ISSN 1364-0380 (on line) 1465-3060 (printed)

Geometry & Topology
Volume 7 (2003) 33–90
Published: 1 February 2003

Stable Teichmüller quasigeodesics and ending laminations

Lee Mosher

Deptartment of Mathematics and Computer Science
Rutgers University, Newark, NJ 07102

Email: `mosher@andromeda.rutgers.edu`

Abstract

We characterize which cobounded quasigeodesics in the Teichmüller space $\mathcal{T}$ of a closed surface are at bounded distance from a geodesic. More generally, given a cobounded lipschitz path γ in $\mathcal{T}$, we show that γ is a quasigeodesic with finite Hausdorff distance from some geodesic if and only if the canonical hyperbolic plane bundle over γ is a hyperbolic metric space. As an application, for complete hyperbolic 3–manifolds N with finitely generated, freely indecomposable fundamental group and with bounded geometry, we give a new construction of model geometries for the geometrically infinite ends of N, a key step in Minsky's proof of Thurston's ending lamination conjecture for such manifolds.

AMS Classification numbers Primary: 57M50

Secondary: 32G15

Keywords: Teichmüller space, hyperbolic space, quasigeodesics, ending laminations

Proposed: Walter Neumann
Seconded: Benson Farb, David Gabai

Received: 15 November 2001
Revised: 6 January 2003

1 Introduction

1.1 Stable Teichmüller quasigeodesics

If X is a geodesic metric space which is hyperbolic in the sense of Gromov, then quasigeodesics in X are "stable": each quasigeodesic line, ray, or segment in X has finite Hausdorff distance from a geodesic, with Hausdorff distance bounded solely in terms of quasigeodesic constants and the hyperbolicity constant of X. In the case of hyperbolic 2–space, stability of quasigeodesics goes back to Morse [21].

Consider a closed, oriented surface S of genus ≥ 2. Its Teichmüller space $\mathcal{T}$ is often studied by finding analogies with hyperbolic metric spaces. For example, the mapping class group $\mathcal{MCG}$ of S acts on $\mathcal{T}$ properly discontinuously by isometries, with quotient orbifold $\mathcal{M} = \mathcal{T}/\mathcal{MCG}$ known as Riemann's moduli space, and $\mathcal{M}$ is often viewed as the analogue of a finite volume, cusped hyperbolic orbifold. Although this analogy is limited—$\mathcal{T}$ is *not* a hyperbolic metric space in the Teichmüller metric [22], and indeed there is no $\mathcal{MCG}$ equivariant hyperbolic metric on $\mathcal{T}$ with "finite volume" quotient [3]—nonetheless the analogy has recently been strengthened and put to use in applications [17], [19], [10].

Minsky's projection theorem [17] gives a version in $\mathcal{T}$ of stability of quasigeodesic segments, assuming coboundedness of the corresponding geodesic. Given a λ, η quasigeodesic segment γ in $\mathcal{T}$, let g be the geodesic with the same endpoints as γ. If g is *cobounded*, meaning that each point of g represents a hyperbolic structure with injectivity radius bounded away from zero, then according to Minsky's theorem the Hausdorff distance between γ and g is bounded by a constant depending only on λ, η, and ϵ. On the other hand, Masur and Minsky [20] produced a cobounded quasigeodesic line—the orbit of a partially pseudo-Anosov cyclic subgroup—for which there does not exist any geodesic line at finite Hausdorff distance.

In joint work with Benson Farb [10] we formulated a theory of convex cocompact groups of isometries of $\mathcal{T}$, pursuing an analogy with hyperbolic spaces. In the course of this work we needed to consider the following:

Problem Give conditions on a cobounded, lipschitz path γ in $\mathcal{T}$, in terms of the geometry of the canonical hyperbolic plane bundle over γ, which imply that γ is a quasigeodesic with finite Hausdorff distance from a geodesic.

Recall that for any path $\gamma\colon I \to \mathcal{T}$, $I \subset \mathbf{R}$ closed and connected, there is a canonical bundle $\mathcal{H}_\gamma \to I$ of hyperbolic planes over γ on which $\pi_1 S$ acts, so that for each $t \in I$ the marked hyperbolic surface $\mathcal{S}_t = \mathcal{H}_t/\pi_1 S$ represents the point $\gamma(t) \in \mathcal{T}$. After perturbation of γ, the fiberwise hyperbolic metric on $\mathcal{H}_\gamma$ extends to a $\pi_1 S$–equivariant, piecewise Riemannian metric on $\mathcal{H}_\gamma$, whose large scale geometry is independent of the perturbation; see [10], proposition 4.2, recounted below as proposition 2.3. Thus the above problem asks for an invariant of the large-scale geometry of $\mathcal{H}_\gamma$ that characterizes when γ is a stable quasigeodesic.

The first goal of this paper is a solution of the above problem:

Theorem 1.1 (Hyperbolicity implies stability) *If $\gamma\colon I \to \mathcal{T}$ is a cobounded, lipschitz map defined on a closed, connected subset $I \subset \mathbf{R}$, and if $\mathcal{H}_\gamma$ is a hyperbolic metric space, then γ is a quasigeodesic and $\gamma(I)$ has finite Hausdorff distance from some geodesic g in $\mathcal{T}$ sharing the same endpoints as γ. To be precise, for every bounded subset $\mathcal{B} \subset \mathcal{M}$, and every $\rho \ge 1$, $\delta \ge 0$ there exists $\lambda \ge 1$, $\eta \ge 0$, $A \ge 0$ such that the following holds: if $\gamma\colon I \to \mathcal{T}$ is $\mathcal{B}$–cobounded and ρ–lipschitz, and if $\mathcal{H}_\gamma$ is δ–hyperbolic, then γ is a λ,η quasigeodesic, and there exists a geodesic $g\colon I \to \mathcal{T}$, with $\gamma(t) = g(t)$ for all $t \in \partial I$, such that the Hausdorff distance between $\gamma(I)$ and $g(I)$ is $\le A$.*

Remark This theorem, and its application to Minsky's results on the ending lamination conjecture, have been discovered and proved independently by Brian Bowditch [5], using different methods.

Remark In the case of a compact segment $I = [a,b]$, the first sentence of the theorem has no content: $\mathcal{H}_\gamma$ is quasi-isometric to the hyperbolic plane and so is hyperbolic, and γ is Hausdorff equivalent to the Teichmüller geodesic $\overline{\gamma(a)\gamma(b)}$. The quantifiers in the second sentence are therefore necessary in order to say anything of substance when γ is a segment. Even in the case of a line or ray, where the first sentence actually has content, the second sentence gives a lot more information.

Remark While γ is only assumed lipschitz (coarse Lipschitz would do), the conclusion shows that γ is a quasigeodesic. Note that for any lipschitz path γ and geodesic g, to say that γ is a quasigeodesic at finite Hausdorff distance from g is equivalent to saying that γ and g are asynchronous fellow travellers; see section 2.1.

Recent work of Minsky and Rafi provides a strong converse to Theorem 1.1 in the case of a bi-infinite path, and putting these together we get the following result:

Corollary 1.2 *Given a cobounded, lipschitz path $\gamma\colon \mathbf{R} \to \mathcal{T}$, the following are equivalent:*

1. *γ is a quasigeodesic and there is a geodesic g at finite Hausdorff distance from γ.*
2. *$\mathcal{H}_\gamma$ is a hyperbolic metric space.*
3. *$\mathcal{H}_\gamma$ is quasi-isometric to $\mathbf{H}^3$.*

Proof (2) $\implies$ (1) is Theorem 1.1. (1) $\implies$ (3) follows from the Minsky–Rafi theorem [18] which says that $\mathcal{H}_g^{\text{SOLV}}$ is quasi-isometric to $\mathbf{H}^3$, together with the fact that $\mathcal{H}_\gamma$ is quasi-isometric to $\mathcal{H}_g^{\text{SOLV}}$ (see Proposition 2.3). (3) $\implies$ (2) is immediate. □

Application: hyperbolic surface group extensions Theorem 1.1 is applied in [10] as follows. Consider a short exact sequence of finitely generated groups of the form $1 \to \pi_1(S) \to \Gamma_G \to G \to 1$, determined by a group homomorphism $f\colon G \to \mathcal{MCG}(S) \approx \text{Out}(\pi_1(S))$. One of the main results in [10] says that if Γ_G is a word hyperbolic group, then G maps with finite kernel onto a subgroup of $\mathcal{MCG}(S)$ acting quasiconvexly on Teichmüller space; Theorem 1.1 plays a key role in the proof.

We remark that when the group G is free then the converse is also proved in [10]: a free, convex cocompact subgroup $G < \text{Isom}(\mathcal{T})$ has a word hyperbolic extension group $\Gamma_G = \pi_1(S) \rtimes G$.

1.2 Ending laminations

The second goal of this paper is to apply Theorem 1.1 to give a new construction of model manifolds for geometrically infinite ends, used by Minsky in proving a special case of Thurston's ending lamination conjecture [16].

Let N be a complete hyperbolic 3–manifold satisfying the following properties:

1. The fundamental group $\pi_1 N$ is finitely generated, freely indecomposable.
2. The action of $\pi_1 N$ on $\mathbf{H}^3$ has no parabolic elements.

Work of Thurston [25] and Bonahon [4] describes the topology and geometry of N nears its ends. The manifold N is the interior of a compact 3–manifold with boundary denoted $\overline{N}$, and so there is a one-to-one correspondence between ends e of N and components S_e of $\partial\overline{N}$. Associated to each end e of N there is an *end invariant*, a geometric/topological structure on the associated surface S_e, which describes the behavior of N in the end e. The end invariant comes in two flavors. In one case, the end e is *geometrically finite* and the end invariant is a point in the Teichmüller space of S_e. In the case where e is geometrically infinite, it follows that e is *simply degenerate* and the end invariant is the *ending lamination*, an element of the geodesic lamination space of S_e.

Thurston's ending lamination conjecture says that N is determined up to isometry by its topological type and its end invariants: if N, N' are complete hyperbolic 3–manifolds as above, and if $f\colon N \to N'$ is a homeomorphism respecting end invariants, then f is properly homotopic to an isometry. Minsky's theorem says that this is true if N, N' each have *bounded geometry*, meaning that injectivity radii in N and in N' are both bounded away from zero.

The heart of Minsky's result is theorem 5.1 of [16], the construction of a *model manifold* for N, a proper geodesic metric space M equipped with a homeomorphism $f\colon N \to M$, such that the metric on M depends only on the end invariants of N, and such that f is properly homotopic to a map for which any lift $\widetilde{N} \to \widetilde{M}$ is a quasi-isometry. When N, N' as above have the same end invariants, it follows that they have isometric model manifolds, and so any homeomorphism $N \to N'$ respecting end invariants is properly homotopic to a map each of whose lifts $\widetilde{N} = \mathbf{H}^3 \to \mathbf{H}^3 = \widetilde{N}'$ is an equivariant quasi-isometry. An application of Sullivan's rigidity theorem [24] shows immediately that this quasi-isometry is a bounded distance from an equivariant isometry, proving the ending lamination conjecture in the presence of bounded geometry.

Minsky's construction of the model manifold for N is broken into three pieces: a standard construction in the compact core of N; a construction in the geometrically finite ends using results of Epstein and Marden [9]; and a construction in the geometrically infinite ends using Minsky's earlier results [15]. The geometrically infinite construction is described in two results. In the singly degenerate case, Section 5 of [16] describes a model manifold for a neighborhood of the end. In the doubly degenerate case, Corollary 5.10 of [16], recalled below in Theorem 1.3, describes a model manifold for all of N. It is these two results which we will prove anew, by applying Theorem 1.1.

Given a closed surface S, let $\mathcal{T}(S)$ be the Teichmüller space, $\mathcal{PMF}(S)$ the projective measured foliation space, and $\overline{\mathcal{T}}(S) = \mathcal{T}(S) \cup \mathcal{PMF}(S)$ Thurston's compactification. As mentioned above, associated to each geodesic segment, ray, or line $g\colon I \to \mathcal{T}(S)$, where I is a closed connected subset of $\mathbf{R}$, there is a singular SOLV metric on $S \times I$ denoted $\mathcal{S}_g^{\text{SOLV}}$ with universal cover $\mathcal{H}_g^{\text{SOLV}}$, such that the induced conformal structure on $S \times t$ represents the point $g(t) \in \mathcal{T}(S)$.

Recall that a *pleated surface* in the hyperbolic 3–manifold N is a map denoted $f\colon (S,\sigma) \to N$ where S is a closed surface, σ is a hyperbolic structure on S, the map f takes rectifiable paths to rectifiable paths of the same length, and there is a geodesic lamination λ in the hyperbolic structure σ, called the *pleating locus* of f, such that f is totally geodesic on each leaf of λ and on each complementary component of λ. Given a π_1–injective map of a closed surface $S \to N$, let $\Sigma(S \to N)$ be the set of points $\sigma \in \mathcal{T}(S)$ represented by pleated surfaces $(S,\sigma) \to N$ in the homotopy class of the map $S \to N$.

Consider now a hyperbolic 3–manifold N as above. An end e of N has neighborhood $N_e \approx S_e \times [0,\infty)$. By Bonahon's theorem, the end e is *geometrically infinite* if and only if each neighhborhood of e contains the image of a pleated surface in N homotopic to the inclusion $S_e \hookrightarrow N$. If $\widetilde N$ is the covering space of e corresponding to the injection $\pi_1(S_e) \hookrightarrow \pi_1(N)$, we obtain a homeomorphism $\widetilde N \approx S_e \times (-\infty,+\infty)$, with one end corresponding to N_e. Assuming e is geometrically infinite, we say that the inclusion $\pi_1(S_e) \hookrightarrow \pi_1(N)$ is *singly degenerate* if $\widetilde N$ has one geometrically infinite end; otherwise $\widetilde N$ has two geometrically infinite ends, in which case we say that $\pi_1(S_e) \hookrightarrow \pi_1(N)$ is *doubly degenerate*. In the doubly degenerate case, either the covering map $\widetilde N \to N$ is degree 1 and so $N \approx S \times (-\infty,+\infty)$, or $\widetilde N \to N$ has degree 2 and there is an orbifold fibration of N, with generic fiber S_e, whose base space is the *ray orbifold* $[0,\infty)$ with a $\mathbf{Z}/2$ mirror group at the point 0; we'll review these facts in section 4.2.

Theorem 1.3 (Doubly degenerate model manifold) *Let N be a bounded geometry hyperbolic 3–manifold satisfying (1) and (2) above. Let e be an end, and suppose $\pi_1(S_e) \hookrightarrow \pi_1(N)$ is doubly degenerate. There exists a unique cobounded geodesic line g in $\mathcal{T}(S_e)$ such that:*

(1) *$\Sigma(S_e \to N)$ is Hausdorff equivalent to g in $\mathcal{T}(S_e)$.*

Moreover:

(2) *The homeomorphism $S_e \times (-\infty,+\infty) \approx \widetilde N$ is properly homotopic to a map which lifts to a quasi-isometry $\mathcal{H}_g^{\text{SOLV}} \to \mathbf{H}^3$.*

(3) *The ideal endpoints of g in $\mathcal{PMF}(S_e)$ are the respective ending laminations of the two ends of $\widetilde N$.*

In the degree 2 case, the order 2 covering transformation group on $\widetilde{N}$ acts isometrically on $\mathcal{S}_g^{\text{SOLV}}$ so as to commute with the homeomorphism $\mathcal{S}_g^{\text{SOLV}} \approx \widetilde{N}$.

In the degree 2 case, it follows from the theorem that the degree 2 covering map $\mathcal{S}_g^{\text{SOLV}} \to N$ induces a singular SOLV metric on N.

In the doubly degenerate case we have constructed the entire model manifold, and therefore by applying Sullivan's rigidity theorem [24] we obtain a complete proof of the ending lamination conjecture for the case $N \approx S \times (-\infty, +\infty)$ with two geometrically infinite ends and with bounded geometry.

In the singly degenerate case, one still needs the arguments of Minsky to construct the model manifold out of its pieces, before applying Sullivan's rigidity theorem. Here is the theorem describing the piece of the model manifold corresponding to a singly degenerate end.

Theorem 1.4 (Singly degenerate model manifold) *Let N be a hyperbolic 3–manifold of bounded geometry satisfying (1) and (2) above, and let e be a geometrically infinite end. There exists a cobounded geodesic ray g in $\mathcal{T}(S_e)$, unique up to choice of its finite endpoint, such that:*

(1) *$\Sigma(S_e \to N)$ and g are Hausdorff equivalent in $\mathcal{T}(S_e)$.*

Moreover, g satisfies the following:

(2) *The homeomorphism $S_g^{\text{SOLV}} \approx S_e \times [0, \infty) \approx N_e$ is properly homotopic to a map which lifts to a quasi-isometry of universal covers $\mathcal{H}_g^{\text{SOLV}} \to \widetilde{N}_e$.*

(3) *The ideal endpoint of g in $\mathcal{PMF}(S_e)$ is the ending lamination of e.*

Minsky's proofs of these results in [16] depend on pleated surface arguments found in section 4 of that paper, as well as results found in [15] concerning projection arguments in Teichmüller space and harmonic maps from Riemann surfaces into hyperbolic 3–manifolds. In particular, the statements (1) in these two theorems are parallel to harmonic versions stated in Theorem A of [15].

Our proofs are found in section 4; here is a sketch in the doubly degenerate case $N \approx S \times (-\infty, +\infty)$. The first part uses pleated surface arguments, several borrowed from [16] section 4. There is a sequence of pleated surfaces $f_i\colon (S, \sigma_i) \to N$, $(i = -\infty, \ldots, +\infty)$, whose images progress in a controlled fashion from $-\infty$ to $+\infty$ in N. This allows one to construct a cobounded, lipschitz path $\gamma\colon \mathbf{R} \to \mathcal{T}(S)$ with $\gamma(i) = \sigma_i$, and to construct a map $\mathcal{S}_\gamma \to N$ in the correct homotopy class each of whose lifts to the universal covers $\mathcal{H}_\gamma \to \mathbf{H}^3$ is a quasi-isometry.

