Quasigeodesic flows from infinity

Steven Frankel
Magdalene College
University of Cambridge

Supervised by Danny Calegari

This dissertation is submitted for the degree of Doctor of Philosophy.

May 2013

This dissertation is the result of my own work and includes nothing which is the outcome of work done in collaboration except where specifically indicated in the text.

Abstract

A flow is called *quasigeodesic* if each flowline is uniformly efficient at measuring distances on the large scale. We study quasigeodesic flows on closed hyperbolic 3-manifolds from the perspective at infinity.

If M is a closed hyperbolic 3-manifold with a quasigeodesic flow, the orbit space P of the lifted flow on $\mathbb{H}^3$ is planar, and it inherits an action by $\pi_1(M)$. We use Calegari's universal circle S^1_u to compactify P to a closed disc $\overline{P}$ in a way that is compatible with this action. We show that the maps $e^{\pm} : P \to \partial\mathbb{H}^3$ that send each orbit to its forward/backward limit extend continuously to $\overline{P}$, and $e^+ = e^-$ on $\partial\overline{P}$. Consequently, the restriction $e : \partial\overline{P} \to \partial\mathbb{H}^3$ is a π_1-equivariant sphere-filling curve, generalizing the Cannon-Thurston theorem which produces such curves for suspension flows. If the flow has no closed orbits then we show that $\pi_1(M)$ acts on $\partial\overline{P}$ as a hyperbolic Möbius-like group that is not Möbius, which conjecturally cannot occur. Finally, we show that $\pi_1(M)$ acts on $\partial\overline{P}$ with pseudo-Anosov dynamics.

This completes a major part of Calegari's program to show that quasigeodesic flows are homotopic (through quasigeodesic flows) to pseudo-Anosov flows.

Acknowledgements

Foremost, I would like to thank Danny Calegari. While your mathematical influence on this thesis is obvious, I'm most grateful for your encouragement and advice over the last five years. And Therese, thanks for the hot tea and long conversations that kept me going in the winter. The two of you have made me feel like part of the family. A nephew or cousin, I mean; you're far too young to be parents.

Thanks to Sylvain Cappell. You're the reason I wanted to study mathematics.

Thanks to Matt Day, Sergio Fenley, Dave Gabai, Nik Makarov, Curt McMullen, and Dongping Zhuang for helpful conversations and correspondence. Thanks Benson Farb for your advice while in Chicago.

Thanks to Jake Rasmussen and Henry Wilton for their careful examination of this dissertation, especially on such short notice.

Thanks to my climbing buddies for making me work on beaches and mountains instead of my office. In particular, to J and Matt, for holding the rope on Pinnacle Gully while I worked out the proof of the continuous extension theorem.

Thanks to all of my friends at Caltech, Cambridge, and Chicago. Rafa, Maurice, and Bapt, can we go for a pitcher at Amigos? Himanshu, sure, I'll come to that party I wasn't invited to. Dave and Jeannie, you always made LA interesting. Helge, we should watch more British spy movies together. Giulio, Goldie, Jack, Job, John, Steven, Zeb, and the rest of the CUMC, I'd be more specific if I could remember more of the time we spent together. Brad, that's how I always am, outside of Chicago winter.

Finally, thanks to my family for your love and support, even though I'm going to be the wrong kind of doctor. Please send money.

Contents

CHAPTER 1

Introduction

The goal of this thesis is to outline the theory of quasigeodesic flows from the point of view "at infinity." The guiding conjecture for this, due to Danny Calegari, is that every quasigeodesic flow on a closed hyperbolic 3-manifold is homotopic (through quasigeodesic flows) to a pseudo-Anosov flow. Both quasigeodesic and pseudo-Anosov flows give rise to *universal circles* at infinity, and our program attempts to relate quasigeodesic and pseudo-Anosov flows using their dynamics on the universal circle.

1. Overview

Much of this thesis is about how certain 3-dimensional smooth dynamical objects give rise to an array of intricately related 1- and 2-dimensional discrete dynamical objects. The 3-dimensional objects we're concerned with are quasigeodesic and pseudo-Anosov flows, and these give rise to discrete group actions on 1-dimensional universal circles and 2-dimensional orbit spaces that preserve certain laminations, singular foliations, and unbounded decompositions. Some of these objects are built out of others by topological trickery, some come from quotient maps, and at the end they all fit together nicely in a group-equivariant way. We'll start informally in order to get to this picture before plodding through all of the technical ingredients. To preserve continuity we'll refrain from citing sources here, but precise citations will accompany the precise statements in the main chapters.

Let M be a closed hyperbolic 3-manifold. Its universal cover is isometric to $\mathbb{H}^3$, and $\pi_1(M)$ acts on $\mathbb{H}^3$ by loxodromic isometries with M as its quotient. Hyperbolic space has a natural compactification $\overline{\mathbb{H}}^3 = \mathbb{H}^3 \cup S^2_\infty$ by adding a sphere at infinity, and the action of $\pi_1(M)$ extends continuously to the boundary.

A flow on M lifts to a flow on $\mathbb{H}^3$, and the action of $\pi_1(M)$ by deck transformations preserves the foliation by flowlines. In general, flowlines won't behave well with respect to the boundary of $\overline{\mathbb{H}}^3$, and might spiral close to infinitely many points in S^2_∞. An obvious way to avoid this would be to insist that each flowline is a geodesic, but a theorem of Zeghib implies that no such flows exist. It turns out that quasigeodesic flows, where each flowline tracks a geodesic in the large scale, are exactly the flows whose flowlines have well-defined and continuously varying endpoints in S^2_∞.

A flow $\mathfrak{F}$ on a 3-manifold is called Anosov if the tangent bundle splits into three perpendicular sub-bundles: the bundle $T\mathfrak{F}$ tangent to the flow, a stable bundle E^s, and an unstable bundle E^u. Nearby orbits along the stable direction converge exponentially in forwards time, and nearby orbits along the unstable direction converge exponentially in backwards time. The (weak) stable and unstable bundles integrate to a pair of transverse foliations (Λ^s, Λ^u), where $T\Lambda^s$ spans the tangent and stable directions, and $T\Lambda^u$ spans the tangent and unstable directions. Leaves of Λ^s and Λ^u intersect transversely, and the intersection of two leaves is either empty or an orbit of $\mathfrak{F}$. The orbits in a single leaf of Λ^s are all forwards asymptotic, and the orbits in a single leaf of Λ^u are all backwards asymptotic.

A flow is called pseudo-Anosov if it is Anosov except near some isolated singular orbits, where the flow looks like the suspension of a p-prong in the plane (see Figure 1). A pseudo-Anosov flow has a pair of *singular* foliations (Λ^s, Λ^u). These are true transverse foliations away from the singular orbits, and they look like Figure 2 near singular orbits.

1.1. Suspension flows. The prototype example for both quasigeodesic and pseudo-Anosov flows is the suspension of a pseudo-Anosov homeomorphism. Let Σ be a surface of genus at least two, and consider $M' = \Sigma \times I$ with the unit flow $\mathfrak{F}'$ in the positive direction along I. The manifold M built by gluing $\Sigma \times \{0\}$ to $\Sigma \times \{1\}$ with a homeomorphism f is hyperbolic exactly when the mapping class of f is aperiodic and irreducible. In this case, f is isotopic to a pseudo-Anosov homeomorphism, so it preserves a pair (λ^s, λ^u) of 1-dimensional measured singular foliations on Σ. The monodromy f is

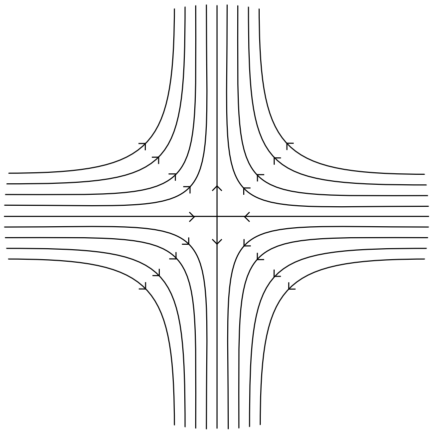

FIGURE 1. A 4-prong singularity in the plane.

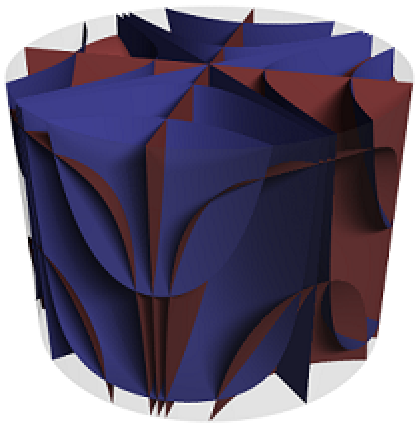

FIGURE 2. Stable and unstable foliations near a singular orbit.

contracting along leaves of λ^s and expanding along leaves of λ^u. The flow $\mathfrak{F}'$ on M' glues up to form a flow $\mathfrak{F}$ on M called the *suspension flow* of f, and $\mathfrak{F}$ is both pseudo-Anosov and quasigeodesic. The 2-dimensional singular foliations of $\mathfrak{F}$ on M restrict to the 1-dimensional singular foliations preserved by f on Σ. That is,

$$\Lambda^s \cap (\Sigma \times \{0\}) = \lambda^s$$

and

$$\Lambda^u \cap (\Sigma \times \{0\}) = \lambda^u.$$

The flow $\mathfrak{F}$ lifts to a pseudo-Anosov flow $\widetilde{\mathfrak{F}}$ on $\mathbb{H}^3$ whose stable and unstable singular foliations are the lifts $(\widetilde{\Lambda}^s, \widetilde{\Lambda}^u)$. A copy of the surface, say

$\Sigma \times \{\frac{1}{2}\}$, lifts to a plane $\widetilde{\Sigma}$ that is transverse to the flow, and $(\widetilde{\Lambda}^s, \widetilde{\Lambda}^u)$ restrict to the lifts $(\widetilde{\lambda}^s, \widetilde{\lambda}^u)$ on $\widetilde{\Sigma}$. See Figure 3.

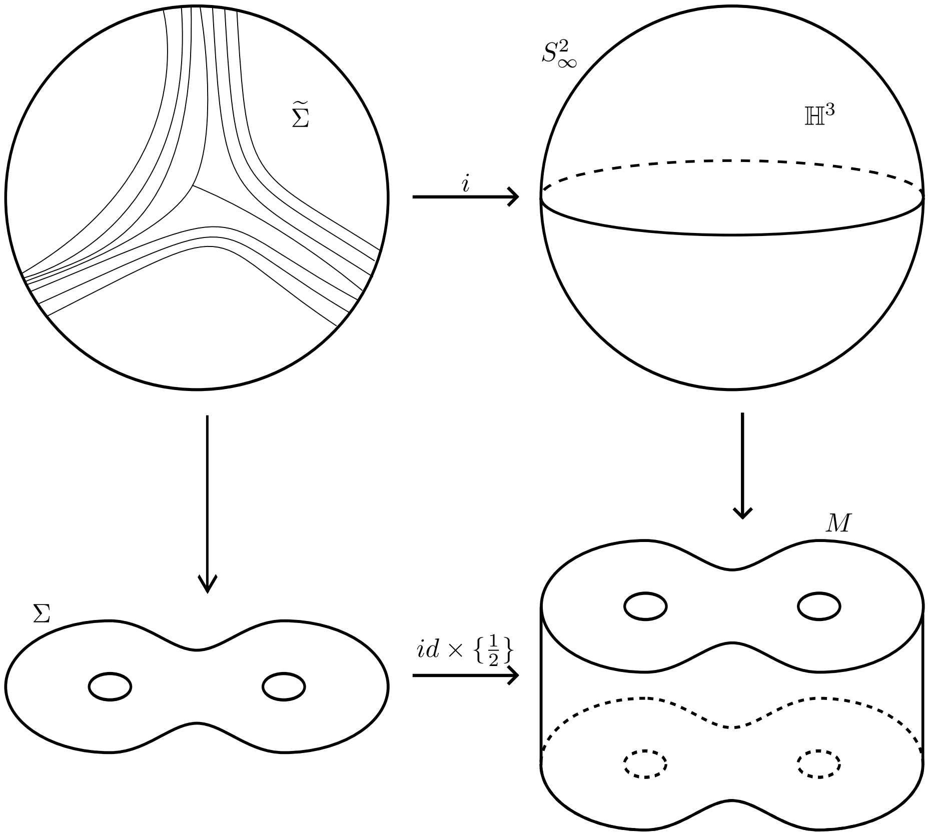

FIGURE 3. A surface bundle and lifted surface.

The surface $\widetilde{\Sigma}$ is transverse to the flow, and it hits every flowline, so there's a natural quotient map

$$\pi : \mathbb{H}^3 \to \widetilde{\Sigma}$$

that sends each flowline to its point of intersection with $\widetilde{\Sigma}$. That is, $\widetilde{\Sigma}$ is naturally homeomorphic to the *orbit space* of the flow $\widetilde{\mathfrak{F}}$. Quasigeodesity of $\mathfrak{F}$ implies that the maps

$$e^{\pm} : \widetilde{\Sigma} \to S^2_\infty$$

that send each orbit to its forward/backward limit are well-defined and continuous. The fact that orbits in a stable leaf are forwards asymptotic means that $e^+(p) = e^+(q)$ if p and q are both contained in a leaf of $\widetilde{\lambda}^s$. Similarly, $e^-(p) = e^-(q)$ if p and q are both contained in a leaf of $\widetilde{\lambda}^u$.

The fundamental group $\pi_1(M)$ acts on $\mathbb{H}^3$ preserving the 1-dimensional foliation by flowlines, as well as the 2-dimensional foliations $\widetilde{\Lambda}^s$ and $\widetilde{\Lambda}^u$. Pushing forward under the map π, we obtain an action of $\pi_1(M)$ on $\widetilde{\Sigma}$ that preserves $\widetilde{\lambda}^s$ and $\widetilde{\lambda}^u$. The $e^{\pm}$ maps are equivariant with respect to this action.

The surprising behavior of suspension flows becomes apparent when we compactify $\widetilde{\Sigma}$. The surface $\widetilde{\Sigma}$ is homeomorphic to the hyperbolic plane $\mathbb{H}^2$, since it's the universal cover of the hyperbolic surface Σ. However, its geometry as a subspace of $\mathbb{H}^3$ is very different from the geometry it gets as a lift of a hyperbolic surface. The hyperbolic plane has a natural compactification to a closed disc $\overline{\mathbb{H}}^2 = \mathbb{H}^2 \cup S^1_\infty$ by adding a circle at infinity, so we can use the homeomorphism $\widetilde{\Sigma} \simeq \mathbb{H}^2$ to compactify our transversal $\widetilde{\Sigma}$. It turns out that this can be done in a way that respects the action of $\pi_1(M)$. The famous Cannon-Thurston theorem states that the inclusion map $i : \widetilde{\Sigma} \to \mathbb{H}^3$ extends to the boundary circle, where it restricts to a π_1-equivariant curve

$$i : S^1_\infty \to S^2_\infty$$

that fills the sphere. In fact, the $e^{\pm}$ maps also extend continuously to the boundary, and $e^+|_{S^1_\infty} = e^-|_{S^1_\infty} = i|_{S^1_\infty}$.

1.2. Pseudo-Anosov and quasigeodesic flows in general. Much of this structure remains in place for a general pseudo-Anosov flow. Fix a closed hyperbolic 3-manifold M and let $\mathfrak{F}$ be a pseudo-Anosov flow on M. The *flowspace* P is the set of flowlines of the lifted flow on $\mathbb{H}^3$. We can endow P with a topology by thinking of it as the quotient of $\mathbb{H}^3$ obtained by collapsing each flowline to a point. It turns out that P is homeomorphic to the plane. If we take a section i of the quotient map $\pi : \mathbb{H}^3 \to P$, we can think of P as a transversal to the flow $\widetilde{\mathfrak{F}}$. That is, P takes the place of $\widetilde{\Sigma}$. As before, the 2-dimensional singular foliations $(\widetilde{\Lambda}^s, \widetilde{\Lambda}^u)$ restrict to 1-dimensional singular foliations $(\widetilde{\lambda}^s, \widetilde{\lambda}^u)$ on P, which are preserved by the induced action of $\pi_1(M)$. However, note that these are not generally lifts of singular foliations on a compact surface, since the quotient of P is not generally a surface. We lose the $e^{\pm}$ maps since $\mathfrak{F}$ is no longer quasigeodesic in general.

Although P no longer has a natural homeomorphism to $\mathbb{H}^2$, it turns out that the pair of foliations $(\widetilde{\lambda}^s, \widetilde{\lambda}^u)$ can be used to build a compactification $\overline{P} = P \cup S^1_u$ by adding a *universal circle* S^1_u at infinity. This universal circle is built by introducing a point at infinity for each end of each leaf of $\widetilde{\lambda}^s$ and $\widetilde{\lambda}^u$ – the nonsingular leaves contribute two points and p-pronged leaves contribute p points. The fact that $\widetilde{\lambda}^s$ and $\widetilde{\lambda}^u$ are π_1-equivariant implies that the action of $\pi_1(M)$ on P extends to S^1_u. However, the section $i : P \to \mathbb{H}^3$ does not generally extend to the boundary of $\overline{P}$.

The action of $\pi_1(M)$ on S^1_u takes a particular dynamical form. Each $g \in \pi_1(M)$ acts on S^1_u with an even number of fixed points that are alternately attracting and repelling.

Now let $\mathfrak{F}$ be a quasigeodesic flow on M. The flowspace P is defined similarly. Once again it turns out to be a plane, and once again it can be interpreted as a transversal to $\widetilde{\mathfrak{F}}$ by choosing a section i of the quotient map $\pi : \mathbb{H}^3 \to P$. We lose the pair of singular foliations $(\widetilde{\Lambda}^s, \widetilde{\Lambda}^u)$, but this time we get to keep the maps $e^{\pm} : P \to S^2_\infty$. And, inspired by the case of a suspension flow, we will use e^+ and e^- to build replacements for $\widetilde{\lambda}^s$ and $\widetilde{\lambda}^u$. Fix a point $p \in S^2_\infty$ and consider $(e^+)^{-1}(p)$, the set of flowlines whose positive endpoints lie at p. Each component of $(e^+)^{-1}(p)$ is closed and unbounded, and we can collate these to form the *positive decomposition* $\mathcal{D}^+$ of P. The *negative decomposition* $\mathcal{D}^-$ is built similarly using the e^- map. The decompositions $\mathcal{D}^{\pm}$ are invariant under the action of $\pi_1(M)$, and we will treat $(\mathcal{D}^+, \mathcal{D}^-)$ as a generalization of $(\widetilde{\lambda}^s, \widetilde{\lambda}^u)$. Where in the pseudo-Anosov case a stable leaf and an unstable leaf intersect transversely in a single orbit, here a positive decomposition element and a negative decomposition element intersect in a compact collection of orbits, all of which share a pair of endpoints in S^2_∞.

