $\mathbf{K}' \in \mathcal{D}$ separates \mathbf{K} from \mathbf{L} if and only if $\partial\mathbf{K}'$ intersects both $I(\mathbf{K}, \mathbf{L})$ and $I(\mathbf{L}, \mathbf{K})$.*

Proposition 6.11. *Let $\mathbf{K}, \mathbf{L} \in \mathcal{D}$ be distinct. Then some $\mathbf{K}' \in \mathcal{D}$ separates \mathbf{K} from \mathbf{L} .*

Proof. See the right side of Figure 6. Since $U(\mathbf{K}, \mathbf{L})$ is path-connected, we can find an arc $c \subset U(\mathbf{K}, \mathbf{L})$ with initial point in $I(\mathbf{K}, \mathbf{L})$ and terminal point in $I(\mathbf{L}, \mathbf{K})$. Let $\mathbf{C} = \mathcal{D}(c)$ be the saturation of c , which is contained in $U(\mathbf{K}, \mathbf{L})$, and let U be the complementary region of \mathbf{C} that contains \mathbf{K} . Then the face $F(U)$ separates \mathbf{K} from \mathbf{L} , since it has points in both $I(\mathbf{K}, \mathbf{L})$ and $I(\mathbf{L}, \mathbf{K})$, and $F(U)$ is contained in a sprig \mathbf{K}' by Lemma 6.8. \square

Note that this provides an alternative proof of Lemma 6.1: simply fix two sprigs and take any sprig that separates them.

6.4. Separation intervals. Let us study the collection of all sprigs that separate two fixed sprigs.

Definition 6.12. Let $(\mathbf{K}_\alpha, \mathbf{K}_\beta)$ be an ordered pair of distinct sprigs. The corresponding *separation interval* is the set

$$\sigma(\mathbf{K}_\alpha, \mathbf{K}_\beta) := \{\mathbf{K} \in \mathcal{D} \mid \mathbf{K} \text{ separates } \mathbf{K}_\alpha \text{ from } \mathbf{K}_\beta\},$$

together with the binary relation \prec defined by setting $\mathbf{K} \prec \mathbf{K}'$ whenever \mathbf{K} separates \mathbf{K}_α from \mathbf{K}' .

We will show that this defines a linear order, and that the separation interval between any pair of distinct sprigs is order-isomorphic to the real line.

Fix an ordered pair of distinct sprigs $(\mathbf{K}_\alpha, \mathbf{K}_\beta)$. For brevity, we write $\sigma_{\alpha\beta} := \sigma(\mathbf{K}_\alpha, \mathbf{K}_\beta)$ for the corresponding separation interval, and $U_{\alpha\beta} := U(\mathbf{K}_\alpha, \mathbf{K}_\beta)$ for the region between \mathbf{K}_α and \mathbf{K}_β . By Lemma 6.9, this is connected, and it intersects S_u^1 in the disjoint union of two intervals. We abbreviate and label these intervals by

$$\begin{aligned} I_{\alpha\beta} &:= I(\mathbf{K}_\alpha, \mathbf{K}_\beta) = (k_\alpha, k_\beta) \\ I_{\beta\alpha} &:= I(\mathbf{K}_\beta, \mathbf{K}_\alpha) = (l_\beta, l_\alpha) \end{aligned}$$

where $k_\alpha, l_\alpha \in \partial\mathbf{K}_\alpha$ and $l_\beta, k_\beta \in \partial\mathbf{K}_\beta$. See the left side of Figure 7.

For each $\mathbf{K} \in \mathcal{D}$, define

$$\begin{aligned} \partial_{\alpha\beta}\mathbf{K} &:= \mathbf{K} \cap I_{\alpha\beta} \\ \partial_{\beta\alpha}\mathbf{K} &:= \mathbf{K} \cap I_{\beta\alpha}. \end{aligned}$$

By Lemma 6.10, $\mathbf{K} \in \sigma_{\alpha\beta}$ if and only if $\partial_{\alpha\beta}\mathbf{K} \neq \emptyset \neq \partial_{\beta\alpha}\mathbf{K}$.

Definition 6.13. Given subsets X, Y of linearly ordered set $(Z, <)$, we take $X < Y$ to mean that $x < y$ for all $x \in X$ and $y \in Y$. We say that X and Y are *comparable* if either $X < Y$ or $Y < X$.

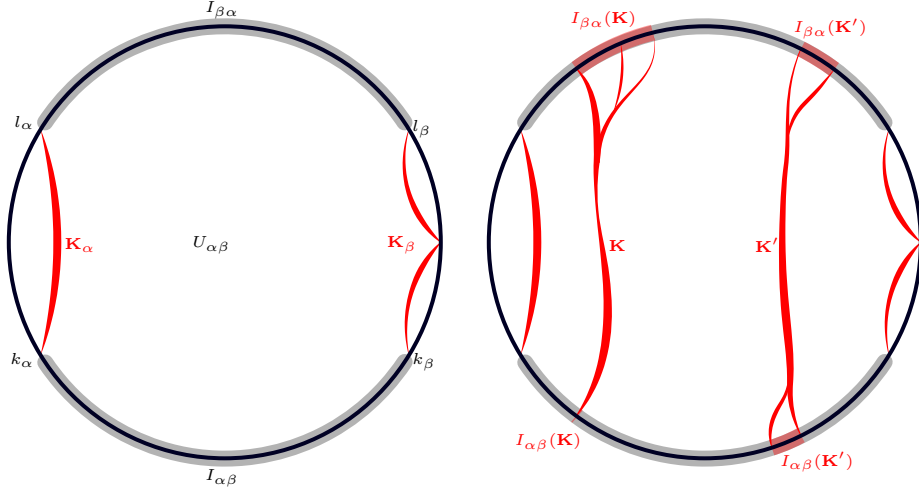


FIGURE 7. The separation interval.

The intervals $I_{\alpha\beta}$ and $I_{\beta\alpha}$ come with linear orders induced by their orientations. For distinct $\mathbf{K}, \mathbf{K}' \in \sigma_{\alpha\beta}$, we would like to say that $\partial_{\alpha\beta}\mathbf{K}$ and $\partial_{\alpha\beta}\mathbf{K}'$ are comparable in $I_{\alpha\beta}$, and that $\partial_{\beta\alpha}\mathbf{K}$ and $\partial_{\beta\alpha}\mathbf{K}'$ are comparable in $I_{\beta\alpha}$.

For each $\mathbf{K} \in \sigma_{\alpha\beta}$, let $I_{\alpha\beta}(\mathbf{K})$ be the minimal sub-interval of $I_{\alpha\beta}$ that contains $\partial_{\alpha\beta}\mathbf{K}$, and let $I_{\beta\alpha}(\mathbf{K})$ be the minimal sub-interval of $I_{\beta\alpha}$ that contains $\partial_{\beta\alpha}\mathbf{K}$. Then $\partial_{\alpha\beta}\mathbf{K}$ and $\partial_{\alpha\beta}\mathbf{K}'$ are comparable if and only if $I_{\alpha\beta}(\mathbf{K})$ and $I_{\alpha\beta}(\mathbf{K}')$ are disjoint, and similarly for $\partial_{\beta\alpha}\mathbf{K}$ and $\partial_{\beta\alpha}\mathbf{K}'$. See the right side of Figure 7.

