

Local connectivity, Kleinian groups and geodesics on the blowup of the torus

Curtis T. McMullen

15 May, 2000

Abstract

Let $N = \mathbb{H}^3/\Gamma$ be a hyperbolic 3-manifold with free fundamental group $\pi_1(N) \cong \Gamma \cong \langle A, B \rangle$, such that $[A, B]$ is parabolic. We show that the limit set Λ of N is always locally connected.

More precisely, let Σ be a compact surface of genus 1 with a single boundary component, equipped with the Fuchsian action of $\pi_1(\Sigma)$ on the circle S_∞^1 . We show that for any homotopy equivalence $f : \Sigma \rightarrow N$, there is a natural continuous map

$$F : S_\infty^1 \rightarrow \Lambda \subset S_\infty^2,$$

respecting the action of $\pi_1(\Sigma)$.

In the course of the proof we determine the location of all closed geodesics in N , using a factorization of elements of $\pi_1(\Sigma)$ into simple loops.

Contents

1	Introduction	1
2	Geodesics on the blowup of the torus	6
3	The model manifold	14
4	Paths in the model	22
5	Model geodesics	25
6	Efficiency: the thick case	31
7	Efficiency: the thin case	41
8	Ambient quasigeodesics	53
9	Conclusion: Continuity	56
10	The case $\Lambda \neq S_\infty^2$	59

Research supported in part by the NSF.

2000 Mathematics Subject Classification: 30F40 (30F60, 57M50, 58F23).

1 Introduction

Let $N = \mathbb{H}^3/\Gamma$ be a complete hyperbolic 3-manifold with free fundamental group

$$\pi_1(N) \cong \Gamma \cong \langle A, B \rangle,$$

such that the commutator $[A, B]$ is parabolic. The *limit set* Λ of N is the locus of chaotic dynamics for the action of $\pi_1(N)$ on $S_\infty^2 = \partial\mathbb{H}^3$; it can be defined as

$$\Lambda = \overline{\Gamma x} \cap S_\infty^2$$

for any $x \in \mathbb{H}^3$.

In this paper we will show that the topological dynamical system $(\Lambda, \pi_1(N))$ is always a quotient of the standard action of a surface group on a circle.

Let Σ be a compact surface of genus one with a single boundary component. Its interior Σ^0 can be endowed with a complete hyperbolic metric of finite volume, providing a natural action of $\pi_1(\Sigma)$ on the circle

$$S_\infty^1 = \partial\tilde{\Sigma}^0 \cong \partial\mathbb{H}.$$

The topology of this action is independent of the choice of metric.

Since $\pi_1(\Sigma)$ is also a free group, we can find a homotopy equivalence or *marking*

$$f : \Sigma \rightarrow N,$$

sending $\partial\Sigma$ to a cusp of N (by the commutator condition). The isomorphism $f_* : \pi_1(\Sigma) \rightarrow \pi_1(N)$ provides an action of $\pi_1(\Sigma)$ on Λ . Let $H(\Sigma)$ denote the set of all such marked hyperbolic 3-manifolds.

Theorem 1.1 *For any $N \in H(\Sigma)$ there is a natural, continuous, surjective map*

$$F : S_\infty^1 \rightarrow \Lambda \subset S_\infty^2,$$

respecting the action of $\pi_1(\Sigma)$.

See Figure 1 for an example in which the limit set is the whole sphere. (This figure was produced with the help of a computer program developed by David Wright.)

Corollary 1.2 *The limit set of any $N \in H(\Sigma)$ is locally connected.*

The boundary of a group. To put these results in context, we formulate:

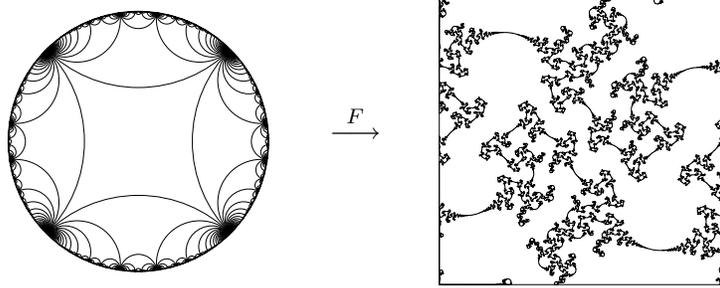


Figure 1. An approximation to the sphere-filling curve $F : S_\infty^1 \rightarrow S_\infty^2$.

Conjecture 1.3 *For any hyperbolic 3-manifold N with finitely-generated fundamental group, there exists a continuous, $\pi_1(N)$ -equivariant map*

$$F : \partial\pi_1(N) \rightarrow \Lambda \subset S_\infty^2.$$

Here $\partial\pi_1(N)$ is constructed by scaling the metric on Cayley graph of $\pi_1(N)$ by the conformal factor of $d(e, x)^{-2}$, then taking the metric completion [Fl].

Theorem 1.1 is a special case of this conjecture. Indeed, let S be a compact orientable surface with $\chi(S) < 0$. Then each boundary subgroup $\pi_1(\partial S, p) \cong \mathbb{Z} \subset \pi_1(S)$ determines a pair of points in $\partial\pi_1(S)$; identifying these pairs, we obtain the circle

$$S_\infty^1 = \partial\pi_1(S) / \sim.$$

For any $N \in H(S)$, the map F must also identify these pairs since $\pi_1(\partial S)$ is parabolic, so the conjecture predicts a natural map

$$F : S_\infty^1 \rightarrow \Lambda.$$

Floyd proved Conjecture 1.3 in the case where N is geometrically finite [Fl]. Cannon and Thurston proved Conjecture 1.3 in the case where S is closed and $N \in H(S)$ is the \mathbb{Z} -covering space of a compact hyperbolic 3-manifold that fibers over the circle [CT]. Minsky proved Conjecture 1.3 in the case where $N \in H(S)$ has injectivity radius bounded below [Min1], generalizing the Cannon–Thurston result. (See also [AM], [Kl] and [ADP].)

Arithmetic and complex dynamics. In contrast to the cases treated by Floyd, Cannon–Thurston and Minsky, Theorem 1.1 above allows N to have a sequence of closed geodesics with $L(\gamma_n) \rightarrow 0$.

It is interesting to compare the limit set $\Lambda(N)$, $N \in H(\Sigma)$, to the Julia set $J(P_\theta)$ of the quadratic map

$$P_\theta(z) = e^{2\pi i\theta} z + z^2.$$

The parameter θ is a kind of end invariant; for example, $J(P_\theta)$ is a quasicircle iff $\theta \in \mathbb{H}$. When $\theta \in \mathbb{R}$ is irrational, but very well-approximated by rational numbers, $z = 0$ belongs to the Julia set and $J(P_\theta)$ is *not* locally connected [Sul, Theorem 8].

Minsky has shown [Min2] that $N \in H(\Sigma)$ is characterized by a pair of end invariants

$$\epsilon_\pm(N) \in \mathbb{H} \cup \mathbb{R} \cup \{\infty\}.$$

When one of its end invariants is well-approximated by rationals, N has arbitrarily short geodesics. By analogy with quadratic polynomials, such manifolds provide natural candidates for badly-behaved limit sets. Theorem 1.1 shows that, despite their short geodesics, the limit sets of these manifolds remain locally connected.

For more on the local connectivity of Julia sets, see [Hub], [Ji], [Lyu], [Mil], [Pet1], [Pet2], [Sor], and [TY].

Factorization into simple curves and model geodesics. We now turn to the main ideas in the proof of Theorem 1.1. For simplicity, we present the proof in the case where $\Lambda = S_\infty^2$. The general case follows the same lines, and all essential features appear in the case where the limit set is the whole sphere. (§10 indicates how to modify the proof to handle the case $\Lambda \neq S_\infty^2$.)

1. To begin with, we regard Σ as the blowup of the flat torus $T = \mathbb{R}^2/\mathbb{Z}^2$ at the origin $p = (0, 0)$. Thus $\partial\Sigma$ is identified with the space of rays in the tangent space $T_p T$. The surface Σ carries a *canonical* affine structure on its interior and projective structure on its boundary.
2. Next, we regard S_∞^1 as the space of geodesic rays $\gamma' : [0, \infty) \rightarrow \Sigma$ with a fixed basepoint $\gamma'(0) = q \in \partial\Sigma$. By *geodesic* we mean simply that γ' is an immersion that follows straight lines in the interior of Σ , and that turns through angle at least π whenever it runs along $\partial\Sigma$.

One can think of these rays as limits of geodesics in a flat metric on $T - B(p, r)$ as $r \rightarrow 0$. They have two advantages over geodesics for a hyperbolic metric on Σ .

First, the affine geodesics are canonical — they depend only on the affine structure on Σ , not on the choice of a metric.

Second, every affine geodesic ray admits a canonical factorization

$$\gamma' = \beta'_0 * \alpha'_1 * \beta'_1 * \alpha'_2 * \beta'_2 * \dots \quad (1.1)$$

where the paths β'_i run along $\partial\Sigma$ and the paths α'_i are straight lines in the interior of Σ . In particular, each α'_i is an embedded *simple geodesic* on Σ .

3. To define the map $F : S^1_\infty \rightarrow S^2_\infty$, we will associate to every ray γ' on Σ as above a geodesic ray $\eta : [0, \infty) \rightarrow \tilde{N} \cong \mathbb{H}^3$, and then set

$$F(\gamma') = \lim_{s \rightarrow \infty} \eta(s) \in \partial\tilde{N} = S^2_\infty.$$

Informally, η will simply be the geodesic representative of γ' under the homotopy equivalence $f : \Sigma \rightarrow N$.

4. To understand the geodesics in N , we use a version of Minsky's *model manifold* M [Min2]. The model is a Riemannian manifold equipped with a quasi-isometry to the thick part,

$$\phi : M \rightarrow N_{\text{thick}} \subset N,$$

with the virtue that M can be explicitly constructed from the Teichmüller geodesic $\tau \subset \text{Teich}(T)$ connecting the ending laminations (ϵ_-, ϵ_+) of N .

Here is a sketch of the construction of M . We regard $\text{Teich}(T)$ as the space of flat metrics g on T of total area one.

Let $\mathcal{S}(T) \cong \mathbb{Q} \cup \{\infty\}$ be the set of slopes of simple loops on T . Let $\mathcal{S}(\tau)$ denote the *short geodesics* for τ , i.e. those $\mu \in \mathcal{S}(T)$ such that μ is very short on (T, g) for all g in a segment $\tau(\mu) \subset \tau$. Let $A(\mu) \subset T$ be a suitable annulus in the homotopy class of μ . Then the model manifold is given by

$$M = T \times \tau - \left(P \cup \bigcup_{\mu} A(\mu) \times \tau(\mu) \right),$$

where P is an infinite cylinder modeling the rank-one cusp of M , and $A(\mu) \times \tau(\mu)$ is a solid torus modeling the Margulis tube about the short geodesic representing μ . We think of the map $M \subset T \times \tau \rightarrow \tau$ as a vertical coordinate or height function on M .

To construct the metric on M , we start with the natural metric on $T \times \tau$ that restricts to g on the subsurface $T \times \{g\}$. We then modify

the metric on the blocks $M(\mu) = M \cap (T \times \tau(\mu))$, $\mu \in \mathcal{S}(\tau)$, making these submanifolds into product regions of bounded height.

5. Given a *simple* loop α' on Σ , it is easy to construct an efficient model α for its geodesic representative in M . Indeed, if $g \in \tau$ is the metric that minimizes the length of α' on (T, g) , then we may take $\alpha \subset M(\mu)$ if $g \in \tau(\mu)$, and $\alpha \subset T \times \{g\}$ otherwise.
6. Now let γ' be a geodesic ray on Σ , factored as in (1.1) above. Letting each *simple* segment α'_i in γ' migrate to its optimal location α_i in M as above, we obtain a *model geodesic ray* $\gamma : [0, \infty) \rightarrow M$ of the form

$$\gamma = \beta_0 * \xi_0 * \alpha_1 * \beta_1 * \xi_1 * \alpha_2 * \beta_2 * \xi_2 \cdots$$

Here the paths β_i are horizontal lifts of the paths β'_i , and the paths ξ_i run vertically to connect pieces at different heights.

7. It is convenient to lift these constructions to the universal cover, and regard a model geodesic as a ray

$$\gamma : [0, \infty) \rightarrow \widetilde{M} \xrightarrow{\widetilde{\phi}} \widetilde{N}_{\text{thick}} = \mathbb{H}^3 - \bigcup \mathcal{H},$$

where \mathcal{H} is a collection of disjoint cylinders and horoballs covering the thin part of M . The main technical result of this paper (Theorem 6.1) then states that γ is an *ambient K -quasigeodesic*. This means that for any segment $\alpha = \gamma|[a, b]$ and any path $A : [a, b] \rightarrow \widetilde{M}$ with the same endpoints, we have $L(A) \leq KL(\alpha) + K$.

Since $\widetilde{\phi}$ is a quasi-isometry, the ray $\widetilde{\phi} \circ \gamma$ is also an ambient K' -quasigeodesic in $\mathbb{H}^3 - \bigcup \mathcal{H}$.

8. Now it is a general fact (§8) that an ambient quasigeodesic ray in $\mathbb{H}^3 - \bigcup \mathcal{H}$ lies in a bounded neighborhood of a hyperbolic geodesic η , modulo \mathcal{H} . We can now finally define $F(\gamma')$ to be the endpoint of η in S_∞^2 . Then continuity of F comes from continuity of the factorization (1.1) of geodesic rays on Σ , completing the proof (§9).

Closed geodesics. The construction of model geodesics can be applied to loops as well as rays, yielding a combinatorial picture of all of the closed geodesics in M .

Higher genus. Many features of the argument above also apply to $N \in H(S)$, where S is a surface of higher genus. For example, simple geodesics are *localized* in N ; one knows:

- All short geodesics $\gamma \subset N \in H(S)$ correspond to simple loops on S ; and
- If a geodesic $\gamma \subset N$ represents a simple loop on S , then γ has bounded diameter modulo the thin part of N .

The key to proving these results is that one can construct a pleated surface through γ whenever γ represents a *simple* loop on S ; see [Th, §8, §9], [Bon]. Using train tracks, one can even show that γ varies continuously on the space of laminations on S .

The factorization of elements of $\pi_1(S)$ into simple curves can also be carried out on surfaces of higher genus. Indeed, if one picks a complex structure on S together with a holomorphic quadratic differential $\phi(z) dz^2$, then every geodesic loop in the metric $|\phi|$ factors into a product of simple curves connecting the zeros of ϕ and/or the punctures of S . This factorization depends only on the real affine structure determined by ϕ .

The main missing ingredient for higher genus is a clear picture of the quasi-isometry type of N . At present such a picture is only available when N is geometrically finite or when the injectivity radius of N is bounded below [Min1].

Acknowledgements. I would like to thank Y. Minsky, H. Miyachi and the referee for many helpful corrections and suggestions regarding the first version of this paper.

Notation. As usual, $A = O(B)$ and $A \asymp B$ mean $A < CB$ and $B/C < A < CB$, for an implicit constant C . We write $\cosh^\alpha x$ for $(\cosh(x))^\alpha$.

2 Geodesics on the blowup of the torus

We begin with a summary of the affine and Euclidean geometry of the torus T and its blowup Σ .

Paths. Let X be a topological space. A *path* $\alpha : I \rightarrow X$ is a continuous map from an interval $I \subset \mathbb{R}$ into X . When (X, g) is a Riemannian manifold, we denote the length of α by $L(\alpha, g)$ or simply $L(\alpha)$.

If $h : I' \rightarrow I$ is an order-preserving homeomorphism, we say $\alpha' = \alpha \circ h$ is obtained from α by *change of parameter*. We generally *identify* paths that differ only by change of parameter.

If $\alpha : [a, b] \rightarrow X$ and $\beta : [b, c] \rightarrow X$ are paths with $\alpha(b) = \beta(b)$, then we can combine them to obtain a new path

$$\alpha * \beta : [a, c] \rightarrow X,$$

the *product* of α and β . A similar construction applies whenever α and β have matching endpoints, so they can be put in the form above by a change of parameter. The notation α^n signifies the n -fold product $\alpha * \cdots * \alpha$ for $n > 0$, and the same with the orientation reversed for $n < 0$.

A path $\alpha : I \rightarrow X$ is a *subpath* of $\beta : I' \rightarrow X$ if $\alpha = \beta \circ h$ for some embedding $h : I \rightarrow I'$.

Loops. A *loop* is a continuous map $\sigma : S^1 \rightarrow X$, where $S^1 = \mathbb{R}/\mathbb{Z}$. As for paths, we generally identify loops that differ by an orientation-preserving change of parameter.

It is sometimes useful to replace a loop by a path $\sigma : [a, b] \rightarrow X$ with $\sigma(a) = \sigma(b)$. For example, when we write $\alpha * \sigma$, we implicitly break σ at one endpoint of α , and then take the product of the resulting paths.

Bilipschitz maps. A bijection $f : X \rightarrow Y$ between metric spaces is *bilipschitz* if there is a constant K such that

$$\frac{1}{K} d(x, x') \leq d(f(x), f(x')) \leq K d(x, x') \quad (2.1)$$

for all $x, x' \in X$. The smallest $K \geq 1$ satisfying (2.1) is the *bilipschitz constant* of f .

Teichmüller space of the torus. Let $T = \mathbb{R}^2/\mathbb{Z}^2$ be the standard closed torus with its Euclidean metric. Let $\text{Teich}(T)$ denote the space of positive-definite quadratic forms g on \mathbb{R}^2 with the same area element as the Euclidean metric. Each $g \in \text{Teich}(T)$ makes the torus into a flat Riemannian manifold (T, g) of area one, so we can regard $\text{Teich}(T)$ as the *Teichmüller space* of T .

There are natural identifications

$$\text{Teich}(T) = \text{SL}_2(\mathbb{R})/\text{SO}_2(\mathbb{R}) = \mathbb{H} \subset \mathbb{C},$$

with $i \in \mathbb{H}$ corresponding to the Euclidean metric on \mathbb{R}^2 . We equip $\text{Teich}(T)$ with the hyperbolic metric of constant curvature -1 , given by $\rho = |dz|/\text{Im } z$ on \mathbb{H} . For any $g, h \in \text{Teich}(T)$, the bilipschitz constant $K(g, h)$ of the identity map $(T, g) \rightarrow (T, h)$ satisfies

$$\log K(g, h) = \frac{1}{2} d(g, h). \quad (2.2)$$

The mapping-class group $\text{Mod}(T) = \text{SL}_2(\mathbb{Z})$ acts on $\text{Teich}(T)$ by hyperbolic isometries.

The boundary of Teichmüller space. Teichmüller space can be naturally compactified by adjoining to it the space

$$\partial \text{Teich}(T) = \mathbb{P}(\mathbb{R}^2)$$

of lines in \mathbb{R}^2 . A sequence of metrics $g_n \in \text{Teich}(T)$ converges to a line $L \in \mathbb{P}(\mathbb{R}^2)$ iff the unit ball $B_n \subset \mathbb{R}^2$ of g_n is an ellipse with eccentricity tending to infinity, and the major axis of B_n tends to L . In terms of the upper halfplane model, we have

$$\text{Teich}(T) \cup \mathbb{P}(\mathbb{R}^2) \cong \mathbb{H} \cup \widehat{\mathbb{R}},$$

where $\widehat{\mathbb{R}} = \mathbb{R} \cup \{\infty\}$, and where $[x : y] \in \mathbb{P}(\mathbb{R}^2)$ corresponds to $x/y \in \mathbb{R}$.

Simple curves and slopes. Let $\mathcal{S}(T)$ denote the set of isotopy classes of unoriented, essential, simple closed loops on T .

