Group invariant Peano curves

JAMES W CANNON
WILLIAM P THURSTON

Our main theorem is that, if M is a closed hyperbolic 3–manifold which fibres over the circle with hyperbolic fibre S and pseudo-Anosov monodromy, then the lift of the inclusion of S in M to universal covers extends to a continuous map of $\mathbf{B}^2$ to $\mathbf{B}^3$, where $\mathbf{B}^n = \mathbf{H}^n \cup S_\infty^{n-1}$. The restriction to S_∞^1 maps onto S_∞^2 and gives an example of an equivariant S^2–filling Peano curve. After proving the main theorem, we discuss the case of the figure-eight knot complement, which provides evidence for the conjecture that the theorem extends to the case when S is a once-punctured hyperbolic surface.

20F65; 57M50, 57M60, 57N05, 57N60

1 Introduction

We will describe doubly degenerate Kleinian groups whose limit sets are beautiful equivariant continuous images of simple closed curves. The possible existence of such curves is suggested by the following considerations:

Thurston has shown that acylindrical Haken 3–manifolds admit a complete hyperbolic structure of finite volume. If S is a hyperbolic 2–manifold and ϕ is a pseudo-Anosov diffeomorphism of S, then the mapping torus $M = M(\phi) = (S \times [0,1])/\{(x,0) = (\phi(x),1) \mid x \in S\}$ is such a 3–manifold. (See Thurston [7] or Sullivan [6].) We may identify the universal cover M' of M with hyperbolic 3–space $\mathbf{H}^3$ and the universal cover S' of S with hyperbolic 2–space $\mathbf{H}^2$ so that S' is a hyperbolic plane embedded in the hyperbolic 3–space $M' = \mathbf{H}^3$. Thus we expect to find the circle $\mathbf{S}^1 = \partial(S') = \partial(\mathbf{H}^2)$ at infinity appearing in 2–sphere $\mathbf{S}^2 = \partial(M') = \partial(\mathbf{H}^3)$. See Figure 1.

But the fundamental group of S is a nontrivial normal subgroup of the nonelementary fundamental group of M, and a classical theorem implies that $\pi_1(S)$ and $\pi_1(M)$ must have the same limit set in $\mathbf{H}^3 \cup \mathbf{S}_\infty^2$. That is, the circle boundary of S' must equal the 2–sphere boundary of M', a paradox.

Published: 20 July 2007 DOI: 10.2140/gt.2007.11.1315

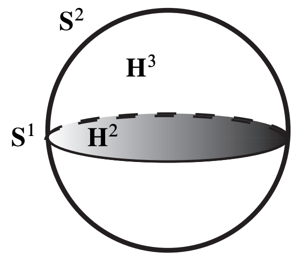

Figure 1

It follows immediately that S' is not embedded in M' as a hyperbolic plane but rather is severely folded in M' so that S' comes arbitrarily close to every point of $\mathbf{S}^2 = \partial \mathbf{H}^3$. This fact suggests that the embedding of S' in M' may extend to take the circle at infinity $\partial(S')$ continuously onto the 2–sphere boundary of M' as a $\pi_1(S)$–invariant 2–sphere-filling Peano curve. Our main theorem states that this is so, at least when S is a closed hyperbolic surface.

To describe these examples, whose limit sets are the entire Riemann sphere we need only the following ingredients:

(i) Closed hyperbolic 3–manifolds fibering over the circle (Thurston [7])

(ii) Analysis of the monodromies of the associated fibrations (Thurston [7] and Section 13, Pseudo-Anosov diffeomorphisms)

(iii) Elementary hyperbolic geometry

There are analogous examples, of interest in the study of Kleinian groups, whose limit sets are nonseparating proper subsets of the Riemann sphere (the totally one-sided degenerate case). To describe these examples, we need, in addition to (i), (ii) and (iii),

(iv) Thurston's deformation theory for quasi-Fuchsian groups (Thurston [7] and Section 7, Approximate metric structure, one-sided degenerate case)

The existence of the latter class of examples contradicts the main result of Abikoff [1], but it has been known for some time that a key argument in [1] is incorrect.

The doubly degenerate case is dealt with in Sections 2–5, with technical background in Sections 8–13. The modifications required for the singly degenerate case are outlined in Sections 6–7.

We analyze the precise topological structure of the 2–sphere-filling Peano curves given by the doubly degenerate case in Sections 14–15. The analysis depends on R L Moore's theorem on cellular upper-semicontinuous decompositions of the 2–sphere [5]: *If G is* *an upper semicontinuous decomposition of $\mathbf{S}^2$ such that each element of G is a proper subcontinuum of $\mathbf{S}^2$ that does not separate $\mathbf{S}^2$, then the decomposition space $\mathbf{S}^2/G$ is homeomorphic with $\mathbf{S}^2$.*

Our examples are all based on hyperbolic 3–manifolds that fiber over a circle, with compact, orientable, hyperbolic surface as fiber and with pseudo-Anosov monodromy. The main result is surely true as well when the fiber is a hyperbolic surface-with-punctures, but our proof is inadequate to deal with the punctured case. The metric and geodesic analysis becomes considerably more difficult in the presence of punctures. We leave the complete analysis of this more complicated setting for others. We simply present the evidence obtained by computer experimentation that approximates the resulting Peano curves when S is a punctured torus. These examples appear in Sections 16 and 17.

We denote hyperbolic n–space by $\mathbf{H}^n$, the $(n-1)$–sphere at infinity by $\mathbf{S}^{n-1}_\infty$, the n–ball that is $\mathbf{H}^n \cup \mathbf{S}^{n-1}_\infty$ by $\mathbf{B}^n$, the hyperbolic metric by d_H and the group of isometries of $\mathbf{H}^n$ by $\mathrm{Isom}(\mathbf{H}^n)$.

A *hyperbolic manifold M* of dimension n is a complete Riemannian n–manifold that is locally isometric with $\mathbf{H}^n$. Since the time of H Poincaré it has been known that each closed, orientable surface of genus $g \geq 2$ admits a hyperbolic Riemannian metric. *Once and for all, we fix such a 2–manifold S* with universal covering projection $p\colon \mathbf{H}^2 \to S$, p a local isometry. The standard identification of $\pi_1(S)$ with the group of covering translations of the universal cover of S gives a discrete faithful representation

$$\rho\colon \pi_1(S) \to \mathrm{Isom}\,(\mathbf{H}^2) \subset \mathrm{Isom}\,(\mathbf{H}^3)$$

A subgroup G of $\mathrm{Isom}(\mathbf{H}^n)$ is *discrete* if, for each compact subset K of $\mathbf{H}^n$, the number of elements g of G for which $K \cap gK \neq 0$ is finite. *The limit set* $\Lambda(G)$ of a discrete group G is the cluster set in $\mathbf{B}^n$ of any orbit Gx, $x \in \mathbf{H}^n$. Discreteness implies that $\Lambda(G) \subset \mathbf{S}^{n-1}_\infty$.

A discrete faithful representation $\rho\colon \pi_1(S) \to \mathrm{Isom}\,(\mathbf{H}^3)$ is called *Fuchsian* if the limit set $\Lambda(\rho(\pi_1(S)))$ is a geometric (round) circle, *quasi-Fuchsian* if $\Lambda(\rho(\pi_1(S)))$ is a topological circle, and *degenerate* otherwise. A degenerate group is called *doubly degenerate* if $\Lambda(\rho(\pi_1(S)))$ is equal to $\mathbf{S}^2_\infty$ and is called *totally one-sided degenerate* if $\mathbf{S}^2 - \Lambda$ is connected and not empty.

At the forefront in our description of totally degenerate groups is the fundamental group of the fiber for any closed hyperbolic 3–manifold that fibers over the circle with fiber homeomorphic with S. (See the next section.)

2 Hyperbolic 3–manifolds fibering over a circle

Our principal object of study is a closed 3–manifold M that fibers over a circle with fiber the closed orientable hyperbolic 2–manifold S of Section 1. Work of W P Thurston (see [7] or [6]) shows that such a manifold M admits a hyperbolic metric if and only if M arises from the following construction: let $\phi\colon S \to S$ denote a pseudo-Anosov diffeomorphism (see Section 13, Pseudo-Anosov diffeomorphisms, or [4]), and let

$$M = (S \times I)/\{(s,0) = (\phi(s),1)\}$$

be the closed manifold formed from the cylinder $S \times I$ over S when the ends $S \times \{0\}$ and $S \times \{1\}$ are identified by the homeomorphism ϕ. The diffeomorphism ϕ is called the *monodromy* of the fibration. An example of such a manifold is constructed in [6].

Once and for all, we fix a pseudo-Anosov diffeomorphism $\phi\colon S \to S$ (see Section 13), the associated 3–manifold $M = M(\phi)$ and covering projection $q\colon \mathbf{H}^3 \to M$, q a local isometry. With the pseudo-Anosov diffeomorphism ϕ are associated geodesic laminations λ_1 and λ_2 of S (Section 11), transverse measures dx and dy that assign a positive value to curves transverse to λ_1 and λ_2, respectively, and a multiplier $k > 1$ satisfying the defining conditions for pseudo-Anosov diffeomorphisms (Section 13):

(1) λ_1 and λ_2 bind S (Section 10).

(2) $\phi(|\lambda_1|) = |\lambda_1|, \quad \phi(|\lambda_2|) = |\lambda_2|$,

(3) $\int_{\phi(\gamma)} dx = (1/k)\int_\gamma dx$ and $\int_{\phi(\gamma)} dy = k\int_\gamma dy$.

The lifts of λ_1 and λ_2 to $\mathbf{H}^2$ are denoted by λ_1' and λ_2'.

3 Doubly degenerate groups

Let $p\colon \mathbf{H}^2 \to S$ and $q\colon \mathbf{H}^3 \to M$ be as in Section 1 (Introduction) and Section 2 (Hyperbolic 3–manifolds fibering over a circle) respectively. Let $G(S)$ and $G(M)$ denote the groups of covering translations of p and q, respectively. Let $h\colon S \to S \times \{0\} \to M$ denote the natural inclusion map. The inclusion h lifts to a map $i\colon \mathbf{H}^2 \to \mathbf{H}^3$ of universal covers.

$$\begin{array}{ccc} \mathbf{H}^2 & \xrightarrow{\;i\;} & \mathbf{H}^3 \\ \downarrow p & & \downarrow q \\ S & \xrightarrow{\;h\;} & M \end{array}$$

Since h induces an embedding $h_*\colon \pi_1(S) \to \pi_1(M) \simeq G(M)$, h induces a discrete faithful representation.

$$\rho\colon G(S) \to \pi_1(S) \to \pi_1(M) \to G(M).$$

We claim that ρ is totally degenerate with limit set $\mathbf{S}^2_\infty$; this claim is true even if S is not compact provided that M has finite volume. In fact, it is well-known and easy to check that a discrete hyperbolic group and any of its nontrivial normal subgroups have the same limit set except in trivial cases. Hence, $\Lambda(\rho(\pi_1(S))) = \Lambda(G(M)) = \mathbf{S}^2_\infty$.

4 Statement of the Main Theorem on Peano curves

The objects of interest are the covering projections $p\colon \mathbf{H}^2 \to S$, $q\colon \mathbf{H}^3 \to M$, the compactifications $\mathbf{B}^2 = \mathbf{H}^2 \cup \mathbf{S}^1_\infty$ of $\mathbf{H}^2$ and $\mathbf{B}^3 = \mathbf{H}^3 \cup \mathbf{S}^2_\infty$ of $\mathbf{H}^3$, the inclusion mapping $h\colon S \to S \times \{0\} \to M$, a lift $i\colon \mathbf{H}^2 \to \mathbf{H}^3$ of h to universal covers, the groups of covering transformations $G(S)$ over p and $G(M)$ over q, the induced mapping $h_*\colon \pi_1(S) \to \pi_1(M)$ induced by h and the totally degenerate group given as the image of the injection

$$\rho\colon G(S) \overset{p_*\ \simeq}{\Longrightarrow} \pi_1(S) \overset{h_*\ 1-1}{\Longrightarrow} \pi_1(M) \overset{q_*^{-1}\ \simeq}{\Longrightarrow} G(M).$$

Main Theorem (Peano curves) *There is a continuous function $j\colon \mathbf{B}^2 \to \mathbf{B}^3$ that renders the following diagram commutative:*

$$\begin{array}{ccc} \mathbf{B}^2 & \xrightarrow{\ j\ } & \mathbf{B}^3 \\ \uparrow & & \uparrow \\ \mathbf{H}^2 & \xrightarrow{\ i\ } & \mathbf{H}^3 \\ \downarrow{\scriptstyle p} & & \downarrow{\scriptstyle q} \\ S & \xrightarrow{\ h\ } & M \end{array}$$

Remark The restriction of j to $\mathbf{S}^1_\infty$ takes $\mathbf{S}^1_\infty$ onto $\mathbf{S}^2_\infty$ and is a $\rho(G(S))$–invariant 2–sphere-filling Peano curve. A more precise description of the topology of j will appear in Section 15. At this point we simply highlight the structures that we will use to prove the existence of the extension j.

In $\mathbf{H}^2$ the important structures will be the leaves of the two foliations λ_1' and λ_2' formed in $\mathbf{H}^2$ by lifting the pseudo-Anosov foliations λ_1 and λ_2 on S to the universal cover. The critical property of λ_1' and λ_2' is summarized in the following theorem which is a consequence of Theorem 10.2.

Theorem 4.1 (Neighborhoods at infinity in the ball) *Each point $x \in \mathbf{S}^1_\infty$ has arbitrarily small neighborhoods in $\mathbf{B}^2 = \mathbf{H}^2 \cup \mathbf{S}^1_\infty$ bounded by the closure in $\mathbf{B}^2$ of a single leaf of $\lambda_1' \cup \lambda_2'$.*

In $\mathbf{H}^3$ there are two important structures. They arise from *viewing* $\mathbf{H}^3$ *as a topological product:* $\mathbf{H}^3 = \mathbf{H}^2 \times (-\infty, \infty)$. The first of the two structures is geometric and consists of those subproducts of $\mathbf{H}^2 \times (-\infty, \infty)$ having the form $L \times (-\infty, \infty)$ where L is a leaf of either λ_1' or λ_2'. The second of our two structures is metric and consists of a Riemannian pseudo-metric-with-singularities ds on $\mathbf{H}^3$ defined in Section 5. The metric ds is obtained as the equivariant lift of a pseudo-metric from M. Hence, in particular it is $G(M)$ invariant. It is carefully chosen to have two properties.