Lemma 6.14. *If $\mathbf{K}, \mathbf{K}' \in \sigma_{\alpha\beta}$ are distinct, then*

$$I_{\alpha\beta}(\mathbf{K}) \cap I_{\alpha\beta}(\mathbf{K}') = \emptyset = I_{\beta\alpha}(\mathbf{K}) \cap I_{\beta\alpha}(\mathbf{K}').$$

Proof. If $I_{\alpha\beta}(\mathbf{K})$ intersects $I_{\alpha\beta}(\mathbf{K}')$, then after possibly switching \mathbf{K} with \mathbf{K}' , there must be some end of \mathbf{K} in the interior of $I_{\alpha\beta}(\mathbf{K}')$. But the ends of \mathbf{K} are contained in a single complementary interval of \mathbf{K}' (Corollary 6.3), so this means that *all* of the ends of \mathbf{K} are contained in the interior of $I_{\alpha\beta}(\mathbf{K}')$. This is impossible, since \mathbf{K} must also have an end in $I_{\beta\alpha}$. Thus we must have $I_{\alpha\beta}(\mathbf{K}) \cap I_{\alpha\beta}(\mathbf{K}') = \emptyset$, and a similar argument shows that $I_{\beta\alpha}(\mathbf{K}) \cap I_{\beta\alpha}(\mathbf{K}') = \emptyset$. \square

Consequently, the linear orders on $I_{\alpha\beta}$ and $I_{\beta\alpha}$ induce linear orders on

$$\begin{aligned} \mathcal{I}_{\alpha\beta} &:= \{I_{\alpha\beta}(\mathbf{K}) \mid \mathbf{K} \in \sigma\} \\ \text{and } \mathcal{I}_{\beta\alpha} &:= \{I_{\beta\alpha}(\mathbf{K}) \mid \mathbf{K} \in \sigma\}. \end{aligned}$$

By the following lemma, we can think of $\mathbf{K} \mapsto I_{\alpha\beta}(\mathbf{K})$ as an order-isomorphism $\sigma_{\alpha\beta} \rightarrow \mathcal{I}_{\alpha\beta}$, and $\mathbf{K} \mapsto I_{\beta\alpha}(\mathbf{K})$ as an anti-order-isomorphism $\sigma_{\alpha\beta} \rightarrow \mathcal{I}_{\beta\alpha}$.

Lemma 6.15. *The following are equivalent for all $\mathbf{K}, \mathbf{K}' \in \sigma$:*

- (1) $I_{\alpha\beta}(\mathbf{K}) < I_{\alpha\beta}(\mathbf{K}')$,
- (2) $I_{\beta\alpha}(\mathbf{K}) > I_{\beta\alpha}(\mathbf{K}')$,
- (3) \mathbf{K} separates \mathbf{K}_{α} from \mathbf{K}' , and
- (4) \mathbf{K}' separates \mathbf{K} from \mathbf{K}_{β} .

Proof. The equivalence (1) \Leftrightarrow (2) follows from Corollary 6.3, while the equivalences (1) + (2) \Leftrightarrow (3) and (1) + (2) \Leftrightarrow (4) follow from Corollary 6.4. \square

In particular:

Corollary 6.16. *Let $\mathbf{K}_\alpha, \mathbf{K}_\beta \in \mathcal{D}$ be distinct. Then the separation interval $\sigma(\mathbf{K}_\alpha, \mathbf{K}_\beta)$ is a linearly ordered set.*

Let

$$\bigcup \mathcal{I}_{\alpha\beta} := \bigcup_{\mathbf{K} \in \sigma_{\alpha\beta}} I_{\alpha\beta}(\mathbf{K})$$

$$\text{and } \bigcup \mathcal{I}_{\beta\alpha} := \bigcup_{\mathbf{K} \in \sigma_{\alpha\beta}} I_{\beta\alpha}(\mathbf{K}).$$

We will use the following lemma to show that $\sigma_{\alpha\beta}$ is order-isomorphic to the real line.

Lemma 6.17. $\bigcup \mathcal{I}_{\alpha\beta} = I_{\alpha\beta}$ and $\bigcup \mathcal{I}_{\beta\alpha} = I_{\beta\alpha}$.

Proof. Let \mathbf{A} be the union of $\mathbf{K}_\alpha, \mathbf{K}_\beta$, and all sprigs in \mathcal{D} that intersect $I_{\beta\alpha}$. This is a saturated continuum, since it is the saturation of the closed interval $\overline{I_{\beta\alpha}}$.

Fix a point $s \in I_{\alpha\beta}$. If $s \in \mathbf{A}$, then the sprig $\mathcal{D}(s)$ is contained in $\sigma_{\alpha\beta}$ since it has ends in both $I_{\alpha\beta}$ and $I_{\beta\alpha}$, and hence $s \in \bigcup \mathcal{I}_{\alpha\beta}$.

If $s \notin \mathbf{A}$, then it is contained in $U \cap S_u^1$ for some complementary region U of \mathbf{A} . By Lemma 6.8, the face $F(U)$ is contained in some sprig $\mathbf{K} \subset \mathbf{A}$. It follows that $\mathbf{K} \in \sigma_{\alpha\beta}$, and $s \in I_{\alpha\beta}(\mathbf{K})$, so we have once again that $s \in \bigcup \mathcal{I}_{\alpha\beta}$. Thus $\bigcup \mathcal{I}_{\alpha\beta} = I_{\alpha\beta}$, and a similar argument shows that $\bigcup \mathcal{I}_{\beta\alpha} = I_{\beta\alpha}$. \square

Proposition 6.18. *Let $\mathbf{K}_\alpha, \mathbf{K}_\beta \in \mathcal{D}$ be distinct. Then the separation interval $\sigma(\mathbf{K}_\alpha, \mathbf{K}_\beta)$ is order-isomorphic to the real line.*

Proof. The preceding lemma implies that $\mathcal{I}_{\alpha\beta}$ is a separable and complete linearly ordered set with no maximum and minimum, which characterizes the linear order on \mathbb{R} . Since $\sigma(\mathbf{K}_\alpha, \mathbf{K}_\beta)$ is order-isomorphic to $\mathcal{I}_{\alpha\beta}$ it is also order-isomorphic to \mathbb{R} . \square

The following properties of separation intervals follow immediately from Lemma 6.15.

Lemma 6.19. *Let $\mathbf{K}_\alpha, \mathbf{K}_\beta \in \mathcal{D}$ be distinct.*

- (1) *Then $\mathbf{K} \mapsto \mathbf{K}$ defines an anti-order-isomorphism $\sigma(\mathbf{K}_\alpha, \mathbf{K}_\beta) \rightarrow \sigma(\mathbf{K}_\beta, \mathbf{K}_\alpha)$.*
- (2) *Let $\mathbf{K}_a, \mathbf{K}_b \in \sigma(\mathbf{K}_\alpha, \mathbf{K}_\beta)$, with $\mathbf{K}_a \prec \mathbf{K}_b$. Then*

$$\sigma(\mathbf{K}_a, \mathbf{K}_b) = \{\mathbf{K} \in \sigma(\mathbf{K}_\alpha, \mathbf{K}_\beta) \mid \mathbf{K}_a \prec \mathbf{K} \prec \mathbf{K}_b\},$$

and the inclusion $\sigma(\mathbf{K}_a, \mathbf{K}_b) \hookrightarrow \sigma(\mathbf{K}_\alpha, \mathbf{K}_\beta)$ is order-preserving.

7. DECOMPOSITIONS II: LINKING

Now that we have some tools to work with the individual sprig decompositions, we can study the relationship between them. In fact, the results in this section apply to any spidery pair in the disc (Definition 5.7).

Let $A, B \subset S^1$ be closed, disjoint, nontrivial subsets of the circle. We define the *linking number* $\text{lk}(A, B)$ to be the number of complementary intervals of A that intersect B , which is finite and symmetric by the following lemma. We say that A and B are *unlinked* if $\text{lk}(A, B) = 1$, and *linked* if $\text{lk}(A, B) \geq 2$. The terminology

comes from the fact that A and B are linked if and only if there are pairs $\{a, a'\} \subset A$ and $\{b, b'\} \subset B$ that are linked as 0-spheres in S^1 .

Lemma 7.1. *Let $A, B \subset S^1$ be closed, disjoint, and nontrivial. Then $\text{lk}(A, B)$ is finite and equal to $\text{lk}(B, A)$.*

Proof. Fix an orientation on S^1 . An AB -interstitial interval is an oriented interval $(a, b) \subset S^1$ that is disjoint from $A \cup B$, with $a \in A$ and $b \in B$. It's easy to see that each complementary interval of A that intersects B contains a unique AB -interstitial interval, and each AB -interstitial interval is contained in a unique complementary interval of A that intersects B . Similarly, each complementary interval of B that intersects A contains a unique AB -interstitial interval, and each AB -interstitial interval is contained in a unique complementary interval of B that intersects A . Thus there is a bijective correspondence between the complementary intervals of B that intersect A , the AB -interstitial intervals, and the complementary intervals of B that intersect A , and hence $\text{lk}(A, B) = \text{lk}(B, A)$.