Each element of $\mathcal{S}(T)$ determines a homology class

$$\pm(a, b) \in H_1(T, \mathbb{Z}) \cong \mathbb{Z}^2$$

up to sign, with a and b relatively prime. By recording the slope $[a : b]$ of this homology class, we obtain a natural identification between $\mathcal{S}(T)$ and the set of rational points on the boundary of Teichmüller space:

$$\mathcal{S}(T) = \mathbb{P}(\mathbb{Q}^2) \subset \mathbb{P}(\mathbb{R}^2) = \partial \text{Teich}(T).$$

This identification respects the action of $\text{Mod}(T) = \text{SL}_2(\mathbb{Z})$. We refer to $[a : b]$ as the *slope* of $\sigma \in \mathcal{S}(T)$.

Length functions, loops and intersections. It is useful to represent each slope $[a : b] \in \mathcal{S}(T)$ by a specific *geodesic loop* $\sigma \subset T$. Choose integers (c, d) with $ad - bc = 1$; then we define $\sigma : \mathbb{R}/\mathbb{Z} \rightarrow T$ by

$$\sigma(t) = (at, bt) + \frac{1}{2}(c, d). \quad (2.3)$$

For any $g \in \text{Teich}(S)$, the loop σ is the unique geodesic on T with slope $[a : b]$ that is as far as possible from the basepoint $p = (0, 0)$.

We denote both elements of $\mathcal{S}(T)$ and the corresponding geodesic loops by σ . The *geodesic length function*

$$L_\sigma : \text{Teich}(T) \rightarrow \mathbb{R}$$

is defined by $L_\sigma(g) = L(\sigma, g)$.

The *intersection number* between a pair of loops $\sigma, \rho \in \mathcal{S}(T)$ with slopes $[a : b]$ and $[c : d]$ is given by

$$i(\sigma, \rho) = |ad - bc|;$$

it is the absolute value of the homological intersection number.

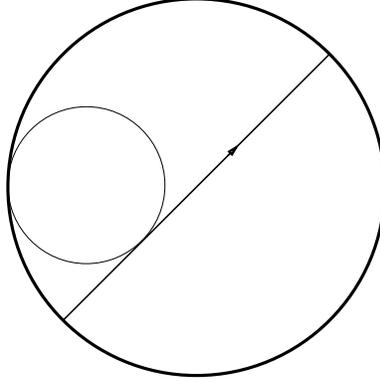


Figure 2. Optimal metrics.

Optimal metrics. Let $\tau \subset \text{Teich}(T) \cong \mathbb{H}$ be a hyperbolic geodesic joining a pair of irrational points (ϵ_-, ϵ_+) in $\widehat{\mathbb{R}} = \partial \text{Teich}(T)$. For each loop $\sigma \in \mathcal{S}(T)$, the length function $L_\sigma(g)$ achieves its minimum along τ at a unique *optimal metric* $g_\sigma \in \tau$. To find the metric g_σ , one expands a horoball centered at σ until it first touches τ (see Figure 2).

The length of σ is exponentially longer in any non-optimal metric.

Lemma 2.1 *For any $g \in \tau$ and $\sigma \in \mathcal{S}(\tau)$, we have*

$$\frac{L_\sigma(g)}{L_\sigma(g_\sigma)} = \cosh^{1/2} d(g, g_\sigma).$$

Proof. It is sufficient to treat the case where σ corresponds to $\infty \in \widehat{\mathbb{R}}$. For each $t \in \mathbb{H}$, the conformal metric of area one on $\mathbb{C}/(\mathbb{Z} \oplus t\mathbb{Z})$ has the form $\rho_t = |dz|/\sqrt{y}$, $y = \text{Im } t$, and thus:

$$L_\sigma(t) = y^{-1/2}.$$

Making the change of coordinates $t \mapsto at + b$, we can assume τ is the geodesic from -1 to 1 , parameterized by arclength as

$$\tau(s) = \frac{1 + ie^{-s}}{1 - ie^{-s}} = \tanh(s) + i \text{sech}(s).$$

Thus $L_\sigma(\tau(s))^2 = \cosh(s)$, which attains its minimum at $s = 0$, and the Lemma follows. \blacksquare

Slopes and optimal metrics. Fix the metric g_σ on \mathbb{R}^2 . Then the lines $L(\epsilon_-)$ and $L(\epsilon_+)$ are orthogonal, so we can take them as coordinate axes adapted to g_σ . The line $L(\sigma)$ makes an angle of 45° with respect to $L(\epsilon_\pm)$, and this property characterizes g_σ . The unit ball of any other metric $g \in \tau$ is an ellipse with major and minor axes along $L(\epsilon_\pm)$.

If we take a sequence of loops $\sigma_i \in \mathcal{S}(T)$ whose optimal metrics g_{σ_i} are evenly spaced along τ , then we obtain a sequence of lines $L(\sigma_i)$ whose slopes form a geometric sequence in the (ϵ_-, ϵ_+) -coordinate system. See Figure 3. In particular, whenever g_ρ is far from g_σ , the line $L(\rho)$ is close to $L(\epsilon_-)$ or $L(\epsilon_+)$.

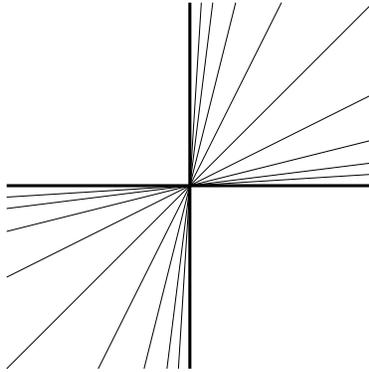


Figure 3. Equally spaced slopes.

Blowing up. Let Σ be the oriented blowup of T at the origin. That is, let Σ be the smooth surface obtained by replacing $p = (0, 0) \in T$ with the circle of rays in the tangent space $T_p T$, forming $\partial\Sigma$. There is a natural map

$$b : \Sigma \rightarrow T,$$

with $b(\partial\Sigma) = p$ and $b|_{\text{int } \Sigma}$ injective.

Topologically, Σ is a compact surface of genus one with one boundary component. The affine structure of \mathbb{R}^2 determines an affine structure on $\text{int}(\Sigma)$ and a real projective structure on $\partial\Sigma$. (The latter structure comes about by regarding $\partial\Sigma$ as a double cover of $\mathbb{P}(T_p T) \cong \mathbb{RP}^1$.)

Any metric $g \in \text{Teich}(T)$ makes the interior and boundary of Σ into Riemannian manifolds $(\text{int}(\Sigma), g)$ and $(\partial\Sigma, g)$. The metric on $\text{int}(\Sigma)$ is pulled back from (T, g) . The metric on $\partial\Sigma$ comes about by identifying $\partial\Sigma$ with $U_p(T, g)$, the unit tangent vectors to $p \in (T, g)$. We have $L(\partial\Sigma, g) = 2\pi$.

them to the corresponding (maximal abelian) covering space $\tilde{\Sigma}^{ab} \rightarrow \Sigma$. One can think of $\tilde{\Sigma}^{ab}$ as \mathbb{R}^2 with the integral lattice \mathbb{Z}^2 blown up; to obtain a reasonable picture, we will draw $\partial\tilde{\Sigma}^{ab}$ as a collection of small circles centered at the integral points.

Figure 4 shows a typical geodesic segment drawn with these conventions.

Factorization. A *simple geodesic segment* $\alpha : [a, b] \rightarrow \Sigma$ is a geodesic with endpoints in $\partial\Sigma$ and interior embedded in the interior of Σ . Simple segments on Σ correspond bijectively (up to orientation) to slopes $\sigma \in \mathcal{S}(T)$.

A *boundary geodesic segment* $\beta : [a, b] \rightarrow \Sigma$ is a geodesic with $\beta(I) \subset \partial\Sigma$.

Every geodesic $\gamma : [a, b] \rightarrow \Sigma$ with endpoints on Σ can be canonically factored as a product

$$\gamma = \beta_0 * \alpha_1 * \beta_1 * \alpha_2 * \beta_2 * \cdots * \alpha_n * \beta_n$$

of simple segments and boundary segments. (We allow β_0 and β_n to be trivial).

Geodesic rays. Let $\tilde{\Sigma} \rightarrow \Sigma$ denote the universal covering space. A *geodesic ray* $\gamma : [0, \infty) \rightarrow \Sigma$ is a geodesic that tends to infinity when lifted to $\tilde{\Sigma}$. We say γ is a *simple ray* (or *boundary ray*) if we have $\gamma(s) \in \text{int}(\Sigma)$ (or $\gamma(s) \subset \partial\Sigma$) for all $s > 0$.

Given a basepoint $q \in \Sigma$, let $\mathcal{G}(\Sigma, q)$ denote the set of all geodesic rays with $\gamma(0) = q$. The rays based at q admit a natural cyclic ordering and topology, making $\mathcal{G}(\Sigma, q)$ into a Cantor set. This Cantor set can be naturally identified with the space of ends of the Cayley graph of $\pi_1(\Sigma, q) \cong \mathbb{Z} * \mathbb{Z}$.

Almost all rays $\gamma \in \mathcal{G}(\Sigma, q)$ can be canonically factored as infinite products

$$\gamma = \beta_0 * \alpha_1 * \beta_1 * \alpha_2 * \beta_2 * \cdots$$

of simple segments and boundary segments. The remaining rays have the form

$$\gamma = \begin{cases} \beta_0 * \alpha_1 * \beta_1 * \alpha_2 * \beta_2 * \cdots * \alpha_n & \text{or} \\ \beta_0 * \alpha_1 * \beta_1 * \alpha_2 * \beta_2 * \cdots * \beta_n, \end{cases}$$

where α_n (or β_n) is a simple ray (or boundary ray). Among these, the complementary intervals of the cyclically ordered Cantor set $\mathcal{G}(\Sigma, q)$ correspond to pairs of rays of the form

$$\gamma(i) = \beta_0(i) * \alpha_1(i) * \beta_1(i) * \cdots * \alpha_n(i) * \beta_n(i),$$

$i = 1, 2$, where $\beta_n(1)$ and $\beta_n(2)$ spiral around $\partial\Sigma$ in opposite directions. In this case we say $\gamma(1)$ and $\gamma(2)$ are *boundary equivalent* and write $\gamma(1) \sim \gamma(2)$.

By identifying these pairs, we obtain a natural circle at infinity:

$$\mathcal{G}(\Sigma, q) / \sim = S_\infty^1(q).$$

A path from q_1 to q_2 determines a natural homeomorphism $S_\infty^1(q_1) \rightarrow S_\infty^1(q_2)$. In particular, there is a natural action of $\pi_1(\Sigma, q)$ on $S_\infty^1(q)$.

Pick a hyperbolic metric of finite volume on $\text{int } \Sigma$; then we can write $\text{int}(\Sigma) = \mathbb{H} / \pi_1(\Sigma, q)$. By straightening rays to hyperbolic geodesics, we obtain a natural identification between $S_\infty^1(q)$ and $\partial\mathbb{H}$, compatible with the action of $\pi_1(\Sigma, q)$.

Bidistance. The notions of a geodesic generalizes in a natural way from Σ to its universal cover $\tilde{\Sigma}$. All geodesics in $\tilde{\Sigma}$ are *embedded*, so it is natural and convenient to treat geodesics as *closed arcs* $\gamma \subset \tilde{\Sigma}$ instead of paths.

Any two points $p, q \in \tilde{\Sigma}$ are joined by a *unique* geodesic γ . Let $\gamma = \beta_0 * \alpha_1 * \beta_1 * \cdots * \alpha_n * \beta_n$ be the canonical factorization of γ into simple segments and boundary segments. We define the *bidistance* $d_g(p, q) \in \mathbb{R}^2$ by

$$d_g(p, q) = (d_g^1(p, q), d_g^2(p, q)) = \left(\sum L(\alpha_i, g), \sum L(\beta_i, g) \right).$$

The two terms measure the length of γ in $\text{int } \tilde{\Sigma}$ and $\partial\tilde{\Sigma}$ respectively.

Bidistances will be compared using the lexicographically partially order on \mathbb{R}^2 , defined by

$$(x, y) < (x', y') \iff (x < x') \text{ or } (x = x' \text{ and } y < y').$$

Projection. Fix a geodesic $\gamma \subset \tilde{\Sigma}$. For each metric $g \in \text{Teich}(\tilde{\Sigma})$, the *nearest-point projection*

$$P_g : \tilde{\Sigma} \rightarrow \gamma$$

maps $x \in \tilde{\Sigma}$ to the point $y \in \gamma$ that minimizes $d_g(x, y)$ in terms of the lexicographical ordering. Since γ is convex, it is not hard to show that this nearest point y exists and is unique. Moreover:

- *The projection P_g contracts the first component of the bidistance.*

This means $d_g^1(P_g(x), P_g(y)) \leq d_g^1(x, y)$.

Despite this contraction, P_g is not quite continuous; its value can jump when $P_g(x)$ lies in $\gamma \cap \partial\tilde{\Sigma}$. See Figure 5 for an example (drawn on $\tilde{\Sigma}^{ab}$). Using the fact that the geodesic from x to $P_g(x)$ is always normal to γ , one can show that the only discontinuities are like those in the example:

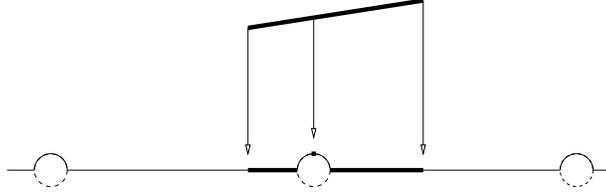


Figure 5. Discontinuity of projection.

- If $x_n \rightarrow x$ and $P_g(x_n) \rightarrow y$ then either $P_g(x) = y$ or

$$d_g(y, P_g(x)) = (0, \pi/2).$$

In particular, $P_g(x) = y$ unless both lie in $\partial\Sigma$.

Almost continuity. For any $x \in \tilde{\Sigma}$ let

$$B(x) = \{y : d_g(x, y) < (0, \pi)\}.$$

For $x \in \partial\tilde{\Sigma}$, $\pi(B(x)) \subset \partial\Sigma$ is the open interval excluding the antipode of x . For $x \notin \partial\tilde{\Sigma}$, $B(x) = \{x\}$. Thus $B(x)$ is independent of the metric g .

Let us say a map $f : X \rightarrow \tilde{\Sigma}$ is *almost continuous* if for any $x \in X$ and any neighborhood U of $B(x)$, the set $f^{-1}(U)$ is open in X . Then the projection $P_g : \tilde{\Sigma} \rightarrow \gamma \subset \tilde{\Sigma}$ is almost continuous.

3 The model manifold

Let $\tau \subset \mathbb{H}$ be a geodesic with irrational endpoints (ϵ_-, ϵ_+) in \mathbb{R} .

In this section we will construct from τ a Riemannian 3-manifold (M^τ, h) , together with a continuous map

$$\pi : M^\tau \rightarrow \Sigma.$$

We will then formulate Minsky's result, stating that if $N \in H(\Sigma)$ is a hyperbolic 3-manifold with ending laminations (ϵ_-, ϵ_+) , then the thick part of N is quasi-isometric to M^τ .

3.1 Teichmüller geodesics

Fix a small number r , $0 < r \ll 1$. The construction we are about to describe works for all r sufficiently small; a concrete value for r can be given in terms of the Margulis constant for hyperbolic 3-manifolds.

Thick/thin. For $g \in \text{Teich}(T)$, let

$$L_{\min}(g) = \inf\{L_{\sigma}(g) : \sigma \in \mathcal{S}(T)\}$$

be the length of the shortest closed geodesic on (T, g) . Clearly $L_{\min}(g)$ is $\text{SL}_2(\mathbb{Z})$ -invariant, and it descends to a proper function $L_{\min} : \mathbb{H}/\text{SL}_2(\mathbb{Z}) \rightarrow (0, \infty)$ on the moduli space of T .

Given $s > 0$, we define the thin and thick parts of Teichmüller space by

$$\begin{aligned} \text{Teich}(T)_{\geq s} &= \{g : L_{\min}(g) \geq s\}, \\ \text{Teich}(T)_{< s} &= \{g : L_{\min}(g) < s\}. \end{aligned}$$

For all s sufficiently small, the thick part of Teichmüller space is connected, while the thin part is the countable unit of disjoint horoballs of the form

$$H_{\mu}(s) = \{g : L_{\mu}(g) < s\}$$

with $\mu \in \mathcal{S}(T)$. For $g \in H_{\mu}(s)$, all loops on (T, g) other than μ are long; indeed, we have $L_{\sigma}(g) \geq 1/s$ for any $\sigma \neq \mu$, since the total area of (T, g) is one.

Short geodesics. Let $\tau \subset \mathbb{H}$ be a hyperbolic geodesic with irrational endpoints. Let

$$\mathcal{S}(\tau) = \{\mu \in \mathcal{S}(T) : L_{\mu}(g) \leq r^2 \text{ for some } g \in \tau\}$$

be the set of *short geodesics* for τ . For each $\mu \in \mathcal{S}(\tau)$, let

$$\tau(\mu) = \tau \cap H_{\mu}(r).$$

Then $\{\tau(\mu) : \mu \in \mathcal{S}(\tau)\}$ is a collection of disjoint open intervals in τ , each with hyperbolic length at least $|\log r| \gg 1$ (by (2.2)). The short geodesics determine a decomposition

$$\tau = \tau(\text{thin}) \cup \tau(\text{thick}),$$

where

$$\tau(\text{thin}) = \bigcup \{\tau(\mu) : \mu \in \mathcal{S}(\tau)\}$$

and $\tau(\text{thick}) = \tau - \tau(\text{thin})$. We have

$$\begin{aligned} L_{\min}(g) &< r \quad \text{for } g \in \tau(\text{thin}), \text{ and} \\ L_{\min}(g) &\geq r^2 \quad \text{for } g \in \tau(\text{thick}). \end{aligned}$$

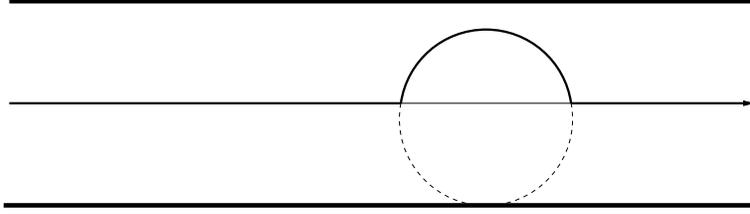


Figure 6. A geodesic τ passing through the thin part of Teichmüller space.

Next we modify τ along $\tau(\text{thin})$ to obtain a new path

$$\eta \subset \text{Teich}(T)_{\geq r^2}.$$

For $\mu \in \mathcal{S}(\tau)$, consider the interval of metrics $\tau(\mu) = (g_-, g_+) \subset \tau$. We have $L_{g_{\pm}}(\mu) = r$ at each endpoint of $\tau(\mu)$. Let $\eta(\mu) \subset \partial H_{\mu}(r)$ be the open arc with the same endpoints as $\tau(\mu)$, but running along the horocycle where the length of μ is constant (see Figure 6). We define

$$\eta = \tau(\text{thick}) \cup \bigcup_{\mu \in \mathcal{S}(T)} \eta(\mu)$$

by replacing each $\tau(\mu)$ in τ with $\eta(\mu)$.

Let $g : \mathbb{R} \rightarrow \eta$ be a piecewise smooth parameterization of η , with hyperbolic speed given by

$$|g'(t)| = \begin{cases} 1 & \text{along } \tau(\text{thick}), \\ |\eta(\mu)| & \text{along } \eta(\mu), \end{cases}$$

where $|\eta(\mu)|$ is the length of $\eta(\mu)$ in the Teichmüller metric. In other words, $g(t)$ traverses each horocycle component of η in unit time, and otherwise travels at unit speed. For convenience, we choose our time parameter such that $g(0)$ lies in the interior of $\tau(\text{thick})$.