(1) The metric ds is near enough to the standard hyperbolic metric that it induces the same topological structure at infinity.

(2) The metric ds respects the product structure and foliations enough so that the sets $L \times (-\infty, \infty)$ are metrically nice.

Property (1) is explained in Theorem 5.1. Property (2) is explained in Theorem 5.2. Properties (1) and (2) are exploited together in Theorem 5.3. And finally all of these properties combine to give the proof of the Main Theorem at the end of Section 5.

5 The approximate metric structure of M

The universal cover $\mathbf{H}^3$ of the fibered manifold M is topologically the product $\mathbf{H}^2 \times \mathbf{R}$ of the universal cover $\mathbf{H}^2$ of the fiber S and the universal cover $\mathbf{R}$ of the base $\mathbf{S}^1$ of the fibration. The monodromy map $\phi\colon S \to S$ and its associated laminations (λ_1, dx) and (λ_2, dy) and multiplier $k > 1$ supply a natural pseudometric $ds_0^2 = dx^2 + dy^2$ for $\mathbf{H}^2 \times \{0\}$ (see Section 12, Measured laminations binding a surface). In order to extend this pseudometric in a natural way to all of $\mathbf{H}^2 \times \mathbf{R}$ so as to be $G(M)$ invariant, we note that invariance requires

$$\int_{\gamma\times\{t\}} k\,dx = \int_{\tilde{\phi}^{-1}(\gamma)\times\{t\}} dx = \int_{\gamma\times\{t+1\}} dx, \quad \text{and}$$

$$\int_{\gamma\times\{t\}} dy = \int_{\tilde{\phi}^{-1}(\gamma)\times\{t\}} k\,dy = \int_{\gamma\times\{t+1\}} k\,dy.$$

where $\widetilde{\phi}\colon \mathbf{H}^2 \to \mathbf{H}^2$ covers $\phi\colon S \to S$. That is, the measure must stretch by a factor of k in the dx direction and by $1/k$ in the dy direction when $t \in \mathbf{R}$ increases by 1. This suggests the $G(M)$ invariant, infinitesimal pseudometric ds on $\mathbf{H}^3 \approx \mathbf{H}^2 \times \mathbf{R}$ given by the formula

$$ds^2 = k^{2t}\,dx^2 + k^{-2t}\,dy^2 + \lambda^2\,dt^2,$$

where λ is a positive constant that may be chosen arbitrarily. We choose $\lambda = \log k$ because the metrics $ds^2 = k^{2t} + (\log k)^2\,dt^2$ and $ds^2 = k^{-2t} + (\log k)^2\,dt^2$ on the plane $\mathbf{R} \times \mathbf{R}$ are isometric with the upper half-plane model of hyperbolic geometry $\mathbf{H}^2$ under the maps $(x,t) \mapsto (x, k^{-t})$ and $(y,t) \mapsto (y, k^t)$, respectively. We denote the associated $G(M)$–invariant global pseudometric by $(s = \int_\gamma ds)\colon \mathbf{H}^3 \times \mathbf{H}^3 \to [0,\infty)$. (See Sections 11-14 for properties of $\int ds$.)

Theorem 5.1 (Quasicomparability of metrics) *The hyperbolic metric d_H and the $G(M)$–invariant global pseudometric $s = \int_\gamma ds$ are quasicomparable.*

Proof The quasicomparability of d_H and s requires the existence of positive numbers $K > 1$ and $K' > 1$ such that

$$\max(s, d_H) > K' \quad \Rightarrow \quad \left(\frac{1}{K}\right) d_H \le s \le K d_H.$$

The facts we need to know about s are the following:

(1) s is $G(M)$ invariant.

(2) s is obtained by integrating a continuous infinitesimal pseudometric ds along paths of $\mathbf{H}^3$. (The continuity of s is convenient but not esssential; a weak kind of continuity like that of Section 11, Measured laminations, would suffice.)

(3) Each compact subset of $\mathbf{H}^3$ has finite s–diameter.

(4) Each pseudometric s–neighborhood

$$N(x,\epsilon;s) = \{y \in \mathbf{H}^3 | s(x,y) < \epsilon, x \in \mathbf{H}^3, \epsilon \in (0,\infty)\},$$

has compact closure in $\mathbf{H}^3$.

Condition (1) is easily checked. The continuity of ds follows from Section 13 (Pseudo-Anosov diffeomorphisms). Conditions (3) and (4) follow easily from the fact proved in Section 12 (Measured laminations binding a surface) that the metrics d_H and $\rho = \int (dx^2 + dy^2)^{1/2}$ on $\mathbf{H}^2$ are quasicomparable.

Realizing as we do that d_H on $\mathbf{H}^3$ also satisfies (1), (2), (3) and (4), *we need not distinguish between* $s_1 = s$ *and* $s_2 = d_H$ *provided we only use properties* (1–4). Hence it suffices to prove the single implication: there exist $K, K' > 1$ such that

$$s_2(x, y) > K' \quad \Rightarrow \quad s_2(x, y) \leq K s_1(x, y).$$

Pick $\alpha > 1$ so large that, if $z \in \mathbf{H}^3$, then the $G(M)$–translates of

$$N(z, \alpha; s_1) = \{x \in \mathbf{H}^3 | s_1(z, x) < \alpha\}$$

cover $\mathbf{H}^3$ (Property (3)). Pick $\beta > 1$ so large that, if $z \in \mathbf{H}^3$, then $N(z, \alpha; s_1) \subset N(z, \beta; s_2)$ (Property (4)).

Take K and $K' > 4\beta$ and suppose that $x, y \in \mathbf{H}^3$ are points satisfying $s_2(x, y) > K'$. Let γ denote a path from x to y satisfying

$$s_1(x, y) \leq \int_\gamma ds_1 < s_1(x, y) + 1.$$

Let $B_1, B_2, \ldots B_n$ denote a minimal covering of γ by neighborhoods of the form $N(z, \beta; s_2)$. It is then clearly true that

$$s_2(x, y) \leq 2\beta \cdot n.$$

But, subdividing γ into segments $\gamma_0, \gamma_1, \ldots, \gamma_k$, satisfying

$$0 \leq \int_{\gamma_0} ds_1 < 1, \quad \int_{\gamma_i} ds_1 = 1 \quad (i = 1 \ldots, k),$$

it becomes obvious from the choice of α and β and the minimality of the covering $B_1, \ldots, B_n$ that $n \leq k$. Hence

$$s_2(x, y) \leq 2\beta \cdot n \leq 2\beta \int_\gamma ds_1 < 2\beta(s_1(x, y) + 1), \quad \text{or}$$

$$s_2(x, y) - 2\beta < 2\beta \cdot s_1(x, y).$$

But $s_2(x, y) > K' > 4\beta$ implies that $s_2(x, y)/2 < s_2(x, y) - 2\beta$, so that

$$s_2(x, y) < 4\beta \cdot s_1(x, y) < K \cdot s_1(x, y),$$

as desired. □

Theorem 5.2 (Sets of the form $L \times (-\infty, \infty)$) *If L is a leaf of λ_1' or λ_2', then with respect to the metric s on $\mathbf{H}^3 = \mathbf{H}^2 \times (-\infty, \infty)$ the product $L \times (-\infty, \infty)$ is totally geodesic.*

Proof We must prove that if $x, y \in L \times (-\infty, \infty)$ and ϵ is any path from x to y in $\mathbf{H}^3$, then $L \times (-\infty, \infty)$ contains a path δ from x to y such that

$$\int_\delta ds \leq \int_\epsilon ds.$$

The standard method for proving the existence of such a δ is to construct a ds–reducing retraction $\rho\colon \mathbf{H}^3 \to L \times (-\infty, \infty)$; then δ may be taken to be $\rho \circ \epsilon$. We shall define ρ as a product map $\rho = \rho_0 \times id\colon \mathbf{H}^2 \times (-\infty, \infty) \to L \times (-\infty, \infty)$. The retraction $\rho_0\colon \mathbf{H}^2 \to L$ is defined as follows.

We may choose the notation so that L is a leaf of λ_1'. We first let C denote the closed subset of $\mathbf{H}^2$ that is the union of those leaves of λ_2' that intersect L. We define $\rho_0|C$ to be the map that sends each leaf L' of λ_2' in C to the point $L' \cap L$. We let U denote a component of $\mathbf{H}^2 \setminus C$. The boundary of U in $\mathbf{H}^2$ is the union of two geodesics $L_1(U)$ and $L_2(U)$, each intersecting L in a single point. The intersection of U with L is an open arc $A(U)$ joining $L_1(U) \cap L$ and $L_2(U) \cap L$. The arc $A(U)$ misses λ_2' entirely. One defines $\rho_0|U$ so that $\rho_0(U) = A(U)$ and so that $(\rho_0|U) \cup (\rho_0|\partial U)$ is continuous. Since the components U form a null sequence, it is easy to check that ρ_0 and also therefore ρ are continuous.

The retraction ρ kills the term in ds^2 involving dx^2, leaves the term involving dt^2 invariant, and, at worst, leaves the term involving dy^2 no larger. Hence the proof of the theorem is complete. □

Theorem 5.3 (The diameter of $L \times (-\infty, \infty)$) *Fix $z \in \mathbf{H}^2$. For each $\epsilon > 0$, there is a positive number N such that if L is a leaf of $\lambda_1' \cup \lambda_2'$ in $\mathbf{H}^2$ and $d_H(z, L) > N$, then the Euclidean diameter of $L \times (-\infty, \infty)$ in $\mathbf{H}^3 = \mathbf{H}^2 \times (-\infty, \infty) \subset \mathbf{B}^3 = \mathbf{H}^3 \cup \mathbf{S}^2_\infty$ is less than ϵ.*

Proof From the quasicomparability of $\int (dx^2 + dy^2)^{1/2}$ and d_H on $\mathbf{H}^2$ and the quasicomparability of $s = \int ds$ and d_H on $\mathbf{H}^3$, it follows that $d_H(z, L) \to \infty$ (as L approaches infinity in $\mathbf{H}^2$). From the quasicomparability of ds and d_H on $\mathbf{H}^3$ it follows (see [3]) that, since $L \times (-\infty, \infty)$ is quasi-totally geodesic in $\mathbf{H}^3$ with respect to ds, it is also quasi-totally geodesic in $\mathbf{H}^3$ with respect to the hyperbolic metric d_H in the following sense.

(∗) There is a $W > 0$ such that if L is any leaf of $\lambda_1' \cup \lambda_2'$, if x and y are any two points of $L \times (-\infty, \infty)$ and g is the hyperbolic geodesic from x to y, then there is a path γ from x to y in $L \times (-\infty, \infty)$ within W of the geodesic g.

But since hyperbolic geodesics g that miss large compact subsets of $\mathbf{H}^3$ must have uniformly small Euclidean diameter in the Poincaré disk model $\mathbf{H}^3 \subset \mathbf{B}^3 = \mathbf{H}^3 \cup \mathbf{S}^2_\infty$, it follows that corresponding s–geodesics γ must have uniformly small Euclidean diameter. Thus if $L \times (-\infty, \infty)$ misses a large compact subset of $\mathbf{H}^3$, then $L \times (-\infty, \infty)$ has small Euclidean diameter in $\mathbf{B}^3$. □

Proof of Main Theorem There exists a natural continuous, possibly multivalued extension $j\colon \mathbf{B}^2 \to \mathbf{B}^3$ of $i\colon \mathbf{H}^2 \to \mathbf{H}^3$ defined as follows for $x \in \mathbf{S}^1 = \partial \mathbf{H}^2$.

Let $C_1, C_2, \dots$ denote a sequence of compact neighborhoods of x in $\mathbf{B}^2$ whose intersection is $\{x\}$. Let $U_1 = C_1 \cap \mathbf{H}^2$, $U_2 = C_2 \cap \mathbf{H}^2$, etc. Let

$$j(x) = \cap_{n=1}^{\infty} \overline{i\,U_n} \subset \mathbf{B}^3.$$

To prove the theorem it suffices to show that j is single valued. We can do this simply by showing that

$$\lim_{n\to\infty} \operatorname{diam}\,(i\,U_n) = 0.$$

Theorem 4.1 (Neighborhoods at infinity in the ball) implies that we lose no generality in assuming that C_n is bounded in $\mathbf{B}^2$ by the closure of a single leaf L_n of $\lambda'_1 \cup \lambda'_2$.

By Theorem 5.3,

$$\lim_{n\to\infty} \operatorname{diam}\, i\,(L_n \times (\infty, \infty)) = 0.$$

But $i\,(L_n \times (-\infty, \infty))$ separates $i\,U_n$ from a large compact subset of $\mathbf{H}^3$. It follows that

$$\lim_{n\to\infty} \operatorname{diam}\,(i\,U_n) = 0,$$

hence that j is a continuous function as desired. □

6 Fully one-sided degenerate groups

A typical discrete faithful action of a surface group in $\operatorname{Isom}(\mathbf{H}^3)$ is a quasi-Fuchsian group, where there are exactly two components of the domain of discontinuity separated by a Jordan curve.

The surface groups coming from fibers of fibrations are completely opposite to these typical actions: in the case of fibers, both components of the complement of the curve have shrunk to nothing. As we have seen, the Jordan curve becomes a Peano curve in these cases, filling the sphere.

There is also an intermediate case, surface groups for which one of the expected components of the domain of discontinuity is completely absent but the other component is healthy. Such a group is called a *fully one-sided degenerate group*. We will show that at least for certain of these groups, the Jordan curve becomes a continuous map of the circle that traces out a fuzzy tree embedded on $\mathbf{S}^2$.

The existence of fully one-sided degenerate surface groups was established by Bers, who analyzed what happens if a quasi-Fuchsian group is varied so that the conformal structure on the quotient surface S_1 of one of the domains of discontinuity is held constant. Here in outline is the proof that they exist.

According to the deformation theory of quasi-Fuchsian groups, the conformal structure on the quotient S_2 of the other domain can be varied arbitrarily over Teichmüller space. If a normalization of the groups is chosen in some appropriate way, this gives rise to a family of conformal embeddings of the universal cover of S_1 into the sphere $\mathbf{S}^2_\infty$. It is easy to see that the closure of this family of embeddings is compact in the topology of pointwise convergence, by the theory of normal families.