To complete the lemma, it suffices to show that there are only finitely many AB -interstitial intervals. Otherwise, we would have an infinite sequence of distinct AB -interstitial intervals $(a_i, b_i)_{i=1}^{\infty}$. These are pairwise disjoint intervals in the circle, so their diameters must go to zero, and we can assume after taking a subsequence that they converge to a single point $s \in S^1$. Then since A and B are closed, we have $s = \lim a_i \in A$ and $s = \lim b_i \in B$, which contradicts the assumption that A and B are disjoint. \square

The following lemma is an immediate consequence of the usual pigeonhole principle; we will use it repeatedly.

Lemma 7.2 (Linking Pigeonhole Principle). *Let $A, B \subset S^1$ be closed, disjoint, and nontrivial, and let $X \subset S^1$ be a subset with $A \subset X$. If X has fewer than $\text{lk}(A, B)$ connected components, then $X \cap B \neq \emptyset$.*

In particular, if $A, B \subset S^1$ are linked, then any connected set that contains A must intersect B .

7.1. The linked region. Let $\mathbf{K} \in \mathcal{D}^+$ and $\mathbf{L} \in \mathcal{D}^-$ be positive and negative sprigs. If $\partial\mathbf{K} \cap \partial\mathbf{L} = \emptyset$, then we define their linking number by

$$\text{lk}(\mathbf{K}, \mathbf{L}) := \text{lk}(\partial\mathbf{K}, \partial\mathbf{L}).$$

We say that \mathbf{K} and \mathbf{L} are *unlinked* when $\text{lk}(\mathbf{K}, \mathbf{L}) = 1$ and *linked* when $\text{lk}(\mathbf{K}, \mathbf{L}) \geq 2$. If $\partial\mathbf{K} \cap \partial\mathbf{L} \neq \emptyset$, then their linking number is undefined, and they are neither linked nor unlinked. One could also define linking numbers for pairs of sprigs of the same kind, but these are always unlinked by Corollary 6.3.

The sprigs through a point $p \in P$ always have disjoint ends (Corollary 4.9), so

$$\text{lk}(p) := \text{lk}(\mathcal{D}^+(p), \mathcal{D}^-(p))$$

is always defined, and we call p a *linked point* or *unlinked point* accordingly. The set of all linked points is called the *linked region* and denoted by $P_l \subset P$. We will show that P_l is closed in P and nontrivial.

7.2. Linked sequences. The following lemma will be useful when considering sequences of linked points.

Lemma 7.3. *Let $(p_i)_{i=1}^\infty$ be a sequence of points in P_l that converge to a point $p \in P$. Then $\{p_i\}$ intersects at most finitely many complementary regions of $\mathcal{D}^+(p)$ and of $\mathcal{D}^-(p)$.*

Proof. Let $\mathbf{K}_i = \mathcal{D}^+(p_i)$ and $\mathbf{L}_i = \mathcal{D}^-(p_i)$ be the positive and negative sprigs through each p_i , and let $\mathbf{K} = \mathcal{D}^+(p)$ and $\mathbf{L} = \mathcal{D}^-(p)$ be the positive and negative sprigs through p . We have $\overline{\lim} \partial \mathbf{K}_i \subset \partial \mathbf{K}$ and $\overline{\lim} \partial \mathbf{L}_i \subset \partial \mathbf{L}$ by upper semicontinuity. Since $p \in P$, Corollary 4.9 implies that $\partial \mathbf{K}$ is disjoint from $\partial \mathbf{L}$, so $\overline{\lim} \partial \mathbf{K}_i$ is disjoint from $\overline{\lim} \partial \mathbf{L}_i$.

Suppose that $\{p_i\}$ intersects infinitely many complementary regions of \mathbf{K} . Then after taking a subsequence we can assume that each p_i is contained in a distinct complementary region U_i . Then $\mathbf{K}_i \subset U_i$ and $\partial \mathbf{K}_i \subset U_i \cap S_u^1 =: J_i$ for each i .

For each i , let k_i be an end of \mathbf{K}_i , which is contained in J_i . By the Linking Pigeonhole Principle, we can also find an end l_i of \mathbf{L}_i that lies in J_i . The J_i are pairwise disjoint intervals in the circle, so their diameters must go to zero, and we can assume after taking a subsequence that they converge to a single point $s \in S_u^1$. Then $\lim k_i = s = \lim l_i$, which contradicts the fact that $\overline{\lim} \partial \mathbf{K}_i$ is disjoint from $\overline{\lim} \partial \mathbf{L}_i$. Thus the p_i can visit only finitely many complementary regions of \mathbf{K} , and the same argument shows that they can visit only finitely many complementary regions of \mathbf{L} . \square

Lemma 7.4. *Let $(p_i)_{i=1}^\infty$ be a sequence of points in P_l that converge to a point $p \in P$, and suppose that there is an integer n such that $\text{lk}(p_i) \geq n$ for all i . Then $\text{lk}(p) \geq n$.*

Proof. We will continue to use the notation and observations in the first paragraph of the preceding proof.

For each i , and choose ends $k_i^0, k_i^1, \dots, k_i^n \in \partial \mathbf{K}_i$ and $l_i^0, l_i^1, \dots, l_i^n \in \partial \mathbf{L}_i$ such that $k_i^0, l_i^0, k_i^1, l_i^1, \dots, k_i^n, l_i^n$ is positively ordered in S_u^1 . After taking a subsequence, we can assume that $\lim_{i \rightarrow \infty} k_i^j$ converges to an end $k^j \in \partial \mathbf{K}$ for each j and $\lim_{i \rightarrow \infty} l_i^j$ converges to an end $l^j \in \partial \mathbf{L}$ for each j .

It suffices to show that $k^0, l^0, k^1, l^1, \dots, k^n, l^n$ are pairwise distinct and positively ordered in S_u^1 . Since $\partial \mathbf{K}$ is disjoint from $\partial \mathbf{L}$, it follows that $k^j \neq l^{j'}$ for all j, j' . Suppose that $k^j = k^{j'}$ for some $j \neq j'$. Then either $\lim_{i \rightarrow \infty} (k_i^j, k_i^{j'}) = k^j$ or $\lim_{i \rightarrow \infty} (k_i^{j'}, k_i^j) = k^j$. Either case is impossible, since each of these intervals contain at least one of the $l^{j''}$'s, which cannot accumulate on $k^j \in \partial \mathbf{K}$. Therefore, $k^j \neq k^{j'}$ for all $j \neq j'$ and a similar argument shows that $l^j \neq l^{j'}$ for all $j \neq j'$.

The fact that $k_i^0, l_i^0, k_i^1, l_i^1, \dots, k_i^n, l_i^n$ is positively ordered for each i easily implies that $k^0, l^0, k^1, l^1, \dots, k^n, l^n$ is positively ordered. \square

Proposition 7.5. *P_l is closed in P .*

Proof. Let (p_i) be a sequence of points in P_l that converge to $p \in P$. Then $\text{lk}(p_i) \geq 2$ for all i . By the preceding lemma, $\text{lk}(p) \geq 2$ which means that $p \in P_l$. \square

The following lemma will be used in the proof of the Homotopy Closing Lemma.

Lemma 7.6. *Let $(p_i)_{i=1}^\infty$ be a sequence of points in P that converge to a point $p \in P$, and suppose that $\text{lk}(p_i) \geq 3$ for all i . Then $p_i \in \mathcal{D}^+(p) \cap \mathcal{D}^-(p)$ for all sufficiently large i .*

Proof. With the notation above, suppose that there are infinitely many i for which $p_i \notin \mathbf{K}$. By Lemma 7.3 we can assume after passing to a subsequence that each \mathbf{K}_i is contained in the same complementary region U of \mathbf{K} .

Then the ends of each \mathbf{K}_i are contained in the corresponding complementary interval, which can be written in the form $U \cap S_u^1 = (k, k')$ for $k, k' \in \partial\mathbf{K}$. Thus $\limsup \partial\mathbf{K}_i \subset \{k, k'\}$, so we can find sequences of points $s_i, s'_i \in (k, k')$ such that $\partial\mathbf{K}_i(k, s_i) \sqcup (s'_i, k')$ for each i , and $\lim s_i = k$ and $\lim s'_i = k'$. Since \mathbf{K}_i and \mathbf{L}_i have linking number at least 3, the Linking Pigeonhole Principle implies that some end $l_i \in \partial\mathbf{L}_i$ is contained in either $\overline{(k, s_i)}$ or $\overline{(s'_i, k')}$. Then $\overline{\lim} l_i$ contains either k or k' , contradicting the fact that $\overline{\lim} \partial\mathbf{K}_i$ is disjoint from $\overline{\lim} \partial\mathbf{L}_i$. Thus $p_i \in \mathbf{K}$ for all but finitely many i , and a similar argument shows that $p_i \in \mathbf{L}$ for all but finitely many i . \square

7.3. Nontriviality.

Proposition 7.7. $P_l \neq \emptyset$.

