

Quasigeodesic flows and sphere-filling curves

STEVEN FRANKEL

Given a closed hyperbolic 3-manifold M with a quasigeodesic flow, we construct a π_1 -equivariant sphere-filling curve in the boundary of hyperbolic space. Specifically, we show that any complete transversal P to the lifted flow on \mathbb{H}^3 has a natural compactification to a closed disc that inherits a π_1 -action. The embedding $P \hookrightarrow \mathbb{H}^3$ extends continuously to the compactification, and restricts to a surjective π_1 -equivariant map $\partial P \rightarrow \partial\mathbb{H}^3$ on the boundary. This generalizes the Cannon–Thurston theorem, which produces such group-invariant space-filling curves for fibered hyperbolic 3-manifolds.

57M60; 57M50, 37C27

1 Introduction

1.1 Background and motivation

Cannon and Thurston [3] showed that a hyperbolic 3-manifold fibering over the circle gives rise to a π_1 -equivariant sphere-filling curve. The lift of a fiber to the universal cover is an embedded plane $P \hookrightarrow \mathbb{H}^3$ which can be compactified to a disc $P \cup S_u^1$ in such a way that the embedding extends continuously to a map $P \cup S_u^1 \rightarrow \mathbb{H}^3 \cup S_\infty^2$ and the restriction $S_u^1 \rightarrow S_\infty^2$ to the boundary is surjective and π_1 -equivariant.

The proof depends on the interaction between the geometry of the surface and the geometry of a certain transverse flow. A surface bundle over the circle is determined up to homeomorphism by its monodromy. If the resulting manifold is hyperbolic, then the monodromy can be taken to be pseudo-Anosov, so its suspension gives rise to a pseudo-Anosov flow. It turns out that this flow is quasigeodesic in the sense that each lifted flowline is a quasi-isometrically embedded line in \mathbb{H}^3 , and one can use this fact to see that leaves of the stable/unstable foliations are quasi-isometrically embedded hyperbolic planes. Cannon and Thurston used the separation properties of these planes to prove the continuous extension property for the lifted surface fibers.

The main goal of this paper is to show that the continuous extension property can be established in far greater generality and assuming far less initial structure. For any quasigeodesic flow on a closed hyperbolic 3-manifold, the lifted flow on the universal

cover admits a complete planar transversal, ie a planar transversal that intersects every orbit. We show for *any* quasigeodesic flow that *any* complete transversal has the continuous extension property; it can be compactified to a closed disk whose boundary circle maps continuously and surjectively to the sphere at infinity of hyperbolic 3-space. And although the transversal itself will typically have trivial stabilizer in π_1 , its boundary admits a natural π_1 -action and the continuous surjection to S_∞^2 is equivariant with respect to this action. Thus one obtains group invariant sphere-filling curves directly from the quasigeodesic property without appealing to a surface bundle structure (which does not in general exist).

Our continuous extension theorem thus generalizes the Cannon–Thurston theorem to arbitrary quasigeodesic flows. Any finite depth taut foliation of a hyperbolic 3-manifold admits an almost transverse quasigeodesic flow, so such flows can be found in abundance (see Fenley and Mosher [5]). Quasigeodesic flows are found even on nonfibered manifolds. Our argument is fundamentally different from the Cannon–Thurston argument, which makes use of 2-dimensional objects — the surface fibers and the 2-dimensional leaves of stable/unstable foliations — to reason about topology in 3 dimensions. By contrast, we work in the orbit space of the flow (a 2-manifold) and need to deal with much more complicated objects, namely arbitrary closed continua, which are derived indirectly from the global geometry and dynamics of the transversal, and which hint at a kind of “pseudo-Anosov structure in the large”. In fact, our main theorem may be thought of as a key step in Calegari’s program to show that every quasigeodesic flow on a hyperbolic 3-manifold is homotopic (through quasigeodesic flows) to a quasigeodesic *pseudo-Anosov* flow. Finally, our main theorem strengthens the main result of Frankel’s [6] by weakening the hypothesis under which a quasigeodesic flow can be shown to have a closed orbit.