For the next part of the proof, since $\mathbf{H}^3$ is Gromov hyperbolic, so is $\mathcal{H}_\gamma$, and so Theorem 1.1 applies to the path γ. The output is a cobounded Teichmüller geodesic g such that γ and g are asynchronous fellow travellers. It is then straightforward to show that g satisfies the desired conclusions of Theorem 1.3.

The singly degenerate case is similar, requiring an additional argument to show that N_e itself has Gromov hyperbolic universal cover, so that Theorem 1.1 can apply to produce the desired Teichmüller ray.

Finally, suppose that N satisfies (1) and (2) but we do not assume that N has bounded geometry. Suppose however that some geometrically infinite end e of N does have bounded geometry, in the sense that for some $\epsilon > 0$ each point of N_e has injectivity radius $\geq \epsilon$. In this case we are still able to construct a model geometry for N_e: there is a geodesic ray $g\colon [0,\infty) \to \mathcal{T}$ and a map $\mathcal{S}_g \to N_e$ in the correct homotopy class which lifts to a quasi-isometry $\mathcal{H}_g \to \widetilde{N}_e$, and the other conclusions of Theorem 1.4 hold as well. This case is not covered by Minsky's results in [16], although probably with some work the original construction could be pushed through. Our construction in this new case is given in Section 4.4. Thanks to Jeff Brock for asking this question.

1.3 An outline of the proof of Theorem 1.1

We shall outline the proof for a path γ whose domain is the whole real line, $\gamma\colon \mathbf{R} \to \mathcal{T}$. As remarked earlier, $\pi_1 S$ acts on each fiber $\mathcal{H}_t$ of the hyperbolic plane bundle $\mathcal{H}_\gamma \to \mathbf{R}$, with quotient a marked hyperbolic surface $\mathcal{S}_t = \mathcal{H}_t/\pi_1 S$. These surfaces fit together to form the canonical marked hyperbolic surface bundle $\mathcal{S}_\gamma \to \mathbf{R}$ over γ. Note that $\mathcal{H}_\gamma$ is the universal cover of $\mathcal{S}_\gamma$, with deck transformation group $\pi_1 S$.

The proof uses the Bestvina–Feighn flaring condition for $\mathcal{H}_\gamma$, a necessary and sufficient condition for Gromov hyperbolicity of $\mathcal{H}_\gamma$. The key idea of the proof is to use hyperbolicity of $\mathcal{H}_\gamma$, via the Bestvina–Feighn flaring condition, to construct ending laminations for $\mathcal{S}_\gamma$, one lamination λ^- for the negative end and another λ^+ for the positive end. These two laminations can then be used to construct the desired Teichmüller geodesic g.

The flaring condition is concerned with quasihorizontal paths in $\mathcal{H}_\gamma$, that is, sections $\ell\colon \mathbf{R} \to \mathcal{H}_\gamma$ of the projection map $\mathcal{H}_\gamma \to \mathbf{R}$ such that ℓ satisfies a coarse Lipschitz condition. The flaring condition says, roughly, that for any two quasihorizontal paths $\ell, \ell'\colon \mathbf{R} \to \mathcal{H}_\gamma$, the sequence of distances $d_i = d_{\mathcal{H}_i}(\ell(i), \ell'(i))$

satisfies an exponential growth property in either forward or backward time, possibly both.

The central concept in the proof is an exponential growth property for measured geodesic laminations. Each element λ in $\mathcal{MF}$, the space of measured foliations on S, can be represented as a measured geodesic lamination λ_s in each hyperbolic surface $\mathcal{S}_s$, $s \in \mathbf{R}$. Although not strictly necessary for the proof, it is nevertheless convenient to arrange that as s varies in $\mathbf{R}$ the lamination λ_s varies nicely in the bundle $\mathcal{S}$, filling out a 2–dimensional lamination on $\mathcal{S}$; the proof of this fact relies on basic tools of partially hyperbolic dynamics to show that the geodesic flows on the surfaces $\mathcal{S}_s$ can be packaged together in a nice manner. The Thurston length of λ_s is obtained by integrating the transverse measure on λ_s against the Lebesgue measure along leaves of λ_s, yielding a number $\ell_s > 0$. We are interested in the growth properties of the sequence ℓ_i, $i \in \mathbf{Z}$. By applying the Bestvina–Feighn flaring property for $\mathcal{H}_\gamma$, together with an argument using Fubini's theorem, we show that ℓ_i satisfies an exponential growth property in either forward or backward time, possibly both. As a consequence, each element $\lambda \in \mathcal{MF}$ satisfies the following trichotomy:

λ **is realized at** $-\infty$**,** meaning that ℓ_i goes exponentially to zero as $i \to -\infty$ and exponentially to infinity as $i \to +\infty$; *or*

λ **is realized at** $+\infty$**,** meaning that ℓ_i goes exponentially to zero as $i \to +\infty$ and exponentially to infinity as $i \to -\infty$; *or*

λ **is finitely realized,** meaning that ℓ_i goes exponentially to infinity as $i \to -\infty$ or $+\infty$, and ℓ_i is minimized on a subinterval of $\mathbf{R}$ of uniformly bounded length.

We also prove that the position at which λ is realized, either at $-\infty$, or at $+\infty$, or on a certain subinterval of $\mathbf{R}$ of uniformly bounded length, is a coarsely continuous function of $\lambda \in \mathcal{MF}$. By considering the shortest geodesic in each $\mathcal{S}_i$, letting $i \to \pm\infty$, and passing to limits in the projective measured foliation space $\mathcal{PMF}$, we exhibit the existence of geodesic laminations λ^-, λ^+ realized at $-\infty$, $+\infty$ respectively. These are the *ending laminations* of the hyperbolic surface bundle $\mathcal{S}_\gamma$.

The next step is to prove that λ^-, λ^+ each fill the surface S. This uses a trick that I remember learning in Thurston's class at Princeton, sometime from 1979 to 1983. If, say, λ^+ does not fill, then there is a simple closed geodesic c that, together with certain boundary leaves of λ^+, cobounds a "crown surface" (see Figure 1). But then one can play off the exponential decay of λ^+ against the exponential growth of c to get a contradiction.

It is easy to show that λ^-, λ^+ are topologically inequivalent, and it follows that they jointly fill the surface, which is precisely the condition needed to exhibit a Teichmüller geodesic g with horizontal and vertical measured foliations equivalent to λ^-, λ^+.

It remains to show that γ and g are at finite Hausdorff distance. Our first proof of this was somewhat laborious, requiring one to go through the whole construction of g and check various additional properties along the way. But then we discovered a compactness argument that establishes finite Hausdorff distance seemingly by magic. The idea is to fix a compact subset $\mathcal{B} \subset \mathcal{M}$ and constants $\rho, \delta \geq 0$, and to consider the space $\Gamma_{\mathcal{B},\rho,\delta}$ consisting of all triples $(\gamma, \lambda^-, \lambda^+)$ such that $\gamma\colon \mathbf{R} \to \mathcal{T}$ is a $\mathcal{B}$–cobounded, ρ–lipschitz path for which $\mathcal{H}_\gamma$ is δ–hyperbolic, and λ^-, λ^+ are ending laminations for $\mathcal{S}_\gamma$, normalized to have Thurston length 1 in $\mathcal{S}_0$. The mapping class group $\mathcal{MCG}$ acts on $\Gamma_{\mathcal{B},\rho,\delta}$, and we prove that this action is cocompact, using compact open topology for γ and the topology on $\mathcal{MF}$ for λ^-, λ^+. The proof of cocompactness uses the Ascoli–Arzela theorem together with the fact that δ–hyperbolicity is closed in the Gromov–Hausdorff topology on metric spaces. We also prove that if g is the Teichmüller geodesic determined by λ^-, λ^+, then the point $g(0) \in \mathcal{T}$ is a continuous function of the data $(\gamma, \lambda^-, \lambda^+)$. It follows that the distance in Teichmüller space between $\gamma(0)$ and $g(0)$ is a continuous function of the data, and since $\mathcal{MCG}$ acts cocompactly on the data, then $d_{\mathcal{T}}(\gamma(0), g(0))$ is uniformly bounded. Suppressed in this exposition is a delicate parameterization issue, which comes up in proving that γ is a quasigeodesic: we actually prove not just that γ, g are at finite Hausdorff distance, but that γ, g have quasigeodesic parameterizations so that $d_{\mathcal{T}}(\gamma(t), g(t))$ is uniformly bounded.

Some additional work is needed to prove the theorem when γ is parameterized by a ray or a finite interval; in the exposition, these cases are smoothly integrated with the case of a line.

Acknowledgements Thanks to Benson Farb, for a rich collaboration from which many offshoots can grow, and in particular for many shared ideas which contributed to this paper.

Thanks also to the Technion in Haifa, Israel, and to Michah Sageev, for hospitality and support during the June 2000 Conference on Geometric and Combinatorial Group Theory, where this work was carried out and was first presented.

Thanks to Amie Wilkinson for telling me how to prove Lemma 3.2, and to Jeff Brock for suggesting the result in Section 4.4.

For useful suggestions in preparing the revised version, thanks to the referee, and to Hossein Namazi and other members of the Complex Analysis and Geometry Seminar at SUNY Stony Brook.

This work was partially supported by grants from the National Science Foundation.

2 Preliminaries

2.1 Coarse geometry

Consider two metric spaces X, Y and a map $f\colon X \to Y$. The map f is K, C *coarse lipschitz* with $K \geq 1, C \geq 0$ if

$$d_Y(f(x), f(x')) \leq K d_X(x, x') + C \quad \text{for} \quad x, x' \in X.$$

If f is coarse lipschitz, then we say in addition that f is *uniformly proper* with respect to a proper, monotonic function $\rho\colon [0, \infty) \to [0, \infty)$ if

$$d_Y(f(x), f(x')) \geq \rho(d_X(x, x')) \quad \text{for} \quad x, x' \in X.$$

The function ρ is called a *properness gauge* for f. We say that f is a K, C *quasi-isometric embedding* if f is K, C coarse lipschitz and uniformly proper with properness gauge $\rho(d) = \frac{1}{K}d - C$. The map f is C*–coarsely surjective* if for all $y \in Y$ there exists $x \in X$ such that $d_Y(f(x), y) \leq C$. The map f is a K, C *quasi-isometry* if it is a C–coarsely surjective, K, C quasi-isometric embedding.

A metric space X is *geodesic* if for any x, y there is a rectifiable path p from x to y such that $\mathrm{len}(p) = d(x, y)$. X is *proper* if closed balls are compact.

A map $\bar f\colon Y \to X$ is a C*–coarse inverse* for f if $d_{\sup}(\bar f \circ f, \mathrm{Id}_X) \leq C$ and $d_{\sup}(f \circ \bar f, \mathrm{Id}_Y) \leq C$.

Here are some basic facts concerning these concepts:

Lemma 2.1

(1) *A coarse lipschitz map is a quasi-isometry if and only if it has a coarse inverse which is coarse lipschitz.*

(2) *Suppose X, Y are geodesic metric spaces. Any coarsely surjective, uniformly proper map $f\colon X \to Y$ is a quasi-isometry.* □

In each of these facts, the constants implicit in the conclusion of the statement depend only on the constants in the hypothesis.

Quasigeodesics Given a geodesic metric space X, a λ, ϵ *quasigeodesic* in X is a λ, ϵ quasi-isometric embedding $\gamma\colon I \to X$, where I is a closed, connected subset of X. With I is a compact interval we have a *quasigeodesic segment*, when I is a half-line we have a *quasigeodesic ray*, and when $I = \mathbf{R}$ we have a *quasigeodesic line.*

Two paths $\gamma\colon I \to X$, $\gamma'\colon I' \to X$ are *asynchronous fellow travellers* with respect to a K, C quasi-isometry $\phi\colon I \to I'$ if there is a constant A such that $d(\gamma'(\phi(t)), \gamma(t)) \le A$ for $t \in I$.

Recall that the Hausdorff distance between two sets $A, B \subset X$ is the infimum of $r \subset \mathbf{R}_+ \cup +\infty$ such that A is contained in the r–neighborhood of B, and B is contained in the r–neighborhood of A.

Consider two paths $\gamma\colon I \to X$, $\gamma'\colon I' \to X$ such that γ is a quasigeodesic. In this situation, the paths γ, γ' are asynchronous fellow travellers if and only if γ' is a quasigeodesic and the sets $\gamma(I)$, $\gamma'(I')$ have finite Hausdorff distance in X; moreover, the constants implicit in these properties are uniformly related. To be precise, suppose that γ is a λ, ϵ quasigeodesic. If γ, γ' are asynchronous fellow travellers with constants K, C, A as above then γ' is a λ', ϵ' quasigeodesic with λ', ϵ' depending only on $\lambda, \epsilon, K, C, A$. Conversely, if γ' is a λ', ϵ' quasigeodesic then there exist constants K, C, A depending only on $\lambda, \epsilon, \lambda', \epsilon'$, such that if $\gamma(I)$, $\gamma'(I')$ have Hausdorff distance $\le A$ then any map $\phi\colon I \to I'$ with the property that $d(\gamma(t), \gamma'(\phi(t)) \le A$ is a K, C–quasi-isometry.

2.2 Surface geometry and topology

Fix a closed, oriented surface S of genus ≥ 2.

We review Teichmüller space and the accompanying structures: the mapping class group; measured foliations and the Thurston boundary; measured geodesic laminations; quadratic differentials, geodesics, and the Teichmüller metric; and canonical bundles over Teichmüller space. Much of the material in this section is covered in more detail in Sections 2 and 4 of [10].

Teichmüller space and mapping class group Let $\mathrm{Homeo}(S)$ be the group of homeomorphisms of S and let $\mathrm{Homeo}_0(S)$ be the normal subgroup of homeomorphisms isotopic to the identity. The *mapping class group* of S is $\mathcal{MCG} = \mathrm{Homeo}(S)/\mathrm{Homeo}_0(S)$. Let $\mathcal{C}$ be the set of essential simple closed curves on S modulo isotopy, that is, modulo the action of $\mathrm{Homeo}_0(S)$. The *Teichmüller* *space* of S, denoted $\mathcal{T}$, is the set of hyperbolic structures on S modulo isotopy, or equivalently the set of conformal structures modulo isotopy. There are natural actions of $\mathcal{MCG}$ on $\mathcal{C}$ and on $\mathcal{T}$. The length pairing $\mathcal{T} \times \mathcal{C} \to \mathbf{R}_+$, associating to each $\sigma \in \mathcal{T}$, $C \in \mathcal{C}$ the length of the unique closed geodesic on the hyperbolic surface σ in the isotopy class C, induces an $\mathcal{MCG}$–equivariant embedding $\mathcal{T} \to [0,\infty)^{\mathcal{C}}$, giving $\mathcal{T}$ the $\mathcal{MCG}$–equivariant structure of a smooth manifold of dimension $6g-6$ diffeomorphic to $\mathbf{R}^{6g-6}$. The action of $\mathcal{MCG}$ on $\mathcal{T}$ is properly discontinuous and noncocompact, and so the *moduli space* $\mathcal{M} = \mathcal{T}/\mathcal{MCG}$ is a smooth, noncompact orbifold of dimension $6g-6$. The action of $\mathcal{MCG}$ on $\mathcal{T}$ is faithful except in genus 2 where there is a $\mathbf{Z}/2$ kernel generated by the hyperelliptic involution.

Measured foliations and Thurston's boundary A *measured foliation* on S is a foliation with finitely many singularities, equipped with a positive transverse Borel measure, such that each singularity is an *n–pronged singularity* for some $n \ge 3$, modelled on the singularity at the origin of the horizontal foliation of the quadratic differential $z^{n-2}dz^2$. Given a measured foliation, collapsing a saddle connection—a leaf segment connecting two singularities—results in another measured foliation. The *measured foliation space* of S, denoted $\mathcal{MF}$, is the set of measured foliations modulo the equivalence relation generated by saddle collapses and isotopies. There is a natural action of $\mathcal{MCG}$ on $\mathcal{MF}$. The *geometric intersection number* pairing $\mathcal{MF} \times \mathcal{C} \to [0,\infty)$ assigns to each $\mathcal{F} \in \mathcal{MF}$, $C \in \mathcal{C}$ the number

$$\langle \mathcal{F}, C \rangle = \inf_{c \in C, f \in \mathcal{F}} \int_c f$$

where $\int_c f$ is the integral of the transverse measure on f pulled back to a measure on c. This pairing induces $\mathcal{MCG}$–equivariant embedding $i\colon \mathcal{MF} \hookrightarrow [0,\infty)^{\mathcal{C}}$. Multiplying transverse measures by positive real numbers defines an action of $(0,\infty)$ on $\mathcal{MF}$, whose orbit space is defined to be $\mathcal{PMF}$. The embedding $i\colon \mathcal{MF} \to [0,\infty)^{\mathcal{C}}$ induces an embedding $\mathcal{PMF} \to \mathcal{P}[0,\infty)^{\mathcal{C}}$, whose image is homeomorphic to a sphere of dimension $6g-5$. The composed map $\mathcal{T} \to [0,\infty)^{\mathcal{C}} \to \mathcal{P}[0,\infty)^{\mathcal{C}}$ is an embedding, the closure of whose image is a closed ball of dimension $6g-6$ with interior $\mathcal{T}$ and boundary sphere $\mathcal{PMF}$, called the *Thurston compactification* $\overline{\mathcal{T}} = \mathcal{T} \cup \mathcal{PMF}$.

Given a simple closed curve c on S, by thickening c to form a foliated annulus with total transverse measure 1, and then collapsing complementary components of the annulus to a spine, we obtain a measured foliation on S well-defined in $\mathcal{MF}$. This gives an embedding $\mathcal{C} \hookrightarrow \mathcal{MF}$, whose induced map

$\mathbf{R}_+ \times \mathcal{C} \to \mathcal{MF}$ is also an embedding with dense image. With respect to this embedding, the geometric intersection number function $\mathcal{MF} \times (\mathbf{R}_+ \times \mathcal{C}) \to [0,\infty)$ extends continuously to an intersection number $\mathcal{MF} \times \mathcal{MF} \to [0,\infty)$, denoted $\langle \mathcal{F}_1, \mathcal{F}_2 \rangle$, $\mathcal{F}_1, \mathcal{F}_2 \in \mathcal{MF}$.