Dehn twists. Let $\text{tw}_{\mu} \in \text{Mod}(T)$ denote a right Dehn twist about μ ; it stabilizes $\partial H_{\mu}(r)$. Then we have

$$g_+ \approx \text{tw}_{\mu}^n(g_-)$$

for some n , in the sense that the Teichmüller distance between the two sides is $O(1)$. To make n unique, choose $|n|$ minimal such that $\text{tw}_{\mu}^n(g_-)$ lies past g_+ on $\partial H_{\mu}(r)$.

We call $D(\mu) = n$ the *Dehn twisting number* of μ ; it records the number of times $\eta(\mu)$ wraps around the cusp, when projected to the moduli space $\text{Teich}(T)/\text{Mod}(T) = \mathbb{H}/\text{SL}_2(\mathbb{Z})$.

3.2 The model manifold

We can now construct the model manifold (M, h) . The manifold M is built from $T \times \mathbb{R}$ by introducing a metric h_0 determined by $g(t)$, then cutting away the model thin part, and finally modifying the metric near the short geodesics.

A metric on $T \times \mathbb{R}$. Let $T_t = T \times \{t\}$ and $R_q = \{q\} \times \mathbb{R}$. Let h_0 be the unique Riemannian metric on $T \times \mathbb{R}$ such that

- (a) $h_0|_{T_t} = g(t)$ for all $t \in \mathbb{R}$,
- (b) $h_0|_{R_q} = |dt|^2$ for all $q \in T$, and
- (c) T_t and R_q are orthogonal for all t, q .

The model thin parts. To construct M , we will cut away the parts of $T \times \mathbb{R}$ that model the rank one cusp and tubes about short geodesics. Let

$$\begin{aligned} T(\mu) &= \{t \in \mathbb{R} : g(t) \in \eta(\mu)\}, \\ T(\text{thin}) &= \bigcup_{\mu \in \mathcal{S}(\tau)} T(\mu) \text{ and} \\ T(\text{thick}) &= \{t \in \mathbb{R} : g(t) \in \tau(\text{thick})\} = \mathbb{R} - T(\text{thin}). \end{aligned}$$

Let $B_t = B(p_t, r^2/10)$ be a small ball about the basepoint

$$p_t = ((0, 0), t) \in T_t$$

in the metric $h_0|_{T_t} = g(t)$. Since $L_{\min}(g(t)) \geq r^2$ for all t , the ball B_t is embedded. Let

$$P = \bigcup_t B_t \subset T \times \mathbb{R}$$

be the *model cusp* for M .

For $\mu \in \mathcal{S}(\tau)$ and $t \in T(\mu)$, let $\mu_t \subset T_t$ be the geodesic representative of μ that is as far away from p_t as possible (defined as in equation (2.3)). Then $T_t - \mu_t$ is an annulus with p_t midway between its boundary components. Let $A_t = B(\mu_t, r^2/10) \subset T_t$ be an open annular neighborhood of μ_t . We define the *model Margulis tube* for μ by

$$C(\mu) = \bigcup \{A_t : t \in T(\mu)\}$$

The tube $C(\mu) \subset T$ is an open solid torus, disjoint from P .

The *model manifold* $M = M^\tau$ is defined by

$$M = (T \times \mathbb{R} - P) - \bigcup_{\mu} C(\mu). \quad (3.1)$$

The height function. The space M can be thought of as a stack of surfaces

$$\Sigma_t = T_t \cap M,$$

each of which is either a flat torus with a ball removed (for $t \in T(\text{thick})$), or an annulus with a ball removed (for $t \in T(\mu)$). The *height function*

$$H : M \rightarrow \mathbb{R}$$

is defined by $H(x, t) = t$, where $(x, t) \in M \subset T \times \mathbb{R}$.

In terms of the height function, define

$$\begin{aligned} M(\text{thick}) &= H^{-1}(T(\text{thick})), \\ M(\text{thin}) &= H^{-1}(T(\text{thin})), \text{ and} \\ M(\mu) &= H^{-1}(T(\mu)), \quad \mu \in \mathcal{S}(\tau). \end{aligned}$$

Then $M(\text{thin}) = \bigcup_{\mu} M(\mu)$.

The model metric. We now introduce a Riemannian metric h on M .

First, we define $h = h_0$ on $M(\text{thick}) \subset M \subset T \times \mathbb{R}$. Next, consider $\mu \in \mathcal{S}(\tau)$. To define h on $M(\mu)$, note that the surfaces $\Sigma_t \subset (M(\mu), h_0)$ are all isometric. Indeed,

$$\Sigma_t = (T_t - B(\mu_t, r^2/10)) - B(p_t, r^2/10)$$

is an annulus with a ball removed, and its shape depends only the length of μ on T_t , which is constant along $\eta(\mu)$.

Pick one surface $\Sigma_u \subset M(\mu)$ for reference, and give $\Sigma_u \times T(\mu)$ the product metric. Then there is a unique diffeomorphism

$$M(\mu) \rightarrow \Sigma \times T(\mu) \quad (3.2)$$

that sends Σ_t isometrically to $\Sigma_u \times \{t\}$ and restricts to the identity on Σ_u .

Define the metric h on $M(\mu)$ by the requirement that (3.2) is an *isometry*.

Recall that the vertical lines $R_q = \{q\} \times \mathbb{R} \subset T \times \mathbb{R}$ are normal to the surfaces $\Sigma_t \subset M(\text{thick})$. In $(M(\mu), h)$ this normality breaks down,

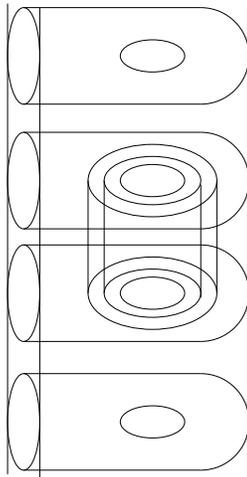


Figure 7. The model manifold M .

because the product structure on $M(\mu)$ coming from (3.2) is different from the product structure coming from $T \times \mathbb{R}$. Instead, the normals in $M(\mu)$ have the form

$$\gamma_q(t) = (f_t(q), t) \in T \times \mathbb{R},$$

where $f_t : T \rightarrow T$ is part of a 1-parameter group of shears (or continuous Dehn twists) along the geodesics parallel to μ .

The projection π . Finally we define the projection $\pi : M \rightarrow \Sigma$ by giving its restriction

$$\pi_t : \Sigma_t \rightarrow \Sigma$$

for each $t \in \mathbb{R}$.

Let $b : \Sigma \rightarrow T$ be the blowup map. We have a natural isometric inclusion

$$i : (\Sigma_t, h) \hookrightarrow (\text{int}(\Sigma), g(t))$$

via the composition $M \subset T \times \mathbb{R} \rightarrow T \xrightarrow{b^{-1}} \Sigma$. Let (ρ, θ) denote polar coordinates near the boundary of $(\Sigma, g(t))$, so that $(0, \theta)$ parameterizes $\partial\Sigma$. Let $A[r_1, r_2] \subset \Sigma$ denote the annulus with $\rho \in [r_1, r_2]$.

Since Σ_t omits $B(p_t, r^2/10)$, its image $i(\Sigma_t)$ omits $A[0, r^2/10]$. To make the image fit more snugly into Σ , define

$$j : A[r^2/10, r^2/5] \rightarrow A[0, r^2/5]$$

by expanding each radius by a factor of two; that is, by

$$j(\rho, \theta) = (2\rho - r^2/5, \theta). \quad (3.3)$$

We then define

$$\pi_t(x) = \begin{cases} j \circ i(x) & \text{if } \rho(i(x)) \leq r^2/5, \\ i(x) & \text{otherwise.} \end{cases}$$

For $t \in T(\text{thick})$, the map $\pi_t : \Sigma_t \rightarrow \Sigma$ is a piecewise smooth homeomorphism, while for $t \in I(\mu)$, $\pi_t(\Sigma_t)$ omits an annulus in Σ of slope μ .

Strictly speaking, only the homotopy class of the projection π is required for the definition of the model. The specific form we have chosen for π , however, will be useful for the discussion that follows. Combining π with the height function, we obtain an embedding

$$(\pi, H) : M \rightarrow \sigma \times \mathbb{R}.$$

3.3 Hyperbolic 3-manifolds

We now turn to the hyperbolic manifolds N modeled by M^τ .

Let $H(\Sigma)$ denote the space of oriented hyperbolic 3-manifolds marked by Σ . A point in $H(\Sigma)$ is specified by the data

$$f : \Sigma \rightarrow N,$$

where $N = \mathbb{H}^3/\Gamma$ is a hyperbolic manifold and f is a homotopy equivalence or *marking*, such that $f(\partial\Sigma)$ is homotopic to a cusp of N . Two marked manifolds $f_i : \Sigma \rightarrow N_i$, $i = 1, 2$, represent the same point in $H(\Sigma)$ if there is an isometry $\alpha : N_1 \rightarrow N_2$ such that the diagram

$$\begin{array}{ccc} \Sigma & \xrightarrow{f_1} & N_1 \\ \parallel & & \downarrow \alpha \\ \Sigma & \xrightarrow{f_2} & N_2 \end{array}$$

commutes up to homotopy.

A point in $H(\Sigma)$ is specified equivalently by a discrete group $\Gamma \subset \text{Isom}^+(\mathbb{H}^3)$ together with an isomorphism $f_* : \pi_1(\Sigma) \rightarrow \Gamma$ such that $f_*(\pi_1(\partial\Sigma))$ is parabolic. Conjugate representations represent the same point.

Thick parts of N^3 . Let $N(\epsilon) \subset N$ denote the points at which the hyperbolic metric on N has injectivity radius $\geq \epsilon$.

A submanifold $N_{\text{thick}} \subset N$ is a *thick part* for N if:

1. We have $N(\epsilon_1) \subset N_{\text{thick}} \subset N(\epsilon_2)$ for some $\epsilon_1 > \epsilon_2 > 0$, and
2. Each component U of $N - N_{\text{thick}}$ is either an open horoball neighborhood of a cusp, or an open solid torus neighborhood $U = B(\mu, r)$ of a simple geodesic $\mu \subset N$.

By the Margulis lemma, $N(\epsilon_0)$ is a thick part for N for some $\epsilon_0 > 0$. It is useful to also have the more flexible kinds of thick parts defined above.

Models. Let $\pi : (M, h) \rightarrow \Sigma$ be a Riemannian manifold equipped with a map to Σ . Let $\tilde{\Sigma} \rightarrow \Sigma$ be the universal cover, and let $\tilde{M} \rightarrow \tilde{\Sigma}$ be the induced covering of M , equipped with the natural action of $\pi_1(\Sigma)$.

We say M provides a *model* for $N \in H(\Sigma)$ if there is a thick part $N_{\text{thick}} \subset N$, and a $\pi_1(\Sigma)$ -equivariant quasi-isometry

$$\phi : \tilde{M} \rightarrow \tilde{N}_{\text{thick}},$$

where $\tilde{N}_{\text{thick}} \subset \tilde{N} = \mathbb{H}^3$ is the preimage of the thick part in the universal cover of N . The quasi-isometry condition means that ϕ is surjective and that for some constant $K \geq 1$ we have

$$\frac{1}{K}d(f(x), f(y)) - K < d(x, y) < Kd(f(x), f(y))$$

for all $x, y \in \tilde{M}$. Here distances are measured in the path metrics on \tilde{M} and \tilde{N}_{thick} .

Theorem 3.1 (Minsky) *Let $N \in H(\Sigma)$ be a hyperbolic manifold whose limit set is the whole sphere. Then there is a unique geodesic $\tau \subset \mathbb{H}$ with irrational endpoints (ϵ_-, ϵ_+) such that M^τ is a model for N .*

See [Min2, Theorem 11.1]. Minsky's description of the model manifold is different from ours, but easily seen to be equivalent to it.

To see the correspondence, recall that Minsky's model is based the *Farey triangulation* of \mathbb{H} , whose edges E are the hyperbolic geodesics between pairs of rational points $(p/q, r/s)$ in $\partial\mathbb{H} = \mathbb{R} \cup \{\infty\}$ with $|ps - qr| = 1$. Rather than working with the geodesic $\tau = (\epsilon_-, \epsilon_+)$, Minsky works with an infinite chain of edges $(a_i, a_{i+1})_{i \in \mathbb{Z}}$ in E connecting ϵ_- to ϵ_+ . The edges in E separating ϵ_- from ϵ_+ are fans of the form (a_i, b) , where b ranges over n_i rational points in the interval $(a_{i-1}, a_{i+1}) \subset \partial\mathbb{H}$. The bases of the fans (a_i) are called 'pivots'. Each edge (a_i, a_{i+1}) has a natural midpoint τ_{i+1} lying in the thick part of Teichmüller space (and corresponding to a square torus). It is easy to see $d(\tau_i, \tau) = O(1)$ in the Teichmüller metric.

Minsky constructs a Riemannian manifold

$$M = \bigcup B^i = \bigcup \Sigma \times [i, i + 1],$$

quasi-isometric to N , where the metric on $\Sigma \times \{i\}$ is comparable to τ_i . Thus M_{thick} is a model for N_{thick} . The region B^i contains a model Margulis tube whose geometry is determined by the number of Dehn twists n_i about a_i needed to transport a_{i-1} to a_{i+1} .

To compare M_{thick} to our model M^τ , first suppose $\tau = \tau(\text{thick})$. In this case we have $d(\tau_i, \tau_{i+1}) \asymp 1$, so we can find a sequence $g(s_i) \in \tau$, $|s_i - s_{i+1}| \asymp 1$, such that $d(g(s_i), \tau_i) = O(1)$. The constants n_i are also bounded in this case, and therefore B_{thick}^i is quasi-isometric to the submanifold $M_i \subset M^\tau$ between heights s_i and s_{i+1} . Combining these quasi-isometries gives a quasi-isometry between M_{thick} and M^τ . On the other hand, when τ has thin parts $\tau(\mu)$, each $\mu \in \mathcal{S}(\tau)$ coincides with a pivot a_i such that n_i is large. In this case the B_{thick}^i is constructed in much the same way as our $M(\mu)$, so again there is a quasi-isometry from M_{thick} to M^τ . For more details see [Min2], especially §4, §9 and §10.

Remark. The endpoints (ϵ_-, ϵ_+) of τ correspond to the ending laminations of N . Minsky also shows every M^τ actually arises as the model of a unique hyperbolic manifold $N^\tau \in H(\Sigma)$. These results settle Thurston's ending lamination conjecture for the case of cusped tori.

4 Paths in the model

Next we define 3 types of *standard paths* in (M, h) . These paths are well-adapted to structure of (M, h) as a stack of surfaces, and at the same time versatile enough to allow for efficient navigation.

1. Simple paths. We say

$$\alpha : I \rightarrow \Sigma_t$$

is a *simple path* if $t \in T(\text{thick})$ and $\alpha = \pi_t^{-1} \circ \alpha'$ for some simple geodesic segment $\alpha' : I \rightarrow \Sigma$. A simple path is uniquely determined by its slope, orientation and the height t of the surface on which it resides. The diffeomorphism $\pi_t : (\Sigma_t, h) \rightarrow (\Sigma, g(t))$ is not quite an isometry near $\partial\Sigma$, so α need not be a geodesic on (M, h) .

2. Boundary paths. We say

$$\beta : [a, b] \rightarrow \partial\Sigma_t$$

is a *boundary path* if $t \in T(\text{thick})$ and $\beta = \pi_t^{-1} \circ \beta'$ for some boundary geodesic segment $\beta' : I \rightarrow \partial\Sigma$. We can always prolong β to obtain a closed loop, i.e. a path with $\beta(a) = \beta(b)$. The degree of the shortest such prolongation determines the *winding number* $w(\beta) \in \mathbb{Z} - \{0\}$.

3. Vertical paths. For the last type of path, we use the embedding $(\pi, H) : M \rightarrow \Sigma \times \mathbb{R}$ to give the cuspidal part of ∂M a product structure,

$$\partial P \cong \partial\Sigma \times \mathbb{R}. \quad (4.1)$$

We refer to the sets of the form

$$L_q = \pi^{-1}(q) \subset \partial P$$

as *vertical lines*. A *vertical path* is an embedding

$$\xi : [a, b] \rightarrow \partial P \subset \partial M$$

whose image lies along a vertical line, and whose endpoints lie on surfaces Σ_s, Σ_t with $s, t \in T(\text{thick})$.

Note. Unlike simple paths and boundary paths, a vertical path is permitted to pass through one or more of the open submanifolds $M(\mu)$ corresponding to short geodesics. It cannot, however, begin or end in such a region.

Theorem 4.1 *The length of a standard path in (M, h) satisfies:*

- $L(\alpha) \asymp L(\alpha', g(t))$ for a simple path with $\pi \circ \alpha = \alpha'$;
- $L(\beta) \asymp |w(\beta)| + O(1)$ for a boundary path β ; and
- $L(\xi) \asymp |s - t|$ for a vertical path ξ running from Σ_s to Σ_t .

Proof. The map $\pi : (\Sigma_t, h) \rightarrow (\Sigma, g(t))$ is an isometry outside $2B_t = B(p_t, r^2/5)$, and 2-Lipschitz within the ball. Since $\alpha' = \pi \circ \alpha$, we have

$$L(\alpha') \leq 2L(\alpha).$$

On the other hand, π shrinks each arc of $\alpha \cap 2B_t$ by a factor of at most π (because $j(\rho, \theta)$ preserves θ ; see (3.3)), giving the reverse bound $L(\alpha) \leq \pi L(\alpha', g(t))$ as desired.

As for β , the boundary arc makes $|w(\beta)| + O(1)$ revolutions around $\partial\Sigma_t$, whose length is $2\pi r^2/10$. Therefore $L(\beta)$ and $|w(\beta)| + O(1)$ are comparable.

The final estimate, for the length of a vertical arc ξ , is straightforward if ξ lies outside the submanifolds $M(\mu)$ associated to short geodesics $\mu \in \mathcal{S}(\tau)$.

To complete the proof, we consider the case where ξ runs from the top to the bottom of a single submanifold $M(\mu)$.

The metric on $M(\mu)$ is given by a product

$$M(\mu) = \Sigma_u \times T(\mu), \quad u \in T(\mu).$$

Recall there is a second product structure

$$P \cap M(\mu) \cong \partial\Sigma \times T(\mu)$$

coming from (4.1). The transition function

$$h : \partial\Sigma \times T(\mu) \rightarrow \partial\Sigma_u \times T(\mu)$$

has the form $h(v, t) = (j \circ f_t(v), t)$, where $j : \partial\Sigma \rightarrow \partial\Sigma_u$ is a diffeomorphism and where $f_t \subset \mathrm{SL}_2(\mathbb{R})$ is part of a 1-parameter group of parabolic transformations, fixing μ and acting on $\partial\Sigma$ via its identification with the double cover of $\widehat{\mathbb{R}}$.

Since ξ runs along a vertical line $\{v_0\} \times T(\mu) \subset \Sigma \times T(\mu)$, it can be expressed as a parameterized path

$$\xi(t) = (g \circ f_t(v_0), t)$$

in $\Sigma_u \times T(\mu)$. As t ranges through $T(\sigma)$, its image $g \circ f_t(v_0)$ passes through each point on $\partial\Sigma_u$ at most once, showing

$$1 \leq L(\xi) \leq 1 + 2\pi r^2/10$$

and therefore $L(\xi) \asymp 1 = |T(\sigma)| = |t_1 - t_2|$ as claimed. ■

Navigation by standard paths. All standard paths have endpoints in

$$\partial P(\text{thick}) = (\partial P) \cap M(\text{thick}).$$

The next result shows that standard paths are efficient for connecting points in $\partial P(\text{thick})$.