1.2 Statement of results

A quasigeodesic is a quasi-isometric embedding of \mathbb{R} in a metric space X . That is, distances as measured in X and \mathbb{R} are comparable on the large scale up to a multiplicative constant. (See Section 2.1 for a precise definition.)

A nonsingular flow on a manifold is said to be quasigeodesic if the lift of each flowline to the universal cover is a quasigeodesic. Fix a quasigeodesic flow Φ on a closed hyperbolic 3-manifold M . We can lift it to a flow $\tilde{\Phi}$ on the universal cover $\tilde{M} \simeq \mathbb{H}^3$. The flowspace P is the image of the quotient map $\pi: \mathbb{H}^3 \rightarrow P$ that collapses each lifted flowline to a point. This space is a plane and it comes with natural maps $e^\pm: P \rightarrow S_\infty^2$ that send each orbit to its forward/backward limit in the sphere at infinity. There is a natural compactification \bar{P} of the orbit space as a closed disc, called the *end compactification*, where the boundary is Calegari’s universal circle S_u^1 .

Continuous extension theorem *Let Φ be a quasigeodesic flow on a closed hyperbolic 3-manifold M . There are unique continuous extensions of the endpoint maps e^\pm to the compactified orbit space \bar{P} , and e^+ agrees with e^- on the boundary.*

A complete transversal to the flow $\tilde{\Phi}$ can be identified with the orbit space P , so we'll think of it as a section $P \rightarrow S_\infty^2$ of the quotient map π . Thus it is immediate that any transversal can be compactified by adding the universal circle S_u^1 . It is easy to show that we can extend ψ by setting its values on the boundary equal to those of the e^\pm maps, and our generalization of the Cannon–Thurston theorem follows.

Cannon–Thurston theorem *Let $\tilde{\Phi}$ be a quasigeodesic flow on a closed hyperbolic 3-manifold M . A complete transversal P to the lifted flow has a natural compactification to a closed disc $P \cup S_u^1$. The action of π_1 on P extends to $P \cup S_u^1$ and the embedding $P \hookrightarrow \mathbb{H}^3$ extends to a map $P \cup S_u^1 \rightarrow \mathbb{H}^3 \cup S_\infty^2$ that is equivariant with respect to this action. The restriction to the boundary is a π_1 -equivariant sphere-filling curve $S_u^1 \rightarrow S_\infty^2$.*

A group of homeomorphisms of the circle is said to be *Möbius-like* if each element is topologically conjugate to a Möbius transformation, ie an element of $\text{PSL}(2, \mathbb{R})$ acting in the standard way on \mathbb{RP}^1 . It is said to be *hyperbolic Möbius-like* if each element is conjugate to a hyperbolic Möbius transformation. We can use the extension theorem to improve on the main result in [6].

Closed orbits theorem *Let Φ be a quasigeodesic flow on a closed hyperbolic 3-manifold M . Suppose Φ has no closed orbits. Then $\pi_1(M)$ acts on the universal circle S_u^1 as a hyperbolic Möbius-like group.*

Previously, we could only conclude that such a flow would have a “rotationless Möbius-like” universal circle action, meaning that each element is conjugate to either a hyperbolic or a parabolic Möbius transformation.

1.3 Future directions

Given a closed hyperbolic 3-manifold M with a quasigeodesic flow Φ , take two copies \bar{P}^+ and \bar{P}^- of the compactified orbit space and glue them together along the boundary. We'll call the resulting space the *universal sphere* S_u^2 . Since the e^+ and e^- maps agree on the boundary one can define a map

$$e: S_u^2 \rightarrow S_\infty^2$$

that agrees with e^+ on \bar{P}^+ and e^- on \bar{P}^- .

In addition, one can fill in the universal sphere to form a closed ball B_u^3 with a 1–dimensional foliation \mathfrak{F}_u by connecting each point in P^+ to the corresponding point in P^- with a line. This parallels Cannon and Thurston’s construction for fibered hyperbolic manifolds in [3] and Fenley’s “model precompactification” for hyperbolic manifolds with a pseudo-Anosov flow in [4].

There is a natural decomposition of S_u^2 by point preimages $\{e^{-1}(p)\}$. The fundamental group acts on B_u^3 preserving this decomposition and the foliation \mathfrak{F}_u . After collapsing each decomposition element to a point, one recovers the *topological* structure of the fundamental group acting on $\mathbb{H}^3 \cup S_\infty^2$ and the foliation by lifted flowlines.