A pair of measured foliations is *transverse* if they have the same singular set, they are transverse in the usual sense at each nonsingular point, and for each singularity s there exists $n \geq 3$ such that the two foliations are modelled on the horizontal and vertical foliations of $z^{n-2}dz^2$. Given two points $\mathcal{F}_1, \mathcal{F}_2 \in \mathcal{MF}$, we say that $\mathcal{F}_1, \mathcal{F}_2$ *jointly fill* if for each $\mathcal{G} \in \mathcal{MF}$ we have $\langle \mathcal{F}_1, \mathcal{G} \rangle \neq 0$ or $\langle \mathcal{F}_2, \mathcal{G} \rangle \neq 0$. A pair $\mathcal{F}_1, \mathcal{F}_2$ jointly fills if and only if they are represented, respectively, by a transverse pair f_1, f_2; moreover, the pair f_1, f_2 is unique up to joint isotopy, meaning that if f_1', f_2' is any other transverse pair representing $\mathcal{F}_1, \mathcal{F}_2$ then there exists $h \in \mathrm{Homeo}_0(S)$ such that $f_1' = h(f_1)$ and $f_2' = h(f_2)$.

The set of jointly filling pairs forms an open subset $\mathcal{FP} \subset \mathcal{MF} \times \mathcal{MF}$. The image of this set in $\mathcal{PMF} \times \mathcal{PMF}$ we denote $\mathcal{PFP}$.

Measured geodesic laminations For details of measured geodesic laminations see [7], [25]. Here is a brief review.

Given a hyperbolic structure σ on S, a *geodesic lamination* on σ is a closed subset of S decomposed into complete geodesics of σ. A *measured geodesic lamination* is a geodesic lamination equipped with a positive, transverse Borel measure. The set of all measured geodesic laminations on σ is denoted $\mathcal{ML}(\sigma)$. A measured geodesic lamination, when lifted to the universal cover $\widetilde{\sigma} \approx \mathbf{H}^2$ with boundary circle S^1, determines a positive Borel measure on the complement of the diagonal in $S^1 \times S^1$. This embeds $\mathcal{ML}(\sigma)$ into the space of positive Borel measures, allowing us to impose a topology on $\mathcal{ML}(\sigma)$ using the weak* topology on measures. For another description of the same topology, given any simple closed curve c on S and measured geodesic lamination μ we define the intersection number $\langle c, \mu \rangle$ by pulling the transverse measure on μ back to the domain of c and integrating, and we obtain an embedding $\mathcal{ML}(\sigma) \to [0,\infty)^{\mathcal{C}}$ which is a homeomorphism onto its image, using the product topology on $[0,\infty)^{\mathcal{C}}$.

The space $\mathcal{ML}(\sigma)$ depends naturally on the hyperbolic structure σ in the following sense. For any two hyperbolic structures σ, σ' we have a homeomorphism $\mathcal{ML}(\sigma) \to \mathcal{ML}(\sigma')$, obtained by using the natural identification of the circles at infinity of the universal covers $\widetilde{\sigma}$ and $\widetilde{\sigma}'$, via the Gromov boundary of the group $\pi_1 S$. To visualize this homeomorphism, if μ is a measured geodesic lamination on σ then we may regard μ as a measured *non*-geodesic lamination on σ', which may be straightened to form the corresponding measured geodesic lamination on σ'. When $\sigma = \sigma'$ this homeomorphism is the identity; and the composition $\mathcal{ML}(\sigma) \to \mathcal{ML}(\sigma') \to \mathcal{ML}(\sigma'')$ agrees with the map $\mathcal{ML}(\sigma) \to \mathcal{ML}(\sigma'')$ for three hyperbolic structures $\sigma, \sigma', \sigma''$. We may therefore identify all of the spaces $\mathcal{ML}(\sigma)$ to a single space denoted $\mathcal{ML}$.

There is a natural isomorphism $\mathcal{ML} \approx \mathcal{MF}$, obtained by taking a measured geodesic lamination λ, collapsing the components of $S - \lambda$ to get a foliation f, and pushing the transverse measure on λ forward under the collapse map to get a transverse measure on f; the inverse map takes a measured foliation f and straightens its leaves to get a geodesic lamination λ which collapses back to f, and the transverse measure on f is pulled back under the collapse to define the transverse measure on λ. Under this isomorphism, the embeddings of $\mathcal{MF}$ and $\mathcal{ML}$ into $[0,\infty)^{\mathcal{C}}$ agree.

Since each element of $\mathcal{MF}$ is uniquely represented by a measured geodesic lamination once the hyperbolic structure is chosen, we will often use the same notation to represent either an element of $\mathcal{MF}$ or its representative measured geodesic lamination, when the hyperbolic structure is clear from the context.

Given a hyperbolic structure σ on S and a measured geodesic lamination μ on σ, let $d\mu^\perp$ denote the transverse measure on μ, let $d\mu^{||}$ denote the leafwise Lebesgue measure on leaves, and let $d\mu = d\mu^\perp \times d\mu^{||}$ denote the measure on S obtained locally as the Fubini product of $d\mu^\perp$ with $d\mu^{||}$. The support of $d\mu$ is μ, and the *length of μ with respect to σ* is defined to be

$$\mathrm{len}_\sigma(\mu) = \int d\mu.$$

We need the well known fact that length defines a continuous function:

$$\begin{aligned} \mathcal{T} \times \mathcal{MF} &\mapsto (0,\infty) \\ (\sigma,\mu) &\to \mathrm{len}_\sigma(\mu) \end{aligned}$$

Consider now a hyperbolic structure σ on S, a pair $\mathcal{F}_1, \mathcal{F}_2 \in \mathcal{MF}$, and measured geodesic laminations λ_1, λ_2 on σ representing $\mathcal{F}_1, \mathcal{F}_2$ respectively. The intersection number $\langle \mathcal{F}_1, \mathcal{F}_2 \rangle$ has the following interpretation. There exist unique maximal closed sublaminations $\lambda_1' \subset \lambda_1, \lambda_2' \subset \lambda_2$, possibly empty, with the property that λ_i' is transverse to λ_j, $i \ne j \in \{1,2\}$. By taking the Fubini product of the transverse measures on λ_1', λ_2' and integrating over S, we obtain $\langle \mathcal{F}_1, \mathcal{F}_2 \rangle$. Joint filling also has an interpretation: the pair $(\mathcal{F}_1, \mathcal{F}_2)$ jointly fills if and only if λ_1, λ_2 are transverse and each component of $S - (\lambda_1 \cup \lambda_2)$ is simply connected.

Associated to a hyperbolic structure σ on S we also have the space $\mathcal{GL}(\sigma)$ of (unmeasured) geodesic laminations with the Hausdorff topology. These spaces also depend naturally on σ, and hence we may identify them to obtain a single space $\mathcal{GL}$ depending only on S.

Quadratic differentials Given a conformal structure on S, a *quadratic differential* associates to each conformal coordinate z an expression $q(z)dz^2$ with f holomorphic, such that whenever z, w are two overlapping conformal coordinates we have $q(z) = q(w)\left(\frac{dw}{dz}\right)^2$. The *area form* of q is expressed in a conformal coordinate $z = x + iy$ as $|q(z)|\,|dx|\,|dy|$, and the integral of this form is a positive number $\|q\|$ called the *area*. We say that q is *normalized* if $\|q\| = 1$.

Given a conformal structure σ, the Riemann–Roch theorem says that the quadratic differentials on σ form a vector space QD_σ of dimension $6g-6$, and as σ varies over $\mathcal{T}$ these vector spaces fit together to form a vector bundle $\mathrm{QD} \to \mathcal{T}$. The normalized quadratic differentials form a sphere bundle $\mathrm{QD}^1 \to \mathcal{T}$.

Given a quadratic differential q, for each point $p \in S$, there exists a conformal coordinate z in which p corresponds to the origin, and a unique $n \ge 2$, such that $q(z) = z^{n-2}$; if $n = 2$ then p is a *regular point* of q, and otherwise p is a zero of order $n-2$. The coordinate z is unique up to multiplication by n^{th} roots of unity, and is called a *canonical coordinate* at p.

Associated to each quadratic differential q is a transverse pair of measured foliations, the *horizontal measured foliation* $f_x(q)$ and the *vertical measured foliation* $f_y(q)$, characterized by the property that for each regular canonical coordinate $z = x + iy$, $f_x(q)$ is the foliation by lines parallel to the x–axis with transverse measure $|dy|$, and $f_y(q)$ is the foliation by lines parallel to the y–axis with transverse measure $|dx|$. As mentioned earlier, for the quadratic differential $z^{n-2}dz^2$ on the complex plane the horizontal and vertical measured foliations have singularities at the origin which are models for a transverse pair of n–pronged singularities.

Conversely, if f_x, f_y is a transverse pair of measured foliations then there exists a unique conformal structure $\sigma(f_x, f_y)$ and quadratic differential $q = q(f_x, f_y)$ such that $f_x = f_x(q)$ and $f_y = f_y(q)$. In particular, associated to each jointly filling pair $\mathcal{F}_1, \mathcal{F}_2 \in \mathcal{MF} \times \mathcal{MF}$ are uniquely defined points $\sigma = \sigma(\mathcal{F}_1, \mathcal{F}_2) \in \mathcal{T}$, $q(\mathcal{F}_1, \mathcal{F}_2) \in \mathrm{QD}_\sigma$.

We obtain an injective map $\mathrm{QD} \to \mathcal{MF} \times \mathcal{MF}$ with image $\mathcal{FP}$, given by $q \mapsto \big([f_x(q)], [f_y(q)]\big)$, and this map is a homeomorphism between QD and $\mathcal{FP}$ [13].

Geodesics and metric on $\mathcal{T}$ Associated to each quadratic differential q is a *Teichmüller geodesic* $g\colon \mathbf{R} \to \mathcal{T}$ defined as follows:

$$g(t) = \sigma(e^{-t} f_x(q), e^t f_y(q)), \quad t \in \mathbf{R}$$

Teichmüller's theorem says that any two points $p \neq q \in \mathcal{T}$ lie on a Teichmüller line g, and that line is unique up to an isometry of the parameter line $\mathbf{R}$. Moreover, if $p = g(s)$ and $q = g(t)$, then the formula $d(p,q) = |s - t|$ defines a proper, geodesic metric on $\mathcal{T}$, called the *Teichmüller metric*. The *positive ending lamination* of the Teichmüller geodesic g is defined to be the point $\mathcal{P}f_y(q) \in \mathcal{PMF}$, and the *negative ending lamination* of g is $\mathcal{P}f_x(q)$. By uniqueness of representing transverse pairs as described above, it follows that image(g) is completely determined by the pair of points $\mathcal{P}\xi = \mathcal{P}f_x(q), \mathcal{P}\eta = \mathcal{P}f_y(q) \in \mathcal{PMF}$, and we write

$$\text{image}(g) = \overleftrightarrow{(\mathcal{P}\xi, \mathcal{P}\eta)}.$$

Given any $\sigma \in \mathcal{T}$ and $\mathcal{P}\eta \in \mathcal{PMF}$, there is a unique geodesic ray denoted $\overrightarrow{[\sigma, \mathcal{P}\eta)}$ with finite endpoint σ, given by the formula above with $t \geq 0$.

The group $\mathcal{MCG}$ acts isometrically on $\mathcal{T}$, and so the metric on $\mathcal{T}$ descends to a proper geodesic metric on the moduli space $\mathcal{M}$. A subset $A \subset \mathcal{T}$ is said to be *cobounded* if there exists a bounded subset $B \subset \mathcal{T}$ whose translates under $\mathcal{MCG}$ cover A; equivalently, the projection of A to $\mathcal{M}$ has bounded image. In a similar way we define *cocompact* subsets of $\mathcal{T}$. A subset of $\mathcal{T}$ is cocompact if and only if it is closed and cobounded. Usually we express coboundedness in terms of some bounded subset $\mathcal{B} \subset \mathcal{M}$; a subset $A \subset \mathcal{T}$ is said to be *$\mathcal{B}$–cobounded* if the projection of A to $\mathcal{M}$ is contained in $\mathcal{B}$.

Mumford's theorem provides a gauge for coboundedness. Given $\epsilon > 0$, define $\mathcal{T}_\epsilon$ to be the set of hyperbolic structures σ whose shortest closed geodesic has length $\geq \epsilon$, and define $\mathcal{M}_\epsilon$ to be the projected image of $\mathcal{T}_\epsilon$. Mumford's theorem says that the sets $\mathcal{M}_\epsilon$ are all compact, and their union as $\epsilon \to 0$ is evidently all of $\mathcal{M}$. It follows that a subset $A \subset \mathcal{T}$ is cobounded if and only if it is contained in some $\mathcal{T}_\epsilon$. We generally will not rely Mumford's gauge, instead relying on somewhat more primitive compactness arguments.

Canonical bundles over Teichmüller space There is a smooth fiber bundle $\mathcal{S} \to \mathcal{T}$ whose fiber $\mathcal{S}_\sigma$ over $\sigma \in \mathcal{T}$ is a hyperbolic surface representing the point $\sigma \in \mathcal{T}$. To make this precise, as a smooth fiber bundle we identify $\mathcal{S}$ with $S \times \mathcal{T}$, and we impose smoothly varying hyperbolic structures on the fibers $\mathcal{S}_\sigma = S \times \sigma$, $\sigma \in \mathcal{T}$, such that under the canonical homeomorphism $\mathcal{S}_\sigma \to S$ the hyperbolic structure on $\mathcal{S}_\sigma$ represents the point $\sigma \in \mathcal{T}$. The action of $\mathcal{MCG}$ on $\mathcal{T}$ lifts to a fiberwise isometric action of $\mathcal{MCG}$ on $\mathcal{S}$. Each fiber $\mathcal{S}_\sigma$ is therefore a *marked* hyperbolic surface, meaning that it comes equipped with an isotopy class of homeomorphisms to S. The bundle $\mathcal{S} \to \mathcal{T}$ is called the *canonical marked hyperbolic surface bundle* over $\mathcal{T}$.

Structures living on $\mathcal{S}_\sigma$ such as measured foliations or measured geodesic laminations can be regarded as living on S, via the identification $\mathcal{S}_\sigma \approx S$. We may therefore represent an element of $\mathcal{MF}$, for example, as a measured foliation on a fiber $\mathcal{S}_\sigma$. The same remark will apply below, without comment, when we discuss pullback bundles of $\mathcal{S}$.

The *canonical hyperbolic plane bundle* $\mathcal{H} \to \mathcal{T}$ is defined as the composition $\mathcal{H} \to \mathcal{S} \to \mathcal{T}$ where $\mathcal{H} \to \mathcal{S}$ is the universal covering map. Each fiber $\mathcal{H}_\sigma$, $\sigma \in \mathcal{T}$, is isometric to the hyperbolic plane, with hyperbolic structures varying smoothly in σ. The group $\pi_1 S$ acts as deck transformations of the covering map $\mathcal{H} \to \mathcal{S}$ is $\pi_1 S$, acting on each fiber $\mathcal{H}_\sigma$ by isometric deck transformations with quotient $\mathcal{S}_\sigma$. The action of $\pi_1 S$ on $\mathcal{H}$ extends to a fiberwise isometric action of $\mathcal{MCG}(S,p)$ on $\mathcal{H}$, such that the covering map $\mathcal{H} \to \mathcal{S}$ is equivariant with respect to the group homomorphism $\mathcal{MCG}(S,p) \to \mathcal{MCG}(S)$. By a result of Bers [1], $\mathcal{H}$ can be identified with the Teichmüller space of the once-punctured surface (S,p), and the action of $\mathcal{MCG}(S,p)$ on $\mathcal{H}$ is identified with the natural action of the mapping class group on Teichmüller space.

Let $T\mathcal{S}$ denote the tangent bundle of $\mathcal{S}$. Let $T_v\mathcal{S}$ denote the vertical sub-bundle of $T\mathcal{S}$, that is, the kernel of the derivative of the fiber bundle projection $\mathcal{S} \to \mathcal{T}$, consisting of the tangent planes to the fibers of $\mathcal{S} \to \mathcal{T}$. There exists a smoothly varying $\mathcal{MCG}$–equivariant *connection* on the bundle $\mathcal{S} \to \mathcal{T}$, which means a smooth sub-bundle $T_h\mathcal{S}$ of $T\mathcal{S}$ which is complementary to $T_v\mathcal{S}$, that is, $T\mathcal{S} = T_h\mathcal{S} \oplus T_v\mathcal{S}$; see [10] for the construction of $T_h\mathcal{S}$. Lifting to $\mathcal{H}$ we have a connection $T_h\mathcal{H}$ on the bundle $\mathcal{H} \to \mathcal{T}$, equivariant with respect to $\mathcal{MCG}(S,p)$.

Hyperbolic surface bundles over lines A *closed interval* is a closed, connected subset of $\mathbf{R}$, either a closed segment, a closed ray, or the whole line. The domains of all of our paths will be closed intervals.

Given a closed interval $I \subset \mathbf{R}$, a path $\gamma\colon I \to \mathcal{T}$ is *affine* if it satisfies $d(\gamma(s),\gamma(t)) = K\,|s-t|$ for some constant $K \geq 0$, and γ is *piecewise affine* if there is a decomposition of I into subintervals on each of which γ is affine.

Given an affine path $\gamma\colon I \to \mathcal{T}$, by pulling back the canonical marked hyperbolic surface bundle $\mathcal{S} \to \mathcal{T}$ and its connection $T_h\mathcal{S}$, we obtain a marked hyperbolic surface bundle $\mathcal{S}_\gamma \to I$ and a connection $T_h\mathcal{S}_\gamma$. This connection canonically determines a Riemannian metric on $\mathcal{S}_\gamma$, as follows. Since $T_h\mathcal{S}_\gamma$ is 1–dimensional, there is a unique vector field V on $\mathcal{S}_\gamma$ parallel to $T_h\mathcal{S}_\gamma$ such that the derivative of the map $\mathcal{S}_\gamma \to I \subset \mathbf{R}$ takes each vector in V to the positive unit vector in $\mathbf{R}$. The fiberwise Riemannian metric on $\mathcal{S}_\gamma$ now extends uniquely to a Riemannian metric on $\mathcal{S}_\gamma$ such that V is everywhere orthogonal to the fibration and has unit length.