Remarkably, there's still enough structure to build a compactification for P. The decomposition elements are arbitrarily complicated closed unbounded sets instead of just lines and p-prongs, so we need to be more sophisticated when talking about ends. Still, there is a workable notion of an end, and the ends of each decomposition element appear as points in a universal circle S^1_v. Here's where the new part of the story begins. We'll show that this universal circle compactifies the orbit space, and the maps i, e^+, and e^- all extend continuously to (and agree on) the boundary. This vastly generalizes the Cannon-Thurston theorem.

We'll also show that the dynamics of $\pi_1(M)$ on S^1_v looks just like the pseudo-Anosov case. Each $g \in \pi_1(M)$ acts on S^1_u with an even number of fixed points that are alternately attracting and repelling.

Finally, we'll show that the action of $\pi_1(M)$ on S^1_v contains a surprising amount of information about $\mathfrak{F}$, even though it's just a discrete 1-dimensional image of the smooth 3-dimensional dynamics of a quasigeodesic flow. In particular, we can use it to find closed orbits in $\mathfrak{F}$. We'll show that if an element $g \in \pi_1(M)$ acts, up to conjugacy, as anything other than a hyperbolic Möbius transformation, then the flow $\mathfrak{F}$ contains a closed orbit in the free homotopy class of g. Consequently, if $\mathfrak{F}$ has no closed orbits then the action of $\pi_1(M)$ on S^1_v is an example of a Möbius-like group that is not conjugate into $PSL(2, \mathbb{R})$.

2. Statement of results

For convenience, we'll collect our main original theorems. If $\mathfrak{F}$ is a quasigeodesic flow on a 3-manifold M then P is the orbit space of $\widetilde{\mathfrak{F}}$. Calegari [**3**] constructs a universal circle S^1_u, which turns out to be larger than we need. We'll work with a quotient S^1_v on which it's more natural to study the dynamics of $\pi_1(M)$. Still, our first four theorems hold for both S^1_u and S^1_v.

The following idea is due to Calegari [**3**]. Calegari's original construction of the universal circle involves approximating the ends of sets in the plane by proper rays. Our more direct handling of ends allows us to complete the proof.

Compactification Theorem. Let $\mathfrak{F}$ be a quasigeodesic flow on a closed hyperbolic 3-manifold M, and let S^1_v be the universal circle. Then $\overline{P} = P \cup S^1_v$ has a natural topology making it into a closed disc with interior P and boundary S^1_v. The action of $\pi_1(M)$ on P extends to $\overline{P}$ and restricts to the universal circle action on $\partial\overline{P}$.

The space $\overline{P}$ is called the *end compactification* of P.

Continuous Extension Theorem. Let $\mathfrak{F}$ be a quasigeodesic flow on a closed hyperbolic 3-manifold M. The endpoint maps $e^{\pm} : P \to S^2_\infty$ admit unique continuous extensions to $\overline{P}$, and e^+ agrees with e^- on the boundary.

Consequently, the restriction $e : S^1_v \to S^2_\infty$ of $e^{\pm}$ is a π_1-equivariant sphere-filling curve, generalizing the Cannon-Thurston theorem.

A group Γ acting on the circle is called *hyperbolic Möbius-like* if each element is conjugate to a hyperbolic Möbius transformation.

Möbius-like Theorem. Let $\mathfrak{F}$ be a quasigeodesic flow on a closed hyperbolic 3-manifold M. Suppose that $\mathfrak{F}$ has no closed orbits. Then $\pi_1(M)$ acts on the universal circle S^1_v as a hyperbolic Möbius-like group.

In contrast:

Conjugacy Theorem. Let $\mathfrak{F}$ be a quasigeodesic flow on a closed hyperbolic 3-manifold M. Suppose that $\mathfrak{F}$ has no closed orbits. Then the action of $\pi_1(M)$ on the universal circle is not conjugate into $PSL(2, \mathbb{R})$.

Finally:

Pseudo-Anosov Dynamics Theorem. Let $\mathfrak{F}$ be a quasigeodesic flow on a closed hyperbolic 3-manifold M. Then each $g \in \pi_1(M)$ acts on S^1_v with an even number of fixed points (possibly zero) that are alternately attracting and repelling.

CHAPTER 2

Flows

A *flow* $\mathfrak{F}$ on a manifold M is a continuous $\mathbb{R}$-action. That is, for each point $x \in M$ and $t \in \mathbb{R}$ there is a point $t \cdot x \in M$, and this satisfies the following conditions:

(1) If $x \in M$ and $t, s \in \mathbb{R}$, then $(t+s) \cdot x = t \cdot (s \cdot x)$, and
(2) the map $\mathbb{R} \times M \to M$ defined by $(t, x) \mapsto t \cdot x$ is continuous.

We will call a flow *nonsingular* if for each $x \in M$ there is some t such that $t \cdot x \neq x$.

We will restrict attention to nonsingular flows. The orbit of each point $x \in M$ is the *flowline* $\mathbb{R} \cdot x$, and we'll use the same symbol to refer to both the flow and the associated oriented foliation

$$\mathfrak{F} = \{\mathbb{R} \cdot x | x \in M\}$$

by flowlines.

1. The orbit space

Given a nonsingular flow $\mathfrak{F}$ on a manifold M, the *orbit space* $\mathcal{O}_{\mathfrak{F}}$ is the set of flowlines in M, with the topology obtained as a quotient of M by collapsing each flowline to a point.

We denote the universal cover of M by $\widetilde{M}$. The flow $\mathfrak{F}$ on M lifts to a flow $\widetilde{\mathfrak{F}}$ on $\widetilde{M}$. The *flowspace* of $\mathfrak{F}$ is the orbit space $\mathcal{O}_{\widetilde{\mathfrak{F}}}$ of the lifted flow. The action of the fundamental group $\pi_1(M)$ on $\widetilde{M}$ by deck transformations preserves the foliation $\widetilde{\mathfrak{F}}$ by flowlines, so the quotient map

$$\pi : \widetilde{M} \to \mathcal{O}_{\widetilde{\mathfrak{F}}}$$

induces an action of $\pi_1(M)$ on the flowspace $\mathcal{O}_{\widetilde{\mathfrak{F}}}$.

While the orbit space of a flow on a closed 3-manifold is often not very nice, in general not even Hausdorff, the flowspace is often much nicer. In particular, the flows we're concerned with will be product covered:

DEFINITION 2.1. A flow $\mathfrak{F}$ on a 3-manifold M is called *product covered* if the foliation $\widetilde{\mathfrak{F}}$ by lifted flowlines is topologically equivalent to the product foliation of $\mathbb{R}^3$ as $\mathbb{R}^2 \times \mathbb{R}$. Equivalently, if the flowspace $\mathcal{O}_{\widetilde{\mathfrak{F}}}$ is homeomorphic to $\mathbb{R}^2$.

2. Quasigeodesics

DEFINITION 2.2. Let k, ϵ be non-negative constants. A curve $\gamma : \mathbb{R} \to X$ in a metric space (X, d) is a (k, ϵ)-*quasigeodesic* if it satisfies

$$1/k \cdot d(\gamma(x), \gamma(y)) - \epsilon \leq |x - y| \leq k \cdot d(\gamma(x), \gamma(y)) + \epsilon$$

for all $x, y \in \mathbb{R}$. A curve is called a *quasigeodesic* if it is a (k, ϵ)-quasigeodesic for some constants k, ϵ.

If X is a geodesic space and γ is a curve parametrized by arc length then the left-hand inequality always holds, so we can think of the quasigeodesic condition as ensuring that a curve makes definite progress on a large scale.

Recall that hyperbolic space $\mathbb{H}^n$ has a natural compactification as a closed disc by adding a sphere at infinity S^{n-1}_∞. Geodesics in $\mathbb{H}^n$ have well-defined endpoints in S^{n-1}_∞. In hyperbolic space, quasigeodesics are qualitatively similar to geodesics:

PROPOSITION 2.3 (see [**1**], pp. 399–404 or [**18**]).
1. *If γ is a (k, ϵ)-quasigeodesic in $\mathbb{H}^n$, then γ has well-defined, distinct endpoints in S^{n-1}_∞.*
2. *There is a constant C, depending only on k, ϵ, and n, such that every (k, ϵ)-quasigeodesic in $\mathbb{H}^n$ is contained in the C-neighborhood of the geodesic with the same endpoints.*

Quasigeodesity can also be formulated as a local condition. A curve $\gamma : \mathbb{R} \to \mathbb{H}^3$ is called a *c-local k-quasigeodesic* if $d(\gamma(x), \gamma(y)) \geq \frac{|x-y|}{k} - k$ for all $x, y \in \mathbb{R}$ with $|x - y| < c$.

LEMMA 2.4. *(Gromov, [**16**]) For every $k \geq 1$, there is a universal constant $c(k)$ such that every $c(k)$-local k-quasigeodesic is a $(2k, 2k)$-quasigeodesic.*

3. Quasigeodesic flows

DEFINITION 2.5. A nonsingular flow $\mathfrak{F}$ on a manifold M is said to be *quasigeodesic* if each flowline lifts to a quasigeodesic in $\widetilde{M}$. It is *uniformly quasigeodesic* if there are uniform constants k, ϵ such that each flowline lifts to a (k, ϵ)-quasigeodesic.

It turns out that we don't need to distinguish between quasigeodesic and uniformly quasigeodesic flows:

LEMMA 2.6 (Calegari. [**3**], Lemma 3.10). *Let M be a closed hyperbolic 3-manifold. Then every quasigeodesic flow on M is uniformly quasigeodesic.*

THEOREM 2.7 (Calegari. [**3**], Theorem 3.12). *Let M be a closed hyperbolic 3-manifold. Then every quasigeodesic flow on M is product covered.*

In particular, if $\mathfrak{F}$ is a quasigeodesic flow on a closed hyperbolic 3-manifold M, then $\mathcal{O}_{\widetilde{\mathfrak{F}}}$ is homeomorphic to the plane. From now on we'll use the notation $P := \mathcal{O}_{\widetilde{\mathfrak{F}}}$ for the flowspace of a quasigeodesic flow.

EXAMPLE 2.8 (Surface bundles). Let M be a closed surface bundle over the circle. Zeghib shows in [**37**] that any flow that is transverse to the foliation by surfaces is quasigeodesic.

EXAMPLE 2.9 (Mosher's examples). In [**25**], Mosher constructs a class of quasigeodesic flows that do not come from the surface bundle construction.

Start with the flow on $\mathbb{R}^3$ defined by

$$(x, y, z) \cdot t = (x \cdot \lambda^t, y \cdot \lambda^{-t}, z + t)$$

for some $\lambda > 1$. Cut along the eight planes defined by $xy = \pm\frac{1}{2}$, $x = \pm 1$, and $y = \pm 1$ and mod out by a translation in the z direction. The result is a flow $\mathfrak{F}_T$ on an octagonal torus T. Actually, this is just a *semi-flow*, i.e. the flow is not defined for all time, as some orbits leave T. The boundary of T consists of eight annuli, and the flow points inward on two of these, outward on two, and is tangent to the remaining four. There is a single closed orbit σ, the core of T. Every other flowline either spirals around σ or runs from one of the inward-pointing annuli to one of the outward-pointing annuli.

Now let Σ be a compact surface with four boundary components, and consider the unit flow on $\Sigma \times I$ that moves along the positive I direction. The boundary of $\Sigma \times I$ consists of two copies of Σ and four annuli. The flow points inward along $\Sigma \times \{0\}$, outward along $\Sigma \times \{1\}$, and is tangent along the four annuli. We can glue the four boundary annuli of $\Sigma \times I$ to the four annuli of T along which $\mathfrak{F}_T$ is tangent. The result is a 3-manifold M' with a flow $\mathfrak{F}'$ that points inward along one boundary surface ∂_- and outward along another boundary surface ∂_+. Glue these two surfaces together to build a closed 3-manifold M with a flow $\mathfrak{F}$, and let S be the embedded surface that is the image of ∂_+ and ∂_-.

If we make the right choices during this construction, we can ensure that M is a closed hyperbolic 3-manifold and S is an incompressible surface. The flow $\mathfrak{F}$ has one closed orbit, σ. If we follow any flowline γ, we find that it either eventually tracks σ, or it crashes through S with definite frequency. If it tracks σ then quasigeodesity follows easily. Otherwise, we can see that the lifts of S in the universal cover $\mathbb{H}^3$ obey a certain uniform separation property. This is enough to see that the orbits that crash through S makes definite progress towards a single point in S^2_∞, and quasigeodesity follows.

Example 2.10 (Finite depth foliations). Mosher's flows are particular examples of a more general phenomenon. The flow above is transverse to a depth-1 foliation. If M is a closed, orientable, irreducible 3-manifold and ξ is a nontrivial class in $H_2(M)$ then Gabai [**15**] constructs a taut, finite-depth foliation whose compact leaves represent ξ. Given such a foliation on a hyperbolic manifold, Mosher [**26**] constructs an *almost transverse* pseudo-Anosov flow, where almost transverse means that it is transverse after blowing up finitely many closed orbits. In [**9**], Fenley and Mosher show that these flows are also quasigeodesic.

3.1. Endpoint maps. Let M be a closed hyperbolic 3-manifold with a quasigeodesic flow $\mathfrak{F}$. Since each flowline is a quasigeodesic, there are well-defined *endpoint maps*

$$e^{\pm} : P \to S^2_\infty$$

that send each flowline to its positive/negative endpoint.

These endpoint maps are continuous. In fact, this characterizes quasi-geodesic flows:

THEOREM 2.11. *Let M be a closed hyperbolic 3-manifold with a flow $\mathfrak{F}$. Then $\mathfrak{F}$ is quasigeodesic if and only if the maps $e^{\pm}$ are well defined and continuous and satisfy $e^{+}(p) \neq e^{-}(p)$ for all $p \in \widetilde{M}$.*

The "if" direction is [**9**], Theorem B and the "only if" direction is [**3**], Lemma 4.3. Note that this theorem does not presuppose that the orbit space is planar.

Much of the peculiar behavior of quasigeodesic flows is contained in the following two observation.

LEMMA 2.12 (Calegari, [**3**], Lemma 4.4). *Let M be a closed hyperbolic 3-manifold with a quasigeodesic flow $\mathfrak{F}$. Then $e^{+}(P)$ and $e^{-}(P)$ are both dense in S^2_∞.*

This is true simply because $\pi_1(M)$ acts almost transitively on S^2_∞, so any nontrivial invariant subset is dense.

LEMMA 2.13 (Calegari, [**3**], Lemma 4.5). *Let M be a closed hyperbolic 3-manifold with a quasigeodesic flow $\mathfrak{F}$. If $D \subset P$ is an embedded disc in the flowspace, then*

$$e^{\pm}(D) = e^{\pm}(\partial D).$$

PROOF. We'll sketch this. Picture D embedded in $\mathbb{H}^3$ as a transversal to $\mathfrak{F}$. That is, choose a section $\sigma : D \to \mathbb{H}^3$ of the quotient map $\pi : \mathbb{H}^3 \to P$. Let S be the surface in $\mathbb{H}^3$ consisting of $\sigma(D)$, together with the positive half-flowlines emanating from the boundary of $\sigma(D)$. Choosing the correct orientation on S, if $x \in \mathbb{H}^3$ is a point on the positive side of S then the positve half-flowline emanating from x stays on the positive side of S and the negative half-flowline intersects $\sigma(D)$. Now suppose, contradicting our hypothesis, that there is some point $p \in D$ such that $e^+(p) \notin e^+(\partial D)$. Then $e^+(\partial D)$ encloses some non-empty region $U \ni e^+(p)$. Since $e^-(P)$ is dense, there is some flowline γ whose negative endpoint is in U. But then the *negative* half of γ is on the *positive* side of S, a contradiction. □

We can think of the endpoints of lifted flowlines as a subspace of $(S^2_\infty, S^2_\infty) \setminus \Delta$, where $\Delta = \{(p,p) | p \in S^2_\infty\}$ is the diagonal. That is, there is a map

$$e^+ \times e^- : P \to (S^2_\infty, S^2_\infty) \setminus \Delta$$

defined by

$$x \in P \mapsto (e^+(x), e^-(x)).$$

Recall that a map is called *proper* if preimages of compact sets are compact.

LEMMA 2.14. *Let $\mathfrak{F}$ be a quasigeodesic flow on a closed hyperbolic 3-manifold M. Then*

$$e^+ \times e^- : P \to (S^2_\infty, S^2_\infty) \setminus \Delta$$

is a proper map.

PROOF. Let $A \subset (S^2_\infty, S^2_\infty) \setminus \Delta$ be compact, and set $B = (e^+ \times e^-)^{-1}(A)$. Then B is closed since it's the preimage of a closed set.

Let X be the set of all geodesics with endpoints in A. Then we can find a compact set $K \subset \mathbb{H}^3$ that intersects every element of X. Recall that there is a constant C such that each flowline has Hausdorff distance at most C from the geodesic connecting its endpoints (Proposition 2.3). Therefore, each flowline with endpoints in A intersects the closed C-neighborhood $K' = N_C(K)$. The image $\pi(K')$ is compact, and $B \subset \pi(K')$, so B is bounded. Hence B is compact. □

3.2. Decompositions. The following is a consequence of Lemma 2.13.

LEMMA 2.15 (Calegari, [**3**], Lemma 4.8). *Let M be a closed hyperbolic 3-manifold with a quasigeodesic flow $\mathfrak{F}$. For any point in $p \in S^2_\infty$ in the image of e^+, every connected component of $(e^+)^{-1}(p)$ is unbounded, and similarly for e^-.*

Roughly, if some component $K \subset (e^+)^{-1}(p)$ were bounded, then we could surround it by a simple closed curve γ disjoint from $(e^+)^{-1}(p)$, which contradicts the previous lemma.

DEFINITION 2.16. Let $\mathfrak{F}$ be a quasigeodesic flow on a closed hyperbolic 3-manifold M. The *positive* and *negative decompositions* of the flowspace P are

$$\mathcal{D}^{\pm} = \{\text{components of } (e^{\pm})^{-1}(p) | p \in S^2_\infty\}.$$

By the previous lemma, the elements of $\mathcal{D}^+$ and $\mathcal{D}^-$ are closed, unbounded subsets of the plane. The action of $\pi_1(M)$ on P preserves $\mathcal{D}^+$ and $\mathcal{D}^-$.

We will draw an analogy between the pair of decompositions that come from a quasigeodesic flow and the pair of transverse singular foliations that come from a pseudo-Anosov flow. There's no notion of transversality for general closed unbounded subsets of the plane, but the following will suffice:

LEMMA 2.17. *Let $\mathfrak{F}$ be a quasigeodesic flow on a closed hyperbolic 3-manifold M. If $K \in \mathcal{D}^+$ and $L \in \mathcal{D}^-$ then $K \cap L$ is compact.*

PROOF. Let $p = e^+(K)$ and $q = e^-(L)$. Then $K \cap L$ is contained in $(e^+ \times e^-)^{-1}((p, q))$, which is compact by Lemma 2.14. □

Our study of quasigeodesic flows centers around the following theorem.