Proof. By Lemma 6.1, we can find a positive sprig \mathbf{K} with two distinct ends k_α and k_β . By Lemma 5.6, the negative sprigs $\mathbf{N}_\alpha, \mathbf{N}_\beta$ through these points are distinct, so $\sigma(\mathbf{N}_\alpha, \mathbf{N}_\beta)$ is nontrivial. Each $\mathbf{N} \in \sigma(\mathbf{N}_\alpha, \mathbf{N}_\beta)$ separates k_α from k_β , but \mathbf{N} need not be linked with \mathbf{K} since we may have $\partial\mathbf{N} \cap \partial\mathbf{K} \neq \emptyset$. However, we can use the following lemma to find an $\mathbf{N} \in \sigma(\mathbf{N}_\alpha, \mathbf{N}_\beta)$ that intersects \mathbf{K} . Then $\partial\mathbf{N} \cap \partial\mathbf{K} = \emptyset$ (Corollary 4.9), so \mathbf{N} is linked with \mathbf{K} . \square

Lemma 7.8. *Let k_α and k_β be distinct ends of a positive sprig \mathbf{K} , and let $\mathbf{N}_\alpha = \mathcal{D}^+(k_\alpha)$ and $\mathbf{N}_\beta = \mathcal{D}^+(k_\beta)$. Then the separation interval $\sigma(\mathbf{N}_\alpha, \mathbf{N}_\beta)$ contains a dense set of sprigs that intersect \mathbf{K} .*

Proof. See Figure 8. Given a pair of sprigs $\mathbf{N}_a \prec \mathbf{N}_b \in \sigma(\mathbf{N}_\alpha, \mathbf{N}_\beta)$, we must find a sprig $\mathbf{N} \in \sigma(\mathbf{N}_a, \mathbf{N}_b)$ that intersects \mathbf{K} nontrivially. To see this, we will show that there are sprigs $\mathbf{N}'_a \prec \mathbf{N}'_b \in \sigma(\mathbf{N}_a, \mathbf{N}_b)$ such that $I(\mathbf{N}'_a, \mathbf{N}'_b)$ is disjoint from $\partial\mathbf{K}$. By the same argument, we can find sprigs $\mathbf{N}''_a \prec \mathbf{N}''_b \in \sigma(\mathbf{N}'_a, \mathbf{N}'_b)$ such that $I(\mathbf{N}''_a, \mathbf{N}''_b)$ is disjoint from $\partial\mathbf{K}$; note that $I(\mathbf{N}''_a, \mathbf{N}''_b) \subset I(\mathbf{N}'_a, \mathbf{N}'_b)$ is already disjoint from $\partial\mathbf{K}$. Then it suffices to take any $\mathbf{N} \in \sigma(\mathbf{N}''_a, \mathbf{N}''_b)$: such a sprig must intersect \mathbf{K} because it separates k_α from k_β , and the intersection is in \mathbf{K} as opposed to $\partial\mathbf{K}$ because the ends of \mathbf{N} are contained in $I(\mathbf{N}''_a, \mathbf{N}''_b) \cup I(\mathbf{N}''_b, \mathbf{N}''_a)$.

Fix $\mathbf{N}_a \prec \mathbf{N}_b \in \sigma(\mathbf{N}_\alpha, \mathbf{N}_\beta)$, and choose a point $x \in I_{ab} \setminus \partial\mathbf{K}$, where we are using the abbreviation $I_{ab} := I(\mathbf{N}_a, \mathbf{N}_b)$. By Lemma 6.17, $x \in I_{ab}(\mathbf{N}')$ for some sprig $\mathbf{N}' \in \sigma(\mathbf{N}_a, \mathbf{N}_b)$. Note that the initial and terminal points of $\partial_{ab}\mathbf{N}'$, which we will denote by $\partial_{ab}^\mp \mathbf{N}'$, cannot both be contained in $\partial\mathbf{K}$. Indeed, if $I_{ab}(\mathbf{N}')$ is a point, then $\partial_{ab}^- \mathbf{N}' = \partial_{ab}^+ \mathbf{N}' = x$, which was chosen outside of $\partial\mathbf{K}$; if it is an interval, then $\partial_{ab}^- \mathbf{N}'$ and $\partial_{ab}^+ \mathbf{N}'$ are distinct ends of \mathbf{N}' , which can share at most one end with \mathbf{K} .

If $\partial_{ab}^+ \mathbf{N}' \notin \partial\mathbf{K}$, then we take \mathbf{N}'_a to be \mathbf{N}' , and \mathbf{N}'_b to be any sprig in $\sigma(\mathbf{N}'_a, \mathbf{N}_b)$ that is close to \mathbf{N}'_a , in the sense that $\partial\mathbf{K}$ does not separate $\partial_{ab}^+ \mathbf{N}'_a$ from $\partial_{ab}^- \mathbf{N}'_b$. Then no end of \mathbf{K} lies in $I(\mathbf{N}'_a, \mathbf{N}'_b)$ as desired. If $\partial_{ab}^+ \mathbf{N}' \in \partial\mathbf{K}$, then $\partial_{ab}^- \mathbf{N}' \notin \partial\mathbf{K}$, and we take \mathbf{N}'_b to be \mathbf{N}' , and \mathbf{N}'_a to be any sprig in $\sigma(\mathbf{N}_a, \mathbf{N}'_b)$ that is close to \mathbf{N}'_b .

As noted above, we can complete the proof by repeating this argument. \square

Although the proposition does not need the full power of this lemma, we will need it for the Homotopy Closing Lemma.




FIGURE 8. A dense set of sprigs intersecting the bounded part of \mathbf{K} .

8. COARSE TRANSVERSE HYPERBOLICITY

In this section, we will see that the coarsely hyperbolic behavior of our flow is reflected in the action of an ω - or α -sequence on the flowspace.

8.1. Straightening flowlines. Given a flowline $x \cdot \mathbb{R}$, let $(x \cdot \mathbb{R})^g$ denote the corresponding oriented geodesic, running from $E^-(x)$ to $E^+(x)$. The nearest-point projection restricts to a proper surjective map

$$\rho_{x \cdot \mathbb{R}} : x \cdot \mathbb{R} \rightarrow (x \cdot \mathbb{R})^g$$

that moves each point a uniformly bounded distance, independent of x . To see this, recall that $x \cdot \mathbb{R}$ is contained in the C -neighborhood of $(x \cdot \mathbb{R})^g$, for a uniform constant C , which we can picture as a “banana” foliated by radius- C hyperbolic discs. The nearest-point projection takes each point in $x \cdot \mathbb{R}$ to the center of the corresponding disc, so $d(y, \rho_{x \cdot \mathbb{R}}(y)) < C$ for all $y \in x \cdot \mathbb{R}$.

Since the endpoints of flowlines vary continuously, we can define a continuous, $\pi_1(M)$ -equivariant *straightening map*

$$\begin{aligned} G : \widetilde{M} &\rightarrow \widetilde{M} \\ x &\mapsto \rho_{x \cdot \mathbb{R}}(x) \end{aligned}$$

that takes each flowline onto its corresponding geodesic, while moving each point by a distance of at most C .

Lemma 8.1. *There is a constant $D > 0$ such that:*

- (1) *If $x_0, x_1 \in \widetilde{M}$ are contained in the same weak positive leaf then there are times $t_0, t_1 \in \mathbb{R}$ such that $x_0 \cdot [t_0, \infty)$ and $x_1 \cdot [t_1, \infty)$ have Hausdorff distance at most D .*
- (2) *If $y_0, y_1 \in \widetilde{M}$ are contained in the same weak negative leaf then there are times $s_0, s_1 \in \mathbb{R}$ such that $y_0 \cdot (-\infty, s_0]$ and $y_1 \cdot (-\infty, s_1]$ have Hausdorff distance at most D .*

Proof. Let $D = 2C + \epsilon$ for an arbitrary constant $\epsilon > 0$. We will prove (1), and (2) follows from a similar argument.