Given a piecewise affine path $\gamma\colon I \to \mathcal{T}$, the above construction of a Riemannian metric can be carried out over each affine subpath, and at any point $t \in I$ where two such subpaths meet the metrics agree along the fibers, thereby producing a piecewise Riemannian metric on $\mathcal{S}_\gamma$.

Given a (piecewise) affine path $\gamma\colon I \to \mathcal{T}$, the above constructions can be carried out on $\mathcal{H}_\gamma$, producing a (piecewise) Riemannian metric, equivariant with respect to $\pi_1 S$, such that the covering map $\mathcal{H}_\gamma \to \mathcal{S}_\gamma$ is local isometry.

Because our paths all have domains which are closed intervals, it follows that the induced path metric of each Riemannian metric constructed above is a proper geodesic metric.

A *connection path* in either of the bundles $\mathcal{S}_\gamma \to I$, $\mathcal{H}_\gamma \to I$ is a piecewise smooth section of the projection map which is everywhere tangent to the connection. By construction, given $s, t \in I$, a path p from a point in the fiber over s to a point in the fiber over t has length $\le |s-t|$, with equality only if p is a connection path. It follows that the min distance and the Hausdorff distance between fibers are both equal to $|s-t|$. By moving points along connection paths, for each $s, t \in I$ we have well-defined maps $\mathcal{S}_s \to \mathcal{S}_t$, $\mathcal{H}_s \to \mathcal{H}_t$, both denoted h_{st} when no confusion can ensue. The following result gives some regularity for the maps h_{st}; it is closely related to a basic fact in dynamical systems, that if Φ is a smooth flow on a compact manifold then there is a constant $K \ge 1$ such that $\|D\Phi_t\| \le K^{|t|}$ for all $t \in \mathbf{R}$.

Lemma 2.2 ([10], Lemma 4.1) *For each bounded set $\mathcal{B} \subset \mathcal{M}$ and $\rho \ge 1$ there exists K such that if $\gamma\colon I \to \mathcal{T}$ is a $\mathcal{B}$–cobounded, ρ–lipschitz, piecewise affine path, then for each $s, t \in I$ the connection map h_{st} is $K^{|s-t|}$ bilipschitz.* ◻

We have associated natural geometries $\mathcal{S}_\gamma, \mathcal{H}_\gamma$ to any piecewise affine path $\gamma\colon I \to \mathcal{T}$. When γ is a geodesic there is another pair of natural geometries, the singular SOLV space $\mathcal{S}_\gamma^{\text{SOLV}}$ and its universal cover $\mathcal{H}_\gamma^{\text{SOLV}}$. To define these, there is a quadratic differential q such that

$$\gamma(t) = \sigma(e^{-t} f_x(q), e^t f_y(q)), \quad t \in I.$$

Let $|dy|$ denote the transverse measure on the horizontal foliation $f_x(q)$ and $|dx|$ the transverse measure on $f_y(q)$. On the conformal surface $\mathcal{S}_t \subset \mathcal{S}_\gamma$ we have the quadratic differential $q(e^{-t} f_x(q), e^t f_y(q))$ with horizontal transverse measure $e^{-t}|dy|$ and vertical transverse measure $e^t|dx|$. This allows us to define the *singular* SOLV *metric* on $\mathcal{S}_\gamma$ by the formula

$$ds^2 = e^{2t}|dx|^2 + e^{-2t}|dy|^2 + dt^2$$

and we denote this metric space by $\mathcal{S}_\gamma^{\text{SOLV}}$. The lift of this metric to the universal cover $\mathcal{H}_\gamma$ produces a metric space denoted $\mathcal{H}_\gamma^{\text{SOLV}}$.

The following result says that the metric on $\mathcal{H}_\gamma$ is quasi-isometrically stable with respect to perturbation of γ. Moreover, if γ fellow travels a geodesic γ' then the singular SOLV geometry $\mathcal{H}_{\gamma'}^{\text{SOLV}}$ serves as a model geometry.

Proposition 2.3 ([10], Proposition 4.2) *For any $\rho \ge 1$, any bounded subset $\mathcal{B} \subset \mathcal{M}$, and any $A \ge 0$ there exists $K \ge 1$, $C \ge 0$ such that the following holds. If $\gamma, \gamma' \colon I \to \mathcal{T}$ are two ρ–lipschitz, $\mathcal{B}$–cobounded, piecewise affine paths defined on a closed interval I, and if $d(\gamma(s), \gamma'(s)) \le A$ for all $s \in I$, then there exists a map $\mathcal{S}_\gamma \to \mathcal{S}_{\gamma'}$ taking each fiber $\mathcal{S}_{\gamma(t)}$ to the fiber $\mathcal{S}_{\gamma'(t)}$ by a homeomorphism in the correct isotopy class, such that any lifted map $\mathcal{H}_\gamma \to \mathcal{H}_{\gamma'}$ is a K,C–quasi-isometry.*

If γ' is a geodesic, the same is true with $\mathcal{S}_{\gamma'}$, $\mathcal{H}_{\gamma'}$ replaced by the singular SOLV *spaces $\mathcal{S}_{\gamma'}^{\text{SOLV}}$, $\mathcal{H}_{\gamma'}^{\text{SOLV}}$.*

Remark Given any ρ–lipschitz path $\gamma \colon I \to \mathcal{T}$, the bundles $\mathcal{S} \to \mathcal{T}$, $\mathcal{H} \to \mathcal{T}$ can be pulled back to define bundles $\mathcal{S}_\gamma \to I$, $\mathcal{H}_\gamma \to I$. Despite the paucity of smoothness, one can extend the fiberwise hyperbolic metrics on these bundles to measurable Riemannian metrics which then determine proper geodesic metrics in a canonical manner, and hence we would associate a geometry $\mathcal{H}_\gamma$ to γ. However, it is easier to proceed by approximating γ with a ρ–lipschitz piecewise affine path, and to apply Proposition 2.3 to show that the resulting geometry on $\mathcal{H}_\gamma$ is well-defined up to quasi-isometry. This allows us to reduce the proof of Theorem 1.1 to the case where γ is piecewise affine, a technical simplification.

3 Proof of Theorem 1.1

For the proof we fix the closed, oriented surface S of genus ≥ 2, with Teichmüller space $\mathcal{T}$, mapping class group $\mathcal{MCG}$, measured foliation space $\mathcal{MF}$, etc.

3.1 Setting up the proof

Throughout the proof we fix a compact subset $\mathcal{B} \subset \mathcal{M}$, and numbers $\rho \geq 1$, $\delta \geq 0$.

Given $\gamma\colon I \to \mathcal{T}$ a $\mathcal{B}$–cobounded, ρ–Lipschitz, path, Proposition 2.3 says that if we perturb γ to be a piecewise affine path, then the large scale geometry of $\mathcal{H}_\gamma$ is well-defined up to quasi-isometry. The truth of the hypothesis of Theorem 1.1 is therefore unaffected by perturbation, as is the conclusion, with uniform changes in all constants depending only on the size of the perturbation. We shall choose a particular perturbation which will be technically useful in what follows, particularly in Section 3.9.

A $\mathbf{Z}$*–piecewise affine path* is a path $\gamma\colon I \to \mathcal{T}$, defined on a closed interval I whose endpoints, if any, are in $\mathbf{Z}$, such that γ is affine on each subinterval $[n, n+1]$ with $n, n+1 \in I \cap \mathbf{Z}$. We shall often denote $J = I \cap \mathbf{Z}$. If γ is a ρ–lipschitz path defined on an interval I, we can perturb γ to be $\mathbf{Z}$–piecewise affine as follows: first remove a bit from each finite end of I of length less than 1 so that the endpoints are in $\mathbf{Z}$; then replace $\gamma \bigm| [n, n+1]$ by an affine path with the same endpoints whenever $n, n+1 \in J$; note that $d(\gamma(t), \gamma'(t)) \leq \rho$ for all $t \in I$.

Given a $\mathcal{B}$–cobounded, ρ–lipschitz, $\mathbf{Z}$–piecewise affine path $\gamma\colon I \to \mathcal{T}$ whose canonical hyperbolic plane bundle $\mathcal{H}_\gamma \to I$ is δ–hyperbolic, our goal is to construct a Teichmüller geodesic g, sharing any endpoints with γ, and to show that γ and g are at bounded Hausdorff distance and γ is a quasigeodesic.

Motivation The construction of g is inspired by ending laminations methods ([25], chapter 9; [4]), uniform foliations methods [26], and the flaring concepts from [2].

In the case of a line $\gamma\colon \mathbf{R} \to \mathcal{T}$, the idea for finding g is to keep in mind an analogy between the hyperbolic surface bundle $\mathcal{S}_\gamma \to \mathbf{R}$ and a doubly degenerate hyperbolic structure on $S \times \mathbf{R}$. The hyperbolic surfaces $\mathcal{S}_t \subset \mathcal{S}_\gamma$ $(t \in \mathbf{R})$ approach both ends of $\mathcal{S}_\gamma$ as $t \to \pm\infty$ and so, following the analogy, $\mathcal{S}_\gamma$ is "geometrically tame" in the sense of [4] and [25] chapter 9. This suggests that we produce an ending lamination in $\mathcal{PMF}$ for each of the two ends. This pair of laminations will jointly fill the surface and so will determine a geodesic g in $\mathcal{T}$.

Despite the analogy, our construction of ending laminations for $\mathcal{S}_\gamma$ is entirely self-contained and new. The construction is inspired by uniform foliations methods [26], in which large-scale geometry of a foliation is used to determine laminations on leaves of that foliation. The main new idea is that, in the presence of δ–hyperbolicity, laminations can be constructed using flaring concepts from [2]. The only prerequisites for the construction are basic facts about measured geodesic laminations and about flaring.

3.2 Flaring

In order to get the proof off the ground, the key observation needed is that hyperbolicity of $\mathcal{H}_\gamma$ is equivalent to the "rectangles flare" property introduced by Bestvina and Feighn [2]. Sufficiency of the rectangles flare property was proved by Bestvina and Feighn, and necessity was proved by Gersten [11]. In [10] the rectangles flare property is recast in a manner which is followed here.

Consider a sequence of positive numbers r_j $(j \in J)$ indexed by a subinterval J of the integers $\mathbf{Z}$, which can be finite, half-infinite, or all of $\mathbf{Z}$. Given $\kappa > 1$, $n \in \mathbf{Z}_+$, $A \geq 0$, the κ, n, A*–flaring property* says that if $j-n, j, j+n \in J$, and if $r_j > A$, then $\max\{r_{j-n}, r_{j+n}\} \geq \kappa r_j$. Given $L \geq 1$, the L*–lipschitz growth condition* says that for every $j,k \in J$ we have $r_j/r_k \leq L^{|j-k|}$ for $j,k \in J$; equivalently, if $|j-k| = 1$ then $r_j \leq L r_k$. Given $L \geq 1$, $D \geq 0$, the (L,D) *coarse lipschitz growth condition* says that if $|j-k| = 1$ then $r_j \leq L r_k + D$.

The flaring property and the coarse lipschitz growth condition work together. The κ, n, A flaring property alone only controls behavior on arithmetic subsequences of the form $j_0 + kn$, e.g. if $r_{j_0} > A$ then there exists $\epsilon = \pm 1$ such that

$$r_{j_0+\epsilon kn} > r_j \kappa^k \quad \text{for all} \quad k \geq 1. \tag{3.1}$$

But in company with, say, the L–lipschitz growth condition, growth of disjoint arithmetic subsequences is conjoined, in that

$$r_{j_0+\epsilon m} > r_j L^{-n} \kappa'^m, \quad \text{for all} \quad m \geq 1. \tag{3.2}$$

where $\kappa' = \kappa^{1/n}$. A similar result holds for a coarse lipschitz growth condition.

The number A is called the *flaring threshold*. A κ, n, A flaring sequence can stay $\leq A$ on a subinterval of arbitrary length, but as we have seen, at any point where the sequence gets above the flaring threshold then exponential growth is inexorable in at least one of the two directions. In particular, if the flaring threshold is zero then exponential growth is everywhere.

Consider now a cobounded, lipschitz, piecewise affine path $\gamma\colon I \to \mathcal{T}$ and the hyperbolic plane bundle $\mathcal{H}_\gamma \to I$. Given $s \in I$ let d_s denote the distance function on the fiber $\mathcal{H}_s$. A λ*–quasihorizontal path* in $\mathcal{H}_\gamma$, $\lambda \geq 1$, is a section $\alpha\colon I \to \mathcal{H}_\gamma$ of the bundle projection which is λ–lipschitz. A 1–quasihorizontal path is the same thing as a connection path. For any pair of λ–quasihorizontal paths $\alpha, \beta\colon I \to \mathcal{H}_\gamma$, the sequence of distances

$$d_j(\alpha(j), \beta(j)), \quad j \in J = I \cap \mathbf{Z}$$

automatically satisfies an (L, D) coarse Lipschitz growth condition, where $L = K$ is the constant given by Lemma 2.2 and hence depends only on the coboundedness and the lipschitz constant of γ, and where $D = 2K(\lambda+1)$. To see this, given $j, k \in J$ with $|j-k| = 1$, let $a = \alpha(j)$, $a' = \alpha(k)$, $a'' = h_{jk}(a)$, $b = \beta(j)$, $b' = \beta(k)$, $b'' = h_{jk}(b)$. The points a', a'' are connected by a path of length $\leq \lambda + 1$ staying between $\mathcal{H}_j$ and $\mathcal{H}_k$, consisting of a segment of α from a' to a and a connection path from a to a''; projection of this path onto $\mathcal{H}_k$ is K–lipschitz, and similarly for the b's, and so:

$$\begin{aligned}
d_k(a', a'') &\leq K(\lambda+1) \\
d_k(b', b'') &\leq K(\lambda+1) \\
d_k(a', b') &\leq d_k(a', a'') + d_k(a'', b'') + d_k(b'', b') \\
&\leq K(\lambda+1) + K d_j(a, b) + K(\lambda+1)
\end{aligned}$$

Given constants $\kappa \geq 1$, $n \in \mathbf{Z}_+$, and a function $A(\lambda) \geq 0$ defined for $\lambda \geq 1$, we say that $\mathcal{H}_\gamma$ satisfies the $\kappa, n, A(\lambda)$ *horizontal flaring property* if for any $\lambda \geq 1$ and any λ–quasihorizontal paths $\alpha, \beta\colon I \to \mathcal{H}_\gamma$, the sequence of distances

$$d_j(\alpha(j), \beta(j)), \quad j \in J = I \cap \mathbf{Z}$$

satisfies $\kappa, n, A(\lambda)$ flaring.

The following result is an almost immediate consequence of Gersten's theorem [11] which gives the converse to the Bestvina–Feighn combination theorem; the interface with Gersten's theorem is explained in [10]. See also Lemma 5.2 of [10] for an alternative proof following the lines of the well-known fact that in a hyperbolic metric space, geodesic rays satisfy an exponential divergence property [6].

Proposition 3.1 *Given a ρ–lipschitz, $\mathcal{B}$–cobounded map $\gamma\colon I \to \mathcal{T}$, if $\mathcal{H}_\gamma$ is δ–hyperbolic, then $\mathcal{H}_\gamma$ satisfies a horizontal flaring property, with flaring data $\kappa, n, A(\lambda)$ dependent only $\rho, \mathcal{B}, \delta$.* □

In the context of the proof of Theorem 1.1, for each $\mu \in \mathcal{MF}$, each marked hyperbolic surface $\mathcal{S}_t$ in the bundle $\mathcal{S}_\gamma$ has a measured geodesic lamination μ_t in the class μ. Denote $\mathrm{len}_t(\mu) = \mathrm{len}_{\gamma(t)}(\mu_t)$. The key to our proof of

Theorem 1.1 is to study how the function $\text{len}_t(\mu)$ varies in t and in μ. What gets us off the ground is Lemma 3.6 which shows that for each $\mu \in \mathcal{MF}$ the function $t \mapsto \text{len}_t(\mu)$ flares, with uniform flaring data independent of μ, and with a flaring threshold of zero; intuitively this follows from the flaring of $\mathcal{H}_\gamma$. The technical step of verifying uniform flaring data depends on the construction of a flow preserving connection on the leafwise geodesic flow bundle of $\mathcal{S}_\gamma$, to which we now turn.

3.3 Connections on geodesic flow bundles

Let $\mathcal{GFL}(F)$ denote the geodesic flow of a hyperbolic surface F, defined on the unit tangent bundle T^1F of F. The foliation of $\mathcal{GFL}(F)$ by flow lines, being invariant under the antipodal map $\vec{v} \mapsto -\vec{v}$ of T^1F, descends to a foliation on the tangent line bundle $\mathcal{P}T^1F = T^1F/\pm$, the *geodesic foliation* $\mathcal{GF}(F)$. The Liouville metric on T^1F descends to a metric on $\mathcal{P}T^1F$ also called the Liouville metric.