THEOREM 2.18 (Calegari, [**3**]). *Let M be a closed hyperbolic 3-manifold with a quasigeodesic flow $\mathfrak{F}$. Then $\pi_1(M)$ acts faithfully by homeomorphisms on a topological circle S^1_u, called the* universal circle *of $\mathfrak{F}$.*

We will prove this theorem in Chapter 5, where it follows easily from the technical machinery in Chapters 3 and 4. Each (positive or negative) decomposition element is unbounded, and we will think of the *ends* of a decomposition element as directions towards infinity. The universal circle is built by collating the ends of decomposition elements, taking a sort of completion, and collapsing some ends together.

CHAPTER 3

Circular Orders

An ordered triple of distinct points x, y, z in the circle is said to be positively ordered if it spans a positively oriented simplex in the disc, and negatively ordered otherwise. Four points $x, y, z, w \in S$ span four unoriented simplices, and the orientation on one pair of adjacent simplices determines the orientation on the other pair. This structure, in abstract, is called a circular order.

DEFINITION 3.1. A *circular order* on a set S is a map

$$\langle \cdot, \cdot, \cdot \rangle : S \times S \times S \to \{+1, 0, -1\}$$

such that

- $\langle x, y, z \rangle$ is nonzero exactly when x, y, and z are distinct,
- if τ is a permutation of $x, y, z \in S$, then $\langle \tau x, \tau y, \tau z \rangle = -1^{sgn(\tau)} \langle x, y, z \rangle$, and
- (cocycle condition) if $x, y, z, w \in S$ are distinct then $\langle y, z, w \rangle - \langle x, z, w \rangle + \langle x, y, w \rangle - \langle x, y, z \rangle = 0$.

A circularly ordered set is a set together with a circular order. An ordered triple of points x, y, z in a circularly ordered set is said to be *positively ordered* if $\langle x, y, z \rangle = +1$ and *negatively ordered* if $\langle x, y, z \rangle = -1$. In line with tradition, we also refer to these, respectively, as being in counterclockwise and clockwise order. An ordered n-tuple of points $x_0, x_2, x_3, ..., x_{n-1}$ in a circularly ordered set is said to be positively/negatively ordered if $\langle x_{i-1}, x_i, x_{i+1} \rangle = \pm 1$ for all i *mod* n.

EXAMPLE 3.2. The oriented circle S^1 has a natural circular order determined by the choice of orientation. One way to think of an orientation of is a choice of generator τ for $H^1(S^1; \mathbb{Z}) \simeq \mathbb{Z}$. A ordered triple of distinct points $x, y, z \in S^1$ determines a 1-cycle Δ_{xyz} consisting of three 1-simplices, the ordered segments between x and y, y and z, and z and x. We can think of $\langle x, y, z\rangle$ as the value of τ on Δ_{xyz}.

REMARK 3.3. If S is circularly ordered and $s \in S$, then one can construct a linear order $<_s$ on $S \setminus \{s\}$ by defining $q <_s p$ if $\langle s, p, q\rangle = +1$. For distinct $s, t \in S$, the orders $<_s$ and $<_t$ differ by a *cut*: For any p, q distinct from s, t, $p <_s q$ implies $p <_t q$ unless $p <_s t <_s q$ in which case $q <_t p$.

This provides an alternate definition of a circular order on S as a family of linear orderings $\{<_s\}_{s\in S}$, where each $<_s$ is defined on $S \setminus \{s\}$, and any two differ by a cut.

Let x and z be points in a circularly ordered set S. We define the *open interval*

$$(x, z) := \{y \in S | \langle x, y, z\rangle = +1\}$$

and the *closed interval*

$$[x, z] := (x, z) \cup \{x, z\},$$

and note that $[x, z] = (z, x)$.

The open intervals in a circularly ordered set form a basis for a topology, called the *order topology*. Note that the statement

$$y \in (x, z)$$

is the same as

$$\langle x, y, z\rangle = +1,$$

so we will often use the former, more familiar notation to work with circular orders.

A map

$$f : S_1 \to S_2$$

between circularly ordered sets S_1 and S_2 is called *order-preserving* if

$$\langle x, y, z\rangle = \langle f(x), f(y), f(z)\rangle$$

for all $x, y, z \in S_1$. Two circularly ordered sets are called *order-isomorphic* is there is an order-preserving bijection between them.

Given a circularly ordered set, we would like to know whether it embeds in the circle. This is immediate if our set is countable.

LEMMA 3.4. *Let S be a countable circularly ordered set. Then S is order-isomorphic to a subset of the circle.*

PROOF. We will construct a map $f : S \hookrightarrow S^1$ directly.

Enumerate $S = \{s_i\}_{i=1}^{\infty}$. Map s_1 and s_2 to antipodal points. Once we have specified f on $s_1, s_2, ..., s_{i-1}$ there are unique $a, b < i$ such that $s_i \in (s_a, s_b)$, and no other s_j, $j < i$ lies in (s_a, s_b). Map s_i to the midpoint of the interval $(f(s_a), f(s_b))$. □

REMARK 3.5. There are other ways to embed a countable circularly ordered set in the circle. For example, if $f : S \to S^1$ is an order-preserving map, one can blow up any point $p \in S^1 \setminus f(S)$ to an open interval. Our construction is more natural since it does not add such extraneous intervals.

More precisely, a *gap* in a circularly ordered set S is an ordered pair of points $x, y \in S$ such that (x, y) is empty. Let S be a countable circularly ordered set, and let $f : S \hookrightarrow S^1$ be the order-isomorphism constructed in the proof of Lemma 3.4. We will show that if $I \subset S^1$ is a maximal open interval in the complement of $f(S)$, then there are $x, y \in S$ such that (x, y) is a gap and $I = (f(x), f(y))$. Indeed, otherwise we could find i and j such that $f(s_i)$ and $f(s_j)$ are arbitrarily close to, but not both equal to, the endpoints of I. Then for k large enough, $f(s_k)$ would be the midpoint of $(f(s_i), f(s_j))$, and hence contained in I, a contradiction.

1. Order completion

DEFINITION 3.6. A circularly ordered set S is called *order complete* if every nested sequence of closed intervals $I_1 \subset I_2 \subset I_3 ...$ has non-empty intersection.

Every circularly ordered set S has a canonical *order completion.*

LEMMA 3.7. *Let S be a circularly ordered set. There is a unique, minimal, order complete set $\overline{S}$ containing S, and S is dense in $\overline{S}$ with the order topology.*

By minimal, we mean that if T is an order complete set and $i : S \hookrightarrow T$ is an order-preserving injection, then i extends to $\overline{S}$ (though the extension may not be unique).

Proof. We will construct $\overline{S}$ directly.

Call a sequence $(I_i)_{i=1}^{\infty}$ of closed intervals in S *admissible* if $I_i \subset I_{i+1}$ for each i, and $\bigcap_i I_i = \emptyset$. That is, admissible sequences are the ones that make S fail the test of order completeness. If (I_i) and (J_j) are admissible sequences, define $(I_i) \sim (J_j)$ if for every $n > 0$ there exists $k > 0$ such that $I_k \subset J_n$.

Let's check that this is an equivalence relation. For symmetry, suppose $(I_i) \sim (J_j)$. Given $n' > 0$, the fact that $\bigcap I_i = \emptyset$ means that for k' large enough, $I_{k'}$ is either contained in or disjoint from $J_{n'}$. But it cannot be disjoint, since by hypothesis (setting $n = k'$) there is a k such that $I_k \subset J_{k'}$. Thus $(J_j) \sim (I_i)$. Transitivity and reflexivity are obvious.

Now let S' be the set of admissible sequences modulo this equivalence, and set $\overline{S} = S \cup S'$. To define the circular order on $\overline{S}$, we can represent each point $x \in \overline{S} \setminus S'$ by a constant sequence $([x, x])$. If (I_i), (J_j), and (K_k) represent distinct elements of $\overline{S}$ then for n large enough, I_n, J_n, and K_n are disjoint. Choose $x \in I_n$, $y \in J_n$, and $z \in K_n$ and set $\langle (I_i), (J_j), (K_k) \rangle = \langle x, y, z \rangle$. It is easy to check that this is well-defined, order complete, and minimal, and that S is dense in $\overline{S}$; uniqueness follows immediately using minimality. □

2. Universal circles

We would like to know which circularly ordered sets embed in the circle. 2nd countability is certainly necessary, and this turns out to be the only requirement.

In practice, we'll use the following criterion to check for 2nd countability. Recall that a gap in a circularly ordered set S is an ordered pair of points $x, y \in S$ with $(x, y) = \emptyset$.

Lemma 3.8. *Let S be a circularly ordered set. Then S is 2nd countable if and only if it is separable and has countably many gaps.*

PROOF. Suppose that S is separable and has countably many gaps. Let $S_0 \subset S$ be a countable, dense set that contains the endpoints of all gaps, and let $\mathcal{U}$ be the collection of open intervals with endpoints in S_0. We will show that $\mathcal{U}$ is a basis for the order topology on S.

Suppose $U \subset S$ is open, and let $x \in U$. Then x is contained in some open interval $(a, b) \subset U$. We will find $a', b' \in S_0$ such that $x \in (a', b') \subset (a, b)$. If (a, x) is nonempty then it contains some point in S_0, so let a' be this point. Otherwise (a, x) is a gap, so $a \in S_0$ and we can take $a' = a$. A suitable b' can be found similarly, so $\mathcal{U}$ is indeed a basis.

For the converse, suppose that S is 2nd countable. It is well-known that 2nd countable implies separable. To see that S has countably many gaps, let $\mathcal{U}$ be a basis, which we may assume consist of open intervals. We will show that for each gap (a, b) in S, there is a point b' so that (a, b') is contained in $\mathcal{U}$. Consequently, if there were uncountably many gaps then we will have produced uncountably many distinct elements in $\mathcal{U}$.

Indeed, suppose (a, b) is a gap, and let $c \in S \backslash \{a, b\}$. Note that (a, c) is an open set that contains b, so there is some interval $I \in \mathcal{U}$ with $b \in I \subset (a, c)$. Since $(a, b) = \emptyset$, such an interval must be of the form $I = (a, b')$ for some $b' \in S$. □

EXAMPLE 3.9. Consider the *double circle* $S := S^1 \times \{0, 1\}$, where each of the subspaces $S^1 \times \{0\}$ and $S^1 \times \{1\}$ have the usual circular order, and each point in $S^1 \times \{1\}$ is immediately counterclockwise to the corresponding point in $S^1 \times \{0\}$. This set has uncountably many gaps; for each $s \in S^1$, $(s, 0)$ and $(s, 1)$ form a gap.

PROPOSITION 3.10. *Let S be a 2nd countable, order complete, circularly ordered set. Then S is order-isomorphic to a compact subset of the circle.*

PROOF. We'll construct a map $f : S \to S^1$.

By the previous lemma we can find a set $S' \subset S$ that is countable, dense, and contains the endpoints of gaps. Enumerate this set as $\{s_i\}_{i=1}^{\infty} = S'$ and define f on S' as in Lemma 3.4.

If $x \in S \setminus S'$, choose a sequence of points $x_i \in S'$ converging to x, and define $f(s) = \lim_{i\to\infty} f(x_i)$. To see that this is well-defined, recall that S is 2nd countable, so we can find a nested sequence of intervals $(a_i, b_i) \subset S$ with $a_i, b_i \in S'$ that form a neighborhood basis for x. Since these are nested, the sequence of intervals $(f(a_i), f(b_i))$ has a well-defined Hausdorff limit. Let I be this limit, and note that I is in the complement of $f(S')$ since $\bigcap(a_i, b_i) = x$. Therefore, by Remark 3.5, I must be a point.

It follows that f is well-defined and order-preserving, and injectivity is obvious. We'll show that $f(S)$ is closed. Suppose $a \in S^1$ lies in the closure of $f(S)$. If $f(S)$ approaches a from only one side, then a is an endpoint of some complementary open interval, hence it lies in $f(S)$ by the Remark 3.5. If $f(S)$ approaches a from both sides, let x_i and y_i be sequences so that $f(x_i)$ and $f(y_i)$ approach a from the counterclockwise and clockwise direction, respectively. Since S is order complete, there is a point x in $\bigcap_i [x_i, y_i]$, and $a = f(x)$.

Finally, continuity of f follows from the fact that $f(S)$ is closed. If $(a, b) \subset S^1$ is an open interval, let a' be the point in $f(S)$ closest to a on its clockwise side, and let b' be the point in $f(S)$ closest to b on its counterclockwise side. Then $f^{-1}((a, b)) = (f^{-1}(a'), f^{-1}(b'))$. □

Given an uncountable compact subset of the circle, we would like to collapse complementary intervals to obtain the circle itself. Naively, one might take the closure of each complementary interval and collapse this to a point. However, some of these may intersect, so we'd have to combine them before collapsing. We'd then have to combine any of these larger intervals that intersect, etc., and it's not clear *a priori* that this process terminates.

We'll be more clever and use the Cantor-Bendixson theorem. See [**19**] for a proof.

THEOREM 3.11 (Cantor-Bendixson). *Let Y be a closed subset of a separable, completely metrizable topological space X. Then $Y = T \cup U$ where T is closed and perfect and U is countable.*

CONSTRUCTION 3.12 (Universal circles). Let S be an uncountable, 2nd countable circularly ordered set. Take the order completion $\overline{S}$, which is an uncountable compact subset of S^1, and decompose it as $\overline{S} = T \cup U$ as in the Cantor-Bendixson theorem. The *universal circle* S^1_u is the image of the map

$$\phi : S \to S^1_u$$

that collapses the closure of each interval in $S^1 \setminus T$.

Note that points $x, y \in S$ are collapsed in the universal circle if and only if there are countably many points between them, i.e. either (x, y) or (y, x) is countable.

3. Circularly ordered groups

The following section is not integral to our discussion, but it provides examples of manifolds that admit no quasigeodesic flows.

DEFINITION 3.13. Let G be a group. A *(left) circular order* on G is a circular order on the underlying set, such that the left action of any element induces an order isomorphism $G \to G$. A group is *(left) circularly orderable* if it admits a left circular order.

EXAMPLE 3.14. The usual order on the circle is a (left- and right-invariant) circular order for the multiplicative group $S^1 = \{z \in \mathbb{C} \| z| = 1\}$.

The following well-known theorem is a useful characterization of circularly orderable groups. A proof can be found in [**4**], Section 2.6.

THEOREM 3.15. *A countable group is circularly orderable if and only if it admits a faithful representation to* $Homeo^+(S^1)$*, the group of orientation-preserving homeomorphisms of the circle.*

One can check whether a group is circularly ordered in some cases and for certain classes of groups, such as automatic groups [**2**]. If a group is not circularly ordered then it cannot have any quasigeodesic or pseudo-Anosov flows, since the existence of such a flow implies that the fundamental group acts faithfully on the circle by Theorem 2.18.

Example 3.16 ([**4**], Example 2.102). The Weeks manifold is the smallest-volume closed, orientable hyperbolic 3-manifold. One can show that its fundamental group is not circularly orderable, so it has no quasigeodesic flows.

CHAPTER 4

Decompositions

A quasigeodesic flow gives rise to decompositions $\mathcal{D}^{\pm}$ of the planar orbit space P by closed, unbounded subsets (see Section 3.1). These decompositions have some extra properties which we will discuss in Chapter 5, but much of the machinery we'll need to analyze them makes sense in a more general context.

Definition 4.1. An *unbounded continuum* is a closed, connected, unbounded subset of the plane.

Definition 4.2. An *unbounded decomposition* $\mathcal{D}$ of the plane P is a collection of unbounded continua such that

(1) for any $K, L \in \mathcal{D}$ distinct, $K \cap L = \emptyset$, and
(2) the union of all $K \in \mathcal{D}$ covers P.

Sierpinski proved [**30**] that no Hausdorff continuum can be decomposed into countably many disjoint closed sets. If $\mathcal{D}$ is a nontrivial (i.e. containing more than one element) unbounded decomposition of the plane, then an arc γ between any two decomposition elements is a Hausdorff continuum. We obtain a partition of γ by taking the intersection of each decomposition element with γ. Applying Sierpinski's theorem:

Lemma 4.3. *Every nontrivial unbounded decomposition of the plane contains uncountably many elements.*

The decompositions $\mathcal{D}^+$ and $\mathcal{D}^-$ that come from a quasigeodesic flow are unbounded decompositions. Recall that if $K \in \mathcal{D}^+$ and $L \in \mathcal{D}^-$, then K and L are *eventually disjoint*, i.e. the intersection $K \cap L$ is compact (Lemma 2.17). This is all we need to work with ends, so we'll expand our notion of a decomposition to handle the union $\mathcal{D} = \mathcal{D}^+ \cup \mathcal{D}^-$.

DEFINITION 4.4. A *generalized unbounded decomposition* of the plane P is a collection of unbounded continua such that

(1') for any $K, L \in \mathcal{D}$ distinct, $K \cap L$ is compact, and
(2) the union of all $K \in \mathcal{D}$ covers P.

Unlike unbounded decompositions, generalized unbounded decompositions need not be uncountable. For example, the plane may be decomposed into elements K_i, where each K_i is a disc of radius i together with a proper ray. Of course, the generalized unbounded decomposition $\mathcal{D} = \mathcal{D}^+ \cup \mathcal{D}^-$ coming from a quasigeodesic flow is uncountable, since $\mathcal{D}^+$ and $\mathcal{D}^-$ are uncountable.

Throughout this chapter, P will denote the plane and $\mathcal{D}$ a generalized unbounded decomposition of P. The reader may think of P as the flowspace of a quasigeodesic flow with $\mathcal{D}$ the associated decomposition.

We'll summarize the results of this chapter. Suppose that $\mathcal{D}$ is a generalized unbounded decomposition of the plane P. Each decomposition element $K \in \mathcal{D}$ has a natural set of ends, $\mathcal{E}(K)$, which we'll define formally in Section 1. In Section 2, we'll show that the set $\mathcal{E}(\mathcal{D})$ of all ends of decomposition elements has a natural circular order. In Section 3 we'll show that $\mathcal{E}(\mathcal{D})$ has the necessary topological properties to construct a universal circle S^1_u. Finally, in Section 4 we'll use the universal circle to construct the end compactification of P. This is a closed disc $\overline{P}$ with interior P and boundary S^1_u that behaves well with respect to the decomposition $\mathcal{D}$.

1. Ends

Freudenthal [**13**] introduced the notion of the ends of a topological space. We'll state the definition for metric spaces, since we'd like to be able to use bounded sets instead of just compact sets.

DEFINITION 4.5. Let X be a metric space. An *end* of X is a map ξ that assigns to each bounded set $A \subset X$ a component $\xi(A)$ of $X \setminus A$, with the condition that if $A' \supset A$ then $\xi(A') \subset \xi(A)$. The set of all ends of X is denoted $\mathcal{E}(X)$.

In the sequel we'll specialize to subsets of the plane. Suppose $X \subset$, and $\xi \in \mathcal{E}(X)$. Then $\xi(A)$ is defined for bounded sets $A \subset P$ by setting $\xi(A) = \xi(A \cap X)$.