All the limit embeddings give rise to discrete faithful group actions. Those actions for which there is more than one component of the domain of discontinuity are contained in a countable union of complex subvarieties of the deformation space that have less than full dimension. Therefore, most of the limit groups have to be fully one-sided degenerate.

Unfortunately, this construction does not give much of a cluse as to the nature of the fully one-sided degenerate groups. In [8] and [7], a good deal of theory for these groups is developed, expressed in terms of geodesic laminations. To each one-sided degenerate group $\rho(\pi_1(S))$ is associated a certain geodesic lamination on S, called the *ending lamination* $\epsilon(\rho)$ for the group. The ending lamination can be defined in terms of the rough location of geodesics in the hyperbolic three-manifold $M_\rho = \mathbf{H}^3/\rho(\pi_1(S))$. For any compact set $K \subset M_\rho$, let SC_K denote the set of simple closed curves on S whose geodesics in M_ρ do not intersect K. Thurston showed that for a limit of quasi-Fuchsian groups that is a one-sided degenerate group, this set is nonempty no matter how large the compact set K. As K increases toward all of M, SC_K decreases. The limit set for SC_K in the space of measured laminations on S consists of all measures supported on a certain geodesic lamination $\epsilon(M_\rho)$.

If a limit of quasi-Fuchsian groups degenerates at both ends, then there are two ending laminations, one for each end of the limit manifold. For more details, see Thurston, [8] and [7]. More recently, Francis Bonahon [2] has proven that the theory of ending laminations extends to the case when a surface group is not known to be the limit of quasi-Fuchsian groups.

Our main theorem in the case of fully one-sided degenerate groups is the following:

Theorem 6.1 (Ending stable gives fuzzy tree) *If ϕ is any pseudo-Anosov homeomorphism of S, then there is a fully one-sided degenerate group $\rho(\pi_1 S)$ whose ending lamination is the stable lamination for ϕ such that any π_1–equivariant map from $\mathbf{H}^2 = \widetilde{S}$ onto $\mathbf{H}^3$ extends to a continuous map of $\mathbf{B}^2 = \mathbf{H}^2 \cup \mathbf{S}^1_\infty$ to $\mathbf{B}^3 = \mathbf{H}^3 \cup \mathbf{S}^2_\infty$.*

It would be nice to be able to say what happens for other ending laminations. The problem in general is equivalent to knowing whether or not the limit set for the group is locally connected. No examples are known where the limit set for a surface group is not locally connected. Our method depends on knowing the approximate metric structure of the associated 3–manifold. We do not have a good analysis of this metric structure in the general case, although our method would probably extend before encountering insurmountable obstacles to the case of ending laminations belonging to an uncountable but rare set of laminations having the regularity properties enjoyed by the stable and unstable laminations for pseudo-Anosov homeomorphisms. The success of Bonahon's proof that a general surface group is geometrically tame gives more hope for understanding the topology of the general limit set.

The analysis in this case is inevitably less elementary and less self-contained than in the case of limit sets for surface groups that are fibers of fibrations, since the hypothesis already involves the idea of the ending lamination of an end. Basic references are Thurston, [8] and [7].

7 Approximate metric structure, one-sided degenerate case

The three-manifold M_ρ associated with a fully one-sided degenerate surface group $\rho\colon \pi_1(S) \to \operatorname{Isom}(\mathbf{H}^3)$ has two ends, one of which flares out exponentially and develops out toward the domain of discontinuity, the other of which has bounded diameter and somehow controls the topological and geometrical structure of the limit set. In order to analyze the topology of the limit set as we did in the case of a surface group coming from a fiber, we have to analyze the approximate metric structure of this latter end.

The idea is that when the ending lamination is the stable or unstable manifold of a pseudo-Anosov homeomorphism ϕ of S, we should expect the end to look metrically similar to one of the two ends for the doubly degenerate group that is the fiber of the mapping torus M_ϕ. We will prove this by a trick of passing to a limit of a sequence of representations of the surface group differing from the original by an automorphism of the surface group.

There is a compactification $\overline{T}$ of the Teichmüller space T that was discovered by Thurston in conjunction with the theory of pseudo-Anosov homeomorphisms, in which a sphere of projective classes of measured laminations on S is adjoined to the open ball of hyperbolic structures on S to form a closed ball.

Any homeomorphism of S acts as a homeomorphism of this ball. If ϕ is a pseudo-Anosov homeomorphism, then ϕ fixes exactly two points in the ball, both of which are on the boundary. These two points are the stable lamination λ_1 and unstable lamination λ_2 of ϕ. On S, the transverse measure of λ_1 is increased by ϕ while the transverse measure of λ_2 is decreased by ϕ. Simple closed curves on S tend to look closer to λ_1 after applications of ϕ. On $\overline{T}$, the stable lamination λ_1 is an attracting fixed point for the action of ϕ, while the unstable lamination is a repelling fixed point. In fact, under iteration of ϕ, the entire ball except for λ_2 converges toward λ_1.

The ending lamination is unfortunately not defined as a boundary point of $\overline{T}$—it is only a topological lamination that admits a transverse measure, not a measured lamination. If a lamination is uniquely ergodic, it defines a unique point on $\partial\overline{T}$. The stable and unstable laminations of ϕ are uniquely ergodic.

Let R be the set of conjugacy classes of representations of $\pi_1(S)$ up to conjugacy that are quasi-Fuchsian, with conformal structures on the two domains of discontinuity of the form $(\phi^{-n}(g_0), \phi^m(g_0))$, where g_0 is a fixed conformal structure and $n, m \geq 0$. The double limit theorem of Thurston [8] implies that the closure of R is compact. The added elements of the closure consist of:

(1) if $m \to \infty$ but n stays bounded, fully one-sided degenerate representations such that one ending lamination is λ_1, while the remaining component of the domain of discontinuity has conformal structure $\phi^{-n}(g_0)$;

(2) if $n \to \infty$ but m stays bounded, fully one-sided degenerate representations such that one ending lamination is λ_2, while the remaining component of the domain of discontinuity has conformal structure $\phi^m(g_0)$;

(3=1+2) if $n \to \infty$ and $m \to \infty$, doubly degenerate representations with ending laminations λ_1 and λ_2.

A hyperbolic structure for the mapping torus of ϕ was constructed from limits of type (3) by adjoining an additional transformation T_ϕ that induces the automorphism ϕ of $\pi_1(S)$. Since T_ϕ acts as a quasiconformal automorphism of each group in R, with uniformly bounded quasiconformal constant, it acts also as a quasiconformal automorphism of any of the limiting groups. But in case 3, there are no quasiconformal automorphisms except isometries, by a theorem of Sullivan or a theorem of Thurston. Thus, any limit of type (3) gives rise to a hyperbolic 3–manifold for M_0. Mostow's rigidity theorem asserts that such a hyperbolic structure is unique, so there is a unique limit of type (3).

We can also obtain the limit of type (3) by first forming the limits of type (1), then letting n go to infinity. The map T_ϕ^{-1} acts on the set of limits of type (1) by incrementing n. This means that we can choose a fixed representation ρ_0 of type (1), and then the limit of

$$\rho_0 \circ T_\phi^{-k}$$

must exist and equal the unique limit of type (3). (Every subsequence must have a convergent subsequence, but there is only one possible limit for these subsequences, so the entire sequence must converge.)

There is also a geometric form of the statement about algebraic convergence. The quotient manifold $M_0 = \mathbf{H}^3/\rho_0(\pi_1(S))$ has one geometrically infinite tame end. If we take a sequence of points tending toward this end, this defines a sequence of based hyperbolic manifolds. There is necessarily a convergent subsequence, and the limit of the sequence is isometric to the quotient manifold by the limit of type (3), with some choice of base point.

This fact can be proven from the existence of the algebraic limit above, by first considering the sequence of base points $\{p_k\}$ where p_k comes from the point in $\mathbf{H}^3$ that minimizes the total translation length of a fixed set of generators for $\pi_1(S)$ under the representation $\rho_0 \circ T_\phi^{-k}$. The normalization forces the actual sequence of representations to have a convergent subsequence, not just a subsequence convergent up to conjugacy. The geometric limit of the manifolds based at p_k is covered by the quotient of the group coming from the algebraic limit representation. But this manifold is only a covering space in very special ways: anything it covers is either compact, or is covered with finite degree. These kinds of covering spaces cannot arise in a geometric limit, so the geometric limit is simply the quotient of the representation of type (3).

A knowledge of all possible geometric limits as the base point moves toward the end gives a description of the approximate geometry of the geometrically infinite tame end: it is exactly the same approximate geometry as the previous case, the infinite cyclic covering of the mapping torus of ϕ. In the geometrically finite direction, the approximate geometry is easy to understand: it is just an exponentially flaring collar.

We can write down an expression for the approximate pseudometric: if dx and dy are the transverse measures for λ_1 and λ_2, then on $\widetilde{S}\times\mathbf{R}$, the metric $ds^2 = \exp(2|t|)dx^2 + \exp(2t)dy^2 + dt^2$ is quasicomparable to the hyperbolic metric, where t is a parameter for $\mathbf{R}$.

In the ds pseudometric for $\widetilde{S} \times \mathbf{R}$, any leaf of the lamination λ_1 sweeps out a surface that is locally isometric to a hyperbolic plane. The retraction of $\widetilde{S}\times\mathbf{R}$ to this hyperbolic plane decreases distances, as before. Therefore, each of these hyperbolic planes is isometrically embedded in the pseudometric, hence quasi-isometrically embedded in $\mathbf{H}^3$. Its boundary forms a circle in $\mathbf{S}^2_\infty$, with the leaves of λ_1 at the geometrically infinite end being collapsed to points on the limit set for the group.

8 Hyperbolic surfaces and limit geodesics

A *hyperbolic surface* is a 2–manifold S with a complete Riemannian metric ds of constant negative curvature -1. The universal cover of S is hyperbolic 2–space $\mathbf{H}^2$. The group $G(S)$ of covering transformations is a group of hyperbolic isometries of $\mathbf{H}^2$. The covering map $p\colon \mathbf{H}^2 \to S$ is a local isometry.

We consider only the case where S is orientable and closed of genus $g \geq 2$. The metric ds then defines an area element dA on S with respect to which the area $\int_S dA$ of S is finite and equal to $2\pi g$.

Theorem 8.1 (Limit geodesics) *Every geodesic ray $r\colon [0,\infty) \to S$ has a (possibly nonunique) limit geodesic $r'\colon (-\infty,\infty) \to S$ having the following property.*

> *Given $\epsilon > 0$ and any geodesic path R' in r', there are infinitely many geodesic paths R in r pointwise within ϵ of R'.*

Proof By the compactness of S, some sequence of segments $r \mid [n_1, n_1+1]$, $r \mid [n_2, n_2+1]$, $\ldots$ $(n_1 < n_2 \ldots)$ converges to a geodesic segment $r'\colon [0,1] \to S$. The segment r' extends to a geodesic $r'\colon (-\infty,\infty) \to S$. Let $\epsilon > 0$ and any geodesic path R' in r' be given. We may assume R' has the form $R' = r' \mid [-n,n]$ for some n. Then for $r \mid [n_i, n_i+1]$ close enough to $r' \mid [0,1]$, a requirement that can be realized simply by choosing $n_i > n$ very large, $r \mid [n_i - n, n_i + n]$ will be pointwise close to $R' = r' \mid [-n,n]$, as desired. □

9 Geodesic laminations

A *geodesic lamination* λ on S is a collection of disjoint, simple (that is, nonsingular) geodesics on S called *leaves* whose union $|\lambda|$ is closed in S.

Theorem 9.1 (Lamination complement) *The complement $S - |\lambda|$ of a lamination is nonempty (since no closed surface of genus > 1 admits a foliation without singularities).*

Each component C of $S - |\lambda|$ is (clearly) a convex hyperbolic polyhedron (that is, C has convex universal cover) that has only ideal vertices and is bounded by leaves of λ. The set $S - |\lambda|$ has only finitely many sides.

Proof Augment λ by adding to λ disjoint, simple closed geodesics $J_1, J_2, \dots$ in $S-|\lambda|$ until it is impossible to add more. The process stops after finitely many additions since (i) no closed geodesic is homotopically trivial. (ii) no compact surface contains infinitely many disjoint, nontrivial, non-freely-homotopic simple closed curves, and (iii) freely homotopic closed geodesics coincide.

We recall that the area of a hyperbolic triangle with angles α, β, and γ is $\pi-\alpha-\beta-\gamma$. Area calculations will allow us to prove that, with λ so augmented, the only forms assumed by components C of $S \setminus |\lambda|$ are the following:

(1) C is simply connected and has area $0 < \text{Area}(C) = (n-2)\pi \le 2\pi g$, where n is the number of sides of C. (The polygon C has its n vertices at ∞, hence can be divided into $n-2$ ideal triangles, each of area π.)

(2) C is an open annulus, one boundary component is a simple closed geodesic J, the other boundary has $n \ge 1$ sides formed from non-closed geodesics, and the area of C is $0 < \text{Area}(C) = n\pi \le 2\pi g$. (The polygon C cut along a geodesic joining J to some other boundary component can be divided into $n+2$ triangles of angle sum 2π.)

(3) C is a pair of pants (an open disk with two holes) and each of the three boundary components consists of a single simple closed geodesic. The area of C is $0 < \text{Area}(C) = 2\pi \le 2\pi g$. (The polygon C can be divided into two right-angled hexagons. Each hexagon can be divided into 4 triangles with angle sum $6 \cdot \pi/2 = 3\pi$.)

Indeed, any component C of $X-|\lambda|$ not having form (1), (2), or (3) contains a simple closed curve J' not freely homotopic into the boundary of C. Since C is convex, the geodesic J freely homotopic to J' must lie in $C \cup \partial C$, hence in C. But then J' could be used to further augment λ, a contradiction.

From the area calculations of (1), (2) and (3) it follows immediately, as asserted, that $S-|\lambda|$ has only finitely many components, each having only finitely many sides. □

Theorem 9.2 (Lamination area) *Area* $(|\lambda|) = 0$.