A goal of Calegari’s program is to show that one can deduce the pseudo-Anosov property of such a flow from the dynamics on the universal sphere. We would like to isotope a quasigeodesic flow in such a way that each decomposition element, restricted to P^+ or P^- , is a union of leaves of a singular foliation with “pseudo-Anosov dynamics” and conclude that the resulting flow is both quasigeodesic and pseudo-Anosov.

Acknowledgements The arguments in this paper have gone through several iterations, each much longer and more technical than the present, and I am indebted to Danny Calegari for his careful and patient critique of each one. I would also like to thank André de Carvalho, Sergio Fenley, Jake Rasmussen, Henry Wilton, Dongping Zhuang and the referee for helpful discussions and corrections.

2 Quasigeodesic flows and sphere-filling curves

2.1 Quasigeodesic flows

Fix a closed hyperbolic 3–manifold M with a quasigeodesic flow Φ .

We recall that this means that the lift $\gamma: \mathbb{R} \rightarrow \mathbb{H}^3$ of each flowline in the universal cover satisfies

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

for some constants $k, \epsilon > 0$, for all $x, y \in \mathbb{R}$. We’ll not work directly with this metric property but with its topological consequences, reviewed below. See [6, Sections 2 and 3] or [2] for a more detailed discussion.

We can lift Φ to a flow $\tilde{\Phi}$ on the universal cover $\tilde{M} \simeq \mathbb{H}^3$. Each flowline of $\tilde{\Phi}$ has well-defined and distinct positive and negative endpoints in $S_\infty^2 = \partial\mathbb{H}^3$. Furthermore there is a constant C such that each flowline has Hausdorff distance at most C from the geodesic with the same endpoints; see [1, pages 399–404].

The *flowspace* P is the image of the map

$$\pi: \mathbb{H}^3 \rightarrow P$$

that collapses each flowline of $\tilde{\Phi}$ to a point. The flowspace is homeomorphic to the plane [2, Theorem 3.12], and the action of $\pi_1(M)$ on \mathbb{H}^3 by deck transformations descends to an action on P . The maps

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

that send each point to the positive/negative endpoint of the corresponding flowline are continuous and have dense image; see [2, Lemmas 4.3 and 4.4]. Since the endpoints of each flowline are distinct, $e^+(p) \neq e^-(p)$ for every $p \in P$.

2.2 Sequences in P

If we think about \mathbb{H}^3 in the unit ball model, we can endow $\overline{\mathbb{H}^3} = \mathbb{H}^3 \cup S_\infty^2$ with the Euclidean metric. This will be useful when considering sets that are “close to infinity”.

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

Lemma 2.1 *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 If $B \in \mathbb{H}^3$ is a large compact ball, then each geodesic not intersecting B has small diameter in the Euclidean metric. Since each flowline has Hausdorff distance at most C from a geodesic with the same endpoints, the same applies to flowlines.

Given $\epsilon > 0$, choose a compact ball $B \in \mathbb{H}^3$ such that each flowline not intersecting B has diameter less than ϵ and set $A = \pi(B)$. □

In particular, points close to infinity in P correspond to flowlines with endpoints uniformly close together, ie for each $\epsilon > 0$ there is a compact set $A \subset P$ such that, for each $p \in P \setminus A$, the distance between $e^+(p)$ and $e^-(p)$ is at most ϵ .

Recall that if $(A_i)_{i=1}^\infty$ is a sequence of sets in a topological space, then $\limsup A_i$ is the set of points x such that every neighborhood of x intersects infinitely many of the A_i .

If Y is a subset of a topological space X , then the frontier of Y is defined by

$$\text{Fr } Y := \overline{Y} \cap \overline{X \setminus Y}.$$

Lemma 2.2 Let $(U_i)_{i=1}^\infty$ be a sequence of disjoint open sets in P with frontiers $A_i = \text{Fr } U_i$. Then

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

Proof It suffices to show that if $(x_i \in U_i)_{i=1}^\infty$ is a sequence of points such that $e^+(x_i)$ converges, where we might have already replaced the U_i with a subsequence, 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 , noting 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 converge to x . But then $\lim e^+(y_i) = p$, so $p \in Q$ as desired.