Starting from the canonical marked hyperbolic surface bundle $\mathcal{S} \to \mathcal{T}$, by taking the geodesic flow on each fiber we obtain the *fiberwise geodesic flow bundle* $\mathcal{GFL}(\mathcal{S}) \to \mathcal{T}$ whose fiber over $\sigma \in \mathcal{T}$ is $\mathcal{GFL}(\mathcal{S}_\sigma)$, and we similarly obtain the fiberwise geodesic foliation bundle $\mathcal{GF}(\mathcal{S}) \to \mathcal{T}$. Note that the geodesic flows on the fibers of $\mathcal{GFL}(\mathcal{S})$ fit together to form a smooth flow on $\mathcal{GFL}(\mathcal{S})$, and similarly for the geodesic foliations on fibers of $\mathcal{GF}(\mathcal{S})$; smoothness follows from the fact that the leafwise hyperbolic metrics on $\mathcal{S} \to \mathcal{T}$ vary smoothly, together with the fact that the geodesic differential equation has coefficients and also solutions varying smoothly with the metric. We have a fiber bundle $\mathcal{GFL}(\mathcal{S}) \to \mathcal{S}$ whose fiber over $x \in \mathcal{S}_\sigma \subset \mathcal{S}$ is the unit tangent space $T^1_x\mathcal{S}_\sigma$, and similarly a fiber bundle $\mathcal{GF}(\mathcal{S}) \to \mathcal{S}$ whose fiber is the space of tangent lines $\mathcal{P}T^1_x\mathcal{S}_\sigma$.

For any piecewise affine path $\gamma\colon I \to \mathcal{T}$, we have pullback bundles $\mathcal{GFL}(\mathcal{S}_\gamma) \to I$ and $\mathcal{GF}(\mathcal{S}_\gamma) \to I$, whose fibers over $t \in I$ are $\mathcal{GFL}(\mathcal{S}_t)$ and $\mathcal{GF}(\mathcal{S}_t)$ respectively.

In the case where γ is affine, by a *connection* on the bundle $\mathcal{GF}(\mathcal{S}_\gamma) \to I$ we mean a connection which preserves the geodesic foliations, that is: a 1–dimensional sub-bundle of the tangent bundle of $\mathcal{GF}(\mathcal{S}_\gamma)$ which is complementary to the vertical sub-bundle of $\mathcal{GF}(\mathcal{S}_\gamma) \to I$, such that for any $s,t \in I$, the map $H_{st}\colon \mathcal{GF}(\mathcal{S}_s) \to \mathcal{GF}(\mathcal{S}_t)$ obtained by moving points along connection paths takes leaves of $\mathcal{GF}(\mathcal{S}_s)$ to leaves of $\mathcal{GF}(\mathcal{S}_t)$. The *connection flow* Φ on $\mathcal{GF}(\mathcal{S}_\gamma)$ is defined by $\Phi_r(l) = H_{s,s+r}(l)$ whenever $l \in \mathcal{GF}(\mathcal{S}_s)$.

In the general case where γ is only piecewise affine, the connection flows defined over the intervals where γ is affine piece together to define a connection flow Φ on all of $\mathcal{GF}(\mathcal{S}_\gamma)$, with corresponding connection maps $H_{st}(l) = \Phi_{t-s}(l)$ for all $s, t \in I$, $l \in \mathcal{GF}(\mathcal{S}_s)$.

A connection on $\mathcal{GF}(\mathcal{S}_\gamma)$ is *leafwise* L*–bilipschitz* if H_{st} restricts to a $L^{|s-t|}$–bilipschitz homeomorphism between leaves of the geodesic foliation.

All the above concepts apply as well to the fiberwise geodesic flow bundle and geodesic foliation bundle $\mathcal{GFL}(\mathcal{H}) \to \mathcal{T}$, $\mathcal{GF}(\mathcal{H}) \to \mathcal{T}$, and any associated pullback bundles.

Lemma 3.2 *For each bounded set $\mathcal{B} \subset \mathcal{M}$ and each $\rho \geq 1$ there exists $L \geq 1$ and $\lambda \geq 1$ such that if $\gamma\colon I \to \mathcal{T}$ is a $\mathcal{B}$–cobounded, ρ–lipschitz, piecewise affine path, then there is a leafwise L–bilipschitz connection on the bundle $\mathcal{GF}(\mathcal{S}_\gamma) \to I$, and each connection line in $\mathcal{GF}(\mathcal{S}_\gamma)$ projects to a path in $\mathcal{S}_\gamma$ which is λ–quasihorizontal. By lifting that we obtain a $\pi_1(S)$–equivariant leafwise L–bilipschitz connection on $\mathcal{GF}(\mathcal{H}_\gamma)$ whose connection lines project to λ–quasihorizontal paths in $\mathcal{H}_\gamma$. Moreover the connection satisfies the following uniform continuity condition: for every bounded $\mathcal{B} \subset \mathcal{M}$, $\rho \geq 1$, $M > 0$, and $\epsilon > 0$, there exists $\delta > 0$ such that if $\gamma\colon I \to \mathcal{T}$ is $\mathcal{B}$–cobounded, ρ–lipschitz, and piecewise affine, if $s, t \in I$ with $|s-t| < M$, and if $l, m \in \mathcal{GF}(\mathcal{S}_s)$ with $d(l,m) < \delta$, then $d(H_{st}(l), H_{st}(m)) < \epsilon$.*

I am grateful to Amie Wilkinson for the proof of this lemma, particularly for explaining how to apply partially hyperbolic dynamics.

Proof The intuition behind the proof is that the geodesic flows on $\mathcal{S}_s$ and $\mathcal{S}_t$ are topologically conjugate; to put it another way, for each geodesic ℓ on $\mathcal{S}_s$, the connection on $\mathcal{S}_\gamma$ moves ℓ to a $K^{|s-t|}$–bilipschitz path in $\mathcal{S}_t$ and hence that path is close to a geodesic in $\mathcal{S}_t$. In order to carry this out uniformly up in the geodesic foliation bundle $\mathcal{GF}(\mathcal{S}_\gamma)$ we shall apply structural stability tools from the theory of hyperbolic dynamical systems, as encapsulated in the Sublemma 3.3.

A k–dimensional foliation of a subset of $\mathbf{R}^n$ has *uniformly smooth leaves* if the leaves are defined locally by immersions from open subsets of $\mathbf{R}^k$ into $\mathbf{R}^n$ such that for each $r \geq 0$ the partial derivatives up to order r are uniformly bounded away from zero and from infinity. A foliation of a smooth manifold M has *locally uniformly smooth leaves* if M is covered by coordinate charts in each of which the leaves are uniformly smooth. The property of *(locally) uniformly C^r leaves* is similarly defined by omitting the words "for each $r \geq 0$".

Let D be the dimension of $\mathcal{T}$.

Sublemma 3.3 *There exists a unique, $D+1$ dimensional foliation $\mathcal{G}$ of $\mathcal{GFL}(\mathcal{S})$ with locally uniformly smooth leaves such that $\mathcal{G}$ is transverse to the fibers of $\mathcal{GFL}(\mathcal{S}) \to \mathcal{T}$, and the foliation of $\mathcal{GFL}(\mathcal{S})$ obtained by intersecting $\mathcal{G}$ with the fibers of $\mathcal{GFL}(\mathcal{S}) \to \mathcal{T}$ is identical to the foliation by geodesic flow lines.*

Before proving the claim we apply it to prove Lemma 3.2.

Since the conclusion of the lemma is local, it suffices to prove it when γ is a ρ–lipschitz, $\mathcal{B}$–cobounded affine arc $\gamma\colon [0,1] \to \mathcal{T}$.

Choose a compact set $\mathcal{A} \subset \mathcal{T}$ such that any ρ–lipschitz, $\mathcal{B}$–cobounded affine arc $[0,1] \to \mathcal{T}$ may be translated by $\operatorname{Isom}(\mathcal{T})$ to lie in $\mathcal{A}$. Let $C_\rho(\mathcal{A})$ be the space of all ρ–lipschitz affine arcs $[0,1] \to \mathcal{A}$. The conclusion of the lemma is invariant under the action of $\mathcal{MCG}$ and so we may assume $\gamma \in C_\rho(\mathcal{A})$. By enlarging $\mathcal{A}$ we may assume that $\mathcal{A}$ is a smooth codimension 0 submanifold of $\mathcal{T}$. By restricting $\mathcal{S}$ to $\mathcal{A}$ we obtain a hyperbolic surface bundle $\mathcal{S}_\mathcal{A}$ and its geodesic flow bundle $\mathcal{GFL}(\mathcal{S}_\mathcal{A})$; fix a smooth Riemannian metric on $\mathcal{GFL}(\mathcal{S}_\mathcal{A})$. Let $\mathcal{G}_\mathcal{A}$ be the restriction of the foliation $\mathcal{G}$ to $\mathcal{GFL}(\mathcal{S}_\mathcal{A})$. For each $\gamma \in C_\rho(\mathcal{A})$, the foliation $\mathcal{G}_\mathcal{A}$ restricts to a 2–dimensional foliation $\mathcal{G}_\gamma$ of $\mathcal{GFL}(\mathcal{S}_\gamma)$ with uniformly smooth leaves, transverse to the fibers of $\mathcal{GFL}(\mathcal{S}_\gamma) \to [0,1]$. Also, the Riemannian metric on $\mathcal{GFL}(\mathcal{S}_\mathcal{A})$ restricts to a smooth Riemannian metric on $\mathcal{GFL}(\mathcal{S}_\gamma)$. There is a unique vector field V_γ on $\mathcal{GFL}(\mathcal{S}_\gamma)$ which is tangent to $\mathcal{G}_\gamma$ and is perpendicular to the geodesic flow lines, such that each $v \in V_\gamma$ projects to a positive unit tangent vector in $[0,1] = \operatorname{domain}(\gamma)$. By uniqueness the foliation $\mathcal{G}$ is invariant under the antipodal map on $\mathcal{GFL}(\mathcal{S})$, implying that $\mathcal{G}_\gamma$ is antipode invariant on $\mathcal{GFL}(\mathcal{S}_\gamma)$. Assuming as we may that the Riemannian metric on $\mathcal{GFL}(\mathcal{S}_\mathcal{A})$ is also antipode-invariant, it follows that V_γ is antipode-invariant, and so descends to a vector field on $\mathcal{GF}(\mathcal{S}_\gamma)$. This vector field spans the desired connection on $\mathcal{GF}(\mathcal{S}_\gamma)$.

Note that the connection on $\mathcal{GF}(\mathcal{S}_\gamma)$ is uniformly smooth along leaves of $\mathcal{G}_\gamma$, and as γ varies over $C_\rho(\mathcal{A})$ the connection varies continuously; this follows from leafwise uniform smoothness of $\mathcal{G}$.

Let $\psi_{\gamma,t}$ be the connection flow on $\mathcal{GF}(\mathcal{S}_\gamma)$, which has connection maps $H_{st}\colon \mathcal{GF}(\mathcal{S}_s) \to \mathcal{GF}(\mathcal{S}_t)$, that is, $H_{st}(\ell) = \psi_{\gamma,t-s}(\ell)$ for $\ell \in \mathcal{GF}(\mathcal{S}_s)$. To prove that H_{st} is $L^{|s-t|}$ bilipschitz, it suffices by a standard result in differential equations to prove that

$$S(\gamma) = \left| \frac{d}{dt} \left\| D\psi_{\gamma,t} \right\| \right|_{t=0} \le \log(L).$$

Note that $S(\gamma)$ is continuous as a function of $\gamma \in C_\rho(\mathcal{A})$. Since $C_\rho(\mathcal{A})$ is compact, $S(\gamma)$ has a finite upper bound $l \geq 0$, and so $L = e^l$ is the desired bilipschitz constant.

Further compactness arguments show that the connection lines project to λ–quasihorizontal lines in $\mathcal{S}_\gamma$ for uniform λ, and that the uniform continuity clause holds. $\square$

Proof of Sublemma 3.3 Uniqueness of $\mathcal{G}$ follows because, for any $\sigma \in \mathcal{T}$ and any closed geodesic c in $\mathcal{S}_\sigma$, and for any $\tau \in \mathcal{T}$, the leaf of $\mathcal{G}$ containing c must contain the closed geodesic in $\mathcal{S}_\tau$ that is in the isotopy class of c with respect to the canonical homeomorphism $\mathcal{S}_\sigma \approx \mathcal{S}_\tau$. The non-simply connected leaves of $\mathcal{G}$ are therefore determined, but they are dense in $\mathcal{GFL}(\mathcal{S})$ and so $\mathcal{G}$ is determined.

Existence of $\mathcal{G}$ is a purely local phenomenon, because if we have open subsets $U, V \subset \mathcal{T}$ and foliations $\mathcal{G}_U, \mathcal{G}_V$ on $\mathcal{GFL}(\mathcal{S}_U), \mathcal{GFL}(\mathcal{S}_V)$ respectively, so that $\mathcal{G}_U, \mathcal{G}_V$ each satisfy the conclusions of the sublemma, then the uniqueness argument above can be applied locally to show that the restrictions to $\mathcal{GFL}(\mathcal{S}_{U\cap V})$ of $\mathcal{G}_U, \mathcal{G}_V$ are identical, and so $\mathcal{G}_U, \mathcal{G}_V$ can be pasted together over $U \cup V$ to give a foliation of $\mathcal{GFL}(\mathcal{S}_{U\cup V})$ satisfying the conclusions of the sublemma. Arguing similarly with respect to some locally finite open cover of $\mathcal{T}$ allows one to construct $\mathcal{G}$.

For each $\sigma \in \mathcal{T}$ it therefore suffices to find some closed ball B in $\mathcal{T}$ around σ and a $D+1$ dimensional foliation $\mathcal{G}_B$ of $\mathcal{GFL}(\mathcal{S}_B)$ satisfying the conclusions of the sublemma, namely, that $\mathcal{G}_B$ has uniformly smooth leaves and is transverse to the fibers of $\mathcal{GFL}(\mathcal{S}_B) \to B$, and the intersection of $\mathcal{G}_B$ with each fiber of is the foliation of that fiber by geodesic flow lines.

We review some elements of partially hyperbolic dynamical systems from [14]. Let M be a smooth compact Riemannian manifold. A C^r–flow $\phi\colon M \times \mathbf{R} \to M$ is *r–normally hyperbolic* at a foliation $\mathcal{F}$ if there is a splitting $TM = E^u \oplus T\mathcal{F} \oplus E^s$, invariant under the flow ϕ, and there exists $a > 0$, such that for each $t > 0$ we have

- $\left\|T\phi_t \mid E^s\right\| < e^{-at}$,
- $\left\|T\phi_{-t} \mid E^u\right\| < e^{-at}$,
- $\left\|T\phi_{-t} \mid T\mathcal{F}\right\|^r \cdot \left\|T\phi_t \mid E^s\right\| < e^{-at}$,
- $\left\|T\phi_t \mid T\mathcal{F}\right\|^r \cdot \left\|T\phi_{-t} \mid E^u\right\| < e^{-at}$.

We need the following results from [14]:

Theorem 3.4 *Suppose that the flow ϕ is r–normally hyperbolic at $\mathcal{F}$. If the foliation $\mathcal{F}$ is C^1, then for every C^r flow ψ, sufficiently C^1–close to ϕ, there exists a foliation $\mathcal{G}$ such that ψ is r–normally hyperbolic at $\mathcal{G}$, and such that the dimensions of corresponding summands in the splittings of TM associated to $\mathcal{F}$ and to $\mathcal{G}$ are identical.*

Theorem 3.5 *If the flow ψ is r–normally hyperbolic at the foliation $\mathcal{G}$, then the leaves of $\mathcal{G}$ are uniformly C^r.*

Given a smooth closed ball B in $\mathcal{T}$, consider the $D+3$ dimensional manifold $M = T^1\mathcal{S}_B$, the fiberwise unit tangent bundle of $\mathcal{S}_B$, whose fiber over the point $x \in \mathcal{S}_\tau \subset \mathcal{S}_B$, where $\tau \in B$, is $T^1_x(\mathcal{S}_\tau)$. There is also a fibration of $T^1\mathcal{S}_B$ over B, whose fiber over $\tau \in B$ is $T^1\mathcal{S}_\tau$. As a manifold, the space $T^1\mathcal{S}_B$ is identified with the underlying space of the geodesic flow bundle $\mathcal{GFL}(\mathcal{S}_B)$. Let ψ be the fiberwise geodesic flow on $T^1\mathcal{S}_B$ (i.e. the flow on $\mathcal{GFL}(\mathcal{S}_B)$).

We claim that if $\mathcal{G}$ is a codimension–2 foliation of $T^1\mathcal{S}_B$ at which ψ is r–normally hyperbolic, so that the bundles E^s, E^u are 1–dimensional, then $\mathcal{G}$ is transverse to the fibers of the fibration $T^1\mathcal{S}_B \to B$, and the intersection of $\mathcal{G}$ with each fiber $T^1\mathcal{S}_\tau$, $\tau \in B$, is the geodesic flow of the hyperbolic surface $\mathcal{S}_\tau$. Uniform smoothness of leaves of $\mathcal{G}$ follows from Theorem 3.5, thereby proving Sublemma 3.3.

To prove the claim, note that the 1–dimensional line bundle $T\psi$ is a sub-bundle of $T\mathcal{G}$, and moreover $T\psi$ is tangent to each the fibers $T^1\mathcal{S}_\tau$. Each of these fibers is 3–dimensional, and so transversality of $\mathcal{G}$ to these fibers will follow by proving that in the splitting $TM = E^u \oplus T\mathcal{G} \oplus E^s$, the sub-bundle $E^u \oplus T\psi \oplus E^s$ is identical to the fiberwise tangent bundle of the fibration $T^1\mathcal{S}_B \to B$. This will follow in turn by proving that each of the 1–dimensional bundles E^u, E^s is tangent to the fibers $T^1\mathcal{S}_\tau$. The restriction of ψ to each fiber $T^1\mathcal{S}_\tau$ is the geodesic flow, which is known to be an Anosov flow with 1–dimensional stable and unstable bundles. Moreover, for any vector $v \in TM$ which is not tangent to a fiber $T^1\mathcal{S}_\tau$, the component of v transverse to the fibers is preserved in norm by the flow ψ, with respect to a Riemannian metric on M that assigns constant distance from any point in one fiber to any other fiber, and so v cannot be in the stable or unstable bundle of the splitting $E^s \oplus T\mathcal{G} \oplus E^u$. It follows that E^u is the same as the Anosov unstable bundle of the fiberwise geodesic flow on $T^1\mathcal{S}_B$, and E^s is the same as the Anosov stable bundle, and so E^s, E^u are indeed tangent to the fibers $T^1\mathcal{S}_\tau$. And since $T\psi$ is a sub-bundle of $T\mathcal{G}$ and E^s, E^u are independent of $T\mathcal{G}$, it follows that the restriction of $\mathcal{G}$ to each fiber $T^1\mathcal{S}_\tau$ is indeed the geodesic flow.