Proof Adding leaves in $S-|\lambda|$, one can complete λ to a foliation-with-singularities completely filling S. The index theorem allows one to calculate the Euler characteristic in terms of these singularities. The estimate thus obtained on the singularities alllows one to calculate via (1), (2) and (3) the total area of $S-|\lambda|$. It equals Area (S). Hence Area $(|\lambda|) = 0$. □

10 Laminations binding a surface

A pair of geodesic laminations λ_1 and λ_2 is said to *bind* S if they satisfy the four conditions of the following theorem.

Theorem 10.1 (Laminations binding a surface) *For geodesic laminations λ_1 and λ_2 on S the first three of the following conditions are equivalent and imply the fourth.*

(1) *Each simple geodesic on S crosses a leaf of $\lambda_1 \cup \lambda_2$.*

(2) *Each simple geodesic ray on S crosses a leaf of $\lambda_1 \cup \lambda_2$.*

(3) *Each geodesic ray on S crosses a leaf of $\lambda_1 \cup \lambda_2$.*

(4) *Let λ' denote the lift of λ to the universal cover $\mathbf{H}^2$ of S. Then each component of $\mathbf{H}^2 - (|\lambda'_1| \cup |\lambda'_2|)$ is the interior of a compact, convex hyperbolic polyhedron in $\mathbf{H}^2$.*

Proof That (3) $\Rightarrow$ (2) and (2) $\Rightarrow$ (1) is clear.

(1) $\Rightarrow$ (2): Let r denote a simple geodesic ray on S and r' a limit geodesic of r (Section 8, Hyperbolic surfaces and limit geodesics). By (1), r' crosses a leaf of $\lambda_1 \cup \lambda_2$. Hence, by Section 8, r crosses the same leaf of $\lambda_1 \cup \lambda_2$.

(2) $\Rightarrow$ (4): Each component C of $S - (|\lambda_1| \cup |\lambda_2|)$ clearly has convex lifts in $\mathbf{H}^2$. We claim that C is simply connected. For otherwise $C \cup \partial C$ would contain a simple closed geodesic J; by (2) J would cross a leaf of $\lambda_1 \cup \lambda_2$ and J would not lie in $C \cup \partial C$, a contradiction. If a lift C' of C did not have compact closure in $\mathbf{H}^2$, then C' would contain a geodesic ray r'. Since C is simply connected, the projection $C' \to C$ is a homeomorphism. Hence the image r of r' would be a simple geodesic ray in S not crossing any leaf of $\lambda_1 \cup \lambda_2$. Hence C' has compact closure in $\mathbf{H}^2$, and (4) follows.

(2) $\Rightarrow$ (3): Let r denote a geodesic ray on S. Since (2) $\Rightarrow$ (4), any lift r' of r to $\mathbf{H}^2$ intersects a leaf of $\lambda'_1 \cup \lambda'_2$. Hence r intersects a leaf of $\lambda_1 \cup \lambda_2$. If r does not cross that leaf, then r coincides with that leaf, hence is simple, hence crosses some leaf by (2). □

Remark Easy examples, where λ'_1 and λ'_2 share a leaf that does not appear in the boundary of any component of $\mathbf{H}^2 \setminus (|\lambda'_1| \cup |\lambda'_2|)$, show that (4) does not imply (1).

Theorem 10.2 (Neighborhoods at infinity) *Suppose laminations λ_1 and λ_2 bind S. Then each point $x \in \mathbf{S}^1_\infty$ at infinity in $\mathbf{B}^2 = \mathbf{H}^2 \cup \mathbf{S}^1_\infty$ has arbitrarily small Euclidean neighborhoods in $\mathbf{B}^2$ bounded by the closure in $\mathbf{B}^2$ of a single leaf of $\lambda'_1 \cup \lambda'_2$.*

Proof Since λ_1 and λ_2 bind S, λ_1 and λ_2 have no leaf in common. Hence no leaf L_1 of λ'_1 and leaf L_2 of λ'_2 have a common endpoint at infinity; for otherwise geodesic rays r_1 and r_2 in the projections of L_1 and L_2 to S would have a common limit geodesic L (Section 8, Hyperbolic surfaces and limit geodesics) and L would be a leaf common to λ_1 and λ_2.

Let r denote a geodesic ray in $\mathbf{H}^2$ with infinite endpoint x. If possible, take r to lie in $|\lambda'_1|$ or $|\lambda'_2|$. Since λ_1 and λ_2 bind S, there exist leaves $L_1, L_2, L_3, \dots$ in $\lambda'_1 \cup \lambda'_2$ such that L_i crosses r in the Euclidean $(1/i)$–neighborhood of x. We claim that $L_i \to x$ as $i \to \infty$. Suppose not. Then some subsequence converges to a leaf L of $\lambda'_1 \cup \lambda'_2$ with infinite endpoint x, say L in $|\lambda'_1|$. By the choice of r it follows that r is contained either in $|\lambda'_1|$ or in $|\lambda'_2|$. By the previous paragraph, since L is contained in $|\lambda'_1|$, it is impossible that r be contained in $|\lambda'_2|$. Hence r is contained in $|\lambda'_1|$. But then each L_i is a leaf of λ'_2, hence L is contained in $|\lambda'_2|$. But this contradicts the previous paragraph. Hence $L_i \to x$ as asserted, and the L_i cut off small neighborhoods of x in $\mathbf{B}^2$ as desired. □

11 Measured laminations

A *transverse measure* on a geodesic lamination λ is a positive measure dm defined on local transversals to the leaves of λ, invariant under local projection along the leaves of λ and positive and finite on nontrivial compact transversals to the leaves of λ. Such a measure lifts to a $G(S)$–invariant transverse measure on the inverse λ' of λ in $\mathbf{H}^2$, the lifted measure also denoted by dm. A geodesic lamination with a transverse measure is called a *measured lamination*.

If dm is a transverse measure on a lamination λ and $\gamma\colon [a,b] \to \mathbf{H}^2$ is any path, one may define the *integral* $\int_\gamma dm$: if X is a closed subset of $\mathbf{H}^2$, define $m(X)$ to be the measure of the projection of $X \cap |\lambda'|$ into the leaf space of λ'; let $P = a = a_0 < a_1 < \cdots < a_n = b$ denote a partition of $[a,b]$; define $m(\gamma, P) = \sum_{i=1}^{n} m(\gamma[a_{i-1}, a_i]) - \sum_{i=1}^{n-1} m(\gamma(a_i))$; take $\int_\gamma dm$ as the supremum of $m(\gamma, P)$ over all partitions P of $[a,b]$. The negative term in $m(P)$ avoids doubling the contribution to $\int_\gamma dm$ of a point lying on a leaf that supports an atom of dm.

Given two points x and y of $\mathbf{H}^2$, if γ is the geodesic segment joining x and y, and δ is any path joining x and y, then $\int_\gamma dm \le \int_\delta dm$.

We present four theorems. The first two describe limitations on the leaves of a measured lamination. The second two describe weak continuity properties of $\int_\gamma dm$.

Theorem 11.1 (Atoms and closed curves) *If (γ, dm) is a measured lamination on S and L is a leaf of λ carrying an atom of dm, then L is a simple closed curve.*

Proof Otherwise, let R denote a limit geodesic of L in the sense of Section 8 (Hyperbolic surfaces and limit geodesics) and let X denote a compact transversal to R. Then X intersects L infinitely often and $\int_X dm = \infty$ a contradiction. $\square$

Theorem 11.2 (Zero-angle) *Suppose that (λ, dm) is a measured lamination on S, and suppose that r_1 and r_2 are distinct geodesic rays in $\mathbf{H}^2$ each compatible with λ' in the sense that either $r_i \cap |\lambda'| = r_i$ or $r_i \cap |\lambda'| = \varnothing$. Then if r_1 and r_2 have the same endpoint x at infinity, and there is a neighborhood N of x in $\mathbf{B}^2 = \mathbf{H}^2 \cup \mathbf{S}^1_\infty$ such that either*

(1) *no point of $N \cap |\lambda'|$ lies in the open angle of zero measure between r_1 and r_2; or*

(2) *exactly one leaf L of λ' intersects the closed angle of zero measure bounded by r_1 and r_2 in N; L has infinite endpoint x; L is isolated and carries an atom of dm: and the projection of L in S is a simple closed curve.*

Proof Since r_1 and r_2 are compatible with λ', N may be chosen so small that if L is any leaf of λ' that intersects the closed angle between r_1 and r_2 in N, then L has x as an infinite endpoint. If no leaf L of λ' intersects the open angle between r_1 and r_2 in N, (1) is satisfied and we are done. Hence we suppose the existence of a leaf L with infinite endpoint x containing an infinite ray r that lies strictly between r_1 and r_2 in N.

The image $p(r)$ in S has a limit geodesic r' in the sense of Section 8 (Hyperbolic surfaces and limit geodesics). Let X denote a short compact geodesic arc transverse to some short geodesic subarc Y of r' in S. We may assume $X \cap Y$ is a single point. Arbitrarily short subsegments of X near $X \cap Y$ lift to segments in $\mathbf{H}^2$ crossing the entire angle from r_1 to r_2 in N. The integral of dm along each of these subsegments is the same strictly positive number. If these segments may be chosen to be disjoint, then $\int_X dm = \infty$, a contradiction. Otherwise we conclude that $r' = p(r)$, that r' is isolated, and that r' carries an atom of dm. Hence, by Theorem 11.1 (Atoms and closed curves), $r' = p(L)$ is a simple closed leaf. Any other leaf L' satisfying the properties of L must have projection $p(L')$ with the same limit geodesic r'; as above, $p(L') = r'$. Two lifts of the same simple closed curve with a common endpoint at infinity are equal. Hence $L = L'$ and (2) is satisfied. $\square$

Measured laminations satisfy the following two weak continuity properties.

Theorem 11.3 (Continuity property 1) *If (λ, dm) is a measured lamination and A a positive constant, then there is a positive constant B such that each geodesic segment γ in $\mathbf{H}^2$ of length $\le A$ has integral $\int_\gamma dm \le B$.*

Proof Otherwise, take segments γ_i of length $\le A$ converging to a segment in the interior of a segment Y, $\int_{\gamma_i} dm \to \infty$. Let X denote a segment transverse to γ. For all i sufficiently large, each leaf of λ' hitting γ_i hits γ or X. Hence

$$\infty = \lim \int_{\gamma_i} dm \le \int_{\gamma \cup X} dm < \infty,$$

a contradiction. □

Theorem 11.4 (Continuity property 2) *If (λ, dm) is a measured lamination, then there is a positive number C such that, if γ_i is any sequence of geodesic segments converging to any geodesic segment γ, then*

$$\left(\limsup \int_{\gamma_i} dm\right) - C \le \int_\gamma dm \le \left(\liminf \int_{\gamma_i} dm\right) + C.$$

Proof Use property 1 to find a positive number B such that each geodesic segment of length ≤ 1 has integral $\le B$. Take $C = 2B$.

Let γ_- be a geodesic segment in $\mathrm{Int}\gamma$ with length $(\gamma_-) >$ length$(\gamma) - 1$. If $\mathrm{Int}\gamma$ intersects $|\lambda'|$, choose γ_- to contain a point of $|\lambda'|$.

Let γ_+ be a geodesic segment containing γ in its interior with length $(\gamma_+) <$ length$(\gamma) + 1$.

Let X denote a geodesic segment transverse to γ at a point of γ_- with length$(X) < 1$. If γ_- intersects $|\lambda'|$ and does not lie in $|\lambda'|$, take $X \subset |\lambda'|$.

Case 1 $\gamma \subset |\lambda'|$ Note that $\int_\gamma dm \le \int_X dm < B < C$. For all i sufficiently large, each leaf of λ' intersecting γ_i also intersects X. Hence $\int_{\gamma_i} dm \le \int_X dm < B < C$. Hence $\left|\int_\gamma dm - \int_{\gamma_i} dm\right| < C$.

Case 2 $(\mathrm{Int}\gamma) \cap |\lambda'| = \varnothing$ Then for γ and for all γ_i, i sufficiently large, $|\lambda'|$ intersects each of γ and γ_i in two segments of total length < 1. Hence by choice of B, $\int_\gamma dm \le 2B = C$ and $\int_{\gamma_i} dm \le 2B = C$.

Case 3 $(\mathrm{Int}\gamma) \cap |\lambda'| \ne \varnothing$ and $X \subset |\lambda'|$ For all i sufficiently large, each leaf of λ' that intersects γ_- also intersects γ_i. Hence

$$\int_\gamma dm \le \left(\int_{\gamma_-} dm\right) + 2B \le \left(\int_{\gamma_i} dm\right) + C.$$

Also, for all i sufficiently large, each leaf of λ' that intersects γ_i intersects γ_+. Hence

$$\int_{\gamma_i} dm \leq \int_{\gamma_+} dm \leq \int_{\gamma} dm + 2B = \left(\int_{\gamma} dm\right) + C.$$

Cases 1, 2 and 3 exhaust the possibilities and complete the proof of Property 2. □

12 Measured laminations binding a surface

Two measured laminations (λ_1, dx) and (λ_2, dy) are said to *bind* S if the associated geodesic laminations λ_1 and λ_2 bind S in the sense of Section 10 (Laminations binding a surface).

Theorem 12.1 (Laminations give good pseudometric on surface) *Suppose* (λ_1, dx) *and* (λ_2, dy) *denote measured laminations binding* S*, and define*

$$d\rho^2 = dx^2 + dy^2$$

Then $d\rho$ *lifted to* $\mathbf{H}^2$ *defines a* $G(S)$*–invariant pseudo-metric* ρ *on* $\mathbf{H}^2$ *quasicomparable with the hyperbolic metric* d_H*.*

Proof If $x \in \mathbf{H}^2$, define $\rho(x,x) = 0$. If x and y are distinct points of $\mathbf{H}^2$ and γ is the geodesic segment joining x to y, define

$$\rho(x,y) = \int_{\gamma} d\rho.$$

Then ρ clearly satisfies the conditions for a pseudometric: $0 = \rho(x,x), 0 \leq \rho(x,y) < \infty, \rho(x,y) = \rho(y,x)$, and $\rho(x,z) \leq \rho(x,y) + \rho(y,z)$.

The quasicomparability of ρ and d_H entails the existence of positive constants K and K' such that

$$(*) \qquad \max(\rho, d_H) > K' \quad \Rightarrow \quad (1/K) d_H \leq \rho \leq K d_H.$$

We prove the existence of K and K' as follows.