Theorem 4.2 *Let $\gamma : [0, 1] \rightarrow M$ be a path with endpoints in $\partial P(\text{thick})$. Then γ is homotopic (rel endpoints) to a product of standard paths $\gamma_1 * \cdots * \gamma_n$ with*

$$\sum L(\gamma_i) = O(L(\gamma)).$$

Proof. Every point of M is distance $O(1)$ from $\partial P(\text{thick})$. Any long path γ can thus be replaced by one of comparable length that returns periodically to $\partial P(\text{thick})$. It suffices therefore to establish the result in the case $L(\gamma) \leq R$ for some fixed constant R .

Let $\Sigma_s \subset M(\text{thick})$ be the surface containing $\gamma(0)$. Then γ is contained in $T \times I$, where $I = [s - R, s + R]$. The portion of M containing γ admits a height-preserving bilipschitz embedding

$$M \cap (T \times I) \rightarrow \Sigma_s \times I$$

with the product metric. (To check this, recall that the metric $g(t)$ moves at bounded speed for $t \in T(\text{thick})$, while the regions $M(\mu)$ are already products.)

Using the product structure on $\Sigma_s \times I$, we can replace γ with a path $\gamma_1 * \gamma_2$ of comparable length, such that γ_1 is contained in Σ_s and γ_2 is contained in ∂P . We can also assume γ_1 begins and ends at a single basepoint $q \in \partial \Sigma_s$.

Now the metric $g(s)$ on Σ_s lies in the thick part of Teichmüller space, so the free group $\pi_1(\Sigma_s, q)$ is generated by a pair of simple loops δ_1, δ_2 of bounded length. These loops can in turn be chosen as products $\delta_i = \alpha_i * \beta_i$ of surface and boundary paths on Σ_s . Since $L(\gamma_1)$ is bounded below by its word length in $\langle \delta_1, \delta_2 \rangle$, the path γ_1 can be replaced by a product of standard paths of comparable length.

The argument for γ_2 is similar, using boundary paths and vertical paths to navigate in ∂P . ■

5 Model geodesics

In this section we describe a process for lifting a geodesic

$$\gamma' : I \rightarrow \Sigma,$$

with endpoints in $\partial \Sigma$, to a *model geodesic*

$$\gamma : I \rightarrow M$$

with endpoints in ∂P . By *lifting* we mean that γ' and $\pi \circ \gamma$ are homotopic *rel* endpoints.

Our central result, Theorem 6.1, states that these model geodesics are ambient K -quasigeodesics in (M, h) , justifying the terminology. The proof will be given in §6 and §7.

Shortening around model Margulis tubes. As a preliminary, let $\mu \in \mathcal{S}(\tau)$ be a short geodesic, associated to a model Margulis tube $C_\mu \subset T \times \mathbb{R}$. (Recall from (3.1) that $M = (T \times \mathbb{R} - P) - \bigcup_\mu C(\mu)$.)

Let $T(\mu) = (t_-, t_+)$, $g_\pm = g(t_\pm)$, $\Sigma_\pm = \Sigma_{t_\pm}$ and $\pi_\pm = \pi_{t_\pm}$. Let $n = D(\mu)$ be the Dehn twisting number of μ , so that

$$g_+ \approx \text{tw}_\mu^n(g_-).$$

We will show that the multiple loop

$$\mu^n : S^1 \rightarrow \Sigma$$

admits a lift

$$\delta : S^1 \rightarrow \overline{M(\mu)}$$

with $L(\delta) = O(1)$. The fact that μ^n , which may be very long on Σ , can be realized by a loop of bounded length in M , is a key feature which we must take into account to construct efficient model geodesics.

Since the metrics g_\pm lie in the thick part of Teichmüller space, we can find slopes ρ_\pm satisfying

$$\begin{aligned} i(\rho_\pm, \mu) &= 1, \\ L_{\rho_\pm}(g_\pm) &= O(1), \\ \rho_+ &= \text{tw}_\mu^n(\rho_-). \end{aligned}$$

With suitable choice of orientations, μ and ρ_\pm determine homology classes in $H_1(T, \mathbb{Z})$ satisfying

$$[\rho_+] - [\rho_-] = n[\mu]. \tag{5.1}$$

The construction of δ is motivated by equation (5.1). Let $\alpha'_\pm : I \rightarrow \Sigma$ be the geodesic segments with slopes ρ_\pm , oriented so that

$$[\alpha'_\pm] = \pm[\rho_\pm]$$

in $H_1(\Sigma, \partial\Sigma) \cong H_1(T)$. Since ρ_+ and ρ_- differ by an n -fold Dehn twist around μ , we can choose boundary paths β_\pm such that the loop

$$\delta' = \alpha'_+ * \beta'_+ * \alpha'_- * \beta'_-$$

is freely homotopic to μ^n . The boundary paths serve to connect the endpoints of α'_\pm and satisfy $|w(\beta'_\pm)| \leq 1$. See Figure 8 for an example with $n = 4$.

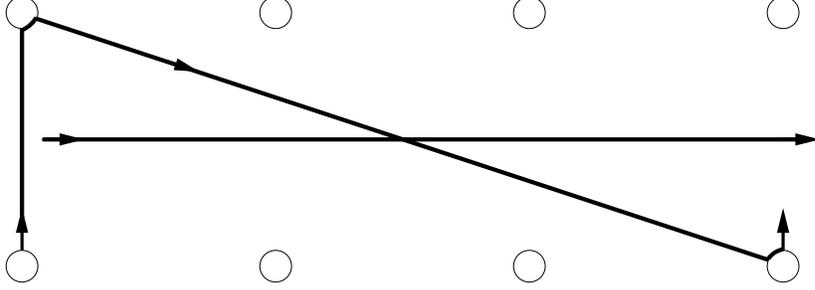


Figure 8. The loop $\alpha'_+ * \beta'_+ * \alpha'_- * \beta'_-$ is freely homotopic to μ^4 .

Now let $\alpha_+ * \beta_+ = \pi_+^{-1} \circ (\alpha'_+ * \beta'_+)$ be a lifting of the first part of δ' to a path on Σ_+ , and let $\alpha_- * \beta_-$ be a similar lifting of the second part to a path on Σ_- . Finally define

$$\delta = \alpha_+ * \beta_+ * \xi_+ * \alpha_- * \beta_- * \xi_-,$$

where ξ_{\pm} are vertical paths joining $\alpha_+ * \beta_+$ to $\alpha_- * \beta_-$. Then $\delta : S^1 \rightarrow \overline{M(\mu)}$ satisfies

$$L(\delta) = O(L_{\rho_+}(g_+) + L_{\rho_-}(g_-) + 1 + |t_+ - t_-|) = O(1),$$

and $\pi \circ \delta = \delta'$ is homotopic to μ^n as desired.

Lifting simple geodesic segments. Next we describe the lifting of a simple geodesic segment $\alpha' : I \rightarrow \Sigma$. Let $\sigma \in \mathcal{S}(T)$ be the slope of α' , and let $g_\sigma \in \tau$ be the optimal metric for σ (minimizing $L_\sigma(g)$ along τ).

We now distinguish several cases.

I. Thick case. If g_σ lies in $\tau(\text{thick})$, then $g_\sigma = g(t)$ for a unique $t \in T(\text{thick})$, and we define

$$\alpha = \pi_t^{-1} \circ \alpha' : I \rightarrow \Sigma_t.$$

Then α is a simple path in M , and by Theorem 4.1 we have

$$L(\alpha, h) \asymp L(\alpha', g(t)).$$

II. Thin case. The case where g_σ lies in $\tau(\text{thin})$ is more interesting. In this case g_σ belongs to $\tau(\mu) = (g_-, g_+)$ for some short geodesic $\mu \in \mathcal{S}(\tau)$, and we adopt notation from the discussion of the Margulis tube C_μ above.

Suppose $L_\sigma(g_+) \leq L_\sigma(g_-)$. We will lift α' to a path α with endpoints in $\partial\Sigma_+$. (In the case $L_\sigma(g_+) > L_\sigma(g_-)$, we use a completely analogous construction to lift α' to a path α with endpoints in $\partial\Sigma_-$.)

Let $[\sigma] = a[\mu] + b[\rho_+]$ in $H_1(T, \mathbb{Z})$.

IIa. $|b| \leq 1$. For simplicity, first assume $|b| \leq 1$. Then α' is homotopic, *rel* endpoints, to a path of the form

$$\gamma'_1 * \mu^a * \gamma'_2,$$

where the arcs γ'_i run along $\partial\Sigma$ and α'_+ , and where the loop μ is broken at its intersection with α'_+ to make it into a path (suitable for concatenation). Since $|b| \leq 1$, we may assume γ'_i has length $O(1)$ in the metric g_+ , and thus it admits a lift $\gamma_i = \pi_+^{-1} \circ \gamma'_i$ to Σ_+ with $L(\gamma_i) = O(1)$.

Let $\mu_+ = \pi_+^{-1} \circ \mu$, and write

$$a = a_1n + a_2$$

with $|a_2| \leq |n|/2$, so that $\mu^a = \mu^{a_1n} * \mu^{a_2}$. Using the fact that μ^n lifts to δ , we lift α' to a path of the form

$$\alpha = \gamma_1 * \delta^{a_1} * \mu_+^{a_2} * \gamma_2, \quad (5.2)$$

where the loop δ (like μ above) is broken at the midpoint of α_+ to make it into a path. Since $L(\delta) = O(1)$ we have

$$L(\alpha) = O(1 + |a_1| + |a_2|).$$

IIb. $|b| > 1$. In general, α' is homotopic to a product of $|b|$ paths of the type treated in case (IIa). By lifting each of these paths as above, we obtain a lift with

$$L(\alpha) = O(|b| + |ba_1| + |ba_2|),$$

where $a/b = a_1n + a_2$ with $|a_2| \leq |n|/2$ as before.

III. Multiple short paths. The last type of path that needs to be lifted specially is one of the form

$$\gamma' = (\alpha' * \beta')^a * \alpha' \quad (5.3)$$

with $a > 0$, where α' is a simple geodesic of slope $\mu \in \mathcal{S}(\tau)$ and where β' is a boundary geodesic of length π . We refer to γ' as a *multiple short path*. In this case γ' is homotopic to a path of the form

$$\gamma' = \gamma_1 * \mu^a * \gamma_2$$

as in case (IIa) above, so it admits a lift with

$$L(\gamma) = O(|a_1| + |a_2|)$$

where $a = a_1n + a_2$, $|a_2| \leq |n|/2$.

Reduced factorization. Let $\gamma' : [a, b] \rightarrow \Sigma$ be a geodesic segment with endpoints in $\partial\Sigma$, with canonical factorization

$$\gamma' = \beta'_0 * \alpha'_1 * \beta'_1 * \cdots * \alpha'_m * \beta'_m$$

into a product of simple geodesics and boundary geodesics. To take advantage of case (III) above, consider a maximal sequence $(i, i+1, \dots, i+a)$ such that the subproduct

$$\alpha' = \alpha'_i * \beta'_i * \cdots * \alpha'_{i+a} = (\alpha'_i * \beta'_i)^a * \alpha'_i$$

is a multiple short path, with slope $\mu \in \mathcal{S}(\tau)$ as in (5.3). Absorbing these subproducts into single paths, we obtain the *reduced factorization*

$$\gamma' = \beta'_0 * \alpha'_1 * \beta'_1 * \cdots * \alpha'_n * \beta'_n \tag{5.4}$$

into a product of simple geodesics, boundary geodesics and multiple short paths.

Model geodesic segments. We can now describe the lifting of a general geodesic $\gamma' : I \rightarrow \Sigma$ to a *model geodesic segment*

$$\gamma : I \rightarrow M.$$

Let $p = (0, 0) \in T$ be the natural basepoint, let $q \in \partial\Sigma$ be the basepoint corresponding to the horizontal ray through $(1, 0) \in \mathbb{T}_p T = \mathbb{R}^2$, and let $q_0 = \pi_0^{-1}(q) \in \Sigma_0 \subset M(\text{thick})$. Let

$$\gamma' : I \rightarrow \Sigma$$

be a geodesic, starting at q and ending on $\partial\Sigma$. We will construct a lift of γ' to a path $\gamma : I \rightarrow M$ starting at q_0 and ending on ∂P .

To construct γ , consider the reduced factorization (5.4) of γ' . Lift each path α'_i in the factorization to a path α_i in M , using the constructions (I), (II) and (III) above. The endpoints of α_i reside in $\partial\Sigma_{t_i}$ for some $t_i \in T(\text{thick})$. Let $t_0 = 0$, and define the lifts of the boundary geodesics by

$$\beta_i = \pi_{t_i}^{-1} \circ \beta'_i, \tag{5.5}$$

so that β_i is contained in $\partial\Sigma_{t_i}$. Let ξ_i be the unique vertical path in ∂P running from height t_i to t_{i+1} and connecting β_i to α_{i+1} . Finally, let

$$\gamma = \beta_0 * \xi_0 * \alpha_1 * \beta_1 * \xi_1 * \cdots * \xi_{n-1} * \alpha_n * \beta_n.$$

Choice of heights. Intuitively, the model geodesic γ is obtained by letting each simple segment α'_i in γ' migrate to the height t_i in M at which it is most efficiently realized, and then connecting these migrated segments together with paths along the parabolic locus. The height t_i is characterized by the condition that

$$g(t_i) = \begin{cases} g_\sigma & \text{if } g_\sigma \in \tau(\text{thick}), \\ g_+ & \text{if } g_\sigma \in \tau(\mu) = (g_-, g_+) \text{ and } L_\sigma(g_+) \leq L_\sigma(g_-), \text{ and} \\ g_- & \text{if } g_\sigma \in \tau(\mu) = (g_-, g_+) \text{ and } L_\sigma(g_-) < L_\sigma(g_+), \end{cases} \quad (5.6)$$

where σ is the slope of α'_i and $g_\sigma \in \tau$ is its optimal metric.

Model geodesic rays. A *model geodesic ray* $\gamma : [0, \infty) \rightarrow M$ is a path based at q_0 obtained by the following construction.

Consider a geodesic ray

$$\gamma' : [0, \infty) \rightarrow \Sigma$$

with $\gamma'(0) = q$. The lifting of geodesic segments generalizes in a straightforward way to the lifting such rays. First we form the reduced factorization

$$\gamma' = \beta'_0 * \alpha'_1 * \beta'_1 * \alpha'_1 \cdots,$$

which may be finite or infinite. We then lift each term and define

$$\gamma = \beta_0 * \xi_0 * \alpha_1 * \beta_1 * \xi_1 * \cdots$$

as above.

When the reduced factorization of γ' is finite, the last term is either a boundary ray, a simple ray or an infinite multiple short path. These three cases are treated specially, as follows.

1. If the last term is a boundary ray β'_n , then we take the last term of γ to be β_n , the lift of β'_n defined by (5.5) as above.
2. If the last term is a simple ray α'_n , then the slope σ of α'_n is an irrational real number. Associated to σ is a length function $L_\sigma : \text{Teich}(T) \rightarrow \mathbb{R}$, well-defined up to scale; it gives the length of a measured foliation on T with leaves of slope σ . We distinguish 3 subcases.

- (a) $L_\sigma(g)$ is minimized along τ at $g_\sigma = g(t) \in \tau(\text{thick})$. Then we define the last term in γ to be the lift $\alpha_n = \pi_t^{-1} \circ \alpha'_n$, as in the thick case (I) above.
 - (b) $L_\sigma(g)$ is minimized at $g_\sigma \in \tau(\mu)$. Then we treat α'_n as an infinite product of paths to which construction (IIa) above applies, and obtain the last term α_n of γ by lifting the infinite product term by term.
 - (c) $L_\sigma(g)$ does not achieve its minimum on τ . Then $\sigma = \epsilon_\pm$ is one of the ending laminations. In this case, we define $t_n = \pm\infty$ to have the same sign as $\sigma = \epsilon_\pm$, and we let ξ_{n-1} be the final term of γ : an infinite vertical path running from height t_{n-1} to height t_n .
3. Finally the last term may be an infinite multiple short path, $\alpha'_n = \alpha * \beta * \alpha * \beta * \dots$, associated to a short geodesic μ . In this case we follow construction (III) above, and lift α'_n to a ray α_n terminating with the infinite product $\delta * \delta * \dots$.

The lifts in these special cases are chosen so that γ varies ‘continuously’ as a function of γ' , as we will see in §9.

6 Efficiency: the thick case

Our next task is to establish that model geodesic rays are close to hyperbolic geodesics. The proof will be presented in 3 stages.

- §6. This section will show that model geodesic rays are ambient K -quasigeodesics in the case where $M = M(\text{thick})$. We treat the thick case first to clarify the main features of the argument.
- §7. The next section will handle the general case, where Margulis tubes and short geodesics are allowed.
- §8. Finally we will show that ambient K -quasigeodesics are close to hyperbolic geodesics.

6.1 Ambient quasigeodesics

Let $\tilde{\Sigma} \rightarrow \Sigma$ be the universal cover, and let

$$\tilde{M} \subset \tilde{\Sigma} \times \mathbb{R} \xrightarrow{\tilde{\pi}} \tilde{\Sigma}$$

the corresponding covering space of $M \subset \Sigma \times \mathbb{R}$ (the latter embedded via (π, H)).

It will be convenient to lift all constructions to these covering spaces. We define various types of paths on \widetilde{M} and $\widetilde{\Sigma}$ (geodesic, simple, boundary, vertical, etc.) to be lifts of the corresponding types of paths on M and Σ . The lengths of all paths on \widetilde{M} will be measured using the pullback \widetilde{h} of the model Riemannian metric h on M .

A path $\gamma : [a, b] \rightarrow \widetilde{M}$ is *K-efficient*, $K \geq 1$, if we have

$$L(\gamma) \leq KL(G) + K \tag{6.1}$$

for any path $G : [a, b] \rightarrow \widetilde{M}$ with the same endpoints. A path $\gamma : I \rightarrow \widetilde{M}$ is an *ambient K-quasigeodesic* if $\gamma|_{[a, b]}$ is *K-efficient* for every $[a, b] \subset I$.

To explain the terminology, suppose $M = N_{\text{thick}} \subset N$ actually coincides with the thick part of a hyperbolic 3-manifold N . Then we have

$$\widetilde{M} = \widetilde{N}_{\text{thick}} \subset \widetilde{N} = \mathbb{H}^3.$$

A path $\gamma : I \rightarrow \widetilde{M}$ is an ambient quasigeodesic if it competes well against other paths $G : I \rightarrow \widetilde{M}$ with the same endpoints. The competing paths G are homotopic to γ in \mathbb{H}^3 , but *not* necessarily in \widetilde{M} . Indeed, $\mathbb{H}^3 - \widetilde{M}$ may contain tubes around lifts of short geodesics, and G and γ can pass around these tubes in different ways. Thus γ models a geodesic in the *ambient manifold* $\widetilde{N} = \mathbb{H}^3$, subject to the constraint that it lies in the thick part $\widetilde{N}_{\text{thick}}$.