Figure 1: Case 1 in the proof of Lemma 2.2

So assume that the x_i escape to infinity.

Case 2 Suppose for the moment that the A_i escape to infinity, ie that they are eventually disjoint from every bounded set. If $p \notin Q$, then we can find disjoint open sets U and V in \mathbb{H}^3 that contain p and Q , respectively. Since sets in P close to infinity correspond to small sets in S_∞^2 (Lemma 2.1), it follows that $\pi^{-1}(x_i)$ 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 . Thus $p \in Q$.

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 $Q = \limsup e^+(A_i)$, but by doing it carefully we can ensure that Q shrinks.




Figure 2: Case 2 in the proof of Lemma 2.2

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

Then we can find a “diagonal sequence”: For each j choose an integer $i(j)$ sufficiently large that (a) $x_{i(j)}$ is outside of D_j and such that (b) $\limsup_{j \rightarrow \infty} e^+(D_j \cap U_{i(j)}) \subset Q$. Now 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 $B_j := \text{Fr } V_j$. Also, set $Q' := \limsup_{j \rightarrow \infty} B_j$ and 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.




Figure 3: Case 3 in the proof of Lemma 2.2. Fixing the A_i .

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

Corollary 2.3 *Let $B, C \subset S_\infty^2$ be disjoint compact sets. Then $(e^+)^{-1}(B)$ intersects at most finitely many components of $P \setminus (e^+)^{-1}(C)$.*

Proof Otherwise we could find a sequence of points b_i in $(e^+)^{-1}(B)$ each of which is contained in a different component of $P \setminus (e^+)^{-1}(C)$. Then $\limsup e^+(b_i) \in C$ by the preceding lemma. But $\limsup e^+(b_i) \in B$, so $B \cap C \neq \emptyset$, a contradiction. \square

2.3 Closures of decomposition elements

Recall from [6, Section 3] that the *positive decomposition* of P is defined as

$$\mathcal{D}^+ := \{\text{connected components of } (e^+)^{-1}(p) \mid p \in S_\infty^2\}.$$

Each decomposition element $K \in \mathcal{D}^+$ is an *unbounded continuum* in P , ie a noncompact, closed and connected set. The action of $\pi_1(M)$ on P permutes the elements of \mathcal{D}^+ . The *negative decomposition* \mathcal{D}^- is defined similarly using the e^- map, and we set $\mathcal{D} := \mathcal{D}^+ \cup \mathcal{D}^-$.

An unbounded continuum K in the plane has a naturally defined set of *ends* $\mathcal{E}(K)$, where an end $\kappa \in \mathcal{E}(K)$ is a map such that for each bounded set $D \in P$, $\kappa(D)$ is an unbounded component of $K \setminus D$ and if $D \subset D'$ then $\kappa(D') \subset \kappa(D)$.

If $K \in \mathcal{D}^+$ is a positive decomposition element and $L \in \mathcal{D}^-$ is a negative decomposition element, then $K \cap L$ is compact; see [6, Lemma 3.3].

The set of *ends*

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

comes with a natural circular order. For a triple $\kappa, \lambda, \mu \in \mathcal{D}$, choose a bounded open disc $D \subset P$ such that $\kappa(D)$, $\lambda(D)$ and $\mu(D)$ are disjoint and mutually nonseparating (ie no one separates the other two) and let γ be an oriented arc from $\kappa(D)$ to $\mu(D)$. Then κ, λ, μ is positively ordered if $\lambda(D)$ is on the positive side of γ and negatively ordered otherwise. If $\kappa, \mu \in \mathcal{E}$, then the *open interval* (κ, μ) is the set of $\lambda \in \mathcal{E}$ such that κ, λ, μ is positively ordered.

The action of $\pi_1(M)$ on P permutes the elements of \mathcal{E} preserving the circular order.