REMARK 4.6. The data of an end can be expressed in a countable manner. In order to specify an end it is enough to keep track of its values on any exhaustion of P by bounded sets. Indeed, fix such an exhaustion $(A_i)_{i=0}^{\infty}$ and let $X_i = \xi(A_i)$ for each i. If $A \subset P$ is any bounded set, then for i sufficiently large, X_i is disjoint from A. Hence $\xi(A)$ is the component of $X \setminus A$ containing X_i.

In practice, we will use this to explicitly specify an end. Let $(A_i)_{i=0}^{\infty}$ be a bounded exhaustion of P, and suppose we have a sequence $(X_i)_{i=0}^{\infty}$ where X_i is a component of $X \setminus A_i$ and $X_{i+1} \subset X_i$ for each i. Then there is a unique end $\xi \subset \mathcal{E}(X)$ with $\xi(A_i) = X_i$ for all i.

EXAMPLE 4.7. Suppose $K \subset P$ is a proper ray in the plane. There is exactly one end $\kappa \in \mathcal{E}(K)$. Indeed, let $\gamma : [0, \infty) \to P$ be a parametrization of K. If $D \subset P$ is a bounded open disc, let t be the last time that γ intersects $\overline{D}$. Then $\gamma([t, \infty))$ is the only unbounded component of $K \setminus D$, so $\kappa(D) = \gamma([t, \infty))$.

EXAMPLE 4.8. Similarly, if $K \subset P$ is the union of n proper rays that are eventually disjoint, then K has exactly n ends, each corresponding to one of the rays.

LEMMA 4.9. *Let $K \subset P$ be an unbounded continuum in the plane and let $A \subset P$ be a bounded set. Then some component of $K \setminus A$ is unbounded.*

PROOF. Let $\hat{P} = P \cup \{\infty\}$ be the one-point compactification of P, and let $\hat{K}$ be the closure of K in $\hat{P}$, i.e. $\hat{K} = K \cup \{\infty\}$. We'll start by showing that the connected component of $\hat{K} \setminus A$ containing ∞ is not just $\{\infty\}$. Let $\{U_i\}_{i=1}^{\infty}$ be a sequence of connected open neighborhoods of $\hat{K}$ with intersection $\cap_i U_i = \hat{K}$, and let D be a compact disc in P that contains the closure of A. Fix a point $p \in (K \cap D)$, and for each i let $\gamma_i : [0, 1] \to P$ be an arc from p to ∞ contained in U_i. For each i let t_i be the last time that γ_i intersects D, and set $\gamma_i' = \gamma([t_i, 1])$. After taking a subsequence, the γ_i Hausdorff converge (in $\hat{P}$) to some compact, connected set $Z \subset \hat{K} \setminus A$ that contains both p and ∞.

Now we'll show that some component of $Z \setminus \{\infty\}$ accumulates on ∞, completing the proof. Fix a point $q \in Z \setminus \{\infty\}$. The collection of subcontinua of Z that contain $\{q, \infty\}$ is inductively ordered by inclusion, so we can use Zorn's lemma to produce a minimal element Z'. We claim that $Z' \setminus \{\infty\}$ is connected. Suppose on the contrary that $Z' \setminus \{\infty\} = B \cup C$, where B and C are separated. If $q \in B$ then $B \cup \{\infty\}$ is a continuum containing $\{q \cup \infty\}$, contradicting the minimality of Z'. So we see that $Z' \setminus \{\infty\}$ is a connected, unbounded subset of K that is disjoint from A. □

COROLLARY 4.10. *Let $K \subset P$ be an unbounded continuum in the plane. Then K has at least one end.*

PROOF. Let $(A_i)_{i=1}^{\infty}$ be an exhaustion of the plane by nested, bounded open sets. By the preceding lemma we can find a sequence of sets K_i, where K_i is an unbounded component of $K_{i-1} \setminus A_i$ for each i. By Remark 4.6 this determines an end $\kappa \in \mathcal{E}(K)$. □

The notion of an end behaves a bit less nicely for sets that are not closed. For example, we'll produce a connected, unbounded open set with no ends.

EXAMPLE 4.11. Let X be the set in $\mathbb{R}^2$ consisting of the line segment $X_0 = [0, 1] \times \{0\}$ together with the line segments $X_n = \{1/n\} \times [0, n]$ for all integers $n \geq 1$. Then every component of $X \setminus X_0$ is bounded, so $\mathcal{E}(X) = \emptyset$. By thickening each X_n we can make X open.

Let $X \subset Y \subset P$. If ξ is an end of X, there's a corresponding end ν of Y such that $\nu(A) \supset \xi(A)$ for every bounded set $A \subset P$. Therefore, we can define a natural map

$$\mathcal{E}(X) \to \mathcal{E}(Y),$$

that sends each $\xi \in \mathcal{E}(X)$ to its corresponding $\nu \in \mathcal{E}(Y)$.

We'll say that the set $X \subset P$ is a *tail* of $Y \subset P$ if X is an unbounded component of $Y \setminus A$ for some bounded set $A \subset P$. In this case, the map $\mathcal{E}(X) \to \mathcal{E}(Y)$ is injective.

In particular, any set that contains an unbounded continuum has at least one end. The following lemma provides a converse of this for open sets.

LEMMA 4.12. *Let $U \subset P$ be an open, connected set in the plane with at least one end. If $\mu \in \mathcal{E}(U)$, then there is properly embedded ray $K \subset U$ such that the end $\kappa \in K$ maps to μ under the canonical map $\mathcal{E}(K) \to \mathcal{E}(U)$.*

PROOF. Let $(A_i)_{i=1}^{\infty}$ be an exhaustion of P by nested, bounded open sets. Fix a point $k_0 \in U$, and for each i let k_i be a point in $\mu(A_i)$. For each i, let K_i be an arc contained in $\mu(A_{i-1})$ that connects k_{i-1} to k_i. Concatenating these arcs, we obtain a ray K. This ray is proper since it stays outside of A_{i-1} after it hits k_i. We can turn it into an embedded ray by taking a shortcut whenever it intersects itself. By construction, it is immediate that $\kappa \in \mathcal{E}(K)$ maps to μ. □

We're particularly interested in unbounded continua and their complements. Note that if $K \subset P$ is an unbounded continuum then each component U of $P \setminus K$ is a disc. Indeed, if γ is a simple closed curve in U then the compact region bounded by γ must be disjoint from K. So U is simply connected, hence a disc by the uniformization theorem.

LEMMA 4.13. *Let $K \subset P$ be an unbounded continuum and let U be a connected component of $P \setminus K$. Then U has at most one end.*

PROOF. Let $\mu_1, \mu_2 \in \mathcal{E}(U)$. By Lemma 4.12, there are proper rays L_1 and L_2 whose ends λ_1 and λ_2 map to μ_1 and μ_2 respectively. Let L be a curve consisting of L_1, L_2, and an arc connecting them, and make L into an embedded curve by taking shortcuts if necessary. By the Jordan-Schönflies theorem there is a homeomorphism of P taking L to the x-axis. Now K is contained in, say, the lower half plane, so the entire upper half plane is contained in U. It is now obvious that, for any bounded set $A \subset P$, $\lambda_1(A)$ and $\lambda_2(A)$ are contained in the same component of $U \setminus A$. Indeed, we can connect $\lambda_1(A)$ to $\lambda_2(A)$ with an arc in the upper half plane. Therefore, $\mu_1 = \mu_2$. □

It is possible for U to be unbounded, but still have no ends. For example, let $U \subset P$ be the no-ended open set in Example 4.11 and let $K = P \setminus U$.

On the other hand, if K is an element of an unbounded decomposition $\mathcal{D}$, then every component U of $P \setminus K$ is 1-ended, since some other decomposition element L is contained in U.

2. Circular orders

Let $\mathcal{C}$ be a collection of unbounded continua in the plane that are eventually disjoint. We'll define the set of ends of $\mathcal{C}$ to be

$$\mathcal{E}(\mathcal{C}) := \bigcup_{K \in \mathcal{C}} \mathcal{E}(K).$$

In this section we'll define a circular order on $\mathcal{E}(\mathcal{C})$.

Let's sketch the definition of the circular order before going into details. Refer to Figure 1. If $\kappa, \lambda, \mu \in \mathcal{E}(\mathcal{C})$ we can find a bounded open disc $D \subset P$ such that $K = \kappa(D)$, $L = \lambda(D)$, and $M = \mu(D)$ are disjoint unbounded continua, and we can ensure that no one separates the other two. Decide whether the triple K, L, M is positively or negatively ordered as follows: Let γ be an arc from K to M that avoids L. If L is on the positive (right) side of γ, we'll set $\langle K, L, M \rangle = +1$, and if it is on the negative side we'll set $\langle K, L, M \rangle = -1$. Finally, set $\langle \kappa, \lambda, \mu \rangle = \langle K, L, M \rangle$.

Our task in the rest of this section is to define this more carefully. In order to see that $\langle , , \rangle$ defines a circular order on $\mathcal{E}(\mathcal{C})$, we must show that it does not depend on the choice of γ and D, and that it satisfies the cocycle condition on triples (see Definition 3.1).

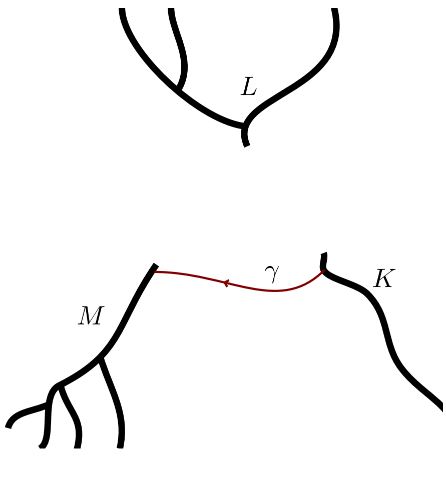

FIGURE 1. $\langle K, L, M \rangle = +1$

2.1. For collections of unbounded continua. A collection $\mathcal{C}$ of disjoint unbounded continua in the plane is said to be *mutually nonseparating* if there are no $K, L, M \in \mathcal{C}$ such that K separates L from M in the plane. We'll start by defining a circular order on such a collection.

LEMMA 4.14. *Let $K_1, K_2, ..., K_n \subset P$ be a disjoint, mutually nonseparating, unbounded continua in the plane. There is exactly one component $C(K_1, ..., K_n)$ of $P \setminus (\bigcup_i K_i)$ that limits on every one of the K_i. Every other component limits only one of the K_i.*

PROOF. The K_i are presumed to be mutually nonseparating, so for each i there is some component U_i of $P \setminus K_i$ that contains *every* K_j, for $j \neq i$. We'll expand the K_i by setting $K_i' := P \setminus U_i$ for each i. By [**36**], Thm. I.9.11 [1], each K_i' is connected.

Let

$$C := \bigcap_{i=1}^{n} U_i = P \setminus (\bigcup_{i=1}^{n} K_i'.$$

Observe that C is connected, since each K_i' is nonseparating, and the union of finitely many disjoint nonseparating sets in the plane is nonseparating ([**36**], Thm. II.5.28a). It is maximal among connected subsets of $P \setminus (\bigcup_i K_i)$ since any $x \in P \setminus (\bigcup_i K_i)$ that is not in C is separated from C by one of the K_i. Therefore, C is a component of $P \setminus (\bigcup_i K_i)$, and it clearly limits on every one of the K_i, so we define

$$C(K_1, ..., K_n) := C.$$

It is easy to see that every other component of $P \setminus (\bigcup_i K_i)$ is just one of the non-U_i components of $P \setminus K_i$ for some i. $\square$

We'll use the notation $C(\cdots)$ regularly in what follows.

DEFINITION 4.15. Let K and M be closed subsets of the plane P. An *arc from K to M* is an embedded, oriented arc γ with initial point in K, terminal point in M, and whose interior $\mathring{\gamma}$ is disjoint from both K and M.

[1] If C is a connected subset of a connected space M and A is a component of $M \setminus C$, then $M \setminus A$ is connected.

If K and M are disjoint unbounded continua, then any arc from K to M has interior in $C(K, M)$, and it separates $C(K, M)$ into two discs.

DEFINITION 4.16. Let $K, M \subset P$ be disjoint unbounded continua in the plane, and let γ be an arc from K to M. We define $C^+(\gamma; K, M)$ to be the component of $C(K, M) \setminus \gamma$ on the positive side of γ and $C^-(\gamma; K, M)$ to be the component on the negative side. If the sets K and M are implicit then we'll use the abbreviation $C^\pm(\gamma)$ for $C^\pm(\gamma; K, M)$.

We note:

LEMMA 4.17. *Let $K, M \subset P$ be disjoint unbounded continua in the plane, and let γ be an arc from K to M. Then $C^+(\gamma)$ and $C^-(\gamma)$ are both unbounded.*

PROOF. Suppose that $C^+(\gamma)$ is bounded. Then Fr $C^+(\gamma) = K' \cup \gamma \cup M'$, where K' and M' are bounded subsets of K and M respectively. [2] Choose a point p in the interior of $C^+(\gamma)$. Neither K' nor γ separate p from ∞ in $\hat{P} = P \cup \{\infty\}$, and $K' \cap \gamma$ is connected. So $K' \cup \gamma$ does not separate p from ∞ by **[36]**, Thm II.5.29. [3] Similarly, neither $K' \cup \gamma$ nor L' separate p from ∞ so $(K' \cup \gamma) \cap L' =$ Fr $C^+(\gamma)$ does not separate p from ∞. This is a contradiction, so $C^+(\gamma)$ must be unbounded. The same argument works for $C^-(\gamma)$. □

LEMMA 4.18. *Let $K, L, M \subset P$ be disjoint, mutually nonseparating unbounded continua. Then $L \subset C(K, M)$.*

PROOF. Note that $C(K, L, M)$ limits on both K and M, so it must be contained in $C(K, M)$. But $C(K, L, M)$ also limits on L, so L must also be in $C(K, M)$. □

Therefore, if K, L, M are disjoint unbounded continua and γ is an arc from K to M that avoids L, then L is contained in either $C^+(\gamma; K, M)$ or

[2]Recall that if A is a subspace of a topological space X, the *frontier* of A is Fr $A = \mathrm{cl}(A) \cap \mathrm{cl}(X \setminus A)$.

[3]If x and y are points in S^2 which are not separated by either of the closed sets A and B, and $A \cap B$ is connected, then x and y are not separated by $A \cup B$.

$C^-(\gamma; K, M)$. We now have all the notation we need to define a circular order on a triple of unbounded continua.

Definition 4.19. Let $K, L, M \subset P$ be disjoint, mutually nonseparating, unbounded continua in the plane. Choose an arc γ from K to M that avoids L. Define

$$\langle K, L, M \rangle = +1$$

if $L \subset C^+(\gamma; K, M)$ and

$$\langle K, L, M \rangle = -1$$

if $L \subset C^-(\gamma; K, M)$

Before we prove that this is well-defined, we need to understand what it looks like when we have multiple arcs between two unbounded continua.

Lemma 4.20. *Let $K, M \subset P$ be disjoint unbounded continua in the plane and let γ_1, γ_2 be disjoint arcs from K to M. We can relabel γ_1 and γ_2 as the "outer arc" γ_+ and "inner arc" γ_- in such a way that*

(1) $\gamma_+ \subset C^+(\gamma_-)$ *and* $\gamma_- \subset C^-(\gamma_+)$;
(2) $C(K, M) \setminus (\gamma_+ \cup \gamma_-)$ *consists of three components:*

$$C^+(\gamma_+),$$

$$C^-(\gamma_+) \cap C^+(\gamma_-),$$

and

$$C^-(\gamma_-);$$

(3) $C^+(\gamma_+) \subset C^+(\gamma_-)$ *and* $C^-(\gamma_-) \subset C^-(\gamma_+)$; *and*
(4) $C^-(\gamma_+) \cap C^+(\gamma_-)$ *has no ends.*

Proof. If $\gamma_1 \subset C^+(\gamma_2)$, then label $\gamma_+ := \gamma_1$ and $\gamma_- := \gamma_2$. If $\gamma_2 \subset C^+(\gamma_1)$, then label $\gamma_+ := \gamma_2$ and $\gamma_- := \gamma_1$. We need to show that

$$\gamma_- \subset C^-(\gamma_+).$$

If on the contrary $\gamma_- \subset C^+(\gamma_+)$, then there is an oriented arc λ from γ_- to γ_+ whose interior $\mathring{\lambda}$ lies on the positive side of both γ_+ and γ_-. Let λ' be an arc from $\gamma_+ \cap K$ to $\gamma_- \cap K$ that avoids M. Let c be the simple closed

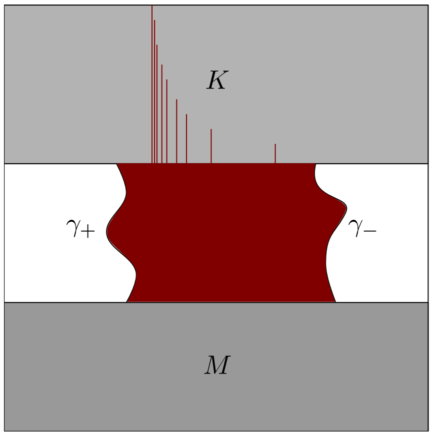

FIGURE 2. An unbounded middle region.

curve that follows λ, then γ_+ backwards to where it meets K, then λ', and finally γ_- to where it meets λ. Then c separates $M \cap \gamma_+$ from $M \cap \gamma_-$, which is impossible.

This completes the proof of (1). Statements (2) and (3) follow easily.

Finally, suppose that $C^-(\gamma_+) \cap C^+(\gamma_-)$ has an end. Then by Lemma 4.12 it contains a properly embedded ray $\lambda : [0, \infty) \to P$. Let ν be an arc that runs from γ_+ to $\lambda(0)$, and then to γ_-. Either K or M is on the same side of ν as λ, say K is. Let λ' be an arc from $\lambda(0)$ to M, and observe that $M \cup \lambda \cup \lambda'$ separates γ_+ from γ_-. Therefore, it separates $\gamma_+ \cap K$ from $\gamma_- \cap K$, which is impossible. □

REMARK 4.21. Let $K, M \subset P$ be disjoint unbounded continua in the plane, and let Γ be a collection of disjoint arcs from K to M. Then Γ is naturally endowed with a linear order: for $\gamma, \gamma' \in \Gamma$, define $\gamma' > \gamma$ if $\gamma' \subset C^+(\gamma)$. This is well-defined by the preceding lemma.

REMARK 4.22. While the middle section $C^-(\gamma_+) \cap C^+(\gamma_-)$ in the preceeding lemma has no ends, it may be unbounded. For example, see Figure 2 where K is upper rectangle with the no-ended open set of Example 4.11 cut out.

PROPOSITION 4.23. *Let $\mathcal{C}$ be a collection of disjoint, mutually nonseparating unbounded continua in the plane. Then $\langle , , \rangle$ defines a circular order on $\mathcal{C}$.*

Proof. Let $K, L, M \in \mathcal{C}$ be distinct. We will start by showing that $\langle K, L, M\rangle$ does not depend on the choice of γ. Note that if γ_1 and γ_2 are arcs from K to M then we can find another arc from K to M disjoint from both of them. So we may assume without loss of generality that γ_1 and γ_2 are disjoint. Relabel the arcs according to Lemma 4.20 and suppose that $L \subset C^+(\gamma_+)$. By part (4) of the lemma, L is contained in either $C^+(\gamma_+)$ or $C^-(\gamma_-)$, so $\langle K, L, M\rangle$ is well-defined by part (3) of the lemma.