Since x_0 and x_1 are contained in the same weak positive leaf, the geodesics $G(x_0 \cdot \mathbb{R})$ and $G(x_1 \cdot \mathbb{R})$ have the same positive endpoint $z = E^+(x_0) = E^+(x_1)$. Therefore, we can find a horosphere S centered at z such that the distance between $G(x_0 \cdot \mathbb{R}) \cap S'$ and $G(x_1 \cdot \mathbb{R}) \cap S'$ is less than ϵ for every horosphere S' centered at z that lies forward of S . Since G moves each point by a distance of at most C , it suffices to take $t_i = \sup\{t \in \mathbb{R} \mid G(x_i \cdot t) \in S\}$. \square

Let $U\widetilde{M}$ be the unit tangent bundle of \widetilde{M} , thought of as the space of pairs (γ, y) consisting of an oriented geodesic $\gamma \subset \widetilde{M}$ together with a point $y \in \gamma$. This comes with a natural flow, the *geodesic flow* $\widetilde{\Theta}$, which takes a vector (γ, y) , after time $\tau \in \mathbb{R}$, to the vector (γ, y') such that y' lies at signed distance τ from y .

Our straightening map G has a natural lift

$$F : \widetilde{M} \rightarrow U\widetilde{M}$$

$$x \mapsto ((x \cdot \mathbb{R})^g, G(x)),$$

which takes each orbit of $\widetilde{\Phi}$ onto an orbit of $\widetilde{\Theta}$. In particular, if we fix $x \in \widetilde{M}$, then for each time $t \in \mathbb{R}$, there is a time $\tau \in \mathbb{R}$ such that $F(x \cdot t) = \widetilde{\Theta}_{\tau(t)}F(x)$. Since G moves each point a uniformly bounded distance, $\tau \rightarrow \pm\infty$ as $t \rightarrow \pm\infty$.

Each horosphere $S \subset \widetilde{M}$ has two natural lifts to the unit tangent bundle, a *stable horosphere* S^+ consisting of inward-pointing normal vectors and an *unstable horosphere* S^- consisting of outward-pointing normal vectors. If S is centered at $z \in S_\infty^2$, then

$$S^+ = \{(\gamma, y) \mid E^+(\gamma) = z, y \in S\},$$

$$\text{and } S^- = \{(\gamma, y) \mid E^-(\gamma) = z, y \in S\}.$$

Here, we are abusing notation and writing $E^\pm(\gamma)$ for the positive/negative endpoint of γ . We say that S^+ *points towards* z , and S^- *points away from* z . The geodesic flow is an Anosov flow whose strong stable and unstable leaves are exactly the stable and unstable horospheres [1].

For each $x \in \widetilde{M}$, the point $F(x)$ is contained in a unique stable horosphere $S^+(x)$, which points towards $E^+(x)$. Since F takes the flowline through x surjectively onto the corresponding lifted geodesic, every stable horosphere that points towards $E^+(x)$ is of the form $S^+(x \cdot t)$ for some t . Similar observations hold for unstable horospheres.

8.2. Coarse contraction for ω -sequences. We will use the map F , together with the Anosov behavior of $\widetilde{\Theta}$, to understand the way that ω -sequences act on positive sprigs.

For each $q \in \mathbf{P}$, we define

$$\widehat{\Omega}(q) := \{q' \in \mathbf{P} \mid \mathbf{e}^+(q') = \mathbf{e}^+(q) \text{ and } \mathbf{e}^-(q') = \mathbf{e}^-(q)\}.$$

Lemma 8.2. *If $q \in P$, then $\widehat{\Omega}(q)$ is a compact subset of P .*

Proof. Let $z = \mathbf{e}^+(q)$ and $w = \mathbf{e}^-(q)$, and note that $z \neq w$. Then

$$\begin{aligned} \widehat{\Omega}(q) &= (\mathbf{e}^+)^{-1}(z) \cap (\mathbf{e}^-)^{-1}(w) \\ &= (e^+)^{-1}(z) \cap (e^-)^{-1}(w) \end{aligned}$$

where the latter equality comes from the fact that \mathbf{e}^+ agrees with \mathbf{e}^- on S_u^1 . Thus $\widehat{\Omega}(q) \subset P$, and compactness follows from Lemma 3.3. \square

Proposition 8.3 (Coarse contraction). *Let (g_i) be an ω -sequence for $q \in \omega(p)$. Then*

$$\overline{\lim}_{i \rightarrow \infty} g_i(p') \subset \widehat{\Omega}(q)$$

for every $p' \in P$ with $e^+(p') = e^+(p)$.

Proof. It suffices to show that $e^\pm(g_i(p'))$ converges to $e^\pm(q)$. Equivalently, we will show that the geodesics $F(\langle g_i(p) \rangle)$ converge to $F(\langle q \rangle)$.

By the definition of an ω -sequence, we can find a sequence of points $x_0, x_1, \dots \in \langle p \rangle$ such that $\lim x_i = e^+(p)$ and $\lim g_i(x_i) = y \in \langle q \rangle$. For each i , let $S_i^+ := S^+(x_i)$ be the stable horosphere that contains $F(x_i)$. Since $e^+(p) = e^+(p')$, each S_i^+ also contains $F(x'_i)$ for some point $x'_i \in \langle p' \rangle$. See Figure 9.




FIGURE 9. Coarse contraction.

For each i , $F(x_i)$ and $F(x'_i)$ are obtained by flowing $F(x_0)$ and $F(x'_0)$ for some time τ_i under the geodesic flow, where $\tau_i \rightarrow \infty$. The geodesic flow exponentially contracts stable horospheres, so the distance between $F(x_i)$ and $F(x'_i)$ goes to 0 as $i \rightarrow \infty$. Each g_i acts as an isometry on $\mathbb{U}\widetilde{M}$, so we have $\lim F(g_i(x'_i)) = \lim F(g_i(x_i)) = F(y)$, and hence $\lim F(\langle g_i(p') \rangle) = F(\langle q \rangle)$ as desired. \square

For each $q \in \mathbf{P}$, define

$$\Omega(q) := \widehat{\Omega}(q) \cap \mathring{\mathcal{D}}^+(q),$$

which is a compact subset of $\mathring{\mathcal{D}}^+(q)$.

Proposition 8.4 (Coarse contraction for sprigs). *Let (g_i) be an ω -sequence for $q \in \omega(p)$. Then*

$$\overline{\lim}_{i \rightarrow \infty} g_i(p') \subset \Omega(q)$$

for every $p' \in \mathring{\mathcal{D}}^+(p)$

Proof. This follows from the preceding proposition, together with the fact that $\overline{\lim} \mathring{\mathcal{D}}^+(g_i(p)) \subset \mathring{\mathcal{D}}^+(q)$. \square

This is the final ingredient in our proof of the Homotopy Closing Lemma. The corresponding result for α -sequences is obtained by switching $+$ and $-$ in the proposition and in the definition of $\Omega(q)$.

9. CLOSED ORBITS FOR QUASIGEODESIC FLOWS

We turn to our main results.

Recurrent Links Lemma. *The linked region P_l is closed, nontrivial, $\pi_1(M)$ -invariant, and contains an ω -recurrent point.*

Proof. P_l is closed by Proposition 7.5 and nontrivial by Proposition 7.7. It is obviously $\pi_1(M)$ -invariant, so it corresponds to a closed, flow-invariant subset $M_l = \pi(\langle P_l \rangle) \subset M$. Since M is compact, M_l is compact, so it must contain some minimal set. A minimal set is the closure of an almost-periodic orbit (see, e.g., [2, Theorem 1.7]), which is *a fortiori* ω -recurrent. \square

This, together with the Homotopy Closing Lemma, implies the Closed Orbits Theorem.

9.1. Proof of the Homotopy Closing Lemma. Fix an ω -recurrent point $p \in P_l$, and a corresponding ω -sequence $(g_i)_{i=1}^\infty$. We must show that g_i eventually fixes some point in P . By Lemma 5.8, it suffices to show that g_i eventually fixes some nontrivial sprig.

Let $\mathbf{K} = \mathcal{D}^+(p)$ and $\mathbf{L} = \mathcal{D}^-(p)$ be the positive and negative sprigs through p . By Corollary 4.9, $\partial\mathbf{K} \cap \partial\mathbf{L} = \emptyset$. Since $\lim g_i(p) = p$, upper semicontinuity implies that $\lim g_i(\mathbf{K}) \subset \mathbf{K}$ and $\lim g_i(\mathbf{L}) \subset \mathbf{L}$.