We have therefore reduced the proof of Sublemma 3.3 to the construction of a foliation $\mathcal{G}$ on $T^1\mathcal{S}_B$ at which the fiberwise geodesic flow is r–normally hyperbolic, for some open neighborhood B of any point $\sigma \in \mathcal{T}$. We carry out this construction using Theorem 3.4.

Fix a point σ of $\mathcal{T}$. To start with let B be any smooth closed ball in $\mathcal{T}$ whose interior contains σ. Take $M = T^1\mathcal{S}_B$ as above. Pick a diffeomorphism $\mathcal{S}_B \approx \mathcal{S}_\sigma \times B$ respecting projection to B. This induces a diffeomorphism $\Theta\colon T^1\mathcal{S}_B \to T^1\mathcal{S}_\sigma \times B$. The geodesic flow on $T^1\mathcal{S}_\sigma \times B$ pulls back via Θ to a flow ϕ on $T^1\mathcal{S}_B$. Also, there is a $D+1$ dimensional foliation of $T^1\mathcal{S}_\sigma \times B$ which is the product of the geodesic flow on $T^1\mathcal{S}_\sigma$ with B; pulling this foliation back via Θ we obtain a foliation of $T^1\mathcal{S}_B$ denoted $\mathcal{F}$. Noting that $\left\| T\phi_t \mid T\mathcal{F} \right\| = 1$ for all t, from the fact that the geodesic flow on $\mathcal{S}_\sigma$ is Anosov it follows that ϕ is r–normally hyperbolic at $\mathcal{F}$ for all r, with 1–dimensional stable and unstable bundles E^s, E^u.

For the flow ψ on $T^1\mathcal{S}_B$ we would like to take the leafwise geodesic flow. However, we have no control on the C^1 distance between ϕ and ψ as required to apply Theorem 3.4. To fix this, we want to "choose the ball B to be sufficiently small", but we must do this in a way that does not change the domain manifold $M = T^1\mathcal{S}_B$. Choose a diffeomorphism between (B,σ) and the unit ball in Euclidean space centered at the origin. With respect to this diffeomorphism let $(0,1)\times B \to B$ denote scalar multiplication, and so $s\cdot B$ corresponds to the ball of radius s in Euclidean space. As s approaches zero, lift the maps $B \to s\cdot B$ to a smooth family of diffeomorphisms $T^1\mathcal{S}_B \to T^1\mathcal{S}_{s\cdot B}$ so that for each $b \in B$ the diffeomorphisms $T^1\mathcal{S}_b \to T^1\mathcal{S}_{s\cdot b}$ converge uniformly to $\Theta\colon T^1\mathcal{S}_b \to T^1\mathcal{S}_\sigma$ as $s \to 0$. Pulling back the leafwise geodesic flow on $T^1\mathcal{S}_{s\cdot B}$ we obtain a smooth family of flows ψ_s on $T^1\mathcal{S}_B$ converging to the flow ϕ in the C^1 topology, as $s \to 0$.

We may now apply Theorem 3.4 to ψ_s for s sufficiently close to 0, to obtain a foliation $\mathcal{G}_s$ on $T^1\mathcal{S}_B$ at which ψ_s is r–normally hyperbolic. Pulling back to the ball $B' = s\cdot B$ around σ we have constructed a foliation $\mathcal{G}_{B'}$ at which the leafwise geodesic flow is r–normally hyperbolic. ◻

3.4 Flaring of geodesic laminations

Until further notice we shall fix a $\mathbf{Z}$–piecewise affine path $\gamma\colon I \to \mathcal{T}$ which is $\mathcal{B}$–cobounded and ρ–lipschitz, such that $\mathcal{H}_\gamma$ is δ–hyperbolic. Let $J = I \cap \mathbf{Z}$.

For each $\mu \in \mathcal{MF}$, let μ_t denote the measured geodesic lamination on $\mathcal{S}_t$ representing μ. For $i \in J$ let $\operatorname{len}_i(\mu)$ be the length of μ_i in the hyperbolic surface $\mathcal{S}_i$. The following lemma says that the sequence $\operatorname{len}_i(\mu)$ flares with uniform flaring constants, and with a flaring threshold of *zero*.

Lemma 3.6 (Length flares) *There exist constants $L \ge 1$, $\kappa > 1$, $n \in \mathbf{Z}_+$ depending only on $\mathcal{B}, \rho, \delta$ such that the following holds. For any $\mu \in \mathcal{MF}$, the sequence $i \mapsto \operatorname{len}_i(\mu)$, $(i \in J)$, satisfies the L–lipschitz, $(\kappa, n, 0)$ flaring property.*

Proof From Lemma 2.2 we have a K–bilipschitz connection on $\mathcal{S}_\gamma$, and from Lemma 3.2 we have an L–bilipschitz connection on $\mathcal{GF}(\mathcal{S}_\gamma)$ whose connection lines project to λ–quasihorizontal lines in $\mathcal{S}_\gamma$. By lifting we obtain similar connections in $\mathcal{H}_\gamma$ and $\mathcal{GF}(\mathcal{H}_\gamma)$. The constants K, L, λ depend only on $\mathcal{B}$, ρ.

Define the *suspension* of μ to be the following 2–dimensional measured lamination in $\mathcal{S}_\gamma$:

$$\operatorname{Susp}(\mu) = \bigcup_{t \in \mathbf{R}} \mu_t$$

By Theorems 3.4 and 3.5, it follows that each leaf of $\operatorname{Susp}(\mu)$ is piecewise smooth, being uniformly smooth over each affine segment of γ. Restricting the projection $\mathcal{S}_\gamma \to \mathbf{R}$ to $\operatorname{Susp}(\mu)$, we may think of $\operatorname{Susp}(\mu)$ as a μ–bundle over I.

By restricting the connection on $\mathcal{GF}(\mathcal{S}_\gamma)$, we obtain a connection on $\operatorname{Susp}(\mu)$ whose connection lines respect the leaves of $\operatorname{Susp}(\mu)$ and are transverse to the fibers μ_t. By lifting to the universal cover $\mathcal{H}_\gamma$ of $\mathcal{S}_\gamma$ we obtain the suspension $\operatorname{Susp}(\widetilde{\mu})$ of $\widetilde{\mu}$, whose fiber in $\mathcal{H}_t$ is $\widetilde{\mu}_t$, and we obtain a connection on $\operatorname{Susp}(\widetilde{\mu})$. We will use h_{st} to denote either of the connection maps $\mu_s \to \mu_t$ or $\widetilde{\mu}_s \to \widetilde{\mu}_t$; the context should make the meaning clear. Note in particular that h_{st} is $L^{|s-t|}$ bilipschitz from leaves of $\widetilde{\mu}_s$ to leaves of $\widetilde{\mu}_t$, preserves transverse measure, and has connection lines which are λ quasihorizontal in $\mathcal{H}_\gamma$.

Since $\mathcal{H}_\gamma$ is δ–hyperbolic, it satisfies horizontal flaring with data depending only on $\mathcal{B}$, ρ, δ, and so we have:

Lemma 3.7 (Uniform Flaring) *There exist constants (κ, n, A) depending only on $\mathcal{B}$, ρ, δ with the following property. For any $\mu \in \mathcal{MF}$, and for any two connection lines α, α' of $\operatorname{Susp}(\widetilde{\mu})$, the sequence of distances $d_{\mathcal{H}_i}(\alpha(i), \alpha'(i))$ satisfy the L–lipschitz, (κ, n, A) flaring property. In particular, for any $\mu \in \mathcal{MF}$, any $s \in \mathbf{R}$, and any leaf segment ℓ of μ_s, the sequence of lengths $\operatorname{len}_{s+i}(h_{s,s+i}(\ell))$ satisfies the L–lipschitz, (κ, n, A) flaring property.* □

Remark Connection lines of $\operatorname{Susp}(\widetilde{\mu})$ are only quasihorizontal, not horizontal, so they do *not* necessarily coincide with connection lines in $\mathcal{H}_\gamma$.

Continuing the proof of Lemma 3.6, recall that:

$$\operatorname{len}_t(\mu) = \int_{\mu_t} d\mu_t = \int_{\mu_t} d\mu_t^{\perp} \times d\mu_t^{||}$$

Using this formula we can express the length change from $\operatorname{len}_s(\mu)$ to $\operatorname{len}_t(\mu)$ as an integral of a derivative. To be precise, as a measurable function on μ_s we have a well-defined Radon–Nykodym derivative:

$$h'_{st} = \frac{h^{-1}_{st*}(d\mu_t)}{d\mu_s} = \frac{h^{-1}_{st*}(d\mu_t^{||})}{d\mu_s^{||}}$$

It follows that:

$$\operatorname{len}_t(\mu) = \int_{\mu_t} d\mu_t = \int_{\mu_s} h'_{st}\, d\mu_s$$

For each $x \in \mu_s$ let $I_a(x)$ be the lamination segment of length a centered on x. By applying Fubini's theorem and using a change of variables, we get:

$$\operatorname{len}_t(\mu) = \int_{\mu_s} \left(\int_{\tau \in I_a(x)} \frac{h'_{st}(\tau)}{a}\, d\tau \right) d\mu_s(x)$$

where $d\tau$ is simply leafwise Lebesgue measure, that is, $d\tau = d\mu_s^{||}$. We rewrite this as:

$$\begin{aligned} \operatorname{len}_t(\mu) &= \int_{\mu_s} \frac{\operatorname{len}(h_{st}(I_a(x)))}{\operatorname{len}(I_a(x))}\, d\mu_s(x) \\ &= \int_{\mu_s} S_a^{t-s}(x)\, d\mu_s(x) \end{aligned}$$

where $S_a^r(x)$, the "stretch" of the segment $I_a(x)$ with displacement r, is defined to be:

$$S_a^r(x) = \frac{\operatorname{len}(h_{s,s+r}(I_a(x)))}{\operatorname{len}(I_a(x))} = \int_{\tau \in I_a(x)} \frac{h'_{s,s+r}(\tau)}{a}\, d\tau$$

Now apply this for $r = \pm n$, and we have two versions:

$$\operatorname{len}_{s+n}(\mu) = \int_{\mu_s} S_a^{+n}(x)\, d\mu_s(x)$$

and

$$\operatorname{len}_{s-n}(\mu) = \int_{\mu_s} S_a^{-n}(x)\, d\mu_s(x)$$

Applying Lemma 3.7, for each $x \in \mu_s$ the sequence $\operatorname{len}(h_{s,s+i}(I_a(x)))$ satisfies (κ, n, A) flaring. Taking $a = A$ it follows that for every $x \in \mu_s$, either $S_A^{-n}(x) \geq \kappa$ or $S_A^{+n}(x) \geq \kappa$. Now we can define two subsets:

$$\begin{aligned} \mu_s^+ &= \{x \in \mu_s \mid S_A^{+n}(x) \geq \kappa\} \\ \mu_s^- &= \{x \in \mu_s \mid S_A^{-n}(x) \geq \kappa\} \end{aligned}$$

Each of these subsets is measurable, and $\mu_s = \mu_s^+ \cup \mu_s^-$. It follows that one of the two sets μ_s^+, μ_s^- contains at least half of the total $d\mu_s$ measure (maybe they both do). Choose $\epsilon \in \{+, -\}$ so that μ_s^ϵ has more than half of the measure.

By increasing n to $\lceil n \log_\kappa 2 \rceil$ if necessary, we may assume that $\kappa > 2$. It follows that:

$$\begin{aligned} \operatorname{len}_{s+\epsilon n}(\mu) &\geq \int_{\mu_s^\epsilon} S_a^{\epsilon n}(x)\, d\mu_s(x) \\ &\geq \kappa \int_{\mu_s^\epsilon} d\mu_s \\ &\geq \frac{\kappa}{2} \operatorname{len}_s(\mu) \end{aligned}$$

Since $\kappa/2 > 1$, this proves the lemma. □

3.5 Growth of measured laminations

In addition to the objects fixed at the beginning of Section 3.4, until further notice we shall fix numbers $L \geq 1$, $\kappa > 1$, $n \in \mathbf{Z}_+$, and $A \geq 0$, depending only on $\mathcal{B}$, ρ, δ, such that the conclusions of Lemmas 3.6 and 3.7 both hold. In particular, for each $\mu \in \mathcal{MF}$ the sequence $\operatorname{len}_i(\mu)$, parameterized by $J = I \cap \mathbf{Z}$, is L–lipschitz and satisfies $(\kappa, n, 0)$ flaring.

Consider now any sequence ℓ_i, $(i \in J)$, satisfying $(\kappa, n, 0)$ flaring. Let $\kappa' = \kappa^{1/n}$, and recall inequalities 3.1 and 3.2. Given $i, j = i + n \in J$, either $\ell_j \geq \ell_i$ or $\ell_i \geq \ell_j$. In the case $\ell_j \geq \ell_i$ it follows from flaring that $\ell_{j+n} \geq \kappa \ell_j$, and inequalities 3.1 and 3.2 hold with $j_0 = j$ and $\epsilon = +1$. In the case $\ell_i \geq \ell_j$ it follows that $\ell_{i-n} \geq \kappa \ell_i$, and the same inequalities hold but with $j_0 = i$ and $\epsilon = -1$. It immediately follows that:

Proposition 3.8 *There exist a constant ω depending only on κ, n, L (and hence depending only on $\mathcal{B}$, ρ, δ) such that if $(\ell_i)_{i \in J}$ is an L–lipschitz sequence exhibiting $(\kappa, n, 0)$ flaring then the following hold.*

- *If ℓ_i has no minimum for $i \in J$ then J is infinite and ℓ_i approaches zero as $i \to -\infty$ or as $i \to +\infty$, but not both; in this case we say that ℓ_i achieves its minimum at $-\infty$ or at $+\infty$, respectively.*
- *If ℓ_i has a minimum for $i \in J$, then the smallest subinterval of J containing all minima, called the* trough *of ℓ_i, has length at most ω.*
- *The sequence grows exponentially as it moves away from the minima, in the following sense:*
  - *If $i_0 \in J$ is to the right of all minima then for $i \in J$, $i \geq i_0$ we have*
  $$\ell_i \geq L^{-n} \kappa'^{i-i_0} \ell_{i_0}$$
  and moreover if $i = i_0 + kn$ for $k \in \mathbf{Z}_+$ then we have $\ell_i \geq \kappa^n \ell_{i_0}$.
  - *If $i_0 \in J$ is to the left of all minima then for $i \in J$, $i \leq i_0$ we have*
  $$\ell_i \geq L^{-n} \kappa'^{i_0-i} \ell_{i_0}$$
  and moreover if $i = i_0 - kn$ for $k \in \mathbf{Z}_+$ then we have $\ell_i \geq \kappa^n \ell_{i_0}$.

For each $\mu \in \mathcal{MF}$, Proposition 3.8 applies to the sequence $\text{len}_i(\mu)$, $i \in J = I \cap \mathbf{Z}$. If $\text{len}_i(\mu)$ achieves its minimum at $\pm\infty$, then we say that μ is *realized at* $\pm\infty$. On the other hand, if the length sequence $\text{len}_i(\mu)$ achieves its minimum at a finite value then we say that μ is *realized* at that value; μ may be realized at several values, and the *trough* of μ is defined to be the trough of the length sequence $\text{len}_i(\mu)$. Since length is a homogeneous function on $\mathcal{MF}$, these concepts apply as well to elements of $\mathcal{PMF}$.

Corollary 3.9 *Each measured lamination consisting of a simple closed geodesic is realized at a finite value.*

Proof Since $\mathcal{B}$ is compact and γ is $\mathcal{B}$–cobounded, there exists $m > 0$ depending only on $\mathcal{B}$ such that for each $t \in I$, every simple closed geodesic in $\mathcal{S}_t$ has length $\geq m$. If c is a simple closed geodesic equipped with a transverse measure $r \in \mathbf{R}$, it follows that $\text{len}_i(c) \geq rm$ for all i, and so c cannot be realized at $+\infty$ or at $-\infty$. ◻

The next result shows that the position of realizability, either $-\infty$, a finite set, or $+\infty$, is a coarsely continuous function of $\mathcal{P}\mu \in \mathcal{PMF}$. It is a consequence of the fact that as λ varies in $\mathcal{MF}$, the length function $(\text{len}_i(\mu))_{i \in J}$ varies continuously in the topology of pointwise convergence.

Proposition 3.10 *There exists a constant $\Delta = \Delta(\kappa, n, L)$ such that the following holds for each $\mathcal{P}\mu \in \mathcal{PMF}$.*

(1) *If $\mathcal{P}\mu$ has trough $W \subset \mathbf{Z}$ then there is a neighborhood $U \subset \mathcal{PMF}$ of $\mathcal{P}\mu$ such that each $\mathcal{P}\mu' \in U$ has a trough W' with* $\operatorname{diam}(W \cup W') \le \Delta$.

(2) *If $\mathcal{P}\mu$ is realized at $+\infty$ then for each $i_0 \in J$ there exists a neighborhood $U \subset \mathcal{PMF}$ of $\mathcal{P}\mu$ such that if $\mathcal{P}\mu' \in U$ then either $\mathcal{P}\mu'$ is realized at $+\infty$ or $\mathcal{P}\mu'$ is finitely realized with trough contained in $[i_0, +\infty)$.*

(3) *If $\mathcal{P}\mu$ is realized at $-\infty$ then a similar statement holds.*

Proof By homogeneity of length it suffices to prove the analogous statement for each $\mu \in \mathcal{MF}$.