Next we'll show that

$$\langle K, L, M\rangle = (-1)^{sgn(\tau)}\langle \tau(K), \tau(L), \tau(M)\rangle,$$

for a permutation τ of (K, L, M).

It is immediate that $\langle K, L, M\rangle = -\langle M, L, K\rangle$, so we just need to show that $\langle K, L, M\rangle = -\langle K, M, L\rangle$. Indeed, assume that $\langle K, L, M\rangle = +1$, i.e. that $L \subset C^+(\gamma)$ for some arc γ from K to M. We can find an arc γ' from K to M that intersects L and lies on the positive side of γ. Let γ'' be the sub-arc of γ' that runs from K to L. Then γ is on the negative side of γ'', hence so is M. Therefore, $\langle K, M, L\rangle = -1$ as desired.

It remains to show that the circular order is compatible on quadruples, i.e. to verify the cocycle condition. Let $K, L, M, N \in \mathcal{C}$. Suppose that $\langle K, L, M\rangle = +1$ and $\langle K, N, M\rangle = -1$. Choose an arc γ from K to M that avoids L and N. Then $L \in C^+(\gamma)$ while $N \in C^-(\gamma)$. Let γ' be an arc from L to N, which we can choose to intersect γ only once and transversely. Then the initial segment of γ is on the negative side of γ' hence so is K. Therefore we have $\langle L, K, N\rangle = -1$ and $\langle K, L, N\rangle = +1$ as desired, and similarly $\langle L, M, N\rangle = +1$. □

2.2. For ends. We can now define a circular order on the ends of a generalized unbounded decomposition.

Definition 4.24. Let $\mathcal{C}$ be a collection of eventually disjoint unbounded continua in the plane, and let $\kappa_1, \kappa_2, \kappa_3, \ldots \in \mathcal{E}(\mathcal{C})$ be distinct ends. A bounded set $A \in P$ is said to *distinguish* the κ_i if $\kappa_1(A), \kappa_2(A), \kappa_3(A), \ldots$ are disjoint and mutually nonseparating.

We can always find a distinguishing set for a finite collection of ends. Indeed, if $\kappa_1, \kappa_2, ..., \kappa_n$ are distinct, one can choose a bounded open disc D' so that the $\kappa_1(D'), \kappa_2(D'), ..., \kappa_n(D')$ are disjoint. Then any disc $D \supset D'$ that intersects all of the $\kappa_i(D')$ distinguishes.

Definition 4.25. Let $\mathcal{C}$ be a collection of eventually disjoint unbounded continua in the plane P, and let $\kappa, \lambda, \mu \in \mathcal{E}(\mathcal{C})$ be distinct ends. Choose a disc D that distinguishes these ends and define

$$\langle \kappa, \lambda, \mu \rangle = \langle \kappa(D), \lambda(D), \mu(D) \rangle.$$

Proposition 4.26. *Let $\mathcal{C}$ be a collection of eventually disjoint unbounded continua in the plane P. Then $\langle \cdot, \cdot, \cdot \rangle$ defines a circular order on $\mathcal{E}(\mathcal{C})$.*

Proof. We only need to check that $\langle \kappa(D), \lambda(D), \mu(D) \rangle$ does not depend on the choice of D.

Let D and D' be bounded open discs that distinguish the ends $\kappa, \lambda, \mu \in \mathcal{E}(\mathcal{D})$. Let $K = \kappa(D)$, $L = \lambda(D)$, and $M = \mu(D)$ and define K', L', and M' similarly. Without loss of generality we may assume that $D \subset D'$, which implies that $K' \subset K$, etc.

Suppose that $\langle K, L, M \rangle = +1$. Let γ be an arc from K to M that avoids L. Then L is on the positive side of γ, hence so is L'. So we can find an arc l from γ to L' whose interior lies on the positive side of γ. Let γ' be an arc that runs from K' to the initial point of γ, follows γ, and then runs to M'. We can choose this so that the first and last segments avoid $L \cup l$. Then γ' is an arc from K' to M', and the fact that l is on the positive side of γ' exhibits that L' is as well. Therefore, $\langle K', L', M' \rangle = +1$ as desired. □

Remark 4.27. Suppose that K, L, and M are disjoint, mutually non-separating unbounded continua. Then $\langle K, L, M \rangle = \langle \kappa, \lambda, \mu \rangle$ for any ends $\kappa \in \mathcal{E}(K)$, $\lambda \in \mathcal{E}(L)$, and $\mu \in \mathcal{E}(M)$.

2.3. Properties of the circular order. Two pairs x, y and z, w of points in a circularly ordered set S are *linked* if either $z \in (x, y)$ and $w \in (y, x)$, or $w \in (x, y)$ and $z \in (y, x)$. Two subsets $A, B \subset S$ are linked if there are $x, y \in A$ and $z, w \in B$ that are linked. A subset $A \subset S$ *separates* the subsets $B, C \subset S$ if there are points $a, a' \in A$ such that $B \subset (a, a')$ and $C \subset (a', a)$. Note that this is not the same as topological separation in S with the order topology.

The separation properties of disjoint unbounded continua can be detected by their ends:

PROPOSITION 4.28. *Let $K, L \subset P$ be disjoint unbounded continua in the plane. Then $\mathcal{E}(K)$ and $\mathcal{E}(L)$ do not link in the canonical circular order on $\mathcal{E}(\{K, L\})$.*

Let $K, L, M \subset P$ be disjoint unbounded continua in the plane. Then K separates L from M if and only if $\mathcal{E}(K)$ separates $\mathcal{E}(L)$ from $\mathcal{E}(M)$ in the canonical circular order on $\mathcal{E}(\{K, L, M\})$.

PROOF. For the first statement, let $\kappa_1, \kappa_2 \in \mathcal{E}(K)$ and $\lambda_1, \lambda_2 \in \mathcal{E}(L)$. Let D be a bounded open disc distinguishing the $\kappa_i(D)$ and $\lambda_j(D)$ and choose an oriented arc γ from $\kappa_1(D)$ to $\kappa_2(D)$ that avoids L. Then $\lambda_1(D)$ and $\lambda_2(D)$ are on the same side of γ since they are both contained in L, and L is connected and disjoint from γ. This applies for any $\kappa_1, \kappa_2 \in \mathcal{E}(K)$ and $\lambda_1, \lambda_2 \in \mathcal{E}(L)$, so $\mathcal{E}(K)$ and $\mathcal{E}(L)$ do not link.

One direction in the second statement follows from the first. Suppose that K does not separate L from M. Then we can choose an arc γ from L to M disjoint from K, and let $X = L \cup \gamma \cup M$. The inclusion $(L \cup M) \hookrightarrow X$ induces an order-preserving bijection between $\mathcal{E}(L) \cup \mathcal{E}(M)$ and $\mathcal{E}(X)$. Now X and K are disjoint unbounded continua, so by the first statement their ends do not link, hence $\mathcal{E}(K)$ does not separate $\mathcal{E}(L)$ from $\mathcal{E}(M)$.

It remains to show that if K separates L from M then there are ends $\kappa_1, \kappa_2 \in \mathcal{E}(K)$, $\lambda \in \mathcal{E}(L)$, and $\mu \in \mathcal{E}(M)$ such that $\lambda \in (\kappa_1, \kappa_2)$ and $\mu \in (\kappa_2, \kappa_1)$. Let γ be an arc from L to M, and let γ' be the minimal connected sub-arc that contains $\gamma \cap K$. Set

$$K' := K \cup \gamma'$$

and

$$K'_{\pm} := K' \cap \overline{C^{\pm}(\gamma; L, M)}.$$

Note that the $K'_\pm$ are both closed and connected and neither one separates L from M. Also, $K'_+ \cap K'_-$ is connected, so in order for $K' = K'_+ \cup K'_-$ to separate L from M, both K'_+ and K'_- must be unbounded by [**36**], Theorem II.5.23. [4] So we can choose κ_1 to be any end of K'_+ and κ_2 to be any end of K'_-, where we're using the bijection $\mathcal{E}(K'_+ \cup K'_-) \to \mathcal{E}(K)$. □

LEMMA 4.29. *Let $K \subset P$ be an unbounded continuum in the plane, and let $\kappa_1, \kappa_2, \kappa_3 \in \mathcal{E}(K)$ be distinct ends, ordered counterclockwise. Suppose $U \subset P$ is a connected, bounded open set that distinguishes κ_1, κ_2, and κ_3. Then U distinguishes κ_2 from any end of K in (κ_3, κ_1).*

PROOF. Let $\kappa_4 \in (\kappa_3, \kappa_1)$ and suppose that U does not distinguish κ_2 from κ_4. That is, $\kappa_2(U) = \kappa_4(U)$. Then by Proposition 4.28, $\kappa_2(U)$ separates $\kappa_1(U)$ from $\kappa_3(U)$, contradicting the assumption that U distinguishes. □

REMARK 4.30 (Singular foliations). A singular foliation of the plane is an unbounded decomposition, where each decomposition element consists of finitely many properly embedded rays that share an initial point. It's easier to work with the circular order on the ends of a singular foliation, since ends are represented by properly embedded rays.

Let K, L, and M be properly embedded rays in the plane P. Choose an embedded circle $S \subset P$ that intersects all three, and orient S consistently with P. Let k, l, and m be, respectively, the last intersections of K, L, and M with S. Then k, l, m has a circular order inherited from S, and this determines the circular order of K, L, M.

3. Universal circles

We now have a canonical circular order on the ends of a generalized unbounded decomposition, and we'd like to form a universal circle using Construction 3.12. First, we'll need to know a couple of things about the topology of the set of ends.

PROPOSITION 4.31. *Let $\mathcal{C}$ be a collection of eventually disjoint unbounded continua in the plane. Then $\mathcal{E}(\mathcal{C})$ is separable.*

---

[4] If x and y are points of S^2 which are not separated by either of the closed sets A and B, and $A \cap B$ is connected, then x and y are not separated by $A \cup B$.

PROOF. Fix an exhaustion of the plane P by a nested sequence of bounded open discs D_i. That is, $D_{i-1} \subset D_i$ for each i and $\bigcup_i D_i = P$. For each i, consider

$$P^i := \bigcup_{K \in \mathcal{C}} K_{D_i},$$

where K_{D_i} is the union of the unbounded components of $K \setminus D_i$. Recall that a subspace of a separable space is separable. So P^i is separable for each i, and we can choose a countable set $\{p^i_j\}_{j=1}^{\infty}$ that is dense in P^i. For each i and j let $K^i_j \in \mathcal{C}$ be the element containing p^i_j, and choose an end κ^i_j so that $\kappa^i_j(D_i)$ is the component of $K^i_j \setminus D_i$ containing p^i_j. We will show that

$$\mathcal{E}' = \{\kappa^i_j\}_{i,j}$$

is dense in $\mathcal{E}(\mathcal{C})$.

Let $\kappa, \mu \in \mathcal{E}(\mathcal{C})$, and assume that (κ, μ) contains at least one end, λ. We will show that $\kappa^i_j \in (\kappa, \mu)$ for some i, j. Choose a bounded disc D that distinguishes κ, λ, and μ and set

$$K = \kappa(D),$$

$$M = \mu(D).$$

Let γ be an arc from K to M and choose i large enough so that D_i contains D and γ. Set $C^+ := C^+(\gamma; K, M)$. Now $C^+ \cap P^i$ is open as a subset of P^i, and it is nonempty because it contains $\lambda(D_i)$. So for some j there is a $p^i_j \in C^+$. The corresponding end κ^i_j is contained in (κ, μ), since $\kappa^i_j(D_i)$ is contained in C^+. $\square$

If our unbounded decomposition were actually a singular foliation, then we could work with the circular orders more concretely using embedded circles as in Remark 4.30. Suppose K, L, and M are properly embedded rays. Choose an embedded circle $S \subset P$ that intersects all three, and let $k, l, m \in S$ be the last intersections of K, L, M with S. If $\langle K, L, M \rangle = +1$, then the oriented sub-arc $(k, l) \subset S$ is outermost, in the sense that S never crosses between K and L on the positive side of (k, l). Similarly, (l, m) is the outermost sub-arc between L and M, and (m, n) is the outermost sub-arc between M and N.

In general, there's no natural way to define the last intersection between an unbounded continuum and a circle, but we can still find outermost sub-arcs. The first of the following two lemmas can be taken as a definition of outermost sub-arcs, and the second shows that they behave as expected. We'll use them in the proof of Lemma 4.34.

LEMMA 4.32. *Let $K, M \subset P$ be disjoint unbounded continua in the plane, and let $S \subset P$ be a positively oriented embedded circle that intersects both K and M. There is a unique sub-arc $\gamma_o \subset S$ running from K to M (with the orientation inherited from S) such that every component of $S \cap C^+(\gamma_o)$ has both ends in either K or M.*

We'll call γ_o the *outermost* sub-arc from K to M.

PROOF. The components of $S \setminus (K \cup M)$ form a countable collection of oriented open arcs. Let Γ be the collection of closures of these arcs. We will partition

$$\Gamma = \Gamma_K \cup \Gamma_M \cup \Gamma_{K,M} \cup \Gamma_{M,K}$$

where

$$\Gamma_K = \{\text{sub-arcs from } K \text{ to itself}\},$$

$$\Gamma_M = \{\text{sub-arcs from } M \text{ to itself}\},$$

$$\Gamma_{K,M} = \{\text{sub-arcs from } K \text{ to } M\},$$

and

$$\Gamma_{M,K} = \{\text{sub-arcs from } M \text{ to } K\}.$$

Let $\Gamma_s = \Gamma_{K,M} \cup \Gamma^-_{M,K}$, where the minus means to reverse the orientation of each arc (the "s" stands for "spanning"). As in Remark 4.21, Γ_s is endowed with a linear order. We will show that Γ_s has a maximal element, and that this maximal element lies in $\Gamma_{K,M}$ rather than $\Gamma^-_{M,K}$.

Suppose Γ_s has no maximal element. Then we can find a sequence $\gamma_0 < \gamma_1 < \gamma_2 ...$ of arcs in Γ_s that dominates Γ_s, in the sense that for every $\gamma \in \Gamma_s$ there exists j such that $\gamma < \gamma_j$. Take a Hausdorff convergent subsequence of the γ_i and let γ_∞ be the limit. Although γ_∞ is a sub-arc of S, it is not necessarily an element of Γ since $\mathring{\gamma}_\infty$ might intersect K or M. However there is at least one sub-arc $\gamma \subset \gamma_\infty$ from K to M or M to K. But then $\gamma \in \Gamma_s$, and $\gamma_i < \gamma$ for all i, a contradiction. Hence we can find a maximal element $\gamma_o \in \Gamma_s$.

We now turn to showing that $\gamma_o \in \Gamma_{K,M}$ rather than $\Gamma^-_{M,K}$, i.e. that γ_o runs from K to M with the orientation inherited from S. Enlarge K and M to

$$K' = K \cup (\bigcup_{\gamma \in \Gamma_K} \gamma)$$

and

$$M' = M \cup (\bigcup_{\gamma \in \Gamma_M} \gamma).$$

Note that K' and M' are still disjoint unbounded continua, so $C := C^+(\gamma_o; K', M')$ is unbounded by Lemma 4.17. Note that C is disjoint from S, so if γ_o were oriented from M to K then C would be contained in the bounded component of $P \setminus S$, which is impossible. $\square$

LEMMA 4.33. *Let $K, L, M, N \subset P$ be disjoint, mutually nonseparating unbounded continua in the plane, ordered counterclockwise. If $S \subset P$ is an embedded circle intersecting all four, then the outermost sub-arc from K to L is disjoint from the outermost sub-arc from M to N.*

PROOF. Suppose otherwise. Then without loss of generality we have an arc $\gamma : [0,1] \to P$ with an initial segment $\gamma_1 := \gamma([0,x])$ that is the outermost sub-arc from K to L, and a terminal segment $\gamma_2 := \gamma([y,1])$ that is the outermost sub-arc from M to N, and these two intersect, i.e. $y < x$.

Let's replace M with a tail M' that is disjoint from γ. For example, choose any $\mu \in \mathcal{E}(M)$ and let $M' = \mu(D)$ for a bounded open disc D that contains γ. By hypothesis, $\langle K, M, L\rangle = -1$, so M' is on the negative side of γ_1.

On the other hand, $\langle K, M, N\rangle = +1$ by hypothesis, and γ is an arc from K to N, so M' must be on the positive side of γ. But M' can't be on the positive side of γ and the negative side of $\gamma_1 \subset \gamma$. $\square$

PROPOSITION 4.34. *Let $\mathcal{C}$ be a collection of eventually disjoint unbounded continua in the plane. Then $\mathcal{E}(\mathcal{C})$ has countably many gaps.*

PROOF. Fix an exhaustion $\{D_i\}_{i=1}^{\infty}$ of P by nested bounded open discs and let $S_i^j = \partial D_i$ for all $i > j$ (j is a dummy variable, i.e. for each j we want to keep a copy of every circle outside of D_j). If (κ, λ) is a gap then let n be the first integer such that D_n distinguishes κ from λ, and let k be the first integer such that S_k^n intersects both $\kappa(D_n)$ and $\lambda(D_n)$. We will associate the gap (κ, λ) with the open interval $U_{\kappa,\lambda} \subset S_k^n$ whose closure is the outermost arc from $\kappa(D_n)$ to $\lambda(D_n)$.

The disjoint union

$$\bigcup_{i>j} S_i^j$$

is second countable, so any collection of disjoint open subsets is countable. Distinct gaps correspond to disjoint open intervals in $\bigcup S_i^j$ by Lemma 4.33, hence there are only countably many gaps. □

Applying Lemma 3.8, we have:

COROLLARY 4.35. *Let $\mathcal{C}$ be a collection of eventually disjoint unbounded continua in the plane. Then $\mathcal{E}(\mathcal{C})$ is 2nd countable.*

This fills a gap in the literature, as 2nd countability had not previously been shown.

CONSTRUCTION 4.36 (Universal circles). Let $\mathcal{D}$ be a generalized unbounded decomposition of the plane P, and suppose that $\mathcal{E}(\mathcal{D})$ is uncountable. Recall that this is immediate if $\mathcal{D}$ is a (non-generalized) unbounded decomposition by Lemma 4.3. The *universal circle* S_u^1 of $\mathcal{D}$ is simply the universal circle for $\mathcal{E}(\mathcal{D})$ that we built in Construction 3.12.

EXAMPLE 4.37. It is appealing to think that the ends of a decomposition element would map to a closed set in the universal circle, but this is not generally true. For example, the set in Figure 3 has no rightmost end.

4. The end compactification

Let's fix a generalized unbounded decomposition $\mathcal{D}$ of the plane P, and let S_u^1 be its universal circle. We will show that there is a natural topology on the set $\overline{P} := P \cup S_u^1$ that makes it into a closed disc with interior P and boundary S_u^1.