By Theorem 11.3 (Continuity property 1) there is a constant D such that each geodesic segment of d_H–length not exceeding 1 has ρ–length not exceeding D. Thus if $d_H(x,y) > 1$, with $d_H(x,y) = n + k$, $0 \leq k < 1$, $n \leq 1$,

$$\rho(x,y) \leq (n+1)D = \frac{n+1}{n+k} D d_H(x,y) \leq 2D d_H(x,y).$$

This proves the second inequality of formula $(*)$. The first inequality of $(*)$ follows in an exactly analogous way from the following assertion: there is a positive integer N

and a positive integer i such that each geodesic segment γ of d_H–length exceeding N has ρ–length exceeding $1/i$. Indeed, if the assertion is correct, let $K' > K > 2Ni$. If $\rho(x,y) \le d_H(x,y) > K'$, then $d_H(x,y) = (n+k)N$, $k \le 1$ positive and $n \ge 1$ integral. Since $d_H(x,y) > n \cdot N$, our assertion implies that $\rho(x,y) > n \cdot (1/i)$ so that

$$\frac{d_H(x,y)}{K} < \frac{nN + kN}{2Ni} = \frac{n+k}{2i} \le \frac{n}{i} < \rho(x,y).$$

This would prove the first inequality of formula $(*)$.

There are a number of steps in the proof of the assertion. Assume the assertion false.

(1) Fix N and for each i let $\gamma(N,i)$ denote a geodesic segment of d_H–length N and ρ–length $\le 1/i$. Without loss, the segments $\gamma(N,i)$ converge to a segment $\gamma(N)$ of d_H–length N. By Theorem 11.4 (Continuity property 2) $\gamma(N)$ has ρ–length $\le C$, C a suitably chosen universal constant.

(2) Now let N vary. Without loss $\gamma(1), \gamma(2), \dots$ converge to a geodesic ray γ of ρ–length $\le 2C$, by Theorem 11.4 (Continuity property 2).

(3) We complete the proof by showing that every geodesic ray γ has infinite ρ–length. Indeed, let δ denote a limit geodesic for γ. (Section 8, Hyperbolic surfaces and limit geodesics). By Theorem 10.1 (Laminations binding a surface) item (3), δ crosses a leaf of $\lambda_1' \cup \lambda_2'$. By the definition of limit geodesic, γ crosses a fixed neighborhood of δ infinitely often. Since the support of dx is $|\lambda_1'|$ and the support of dy is $|\lambda_2'|$, it follows easily that $\int_\gamma d\rho = \infty$. This contradiction completes the proof of the theorem. □

13 Pseudo-Anosov diffeomorphisms

A diffeomorphism $\phi\colon S \to S$ is termed *pseudo-Anosov* provided that there exist measured laminations (λ_1, dx) and (λ_2, dy) binding S and a positive number $k > 1$ such that $\phi(|\lambda_i|) = |\lambda_i|$, $\int_{\phi(\gamma)} dx = 1/k \int_\gamma dx$ and $\int_{\phi(\gamma)} dy = k \int_\gamma dy$ for each path γ on S. The lamination λ_1 is called the *stable lamination* and λ_2 the *unstable lamination* for ϕ.

For the remainder of this section we fix a pseudo-Anosov diffeomorphism ϕ with associated laminations (λ_1, dx) and (λ_2, dy) and multiplier $k > 1$. We define $d\phi^2 = dx^2 + dy^2$ as in Section 12 (Measured laminations binding a surface).

Theorem 13.1 (Stable laminations irreducible) *The laminations λ_1 and λ_2 have no simple closed leaves. The measures dx and dy have no atoms. The pseudometric $\rho = \int d\rho$ is continuous on $\mathbf{H}^2 \times \mathbf{H}^2$.*

Proof The third assertion clearly follows from the second. The second assertion is, by Theorem 11.1 (Atoms and closed curves) a consequence of the first. Hence, it suffices to show that neither λ_1 nor λ_2 has a simple closed leaf.

Suppose, on the contrary, that λ_1 had a simple closed leaf L. The surface S cannot contain infinitely many disjoint, nonparallel, nontrivial simple closed curves. Hence, for some positive integer n, L and $\phi^n L$ are parallel closed geodesics, hence equal. But then

$$0 < \int_L dy = \int_{\phi^n(L)} dy = k^n \int_L dy < \infty,$$

a contradiction. □

14 Cellular decompositions of two-manifolds

A *decomposition* G of a space X is simply a collection of disjoint nonempty sets whose union is X. With each decomposition G of X there is an associated *decomposition space* or *identification space* X/G and *identification map* $\pi\colon X \to X/G$; the elements of G are the points of X/G; $\pi(x)$ is the unique element of G containing x; $U \subset X/G$ is open if and only if $\pi^{-1}(U)$ is open in X.

A decomposition G of X satisfies the *upper semicontinuity property* provided that, given $g \in G$ and V open in X containing g, the union of those $g' \in G$ contained in V is an open set in X. Equivalently, π is a closed map.

A decomposition G of a 2–manifold-without-boundary M is *cellular* provided G is upper semicontinuous and provided each $g \in G$ is compact, connected and has a nonseparating embedding in the Euclidean plane E^2.

The following theorem was proved by R L Moore [5] for the case $M = \mathbf{S}^2$ or E^2 and was extended to arbitrary 2–manifolds-without-boundary by Roberts and Steenrod.

Theorem 14.1 (Approximating cellular maps) *Let G denote a cellular decomposition of a 2–manifold M without boundary. Then the identification map $\pi\colon M \to M/G$ can be approximated by homeomorphisms. In particular, M and M/G are homeomorphic.*

For the remainder of this section we fix two geodesic laminations λ_1 and λ_2 of S satisfying the following three conditions:

(1) λ_1 and λ_2 bind S.

(2) λ_1 and λ_2 have no isolated leaves.

(3) λ_1 and λ_2 each satisfy conclusion (1) of Theorem 11.2 (Zero-angle).

We note that if λ_1 and λ_2 are the laminations associated with a pseudo-Anosov diffeomorphism of S, then λ_1 and λ_2 satisfy (1), (2) and (3) by Theorems 11.2 (Zero-angle) and 13.1 (Stable laminations irreducible).

With λ_1 and λ_2 we associate three decompositions $\mathbf{S}^2(\lambda_1), \mathbf{S}^2(\lambda_2), \mathbf{S}^2(\lambda_1, \lambda_2)$ of $\mathbf{S}^2$; a decomposition $\mathbf{H}^2(\lambda_1, \lambda_2)$ of $\mathbf{H}^2$; and a decomposition $S(\lambda_1, \lambda_2)$ of S, all satisfying the conditions of R L Moore's cellular decomposition theorem, as follows.

$\mathbf{S}^2(\lambda_1)$: Consider the 2–sphere $\mathbf{S}^2$ as the union of two copies $\mathbf{B}^2_-$ and $\mathbf{B}^2_+$ of the closed disk $\mathbf{B}^2 = \mathbf{H}^2 \cup \mathbf{S}^1_\infty$, $\mathbf{B}^2_-$ intersecting $\mathbf{B}^2_+$ in the common boundary $\mathbf{S}^1_\infty$. Lift λ_1 to λ'_1 in $\mathbf{H}^2$ and consider λ'_1 as a subset of the copy $\mathbf{B}^2_-$ of $\mathbf{B}^2$. An element g of the decomposition $\mathbf{S}^2(\lambda_1)$ of $\mathbf{S}^2$ is of one of three types:

(i) g is the closure in $\mathbf{B}^2_- \subset \mathbf{S}^2$ of a component of $\mathbf{H}^2_- - \lambda'_1$;

(ii) g is the closure in $\mathbf{B}^2_- \subset \mathbf{S}^2$ of a leaf of λ'_1 not contained in an element of type (i);

(iii) g is a singleton not contained in an element of types (i) or (ii).

No two elements of type (i) intersect by conditions (2) and (3) on λ_1. No element of type (i) intersects an element of type (ii) by condition (3) on λ_1. Hence $\mathbf{S}^2(\lambda_1)$ is a decomposition of $\mathbf{S}^2$ into disjoint, hyperbolically convex, hence cellular, compacta. Upper semicontinuity is easily checked by means of the fact that $|\lambda|$ is closed in S.

$\mathbf{S}^2(\lambda_2)$: The construction is exactly like that of $\mathbf{S}^2(\lambda_1)$ except that $|\lambda'_2|$ is taken to lie in $\mathbf{B}^2_+$.

$\mathbf{S}^2(\lambda_1, \lambda_2)$: The elements are either of type (i) or (ii) in $\mathbf{B}^2_-$ from $\mathbf{S}^2(\lambda_1)$, or of type (i) or (ii) in $\mathbf{B}^2_+$ from $\mathbf{S}^2(\lambda_2)$, or a singleton from $\mathbf{S}^1_\infty$ not contained in any other element. That no two elements of $\mathbf{S}^2(\lambda_1, \lambda_2)$ intersect follows easily from Theorem 10.2 (Neighborhoods at infinity). Upper semicontinuity is checked as before.

$\mathbf{H}^2(\lambda_1, \lambda_2)$: We now consider both λ'_1 and λ'_2 in the same copy of $\mathbf{H}^2$. An element g of the decomposition $\mathbf{H}^2(\lambda_1, \lambda_2)$ is of one of three types:

(i) g is the closure in $\mathbf{H}^2$ of a component of $\mathbf{H}^2 - \big(|\lambda'_1| \cup |\lambda'_2|\big)$.

(ii) g is the closure in $\mathbf{H}^2$ of a component of $|\lambda'_1| - |\lambda'_2|$ or of a component of $|\lambda'_1| - |\lambda'_2|$ not contained in an element of type (i).

(iii) g is the intersection of a leaf of λ'_1 and a leaf of λ'_2 not contained in an element of types (i) or (ii).

It follows from Theorem 10.1 (Laminations binding a surface) that the elements of $\mathbf{H}^2(\lambda_1, \lambda_2)$ are cellular. It follows from condition (2) on λ_1 and λ_2 that no two elements of $\mathbf{H}^2(\lambda_1, \lambda_2)$ intersect. Again upper semicontinuity is easily checked.

$S(\lambda_1, \lambda_2)$: We claim that each element g of $\mathbf{H}^2(\lambda_1, \lambda_2)$ has a neighborhood in $\mathbf{H}^2$ that is embedded in S by the covering map $p\colon \mathbf{H}^2 \to S$. Since p is a covering map, it suffices to show that the element g itself is embedded by p. The key fact is that no leaf of λ_1 or λ_2 is isolated. Using this fact one checks that, if $\tau\colon \mathbf{H}^2 \to \mathbf{H}^2$ is a covering translation that takes a point x of g to a point y of g, then either τ does not leave λ_1' and λ_2' invariant, or τ shows the existence of a closed geodesic not crossed by any leaf of $\lambda_1' \cup \lambda_2'$, a contradiction.

Hence the decomposition $\mathbf{H}^2(\lambda_1, \lambda_2)$ of $\mathbf{H}^2$ induces a cellular decomposition $S(\lambda_1, \lambda_2)$ of S equivariantly covered by $\mathbf{H}^2(\lambda_1, \lambda_2)$.

Finally, if we make the additional supposition,

(4) λ_1 and λ_2 support transverse measures dx and dy without atoms,

then $\rho = \int (dx^2 + dy^2)^{1/2}$ induces a continuous pseudometric on $\mathbf{H}^2$ and S that pushes down to a true metric on $\mathbf{H}^2/\mathbf{H}^2(\lambda_1, \lambda_2)$ and $S/S(\lambda_1, \lambda_2)$. Indeed, the elements of the decompositions $\mathbf{H}^2(\lambda_1, \lambda_2)$ are precisely the equivalence classes of points at zero distance from one another. The images in $S/S(\lambda_1, \lambda_2)$ of λ_1 and λ_2' are transverse measured foliations on the closed surface $S/S(\lambda_1, \lambda_2)$ in the sense of [4].

We shall use these decompositions in describing the topology of the Peano curve.

15 The topology of the Peano curve

The Peano curve is exactly that induced by the decomposition obtained by collapsing the two laminations. It is technically easier at this point to work with the two foliations obtained by collapsing the complementary domains of the union of the two laminations.

The main theorem supplies the map j in the following commutative diagram:

$$
\begin{array}{ccccccc}
\mathbf{B}^2 & & & \xrightarrow{\quad j \quad} & & & \mathbf{B}^3 \\
\uparrow & & & & & & \uparrow \\
\mathbf{H}^2 & \longrightarrow & \mathbf{H}^2 \times \{0\} & \overset{i}{\hookrightarrow} & \mathbf{H}^2 \times \mathbf{R} & \xrightarrow{Q} & \mathbf{H}^3 \\
\downarrow & & & & \downarrow & & \downarrow \\
S & & \overset{h}{\hookrightarrow} & & M & \equiv & M
\end{array}
$$

Vertical arrows are either inclusions or universal coverings; S is the hyperbolic fiber of the fibered manifold M; Q is the covering-induced quasi-isometry between $\mathbf{H}^2 \times \mathbf{R}$, with the invariant metric

$$ds^2 = (k^t dx)^2 + (k^{-t} dy)^2 + (\log k \cdot dt)^2,$$

and $\mathbf{H}^3$, with its standard hyperbolic metric; $\mathbf{B}^2$ and $\mathbf{B}^3$ are the natural compactifications of hyperbolic space that add the circle or sphere at infinity.

The restriction $f = j|\partial\mathbf{B}^2 \colon \partial\mathbf{B}^2 \to \partial\mathbf{B}^3$ is the Peano curve that we wish to analyze more precisely. We compactify $\mathbf{H}^2 \times \mathbf{R}$ in the following way: $\mathbf{H}^2 \times \mathbf{R} = \operatorname{Int}(\mathbf{B}^2 \times I)$, $\mathbf{H}^2 = \operatorname{Int}\mathbf{B}^2$, $\mathbf{R} = \operatorname{Int} I$, $I = [-\infty, +\infty]$. Then the analysis of f proceeds in the following steps.

(1) Extend $f \colon (\partial\mathbf{B}^2 = \partial(\mathbf{B}^2 \times \{0\})) \to \partial\mathbf{B}^3$ to a map $f^* \colon \partial(\mathbf{B}^2 \times I) \to \partial\mathbf{B}^3$.

(2) Define a cellular, upper semicontinuous decomposition G of the 2–sphere $\partial(\mathbf{B}^2 \times I)$ that shrinks the leaves of the two foliations as asserted.