If $\text{lk}(p) \geq 3$, then Lemma 7.6 implies that $g_i(\mathbf{K}) = \mathbf{K}$ for all i sufficiently large. This completes the proof for the (≥ 3) -linked case, so we will assume from now on that $\text{lk}(p) = 2$.

9.1.1. A single complementary region. After deleting any elements that fix \mathbf{K} , we can assume that each g_i takes \mathbf{K} into one of its complementary regions. Since $\text{lk}(p) = 2$, there are exactly two complementary regions that contain ends of \mathbf{L} .

Claim 9.1. *All but finitely many of the g_i take \mathbf{K} into one of these two complementary regions.*

Let W be a complementary region of \mathbf{K} , and suppose that there is an infinite subsequence $(h_i) \subset (g_i)$ such that $h_i(\mathbf{K}) \subset W$ for all i . Then $\partial h_i(\mathbf{K}) \subset W \cap S_u^1$ for all i , so the Linking Pigeonhole Principle says that $W \cap S_u^1$ also contains ends $l_i \in \partial h_i(\mathbf{L})$. Let $l \in \overline{W \cap S_u^1}$ be an accumulation point of the l_i . This is an end of \mathbf{L} , and $l \in W \cap S_u^1$ since the endpoints of this interval are ends of \mathbf{K} . The $g_i(\mathbf{K})$ can visit only finitely many complementary regions (Lemma 7.3), so this proves the claim.

Consequently, it suffices to prove the Homotopy Closing Lemma with the additional assumption that $g_i(\mathbf{K}) \subset W$ for all i , where W is one of the two complementary regions that contain ends of \mathbf{L} . The corresponding complementary interval is of the form $W \cap S_u^1 = (k_\alpha, k_\beta)$ for ends $k_\alpha, k_\beta \in \partial\mathbf{K}$. See Figure 10.

Let \mathbf{N}_α and \mathbf{N}_β be the negative sprigs through k_α and k_β . These are distinct by Lemma 5.6, so the corresponding separation interval $\sigma_{\alpha\beta} = \sigma(\mathbf{N}_\alpha, \mathbf{N}_\beta)$ is nontrivial and order-isomorphic to \mathbb{R} (Proposition 6.18). We will show that g_i eventually acts




FIGURE 10. Homotopy closing for 2-linked points.

as a contraction on some sub-interval of $\sigma_{\alpha\beta}$ which means that it fixes some sprig in this interval.

9.1.2. *Translation- and rotation-like elements.* To begin, we will divide the elements of our ω -sequence into two classes: We call g_i *translation-like* if $g_i(k_\alpha, k_\beta) \subset (k_\alpha, k_\beta)$, and *rotation-like* if $g_i(k_\beta, k_\alpha) \subset (k_\alpha, k_\beta)$. See Figure 11.




FIGURE 11. Translation- and rotation-like elements.

Claim 9.2. (1) Let $(h_i) \subset (g_i)$ be an infinite sequence of translation-like elements. Then

$$\lim h_i(k_\alpha) = k_\alpha \text{ and } \lim h_i(k_\beta) = k_\beta.$$

(2) Let $(h_i) \subset (g_i)$ be an infinite sequence of rotation-like elements. Then

$$\lim h_i(k_\alpha) = k_\beta \text{ and } \lim h_i(k_\beta) = k_\alpha.$$

We will prove (1); (2) follows from a similar argument.

The ends of $h_i(\mathbf{K})$ are contained in (k_α, k_β) , and accumulate on ends of \mathbf{K} , so it follows that both $h_i(k_\alpha)$ and $h_i(k_\beta)$ accumulate on nontrivial subsets of $\{k_\alpha, k_\beta\}$. Thus it suffices to show that $h_i(k_\alpha)$ does not accumulate on k_β , and $h_i(k_\beta)$ does not accumulate on k_α .

Suppose that $h_i(k_\alpha)$ accumulates on k_β . Since the h_i are translation-like, this implies that the intervals $h_i(k_\alpha, k_\beta)$ accumulate on k_β . But this is impossible, since each $h_i(k_\alpha, k_\beta)$ contains an end of $h_i(\mathbf{L})$, which cannot accumulate on $k_\beta \in \partial\mathbf{K}$. A similar argument shows that $h_i(k_\beta)$ cannot accumulate on k_α , completing the claim.

9.1.3. *Eventual return.* Let us show that each sprig in $\sigma_{\alpha\beta}$ eventually returns to $\sigma_{\alpha\beta}$.

Claim 9.3. *For each $\mathbf{N} \in \sigma_{\alpha\beta}$, we eventually have $g_i(\mathbf{N}) \in \sigma_{\alpha\beta}$.*

First, we note that it suffices to prove this in the special case where \mathbf{N} intersects the bounded part $\mathring{\mathbf{K}}$ of \mathbf{K} . Indeed, in the general case, we can use Lemma 7.8 to find $\mathbf{N}', \mathbf{N}'' \in \sigma_{\alpha\beta}$ that intersect $\mathring{\mathbf{K}}$, with $\mathbf{N}' \prec \mathbf{N} \prec \mathbf{N}''$. Then $g_i(\mathbf{N}'), g_i(\mathbf{N}'') \in \sigma_{\alpha\beta} \Rightarrow g_i(\mathbf{N}) \in \sigma_{\alpha\beta}$ by Lemma 6.19(2).

Let \mathbf{N} be an element of $\sigma_{\alpha\beta}$ that intersects $\mathring{\mathbf{K}}$ nontrivially, and suppose for contradiction that there is an infinite subsequence $(h_i) \subset (g_i)$ such that $h_i(\mathbf{N}) \notin \sigma_{\alpha\beta}$ for all i . Choose a point $q \in \mathring{\mathbf{N}} \cap \mathring{\mathbf{K}}$. Then by Coarse Contraction (Proposition 8.4), we can pass to a subsequence for which $h_i(q)$ converges to some point $q_\infty \in \mathring{\mathbf{K}}$. Then $\overline{\lim} h_i(\mathbf{N}) \subset \mathcal{D}^-(q_\infty)$, and $\mathcal{D}^+(q_\infty) = \mathbf{K}$. Since $q_\infty \in P$, this means means that $\overline{\lim} \partial h_i(\mathbf{N})$ is disjoint from $\partial\mathbf{K}$.

After taking a further subsequence, we can assume that the h_i are all translation-like or all rotation-like. Suppose that the h_i are all translation-like. Since $\mathbf{N} \in \sigma_{\alpha\beta}$, it separates k_α from k_β , so it must have ends $n \in (k_\alpha, k_\beta)$ and $n' \in (k_\beta, k_\alpha)$. Since h_i is translation-like, $h_i(n) \in (h_i(k_\alpha), h_i(k_\beta)) \subset (k_\alpha, k_\beta)$. We must also have $h_i(n') \in (k_\alpha, k_\beta)$, since otherwise $h_i(\mathbf{N})$ would separate k_α from k_β . In particular, $h_i(n')$ must be contained in either $(k_\alpha, h_i(k_\alpha))$ or $(h_i(k_\beta), k_\beta)$. Then $h_i(k_\alpha) \rightarrow k_\alpha$ and $h_i(k_\beta) \rightarrow k_\beta$ (Claim 9.2) implies that $h_i(n')$ accumulates on either k_α or k_β , which contradicts the fact that $\overline{\lim} \partial h_i(\mathbf{N})$ is disjoint from $\partial\mathbf{K}$.

A similar argument handles the case where the h_i are rotation-like, completing the claim.

In addition, each g_i preserves or reverses the order on $\sigma_{\alpha\beta}$ depending on whether it is translation- or rotation-like:

Claim 9.4. *Let $\mathbf{N}, \mathbf{N}' \in \sigma_{\alpha\beta}$, with $\mathbf{N} \prec \mathbf{N}'$, and let j large enough so that $g_j(\mathbf{N}), g_j(\mathbf{N}') \in \sigma_{\alpha\beta}$. Then $g_j(\mathbf{N}) \prec g_j(\mathbf{N}')$ when g_j is translation-like, and $g_j(\mathbf{N}') \prec g_j(\mathbf{N})$ when g_j is rotation-like.*

This follows easily from Lemma 6.15, using the fact that $I(\mathbf{N}_\alpha, \mathbf{N}_\beta) \subset (k_\alpha, k_\beta)$ and $I(\mathbf{N}_\beta, \mathbf{N}_\alpha) \subset (k_\beta, k_\alpha)$.