There is a natural *order completion* $\tilde{\mathcal{E}}$ that is homeomorphic to a closed subset of a circle, and such that \mathcal{E} is dense in $\tilde{\mathcal{E}}$. We can collapse complementary intervals in $\tilde{\mathcal{E}}$ to form the universal circle S_u^1 , and $\tilde{\mathcal{E}}$ surjects onto this; see [6, Construction 7.3]. We

denote by $\phi: \mathcal{E} \rightarrow S_u^1$ the composition of this inclusion and surjection. Both $\phi(\mathcal{E}^+)$ and $\phi(\mathcal{E}^-)$ are dense in S_u^1 , where \mathcal{E}^+ and \mathcal{E}^- are the ends that come from positive and negative decomposition elements respectively, and there are at most countably many ends in the preimage of each point. The action of $\pi_1(M)$ on \mathcal{E} extends to $\tilde{\mathcal{E}}$ and induces an action on S_u^1 by homeomorphisms.

The spaces P and S_u^1 glue together to form the *end compactification* \bar{P} , which is homeomorphic to a closed disc with interior P and boundary S_u^1 . The actions of $\pi_1(M)$ on P and S_u^1 match up to form an action on \bar{P} ; see [6, Theorem 7.9].

We refer the reader to [6] for a more complete discussion of ends, circular orders and universal circles. In order to prove the extension theorem we need to work out the details of how decomposition elements look with respect to the end compactification. In particular, distinct positive/negative decomposition elements are disjoint in P , but their closures in \bar{P} might intersect. We concentrate on \mathcal{D}^+ for notational simplicity, though everything we say applies to \mathcal{D}^- .

First, we recall how the separation properties of decomposition elements in the orbit space reflect the separation properties of their ends. Two subsets A and B of a circularly ordered set S are said to *link* if there are $a, a' \in A$ and $b, b' \in B$ such that $b \in (a, a')$ and $b' \in (a', a)$. The set $C \subset S$ is said to *separate* A and B if there are $c, c' \in C$ such that $A \subset (c, c')$ and $B \subset (c', c)$.

Lemma 2.4 [6, Proposition 6.9] *Let $K, L \in \mathcal{D}^+$. Then $\mathcal{E}(K)$ and $\mathcal{E}(L)$ do not link.*

Let $K, L, M \in \mathcal{D}^+$. Then K separates L from M if and only if $\mathcal{E}(K)$ separates $\mathcal{E}(L)$ from $\mathcal{E}(M)$.

The following two lemmas are the basic results we need about the closures of decomposition elements.

Lemma 2.5 [6, Lemma 7.8] *Let $K \in \mathcal{D}^+$. Then*

$$\bar{K} = K \cup \overline{\phi(\mathcal{E}(K))},$$

where the closure is taken in \bar{P} .

Lemma 2.6 *For each point $x \in \partial \bar{P}$ there are at most countably many decomposition elements $K \in \mathcal{D}^+$ with $\bar{K} \ni x$.*

Proof By Lemma 2.5, each decomposition element K whose closure intersects x either has an end at x (ie there is a $\kappa \in \mathcal{E}(K)$ with $\phi(\kappa) = x$) or has ends approaching x

(ie there are $\kappa_i \in \mathcal{E}(K)$ with $\phi(\kappa_i) \rightarrow x$ as $i \rightarrow \infty$). There are at most countably many decomposition elements with an end at x , since countably many ends are collapsed in the construction of the universal circle. Therefore, it suffices to show that there are countably many decomposition elements with ends approaching x ; in fact there are at most two.

We say that a 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 $\phi(\kappa_i) \rightarrow x$ as $i \rightarrow \infty$ and $\phi(\kappa_{i+1}) \subset (\phi(\kappa_i), x)$ for all i . 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, contradicting Lemma 2.4. Similarly, there is at most one decomposition element with ends approaching x in the negative direction. \square

2.4 Unions of decomposition elements

Lemma 2.7 *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 $P \setminus \bar{U}$ that contains y . Then $\text{Fr } V$ is connected because the plane satisfies the Brouwer property (see [9, 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 $\text{Fr } V$ is connected. But $\text{Fr } V \subset A$, so take the component of A containing V . \square

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 [8]. If the A_i live in the plane and infinitely many intersect some compact set, then some subsequence converges in the Hausdorff sense.