(3) Show that $f^* \colon \partial(\mathbf{B}^2 \times I) \to \partial\mathbf{B}^2$ factors through the decomposition space projection $p \colon \partial(\mathbf{B}^2 \times I) \to \partial(\mathbf{B}^2 \times I)/G$:

$$\begin{array}{ccc} \partial(\mathbf{B}^2 \times I) & \xrightarrow{\ f^*\ } & \partial(\mathbf{B}^3) \\ \big\downarrow p & & \big\uparrow q \\ \partial(\mathbf{B}^2 \times I)/G & \equiv & \partial(\mathbf{B}^2 \times I)/G \end{array}$$

(4) Show that q is a homeomorphism.

These four steps show that the topological model for the Peano curve is as asserted.

Step 1 Definition of the extension f^*

If p is a point of $(\partial\mathbf{B}^2) \times I$, let $r(p)$ denote any ray in $\mathbf{H}^2 \times \mathbf{R}$ with infinite endpoint at p. If p is an element of $\mathbf{H}^2 \times \{-\infty, +\infty\}$, let $r(p)$ denote a vertical ray in $\mathbf{H}^2 \times \mathbf{R}$ with infinite endpoint at p. Then $Q(r(p))$ is a ray in $\mathbf{H}^3$ that clusters at some nonempty set in $\partial\mathbf{B}^3$. Define $f^*(p)$ to be that set.

At this point we know from the proof of the main theorem that f^* is a well-defined continuous function when restricted to $(\partial\mathbf{B}^2) \times I$. We do not know that it is single-valued on $\mathbf{H}^2 \times \{-\infty, +\infty\}$, nor do we know that it is continuous on $\mathbf{B}^2 \times \{-\infty, +\infty\}$. These defects will be cared for as we examine Step 3.

Step 2 The cellular decomposition

The appropriate cellular, upper semicontinuous decomposition G of $\partial(\mathbf{B}^2 \times I)$ is defined as follows. We lift the transverse measured foliations of our pseudo-Anosov diffeomorphism of S to $\mathbf{H}^2 = \mathrm{Int}\mathbf{B}^2$ via the inverse of the universal covering map $\mathrm{Int}\mathbf{B}^2 = \mathbf{H}^2 \to S$ and compactify by adding all cluster points at infinity to each leaf. Let F_1 and F_2 denote these compactified foliations. (We may assume that dx is the measure transverse to F_1 and dy is the measure transverse to F_2.) If L is a leaf of F_1, then a segment of L has zero measure with respect to dx and has length measured entirely by dy. Recalling the invariant metric

$$ds^2 = (k^t dx)^2 + (k^{-t} dy)^2 + (\log k \cdot dt)^2,$$

on $\mathbf{H}^2 \times \mathbf{R}$, where t is the standard variable and dt is the standard measure on the real line, $\mathrm{Int}I$, we note that this measure contracts as this leaf L is pushed toward $t = \infty$ through the product structure on $\mathbf{H}^2$. Let the union

$$(L \times \infty) \cup ((\partial L) \times I)$$

be an element of G. Similarly, if L' is a leaf of F_2, then the leaf contracts in measure as it is pushed toward $t = -\infty$. Let the union

$$(L' \times -\infty) \cup ((\partial L') \times I)$$

be an element of G. Elements of G of these two types fill up all of $(\partial\mathbf{B}^2 \times I)$ with the exception of a dense G_δ subset of $(\partial\mathbf{B}^2) \times I$. Each vertical fiber $p \times I$ of that subset is also to denote an element of G.

Preamble to steps 3 and 4: The quasi-isometric extension principle

Our arguments are based on the principles:

($*$) Quasigeodesic sets near infinity have uniformly small Euclidean diameter. In particular:

($**$) If hyperbolic k–space is mapped quasi-isometrically into hyperbolic n–space, then the mapping may be extended continuously so as to map the $k-1$–sphere at infinity injectively into the $(n-1)$–sphere at infinity.

We outline the easy argument that deduces ($**$) from ($*$):

If p is a point at infinity in hyperbolic k–space, let L denote a geodesic or quasi-geodesic ray with infinite endpoint at p. The images of terminal subrays of L have, by principle ($*$), ever decreasing Euclidean diameter in hyperbolic n–space. Hence they must cluster at a unique point p' of the $(n-1)$–sphere at infinity. We define p' as the image of p.

By examining the quasi-isometric images of decreasing half-k–spaces containing p, one easily deduces from $(*)$ the continuity of the extended function.

In order to prove that the extended function is injective, one only need consider a geodesic L joining any two points, p and q, in the $((k-1)$–sphere at infinity. The quasi-isometric image of L lies in a bounded hyperbolic neighborhood of the geodesic L' in hyperbolic n–space joining the images, p' and q', of p and c respectively. If p' and q' coincide, it follows that the image of L lies at infinity in hyperbolic n–space, a contradiction.

Step 3 Factoring f^* through the decomposition space projection

With the quasi-isometric extension principle in hand, we are ready to show

(i) f^* is well-defined (that is, single-valued);

(ii) f^* factors through the decomposition space projection;

(iii) f^* is continuous.

Proof of (i) Let $p \in \mathbf{H}^2$. Since a vertical line $L = \{p\} \times (-\infty, \infty)$ is isometric with hyperbolic 1–space, $Q(L)$ lies within a bounded distance of a true hyperbolic geodesic in $\mathbf{H}^3$ and, by the quasi-isometric extension principle, induces an embedding of its two points at infinity. This establishes that f^* is single valued at the points where that had not previously been established.

Proof of (ii) Take a point p in $\mathbf{H}^2$. The point p is in a unique fiber L_1 of the compactified foliation F_1, and in a unique fiber L_2 of the compactified foliation F_2. We first consider the point $p^+ = \{p\} \times \infty$. Let L_1' denote a maximal arc in L_1 that passes through p. The vertical disk $L_1' \times I$ has induced metric of the form

$$ds^2 = (k^{-t} dy)^2 + (\log k \cdot dt)^2,$$

since the metric dx is 0 along fibers of F_1. The substitution, $T = k^t$, transforms this metric into the metric

$$ds^2 = (dy^2 + dT^2)/(T^2),$$

the standard metric for hyperbolic 2–space in the upper half-plane model. Hence the quasi-isometric extension principle applies to the restriction of Q to this set. Hence, if C is the complement in this hyperbolic 2–space of a very large circle perpendicular to the boundary of that half-space model, then the image of C in hyperbolic 3–space is quasigeodesic and near infinity, hence of small Euclidean diameter. Letting C approach infinity, we find that there is a unique limit point at infinity in hyperbolic 3–space to which we should map p^+. Furthermore, that point is the same point to which we have already mapped $(\partial L_1') \times I$. For the point $p^- = \{p\} \times -\infty$, one uses the fiber L_2 of F_2 through p. One concludes that each element of G is identified under f^* to a single point of $\partial \mathbf{B}^3$. This proves (ii).

Proof of (iii) The continuity of the extended map f^* of f follows from the facts that

(iv) the decomposition space of $(\partial \mathbf{B}^2) \times \{0\}$ induced by the decomposition G of $\partial(\mathbf{B}^2 \times I)$ is naturally homeomorphic with the decomposition space of $\partial(\mathbf{B}^2 \times I)$ by G; and

(v) f is continuous.

(Fact (iv) is evident since every element of G intersects $(\partial \mathbf{B}^2) \times \{0\}$; fact (v) was the content of the main theorem.)

Step 4 q is a homeomorphism

We recall the argument of the preceding paragraph which considered vertical disks: Take a maximal arc L_1' in a single leaf L_1 of the compactified foliation F_1. The vertical disk $L_1' \times I$ has induced metric of the form

$$ds^2 = (k^{-t} dy)^2 + (\log k \cdot dt)^2,$$

since the metric dx is 0 along fibers of F_1. The substitution, $T = k^t$ transforms this metric into the metric

$$ds^2 = (dy^2 + dT^2)/(T^2),$$

the standard metric for hyperbolic 2–space in the upper half-plane model. The quasi-isometric extension principle applies to the mapping Q restricted to this 2–space and shows that f^* embeds the circle at infinity. This embedding shrinks the entire element

$$((\partial L_1') \times I) \cup (L_1' \times \{+\infty\})$$

of G to a single point in $\partial \mathbf{B}^3$ and identifies precisely the endpoints of the arc

$$L_1' \times \{-\infty\}$$

in forming an embedded simple closed curve in $\partial \mathbf{B}^3$.

This argument shows that we have a large family of disks embedded by the union of the mappings Q and f^*. These embedded disks will suffice to show that the factor mapping q is $1-1$, hence an embedding.

We have defined for every point p of $\partial(\mathbf{B}^2 \times I)$ a ray $r(p)$ in $\mathbf{H}^2 \times \mathbf{R}$ with infinite endpoint at p such that the mapping $Q\colon \mathbf{H}^2 \times \mathbf{R} \to \mathbf{H}^3$ takes the ray to a set in $\mathbf{H}^3$

that clusters only at $f^*(p)$ in the 2–sphere at infinity. We shall find this ray useful in our proof that f^* identifies two points of $\partial(\mathbf{B}^2 \times I)$ only if they lie in the same element of G.

Let a and b be two points of $\partial(\mathbf{B}^2 \times I)$ that are identified by f^*. Since every element of G intersects $(\partial\mathbf{B}^2) \times I$ and since f^* identifies each element of G to a point, we may assume that a and b are in the set $(\partial\mathbf{B}^2) \times I$.

We claim that

Fact 1 There are points a' and b' in $\mathbf{H}^2 \times \{-\infty, +\infty\}$ arbitrarily close (with respect to the hyperbolic metric d_H or the pseudometric ds) to a and b, respectively, such that all of a, a', b and b' are identified by f^*.

Fact 2 The points a' and b' lie in the same element of G.

Assuming these two facts for the moment, then taking limits as a' approaches a and b' approaches b, we find that a and b are in a common element of G. Hence it suffices to establish these two facts.

Proof of Fact 1 Let L be a leaf of one of the two foliations whose corresponding vertical disk lies near the vertical fiber through a and separates a from b. The image of L is a very small embedded disk in the compactification $\mathbf{B}^3$ of $\mathbf{H}^3$ that separates terminal rays of the image of $r(a)$ from terminal rays of the image of $r(b)$. But $f^*(a) = f^*(b)$ is in the closure of both the image of $r(a)$ and of the image of $r(b)$. It follows that $f^*(a)$ lies in the boundary of that separating disk. The preimage of that boundary point may always be taken to lie in the interior of one of the open disks at the end of $\mathbf{B}^2 \times I$. The preimage may be taken as a'. One obtains b' in a similar way.

Proof of Fact 2 Let a_+ be the projection of a' into $\mathbf{H}^2 \times \infty$ and a_- the projection into $\mathbf{H}^2 \times -\infty$. Define b_+ and b_- similarly. We assume that $a' = a_+$. We shall show that:

Fact 3 The fiber of $F_1 \times \infty$ that contains a_+ also contains b_+.

Assuming Fact 3 for the moment, we complete the proof of Fact 2. Fact 3 implies that a_+ and b_+ are identified and lie in the same element of G. If $b' = b_+$, then we are done. If not, then $b' = b_-$, and we discover that b_- is identified to a_+, which is identified to b_+. But that says that b_+ and b_- are identified, a contradiction to the quasi-isometric extension principle.

Proof of Fact 3 We first note that the leaf L_a of $F_2 \times \infty$ through a_+ does not contain b_+, for that would imply that a_+ is identified with no other point of $\partial(L \times I)$ under

f^* in particular is not identified with b', a contradiction. Let L_b denote the leaf of $F_2 \times \infty$ that contains b_+. The leaves of $F_2 \times \infty$ that separate a_+ from b_+ in $\mathbf{B}^2 \times \infty$ are linearly ordered, and form an open arc of leaves with L_a and L_b as endpoints of that arc. The leaf K of $F_1 \times \infty$ that passes through a_+ intersects a last of these leaves in the linear order. We claim that this last leaf intersected is L_b. Suppose not. Let L' be the leaf. The leaf L' has maximal simple subarcs A and B having the property that A separates both b' and $L' \setminus A$ from a', while B separates both a' and $L' \setminus B$ from b'. The argument for Fact 1 shows that both $A \times I$ and $B \times I$ contain a boundary point identified with both a' and b', where the first of these may be taken in K. Since K does not have a point at infinity in common with a leaf of $F_2 \times \infty$, that point must lie $\mathrm{Int} A$. But no point of $\mathrm{Int} A$ is identified with any other point of $L' \times I$. Hence the point in $\mathrm{Int} A$ must coincide with point in B. But the leaf through any point in $\mathrm{Int} B$ continues into every domain of $\mathbf{H}^2 \times \infty$ whose closure contains that point. But the continuation intersects further leaves that separate a_+ and b_+, a contradiction. This completes the proof that K hits L_b, in fact at an interior point b'' of L_b. Hence b' and b'' are identified to a'. Thus, $b' = b''$ and $b'' = b_+$. This proves Fact 3.

16 The figure-eight-knot complement

When the surface S is not a closed, orientable, hyperbolic surface, but rather a punctured, orientable, hyperbolic surface, we conjecture that the main theorem is still true. As experimental evidence, we explain how to create computer approximations to the resulting 2–sphere-filling Peano curve in the case where S is a once-punctured torus T.

We begin with the simplest case, namely, the complement in $\mathbf{S}^3$ of the figure-eight knot K, Figure 2. The punctured torus T arises as the Seifert surface spanning K.