This section and the next are devoted to proving:

Theorem 6.1 *There is a universal K such that every model geodesic ray $\gamma : [0, \infty) \rightarrow \widetilde{M}$ is an ambient K -quasigeodesic.*

6.2 From rays to segments

To prove the result for rays, it is sufficient to show:

Theorem 6.2 (Efficiency) *There is a universal K such that for every model manifold M , every model geodesic segment $\gamma : [a, b] \rightarrow \widetilde{M}$ is K -efficient.*

Proof that Theorem 6.2 \implies Theorem 6.1. Let

$$\gamma = \beta_0 * \xi_0 * \alpha_1 * \beta_1 * \xi_1 * \dots$$

be a model geodesic ray, arising as the lift of the geodesic

$$\gamma' = \beta'_0 * \alpha'_1 * \beta'_1 * \cdots .$$

on $\widetilde{\Sigma}$. Let t_i denote the height at which α_i and β_i are realized in \widetilde{M} .

Let $\gamma_0 = \gamma|[a, b]$. We must show γ_0 is K -efficient for a universal K . Since we are assuming Theorem 6.2, it suffices to find a model geodesic segment γ_1 with the endpoints close to those of γ_0 and with

$$L(\gamma_0) \leq KL(\gamma_1) + K$$

for a universal K .

Consider first a subsegment of the form $\gamma_0 = \gamma|[0, a]$. We treat three cases, depending on whether $\gamma_0(a)$ occurs in a factor of γ of the type α_i , β_i or ξ .

Type α_i . Here we have two subcases, according to whether α_i lies in \widetilde{M} (thick) or $\widetilde{M}(\mu)$, $\mu \in \mathcal{S}(\tau)$.

To treat the thick case, let $\gamma'_0 = \widetilde{\pi} \circ \gamma_0 : [0, a] \rightarrow \widetilde{\Sigma}$, and let η' be a geodesic of length $O(1)$ in the $g(t_i)$ -metric, joining $\gamma'_0(a)$ to $\partial\widetilde{\Sigma}$. Let γ'_1 be the geodesic representative of $\gamma'_0 * \eta'$ (*rel* endpoints), and let γ_1 be the model geodesic lifting γ'_1 .

With a suitable choice of η' , we can arrange that the factorization of γ'_1 has the form

$$\gamma'_1 = \beta'_0 * \alpha'_1 * \cdots * \alpha'_{i-1} * \beta''_{i-1} * \alpha''_i * \beta''_i * \cdots * \alpha''_{i+m}$$

for some $m \geq 0$. Here the terminal segment α'_i of γ'_0 (truncated at $\widetilde{\pi} \circ \gamma_0(a)$) has been replaced with a chain of nearly parallel segments $\alpha''_i, \dots, \alpha''_{i+m}$. Using the fact that the metric $g(t_i)$ is nearly optimal for these new segments, it is then straightforward to check that the lift γ_1 of γ'_1 has endpoints close to those of γ_0 , and that

$$L(\gamma_1) \asymp L(\gamma_0)$$

as desired.

To treat the thin case, suppose α_i lies in $M(\mu)$. By adjusting the value of a (moving $\gamma_0(a)$ a bounded distance), we can assume that the truncation of α_i at $\gamma_0(a)$ has the form $\delta^j * \mu_{\pm}^k$, in the notation of (5.2). Thereafter the construction of γ_1 follows the same pattern as the thick case: namely we let $\gamma'_0 = \widetilde{\pi} \circ \gamma_0$, let η' be a bounded geodesic joining $\gamma'_0(a)$ to $\partial\widetilde{\Sigma}$, let γ'_1 be the geodesic representative of $\gamma'_0 * \eta'$ and let γ_1 be the model geodesic lifting γ'_1 . Applying the algorithm for lifting in the thin part presented in §5, we

find that γ_1 and γ_0 have comparable length. Their endpoints are close in \widetilde{M} because γ'_1 is homotopic to $\gamma'_0 * \eta'$.

Type β_i . This case is similar to the preceding one, but easier; γ'_1 is constructed from γ'_0 by replacing the terminal segment β'_i with a subsegment β''_i .

Type ξ_i . Finally suppose $\gamma_0(a)$ occurs at height t along a vertical path ξ_i running from height t_i to t_{i+1} . Choose a slope $\sigma \in \mathcal{S}(T)$ such that $L_\sigma(g(t))$ is minimized at a height $s \in T(\text{thick})$, the minimum is $O(1)$ and $|s - t| = O(1)$. Let α' be a simple geodesic on $\widetilde{\Sigma}$ of slope σ , and let

$$\gamma'_1 = \beta'_0 * \alpha'_1 * \cdots * \alpha'_i * \beta' * \alpha',$$

where β' is a small modification of β'_i to reach the initial point of α' . Then the corresponding model geodesic has the form

$$\gamma_1 = \beta_0 * \xi_0 * \alpha_1 * \beta_1 * \xi_1 * \cdots * \alpha_i * \beta * \xi * \alpha,$$

where ξ now runs from height t_i to s , and where $L(\alpha) = O(1)$. Then γ_1 is a good approximation to the truncation of γ at $\gamma(a)$, so we have shown all initial segments $\gamma|_{[0, a]}$ are K -efficient.

General segments. This completes the proof for a segment $\gamma_0 = \gamma|_{[0, a]}$. To treat a general segment $\gamma_0 = \gamma|_{[a, b]}$, we first change the normalization of the height function $H : M \rightarrow \mathbb{R}$ so that $\gamma(a)$ resides near $\widetilde{\Sigma}_0$. We then adjust the endpoints $\gamma(a), \gamma(b)$ as above, to obtain a geodesic $\gamma'_1 : I \rightarrow \widetilde{\Sigma}$ whose lift γ_1 (based at $\tilde{q}_0 \in \widetilde{\Sigma}_0 \subset \widetilde{M}$) is a model geodesic segment providing a good approximation to γ_0 . Since the efficiency constant for γ_1 (provided by Theorem 6.2) is independent of the normalization of H , the path γ_0 is K -efficient and for a universal K , and hence γ is a K -quasigeodesic. ■

6.3 The model geodesic and its competitor

Next we fix the notation to be used in the proof of efficiency of model geodesics.

Let $\gamma : I \rightarrow \widetilde{M}$ be a model geodesic segment, constructed from a geodesic $\gamma' : I \rightarrow \widetilde{\Sigma}$ with reduced factorization

$$\gamma' = \beta'_0 * \alpha'_1 * \beta'_1 * \cdots * \alpha'_n * \beta'_n.$$

The geodesic γ' is *embedded* in $\widetilde{\Sigma}$, so we will sometimes identify γ' with its image. The corresponding factorization of γ is:

$$\gamma = \beta_0 * \xi_0 * \alpha_1 * \beta_1 * \xi_1 * \cdots * \alpha_n * \beta_n.$$

The paths β_i and α_i reside at height t_i in \widetilde{M} , and ξ_i runs from height t_i to t_{i+1} .

Let $G : I \rightarrow \widetilde{M}$ be a path with the same endpoints as γ . To prove efficiency, we must show

$$L(\gamma) = O(L(G)).$$

By Theorem 4.2, we can use standard paths to navigate efficiently in \widetilde{M} , and thus we can assume G factors as a product

$$G = A_1 * B_1 * X_1 * \cdots * A_m * B_m * X_m \quad (6.2)$$

of simple paths A_j , boundary paths B_j and vertical paths X_j . The paths A_j and B_j reside at height T_j , and X_j runs from height T_j to T_{j+1} . (We allow some factors to be trivial.)

Projecting G to $\widetilde{\Sigma}$, we obtain a piecewise geodesic path

$$G' = \widetilde{\pi} \circ G : I \rightarrow \widetilde{\Sigma}$$

with the same endpoints as γ' , factoring as a product

$$G' = A'_1 * B'_1 s * A'_m * B'_m$$

of simple geodesics and boundary geodesics. For each vertical path X_j , the composition $\widetilde{\pi} \circ X_j$ assumes a constant value $X'_j \in \widetilde{\Sigma}$.

Adjustments. It is convenient to arrange G so that three additional conditions hold:

- (G.1) For any pair of adjacent heights, we have $|T_j - T_{j+1}| \asymp 1$ (or X_j is trivial).
- (G.2) Either $(T_j, T_{j+1}) \subset T(\text{thick})$ or $(T_j, T_{j+1}) = T(\mu)$ for some $\mu \in \mathcal{S}(\tau)$.
- (G.3) A nontrivial path A_j or X_j comes between each pair of boundary paths B_j and B_{j+1} .

Note that (G.1) implies $L(X_j) \asymp 1$ when X_j is nontrivial.

To achieve (G.2), we simply factor any X_j that traverses a thin part $M(\mu)$ so there is a single vertical path that begins and ends on $\partial M(\mu)$. We insert trivial factors A_j and B_j as needed to maintain the format (6.2).

To achieve (G.1), we then amalgamate short vertical paths in $M(\text{thick})$ to form vertical paths of definite length. In the process it may be necessary to change the heights of some paths A_j and B_j by a bounded amount, but

this changes their lengths by a bounded factor since $g(t)$ is Lipschitz for $t \in T(\text{thick})$.

To achieve (G.3), we simply join B_j and B_{j+1} into a single (possibly shorter) boundary path whenever A_j and X_j are both trivial.

Nearest-point projection. Recall that each $g \in \text{Teich}(T)$ determines a nearest-point projection

$$P_g : \tilde{\Sigma} \rightarrow \gamma' \subset \tilde{\Sigma}.$$

Putting these projections together, we define

$$P : \tilde{M} \rightarrow \tilde{\gamma}' \subset \tilde{\Sigma}$$

by using the metric $g(t)$ at level t ; that is, by setting

$$P(x, t) = P_{g(t)}(x)$$

in the coordinates (x, t) given by $(\tilde{\pi}, H) : \tilde{M} \rightarrow \tilde{\Sigma} \times \mathbb{R}$. Since P_g is almost continuous, so is $P : \tilde{M} \rightarrow \tilde{\Sigma}$ (see §2). That is, P is continuous apart from jumps of angle at most π when $P(x)$ belongs to $\partial\tilde{\Sigma}$.

We will use the almost continuous map

$$P \circ G : I \rightarrow \gamma' \subset \tilde{\Sigma}.$$

to compare the model geodesic γ to its competitor $G : I \rightarrow \tilde{M}$. Since $P \circ G$ and γ' have the same endpoints, $P \circ G$ comes within bidistance $(0, \pi)$ of every point of γ' .

6.4 The thick case

We conclude this section by showing model geodesics are K -quasigeodesics in the case where $M = M(\text{thick})$. The next section will address the new issues that arise in the general case, when M may have short geodesics.

Recall that G is a product of paths A_j , B_j and X_j . Let A_{ij} , B_{ij} and X_{ij} denote the subpaths of each that map under P into $\alpha'_i \subset \gamma'$.

Lemma 6.3 *If $M = M(\text{thick})$, then we have*

$$\sum_i L(P \circ X_{ij}, g(T_j)) = O(1).$$

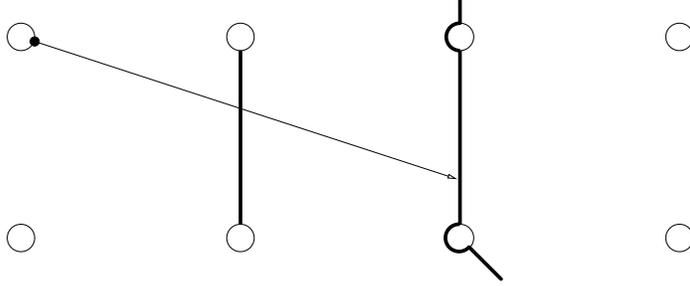


Figure 9. The projection of X_j into γ' has length $O(1)$.

Proof. The sum above is simply the length $L(P \circ X_j, g(T_j)) \subset \gamma'$, ignoring the portion that runs along $\bigcup \beta_i = \gamma \cap \partial \tilde{\Sigma}$. So in the course of the proof we will ignore all paths along the boundary.

We can write the projection of X_j as $P \circ X_j : [T_j, T_{j+1}] \rightarrow \tilde{\Sigma}$, where

$$P \circ X_j(t) = P_{g(t)}(X'_j)$$

and $X'_j \in \tilde{\Sigma}$ is a single point. Assume the path $P \circ X_j(t)$ is nonconstant, and crosses the point $p \in \gamma'$. Then $\delta = \{g \in \mathbb{H} : P_g(X'_j) = p\}$ is a Teichmüller geodesic, consisting of those metrics for which γ' and the geodesic from X'_j to p are perpendicular. Since δ meets the geodesic segment $g[T_j, T_{j+1}] \subset \mathbb{H}$ in a unique point, as t increases, $P \circ X_j(t)$ moves in one direction along γ' . So to bound $L(P \circ X_j, g(T_j))$ from above, we need only bound the distance between its endpoints.

To this end, fix the metric $g(T_j)$ on Σ and $\tilde{\Sigma}$. Choose a pair of slopes σ_1, σ_2 forming a basis for $H_1(T, \mathbb{Z})$, each of length $O(1)$. (This is possible since $g(T_j)$ lies in the thick part of Teichmüller space). Cutting Σ along the simple geodesics parallel to σ_1 and σ_2 , we obtain a parallelogram Q of bounded diameter. The preimages of Q give a tiling of Σ . Since $\pi_1(\Sigma)$ is a free group, the dual graph to this tiling is a tree.

First suppose the tile containing X'_j does not meet γ' . Among the edges of the tiles meeting γ' , let δ denote the edge closest to X'_j . Then δ separates X'_j and γ' , so the nearest-point projection satisfies

$$P_g(X'_j) \subset P_g(\delta) \subset \gamma'$$

for any metric g . See Figure 9.

Recall that the vertical paths X_j have bounded length in M , and therefore $|T_j - T_{j+1}| \asymp 1$ (condition (G.1) above). Thus the metrics $g(T_j)$ and

$g(T_{j+1})$ are uniformly bilipschitz. Now the tiles coming from Q have diameter $O(1)$, so we also have $d_{g(t)}(\delta, \gamma') = O(1)$ for all $t \in [T_j, T_{j+1}]$. Therefore the union of $P_g(\delta)$ over these metrics has bounded $g(T_j)$ -diameter. But this union contains the image of $P \circ X_j$, so the endpoints of $P \circ X_j$ have bounded separation and therefore

$$L(P \circ X_j, g(T_j)) = O(1).$$

Finally, if the tile containing X'_j does meet γ' , then $d(X'_j, \gamma') = O(1)$ already, and we can obtain the preceding bound directly, by letting X'_j take the place of δ . \blacksquare

Remark. When X'_j is very far from γ' , the map $P \circ X_j : [T_j, T_{j+1}] \rightarrow \gamma'$ can be highly expanding at some points. That is why we tame $P_g(X'_j)$ by factoring it through projection from δ , which is close to γ' .

Proof of Theorem 6.2 (Efficiency) – thick case. Assume $M = M(\text{thick})$. To show

$$L(\gamma) = \sum L(\alpha_i) + \sum L(\beta_i) + \sum L(\xi_i) = O(L(G)),$$

we will bound each of the three sums.

Bound on $\sum L(\alpha_i)$. The first bound is based on the observation that if a simple geodesic on Σ is realized on the wrong level in M , then its realization is much longer than it needs to be. Using this idea, we will justify the chain of inequalities:

$$\sum L(\alpha_i) \asymp \sum L(\alpha'_i, g(t_i)) \tag{6.3}$$

$$\leq \sum_{i,j} L(P \circ A_{ij}, g(t_i)) + L(P \circ X_{ij}, g(t_i)) \tag{6.4}$$

$$\leq \sum_{i,j} L(P \circ A_{ij}, g(T_j)) + L(P \circ X_{ij}, g(T_j)) \tag{6.5}$$

$$\leq \sum_j L(A'_j, g(T_j)) + O(L(X_j)) \tag{6.6}$$

$$\asymp \sum L(A_j) + L(X_j) \leq L(G). \tag{6.7}$$

To obtain (6.3), use Theorem 4.1 to replace $\alpha_i \subset \widetilde{M}$ with $\alpha'_i \subset \widetilde{\Sigma}$. For (6.4), note that $P \circ G$ is almost continuous and runs the length of γ' , so it covers every α'_i . On the other hand, $L(P \circ B_{ij}) = 0$ since projections are distance decreasing, and therefore only A_{ij} and X_{ij} contribute to the sum.

Now the length of any subpath of α'_i is minimized along τ by the metric $g(t_i)$, so the sum on the right in (6.4) only increases if we replace $g(t_i)$ by $g(T_j)$. This yields (6.5). To obtain the first term in (6.6), use the fact that for $g = g(T_j)$, we have

$$L(P \circ A_{ij}, g) = L(P_g \circ A'_{ij}, g) \leq L(A'_{ij}, g),$$

since nearest-point projection is distance-decreasing; and that

$$\sum L(A'_{ij}, g) = L(A'_j, g).$$

To obtain the second term, we apply Lemma 6.3 above and use the fact that $L(X_j) \asymp 1$ (when X_j is nontrivial). Finally we use Theorem 4.1 to replace A'_j by A_j , yielding (6.7).

Bound on $\sum L(\xi_i)$. We can sharpen the preceding bound by taking into account the distance between $g(t_i)$ and $g(T_j)$. To make this precise, let $J(i)$ be the least j such that

$$P \circ (A_1 * B_1 * X_1 * \cdots * A_j * B_j * X_j)$$

covers α'_i . Then $J(i)$ is an increasing function of i . Since $P \circ G$ cannot reach the first point of α'_i before covering α'_{i-1} , the path

$$P \circ (A_{J(i-1)} * B_{J(i-1)} * X_{J(i-1)} * \cdots * A_{J(i)} * B_{J(i)} * X_{J(i)})$$

also covers α'_i . Let $\Delta_i = \inf\{|t_i - T_j| : j \in [J(i-1), J(i)]\}$. Using Lemma 2.1, we obtain the bounds

$$L(P \circ A_{ij}, g(t_i)) \cdot \cosh^{1/2} \Delta_i \leq L(P \circ A_{ij}, g(T_j)), \quad (6.8)$$

$$L(P \circ X_{ij}, g(t_i)) \cdot \cosh^{1/2} \Delta_i \leq L(P \circ X_{ij}, g(T_j)), \quad (6.9)$$

for $J(i-1) \leq j \leq J(i)$. Summing over i and over $j \in [J(i-1), J(i)]$, we obtain the improved bound

$$\sum L(\alpha'_i) \cdot \cosh^{1/2} \Delta_i = O\left(\sum L(A_j) + L(X_j)\right) = O(L(G)).$$

This sharpening allows us to bound $\sum L(\xi_i)$. Indeed, $L(\alpha'_i)$ is uniformly bounded below, and $\Delta_i = O(\cosh^{1/2} \Delta_i)$, so we also have

$$\sum \Delta_i = O(L(G)).$$

Now for each i , let $j = j(i)$ minimize $|t_i - T_j|$ in the range $J(i-1) \leq j \leq J(i)$. Then $j(i)$ is an increasing function of i , so we have

$$\begin{aligned}
\sum L(\xi_i) &\asymp \sum |t_i - t_{i+1}| \\
&\leq \sum |t_i - T_{j(i)}| + |T_{j(i)} - T_{j(i+1)}| + |T_{j(i+1)} - t_{i+1}| \\
&\leq 2 \sum \Delta_i + \sum_j |T_j - T_{j+1}| \\
&= O(L(G)) + O\left(\sum L(X_j)\right) = O(L(G)).
\end{aligned}$$

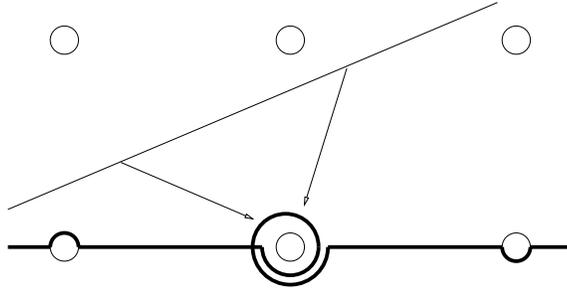


Figure 10. The projection of A'_j to β'_i has length at most π .