9.1.4. *Eventual contraction.* To complete the proof, we will see that the elements of $\sigma_{\alpha\beta}$ not only return to $\sigma_{\alpha\beta}$, but are pulled closer together.

Let $\Omega(p) \subset \mathring{\mathbf{K}}$ be as in §8.2, and let $\Omega := \Omega(p) \cap U(\mathbf{N}_\alpha, \mathbf{N}_\beta)$. For each $p' \in \mathring{\mathbf{K}}$, Proposition 8.4 says that

$$\overline{\lim} g_i(p') \subset \Omega(p).$$

In addition, if the negative sprig through p' is contained in $\sigma_{\alpha\beta}$, then p' is eventually contained in $U(\mathbf{N}_\alpha, \mathbf{N}_\beta)$ by Claim 9.3, so we have

$$\overline{\lim} g_i(p') \subset \Omega.$$

Since Ω is a compact subset of the open set $U(\mathbf{N}_\alpha, \mathbf{N}_\beta)$, we can find negative sprigs $\mathbf{N}_a \prec \mathbf{N}_b$ in $\sigma_{\alpha\beta}$ such that $\Omega \subset U(\mathbf{N}_a, \mathbf{N}_b)$. By Lemma 7.8, we can assume that \mathbf{N}_a and \mathbf{N}_b intersect $\mathring{\mathbf{K}}$, and choose points $p_a \in \mathring{\mathbf{N}}_a \cap \mathring{\mathbf{K}}$ and $p_b \in \mathring{\mathbf{N}}_b \cap \mathring{\mathbf{K}}$. See Figure 12.

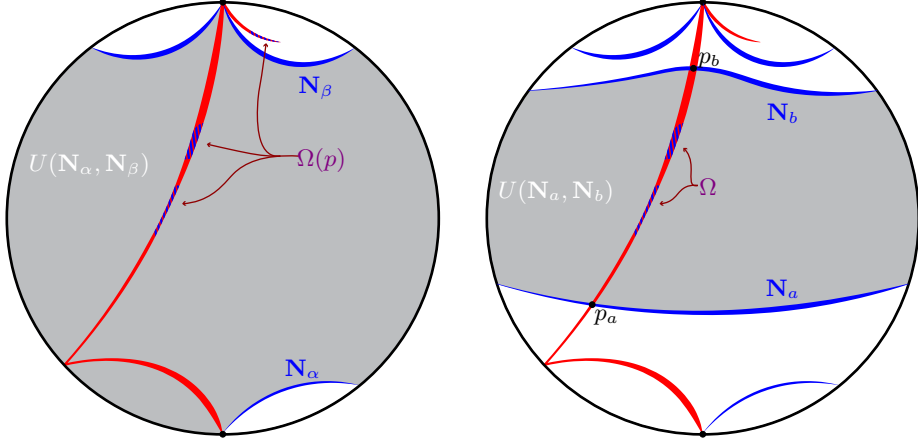


FIGURE 12. Eventual contraction.

By taking i sufficiently large, we can assume $g_i(\mathbf{N}_a)$ and $g_i(\mathbf{N}_b)$ are contained in $\sigma_{\alpha\beta}$. Furthermore, since $\overline{\lim} g_i(p_a)$ and $\overline{\lim} g_i(p_b)$ are contained in Ω , which is a compact subset of the open set $U(\mathbf{N}_a, \mathbf{N}_b)$, we can assume that $g_i(p_a)$ and $g_i(p_b)$ are contained in $U(\mathbf{N}_a, \mathbf{N}_b)$. Then $g_i(\mathbf{N}_a)$ and $g_i(\mathbf{N}_b)$ are contained in $\sigma(\mathbf{N}_a, \mathbf{N}_b) \subset \sigma_{\alpha\beta}$, which means that

$$\mathbf{N}_a \prec g_i(\mathbf{N}_a) \prec g_i(\mathbf{N}_b) \prec \mathbf{N}_b$$

when g_i is translation-like, and

$$\mathbf{N}_a \prec g_i(\mathbf{N}_b) \prec g_i(\mathbf{N}_a) \prec \mathbf{N}_b$$

when it is rotation-like. That is, g_i eventually acts as an orientation-preserving or -reversing contraction on $\sigma(\mathbf{N}_a, \mathbf{N}_a)$, and hence fixes some sprig in this interval. Such a sprig is nontrivial because it has at least two ends (Lemma 4.6). This completes the proof of the Homotopy Closing Lemma.

9.2. Closed orbits and the universal circle. Finally, we will show that certain closed orbits can be seen purely in terms of the universal circle.

An orientation-preserving group action $\Gamma \curvearrowright S^1$ is *pA-like* if each nontrivial element has a power that acts with an even number of fixed points, alternately attracting and repelling.

Theorem 9.5. *The action of $\pi_1(M) \curvearrowright S_u^1$ is pA-like.*

Proof. Fix a nontrivial element $g \in \pi_1(M)$, and let $\alpha, \rho \in S_\infty^2$ be the attracting and repelling fixed points for $g \curvearrowright S_\infty^2$. Let A and R be the master sets rooted at these points, and note that $\partial A = \mathbf{e}^{-1}(\alpha)$ and $\partial R = \mathbf{e}^{-1}(\rho)$ are disjoint, closed, totally disconnected, and g -invariant.

Let $F \subset S_u^1$ be the fixed point set for $g \curvearrowright S_u^1$. Then after replacing g by some power, we can assume that $F \neq \emptyset$. Indeed, since ∂A and ∂R are closed and disjoint, they have a finite linking number $\text{lk}(\partial A, \partial R)$ by Lemma 7.1. Choose a complementary interval (a, r) of $\partial A \cup \partial R$ with $a \in \partial A$ and $r \in \partial R$. By the proof of this lemma, there are finitely many such intervals, which we called $\partial A \partial R$ -interstitial intervals, so $g^n(a, r) = (a, r)$ for some n . Then $a, r \in F$ as desired.

Now that $F \neq \emptyset$, note that $F \subset \partial A \cup \partial R$. Indeed, if $s \in S_u^1 \setminus (\partial A \cup \partial R)$, then $\mathbf{e}(s) \notin \{\alpha, \rho\}$, so $g(\mathbf{e}(s)) \neq \mathbf{e}(s)$ and hence $g(s) \neq s$.

Let I be a complementary interval of F . Then g acts as a translation on I , fixing its endpoints so we can write it as either $I = (a_I, r_I)$ or $I = (r_I, a_I)$, where a_I and r_I are attracting and repelling with respect to points in I . Then we must have $a_I \in \partial A$ and $r_I \in \partial R$. Indeed, take $s \in I \setminus (\partial A \cup \partial R)$. Then $a_I = \lim_{i \rightarrow \infty} g^i(s)$, which is contained in ∂A because $\lim_{i \rightarrow \infty} g_i(\mathbf{e}^+(s)) = \alpha$. Similarly, $r_I = \lim_{i \rightarrow -\infty} g^i(s)$, which is contained in ∂R because $\lim_{i \rightarrow -\infty} g_i(\mathbf{e}(s)) = \rho$.

Let $a \in \partial A$. If $g(a) \neq a$, then a is contained in some complementary interval J of F , and $\lim_{i \rightarrow -\infty} g^i(a) = r_J$. But ∂A is g -invariant, so this means that $r_J \in \partial A$, a contradiction. Hence $\partial A \subset F$. Similarly, $\partial R \subset F$, so $\partial A \cup \partial R \subset F$. We already showed the opposite inclusion, so we have $F = \partial A \cup \partial R$.

Since each complementary interval of F has its attracting and repelling endpoints in ∂A and ∂R , it suffices to show that there are only finitely many such complementary intervals. But these are $\partial A \partial R$ - or $\partial R \partial A$ -interstitial interval, of which there are only finitely many because $\text{lk}(\partial A, \partial R)$ is finite. \square

The following result characterizes the elements of $\pi_1(M)$ that do not represent multiples of closed orbits. Conversely, any nontrivial element that does *not* act as in the proposition represents a multiple of the free homotopy class of some closed orbit.