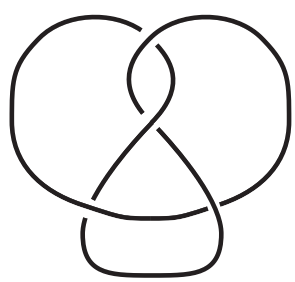

Figure 2: The figure-eight knot

We first describe an explicit hyperbolic structure on the 3–manifold $M = \mathbf{S}^3 \setminus K$. If we push two 3–cells together, one from above K and one from below K, so as to fill

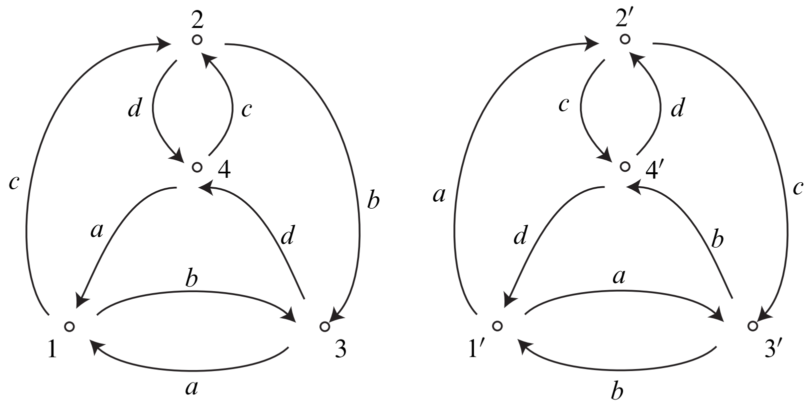

Figure 3: The figure-eight knot complement as an identification space created by sewing together two 3–cells

$\mathbf{S}^3 \setminus K$, we find that M is an identification space of two 3–cells C_1 and C_2, from each of which four vertices (boundary points) have been removed, Figure 3, under the face identifications:

$$\begin{pmatrix} 1 & c & 2 & d & 4 & a \\ 2' & c & 4' & d & 1' & a \end{pmatrix} \quad \begin{pmatrix} 2 & b & 3 & d & 4 & c \\ 3' & b & 4' & d & 2' & c \end{pmatrix}$$

$$\begin{pmatrix} 1 & b & 3 & d & 4 & a \\ 3' & b & 4' & d & 1' & a \end{pmatrix} \quad \begin{pmatrix} 1 & c & 2 & b & 3 & a \\ 2' & c & 3' & b & 1' & a \end{pmatrix}$$

$$\begin{pmatrix} 1 & b & 3 & a \\ 3' & b & 1' & a \end{pmatrix} \quad \begin{pmatrix} 4 & c & 2 & d \\ 2' & c & 4' & d \end{pmatrix}$$

The last two face pairs are digons that may be collapsed to arcs. When this is done, the cells become ideal tetrahedra σ_1 and σ_2, Figure 4, with face identifications:

$$\begin{pmatrix} 1 & 2 & 4 \\ 2' & 4' & 1' \end{pmatrix} \quad \begin{pmatrix} 2 & 3 & 4 \\ 3' & 4' & 2' \end{pmatrix} \quad \begin{pmatrix} 1 & 3 & 4 \\ 3' & 4' & 1' \end{pmatrix} \quad \begin{pmatrix} 1 & 2 & 3 \\ 2' & 3' & 1' \end{pmatrix}$$

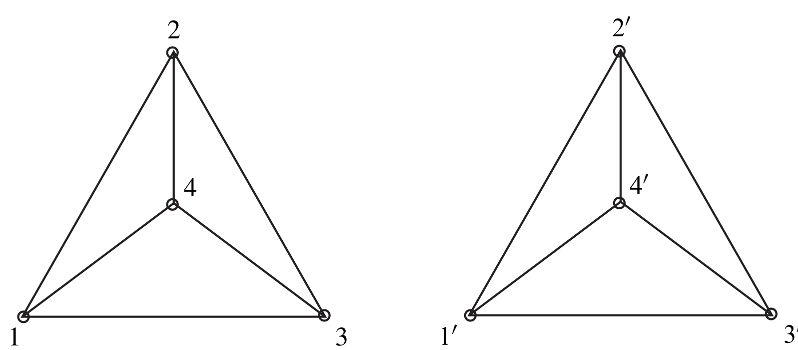

Figure 4: The figure-eight knot complement realized as the union of two ideal tetrahedra

The resulting edge cycles have length 6:

$$12 \to 2'4' \to 43 \to 1'4' \to 42 \to 2'3' \to 12$$
$$14 \to 2'1' \to 13 \to 3'4' \to 23 \to 3'1' \to 14$$

There is an ideal tetrahedron σ in $\mathbf{H}^3$ that is regular: Indeed, the regular Euclidean tetrahedron with vertices

$$(0,0,1) \quad (0, 2\sqrt{2}/3, -1/3) \quad (-\sqrt{6}/3, -\sqrt{2}/3, -1/3) \quad (\sqrt{6}/3, -\sqrt{2}/3, -1/3)$$

is inscribed in the unit ball $\mathbf{B}^3$ and may be considered as regular ideal hyperbolic in the Klein model.

The dihedral angles of σ are $\pi/3$: Indeed, pass to the upper-half-space model for $\mathbf{H}^3$ with one vertex of σ at ∞. The hyperbolic isometry that fixes ∞ and cyclically permutes the other three vertices (existence obvious in the Klein model) must be a Euclidean isometry of order 3, hence a rotation of $2\pi/3$. We conclude that the three planar vertices form an equilateral triangle. The three vertical dihedral angles emanating from ∞ are therefore obviously $\pi/3$.

We may thus realize M as the union of two regular ideal hyperbolic tetrahedra and the face identifications as hyperbolic isometries of ideal triangles. Since the edge cycles have length 6, the total hyperbolic angle at each of the two edges of M is $6 \cdot (\pi/3) = 2\pi$. We thus obtain a smooth, complete hyperbolic structure on M, with no singularities.

We obtain $T \subset M$ as follows: The two faces $\begin{pmatrix}1 & 3 & 4\end{pmatrix} = \begin{pmatrix}3' & 4' & 1'\end{pmatrix}$ and $\begin{pmatrix}1 & 3 & 2\end{pmatrix}$ of the tetrahedron σ_1 form a quadrilateral Q (with ideal vertices) whose opposite sides are identified in pairs to form a punctured torus t. Since there are only two edge classes in M, the diagonal $\begin{pmatrix}1 & 3\end{pmatrix}$ of Q is actually also identified with one of the sides so that t is singular. If $\mathrm{Int}(Q)$ is pushed into $\mathrm{Int}(\sigma_1)$, this singularity of t is removed, and we obtain the punctured torus T. Usually, however, we shall work directly with t.

We identify the universal cover M' of M with the upper-half-space model of $\mathbf{H}^3$ and lift t to $t' \subset \mathbf{H}^3$ as follows: We lift one copy of $\sigma_1 = \begin{pmatrix}1 & 4 & 2 & 3\end{pmatrix}$ to $\sigma_1' = (\infty,\ 0,\ 1,\ \omega = (1/2) - (\sqrt{3}/2)i)$, and we lift one copy of $Q = (1\,3\,4) \cup (1\,2\,3) \subset \partial\sigma_1$ to $Q' = (\infty\,\omega\,0) \cup (\infty\,1\,\omega) \subset \partial\sigma_1'$. These normalizations completely determine the entire triangulations of $M' = \mathbf{H}^3$ by the lifts of σ_1 and σ_2 and the triangulation of the universal cover t' of t containing Q'.

We concentrate on those 3–simplexes of M' and those 2–simplexes of t' containing ∞: The former create a "plane" of 3–simplexes projecting naturally to an equilateral triangulation of the complex plane $C \subset \partial\mathbf{H}^3$, and the latter create a "zig-zag line" of 2–simplexes projecting naturally to a zig-zag line in the 1–skeleton of the triangulation of C; see Figure 5.

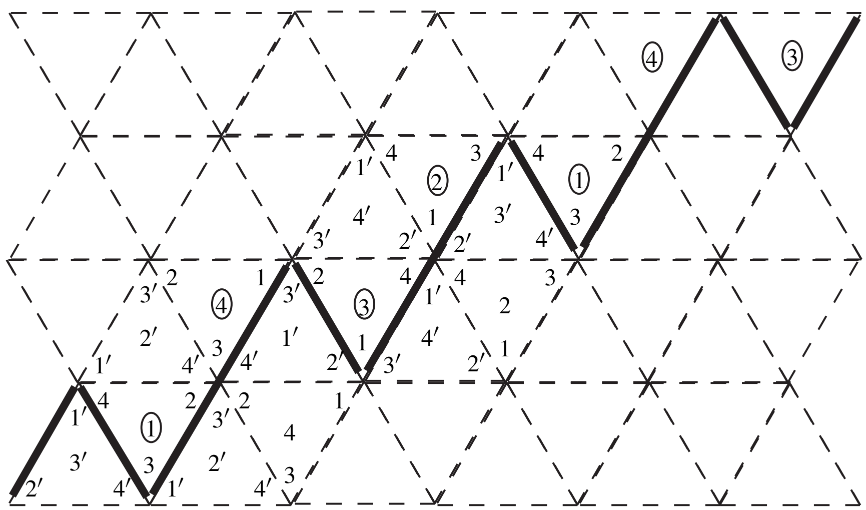

Figure 5: Projections into the complex plane of the tetrahedra and the triangles containing ∞
The zig-zag line associated with the figure-eight knot complement

Each triangle represents a tetrahedron lifted from either σ_1 or σ_2; the central number indicates the vertex at ∞. The zig-zag curve appears to be periodic of period 3 until one examines the labels that reveal the actual period of 6. Each of the six vertices in a period represents one of the six directed edges dual to the 1–skeleton of the triangulation of t; see Figure 6.

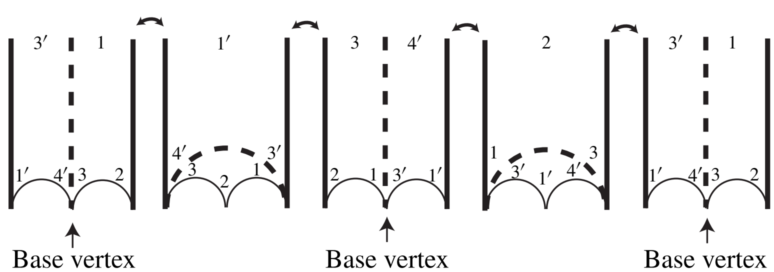

Figure 6: The identifications forming a period of a zig-zag curve (in the abstract)

Theorem 16.1 (Zig-zag curves) *The zig-zag curve determines t'.*

Proof Order the triangles of t' into subcollections t'_0, t'_1, t'_2, $\cdots$ where t'_0 consists of those triangles of t' containing ∞, t'_1 consists of those remaining triangles that share an edge with a triangle from t'_0, and, in general, t'_n consists of those remaining that share an edge with a triangle from t'_{n-1}, etc. [Note that in the abstract zig-zag curve above, in the second and fourth fundamental domains, the triangles $(1\,2\,3)$ and $(1'\,3'\,4')$, respectively, belong to t'_1 and not to t'_0.]

We have already placed each triangle of t'_0. Assume the triangles of t'_{n-1} placed and assume that Δ_1 is a triangle of t'_n sharing an edge with a triangle Δ_0 of t'_{n-1}. There is a pair $\Delta'_0 \cup \Delta'_1$ in the placement of t'_0 (one of six models) corresponding exactly to the union $\Delta_0 \cup \Delta_1$ so that the common edge comes from ∞. There is a unique hyperbolic isometry taking Δ'_0 to Δ_0 that preserves the correspondence. The image of the fourth vertex completes the placement of Δ_1. □

Construction *The image of t'_n has an "outer edge" furthest from the image of t'_0 in t'. It projects naturally to a polygonal curve in the complex plane C. This curve approximates our 2–sphere-filling Peano curve.*

17 Punctured torus bundles over $\mathbf{S}^1$

Let S denote the once-punctured torus, $\phi\colon S \to S$ a pseudo-Anosov diffeomorphism and M the resulting 3–manifold. By Thurston [7] or Sullivan [6], M admits a complete hyperbolic metric of finite volume. We may therefore identify the universal cover M' of M with hyperbolic 3–space $\mathbf{H}^3$. We may also identify the universal cover S' of S with hyperbolic 2–space $\mathbf{H}^2$. As noted before, the lift of $S' = \mathbf{H}^2$ to $M' = \mathbf{H}^3$ has the entire 2–sphere $\mathbf{S}^2 = \partial\mathbf{H}^3$ as its space at infinity. We conjecture that the map $S' = \mathbf{H}^2 \to \mathbf{H}^3$ extends continuously to a map $\mathbf{B}^2 \to \mathbf{B}^3$ so that the circle $\mathbf{S}^1 = \partial\mathbf{H}^2$ maps continuously onto the 2–sphere $\mathbf{S}^2 = \partial\mathbf{H}^3$. [Note added, 2007: This conjecture, in the case discussed in this section, is now known due to the work of Minsky and McMullen.] Assuming this conjecture, we describe the conjectural topological nature of the Peano curve and method for approximating this curve geometrically.

Topological description of the Peano curve Let λ and μ denote the measured foliations associated with the pseudo-Anosov diffeomorphism ϕ. Lift λ to a foliation λ' of $\mathbf{H}^2$ and complete λ' by adding to each leaf ℓ of λ' its limit points in the circle $\mathbf{S}^1 = \partial\mathbf{H}^2$ at infinity. View $\mathbf{B}^2 = \mathbf{H}^2 \cup \partial\mathbf{H}^2$ as the central slice in the 3–ball $\mathbf{B}^3$. Project the completed λ' upward onto the northern hemisphere $\mathbf{B}^2_+$ of $\mathbf{S}^2 = \partial\mathbf{B}^3$ to obtain a foliation λ''. Lift μ to a foliation μ' of $\mathbf{H}^2$, complete μ' as λ' was completed, and project the completed μ downward onto the southern hemisphere $\mathbf{B}^2_-$ of $\mathbf{S}^2$ to obtain a foliation μ'' of $\mathbf{B}^2_-$. Now declare any two points of $\mathbf{S}^2$ to be equivalent if they lie in either a single leaf of λ'' or a single leaf of μ''. Extend this notion of equivalence to an equivalence relation. In contrast to the case where S was a closed surface, corresponding to each ideal vertex of S', there will be an equivalence class that is a union of a countable infinity of completed leaves, put together at the ideal vertex so as to form a spider, with infinitely many arms that lie alternately in $\mathbf{B}^2_+$ and $\mathbf{B}^2_-$. Nevertheless, the equivalence classes will form a cellular upper-semicontinuous decomposition of $\mathbf{S}^2$, each element of which intersects the equator $\mathbf{S}^1$. By R L Moore's theorem, the quotient space is a 2–sphere, and the equator of the original 2–sphere maps onto the image 2–sphere continuously. This quotient map gives the conjectured topological description of the Peano curve.