Lemma 2.8 *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 $\text{Fr } U$ is contained in a single decomposition element.*

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

Suppose that $e^+(\text{Fr } U)$ is not a single 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 them. Note that these are mutually nonseparating, so by Lemma 2.4 $\mathcal{E}(K)$, $\mathcal{E}(L)$ and $\mathcal{E}(M)$ are mutually nonseparating in the circular order on \mathcal{E} .

Assume that K, L, M are ordered so that $\mathcal{E}(L)$ is contained in the positively oriented interval between $\mathcal{E}(K)$ and $\mathcal{E}(M)$ and each decomposition element contained in U has ends contained in the positively oriented interval between $\mathcal{E}(M)$ and $\mathcal{E}(K)$.

Now let $n_i \in P$ be a sequence of points in U that approach some point in L and let N_i be the decomposition element containing n_i for each i . Note that infinitely many of the N_i intersect a compact set (for example, any arc from K to M) so we can take a Hausdorff convergent subsequence. See Figure 4. Let N_∞ be the limit of this subsequence. Then $N_\infty \cup L$ 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 . So by Lemma 2.7 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. So $e^+(\text{Fr } U)$ is a single point as desired. \square




Figure 4: Proof of Lemma 2.8

Combining the last two lemmas, we have the following result.

Corollary 2.9 *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 .*

2.5 Extending the endpoint maps

We can now prove our main theorem.

Theorem 2.10 *The map $e^+ : P \rightarrow S_\infty^2$ has a unique continuous extension to \bar{P} .*

Proof For each $x \in \partial\bar{P}$, 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 $e^+(x_i)$ converges. Set $e^+(x) = \lim e^+(x_i)$. We will show that this is well defined, and continuity and uniqueness follow.

Suppose there are two sequences $(x_i)_{i=1}^\infty$ and $(y_i)_{i=1}^\infty$ with $\lim x_i = x = \lim y_i$, but with $\lim e^+(x_i) = p \neq q = \lim e^+(y_i)$. Let A' be a simple closed curve separating p and q , so $A = (e^+)^{-1}(A')$ separates x_i from y_i for large enough i .

We claim that the x_i are eventually contained in a single component of $P \setminus A$. Otherwise, let U_i be the component of $P \setminus A$ that contains x_i , and take a subsequence so that the U_i are distinct. Then by Lemma 2.2, $\limsup e^+(x_i) \subset \limsup e^+(\text{Fr } U_i)$, which is impossible since $e^+(\text{Fr } U_i) \subset A'$ for all i . Similarly, the y_i are eventually contained in a single component $P \setminus A$.

Now by Corollary 2.9, some decomposition element $K \in \mathcal{D}^+$ separates $\{x_i\}_{i=I}^\infty$ from $\{y_i\}_{i=I}^\infty$ for I large enough. Therefore \bar{K} separates $\{x_i\}_{i=I}^\infty$ from $\{y_i\}_{i=I}^\infty$ in \bar{P} , which means that $x \in \bar{K}$.

We have shown that for each simple closed curve A' separating p and q , there is a decomposition element K lying in $(e^+)^{-1}(A')$ with $x \in \bar{K}$. There are uncountably many disjoint simple closed curves separating p and q , so this argument provides uncountably many distinct decomposition elements K with $x \in \bar{K}$. This contradicts Lemma 2.6, so the extension of e^+ to \bar{P} is well defined. □

We can extend the negative endpoint map similarly.

Proposition 2.11 *The maps $e^\pm : \bar{P} \rightarrow S_\infty^2$ agree on the boundary S_u^1 .*

Proof Let $x \in \partial\bar{P}$ and choose a sequence of points $x_i \in P$ converging to x . By Lemma 2.1, the distance between $e^+(x_i)$ and $e^-(x_i)$ approaches 0 as $i \rightarrow \infty$. □

This completes the extension theorem. Our generalization of the Cannon–Thurston theorem follows easily:

Theorem 2.12 *Let $\psi : P \rightarrow \mathbb{H}^3$ be a section of the quotient map $\mathbb{H}^3 \rightarrow P$. Then there is a unique continuous extension of ψ to \bar{P} .*

Proof Just set ψ to be equal to e^+ on the boundary S_u^1 . That this is continuous follows from Lemma 2.1. □

3 Closed orbits and Möbius-like groups

If M is a closed hyperbolic 3-manifold, then each $g \in \pi_1(M)$ acts on $\mathbb{H}^3 \cup S_\infty^2$ as a translation and rotation along a geodesic axis. The endpoints of this axis in S_∞^2 are the fixed points for the action of g . For each such g we label the attracting fixed point a_g and the repelling fixed point r_g .