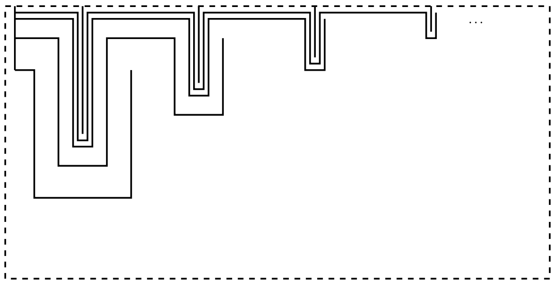

FIGURE 3. An unbounded continuum with no rightmost end.

DEFINITION 4.38. Let $\mathcal{D}$ be a generalized unbounded decomposition of the plane P. A *subordinate set* is a tail of an element of $\mathcal{D}$. That is, K is a subordinate set if there is a $\kappa \in \mathcal{E}(\mathcal{D})$ and a bounded open disc $D \subset P$ so that $K = \kappa(D)$.

If $K = \kappa(D)$ is a subordinate set, let K' be the decomposition element containing K, i.e. $K' = \kappa(\emptyset)$. Then $\mathcal{E}(K)$ is naturally identified with a subset of $\mathcal{E}(K')$.

The "peripheral sets" in the following definition will serve as neighborhoods in $\overline{P}$ of points on the boundary circle S^1_u. Here $\phi : \mathcal{E}(\mathcal{D}) \to S^1_u$ is the natural map that takes ends to their images in the universal circle.

DEFINITION 4.39. Let $\mathcal{D}$ be a generalized unbounded decomposition of the plane P. Fix subordinate sets K and L, and let γ be an arc from K to L. Let I be the maximal open interval in $S^1_u \setminus \phi(\mathcal{E}(K) \cup \mathcal{E}(L))$ that runs from $\phi(\mathcal{E}(K))$ to $\phi(\mathcal{E}(L))$. The *peripheral set* determined by K, L, and γ is defined to be

$$O(K, L, \gamma) := I \cup C^+(\gamma; K, L).$$

REMARK 4.40. It is possible that $O(K, L, \gamma) \cap S = \emptyset$, i.e. that the interval I is empty. For example, let K be the set in Example 4.37 and L its mirror image.

CONSTRUCTION 4.41 (End Compactification). Let $\mathcal{D}$ be a generalized unbounded decomposition of the plane P, and let S^1_u be its universal circle. The *end compactification* of P with respect to $\mathcal{D}$ is the set

$$\overline{P} := P \cup S^1_u$$

with the topology generated by open sets in P and peripheral sets.

We'll start by showing how to construct a peripheral set to particular specifications.

LEMMA 4.42. *Let $\mathcal{D}$ be an unbounded decomposition of the plane P with universal circle S^1_u. Let $a, b, c, d \in S^1_u$ be positively ordered. Then there is a peripheral set O such that $(b, c) \subset (O \cap S^1_u) \subset (a, d)$. In addition, O can be chosen to be disjoint from any bounded set $A \subset P$, and contained in any peripheral set O' as long as $O' \cap S^1_u \supset [b, c]$.*

PROOF. Ends are dense in S^1_u, so we can choose ends $\kappa_1, \kappa, \kappa_2, \lambda_1, \lambda, \lambda_2$ positively ordered such that the κ's have image in (a, b) and the λ's have image in (c, d). Let D be a bounded open disc that distinguishes the κ's and λ's. This way, $K = \kappa(D)$ and $L = \lambda(D)$ have ends in (a, b) and (c, d) respectively by Lemma 4.29. [5] Connect K and L with an arc γ and set

$$O = O(K, L, \gamma).$$

If we're provided with a bounded set A and a peripheral set O', start by shrinking (a, d) so that it's contained in $O' \cap S^1_u$. Suppose $O' = O(K', L', \gamma')$, where $K' = \kappa'(D_{K'})$ and $L' = \lambda'(D_{L'})$. Choose D so that it contains $D_{K'}$, $D_{L'}$, A, and γ', and choose γ to lie inside of O'. Then O will be contained in O' and disjoint from A as desired. □

Now we can get a handle on the topology on $\overline{P}$.

LEMMA 4.43. *Let $\overline{P}$ be the end compactification of the plane P with respect to the generalized unbounded decomposition $\mathcal{D}$. Then*

(1) *the open sets in P and peripheral sets form a basis for $\overline{P}$,*
(2) *$\overline{P}$ is 1st countable,*
(3) *the inclusion maps $P \hookrightarrow \overline{P}$ and $S^1_u \hookrightarrow \overline{P}$ are homeomorphisms onto their images, and*

[5] If we hadn't chosen the "helper ends" κ_1, κ_2 and λ_1, λ_2, then K and L would still have some ends in (a, b) and (c, d) respectively, but they might also have ends outside of these intervals. The helper ends ensure that we choose D large enough to cut off these extra ends.

(4) $\overline{P}$ *is compact.*

PROOF. **(1) – (2):** Let $p \in S^1_u$. We'll start by constructing a neighborhood basis for p that consists of peripheral sets.

Fix an exhaustion of the plane by bounded open discs D_i, and a sequence of open intervals $(a_i, b_i) \subset S^1_u$ such that $[a_{i+1}, b_{i+1}] \subset (a_i, b_i)$ for each i, and $\bigcap_i (a_i, b_i) = p$. Using Lemma 4.42 we can find a sequence of peripheral sets $O_i = O(K_i, L_i, \gamma_i)$ such that $[a_{i+1}, b_{i+1}] \subset (O_i \cap S^1_u) \subset (a_i, b_i)$, $O_{i+1} \subset O_i$, and $O_i \cap D_i = \emptyset$ for each i.

If $O' = O(K', L', \gamma')$ is any peripheral set containing p, then the O_i are eventually contained in O'. Indeed, simply choose i so that $(O_i \cap S^1_u) \subset (O' \cap S^1_u)$ and $D_i \supset (D_{K'} \cup \gamma' \cup D_{L'})$, where $D_{K'}$ and $D_{L'}$ are the discs determining the subordinate sets K' and L'.

To prove (1) and (2) it suffices to note that if U and V are peripheral and $p \in U \cap V \cap S^1_u$ then for large enough i, $p \in O_i \subset (U \cap V)$.

(3): It is clear that $P \hookrightarrow \overline{P}$ is a homeomorphism onto its image. For $S^1_u \hookrightarrow \overline{P}$, it is clear that the preimage of an open set is open. Conversely, if $U \subset S^1_u$ is open then for each $p \in U$ we can find a peripheral set O with $p \in (O \cap S^1_u) \subset U$.

(4): Let $\mathcal{U}$ be an open cover of $\overline{P}$. Since S^1_u is compact we can find a finite subcollection $\{U_1, U_2, ..., U_n\} \subset \mathcal{U}$ that covers S^1_u. After reordering and taking a refinement we can assume that each U_i is a peripheral set $U_i = O(K_i, L_i, \gamma_i)$ such that $K_i, L_{i-1}, K_{i+1}, L_i$ is positively ordered for all i (mod $n+1$). Further, we can assume that the arcs γ_i do not intersect the K_j and L_k for any i, j, k, and that γ_i intersects γ_{i+1} only once, transversely, for all i. See Figure 4.

We can concatenate the sub-arcs between intersections of the γ_i to form a simple closed curve γ. Points on the negative side of γ are contained in at least one of the U_i and the positive side of γ is compact. So we can find a finite sub-cover of $\mathcal{U}$ that covers this compact piece and the rest is covered by the U_i. □

The end compactification behaves well when we take closures of decomposition elements.

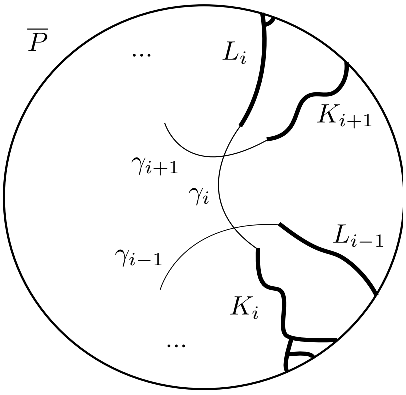

FIGURE 4. A refinement covering S.

LEMMA 4.44. *Let $\overline{P}$ be the end compactification of the plane P with respect to the generalized unbounded decomposition $\mathcal{D}$. If $K \in \mathcal{D}$ then*

$$cl_{\overline{P}}(K) = K \cup cl_{S^1_u}(\phi(\mathcal{E}(K)),$$

where ϕ is the natural map from $\mathcal{E}(\mathcal{D})$ to S^1_u.

PROOF. It is clear that $cl_{\overline{P}}(K) \cap P = K$, so let $p \in S^1_u$. If p is not in $cl_{S^1_u}(\phi(\mathcal{E}(K)))$ then we can find a peripheral set containing p that does not intersect K, so $p \notin cl_{\overline{P}}(K)$. □

We're ready for the punchline of the chapter.

THEOREM 4.45. *Let $\mathcal{D}$ be a generalized unbounded decomposition of the plane P. The end compactification $\overline{P}$ is homeomorphic to a closed disc with interior P and boundary S^1_u. Any homeomorphism of P that preserves $\mathcal{D}$ extends to a homeomorphism of $\overline{P}$.*

PROOF. The second statement follows from the fact that the image of a peripheral set under a $\mathcal{D}$-preserving homeomorphism is again a peripheral set. To prove that $\overline{P}$ is a closed disc we will use the following characterization of the disc. An arc γ with endpoints a, b is said to *span* a set $S \subset X$ if $a, b \in S$ and $\mathring{\gamma} \subset X \setminus S$.

THEOREM 4.46 (Zippin, [**36**], Theorem II.5.1). *Let X be a connected, compact, 1st countable Hausdorff space. Suppose that X contains a 1-sphere*

S such that some arc in X spans S, every arc that spans S separates X, and no closed proper subset of an arc spanning S separates X. Then X is homeomorphic to a closed 2-disc with boundary S.

It is clear that $\overline{P}$ is connected and Hausdorff and we have already shown that it is 1st countable and compact. Let's check the remaining conditions.

Existence of a spanning arc: Fix $p \in S^1_u$. Let $\{O_i\}_{i=1}^{\infty}$ be nested sequence of peripheral sets such that $\cap_i O_i = p$, and $O_i = O(K_i, L_i, \gamma_i)$ for each i. Choose a point $p_i \in \gamma_i$ for each i and let c_i be an arc that connects p_i to p_{i+1} with interior in $O_i \setminus O_{i+1}$. The concatenation of these arcs is a proper ray whose closure is an arc from c_0 to p. Construct two such rays and connect their endpoints with an arc in P.

Spanning arcs separate: Let $\gamma : [0,1] \to \overline{P}$ be an arc spanning S^1_u with initial point a and terminal point b. Note that $\mathring{\gamma}$ is a properly embedded curve in P so it separates P into two unbounded discs by the Jordan curve theorem. We will find subordinate sets K and L that are separated by γ.

Let $O_a = O(K_a, L_a, \gamma_a)$ and $O_b = O(K_b, L_b, \gamma_b)$ be disjoint peripheral sets containing a and b respectively, and let $(x, y) = O_a \cap S^1_u$ and $(z, w) = O_b \cap S^1_u$. There is a t_0 such that $\gamma([0, t_0]) \subset O_a$, since otherwise γ would intersect either K_a or L_a in some proper infinite sequence, implying that γ accumulates on some point in either $cl_{\overline{P}}(K_a)$ or $cl_{\overline{P}}(L_a)$. Similarly, there is a t_1 such that $\gamma([t_1, 1]) \subset U_b$.

Therefore, the compact sub-arc $\gamma' = \gamma([t_0, t_1])$ has the property that $\gamma \setminus \gamma'$ is contained in $O_a \cup O_b$. So we can find a subordinate set K disjoint from γ with ends in (y, z): just choose a subordinate set whose ends lie in (y, z) that is disjoint from $K_a, L_a, \gamma_a, K_b, L_b, \gamma_b$ and γ'. Similarly, find a subordinate set L with ends in (w, x) that is disjoint from γ.

Suppose that γ did not separate K from L. Then we could find an arc μ from K to L disjoint from γ. But then $b \in O(K, L, \mu)$ while $a \notin O(K, L, \mu)$, implying that γ must intersect either K, L, or μ, a contradiction. So γ must separate K from L.

No subset of a spanning arc separates: Let γ be a spanning arc and let $\gamma' = \gamma \setminus \{x\}$ for some point x.

Suppose $x \in S^1_u$. Then γ' separates P into two discs, one containing a subordinate set K and one containing a subordinate set L (as above). Note that $S^1_u \setminus \gamma'$ is connected, and $cl_{\overline{P}}K$ and $cl_{\overline{P}}L$ both intersect $S^1_u \setminus \gamma'$. Hence $\overline{P} \setminus \gamma'$ is connected.

If on the other hand $x \in P$, then $P \setminus \gamma'$ is connected. But $\overline{P} \setminus \gamma'$ is contained in $cl_{\overline{P}}(P \setminus \gamma')$, hence connected. $\square$

CHAPTER 5

Quasigeodesic Flows

Now that the technical machinery of Chapters 3 and 4 is out of the way, we can reprise our main topic. Throughout this chapter, M will be a closed hyperbolic 3-manifold with a quasigeodesic flow $\mathfrak{F}$. The flowspace P is homeomorphic to the plane, and it comes with unbounded decompositions $\mathcal{D}^+$ and $\mathcal{D}^-$. The union of these,

$$\mathcal{D} := \mathcal{D}^+ \cup \mathcal{D}^-,$$

is a generalized unbounded decomposition. We'll write

$$\mathcal{E}^\pm := \mathcal{E}(\mathcal{D}^\pm)$$

and

$$\mathcal{E} := \mathcal{E}(\mathcal{D}).$$

The fundamental group $\pi_1(M)$ acts on $\mathbb{H}^3$ by deck transformations, and this induces an action on P via the quotient map

$$\pi : \mathbb{H}^3 \to P.$$

The decompositions $\mathcal{D}^\pm$ and $\mathcal{D}$ are preserved by this action

Calegari's *universal circle* is defined to be the universal circle S^1_u built from $\mathcal{D}$ (Construction 4.36), which inherits an action of $\pi_1(M)$ via the natural map

$$\phi : \mathcal{E}(\mathcal{D}) \to S^1_u.$$

This completes the proof of Theorem 2.18.

The flowspace P can be compactified to a disc $\overline{P}$ using the universal circle (Construction 4.41), and the action of $\pi_1(M)$ on P extends to $\overline{P}$ (Theorem 4.45. The restriction of this action to the boundary is the usual universal circle action.

This completes one version of our first main theorem.

Compactification Theorem. Let $\mathfrak{F}$ be a quasigeodesic flow on a closed hyperbolic 3-manifold M, and let S^1_u be Calegari's universal circle. Then $\overline{P} = P \cup S^1_u$ has a natural topology making it into a closed disc with interior P and boundary S^1_u. The action of $\pi_1(M)$ on P extends to $\overline{P}$ and restricts to the universal circle action on $\partial\overline{P}$.

We will define another universal circle S^1_v in Section 2 and show that this theorem holds for it too.

Note that one could build universal circles $S^{1,+}_u$ and $S^{1,-}_u$ using only $\mathcal{D}^+$ and $\mathcal{D}^-$. In fact, one could even use these universal circles to build compactifications $\overline{P}^+$ and $\overline{P}^-$. However, the compactification $\overline{P}$ is more useful, as it behaves well with respect to $\mathcal{D}^+$ and $\mathcal{D}^-$ simultaneously.

1. Extending the endpoint maps

In this section we'll show that the endpoint maps

$$e^{\pm} : P \to S^2_\infty$$

extend continuously to $\overline{P}$.

Given a point $x \in S^1_u$, we might hope to find a sequence of (say, positive) decomposition elements that nest down to x. That is, we'd like to find a sequence of $K_i \in \mathcal{D}^+$ that converge to x in the Hausdorff sense, arranged so that K_i separates K_{i-1} from K_{i+1} for each i. If such a sequence exists, then it turns out that $e^+(K_i)$ converges to some point $p \in S^2_\infty$, and we can extend e^+ by setting $e^+(x) = p$. Any sequence of points $x_i \in P$ approaching x gets trapped between the K_i, and it follows that $e^+(x_i)$ converges to p (this is Case 2 in the proof of Lemma 5.3) Therefore, we've found the unique way to extend e^+ continuously to x.

In general, not all points in S^1_u will have such a nice structure. One way to prove continuous extension is to break the analysis into several cases, depending on how the decompositions $\mathcal{D}^{\pm}$ look near each point in S^1_u. The proof that follows is less direct. However, it has the benefit of brevity, and we obtain a concise understanding of $\mathcal{D}^{\pm}$ along the way.

1.1. Sequences in the flowspace. We'll start by understanding how sequences of flowlines look in $\mathbb{H}^3$ as we go to infinity in P.

There's no way to extend hyperbolic metric on $\mathbb{H}^3$ to S^2_∞, since the distance between points will be infinite. However, if we think about $\mathbb{H}^3$ in the unit ball model, then S^2_∞ is the unit sphere, and we can endow $\overline{\mathbb{H}}^3 = \mathbb{H}^3 \cup S^2_\infty$ with the Euclidean metric. This will be useful when we consider sets that are "close to infinity."

REMARK 5.1 (A note on notation). If A is a subset of either $\mathbb{H}^3$ or $\overline{\mathbb{H}}^3$, we'll write $\overline{A}$ for the closure of A in $\overline{\mathbb{H}}^3$. Similarly, if B is a subset of either P or $\overline{P}$, then $\overline{B}$ is the closure of B in $\overline{P}$. In all other cases we'll use $\mathrm{cl}_X(C)$ to mean the closure of the set C in the space X.

Recall that if $p \in P$, then $\pi^{-1}(p) \subset \mathbb{H}^3$ is the flowline corresponding to p, and the closure of a flowline is just the flowline plus its endpoints. In other words,

$$\overline{\pi^{-1}(p)} = e^-(p) \cup \pi^{-1}(p) \cup e^+(p).$$

The following lemma says that sets of flowlines that are close to infinity in P have *uniformly* small diameter in $\overline{\mathbb{H}}^3$.

LEMMA 5.2. *For each $\epsilon > 0$ there is a compact set $A \subset P$ such that for each $p \in P \setminus A$, the diameter of $\pi^{-1}(p)$ is at most ϵ in the Euclidean metric.*

PROOF. Recall that there is a constant C such that every flowline has Hasudorff distance at most C from the geodesic between its endpoints (in the *hyperbolic* metric). Geodesics with endpoints close together in the Euclidean metric have small diameter in the Euclidean metric. So given ϵ there is a constant ϵ' such that all flowlines whose endpoints are ϵ'-close have Euclidean diameter at most ϵ. Let $B \subset (S^2_\infty \times S^2_\infty) \setminus \Delta$ be the set of pairs (p, q) with $d(p, q) \geq \epsilon'$, and let $A = (e^+ \times e^-)^{-1}(B)$, which is compact by Lemma 2.14. $\square$

The following lemma is one of the main technical ingredients in our proof of continuous extension. Recall that if $(A_i)_{i=1}^\infty$ is a sequence of sets in a topological space X then $\limsup A_i$ is the set of points $x \in X$ such that every neighborhood of x intersects infinitely many of the A_i (see [**17**]).