Bound on $\sum L(\beta_i)$. Finally, consider a boundary geodesic β'_i in γ' . Let $J_A(i)$ (respectively, $J_B(i)$, $J_X(i)$) denote the set of j such that the projection $P \circ A_j$ (respectively, $P \circ B_j$, $P \circ X_j$) lies entirely in β'_i . We claim the winding number of β'_i satisfies:

$$w(\beta'_i) = O\left(1 + |J_A(i)| + |J_B(i)| + |J_X(i)| + \sum_{j \in J_B(i)} w(B'_j)\right). \quad (6.10)$$

To see this, fix a metric $g \in \text{Teich}(T)$. Then we have $L(\beta'_i, g) = 2\pi w(\beta'_i) + O(1)$. Consider any x in β'_i with $d_g(x, \partial\beta'_i) > 2\pi$. By almost continuity, the image of $P \circ G$ comes within distance π of x . Thus x is contained within a π -neighborhood of the image δ of $P \circ A_j$, $P \circ B_j$ or $P \circ X_j$, for some j . We claim that $\delta \subset \beta'_i$ and j belongs to the corresponding index set $J_A(i)$, $J_B(i)$ or $J_X(i)$.

Indeed, if δ is the image of $P \circ A_j$, then $L(\delta, g) \leq \pi$, since the visual angle of A'_j as seen from β'_i is at most π — see Figure 10. Thus δ is contained within a 2π neighborhood of x and therefore $\delta \subset \beta'_i$.

If δ is the image of $P \circ X_j$, then δ is a single point: the endpoint of the shortest geodesic from X'_j to β'_i . Thus we have $L(\delta, g) = 0$ and $j \in J_X(i)$.

Finally if δ is the image of $P \circ B_j$, then we have $\delta \subset \beta'_i$ and

$$L(\delta, g) \leq L(B'_j, g) \leq 2\pi w(B'_j)$$

because projection is distance-decreasing. Thus we have $j \in J_B(i)$.

Now the π -neighborhoods of the three types of δ discussed above, together with a 2π -neighborhood of $\partial\beta'_i$, form a covering of δ by open sets. Summing up the lengths of the intervals in this covering, we obtain (6.10).

Summing over i , and using the estimates $L(\beta_i) + O(1) \asymp w(\beta'_i) + O(1)$, $L(B_j) \asymp w(B'_j) + O(1)$ provided by Theorem 4.1, we obtain

$$\sum L(\beta_i) = O\left(\sum 1 + |w(\beta'_i)|\right) = O\left(\sum_j 1 + L(B_j)\right).$$

For every j , either $L(A_j)$ or $L(X_j)$ is bounded below, so we have $\sum_j 1 = O(L(G))$, and therefore $\sum L(\beta_i) = O(L(G))$ as well. \blacksquare

7 Efficiency: the thin case

In this section we show that model geodesics are efficient in general.

There are two new circumstances that arise when we allow M to have thin parts $M(\mu)$. First, some segments α_i of the model geodesic γ may lie in the thin part. To distinguish these, we partition the set of indices i into

$$\begin{aligned} I(\text{thick}) &= \{i : p \circ \alpha_i : I \rightarrow M(\text{thick})\} \quad \text{and} \\ I(\mu) &= \{i : p \circ \alpha_i : I \rightarrow \overline{M(\mu)}\}, \quad \mu \in \mathcal{S}(\tau), \end{aligned}$$

where $p : \widetilde{M} \rightarrow M$ is the universal covering map. We also define

$$I(\text{thin}) = \bigcup_{\mu \in \mathcal{S}(\tau)} I(\mu).$$

Second, while the segments A_j and B_j of the competitor G continue to reside in $M(\text{thick})$, some vertical paths X_j may lie in the thin part. (Recall that by (G.2) of §6.3, each X_j lies entirely in $\widetilde{M}(\text{thick})$ or $\widetilde{M}(\text{thin})$.) To distinguish these, we partition the set of indices j into

$$\begin{aligned} J(\text{thick}) &= \{j : p \circ X_j : I \rightarrow M(\text{thick})\} \quad \text{and} \\ J(\mu) &= \{j : p \circ X_j : I \rightarrow \overline{M(\mu)}\}, \end{aligned}$$

and define

$$J(\text{thin}) = \bigcup_{\mu \in \mathcal{S}(\tau)} J(\mu).$$

7.1 Bookkeeping in the thick part

To show $L(\gamma) = O(L(G))$, we first analyze the interaction between the parts of γ and G realized in *unrelated* thin parts of M . This analysis will show that the part of γ realized in $M(\text{thick})$ is efficient.

Theorem 7.1 *For any model geodesic γ and competitor G , we have*

$$\begin{aligned} |I(\text{thick})| + |I(\text{thin})| &= O(L(G)), \\ \sum_{i \in I(\text{thick})} L(\alpha_i) &= O(L(G)), \\ \sum L(\beta_i) &= O(L(G)) \quad \text{and} \\ \sum L(\xi_i) &= O(L(G)). \end{aligned}$$

Notation. Following the conventions of §6, let $P : \widetilde{M} \rightarrow \gamma' \subset \widetilde{\Sigma}$ be the nearest-point projection, and let A_{ij} , B_{ij} and X_{ij} denote the subpaths of A_j , B_j and X_j projecting into $\alpha'_i \subset \gamma'$. Let $J(i)$ be the least j such that

$$P \circ (A_1 * B_1 * X_1 * \cdots * A_j * B_j * X_j)$$

covers α'_i , let $j = j(i) \in [J(i-1), J(i)]$ minimize $|t_i - T_j|$, and let

$$\Delta_i = |t_i - T_{j(i)}| = \inf\{|t_i - T_j| : j \in [J(i-1), J(i)]\}. \quad (7.1)$$

As before, the projection

$$P \circ (A_{J(i-1)} * B_{J(i-1)} * X_{J(i-1)} * \cdots * A_{J(i)} * B_{J(i)} * X_{J(i)})$$

also covers α'_i .

Let us repartition the set of indices i as

$$I(\text{thick}) \sqcup I(\text{thin}) = I'(\text{thick}) \sqcup I'(\text{thin}),$$

where

$$I'(\text{thin}) = \bigcup_{\mu \in \mathcal{S}(\tau)} \{i \in I(\mu) : \exists j \in J(\mu) \cap [J(i-1), J(i)] \text{ with } X_{ij} \neq \emptyset\}.$$

In other words, $I'(\text{thin})$ consists of those i such that α_i and some X_{ij} are realized in the *same* component $M(\mu)$ of M .

Projection of vertical paths. As a preliminary to the proof of Theorem 7.1, we study the projections $P \circ X_j$ for $j \in J(\mu)$, $\mu \in \mathcal{S}(\tau)$. This study in turn requires a geometric analysis of thin metrics on $\tilde{\Sigma}$.

Given a short geodesic $\mu \in \mathcal{S}(\tau)$, let:

- $\hat{\mu} \subset \Sigma$ denote the (image of the) unique simple geodesic segment with slope μ ; and let
- $\mu' \subset \tilde{\Sigma}$ denote the pre-image of $\hat{\mu}$ under the covering map $\tilde{\Sigma} \rightarrow \Sigma$.

Note that $\Sigma - \hat{\mu}$ is an annulus, and $\tilde{\Sigma} - \mu'$ is a countable union of strips.

Lemma 7.2 *For $i \notin I(\mu)$, any component α' of $\alpha'_i - \mu'$ satisfies*

$$L(\alpha', g(t_i)) = O\left(\cosh^{-1/2} d(t_i, T(\mu))\right).$$

Proof. Let σ be the slope of α'_i , and $g_\sigma \in \tau$ the optimal metric for σ . Let $g_\mu \in \tau(\mu)$ be the optimal metric for μ ; it lies at the midpoint of the long segment $\tau(\mu) \subset \tau$. Since $i \notin I(\mu)$, we have $d(g_\sigma, g_\mu) \gg 1$ in the Teichmüller metric.

Let $g_0 \in \partial\tau(\mu)$ denote the endpoint of $\tau(\mu)$ closest to g_σ . Since $L_\mu(g_0) = r \asymp 1$, we have

$$L_\mu(g_\mu) \asymp \cosh^{-1/2} d(g_\mu, g_0).$$

by Lemma 2.1.

In the metric g_μ , the ending laminations ϵ_\pm make angles of 45° with μ' (see Figure 11). As discussed in §2, the distance between the metrics g_σ and g_μ determines the slope of σ : since $d(g_\sigma, g_\mu)$ is large, σ is nearly parallel to one of the ending laminations. Thus $L(\alpha', g_\mu)$ is comparable to the width of the strip of $\tilde{\Sigma} - \mu'$ containing it. This width is comparable to $L_\mu(g_\mu)^{-1}$, since the total area of the torus is one, showing

$$L(\alpha', g_\mu) \asymp \cosh^{1/2} d(g_\mu, g_0).$$

By Lemma 2.1, the length of α' decreases by a factor comparable to $\cosh^{-1/2} d(g_\mu, g_0)$ as g moves from g_μ to g_0 , so we have

$$L(\alpha', g_0) \asymp 1.$$

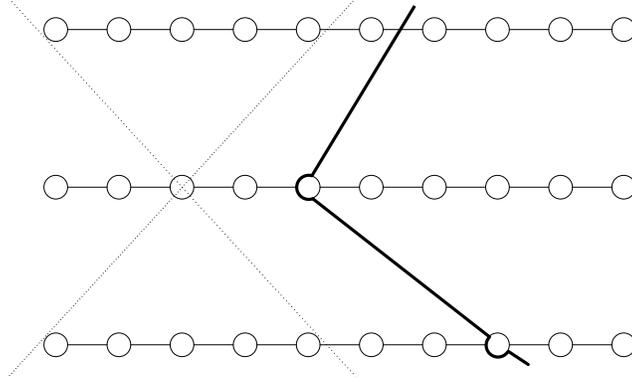


Figure 11. Paths on $\tilde{\Sigma}$ in the g_μ metric. All the short horizontal segments belong to μ' .

Now suppose $i \in I(\text{thick})$, so $g_\sigma = g(t_i)$. Then by Lemma 2.1 we have

$$\begin{aligned} L(\alpha', g(t_i)) &\asymp L(\alpha', g_0) \cdot \cosh^{-1/2} d(g(t_i), g_0) \\ &= O(\cosh^{-1/2} d(t_i, T(\mu))). \end{aligned}$$

Finally suppose $i \in I(\nu)$, $\nu \neq \mu$. Let g_1 denote the endpoint of $\tau(\nu)$ closest to $\tau(\mu)$. Recalling that $g(t_i)$ is the endpoint of $\tau(\nu)$ that minimizes $L_g(\sigma)$ (see equation (5.6)), we again find

$$\begin{aligned} L(\alpha', g(t_i)) &\leq L(\alpha', g_1) \\ &\asymp L(\alpha', g_0) \cdot \cosh^{-1/2} d(g_0, g_1) \\ &= O(\cosh^{-1/2} d(t_i, T(\mu))). \end{aligned}$$

■

Lemma 7.3 *For any vertical path X_j with $j \in J(\mu)$, the projection $P \circ X_j$ meets at most 3 components S_0, S_1, S_2 of $\tilde{\Sigma} - \mu'$.*

Proof. By changing parameters, we can assume the vertical path has the form $X_j : T(\mu) \rightarrow \tilde{M}$, with $X_j(t)$ at height t and

$$P \circ X_j(t) = P_{g(t)}(X'_j).$$

Recall that $g(t)$ moves along the horocycle $\eta(\mu)$ as t moves along $T(\mu)$.

By taking the closures of the components of $\widetilde{\Sigma} - \mu'$, we obtain a tiling of $\widetilde{\Sigma}$ by infinitely many strips, any two of which meet along a component of μ' . The dual graph of this tiling is a tree of infinite degree. Let S_0 denote the strip meeting γ' that is closest to X'_j . Let S_1 and S_2 denote the strips adjacent to S_0 that also meet γ' . (In case γ' begins or ends in S_0 , set $S_1 = S_0$ or $S_2 = S_0$.)

We claim for any metric $g = g(t)$ with $t \in T(\mu)$, the nearest point projection $P_g(X'_j)$ lies in $S_0 \cup S_1 \cup S_2$. To see this, let δ be the shortest geodesic (in the g -metric) from X'_j to γ' . Then δ must enter S_0 to reach γ' . If δ does not terminate in S_0 , then it must enter S_1 or S_2 .

But to reach these adjacent strips, δ must cross the same component of μ' that γ' does. By the definition of $g(t)$, we have $L_\mu(g) = r \ll 1$, so as δ enters S_1 or S_2 it comes within distance r of γ' . Since δ is the shortest path from X'_j to γ' , it can travel at most distance r inside S_1 or S_2 . The width of each strip is at least $1/r$, so δ must terminate on γ' in one of these two adjacent strips. ■

Corollary 7.4 *For $j \in J(\text{thin})$, we have $X_{ij} \neq \emptyset$ for at most 3 values of i .*

Proof. The projection $P \circ X_j$ is contained in $\gamma' \cap (S_0 \cup S_1 \cup S_2)$. By convexity of γ' , each strip S_k meets at most a single geodesic segment α'_i . ■

Lemma 7.5 *We have*

$$\sum_{i \in I'(\text{thick})} \sum_{j \in J(\text{thin}) \cap [J(i-1), J(i)]} L(P \circ X_{ij}, g(t_i)) \cdot \cosh^{1/2} \Delta_i = O\left(\sum L(X_j)\right).$$

(Here Δ_i is defined by (7.1).)

Proof. Fix i and j such that the path X_{ij} contributes to the sum above. Then $j \in J(\mu)$ for some $\mu \in \mathcal{S}(\tau)$. Let S_0, S_1, S_2 be the strips covering $P \circ X_j$ provided by Lemma 7.3. Since $i \in I'(\text{thick})$ and $j \in [J(i-1), J(i)]$, we have $i \notin I(\mu)$. Applying Lemma 7.2 with $\alpha' = \alpha'_i \cap S_k$, we obtain

$$L(\alpha'_i \cap S_k, g(t_i)) \cdot \cosh^{1/2} d(t_i, T(\mu)) = O(1) = O(L(X_j)).$$

Summing over $k = 0, 1, 2$ and using the fact that $\Delta_i \leq d(t_i, T(\mu)) + O(1)$ yields

$$L(P \circ X_{ij}, g(t_i)) \cdot \cosh^{1/2} \Delta_i = O(L(X_j)).$$

By the preceding Corollary, for each j there are at most 3 values of i with $X_{ij} \neq \emptyset$, so when we sum the bound above over i and j we obtain the Lemma. ■

Corollary 7.6 *We have*

$$\sum_{i \in I'(\text{thick})} \sum_{j \in [J(i-1), J(i)]} L(P \circ X_{ij}, g(t_i)) \cdot \cosh^{1/2} \Delta_i = O\left(\sum L(X_j)\right).$$

Proof. The arguments from the thick case (such as the bound (6.9)) control the part of the sum coming from $j \in J(\text{thick})$, and the preceding Lemma controls the part coming from $j \in J(\text{thin})$. ■

Lemma 7.7 *For any model geodesic γ and competitor G , we have*

$$\sum_{I'(\text{thick})} L(\alpha_i) \cdot \cosh^{1/2} \Delta_i = O(L(G)).$$

Proof. First, note that for any i we have

$$L(\alpha_i) = O(L(\alpha'_i, g(t_i))).$$

Indeed, for $i \in I(\text{thick})$ the two sides are comparable by Theorem 4.1, while for $i \in I(\text{thin})$, the efficient lift α_i of α'_i to the boundary of a Margulis tube is, if anything, much shorter than its realization at the top or bottom of the tube, which is at level t_i .

Second, recall from the thick case that we have the bound

$$L(P \circ A_{ij}, g(t_i)) \cosh^{1/2} \Delta_i = O(L(A_{ij})).$$

Also recall that the subpath of $P \circ G$ corresponding to the indices $j \in [J(i-1), J(i)]$ covers α'_i .

From these observations and Corollary 7.6 we obtain:

$$\begin{aligned} \sum_{I'(\text{thick})} L(\alpha_i) \cdot \cosh^{1/2} \Delta_i &= O\left(\sum_{I'(\text{thick})} L(\alpha'_i, g(t_i)) \cdot \cosh^{1/2} \Delta_i\right) \\ &= O\left(\sum_{i \in I'(\text{thick})} \sum_{j \in [J(i-1), J(i)]} (L(P \circ A_{ij}, g(t_i)) + L(P \circ X_{ij}, g(t_i))) \cosh^{1/2} \Delta_i\right) \\ &= O\left(\sum_j L(A_j) + L(X_j)\right) = O(L(G)). \end{aligned}$$

■

Proof of Theorem 7.1. Bound on $\sum_{I(\text{thick})} L(\alpha_i)$. This bound is immediate from Lemma 7.7.

Bound on $\sum L(\beta_i)$. This bound is proved just as in the thick case.

Bound on $|I(\text{thick})| + |I(\text{thin})|$. We will show, equivalently, that $|I'(\text{thick})|$ and $|I'(\text{thin})|$ are both $O(L(G))$. Since $L(\alpha_i)$ is bounded below, the bound on $|I'(\text{thick})|$ is immediate from Lemma 7.7. To bound $|I'(\text{thin})|$, note that for each $i \in I'(\text{thin})$ there exists a $j \in J(\text{thin})$ with $X_{ij} \neq \emptyset$. But by Corollary 7.4, for such j we have $X_{ij} \neq \emptyset$ for at most 3 values of i . Therefore we have

$$|I'(\text{thin})| \leq 3|J(\text{thin})| = O\left(\sum L(X_j)\right) = O(L(G)).$$

Bound on $\sum L(\xi_i)$. As in the thick case, we have

$$\sum L(\xi_i) \asymp \sum |t_i - t_{i+1}| \leq \sum |T_{j(i)} - T_{j(i+1)}| + 2 \sum \Delta_i. \quad (7.2)$$

Now for $i \in I'(\text{thin})$, both t_i and $t_{j(i)}$ belong to the same unit-length interval $T(\mu)$, so $\Delta_i = O(1)$ and thus

$$\sum_{I'(\text{thin})} \Delta_i = O(|I'(\text{thin})|) = O(L(G)).$$

On the other hand, using Lemma 7.7 we have

$$\sum_{I'(\text{thick})} \Delta_i = O\left(\sum_{I'(\text{thick})} L(\alpha_i) \cosh^{1/2} \Delta_i\right) = O(L(G)).$$

Finally we have

$$\sum_i |T_{j(i)} - T_{j(i+1)}| \leq \sum_j |T_j - T_{j+1}| \asymp \sum L(X_j) = O(L(G)).$$

Combining these three bounds with (7.2) shows $\sum L(\xi_i) = O(L(G))$. \blacksquare

7.2 Thin efficiency

We now turn to the heart of the argument, which is to show that the part of γ realized in $M(\text{thin})$ is also efficient.

Theorem 7.8 *For any model geodesic γ and competitor G , we have*

$$\sum_{i \in I(\text{thin})} L(\alpha_i) = O(L(G)).$$

To discuss the proof, fix $\mu \in \mathcal{S}(\tau)$. We will first control $\sum_{i \in I(\mu)} L(\alpha_i)$. For each $i \in I(\mu)$, let S_i denote the union of the closures of the components of $\tilde{\Sigma} - \mu'$ that meet α'_i .

Since G' is homotopic to γ' , it has a subpath G'_i that enters and exits S_i along the same components of μ' that α'_i does (see Figure 12). We say G'_i *tracks* α_i . The strips S_i for different $i \in I(\mu)$ have disjoint interiors, so we choose the subpaths G'_i so they are also disjoint for different $i \in I(\mu)$. Note that G'_i does *not* have to stay inside S_i , although it is homotopic *rel* endpoints to a path which does.