Lemma 3.1 *Let Φ be a quasigeodesic flow on a closed hyperbolic 3-manifold M and let $g \in \pi_1(M)$. Then Φ has a closed orbit in the free homotopy class represented by g if and only if there is a flowline with an endpoint at either a_g or r_g .*

Proof Suppose there is 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 , so the endpoints of γ are at a_g and r_g .

Conversely, suppose that a flowline γ in the lifted flow 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 $n \rightarrow \infty$, so the $g^n(\gamma)$ all intersect a compact set in \mathbb{H}^3 . In other words, the point $\pi(\gamma) \subset P$ in the orbit space corresponding to γ has a bounded forward orbit under g . The Brouwer plane translation theorem [7] 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 . \square

The closed orbits theorem is an immediate consequence of the following.

Theorem 3.2 *Let M be a closed hyperbolic 3-manifold with a quasigeodesic flow Φ . Suppose that the action of $g \in \pi_1(M)$ on S_u^1 is not conjugate to a hyperbolic Möbius transformation. Then Φ has a closed orbit in the free homotopy class represented by g .*

Proof Suppose f is a homeomorphism of the circle and $x \in S^1$ is an isolated fixed point for f . Then x is either attracting, repelling or neither, in which case we call it indifferent. Each homeomorphism of the circle can be classified up to conjugacy by whether it has

- (1) no fixed points,
- (2) one indifferent fixed point (ie it is conjugate to a parabolic Möbius transformation),
- (3) two indifferent fixed points,
- (4) two fixed points in an attracting-repelling pair (ie it is conjugate to a hyperbolic Möbius transformation), or
- (5) more than two fixed points.

In [6], we showed that if the action of g fits into categories 1, 3 or 5, then the flow \mathfrak{F} has a closed orbit in the free homotopy class of g . So we just need to extend this to category 2.

Suppose that g acts on S_u^1 with a single fixed point x . If there are no closed orbits in the homotopy class represented by g , then by Lemma 3.1 there are no flowlines with an endpoint at either a_g or r_g , so $A = (e^+)^{-1}(a_g)$ and $R = (e^-)^{-1}(r_g)$ are both contained in the boundary S_u^1 of \bar{P} . But A and B are invariant sets, so they must both contain the fixed point x , a contradiction. \square

References

- [1] **MR Bridson, A Haefliger**, *Metric spaces of non-positive curvature*, Grundle. Math. Wissen. 319, Springer, Berlin (1999) MR1744486
- [2] **D Calegari**, *Universal circles for quasigeodesic flows*, *Geom. Topol.* 10 (2006) 2271–2298 MR2284058
- [3] **JW Cannon, WP Thurston**, *Group invariant Peano curves*, *Geom. Topol.* 11 (2007) 1315–1355 MR2326947
- [4] **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 MR2872578
- [5] **S Fenley, L Mosher**, *Quasigeodesic flows in hyperbolic 3-manifolds*, *Topology* 40 (2001) 503–537 MR1838993
- [6] **S Frankel**, *Quasigeodesic flows and Möbius-like groups*, *J. Differential Geom.* 93 (2013) 401–429 MR3024301
- [7] **J Franks**, *A new proof of the Brouwer plane translation theorem*, *Ergodic Theory Dynam. Systems* 12 (1992) 217–226 MR1176619
- [8] **JG Hocking, GS Young**, *Topology*, Addison-Wesley Publ., Reading, MA (1961) MR0125557
- [9] **RL Wilder**, *Topology of manifolds*, *AMS Colloq. Publ.* 32, Amer. Math. Soc. (1963) MR0182958

Department of Mathematics, Yale University
PO Box 208283, New Haven, CT 06520-8283, USA

steven.frankel@yale.edu

Proposed: David Gabai
Seconded: Benson Farb, Danny Calegari

Received: 22 July 2013
Revised: 25 February 2014