Lemma 5.3. *Let $(U_i)_{i=1}^{\infty}$ be a sequence of disjoint open sets in P with frontiers $A_i = Fr\, U_i$. Then*

$$\limsup e^+(U_i) \subseteq \limsup e^+(A_i),$$

and similarly for e^-.

Proof. Let $(x_i)_{i=1}^{\infty}$ be a sequence of points with $x_i \in U_i$ for each i, where we might have already taken a subsequence of the U_i. It suffices to show that if $e^+(x_i)$ converges, then $\lim e^+(x_i) \in \limsup e^+(A_i)$.

Fix such a sequence and set $p := \lim e^+(x_i)$ and $Q := \limsup e^+(A_i)$. We need to show that $p \in Q$.

Case 1: Suppose that infinitely many of the x_i are contained in a bounded set. Then after taking a subsequence we can assume that the x_i converge to some point x, and note that $e^+(x) = p$ (see Figure 1). Then since A_i separates x_i from x_{i-1} for each i, we can find points $y_i \in A_i$ that also converges to x. But then $\lim e^+(y_i) = p$, so $p \in Q$ as desired.

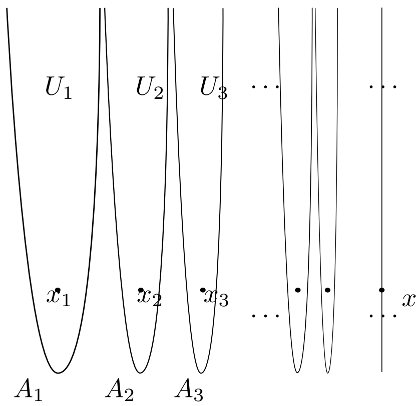

Figure 1. Case 1.

So assume that the x_i escape to infinity.

Case 2: Suppose for the moment that the A_i escape to infinity, i.e. that they are eventually disjoint from every bounded set. If $p \notin Q$ then we can find disjoint open sets U and V in $\overline{\mathbb{H}}^3$ that contain p and Q respectively. By Lemma 5.2, $\pi^{-1}(x_i)$ and $\pi^{-1}(A_i)$ are eventually contained in U and V respectively (see Figure 2). But this contradicts the fact that $\pi^{-1}(A_i)$ separates $\pi^{-1}(x_i)$ from $\pi^{-1}(x_{i-1})$ for each i. So we must have $p \in Q$.

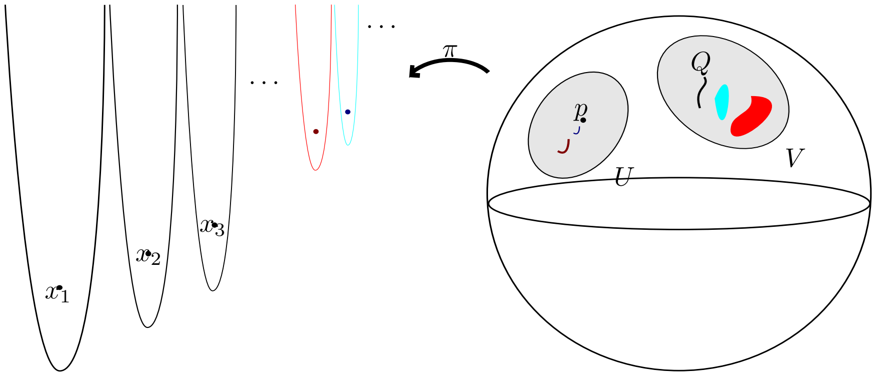

FIGURE 2. Case 2.

Case 3: Now suppose that the A_i do not escape to infinity. We will cut a compact set out of each U_i to make the A_i escape to infinity. In general this might change $Q = \limsup A_i$, but by choosing carefully we can ensure that Q shrinks.

Choose an exhaustion $(D_j)_{j=1}^{\infty}$ of P by compact discs, each of which intersects *all* of the U_i, where we might have to pass to a subsequence. The argument in Case 1 shows that that for a bounded sequence of points $y_i \in U_i$, $e^+(y_i)$ is eventually close to Q. Note that for fixed j, the sets $D_j \cap U_i$ are all contained in the bounded set D_j. So for fixed j, $e^+(D_j \cap U_i)$ is eventually close to Q.

It follows that we can find a diagonal sequence: For each j choose an integer $i(j)$ sufficiently large so that (a) $x_{i(j)}$ is outside of D_j and such that (b) $\limsup_{j\to\infty} e^+(D_j \cap U_{i(j)}) \subset Q$. For each j, let $y_j := x_{i(j)}$, let V_j be the component of $P \setminus (D_j \cup U_{i(j)})$ that contains y_j, and let $B_j := \mathrm{Fr}\ V_j$. Also, set $Q' := \limsup_{j\to\infty} B_j$. Observe that since each point in B_j is contained in either $A_{i(j)}$ or $(D_j \cap U_{i(j)})$, property (b) ensures that $Q' \subset Q$. See Figure 3.

The B_j escape to infinity, so by case 2 $\lim e^+(y_j) \subset Q'$ and therefore $p \in Q$ as desired.

□

COROLLARY 5.4. *Let $B, C \subset S^2_\infty$ be disjoint compact sets. Then $(e^+)^{-1}(B)$ intersects at most finitely many components of $P \setminus (e^+)^{-1}(C)$, and similarly for e^-.*

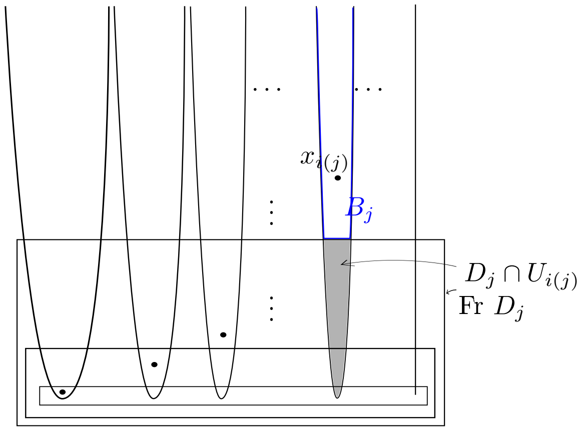

FIGURE 3. Case 3: Fixing the A_i.

PROOF. Otherwise we would have a sequence of points b_i in $(e^+)^{-1}(B)$, each of which is contained in a different component of $P \setminus (e^+)^{-1}(C)$. By the preceding lemma, $\limsup e^+(b_i) \in C$. But $\limsup e^+(b_i) \in B$, so this contradicts the assumption that B and C are disjoint. □

1.2. Unions of decomposition elements. Recall (Lemma 4.44) that if $K \in \mathcal{D}$ then $\overline{K} = K \cup \overline{\phi(\mathcal{E}(K))}$.

LEMMA 5.5. *For each point $x \in S^1_u$, there are at most countably many decomposition elements $K \in \mathcal{D}$ whose closures intersect x.*

PROOF. By Lemma 4.44, each decomposition element $K \in \mathcal{D}$ whose closure intersects x either has an end at x (i.e. there is a $\kappa \in \mathcal{E}(K)$ with $\phi(\kappa) = x$) or has ends approaching x (i.e. there are $\kappa_i \in \mathcal{E}(K)$ with $\phi(\kappa_i) \to x$ as $i \to \infty$). In the construction of the universal circle, at most countably many ends are collapsed into one point, so there are at most countably many decomposition elements with an end at x. Therefore, it suffices to show that there are at most countably many decomposition elements with ends approaching x. In fact there are at most four.

We will say that a positive decomposition element $K \in \mathcal{D}^+$ has ends approaching x in the *positive direction* if there are $\kappa_i \in \mathcal{E}(K)$ such that $\phi(\kappa_i) \to x$ as $i \to \infty$, and the $\kappa_1, \kappa_2, \kappa_3, \ldots$ are ordered counterclockwise. There can be at most one decomposition element with ends approaching x in the positive direction, since if there were two then their ends would link, implying that they intersect (Proposition 4.28). Similarly, there is at most one negative decomposition element with ends approaching x in the positive direction, and at most one positive and one negative decomposition element with ends approaching x in the negative direction. □

LEMMA 5.6. *Let A be a closed subset of the plane that separates the points x and y. Then some component of A separates x and y.*

PROOF. Let U be the component of $P \setminus A$ that contains x, and let V be the component of $P \setminus \overline{U}$ that contains y. Then Fr V separates x and y. It is connected because the plane satisfies the Brouwer property (see [**36**], Definition I.4.1): if M is a closed, connected subset of the plane P and V is a component of $P \setminus M$ then Fr V is connected. Therefore, the component of A containing Fr V separates x and y. □

Recall that if $(A_i)_{i=1}^{\infty}$ is a sequence of sets in a topological space, then $\liminf A_i$ is the set of points x such that every neighborhood of x intersects all but finitely many of the A_i. If $\liminf A_i$ is nontrivial and agrees with $\limsup A_i$, then the A_i are said to converge in the Hausdorff sense, and we write $\lim A_i = \limsup A_i = \liminf A_i$ (see [**17**]). If the A_i live in the plane, and infinitely many intersect some compact set, then we can always find a Hausdorff convergent subsequence.

LEMMA 5.7. *Let $A \subset P$ be a closed, connected set that is a union of positive decomposition elements and let U be a component of $P \setminus A$. Then Fr U is contained in a single decomposition element.*

PROOF. Note that Fr U is connected by the Brouwer property, so it suffices to show that $e^+(\text{Fr } U)$ is a point.

Suppose that $e^+(\text{Fr } U)$ is not a point. Then it is infinite, so we can choose three points $k, l, m \in \text{Fr } U$ whose images under e^+ are distinct. Let $K, L, M \subset A$ be the positive decomposition elements that contain k, l, m respectively. These are mutually nonseparating, so by Proposition 4.28 their ends are mutually nonseparating.

Assume that K, L, M are positively ordered, and if N is a decomposition element contained in U then M, N, M is positively ordered. Let $n_i \in P$ be a sequence of points in U that limit to some point in L, and for each i let N_i be the decomposition element that contains n_i. Note that infinitely many of the N_i intersect a compact set (for example, any arc from K to M) after taking a subsequence they converge to some set $N_\infty \subset P$. See Figure 4.

Observe that $N_\infty \cup L$ separates K from M. Indeed, if γ is any arc from K to M that avoids L then the n_i are eventually on the positive side of γ, and the N_i have ends between $\mathcal{E}(M)$ and $\mathcal{E}(K)$, so the N_i eventually intersect γ. By Lemma 5.6, some component B of $N_\infty \cup L$ separates K from M. But then $\mathcal{E}(B)$ separates $\mathcal{E}(K)$ from $\mathcal{E}(M)$, a contradiction. Therefore, $e^+(\text{Fr } U)$ is a single point as desired.

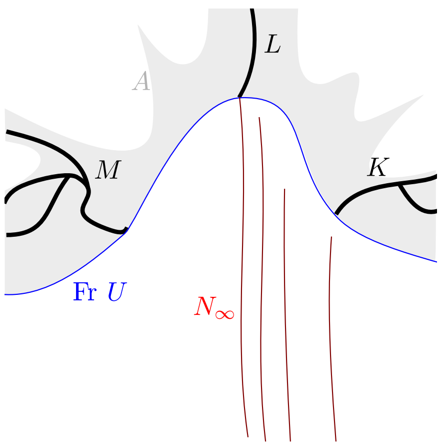

FIGURE 4

□

Combining the last two lemmas, and using the fact that if $X \subset P$ separates $x, y \in P$ then so does Fr X [1]:

COROLLARY 5.8. *Let $A \subset P$ be a closed union of positive decomposition elements. If A separates the points $x, y \in P$ then some decomposition element $K \subset A$ separates x and y.*

1.3. Extending the endpoint maps. We can now prove the main theorem of this section.

[1]Let U be the component of $P \setminus X$ containing x. Any connected set that intersects both x and y must intersect both U and $P \setminus U$, hence it must intersect Fr $U \subset$ Fr X.

Continuous Extension Theorem. Let $\mathfrak{F}$ be a quasigeodesic flow on a closed hyperbolic 3-manifold M. The endpoint maps $e^{\pm} : P \to S^2_\infty$ admit unique continuous extensions to $\overline{P}$, and e^+ agrees with e^- on the boundary.

Proof. Fix $x \in S^1_u$, and let $(x_i)_{i=1}^\infty$ be a sequence of points in P that converges to x. After taking a subsequence we can assume that $e^+(x_i) \to p$ as $i \to \infty$ for some $p \in S^2_\infty$. Set $e^+(x) = p$.

To see that this is well-defined, suppose we have two such sequences $(x_i)_{i=1}^\infty$ and $(y_i)_{i=1}^\infty$ converging to x, with $e^+(x_i) \to p$ and $e^+(x_i) \to q$ for $p \neq q$. Let $A' \subset S^2_\infty$ be a simple closed curve separating p and q. Then $A = (e^+)^{-1}(A')$ separates $\{x_i\}_{i=I}^\infty$ from $\{y_i\}_{i=I}^\infty$ for large enough I. By Lemma 5.3, the x_i are eventually contained in a single component of $P \setminus A$, as are the y_i. So by Corollary 5.8, some decomposition element $K \in \mathcal{D}^+$ separates $\{x_i\}_{i=I}^\infty$ from $\{y_i\}_{i=I}^\infty$. Therefore, $\overline{K}$ separates $\{x_i\}_{i=I}^\infty$ from $\{y_i\}_{i=I}^\infty$ in $\overline{P}$. This implies that $x \in \overline{K}$.

There are uncountably many disjoint simple closed curves separating p and q, so the argument above would produce uncountably many distinct decomposition elements K with $x \in \overline{K}$. This would contradict Lemma 5.5, so the extension of e^+ to $\overline{P}$ is well-defined. The same applies for e^-, and continuity and uniqueness are immediate.

To see that e^+ and e^- agree on S^1_u, let $x \in S^1_u$ and choose a sequence of points $x_i \in P$ converging to x. By Lemma 5.2, the distance between $e^+(x_i)$ and $e^-(x_i)$ approaches 0 as $i \to \infty$, so $\lim e^+(x_i) = \lim e^-(y_i)$. □

From now on,

$$e^{\pm} : \overline{P} \to S^1_u$$

will denote the extended endpoint maps, and

$$e : S^1_u \to S^2_\infty$$

their restriction to the universal circle.

We'll restate the continuous extension theorem in a way that matches the Cannon-Thurston theorem. Let $\mathfrak{F}$ be a quasigeodesic flow on a closed hyperbolic 3-manifold M. A *complete transversal* to the lifted flow is an embedded 2-manifold $N \subset \mathbb{H}^3$ that is transverse to the foliation by flowlines.

Any complete transversal can be identified with the orbit space P via the quotient map $\pi : \mathbb{H}^3 \to P$, so we can think of N as a section $\psi : P \to \mathbb{H}^3$ of π. Therefore, we can compactify N to a closed disc $\overline{N}$ by adding the universal circle S^1_u, and the embedding $i : N \hookrightarrow \mathbb{H}^3$ extends to the boundary by setting $i = e$ on S^1_u. This is continuous since the diameter of flowlines goes to zero as we go to the boundary (Lemma 5.2). We have:

THEOREM 5.9 (Generalized Cannon-Thurston Theorem). *Let $\mathfrak{F}$ be a quasigeodesic flow on a closed hyperbolic 3-manifold M, and let $N \subset \mathbb{H}^3$ be a complete transversal to the lifted flow on the universal cover. There is a natural compactification of N as a closed disc $\overline{N} = N \cup S^1_u$ that inherits a π_1 action, and the embedding $i : N \hookrightarrow \mathbb{H}^3$ has a unique continuous extension $i : N \cup S^1_u \to \mathbb{H}^3 \cup S^2_\infty$. The restriction of i to S^1_u is a π_1-equivariant space-filling curve in S^2_∞.*

2. A more convenient circle

Calegari's universal circle may contain more information than we need. For clarity, we'll momentarily add subscripts to some of our spaces. Let $\overline{P}_u$ be the compactification of P built from S^1_u. Ends of decomposition elements correspond to points in the universal circle via the map $\phi_u : \mathcal{E} \to S^1_u$. The endpoint maps $e^\pm : P \to S^2_\infty$ extend to $e^\pm_u : \overline{P}_u \to S^2_\infty$, and these restrict to $e_u : S^1_u \to S^2_\infty$ on the boundary.

If $p \in S^2_\infty$, it is possible that some components of $e_u^{-1}(p) \subset S^1_u$ are intervals. Let S^1_v be the quotient of S^1_u after collapsing each component of $e_u^{-1}(p)$ to a point, for every $p \in S^2_\infty$. The action of $\pi_1(M)$ on S^1_u descends to an action on S^1_v. Composing the quotient with ϕ_u, we obtain a map $\phi_v : \mathcal{E} \to S^1_v$. The map e_u factors through the quotient to form $e_v : S^1_v \to S^2_\infty$.

We can use Construction 4.41 and Theorem 4.45 to produce a compactification $\overline{P}_v = P \cup S^1_v$ by simply replacing S^1_u with S^1_v and ϕ_u with ϕ_v in the proofs. We can extend the endpoint maps by setting $e^\pm_v : \overline{P} \to S^2_\infty$ to agree with e_v on the boundary.

Note that $e_v^{-1}(p)$ is now totally disconnected for any $p \in S^2_\infty$. Therefore, $\overline{P}_v$ is the smallest compactification of P for which the endpoint maps extend to the boundary. This is exactly the property that we want from the universal circle. In addition, the fact that point preimages are totally disconnected will be useful when we study the dynamics of $\pi_1(M)$ on S^1_v.

In the sequel, we will drop the subscripts and let context determine which version of $\overline{P}$, $e^{\pm}$, e, and ϕ we're referring to.

In many cases, $S^1_u = S^1_v$. For example, if $\mathfrak{F}$ is both quasigeodesic and pseudo-Anosov then the endpoints of leaves of $\widetilde{\lambda}^{\pm}$ determine a dense lamination of S^1_u, and it follows that point preimages are already totally disconnected. In fact, as far as we know, it is possible that $S^1_u = S^1_v$ for every quasigeodesic flow.

2.1. Upper-semicontinuous decompositions. We can also see that S^1_v compactifies P without referring back to the construction of the end-compactification. Let X be a topological space, and let $\mathfrak{D}$ be a decomposition of X into closed, connected subsets. The decomposition is called *upper semicontinuous* if for any open set $U \subset X$, the union of decomposition elements contained in U is open.

Moore [**24**] proved that if $\mathfrak{D}$ is an upper semicontinuous decomposition of the sphere S^2 into non-separating sets, then the quotient obtained by collapsing each decomposition element to a point is still homeomorphic to S^2.