From the argument preceding the statement of Proposition 3.8, we immediately have the following:

Lemma 3.11 *There exists a constant $\zeta \in \mathbf{Z}_+$, depending only on κ, n, L (and so only on $\mathcal{B}$, ρ, δ) such that if ℓ_i $(i \in J)$ is L–lipschitz and $(\kappa, n, 0)$ flaring then*

- *if $j_0, j_0 + n \in J$ and if $\ell_{j_0+n} \ge \ell_{j_0}$ then all minima of ℓ_i lie to the left of $j_0 + \zeta$.*
- *if $j_0, j_0 - n \in J$ and if $\ell_{j_0-n} \ge \ell_{j_0}$ then all minima of ℓ_i lie to the right of $j_0 - \zeta$.* □

To prove (1), let $W = [k, l]$ be the trough of μ. If $k - n \in J$ then $\operatorname{len}_{k-n}(\mu) > \operatorname{len}_k(\mu)$, and if $l + n \in J$ then $\operatorname{len}_{l+n}(\mu) > \operatorname{len}_l(\mu)$. By continuity of $\operatorname{len}\colon \mathcal{T} \times \mathcal{MF} \to (0,\infty)$ we may choose a neighborhood $U \subset \mathcal{MF}$ of μ so that if $\mu' \in U$ then $\operatorname{len}_{k-n}(\mu') > \operatorname{len}_k(\mu')$ and $\operatorname{len}_{l+n}(\mu') > \operatorname{len}_l(\mu')$. It then follows that the trough of μ' is a subset of the interval $[k - \zeta, l + \zeta]$, so (1) is proved with $\Delta = \omega + 2\zeta$.

To prove (2), assuming μ is realized at $+\infty$ it follows that $\operatorname{len}_{i_0+\zeta}(\mu) \ge \kappa \operatorname{len}_{i_0+\zeta+n}(\mu)$, and so we may choose U so that if $\mu' \in U$ we have $\operatorname{len}_{i_0+\zeta}(\mu') > \operatorname{len}_{i_0+\zeta+n}(\mu')$. It follows that all minima of $\operatorname{len}_i(\mu')$ lie to the right of i_0.

The proof of (3) is similar. □

3.6 Construction of ending laminations

We now construct laminations which are realized nearly anywhere one desires, in particular laminations realized at any infinite ends of J. Recall that a measured geodesic lamination is *perfect* if it has no isolated leaves, or equivalently if it has no closed leaves.

Proposition 3.12 *There exists a constant η depending only on $\mathcal{B}$, ρ, δ such that the following holds. For each $k \in J$ there exists $\mu \in \mathcal{MF}$ which is finitely realized and whose trough W satisfies* $\mathrm{diam}(W \cup \{k\}) \le \eta$. *If J is infinite then for each infinite end $\pm\infty$ of J there exists $\mu^\pm \in \mathcal{MF}$ which is realized at $\pm\infty$, respectively; moreover any such $\mu^\pm$ is perfect.*

A lamination $\mu^\pm$ realized at an infinite end $\pm\infty$ is called an *ending lamination* of $\mathcal{S}_\gamma$. Also, for any finite endpoint $k \in J$, we use the term *endpoint lamination* to refer to a lamination μ whose length function $\mathrm{len}_i(\mu)$ has a minimum occuring with distance η of the endpoint k; an alternate definition would require the entire trough of μ to lie within distance η of k, but this does not work out as well, as noted in the remark preceding Proposition 3.18.

Proof As in Corollary 3.9, using compactness of $\mathcal{B}$ and $\mathcal{B}$–coboundedness of γ, there exists $m > 0$ depending only on $\mathcal{B}$ such that for each $t \in I$, every simple closed geodesic in $\mathcal{S}_t$ has length $\ge m$. There also exists $M > 0$ depending only on the topology of S such that for each hyperbolic structure on S the shortest geodesic has length $\le M$; this standard fact follows because the area of any hyperbolic structure is equal to $2\pi\chi(S)$, and if the shortest geodesic had arbitrarily large length then it would have an annulus neighborhood with arbitrarily large area, violating the Gauss–Bonnet theorem.

Given $k \in J$, take a simple closed geodesic c of minimal length in $\mathcal{S}_k$, with the transverse Dirac measure. By Corollary 3.9, c is finitely realized. Consider the subsequence $\mathrm{len}_{k+np}(c)$, and let $p = P$ be the value where it achieves its minimum. Since $\mathrm{len}_{k+nP}(c) \ge m$ and $\mathrm{len}_k(c) \le M$, from $(\kappa, n, 0)$ flaring it follows that $|P| < \left\lceil \log_\kappa \frac{M}{m} \right\rceil$, and so by Lemma 3.11 the trough of c must be located within the interval

$$\left[k - n \left\lceil \log_\kappa \frac{M}{m} \right\rceil - \zeta \ , \ k + n \left\lceil \log_\kappa \frac{M}{m} \right\rceil + \zeta \right]$$

and so we may take $\eta = 2n \left\lceil \log_\kappa \frac{M}{m} \right\rceil + 2\zeta$, proving the first part of the proposition.

Consider now an infinite end of J, say, $+\infty$. For each i choose $\mu_i \in \mathcal{MF}$ to be finitely realized, with trough W_i satisfying $\mathrm{diam}(W_i \cup \{i\}) \le \eta$. Using compactness of $\mathcal{PMF}$, choose $\mu^+ \in \mathcal{MF}$ so that, after passing to a subsequence, $\mathcal{P}\mu_i \to \mathcal{P}\mu^+$ as $i \to +\infty$.

The fact that μ^+ is realized at $+\infty$ is a consequence of the fact that the sequence of length functions $(\mathrm{len}_j(\mu_i))_{j \in J}$ converges pointwise to the length function $(\mathrm{len}_j(\mu^+))_{j \in J}$. To be precise, suppose first that μ^+ is finitely realized with trough $W \subset \mathbf{Z}$. Applying Proposition 3.10 it follows that there is a neighborhood U of $\mathcal{P}\mu^+$ and a larger interval $W' \subset \mathbf{Z}$ such that if $\mathcal{P}\mu' \in U$ then $\mathcal{P}\mu'$ is finitely realized with trough contained in W'. But $\mathcal{P}\mu_i \in U$ for sufficiently large i, and its trough W_i goes off to $+\infty$ as $i \to +\infty$, a contradiction. Suppose next that μ^+ is realized at $-\infty$. By Proposition 3.10 it follows that the trough of $\mathcal{P}\mu_i$ goes to $-\infty$ as $i \to +\infty$, also a contradiction.

The construction of μ^- realized at $-\infty$ is similar.

If $\mu^\pm$ were not perfect it would have a closed leaf c with transverse measure $r \in \mathbf{R}$, but then it would follow that $\mathrm{len}_i(\mu^\pm) \geq rm$ for all i, contradicting that $\mathrm{len}_i(\mu^\pm) \to 0$ as $i \to \pm\infty$. ◻

3.7 Strict decay of ending laminations

In this section and the next we concentrate on properties of ending laminations associated to infinite ends of $\mathcal{S}_\gamma$. The technical lemma 3.13 proved in this section is applied to obtain filling properties for ending laminations. As a consequence, at the end of section 3.8, we will describe the construction of the desired Teichmüller geodesic g in the case where γ is infinite.

Let $\mu^\pm \in \mathcal{MF}$ be an ending lamination realized at an infinite end $\pm\infty$ of J. Let $\mu_i^\pm$ denote the measured geodesic lamination on $\mathcal{S}_i$ representing $\mu^\pm$. We prove a strict decay property for μ^+, say: in any leaf of $\mathrm{Susp}(\mu^+)$, any two connection lines which are sufficiently far apart in that leaf at level i decay exponentially *immediately* in the positive direction—there is no growth *anywhere* in the lamination μ^+ as one approaches $+\infty$, except on uniformly short segments. A similar statement holds for μ^-, flowing in the negative direction along connection lines. We make this precise as follows.

Let $h_{st}^\pm$ be the connection maps on $\mathrm{Susp}(\mu^\pm)$.

Lemma 3.13 *If ℓ is a leaf segment of μ_i^+ with $\mathrm{len}\,\ell \geq A\kappa$ then*

$$\mathrm{len}(h_{i,i+n}^+(\ell)) \leq \frac{1}{\kappa}\,\mathrm{len}\,\ell.$$

Similarly, if ℓ is a leaf segment of μ_i^- with $\mathrm{len}\,\ell \geq A\kappa$ then

$$\mathrm{len}(h_{i,i-n}^-(\ell)) \leq \frac{1}{\kappa}\,\mathrm{len}\,\ell.$$

Proof Borrowing notation from Lemma 3.6, given $x \in \mu_i^+$ let $I_a(x)$ be the leaf segment of μ_i^+ of length a centered on x, and let $S_a^r(x)$ be the stretch of the segment $I_a(x)$ with displacement r, that is:

$$S_a^r(x) = \frac{\mathrm{len}(h_{i+r}^+ I_a(x))}{\mathrm{len}(I_a(x))}$$

The lemma says that $S_a^n(x) \le \frac{1}{\kappa}$ if $a \ge A\kappa$.

If there exists $x \in \mu_i^+$ and $a \ge A\kappa$ such that $S_a^n(x) > \frac{1}{\kappa}$, then letting y be the midpoint of $h_{i,i+n}^+(I_a(x))$, and taking $a' = a \cdot S_a^n(x) > A$ we have $S_{a'}^{-n}(y) < \kappa$. By changing variables it therefore suffices to prove that for all $a \ge A$ and all $x \in \mu_i^+$ we have $S_a^{-n}(x) \ge \kappa$.

Suppose there exists $x \in \mu_i^+$ and $a \ge A$ such that $S_a^{-n}(x) < \kappa$. Applying Lemma 3.7 we conclude that $S_a^n(x) \ge \kappa$. Now $S_a^n(y)$ is a continuous function of $y \in \mu_i^+$ and it follows that there is a neighborhood $U \subset \mu_i^+$ of x such that if $y \in U$ then $S_a^n(y) > 1$. Given $y \in U$, again applying Lemma 3.7 it follows by induction on p that $S_a^{np}(y) \ge \kappa^{p-1}$ for all $p \ge 1$. But U has positive measure $\int_U d\mu_i^+$ in μ_i^+, and so we have

$$\begin{aligned}
\mathrm{len}_{i+np}(\mu^+) &= \int_{\mu_i^+} S_a^{np}(y) d\mu_i^+(y) \\
&\ge \int_U S_a^{np}(y) d\mu_i^+(y) \\
&\ge \kappa^{p-1} \int_U d\mu_i^+ \\
&\to +\infty \quad \text{as} \quad p \to +\infty
\end{aligned}$$

contradicting that $\mathrm{len}_{i+np}(\mu^+) \to 0$ as $p \to +\infty$. □

Recall that two points $\mu, \nu \in \mathcal{MF}$ are *topologically equivalent* if they are represented by measured foliations which have the same underlying nonmeasured foliation. Equivalently, for any hyperbolic structure on S, the straightenings of μ, ν have the same underlying nonmeasured geodesic lamination.

Corollary 3.14 *If J is bi-infinite and if μ^-, μ^+ are the ending laminations realized at $-\infty$, $+\infty$ respectively, then μ^- and μ^+ are not topologically equivalent.*

Proof Suppose they are topologically equivalent, and so on the surface $\mathcal{S}_0$ the laminations μ_0^- and μ_0^+ have the same underlying nonmeasured geodesic lamination. Let ℓ be any leaf segment of this lamination with $\mathrm{len}(\ell) > A\kappa$. Applying Lemma 3.13 twice, from μ_0^+ we conclude that $\mathrm{len}(h_{0,n}(\ell)) < \frac{1}{\kappa}\,\mathrm{len}(\ell)$, and

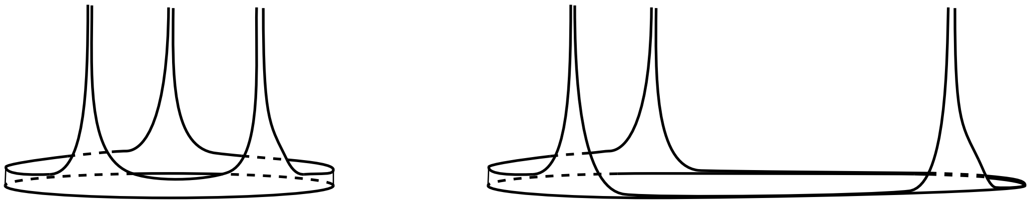

Figure 1: If c is a simple closed geodesic in S which is disjoint from a nonfilling geodesic lamination μ, and if c is peripheral in $S-\mu$, then $S-(\mu\cup c)$ has a component E whose metric completion $\overline{E}$ is a crown. For each $\epsilon>0$, if c is sufficiently long then the set of points in c that are within distance ϵ of $\partial\overline{E}-c$ consists of at least $1-\epsilon$ of the total length of c.

from μ_0^- we conclude that $\mathrm{len}(h_{0,-n}(\ell))\le\frac{1}{\kappa}\mathrm{len}(\ell)$. However, from Lemma 3.7 at least one of $\mathrm{len}(h_{0,n}(\ell))$, $\mathrm{len}(h_{0,-n}(\ell))$ is $\ge\kappa\,\mathrm{len}(\ell)$, a contradiction. □

3.8 Individual filling of the ending laminations

Recall that $\mu\in\mathcal{MF}$ *fills* the surface S if μ has nonzero intersection number with every simple closed curve. Equivalently, for any hyperbolic structure, the realization of μ as a measured lamination has simply connected complementary components. Note that a filling geodesic lamination is necessarily perfect.

Proposition 3.15 *Any ending lamination $\mu^\pm$ fills S.*

Proof of Proposition 3.15 Arguing by contradiction, suppose that, say, μ^+ does not fill S. Consider the straightening μ_0^+ in $\mathcal{S}_0$. Let F be a component of $\mathcal{S}_0-\mu_0^+$ which is not simply connected. Let c be a simple closed geodesic which is peripheral in F. Let E be a component of $F-c$ which is a neighborhood of an end of F. The metric completion $\overline{E}$ is a "crown" surface (see Figure 1), i.e. a complete hyperbolic surface with geodesic boundary homeomorphic to an annulus with $Q\ge 1$ "crown points" removed from one of the boundary components of the annulus. Each removed crown point has a neighborhood isometric to the region in $\mathbf{H}^2$ bounded by two geodesics with a common ideal endpoint in $\partial\mathbf{H}^2$ and a horocycle attached to that endpoint. The compact boundary component of $\overline{E}$ is c. The rest of the boundary $\partial\overline{E}-c$ consists of Q components, each isometric to the real line, each identified with a leaf of μ_0^+.

Let c_i be the straightening of c in S_i, let F_i be the component of $S_i-\mu_i^+$ containing c_i, and let E_i be the component of F_i-c_i corresponding to E_0.

The metric completion $\overline{E}_i$ is also a crown surface with Q crown points, with one compact boundary component c_i, and with $\partial\overline{E}_i - c_i$ consisting of Q noncompact boundary components each identified with a leaf of μ_i^+.

By Corollary 3.9, c is finitely realized, and so $\text{len}_i(c) \to \infty$ as $i \to +\infty$. It follows that c_i has longer and longer subsegments which are closer and closer to subsegments of the opposite boundary of $\overline{E}_i$ (see Figure 1). To be precise:

Claim 3.16 *For each $\epsilon > 0$ there exists i_0 such that if $i \ge i_0$ then the set $c_i^\epsilon = \{x \in c_i \mid d(x, \partial\overline{E}_i - c_i) < \epsilon\}$ consists of at most Q segments, and $\text{len}(c_i^\epsilon) \ge (1-\epsilon)\,\text{len}(c_i)$. If ϵ is sufficiently small then each component of c_i^ϵ is within distance ϵ of at most one component of $\partial\overline{E}_i - c_i$.*

To see why, the hyperbolic surfaces $\overline{E}_i$ have constant topology, and therefore they have constant area, by the Gauss–Bonnet theorem. It follows that $\text{len}(c_i - c_i^\epsilon)$ is bounded above, otherwise the ϵ–neighborhood of $c_i - c_i^\epsilon$ would have too much area. Also, as long as ϵ is sufficiently small, if $[x,y] \subset c_i$ is a segment such that x is within ϵ of one component of $\partial\overline{E}_i - c_i$ and y is within ϵ of a different component, then there must be a point in $[x,y]$ which has distance $> \epsilon$ from $\partial\overline{E}_i - c_i$, and so x, y are in different components of c_i^ϵ. In other words, each component of c_i^ϵ is within distance ϵ of only one of the Q components of $\partial\overline{E}_i - c_i$.

The idea of the proof of Proposition 3.15 is that c_i is growing exponentially, whereas long leaf segments of μ_i are shrinking exponentially, and since most of c_i is very close in the tangent line bundle to μ_i this leads to a contradiction. Now we make this precise.