Proposition 9.6. *Suppose that $g \in \pi_1(M)$ does not fix a point in P . Then it acts on S_u^1 with exactly two fixed points in an attracting-repelling pair.*

Proof. Let $F \subset S_u^1$ be the fixed point set for $g \curvearrowright S_u^1$. Then $F \neq \emptyset$. Indeed, otherwise g would fix a point in P by Brouwer's fixed point theorem.

In the previous proof, we showed that $F = \partial A \cup \partial R$, where A and R are the master sets rooted at α and ρ , and the points in ∂A and ∂R are attracting and repelling. To complete the lemma, it suffices to show that both ∂A and ∂R consists of a single point. If ∂A contained more than one point, then A would be nontrivial, and if ∂R contained more than one point then R would be nontrivial. Either way, g would fix a point in P by Lemma 5.8, contradicting our hypothesis. \square

10. QUESTIONS

Our proof of the Homotopy Closing Lemma holds just as well for pseudo-Anosov flows, even ones that are not quasigeodesic. In fact, it can be applied to a larger class of *coarsely hyperbolic* flows, defined by the existence of a transverse pair of spidery decompositions that are coarsely contracted and expanded.

It's easy to construct coarsely hyperbolic flows that are neither quasigeodesic nor pseudo-Anosov. For example, one can start with a pseudo-Anosov flow that is not quasigeodesic and blow up some orbit to a solid cylinder or torus. These examples are quite trivial since the Homotopy Closing Lemma follows easily from the Anosov Closing Lemma for the original flow. It would be interesting to construct a less trivial class of examples.

10.1. Product-covered flows. A flow on a closed 3-manifold is *product covered* if the lifted flow on the universal cover is conjugate to the vertical flow on \mathbb{R}^3 . By Theorem 3.1, this is equivalent to having a planar flowspace.

Question. *Let Φ be a product-covered flow on a closed hyperbolic 3-manifold M . Is Φ coarsely hyperbolic?*

More generally, we propose the following.

Conjecture. *Every product-covered flow on a closed hyperbolic 3-manifold contains a closed orbit.*

This would be implied by the following conjecture.

Conjecture. *The fundamental group of a closed hyperbolic 3-manifold cannot act freely and cocompactly on the plane.*

Here, a cocompact action is one whose quotient space is compact, but not necessarily Hausdorff.

10.2. Möbius-like groups. In [10] we proposed a very different method for proving the Closed Orbits Theorem.

An action of a group Γ on a circle S^1 is called *Möbius-like* if each $g \in \Gamma$ is conjugate to a Möbius transformation. It is called *hyperbolic Möbius-like* if each $g \in \Gamma$ is conjugate to a hyperbolic Möbius transformation. A Möbius-like or hyperbolic Möbius-like action is called *Möbius* or *hyperbolic Möbius*, respectively, if it is conjugate to an action by Möbius transformations.

The fundamental group of a closed hyperbolic 3-manifold can never act as a hyperbolic Möbius group (see [10]). The only known examples of Möbius-like actions that are not Möbius are found in [17]. We propose the following conjecture.

Conjecture. *The fundamental group of a closed hyperbolic 3-manifold cannot act as a hyperbolic Möbius-like group.*

By Proposition 9.6, this would imply the Closed Orbits Theorem.

REFERENCES

- [1] D. V. Anosov, *Ergodic Properties of Geodesic Flows on Closed Riemannian Manifolds of Negative Curvature*, Dokl. Akad. Nauk SSSR 151, 1250–1252, 1963.
- [2] J. Auslander, *Minimal Flows and their Extensions*, Mathematical Studies **153**, Notas de Matematica, North-Holland Pub. Co., Amsterdam, 1988.
- [3] M. R. Bridson and A. Haefliger, *Metric Spaces of Non-Positive Curvature*, Grundlehren der mathematischen Wissenschaften, vol. 319, Springer-Verlag, Berlin, Germany, 1999.
- [4] D. Calegari, *Universal circles for quasigeodesic flows*, Geom. Topol. **10** (2006), 2271–2298.
- [5] D. Calegari and N. Dunfield, *Laminations and groups of homeomorphisms of the circle*, Invent. Math. **152** (2003), 149–204.
- [6] J.W. Cannon and W.P. Thurston, *Group invariant Peano curves*, Geom. Topol. **11** (2007), 1315–355.

- [7] R. J. M. Dawson, *Paradoxical Connections*, The American Mathematical Monthly **96** (1989), no. 1, 31–33.
- [8] S. Fenley *Ideal boundaries of pseudo-Anosov flows and uniform convergence groups, with connections and applications to large scale geometry*, Geom. Topol. **16** (2012), 1–10.
- [9] S. Fenley and L. Mosher, *Quasigeodesic flows in hyperbolic 3-manifolds*, Topol. **40** (2001), no. 3, 503–537.
- [10] S. Frankel, *Quasigeodesic flows and Möbius-like groups*, J. Diff. Geom. **93** (2013), no. 3, 401–429.
- [11] S. Frankel, *Quasigeodesic flows and sphere-filling curves*, Geom. Topol. **19** (2015), no. 3, 1249–1262.
- [12] J. Franks, *A new proof of the Brouwer plane translation theorem*, Ergod. Th. & Dynam. Sys. **12** (1992), 217–226.
- [13] D. Gabai, *Foliations and the topology of 3-manifolds*, J. Diff. Geom. **18** (1983), 445–503.
- [14] M. Gromov *Hyperbolic groups*, Essays in group theory, Math. Sci. Res. Inst. Publ., vol. 8, Springer, New York, 1987, 75–263.
- [15] J.G. Hocking and G.S. Young, *Topology*, Addison-Wesley Series in Mathematics, Addison-Wesley Publishing Company, Reading, MA, 1961.
- [16] M. Kapovich, *Hyperbolic manifolds and discrete groups*, Progress in Mathematics, vol. 183 (2001), Birkhauser Boston Inc., Boston, MA.
- [17] N. Kovacevic, *Examples of Möbius-like groups which are not Möbius groups*, Trans. Amer. Math. Soc. **351** (1999), no. 12.
- [18] G. Kuperberg, *A volume-preserving counterexample to the Seifert conjecture*, Comm. Math. Helv. **71** (1996), no. 1, 70–79.
- [19] K. Kuperberg, *A smooth counterexample to the Seifert conjecture*, Ann. of Math. (2) **140** (1994), no. 3, 723–732.
- [20] K. Kuratowski, *Topology, Volume I*, Academic Press, New York, NY, 1966.
- [21] K. Kuratowski, *Topology, Volume II*, Academic Press, New York, NY, 1968.
- [22] B. S. Mangum, *Incompressible surfaces and pseudo-Anosov flows*, Topol. Appl. **87** (1998), 29–51.
- [23] D. Montgomery and L. Zippin, *Translation groups of three-space*, Amer. J. Math. **59** (1937), no. 1, 121–128.
- [24] R. L. Moore, *Concerning upper semi-continuous collections of continua*, Trans. Amer. Math. Soc. **27** (1925), no. 4, 416–428.
- [25] A. Rechtman, *Existence of periodic orbits for geodesible vector fields on closed 3-manifolds*, Ergod. Th. & Dynam. Sys. **30** (2010), no. 6, 1817–1841.
- [26] P. A. Schweitzer, *Counterexamples to the Seifert conjecture and opening closed leaves of foliations*, Ann. of Math. **100** (1974), no.2, 386–400.
- [27] H. Seifert, *Closed integral curves in 3-space and isotopic two-dimensional deformations*, Proc. Amer. Math. Soc. **1**, (1950), 287–302.
- [28] C. H. Taubes, *The Seiberg-Witten equations and the Weinstein conjecture*, Geom. Top. **11** (2007), 2117–2202.
- [29] W. P. Thurston, *Geometry and topology of 3-manifolds*, a.k.a. “Thurston’s notes,” (1979).
- [30] R.L. Wilder, *Topology of manifolds*, American Mathematical Society Colloquium Publications, vol. 32, American Mathematical Society, Providence, RI, 1979.
- [31] A. Zeghib, *Sur les feuilletages géodésiques continus des variétés hyperboliques*, Invent. Math. **114** (1993), no. 1, 193–206.

DEPARTMENT OF MATHEMATICS, WASHINGTON UNIVERSITY IN ST LOUIS, ST LOUIS, MO
E-mail address: `steven.frankel@wustl.edu`