Double the disc $\overline{P}_u$ to obtain a sphere S^2. Let $\mathfrak{D}$ be the decomposition whose elements are the components of $e_u^{-1}(p)$ for each $p \in S^2_\infty$, and the points in each copy of P. This is clearly upper semicontinuous, so the quotient is still S^2, consisting of two copies of $\overline{P}_v$ glued along S^1_v.

3. Properties of the decompositions

We'll collect some useful observations concerning the regularity of $\mathcal{D}^{\pm}$ and $\mathcal{D}$. The statements in this section hold for both S^1_u and S^1_v.

LEMMA 5.10. *Let $p \in S^2_\infty$. Then $(e^+)^{-1}(p) \cup (e^-)^{-1}(p) \subset \overline{P}$ is connected.*

PROOF. Let $A_i \subset \overline{\mathbb{H}}^3$ be a nested sequence of closed horoballs centered at p. Set $B_i = \pi(A_i \cap \mathbb{H}^3)$ for each i. The B_i's are connected subsets of P, so their closures $\overline{B_i}$ are connected, compact subset of the compact space $\overline{P}$. Therefore,

$$B = \bigcap_i \overline{B_i}$$

is connected.

Suppose $x \in B$. Then there is a sequence of points $x_i \in B_i$ approaching x (after taking subsequences). The corresponding flowlines $\pi^{-1}(x_i)$ intersect the A_i, so after taking a subsequence either $e^+(x_i)$ or $e^-(x_i)$ approaches p. Therefore, either $e^+(x) = p$ or $e^-(x) = p$ as desired.

If $x \notin B$ then the flowline $\pi^{-1}(x)$ is eventually disjoint from the horoballs B_i, so neither $e^+(x)$ nor $e^-(x)$ is p. □

COROLLARY 5.11. *Let $p \in S^2_\infty$, and suppose that no flowline has p as an endpoint. Then $e^{-1}(p)$ is a point.*

PROOF. Saying that no flowline has p as an endpoint is the same as saying that $(e^+)^{-1}(p) \cup (e^-)^{-1}(p)$ is contained in S^1_u, i.e. $(e^+)^{-1}(p) \cup (e^-)^{-1}(p) = e^{-1}(p)$. By the preceding lemma, $e^{-1}(p)$ is connected, so it is either a point or an interval. Images of ends are dense in S^1_u, so if $e^{-1}(p)$ were an interval, the closure of some $K \in \mathcal{D}$ would intersect it. But then $(e^+)^{-1}(p) \cup (e^-)^{-1}(p)$, which was supposed to be empty, would contain K. □

4. Dynamics on the universal circle

We can now analyze the dynamics of the fundamental group action on the universal circle.

Suppose $f : S^1 \to S^1$ is a homeomorphism. Let $x \in S^1$ be an isolated fixed point, and choose an interval $x \in I \subset S^1$ that isolates x. Then x is called *attracting* if $f(I) \subset I$, *repelling* if $f(I) \supset I$, and *indifferent* otherwise.

Each $g \in \pi_1(M)$ acts on S^2_∞ with two fixed points in an attracting-repelling pair. Let a_g and r_g be the attracting and repelling fixed points, respectively. Set

$$A_g = e_v^{-1}(a_g) \subset S^1_v,$$

and

$$R_g = e_v^{-1}(r_g) \subset S^1_v.$$

Note that A_g and R_g are disjoint, closed, and invariant under g.

LEMMA 5.12. *If $g \in \pi_1(M)$ has a fixed point on S^1_v, then it has at least two.*

PROOF. Suppose $p \in S^1_v$ is the only fixed point of g. The iterates of g take any point to p. Therefore, A_g and R_g must both contain p since they are closed and invariant. This is impossible since they are disjoint. $\square$

Let F_g be the set of fixed points of g on S^1_v. If $x \in S^1_v$ is fixed by g, then $e(x) \in S^2_\infty$ is fixed by g, so

$$F_g \subset A_g \cup R_g.$$

We'll partition F_g into FA_g and FR_g, where

$$FA_g = F_g \cap A_g,$$

and

$$FR_g = F_g \cap R_g.$$

Suppose we have distinct $x, y \in FR_g$. Note that R_g is totally disconnected, so there is some point $z \in (x, y)$ that is not in R_g. Iterating by g^n takes $e(z)$ towards r_g, so after taking a subsequence, $g^n(z)$ approaches some point $z' \in (x, y)$ with $e(z') = a_g$. A similar statement applies replacing R by A and g by g^{-1}. In other words, between any two points in FR_g there is a point in FA_g, and between any two point in FA_g there is a point in FR_g.

It follows that FA_g is finite. Otherwise we could find a sequence of $x_i \in FA_g$ that converges monotonically to some point x, with $x \in FA_g$. For each i let $y_i \in (x_i, x_{i+1})$ be in FR_g, and observe that the y_i also converge to x. This implies that $x \in FR_g$, a contradiction. Similarly, FR_g is finite.

Therefore, F_g is finite, and fixed points alternate between FA_g and FR_g. Note that for any $z \in S^1_v$ that is not a fixed point, $\lim_{i\to\infty} g^n(e(z)) = a_g$, and $\lim_{i\to-\infty} g^n(e(z)) = r_g$. Therefore, $\lim_{i\to\infty} g^n(z) \in FA_g$, and $\lim_{i\to-\infty} g^n(z) \in FR_g$. Hence the points in FA_g are attracting, and the points in FR_g are repelling. In summary:

Pseudo-Anosov Dynamics Theorem. Let $\mathfrak{F}$ be a quasigeodesic flow on a closed hyperbolic 3-manifold M. Then each $g \in \pi_1(M)$ acts on S^1_v with an even number of fixed points (possibly zero) that are alternately attracting and repelling.

5. Closed orbits

We now turn to the question of whether quasigeodesic flows have closed orbits. The general question of finding closed orbits in 3-manifold dates back to 1950, when Seifert asked whether every nonsingular flow on the 3-sphere has a closed orbit [**29**]. Schweizer provided a counterexample in 1974, and showed that every homotopy class of nonsingular flows on a 3-manifold has a C^1 representative without closed orbits. Schweitzer's examples have since been generalized to smooth [**22**] and volume-preserving flows [**21**].

On the other hand, Taubes' 2007 proof of the 3-dimensional Weinstein conjecture shows that Reeb flows on closed 3-manifolds do have closed orbits [**31**]. Reeb flows are *geodesible*, i.e. there is a Riemannian metric in which the flowlines are geodesics. Complementary to this result, though by different methods, Rechtman showed in 2010 that the only geodesible real analytic flows on closed 3-manifolds that contain no closed orbits are on torus bundles over the circle with reducible monodromy [**27**].

Geodesibility is a local condition, and furthermore one that is not stable under perturbations. By contrast, quasigeodesity is a macroscopic condition, and it is a stable condition under C^0 perturbations when the ambient 3-manifold is hyperbolic.

Calegari conjectured in 2006 that quasigeodesic flows on closed hyperbolic 3-manifolds should all have closed orbits. In this section we will provide evidence for this conjecture.

5.1. Möbius-like groups.

A Möbius transformation is a conformal automorphism of the plane that preserves the unit circle. The group of orientation-preserving Möbius transformations is isomorphic to $PSL(2,\mathbb{R})$. Each Möbius transformation is determined by its action on the unit circle, so we can think of $PSL(2,\mathbb{R})$ as a subgroup of $\mathrm{Homeo}^+(S^1)$, the orientation-preserving homeomorphisms of the circle.

An orientation-preserving Möbius transformations is classified as either elliptic, parabolic, or hyperbolic by whether it has no fixed points, one indifferent fixed point, or two fixed points in an attracting-repelling pair.

A group $\Gamma < \mathrm{Homeo}(S^1)$ is called *Möbius-like* if each $g \in \Gamma \setminus \{id\}$ is conjugate to a Möbius transformation. It is called *hyperbolic Möbius-like* if each $g \in \Gamma$ is conjugate to a hyperbolic Möbius transformation. It is called *Möbius* if the whole group is conjugate into $PSL(2,\mathbb{R})$.

Any homeomorphism of S^1 with two fixed points in an attracting-repelling pair is conjugate to a hyperbolic Möbius transformation, so a group $\Gamma < \mathrm{Homeo}(S^1)$ is hyperbolic Möbius-like if and only if each $g \in \Gamma$ has exactly two fixed points, of which one is attracting and the other repelling.

There was a long-standing question whether all Möbius-like groups are conjugate into $PSL(2,\mathbb{R})$. This was shown to be false by Kovacevic [**20**]; however, her counterexamples are still the only ones known.

In fact, work of Casson-Jungreis [**7**], Gabai [**14**] and Tukia [**34**] shows that a group acting on the circle is conjugate into $PSL(2,\mathbb{R})$ if and only if satisfies a certain dynamical criterion.

Definition 5.13. Let $(g_i)_{i=1}^{\infty}$ be a sequence of homeomorphisms of S^1. Then (g_i) is a *convergence sequence* if there are $a, b \in S^1$ such that g_i converges to the constant map to a uniformly on compact subsets of $S^1 \setminus b$. A group Γ acting on the circle is a *convergence group* if every sequence of elements has a subsequence that is convergence.

Theorem 5.14 (Casson-Jungreis, Gabai, Tukia). *A group $\Gamma < Homeo(S^1)$ is Möbius if and only if it is a convergence group.*

We will show:

Möbius-like Theorem. Let $\mathfrak{F}$ be a quasigeodesic flow on a closed hyperbolic 3-manifold M. Suppose that $\mathfrak{F}$ has no closed orbits. Then $\pi_1(M)$ acts on the universal circle S^1_v as a hyperbolic Möbius-like group.

In contrast:

CONJUGACY THEOREM. Let $\mathfrak{F}$ be a quasigeodesic flow on a closed hyperbolic 3-manifold M. Suppose that $\mathfrak{F}$ has no closed orbits. Then the action of $\pi_1(M)$ on S^1_v is not Möbius.

None of the known examples of Möbius-like groups that are not Möbius are 3-manifold groups. This provides evidence for Calegari's conjecture that every quasigeodesic flow on a closed hyperbolic 3-manifold has closed orbits.

The Möbius-like Theorem follows immediately from the finer statement below.

CLOSED ORBITS THEOREM. Let $\mathfrak{F}$ be a quasigeodesic flow on a closed hyperbolic 3-manifold M, and let $g \in \pi_1(M)$. Then $\mathfrak{F}$ has a closed orbit in the free homotopy class represented by g, unless the action of g on S^1_u is conjugate to a hyperbolic Möbius transformation.

These three theorems all hold with S^1_v instead of S^1_u.

We'll need the following lemma. Each $g \in \pi_1(M)$ acts on S^2_∞ as a loxodromic isometry, so it has an attracting fixed point a_g and a repelling fixed point r_g.

LEMMA 5.15. *Let $\mathfrak{F}$ be a quasigeodesic flow on a closed hyperbolic 3-manifold M and let $g \in \pi_1(M)$. Then $\mathfrak{F}$ has a closed orbit in the free homotopy class represented by g if and only if $\widetilde{\mathfrak{F}}$ has a flowline with an endpoint at either a_g or r_g.*

PROOF. Suppose that $\mathfrak{F}$ has a closed orbit in the free homotopy class of g. Then some lift γ of this orbit to the universal cover is fixed by g, hence the endpoints of γ are a_g and r_g.

Conversely, suppose that some flowline γ in $\widetilde{\mathfrak{F}}$ has an endpoint at r_g (replace g by g^{-1} if it has a fixed point at a_g). Then the endpoints of $g^n(\gamma)$ approach a_g and r_g as $n \to \infty$, so the $g^n(\gamma)$ all intersect a compact set in $\mathbb{H}^3$. In other words, the corresponding point $\pi(\gamma) \subset P$ in the flowspace has a bounded forward orbit under g. The Brouwer plane translation theorem [2]

[2] If f is a homeomorphism of the plane P, and some point $x \in P$ has a bounded forward orbit, then f fixes some point in P. See **[12]**.

implies that g must fix a point in P, and this fixed point corresponds to a closed orbit in the free homotopy class of g. □

PROOF OF THE CLOSED ORBITS THEOREM. Suppose that $\mathfrak{F}$ has no closed orbits in the free homotopy class of g. Note that $A = e^{-1}(a_g)$ and $B = e^{-1}(r_g)$ are single points by Corollary 5.11 (which applies to both S^1_u and S^1_v). But these are invariant, so they're fixed points. Every other point in S^1_u is moved, so g has exactly two fixed points.

For every point $x \in S^1_u \setminus \{A \cup B\}$, $g^n(x)$ must limit to A since $g^n(e(x))$ limits to a_g. Therefore, A is attracting and B is repelling. □

PROOF OF THE CONJUGACY THEOREM. Suppose that the action of $\pi_1(M)$ on S^1_v (or S^1_u) were conjugate to $G < PSL(2,\mathbb{R})$. Note that G cannot be discrete, since then it would be isomorphic to a surface group. A closed subgroup of a Lie group is a Lie subgroup by the Cartan theorem ([**35**], Theorem 3.42), so the closure $\overline{G}$ is a Lie subgroup. All proper Lie subgroups of $PSL(2,\mathbb{R})$ are virtually solvable (see below), so $\overline{G} = PSL(2,\mathbb{R})$. But every element of $PSL(2,\mathbb{R})$ with no fixed points has a neighborhood consisting of elements with no fixed points, so some element of G has no fixed points. This contradicts the Möbius-Like Theorem.

Let $H < PSL(2,\mathbb{R})$ be a connected, proper Lie subgroup. We'll show that it's solvable.[3] The exponential map is surjective on $PSL(2,\mathbb{R})$, so H has dimension 0, 1, or 2. If it has dimension 0 or 1 then it's trivially solvable. If it has dimension 2, let $X, Y \in H$ be a basis. Then $[X,Y] = aX + bY$ for some a, b. If either $a = 0$ or $b = 0$ then H is solvable. Otherwise, $[X,Y] = aX + bY = a(X + \frac{b}{a}Y)$, so after a change of basis we're in the previous case and H is solvable. □

[3] Thanks to Danny Calegari for pointing this out.

Bibliography

[1] M. R. Bridson and A. Haefliger, *Metric Spaces of Non-Positive Curvature*, Grundlehren der mathematischen Wissenschafen, vol. 319, Springer-Verlag, Berlin, Germany, 1999.

[2] D. Calegari and N. Dunfield, *Laminations and groups of homeomorphisms of the circle*, Invent. Math. **152** (2003), no. 1, 149–204.

[3] D. Calegari, *Universal circles for quasigeodesic flows*, Geom. Topol. **10** (2006), 2271–2298.

[4] D. Calegari, *Foliations and the geometry of 3-manifolds*, Oxford Mathematical Monographs, Oxford University Press, Oxford, 2007.

[5] J. W. Cannon, W. J. Floyd, R. Kenyon, and W. R. Parry, *Flavors of Geometry: Hyperbolic Geometry*, MSRI Publications **31** (1997).

[6] J.W. Cannon and W.P. Thurston, *Group invariant Peano curves*, Geom. Topol. **11** (2007), 1315–1355.

[7] A. Casson and D. Jungreis, *Convergence groups and Seifert fibered 3-manifolds*, Invent. Math. **118** (1994), no. 3, 441–456.

[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), 1–110.

[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*, preprint.

[12] J. Franks, *A new proof of the Brouwer plane translation theorem*, Ergod. Th. & Dynam. Sys. **12**, (1992), 217–226.

[13] H. Freudenthal, *Über die Enden topologischer Räume und Gruppen*, Math. Z. **33**, (1931), no. 1, 692–713.

[14] D. Gabai, *Convergence groups are Fuchsian groups*, Ann. of Math. **136** (1992), no. 3, 447–510.

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

[16] M. Gromov, *Hyperbolic groups*, Essays in group theory, Math. Sci. Res. Inst. Publ., vol. 8, Springer, New York, 1987, 75–263.

[17] J.G. Hocking and G.S. Young, *Topology* Addison-Wesley Series in Mathematics, Addison-Wesley Publishing Company, Reading, MA, 1961.

[18] M. Kapovich, *Hyperbolic manifolds and discrete groups*, Progress in Mathematics, vol. 183 (2001), Birkhauser Boston Inc., Boston, MA.

[19] A. S. Kechris, *Classical Descriptive Set Theory*, Graduate Texts in Mathematics, vol. 156, Springer-Verlag, New York, NY, 1995.

[20] N. Kovacevic, *Examples of Möbius-like groups which are not Möbius groups*, Trans. Amer. Math. Soc. **351** (1999), no. 12, 4823–4835.

[21] G. Kuperberg, *A volume-preserving counterexample to the Seifert conjecture*, Comm. Math. Helv. **71** (1996), no. 1, 70–79.

[22] K. Kuperberg, *A smooth counterexample to the Seifert conjecture*, Ann. of Math. (2) **140** (1994), no. 3, 723–732.

[23] E. E. Moise, *Geometric Topology in Dimensions 2 and 3*, Springer-Verlag, New York, 1977.

[24] R. L. Moore, *Concerning upper semi-continuous collections of continua*, Trans. Amer. Math. Soc **27** (1925), no. 4, 416–428.

[25] L. Mosher, *Examples of quasi-geodesic flows on hyperbolic 3-manifolds*, Topology '90 (Columbus, OH, 1990), Ohio State Univ. Math. Res. Inst. Publ., vol. 1, de Gruyter, Berlin, 1992, 227–241.

[26] L. Mosher, *Laminations and flows transverse to finite depth foliations, Part I: Branched surfaces and dynamics*, draft.

[27] A. Rechtman, *Existence of periodic orbits for geodesible vector fields on closed 3-manifolds*, Ergod. Th. & Dynam. Sys. **30** (2010), no. 6, 1817–1841.

[28] P. A. Schweitzer, *Counterexamples to the Seifert conjecture and opening closed leaves of foliations*, Ann. of Math. **100** (1974), no.2, 386–400.

[29] H. Seifert, *Closed integral curves in 3-space and isotopic two-dimensional deformations*, Proc. Amer. Math. Soc. **1**, (1950), 287–302.

[30] W. Sierpinski, *Un théoréme sur les continus*, Tohoku Mathematical Journal, **13**, (1918), 300–303.

[31] C. H. Taubes, *The Seiberg-Witten equations and the Weinstein conjecture*, Geom. Top. **11** (2007), 2117–2202.

[32] W. P. Thurston, *Three-manifolds, foliations and circles II: the transverse asymptotic geometry of foliations*, preprint.

[33] W. P. Thurston, *The Geometry and Topology of Three-Manifolds*, electronic notes, (2002).

[34] P. Tukia, *Homeomorphic conjugates of Fuchsian groups*, J. Reine Angew. Math., **391** (1988), 1–54.

[35] F. W. Warner, *Foundations of Differentiable Manifolds and Lie Groups*, Graduate Texts in Mathematics, v. 94, Springer-Verlag, New York, 1983.

[36] R.L. Wilder, *Topology of manifolds*, American Mathematical Society Colloquium Publications, vol. 32, American Mathematical Society, Providence, RI, 1979.

[37] A. Zeghib, *Sur les feuilletages géodésiques continus des variétés hyperboliques*, Invent. Math. **114** (1993), no. 1, 193–206.