Let G_i denote the subpath of G lying over G'_i , and let A_{ij} , B_{ij} and X_{ij} denote the subpaths of A_j , B_j and X_j that are included in G_i . Let

$$\Delta_{\mu j} = d(T(\mu), T_j).$$

Lemma 7.9 (Thin efficiency) *For any $i \in I(\mu)$, we have*

$$L(\alpha_i) = O \left(1 + \sum_j (L(A_{ij}) + L(X_{ij})) \cdot \cosh^{-1/2} \Delta_{\mu j} \right).$$

To present the proof of this key Lemma, we adopt the notation of §5. That is, we let:

- (t_-, t_+) denote the heights of the bottom and top of $M(\mu)$;
- $(g_-, g_+) = (g(t_-), g(t_+))$ denote the endpoints of $\tau(\mu)$;
- (ρ_-, ρ_+) denote slopes complementary to μ , adapted to the metrics (g_-, g_+) ;
- $n = D(\mu)$ denote the Dehn twist number of μ ;
- σ denote the slope of α'_i ; and
- $a, b \geq 0$ denote integral coefficients such that

$$[\sigma] = a[\mu] + b[\rho_+] \tag{7.3}$$

in $H_1(T, \mathbb{Z})$, with suitable orientations.

We will treat in detail the case $b \leq 1$ (case (IIa) of §5). The other cases ((IIb) and (III)) follow the same pattern.

So assume $b \leq 1$. Then S_i is the closure of a single component of $\tilde{\Sigma} - \mu'$. Let $\tilde{\mu} \subset S_i$ denote an infinite geodesic of slope μ , very close to the edge of S_i where α_i exits the strip (see Figure 12). For each $g \in \text{Teich}(T)$, let

$$P_g^\mu : \tilde{\Sigma} \rightarrow \tilde{\mu}$$

denote the nearest-point projection in the g -metric. Define

$$P^\mu : \tilde{M} \rightarrow \tilde{\mu} \subset \tilde{\Sigma}$$

by the formula

$$P^\mu(x, t) = P_{g(t)}^\mu(x),$$

using coordinates coming from the natural inclusion $(\tilde{\pi}, H) : \tilde{M} \hookrightarrow \tilde{\Sigma} \times \mathbb{R}$.

For any path $D : I \rightarrow \tilde{\mu}$, let

$$N(D) = \frac{L(D, g_+)}{L_\mu(g_+)}.$$

The number $N(D)$ measures the number of times $P^\mu \circ D$ wraps around μ , when projected down to Σ .

Let $\tilde{G}_i : I \rightarrow \tilde{\mu}$ denote the embedded subarc of $\tilde{\mu}$ joining the endpoints of $P^\mu \circ G_i$. The projections of A_{ij} and X_{ij} are also embedded subarcs of $\tilde{\mu}$, so we have

$$N(\tilde{G}_i) = \sum_j \pm N(P^\mu \circ A_{ij}) + \sum_j \pm N(P^\mu \circ X_{ij}) \quad (7.4)$$

for suitable choices of signs. (Note that $N(P^\mu \circ B_{ij}) = 0$).

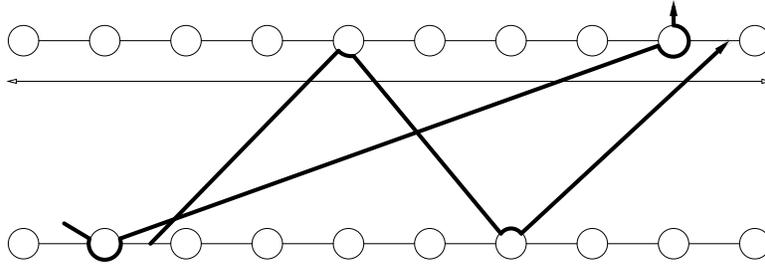


Figure 12. A segment α'_i associated to a short geodesic μ , and the competing path G'_i that tracks it.

Lemma 7.10 *The subarc \tilde{G}_i of $\tilde{\mu}$ satisfies*

$$N(\tilde{G}_i) = |a + \epsilon n| + O(1)$$

with $\epsilon = -1, 0$ or 1 .

(Here $n = D(\mu)$ and a is defined by (7.3).)

Proof. For the course of the proof, we fix the metric g_+ on $\tilde{\Sigma}$, and measure angles and distances using this metric.

Let $\tilde{\alpha}_i = P_{g_+}^\mu \circ \alpha'_i : I \rightarrow \tilde{\mu}$ be the projection of α'_i to $\tilde{\mu}$. Since α'_i has slope σ we have

$$N(\tilde{\alpha}_i) = i(\sigma, \rho_+) + O(1) = a + O(1).$$

Now G'_i and α'_i both exit the strip S_i close to $\tilde{\mu}$, along the same short component of μ' , so the corresponding terminal points of \tilde{G}_i and $\tilde{\alpha}_i$ are close to each other (see Figure 12).

To compare the initial points of \tilde{G}_i and $\tilde{\alpha}_i$, we first remark that the ending lamination ϵ_+ is nearly perpendicular to μ in the metric g_+ . Indeed, the geodesic $\tau = (\epsilon_-, \epsilon_+)$ passes deep into the horocycle $H_\mu(r)$ before reaching g_+ , so ϵ_- and μ are close together on S_∞^1 as seen from $g_+ \in \text{Teich}(T)$; thus ϵ_- and μ are nearly parallel, so ϵ_+ and μ are nearly perpendicular. Similarly, ϵ_- and μ are nearly perpendicular in the g_- -metric. Since we have $g_+ \approx \text{tw}_\mu^n(g_-)$, the projections from the lower end of S_i to the upper edge, along ϵ_- and ϵ_+ , differ by $n + O(1)$ multiples of $L_\mu(g_+)$ (see Figure 13).

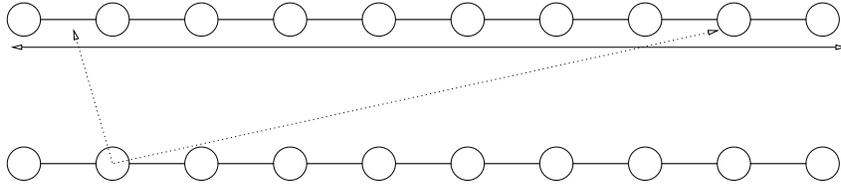


Figure 13. Slopes of the ending laminations in the metric g_+ .

Choose j such that the initial point of G_i occurs along A_{ij} or X_{ij} ; in particular, at height t_j . Then $t_j \notin T(\mu)$, so we have $t_j \geq t_+$ or $t_j \leq t_-$.

Within the strip S_i , the projection P_g^μ to $\tilde{\mu}$ takes the form of a linear projection along lines of constant slope, namely the g -normals to $\tilde{\mu}$.

For $t_j \geq t_+$, the $g(t_j)$ -normals are close to the g_+ -normals, since both are close to ϵ_+ . Since the initial points of G'_i and α'_i are close, their $g(t_j)$

and g_+ projections to $\tilde{\mu}$ are also close, so the initial points of \tilde{G}_i and $\tilde{\alpha}_i$ are close. Therefore we have $N(\tilde{G}_i) = N(\tilde{\alpha}_i) + O(1) = a + O(1)$.

For $t_j \leq t_-$, the $g(t_j)$ -normals to μ are close to ϵ_- . Since projection along ϵ_- shifts by $|n|$ multiples of the length of μ , we have $N(\tilde{G}_i) = |N(\tilde{\alpha}_i) \pm n| + O(1)$. So in either case we obtain the Lemma. ■

Lemma 7.11 *For $j \in J(\mu)$, we have*

$$N(P^\mu \circ X_{ij}) = \epsilon n + O(1),$$

where $\epsilon = 0$ or 1 .

Proof. We may assume $X_{ij} = X_j$, since the entire path X_j appears in G_i whenever the single vertex X'_j appears in G'_i . Let $x \in \partial S_i$ denote a point on the component of μ' separating X'_j from $\text{int}(S_i)$, and define

$$\tilde{X}_j : \tau(\mu) \rightarrow \tilde{\mu}$$

by $\tilde{X}_j(g) = P_g^\mu(x)$. It is easy to see that $N(P^\mu \circ X_{ij}) = N(\tilde{X}_j) + O(1)$, and clearly $N(\tilde{X}_j) = O(1)$ if x lies along the edge of S_i close to $\tilde{\mu}$.

Finally, if x lies along the far edge of S_i , then the endpoints $(P_{g_-}^\mu(x), P_{g_+}^\mu(x))$ of \tilde{X}_j are separated by $|n| + O(1)$ copies of μ , since $g_+ = \text{tw}_\mu^n(g_-)$. ■

Proof of Lemma 7.9 (Thin efficiency). We can now address the key point in the proof — shortening in the Margulis tube. Let

$$f(x) = \left\lfloor \frac{x}{n} \right\rfloor + d(x, n\mathbb{Z}).$$

We will use the following elementary properties of f :

$$\begin{aligned} f(x) &\leq 2|x|, \\ f(x+y) &\leq f(x) + f(y) \quad \text{and} \\ f(x+n) &= f(x) + O(1). \end{aligned}$$

Recall that μ^n admits a lifting to M with length $O(1)$, and that we take advantage of this shortening in the construction of α_i (§5) with the result that

$$L(\alpha_i) = O(f(a)).$$

Now $N(\tilde{G}_i) = |a + \epsilon n| + O(1)$ by Lemma 7.10; since $f(a + n) = f(a) + O(1)$, we have

$$L(\alpha_i) = O(1 + f(N(\tilde{G}_i))). \quad (7.5)$$

Applying f to both sides of equation (7.4) and using the fact that $f(x + y) \leq f(x) + f(y)$, we obtain

$$f(N(\tilde{G}_i)) \leq \sum_j f(N(P^\mu \circ A_{ij})) + f(N(P^\mu \circ X_{ij})). \quad (7.6)$$

Next we establish the bounds

$$f(N(P^\mu \circ A_{ij})) = O(L(A_{ij}) \cdot \cosh^{-1/2} \Delta_{\mu j}), \quad (7.7)$$

$$f(N(P^\mu \circ X_{ij})) = O(L(X_{ij}) \cdot \cosh^{-1/2} \Delta_{\mu j}). \quad (7.8)$$

The first bound follows from the inequality

$$\begin{aligned} N(P^\mu \circ A_{ij}) &= O(L(P^\mu \circ A_{ij}, g_+)) \\ &= O(L(P^\mu \circ A_{ij}, g(T_j)) \cosh^{-1/2} \Delta_{\mu j}) \\ &= O(L(A'_{ij}, g(T_j)) \cosh^{-1/2} \Delta_{\mu j}) \\ &= O(L(A_{ij}) \cdot \cosh^{-1/2} \Delta_{\mu j}). \end{aligned}$$

Here we use the bound $1/L_\mu(g_+) = O(1)$ to compare $N(\cdot)$ and $L(\cdot)$, Lemma 2.1 to compare lengths in the metrics g_+ and $g(T_j)$, and the fact that projections are contracting. Since $f(x) = O(x)$, we obtain (7.7).

Similarly, the proof of Lemma 7.5 can be adapted to show that (7.8) holds for $j \notin J(\mu)$, while Lemma 7.11 and the fact that $f(x + n) = f(x) + O(1)$ shows (7.8) holds $j \in J(\mu)$.

Substituting (7.7) and (7.8) into (7.6), we obtain the desired estimate on $L(\alpha_i)$ from (7.5). \blacksquare

Proof of Theorem 7.8. Recall that the subpaths G_i of G are disjoint for different $i \in I(\mu)$. Thus given μ and j , we have $L(A_{ij}) + L(X_{ij}) > 0$ for at most one $i \in I(\mu)$. Using the preceding Lemma, we then obtain:

$$\begin{aligned} \sum_{i \in I(\text{thin})} L(\alpha_i) &= \sum_{I(\text{thin})} O \left(1 + \sum_j (L(A_{ij}) + L(X_{ij})) \cdot \cosh^{-1/2} \Delta_{\mu j} \right) \\ &= O \left(|I(\text{thin})| + \sum_j (L(A_j) + L(X_j)) \sum_\mu \cosh^{-1/2} \Delta_{\mu j} \right). \end{aligned}$$

By Theorem 7.1, we have $|I(\text{thin})| = O(L(G))$. To control the last term, note that

$$\sum_{\mu} \cosh^{-1/2} \Delta_{\mu j} = O(1),$$

since the sum is bounded by a geometric series; and that $\sum_j L(A_j) + L(X_j) \leq L(G)$. Combining these bounds, we obtain $\sum_{I(\text{thin})} L(\alpha_i) = O(L(G))$. ■

Proof of Theorem 6.2 (Efficiency). Theorems 7.1 and 7.8 together imply

$$L(\gamma) = \sum L(\alpha_i) + L(\beta_i) + L(\xi_i) = O(L(G)).$$

Thus the model geodesic γ is K -efficient for a universal constant K . ■

8 Ambient quasigeodesics

In this section we show any ambient K -quasigeodesic ray follows a hyperbolic geodesic, modulo the thin part, and therefore it converges to a unique point on the sphere at infinity.

Cylinders and horoballs. Since the proof is general, we will formulate the statement so it applies in any dimension.

A *cylinder* $B(\delta, r) \subset \mathbb{H}^n$ is an r -neighborhood of a complete hyperbolic geodesic δ . Let \mathcal{H} be a locally finite collection of disjoint open cylinders and horoballs in \mathbb{H}^n .

Recall that a path $\gamma : I \rightarrow X$ in a Riemannian manifold X is an *ambient K -quasigeodesic* if we have

$$L(\alpha) \leq KL(A) + K$$

for any subsegment $\alpha = \gamma|[a, b]$ and any path $A : [a, b] \rightarrow M$ with the same endpoints as α .

Theorem 8.1 *Let $\gamma : I \rightarrow \mathbb{H}^n - \bigcup \mathcal{H}$ be an ambient K -quasigeodesic segment. Let $\eta \subset \mathbb{H}^n$ be the hyperbolic geodesic segment with the same endpoints as γ , and let $\mathcal{H}(\eta)$ be the union of all $H \in \mathcal{H}$ meeting η . Then we have*

$$\gamma(I) \subset B(\eta \cup \mathcal{H}(\eta), R),$$

where $R = R(K)$ depends only on K .

Proof of Theorem 8.1. For $r > 0$, let $Y_r = \eta \cup \mathcal{H}_r(\eta)$, where

$$\mathcal{H}_r(\eta) = \bigcup \{H \in \mathcal{H} : B(x, r) \subset H \text{ for some } x \in \eta\}.$$

For r sufficiently large, each component H of $\mathcal{H}_r(\eta)$ is locally close to a horosphere normal to η . Using this information, we can construct a closed convex set $C \subset \mathbb{H}^n$ and a map $\rho : \partial C \rightarrow \partial Y_r$ satisfying

1. $Y_r \subset C \subset B(Y_r, 10)$;
2. ρ is M -Lipschitz for a universal constant M ; and
3. $d(x, \rho(x)) \leq 10$.

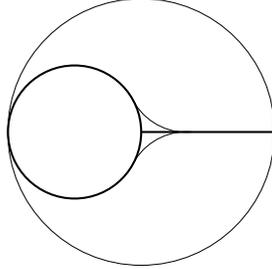


Figure 14. Thickening $\eta \cup H$ to a convex set.

In fact, C can be taken as the convex hull of $\mathcal{H}_r(\eta) \cup B(\eta, 1/100)$. To obtain C from $Y_r = \eta \cup \mathcal{H}_r(\eta)$, one slightly thickens η , then rounds off the corners of the junction between η and ∂H for each component H of $\mathcal{H}_r(\eta)$. Since $B(\eta, 1/100)$ is strictly convex, the rounding off is localized near the junction (see Figure 14), and one finds $C \subset B(Y_r, 10)$. Away from the junctions we can take ρ to be the nearest point projection to Y_r , while near the junctions we can use a local retraction with a uniform Lipschitz constant, to obtain properties (2) and (3) above.

For $s > 0$, consider a maximal subsegment $\alpha = \gamma|[a, b]$ lying outside $B(C, s)$. Let $\pi_C : \mathbb{H}^n \rightarrow C$ be the nearest point projection, and let $A_1 = \rho \circ \pi_C \circ \alpha$; then we have

$$L(A_1) \leq M \cdot L(\alpha) / \cosh(s)$$

because ρ is M -Lipschitz and π_C contracts by at least a factor of $\cosh(s)$ outside $B(C, s)$.

Since α is a maximal interval outside $B(C, s)$, its endpoints lie on $\partial B(C, s)$. Thus the endpoints of A_1 are at most distance $s + 10$ from those of α . The hyperbolic geodesics G_0, G_2 that connect corresponding endpoints may pass through $\bigcup \mathcal{H}$, but deforming these geodesics to the boundaries of the cylinders and horoballs increases their lengths by a controlled amount. Thus we can join corresponding endpoints by paths A_0, A_2 , disjoint from $\bigcup \mathcal{H}$, that satisfy $L(A_0) + L(A_2) \leq M(s)$ for a constant $M(s)$ depending only on s .

Let $A = A_0 * A_1 * A_2$. Then A has the same endpoints as α , it lies in $\mathbb{H}^n - \bigcup \mathcal{H}$ and it satisfies

$$L(A) \leq 2M(s) + M \cdot L(\alpha) / \cosh(s).$$

Since γ is an ambient K -quasigeodesic, α is K -efficient, so we have

$$L(\alpha) \leq KL(A) + K.$$

Choosing $s = s(K)$ such that $M/\cosh(s) < K/2$, the two equations above lead to a bound of the form $L(\alpha) \leq f(K)$. Therefore we have

$$\gamma(I) \subset B(C, s(K) + f(K)) \subset B(Y_r, R(K))$$

where $R(K) = 10 + s(K) + f(K)$. Since Y_r is contained in $\eta \cup \mathcal{H}(\eta)$, the Theorem is proved. \blacksquare

Theorem 8.2 *Let $\gamma : [0, \infty) \rightarrow \mathbb{H}^n - \bigcup \mathcal{H}$ be an ambient K -quasigeodesic ray of infinite length. Then $\gamma(s)$ converges, as $s \rightarrow \infty$, to a unique point $x \in S_\infty^{n-1}$.*

Proof. We may assume γ is parameterized by arclength. By the quasigeodesic property, we have

$$d(\gamma(0), \gamma(s)) \geq s/K - K \rightarrow \infty,$$

so $\gamma : [0, \infty) \rightarrow \mathbb{H}^n$ is proper. Therefore the set of accumulation points A of γ on S_∞^{n-1} is nonempty and connected.

On the other hand, by applying the preceding Theorem to $\gamma|_{[0, s]}$ and taking a limit as $s \rightarrow \infty$, we obtain a geodesic ray η satisfying

$$\gamma[0, \infty) \subset B(\eta \cup \mathcal{H}(\eta), R).$$

Therefore we have

$$A \subset B = \overline{\eta \cup \mathcal{H}(\eta)} \cap S_\infty^{n-1}.$$

Since \mathcal{H} is locally finite, B is a countable set (accumulating only on the endpoint of η). Therefore A consists of a single point $\{x\}$, and $x = \lim \gamma(s)$. \blacksquare

Corollary 8.3 *The ray γ is contained in an $R(K)$ -neighborhood of $\eta \cup \mathcal{H}(\eta)$, where η is the unique geodesic ray connecting $\gamma(0)$ to $x \in S_\infty^{n-1}$.*

Proof. Apply Theorem 8.1 to $\gamma|[0, s]$ and take the limit as $s \rightarrow \infty$. ■

Notes. The proof of Theorem 8.1 is adapted from a standard argument in the proof of Mostow rigidity; see, for example, [BP, Lemma C.1.6]. Related results concerning quasigeodesics outside horoballs appear in [ECH, §11.3] and [La].