The zig-zag curve associated with $S' \subset M' = \mathbf{H}^3$ As with the figure-eight knot, a fundamental domain for S may be realized as an ideal topological quadrilateral that is the union of two ideal triangles. We may assume that these two triangles lift to two ideal triangles in hyperbolic space so that the entire surface S' is triangulated by ideal triangles. (Again, there may be singularities which can be ignored.) We use the upper half-space model of $\mathbf{H}^3$, and we place one of the ideal vertices of S' at infinity. We let T_0 be the collection of triangles in S' that contain that particular vertex. Each projects to a line segment in the plane at infinity. These segments again form a zig-zag curve Z in the plane which has period 6. The curve Z completely determines the surface S' by simple algorithm, as in the case of the figure-eight knot. Hence one may use the curve Z to approximate the Peano curve. Thus it remains only to describe how one might find the curve Z.

Finding the zig-zag curve Z We find the zig-zag curve associated with M by analyzing the monodromy map ϕ. Any two homeomorphisms of S that are homotopic determine the same manifold M. Hence we seek to understand ϕ as an element of the orientation preserving mapping class group of the punctured torus. This mapping class group is isomorphic to $SL(2,Z)$.

The connection of $SL(2,Z)$ with homeomorphisms of the punctured torus S is realized in the following way. Each element of $SL(2,Z)$ is a 2×2 integer matrix of determinant 1. It defines a linear homeomorphism of the plane $\mathbf{R}^2$ which permutes the elements of the integer lattice Z^2 and pushes down to a homeomorphism of the torus $T = \mathbf{R}^2/Z^2$ and the punctured torus S that is the complement in T of the image of the lattice Z^2.

The group $SL(2,Z)$ is generated by two elements, called L and R (left and right) that have the following descriptions as matrices:

$$L = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix} \quad \text{and} \quad R = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$$

Thus we may assume that ϕ is given as a word in L and R.

We may lift a natural two-ideal-triangle triangulation T of S to $\mathbf{H}^3$ to obtain an ideal triangulation T' of one lift of S, and we may lift T to $\mathbf{R}^2$ to obtain an ideal triangulation T'' of $\mathbf{R}^2 \setminus Z^2$. We examine the action of L and R on these triangulation.The matrices L and R act on the triangulation of $\mathbf{R}^2$ as in Figures 7 and 8.

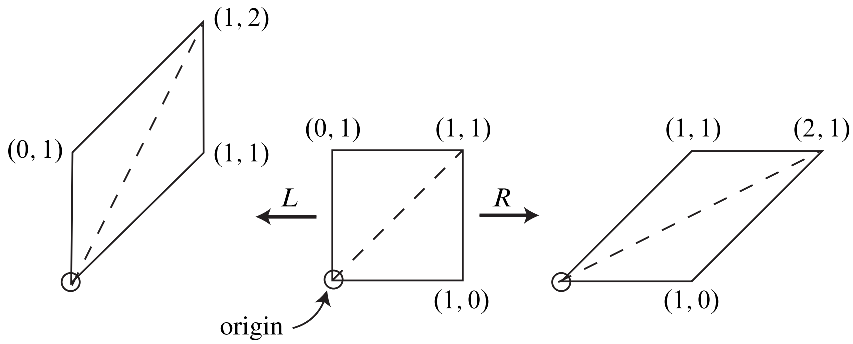

Figure 7: The action of L and R on the ideal triangulation

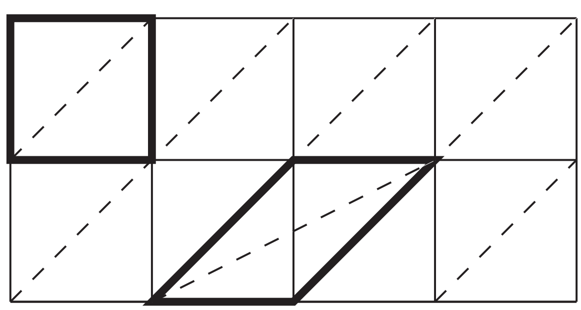

Figure 8: The effect of the transformation R
Left: An old fundamental domain subdivided by dotted line into two ideal triangles
Right: A new fundamental domain subdivided by dotted line into two new ideal triangles Notice the implied tetrahedron.

The vertices of the triangulation T'' remain unchanged, but the fundamental quadrilateral is skewed by R or by L. In $\mathbf{H}^3$, the ideal vertices of T' remain unchanged, but the pair of triangles in T' viewed as fundamental domain changes. The two triangles used as new fundamental domain form two of the four sides of an ideal tetrahedron. One replaces those two triangles by the complementary two sides of the tetrahedron and extends equivariantly to all similar pairs. As a consequence, the surface T' changes, as does the zig-zag curve.

Composing the actions of the various letters $R^{\pm 1}$ and $L^{\pm 1}$ in the expression of ϕ, one moves the zig-zag curve to a new position that is the exact image of the old under ϕ, where ϕ is considered as an element of the Kleinian group $\pi_1(M)$. Since ϕ must be a parabolic element fixing ∞, the two curves must be parallel. In the next paragraph, we describe exactly how one can calculate the new curve after action of R or L. The procedure is outlined in Figure 9.

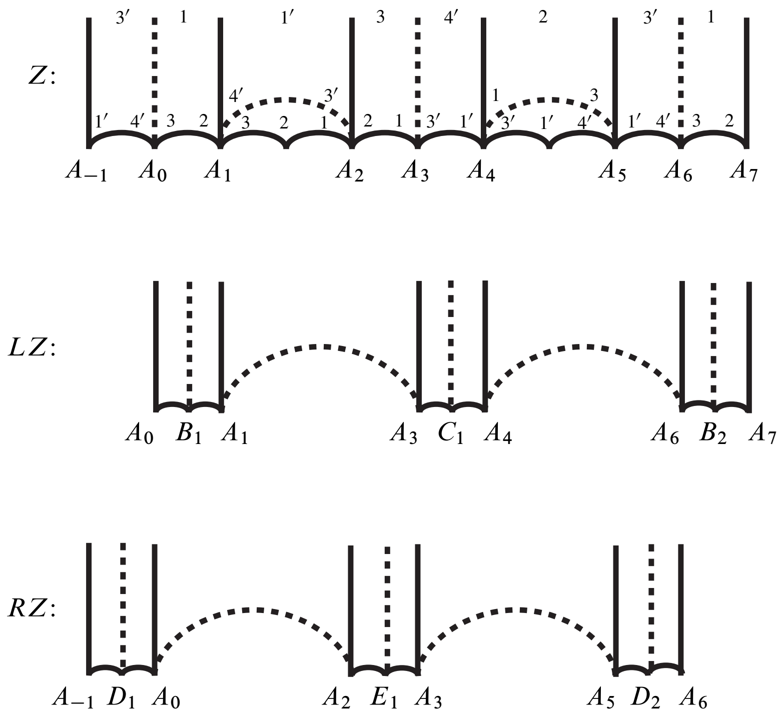

Figure 9: Actions of R and L on a zig-zag curve Z: Letters label the vertices of Z in the plane. Numbers label the vertices of the fundamental domain of the punctured torus. Vertices at the top are at infinity. Solid lines delineate copies of the fundamental domain. Dotted lines divide those copies into triangles.

The zig-zag curve is periodic of period 6. A single period involves four lifts of the quadrilateral that is the fundamental domain. Five lifts are pictured in the diagram. The dotted lines divide the five quadrilaterals into triangles. Notice that the central lower vertex of the second and fourth lift are not vertices of the zig-zag curve.

The action of L removes from a period the vertices labelled A_2 and A_5 and inserts new vertices B_1 and C_1, as pictured. The action of R removes vertices A_1 and A_4 and inserts new vertices D_1 and E_1, as pictured.

The new vertices become the base vertices of the image diagram.

The order in which one composes these actions is important. If $\phi = \phi_1 \circ \phi_2 \circ \cdots \circ \phi_k$, as a standard right-to-left composition of linear maps or multiplication of matrices, where each ϕ_i is either $R^{\pm 1}$ or $L^{\pm 1}$, then, as action on zig-zag curves, one must compose from left to right (column operations on the bases for a planar lattice). Failure to notice this distinction can be very confusing. Thus, for example, with the complement of the figure-eight knot, where $\phi = L \circ R$, the action on zig-zag curves is $\mathrm{Action}(R) \circ \mathrm{Action}(L)$.

Our task is simply to describe how the new vertices B_1, C_1, D_1 and E_1 are found. Since the new vertices are vertices of the original lift t' of t, they may be found by equivariance: one maps one triangle of a fundamental domain by the appropriate linear fractional transformation, and the fourth vertex will be carried to the new vertex. (Then, of course, one pushes the surface across the four-vertex tetrahedron to find the new surface and new zig-zag curve.) Each of the four linear fractional transformations needed has the same form: the image and preimage of ∞ are prescribed, and the image of one other point is prescribed. We describe the transformation as a permutation, the upper row giving the three domain points, the lower row giving the corresponding image points. Here is the common form:

$$\begin{pmatrix} \alpha & \beta & \infty \\ \infty & \gamma & \delta \end{pmatrix} = \frac{\delta z - \delta\beta + \gamma(\beta - \alpha)}{z - \alpha}$$

We require four of these maps:

$$\begin{pmatrix} A_3 & A_2 & \infty \\ \infty & A_1 & A_0 \end{pmatrix} = \frac{A_0 z - A_0 A_2 + A_1(A_2 - A_3)}{z - A_3} \colon A_1 \mapsto B_1$$

$$\begin{pmatrix} A_1 & A_2 & \infty \\ \infty & A_3 & A_4 \end{pmatrix} = \frac{A_4 z - A_4 A_2 + A_3(A_2 - A_1)}{z - A_1} \colon A_3 \mapsto C_1$$

$$\begin{pmatrix} A_2 & A_1 & \infty \\ \infty & A_0 & A_{-1} \end{pmatrix} = \frac{A_{-1} z - A_{-1} A_1 + A_0(A_1 - A_2)}{z - A_2} \colon A_0 \mapsto D_1$$

$$\begin{pmatrix} A_5 & A_4 & \infty \\ \infty & A_3 & A_2 \end{pmatrix} = \frac{A_2 z - A_2 A_4 + A_3(A_4 - A_5)}{z - A_5} \colon A_3 \mapsto E_1$$

The first takes A_1 to B_1. The second takes A_3 to C_1. The third takes A_0 to D_1, the second takes A_3 to E_1.

How does one find the zig-zag curve in general? Periodicity allows one to identify the zig-zag curve with six complex numbers, hence as a point of complex six-dimensional space C^6. Both R and L, hence also the monodromy map ϕ, act on zig-zag curves, hence on the space of all zig-zag curves. The desired zig-zag curve is a fixed point under this action. Starting with the known solution to the fixed point problem for the figure-eight knot, one can work one's way toward longer words in R and L by means of Newton's method.

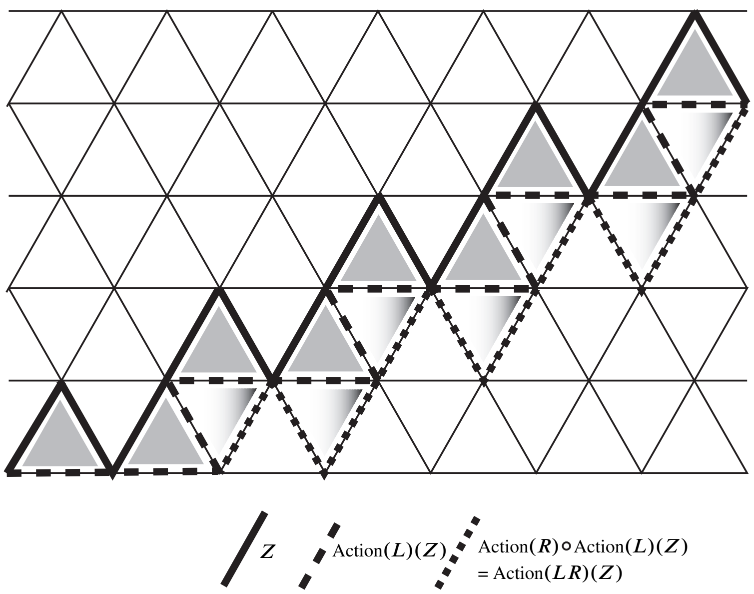

Figure 10: Figure-eight knot – zig-zag curve and its tranforms by the holonomy $\phi = L \circ R$
Note that the final image is parallel to Z. Each move is an equivariant push across a tetrahedron, where the tetrahedron is indicated by a shaded triangle.

References

[1] **W Abikoff**, *Two theorems on totally degenerate Kleinian groups*, Amer. J. Math. 98 (1976) 109–118 MR0396937

[2] **F Bonahon**, *Bouts des variétés hyperboliques de dimension* 3, Ann. of Math. (2) 124 (1986) 71–158 MR847953

[3] **J W Cannon**, *The combinatorial structure of cocompact discrete hyperbolic groups*, Geom. Dedicata 16 (1984) 123–148 MR758901

[4] **A Fathi**, **F Laudenbach**, **V Poenaru**, **et al**, *Travaux de Thurston sur les surfaces*, Séminaire Orsay, Astérisque 66–67, Société Mathématique de France, Paris (1979) MR568308 French with an English summary

[5] **R L Moore**, *Concerning upper semi-continuous collections of continua*, Trans. Amer. Math. Soc. 27 (1925) 416–428 MR1501320

[6] **D Sullivan**, *Travaux de Thurston sur les groupes quasi-fuchsiens et les variétés hyperboliques de dimension* 3 *fibrées sur* S^1, from: "Séminaire Bourbaki 554 (1980)", Lecture Notes in Math. 842, Springer, Berlin (1981) 196–214 MR636524

[7] **W P Thurston**, *Hyperbolic Structures on 3-manifolds, II: Surface groups and 3–manifolds which fiber over the circle*, unpublished

[8] **W P Thurston**, *Three-dimensional geometry and topology. Vol. 1*, Princeton Mathematical Series 35, Princeton University Press, Princeton, NJ (1997) MR1435975 Edited by Silvio Levy

279 TMCB, Brigham Young University
Provo, UT 84602, USA

Department of Mathematics, Cornell University
Ithaca, NY 14853-4201, USA

cannon@math.byu.edu, wpt@math.cornell.edu

Proposed: Dave Gabai
Seconded: Walter Neumann, Joan Birman

Received: 12 August 1999
Revised: 12 April 2007