9 Conclusion: Continuity

Let $N \in H(\Sigma)$ be a hyperbolic 3-manifold with limit set the whole sphere, marked by a homotopy equivalence $f : \Sigma \rightarrow N$. Then we can regard N as a quotient

$$N = \mathbb{H}^3 / \pi_1(\Sigma),$$

using the marking f to obtain an action of $\pi_1(\Sigma) \cong \pi_1(N)$ on $\tilde{N} = \mathbb{H}^3$ by hyperbolic isometries. Passing to the sphere at infinity, we obtain an action of $\pi_1(\Sigma)$ on $S_\infty^2 = \partial\mathbb{H}^3$ by conformal automorphisms.

Fix a basepoint $q \in \partial\Sigma$ and let $S_\infty^1(q)$ be the natural circle at infinity introduced in §2. This circle comes equipped with an action of $\pi_1(\Sigma, q)$, topologically equivalent to the action of the Fuchsian group of a cusped torus on $\partial\mathbb{H}^2$.

In this section we will finally establish:

Theorem 9.1 *For any $N \in H(\Sigma)$ with limit set the whole sphere, there is a natural, continuous, surjective map*

$$F : S_\infty^1(q) \rightarrow S_\infty^2,$$

respecting the action of $\pi_1(\Sigma, q)$.

As remarked in the introduction, the same result holds for any $N \in H(\Sigma)$ if we replace S_∞^2 with the limit set $\Lambda \subset S_\infty^2$.

From the circle to the limit set. We begin by defining the map F . Recall from §2 that $S_\infty^1(q) = \mathcal{G}(q) / \sim$ is obtained from the Cantor set of geodesic rays in $\tilde{\Sigma}$ based at \tilde{q} by identifying boundary equivalent rays. We will first define $F : \mathcal{G}(q) \rightarrow S_\infty^2$, then observe that $F(\gamma') = F(\delta')$ if $\gamma' \sim \delta'$, so F descends to a map on $S_\infty^1(q)$.

Let $M = M^\tau$ be a model for N , as in §3.3. By Minsky's result (Theorem 3.1), there is a $\pi_1(\Sigma, q)$ -equivariant quasi-isometry

$$\phi : \widetilde{M} \rightarrow \widetilde{N}_{\text{thick}} = \mathbb{H}^3 - \bigcup \mathcal{H},$$

where \mathcal{H} is a locally finite set of disjoint cylinders and horoballs covering the thin part of N .

Choose basepoints $\tilde{q} \in \widetilde{\Sigma}$ and $\tilde{q}_0 \in \widetilde{\Sigma}_0 \subset \widetilde{M}$ such that $\tilde{\pi}(\tilde{q}_0) = \tilde{q}$ and \tilde{q} projects to q . Then we can identify $\mathcal{G}(q)$ with the set of geodesic rays

$$\gamma' : ([0, \infty), 0) \rightarrow (\widetilde{\Sigma}, \tilde{q}).$$

Let $\gamma : [0, \infty) \rightarrow \widetilde{M}$ be the unique model geodesic, based at \tilde{q}_0 , lifting γ' . By Theorem 6.1 (the main result of the preceding sections), γ is an ambient K -quasigeodesic. Since ϕ is a quasi-isometry,

$$\phi \circ \gamma : [0, \infty) \rightarrow \mathbb{H}^3 - \bigcup \mathcal{H}$$

is an ambient K' -quasigeodesic, for a universal constant K' . As shown in §8, $\phi \circ \gamma$ converges to a unique point on the sphere at infinity, and we define

$$F(\gamma') = \lim_{s \rightarrow \infty} \phi \circ \gamma(s).$$

Cusps. It is useful to observe that each component $B \subset \partial\widetilde{\Sigma}$ determines a horoball $H(B) \in \mathcal{H}$, satisfying

$$\partial H(B) = \phi(\tilde{\pi}^{-1}(B)).$$

Similarly, B determines a *cuspidal* $c(B) = \overline{H(B)} \cap S_\infty^2$. The boundary component B , the horoball $H(B)$ and the cusp $c(B)$ all have the same stabilizer in $\pi_1(\Sigma, q)$, namely a rank-one parabolic subgroup.

Boundary equivalence. Now consider a pair of boundary equivalent geodesics

$$\gamma'(i) = \beta'_0(i) * \alpha'_1(i) * \beta'_1(i) * \cdots * \alpha'_n(i) * \beta'_n(i),$$

$i = 1, 2$. By definition, $\beta'_n(1)$ and $\beta'_n(2)$ are contained in a single component B of $\partial\widetilde{\Sigma}$. It is then easy to see that $F(\gamma') = F(\delta') = c(B)$, since the lifts $\phi \circ \beta_n(1)$ and $\phi \circ \beta_n(2)$ both converge to the fixed-point of the parabolic subgroup stabilizing B .

Thus F descends to a well-defined map $F : S_\infty^1(q) \rightarrow S_\infty^2$.

Proof of Theorem 9.1. We begin by showing F is continuous. Let $[\gamma'_k] \rightarrow [\gamma']$ be a convergent sequence of rays in $S_\infty^1(q)$. Then $\gamma'_k(s) \rightarrow \gamma'(s)$ uniformly on each compact subset of $[0, \infty)$, with suitable parameterizations.

Let

$$\gamma' = \beta'_0 * \alpha'_1 * \beta'_1 * \alpha'_2 * \beta'_2 * \dots$$

be the canonical factorization of γ' , and let

$$\gamma = \beta_0 * \xi_0 * \alpha_1 * \beta_1 * \xi_1 * \alpha_2 * \beta_2 * \xi_2 \dots$$

denote the model geodesic lifting γ' . We define a sequence of components $\langle B_0, B_1, B_2, \dots \rangle$ of $\partial\tilde{\Sigma}$, visited by γ' and tending to infinity, as follows. It is convenient to distinguish 4 cases:

- (a) The factorization is infinite.
- (b) The factorization terminates with a geodesic ray α'_{n+1} of slope $\sigma \neq \epsilon_\pm$.
- (c) The factorization terminates with a boundary ray β'_n .
- (d) The factorization terminates with a geodesic ray α'_{n+1} of slope $\sigma = \epsilon_+$ or $\sigma = \epsilon_-$.

In all cases we let B_i denote the component containing β'_i . In case (a) this defines an infinite sequence of distinct boundary components. In case (b), we choose the components B_{n+1}, B_{n+2}, \dots so they are at distance $O(1)$ from α'_{n+1} in the metric $g(t) \in \tau(\text{thick})$, where t is the height in M of the initial point of the lifted ray α_{n+1} . In cases (c) and (d) we set $B_i = B_n$ for $i > n$.

We claim that the cusps corresponding to $\langle B_i \rangle$ satisfy

$$c(B_i) \rightarrow F(\gamma') = \lim \phi \circ \gamma(s).$$

Indeed, by the construction of the model geodesic, we have

$$d(H(B_i), \phi \circ \gamma[0, \infty)) = O(1)$$

for all i ; in other words, γ visits the sequence of horoballs $\langle H(B_i) \rangle$. In cases (a) and (b) the horoballs $H(B_i)$ are all distinct, so they must move along $\gamma[0, \infty)$ and converge to its endpoint; while in cases (c) and (d), $\phi \circ \gamma$ converges to the cusp $c(B_n) = \lim c(B_i)$ by construction of γ' .

We can now check continuity in cases (a) and (b). Choose $i \gg 0$ such that the entire horoball $H(B_i)$ is very close to $F(\gamma')$. Since $\gamma'_k \rightarrow \gamma'$ uniformly on

compact sets, γ'_k also visits B_i for all $k \gg 0$. Therefore $d(H(B_i), \phi \circ \gamma_k) = O(1)$ which implies $F(\gamma'_k) = \lim \phi \circ \gamma_k(s)$ is close to $c(B_i)$. Therefore

$$F(\gamma'_k) \rightarrow F(\gamma')$$

as $k \rightarrow \infty$, and F is continuous at γ' .

To prove continuity in case (c), let

$$\gamma'_k = \beta'_0(k) * \alpha'_1(k) * \beta'_1(k) * \alpha'_2(k) * \dots$$

be the canonical factorization of γ'_k . Then for all $k \gg 0$, the factorization of γ'_k agrees with that of γ' up to the term β'_n . Moreover, the winding number of the boundary path $\beta'_n(k)$ tends to infinity as $k \rightarrow \infty$, since it converges to the boundary ray β'_n . It follows easily that $F(\gamma'_k)$ converges to the corresponding cusp $c(B_n) = F(\gamma')$ as $k \rightarrow \infty$.

Finally in case (d), the factorizations of γ'_k and γ' agree up to the term α'_{n+1} . Moreover, $\alpha'_{n+1}(k)$ has slope σ_k converging to one of the ending laminations $\sigma = \epsilon_{\pm}$. Therefore the lifts $\alpha_{n+1}(k)$ in M reside at heights t_k with $|t_k| \rightarrow \infty$. It follows that γ_k contains a factor of the form $\beta_n * \xi_n(k)$ where the length of the vertical path $\xi_n(k)$ tends to infinity. Therefore its endpoint $F(\gamma'_k)$ lies near the cusp $c(B_n)$, showing

$$F(\gamma'_k) \rightarrow c(B_n) = F(\gamma')$$

in this case as well. This completes the proof of continuity of F .

To see that F is equivariant, just note that for any $g \in \pi_1(\Sigma, q)$, the ray $g \cdot \gamma'$ visits the boundary components $\langle g \cdot B_i \rangle$ and therefore

$$F(g \cdot \gamma') = \lim c(g \cdot B_i) = \lim g \cdot c(B_i) = g \cdot F(\gamma').$$

To check that F is surjective, note that fixed-points of elements of $\pi_1(\Sigma) - \{\text{id}\}$ are dense in S^1_{∞} , so the same is true of $F(S^1_{\infty})$; therefore $F(S^1_{\infty}) = \Lambda = S^2_{\infty}$. ■

10 The case $\Lambda \neq S^2_{\infty}$

In this section we briefly indicate how to modify the proof of Theorem 9.1 to handle the case where the limit set is not the whole sphere. Taking into account these modifications, we obtain a continuous, equivariant map

$$F : S^1_{\infty} \rightarrow \Lambda$$

for all hyperbolic manifolds $N \in H(\Sigma)$ (Theorem 1.1); in particular, we find that the limit sets of totally degenerate groups modeled on a cusped torus are locally connected.

It is convenient to discuss the various cases in terms of the *end invariants*

$$\epsilon_{\pm}(N) \in \mathbb{H} \cup \mathbb{R} \cup \{\infty\}.$$

Minsky has shown that N is determined up to isometry by these invariants [Min2], and that all invariants occur, subject to the condition $\epsilon_+ \neq \epsilon_-$ when both lie in $\mathbb{R} \cup \{\infty\}$. Thus there are four possibilities for the end invariants of N .

First, it may be that both invariants ϵ_{\pm} lie in $\mathbb{H} \cup \mathbb{Q} \cup \{\infty\}$. Then N is geometrically finite, and the existence of $F : S_{\infty}^1 \rightarrow \Lambda$ is given by Floyd's theorem [F1].

Second, there is the case $\epsilon_{\pm} \in \mathbb{R} - \mathbb{Q}$. This is the case where the limit set is the whole sphere; it is handled by Theorem 9.1.

Third, there is the case where one end invariant of N , say ϵ_- , belongs to \mathbb{H} , and ϵ_+ belongs to $\mathbb{R} - \mathbb{Q}$. Then one end of N limits on a Riemann surface at infinity, X , while the other end is geometrically infinite; and N corresponds to a totally degenerate group on the boundary of the Bers slice B_X in the space of quasifuchsian groups. The proof of Theorem 9.1 can be adapted to handle this case as follows.

1. We can assume the invariant ϵ_- lies in the thick part of $\text{Teich}(T) \cong \mathbb{H}$. Indeed, the quasi-isometry type of N is the same for all $\epsilon_- \in \mathbb{H}$.
2. The Teichmüller geodesic τ of §3 is then replaced by a ray, starting at $\epsilon_- \in \mathbb{H}$ and tending to $\epsilon_+ \in \partial\mathbb{H}$.
3. The ray τ is modified as before to obtain a path η that skirts the thin part of Teichmüller space. The new path is parameterized by a map $g : [0, \infty) \rightarrow \eta$, determining a family of metrics $g(t)$ on the torus T .
4. The model manifold M now takes the form $M = (T \times [0, \infty) - P) - \bigcup_{\mu} C(\mu)$. Minsky's work implies that M is a model for the thick part of the convex core K of N , with the lowest level M_0 corresponding to ∂K .
5. Given a geodesic $\sigma \in \mathcal{S}(T)$, the length function $L_{\sigma}(g)$ still achieves its minimum at a unique *optimal metric* $g_{\sigma} \in \tau$, as in §2. But now the optimal metric may sometimes coincide with the initial point ϵ_- of the

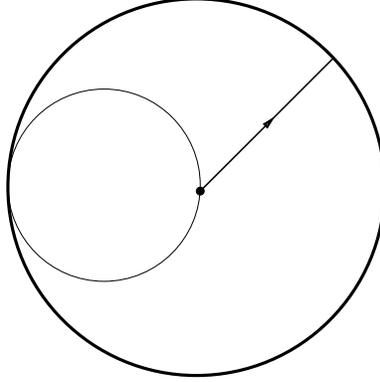


Figure 15. The optimal metric is the end of the ray: $g_\sigma = \epsilon_-$.

ray τ (see Figure 15). Because of this possibility, Lemma 2.1 needs to be modified: instead of the sharp statement

$$L_\sigma(g)/L_\sigma(g_\sigma) = \cosh^{1/2} d(g, g_\sigma)$$

for any $g \in \tau$, we have the qualitative estimate

$$L_\sigma(g)/L_\sigma(g_\sigma) \asymp \exp\left(\frac{1}{2} \cdot d(g, g_\sigma)\right).$$

In any case, $L_\sigma(g)$ still increases exponentially fast as g moves away from g_σ , and this fact all that is needed in the course of the proof.

6. Note that for any short geodesic $\mu \in \mathcal{S}(\tau)$, the optimal metric g_μ lies in the interior of τ . (The metric g_μ cannot coincide with the endpoint ϵ_- of τ because ϵ_- lies in the thick part of Teichmüller space.) Thus the proof of efficiency in the thin case (§7) proceeds as before, and we conclude that model geodesics give ambient quasigeodesics in \widetilde{M} (Theorem 6.1).
7. Since M is a model for the thick part of the convex core $K \subset N$ (instead of N itself), in §9 the model geodesics now give rise the ambient quasigeodesics

$$\phi \circ \gamma : [0, \infty) \rightarrow \widetilde{K} - \bigcup \mathcal{H}.$$

But $\widetilde{K} \subset \mathbb{H}^3$ is convex, so $\phi \circ \gamma$ is also an ambient quasigeodesic in $\mathbb{H}^3 - \bigcup \mathcal{H}$. Thus the shadowing result of §8 can be applied without

change to define $F : S_\infty^1 \rightarrow \Lambda$ and establish its continuity, completing the proof.

In the fourth and final case, one end invariant, say $\epsilon_-(N)$, belongs to $\mathbb{Q} \cup \{\infty\}$, while the other belongs to $\mathbb{R} - \mathbb{Q}$. This case is like the preceding one, except now the geometrically finite end of N has acquired a rank one cusp (an accidental parabolic). The thick part of the convex core of N is still modeled, as above, on a Teichmüller ray $\tau = [\epsilon_0, \epsilon_+)$, namely the subsegment of (ϵ_-, ϵ_+) lying outside the horoball in $\text{Teich}(T)$ where the simple curve μ of slope ϵ_- is short. To handle this case, the proof can be modified as above, taking into account (in §9) the additional components of \mathcal{H} coming from the accidental parabolic.

References

- [ADP] R. C. Alperin, W. Dicks, and J. Porti. The boundary of the Gieseking tree in hyperbolic three-space. *Topology Appl.* **93**(1999), 219–259.
- [AM] J. Anderson and B. Maskit. On the local connectivity of limit sets of Kleinian groups. *Complex Variables Theory Appl.* **31**(1996), 177–183.
- [BP] R. Benedetti and C. Petronio. *Lectures on Hyperbolic Geometry*. Springer-Verlag, 1992.
- [Bon] F. Bonahon. Bouts des variétés hyperboliques de dimension 3. *Annals of Math.* **124**(1986), 71–158.
- [CT] J. W. Cannon and W. P. Thurston. Group invariant Peano curves. *Preprint, Princeton, 1985*.
- [ECH] D. Epstein, J. Cannon, D. Holt, S. Levy, M. Paterson, and W. Thurston. *Word Processing in Groups*. Jones and Bartlett, 1992.
- [Fl] W. J. Floyd. Group completions and limit sets of Kleinian groups. *Invent. math.* **57**(1980), 205–218.
- [Hub] J. H. Hubbard. Local connectivity of Julia sets and bifurcation loci: three theorems of J.-C. Yoccoz. In L. R. Goldberg and A. V. Phillips, editors, *Topological Methods in Modern Mathematics*, pages 467–511. Publish or Perish, Inc., 1993.

- [Ji] Y. Jiang. *Renormalization and Geometry in One-Dimensional and Complex Dynamics*. World Scientific, 1996.
- [Kl] E. Klarreich. Semiconjugacies between Kleinian group actions on the Riemann sphere. *Amer. J. Math.* **121**(1999), 1031–1078.
- [La] U. Lang. Quasigeodesics outside horoballs. *Geom. Dedicata* **63**(1996), 205–215.
- [Lyu] M. Lyubich. Geometry of quadratic polynomials: Moduli, rigidity, and local connectivity. *Stony Brook IMS Preprint 1993/9*.
- [Mil] J. Milnor. Local connectivity of Julia sets: expository lectures. In Tan Lei, editor, *The Mandelbrot Set, Theme and Variations*, pages 67–116. Cambridge University Press, 2000.
- [Min1] Y. Minsky. On rigidity, limit sets, and end invariants of hyperbolic 3-manifolds. *J. Amer. Math. Soc* **7**(1994), 539–588.
- [Min2] Y. Minsky. The classification of punctured torus groups. *Annals of Math.* **149**(1999), 559–626.
- [Pet1] C. L. Petersen. Local connectivity of some Julia sets containing a circle with an irrational rotation. *Acta Math.* **177**(1996), 163–224.
- [Pet2] C. L. Petersen. The Herman–Świątek theorems with applications. In Tan Lei, editor, *The Mandelbrot Set, Theme and Variations*, pages 211–226. Cambridge University Press, 2000.
- [Sor] D. E. K. Sørensen. Accumulation theorems for quadratic polynomials. *Ergod. Th. & Dynam. Sys.* **16**(1996), 555–590.
- [Sul] D. Sullivan. Conformal dynamical systems. In *Geometric Dynamics*, volume 1007 of *Lecture Notes in Mathematics*, pages 725–752. Springer-Verlag, 1983.
- [TY] Tan L. and Yin Y. Local connectivity of the Julia set for geometrically finite rational maps. *Science in China (Series A)* **39**(1996), 39–47.
- [Th] W. P. Thurston. *Geometry and Topology of Three-Manifolds*. Lecture Notes, Princeton University, 1979.

MATHEMATICS DEPARTMENT
HARVARD UNIVERSITY
1 OXFORD ST
CAMBRIDGE, MA 02138-2901