As a consequence of the claim, since $\text{len}(c_i^\epsilon) + \text{len}(c - c_i^\epsilon) = \text{len}(c_i)$, we have

$$\text{len}(c - c_i^\epsilon) \le \frac{\epsilon}{1-\epsilon}\,\text{len}(c_i^\epsilon).$$

Now choose a very small $\epsilon > 0$, and choose i so large that, listing the path components of c_i^ϵ as $\alpha_1, \ldots, \alpha_K$, there are corresponding arcs $\beta_1, \ldots, \beta_K \subset \partial\overline{E}_i - c_i \subset \mu_i^+$, such that for $i = 1, \ldots, K$, the arcs α_k, β_k are ϵ–fellow travellers, and $\text{len}(\beta_k) \ge A\kappa$. Let $\alpha_k' = H_{i,i+n}(\alpha_k) \subset c_{i+n}$, $\beta_k' = H_{i,i+n}(\beta_k) \subset \mu_{i+n}^+$ where $H_{i,i+n}$ is the connection map on geodesic laminations, or equivalently, the connection on the geodesic foliation bundle. Applying Lemma 3.13, we have

$$\text{len}(\beta_k) \ge \kappa\,\text{len}(\beta_k').$$

If ϵ is sufficiently small, each of the pairs of segments α_k, β_k is arbitrarily close when lifted to the tangent line bundle, and so by the uniform continuity property of $H_{i,i+n}$ (see Lemma 3.2) they stretch by very nearly equal amounts:

$$\frac{\mathrm{len}(\alpha'_k)}{\mathrm{len}(\alpha_k)} \bigg/ \frac{\mathrm{len}(\beta'_k)}{\mathrm{len}(\beta_k)} \le 1+\eta$$

for any given $\eta > 0$. We therefore have:

$$\begin{aligned}
\mathrm{len}(c_{i+n}) &= \mathrm{len}(H_{i,i+n}(c_i^\epsilon)) + \mathrm{len}(H_{i,i+n}(c - c_i^\epsilon)) \\
&\le \mathrm{len}(H_{i,i+n}(c_i^\epsilon)) + L^n \,\mathrm{len}(c - c_i^\epsilon) \\
&\le \mathrm{len}(H_{i,i+n}(c_i^\epsilon)) + L^n \frac{\epsilon}{1-\epsilon}\,\mathrm{len}(c_i^\epsilon) \\
&\le \left(1 + \frac{\epsilon L^{2n}}{1-\epsilon}\right) \mathrm{len}(H_{i,i+n}(c_i^\epsilon)) \\
&= \left(1 + \frac{\epsilon L^{2n}}{1-\epsilon}\right) \sum_1^K \mathrm{len}(\alpha'_k) \\
&\le \left(1 + \frac{\epsilon L^{2n}}{1-\epsilon}\right) (1+\eta) \sum_1^K \mathrm{len}(\alpha_k) \cdot \frac{\mathrm{len}(\beta'_k)}{\mathrm{len}(\beta_k)} \\
&\le \left(1 + \frac{\epsilon L^{2n}}{1-\epsilon}\right) \frac{1+\eta}{\kappa} \sum_1^K \mathrm{len}(\alpha_k) \\
&\le \left(1 + \frac{\epsilon L^{2n}}{1-\epsilon}\right) \frac{1+\eta}{\kappa}\,\mathrm{len}(c_i) \\
&\le \left(1 + \frac{\epsilon L^{2n}}{1-\epsilon}\right) \frac{1+\eta}{\kappa^2}\,\mathrm{len}(c_{i+n})
\end{aligned}$$

where the last inequality follows from the fact that $\mathrm{len}(c_{i+n}) \ge \kappa\,\mathrm{len}(c_i)$ for sufficiently large i. When η and ϵ are sufficiently small, the multiplicative constant is arbitrarily close to $1/\kappa^2$, and we obtain a contradiction. □

When γ is bi-infinite we are now in a position to construct the desired Teichmüller geodesic g. For any topologically inequivalent pair of laminations, if at least one of them fills, then the pair jointly fills. We therefore have:

Corollary 3.17 *If γ is bi-infinite then any choice of ending laminations for the two ends jointly fills S.* □

In the bi-infinite case, with ending laminations λ_1, λ_2, we can therefore define a Teichmüller geodesic line $g = \overleftrightarrow{(\mathcal{P}\lambda_1, \mathcal{P}\lambda_2)}$.

In the half-infinite case, associated to the infinite end there is an ending lamination λ_1 which fills, and associated with the finite end there is an endpoint lamination λ_2 which is topologically inequivalent to λ_1, and hence the pair λ_1, λ_2 jointly fills. Unfortunately we cannot yet prove, when γ is finite, that any pair of endpoint laminations jointly fills—indeed it is not always true without an extra condition.

3.9 A compactness property

In this section we prove a compactness property for ending laminations and endpoint laminations associated to cobounded, lipschitz paths in Teichmüller space whose associated hyperbolic plane bundle is a hyperbolic metric space. This will be used in the following section in two key ways: to prove the desired filling property for finite paths; and to prove Theorem 1.1.

For the last several sections we have been fixing a particular path γ, but now we want to let γ vary and investigate convergence of the various pieces of geometric data we have been studying.

Recall that we have fixed a compact subset $\mathcal{B} \subset \mathcal{M}$ and numbers $\rho \ge 1$, $\delta \ge 0$. We also fix a constant $\eta \ge 0$ satisfying the conclusions of Proposition 3.12, in particular each endpoint lamination is realized somewhere within distance η of the endpoint.

Let $\Gamma_{\mathcal{B},\rho,\delta,\eta}$ be the set of all triples $(\gamma, \lambda_-, \lambda_+)$ with the following properties:

(1) $\gamma\colon I \to \mathcal{T}$ is a $\mathcal{B}$–cobounded, ρ–lipschitz, $\mathbf{Z}$–piecewise affine path such that $\mathcal{H}_\gamma$ is δ–hyperbolic,

(2) $0 \in I$, and each of $\lambda_\pm \in \mathcal{MF}$ is normalized to have length 1 in the hyperbolic structure $\gamma(0)$,

(3) The lamination λ^+ is realized in $\mathcal{S}_\gamma$ near the right end, in the following sense:

    (a) If γ is right infinite then λ^+ is realized at $+\infty$.

    (b) If γ is right finite, with right endpoint M, then there exists a minimum of the length sequence $\operatorname{len}_i(\lambda^+)$ lying in the interval $[M-\eta, M]$.

    The lamination λ^- is similarly realized in $\mathcal{S}_\gamma$ near the left end.

We give $\Gamma_{\mathcal{B},\rho,\delta,\eta}$ the product topology, using the usual topology on $\mathcal{MF}$ for the second and third coordinates λ_-, λ_+, and for the γ coordinate we use the compact–open topology. Since the domain interval I may vary, we apply the compact–open topology to the unique, continuous extension $\gamma\colon \mathbf{R} \to \mathcal{T}$ which is constant on each component of $\mathbf{R} - I$.

Remark The existential quantifier in item 3b above is important. In the following proposition, the proof in case 2 would fall apart if item 3b were replaced, say, by the statement that the entire trough of λ^+ lies in the interval $[M-\eta, M]$.

Proposition 3.18 *The action of $\mathcal{MCG}$ on $\Gamma_{\mathcal{B},\rho,\delta,\eta}$ is cocompact.*

Proof Choose a compact subset $\mathcal{A} \subset \mathcal{T}$ such that each $(\gamma, \lambda_-, \lambda_+) \in \Gamma_{\mathcal{B},\rho,\delta,\eta}$ may be translated by the action of $\mathcal{MCG}$ so that

(4) $\gamma(0) \in \mathcal{A}$.

It suffices to prove that the set of $(\gamma, \lambda_-, \lambda_+)$ satisfying (1), (2), (3), and (4) is compact.

By the Ascoli–Arzela theorem, the set of ρ–lipschitz, $\mathbf{Z}$–piecewise affine paths $\gamma \colon I \to \mathcal{T}$ with $\gamma(0) \in \mathcal{A}$ is compact (this is where we use $\mathbf{Z}$–piecewise affine). The subset of those which are $\mathcal{B}$–cobounded is a closed subset, since $\mathcal{B}$ is closed. The subset of those for which $\mathcal{H}_\gamma$ is δ–hyperbolic is closed, because if γ_i converges to γ then $\mathcal{H}_{\gamma_i}$ converges to $\mathcal{H}_\gamma$ in the Gromov–Hausdorff topology, and for fixed δ the property of δ–hyperbolicity is closed in the Gromov–Hausdorff topology [12].

So far we have we have shown that the set of triples satisfying (1) and (4) is compact, and since the length function $\mathcal{T} \times \mathcal{MF} \to (0,\infty)$ is continuous is follows that the set satisfying (1), (2), and (4) is compact. It remains to show that the subset of those satisfying (3) is closed. Let $(\gamma_i, \lambda_i^-, \lambda_i^+)$ be a sequence satisfying (1–4) and converging to a limit $(\gamma, \lambda^-, \lambda^+)$, necessarily satisfying (1), (2), and (4). Let I_i be the domain of γ_i, and I the domain of γ. We must verify (3), and we focus on the proof for λ^+, which will be a consequence of the continuity of length functions. The detailed proof is broken into cases depending on the nature of the positive ends of the domain intervals.

Case 1: I is positive infinite We must prove that λ^+ is realized at $+\infty$ in $\mathcal{S}_\gamma$. If not, then it is realized finitely or at $-\infty$; pick $i_0 \in I$ so that λ^+ is realized to the left of i_0, either at $-\infty$ or with trough to the left of i_0. By Proposition 3.10 there is a $\Delta > 0$, depending only on $\mathcal{B}$, ρ, δ, such that if i is sufficiently large then λ_i^+ is realized to the left of $i_0 + \Delta$. If there exist arbitrarily large i for which I_i is positive infinite then λ_i^+ is realized at $+\infty$, an immediate contradiction. On the other hand, if I_i is positive finite for all sufficiently large i, with right endpoint M_i, then the endpoint lamination λ_i^+ is realized to the right of $M_i - \eta$, but M_i diverges to $+\infty$ and so eventually λ_i^+ is not realized to the left of $i_0 + \Delta$, also a contradiction.

Case 2: I has finite right endpoint M It follows that for sufficiently large i, the interval I_i also has finite right endpoint M, and so each lamination λ_i is realized at some point in the interval $[M-\eta, M]$. By continuity of the length function $\mathcal{T} \times \mathcal{MF} \to (0,\infty)$ it follows that λ^+ is also realized at some point in this interval.

□

3.10 Proof of Theorem 1.1

At the end of section 3.8 we used results about filling to construct the desired Teichmüller geodesic g in the case where the domain I of the path γ is a line. In the case where I is a segment we need the following:

Proposition 3.19 *There exists a constant Λ, depending only on $\mathcal{B}$, ρ, δ, η such that if $I=[m,n]$ is a finite segment with $n-m \ge \Lambda$, and if $(\gamma,\lambda^-,\lambda^+) \in \Gamma_{\mathcal{B},\rho,\delta,\eta}$ with $\gamma\colon I \to \mathcal{T}$, then the λ^-,λ^+ jointly fills S.*

Proof If there exists no such constant Λ, then there is a sequence of examples $(\gamma_i,\lambda_i^-,\lambda_i^+) \in \Gamma_{\mathcal{B},\rho,\delta,\eta}$ with $\gamma_i\colon I_i \to \mathcal{T}$, such that $\mathrm{len}(I_i) \to \infty$, and the pair λ_i^-,λ_i^+ does not jointly fill. After translating the parameter interval I_i we may assume that 0 lies within distance $1/2$ of the midpoint of I_i. After acting appropriately by elements of $\mathcal{MCG}$, we may assume that the sequence $(\gamma_i,\lambda_i^-,\lambda_i^+)$ converges to $(\gamma,\lambda^-,\lambda^+) \in \Gamma_{\mathcal{B},\rho,\delta,\eta}$, and it follows that γ has domain $\mathbf{R}$. By Corollary 3.17, the pair λ^-,λ^+ jointly fills. However, the set of jointly filling pairs in $\mathcal{MF}\times\mathcal{MF}$ is an open subset $\mathcal{FP}$, and so for sufficiently large i the pair λ_i^-,λ_i^+ jointly fills, a contradiction. □

Now we turn to the proof of Theorem 1.1.

Consider a $\mathcal{B}$–cobounded, ρ–lipschitz, $\mathbf{Z}$–piecewise affine path $\gamma\colon I \to \mathcal{T}$ such that $\mathcal{H}_\gamma$ is δ–hyperbolic. By translating the interval I we may assume $0 \in I$. Choose η satisfying Proposition 3.12, and it follows that there are $\lambda_-,\lambda_+ \in \mathcal{MF}$ such that $(\gamma,\lambda_-,\lambda_+) \in \Gamma_{\mathcal{B},\rho,\delta,\eta}$. Recall that λ^-,λ^+ are normalized to have length 1 in the hyperbolic structure $\gamma(0)$.

Fix a constant Λ so that Proposition 3.19 is satisfied.

First we knock off the case where $\gamma\colon I \to \mathcal{T}$ satisfies $\mathrm{len}(I) < \Lambda$. In this case $\mathrm{len}(\gamma) < \rho\Lambda$. Let g be the geodesic segment with the same endpoints as γ, and so $\mathrm{len}(g) \le \mathrm{len}(\gamma) < \rho\Lambda$. It follows that the Hausdorff distance between image(γ) and g is at most $\rho\Lambda$. Also, any ρ–lipschitz segment of length $\le \rho\Lambda$ is a $(1, C)$–quasigeodesic with $C = \max\{1, \rho\Lambda\}$.

We may henceforth assume that $\mathrm{len}(I) \ge \Lambda$.

Now we define the geodesic g. For each infinite end of I we have associated an ending lamination, which determines the corresponding infinite end of g; in particular, in the case where I is bi-infinite we have already defined $g = \overleftrightarrow{(\mathcal{P}\lambda^-, \mathcal{P}\lambda^-)}$. In the case where I is half-infinite or finite, we also have a jointly filling pair λ^-, λ^+ and so we have a geodesic line $g^* = \overleftrightarrow{(\mathcal{P}\lambda^-, \mathcal{P}\lambda^+)}$. We must specify a ray or segment on g^*, and even in the bi-infinite case we must specify how the path γ is synchronized with this ray or segment. These tasks are accomplished as follows.

Recall the notation $\sigma(\lambda, \lambda')$ and $q(\lambda, \lambda')$ for the conformal structure and quadratic differential determined by a jointly filling pair $\lambda, \lambda' \in \mathcal{MF}$. For each $t \in I$ we define:

$$a^-(t) = \frac{1}{\mathrm{len}_t \lambda^-}, \quad a^+(t) = \frac{1}{\mathrm{len}_t \lambda^+}$$

These are continuous functions of $t \in I$, and it follows that we have a continuous function $\Sigma\colon I \to g^*$ defined as follows:

$$\Sigma(t) = \sigma(a^-(t)\lambda^-, a^+(t)\lambda^+)$$

The image of this map is therefore a connected subset of g^* whose closure is the desired geodesic g (we will in fact show that image(Σ) is closed). We also have a continuous family of quadratic differentials $Q\colon I \to \mathrm{QD}$ defined by

$$Q(t) = q(a^-(t)\lambda^-, a^+(t)\lambda^+)$$

where $Q(t)$ is a quadratic differential on the Riemann surface $\Sigma(t)$.

Next we prove that the Teichmüller distance between $\gamma(t)$ and $\Sigma(t)$ is bounded above, by a constant depending only on $\mathcal{B}$, ρ, δ (and η, which depends in turn on $\mathcal{B}$, ρ, δ). For integer values $t = i$ this follows from the compactness result, Proposition 3.18. To see why, defining $\gamma'(s) = \gamma(s + i)$, the ordered triple $(\gamma', a^-(i)\lambda^-, a^+(i)\lambda^+)$ lies in the $\mathcal{MCG}$–cocompact set $\Gamma_{\mathcal{B},\rho,\delta,\eta}$. The map taking $(\gamma', \lambda'^-, \lambda'^+) \in \Gamma_{\mathcal{B},\rho,\delta,\eta}$ to $(\gamma'(0), \sigma(\lambda'^-, \lambda'^+)) \in \mathcal{T} \times \mathcal{T}$ is continuous and $\mathcal{MCG}$–equivariant, and therefore has $\mathcal{MCG}$–cocompact image, and hence the distance function is bounded above as required. If t is not an integer, there exists an integer i such that $|t - i| \le 1$, and recalling the lipschitz constant L for length functions $t \mapsto \mathrm{len}_t(\lambda)$ it follows that

$$\left|\log(a^-(i)/a^-(t))\right|, \left|\log(a^+(i)/a^+(t))\right| \le \log(L)$$

and so $\Sigma(t)$ and $\Sigma(i)$ have Teichmüller distance bounded solely in terms of L (which depends only on $\mathcal{B}$, ρ). Also, $\gamma(t)$ and $\gamma(i)$ have Teichmüller distance at most ρ.

Our final task is to prove that $\gamma\colon I \to \mathcal{T}$ is a quasigeodesic. Since $d(\gamma(t), \Sigma(t))$ is bounded in terms of $\mathcal{B}$, ρ, δ, it suffices to prove that the map $\Sigma\colon I \to \mathcal{T}$ is a quasigeodesic, with constants depending only on $\mathcal{B}$, ρ, δ. Of course the image of Σ is contained in the Teichmüller geodesic g^*, but Σ does not have the geodesic parameterization, which it would have had if we had taken $a^-(t) = e^{-t}$, $a^+(t) = e^t$. Instead, the geodesic parameterization was sacrificed, and $a^-(t)$, $a^+(t)$ were chosen to guarantee synchronization of Σ and γ, that is, so that $d(\Sigma(t), \gamma(t))$ is bounded independent of t. So, even though Σ is not geodesically parameterized, we can nevertheless show that Σ is a quasigeodesic.

Using the fact that image(Σ) is contained in the geodesic g, we can obtain an exact formula for $d(\Sigma(s), \Sigma(t))$, as follows. Note that $Q(t)$ is not necessarily normalized so that $\|Q(t)\| = 1$, but we have:

$$\begin{aligned}\|Q(t)\| &= \left\|q(a^-(t)\lambda^-, a^+(t)\lambda^+)\right\| \\ &= a^-(t)a^+(t)\left\|q(\lambda^-, \lambda^+)\right\| \\ &= a^-(t)a^+(t)\left\|Q(0)\right\|\end{aligned}$$

The ordered pair of measured laminations $a^-(t)\lambda^-$, $a^+(t)\lambda^+$ can be normalized by dividing each of them by $\sqrt{\|Q(t)\|}$, which does not affect $\Sigma(t)$:

$$\begin{aligned}\Sigma(t) &= \sigma\left(\frac{a^-(t)}{\sqrt{\|Q(t)\|}}\,\lambda^-, \frac{a^+(t)}{\sqrt{\|Q(t)\|}}\,\lambda^+\right) \\ &= \sigma\left(\sqrt{\frac{a^-(t)}{a^+(t)}}\,\lambda^-, \sqrt{\frac{a^+(t)}{a^-(t)}}\lambda^+\right)\end{aligned}$$

It follows that for $s, t \in I$ we have:

$$\begin{aligned}d(\Sigma(s), \Sigma(t)) &= \frac{1}{2}\left|\log\left(\frac{a^-(s)}{a^+(s)} \Big/ \frac{a^-(t)}{a^+(t)}\right)\right| \\ &= \frac{1}{2}\left|\log\left(\frac{a^-(s)}{a^-(t)}\right) + \log\left(\frac{a^+(t)}{a^+(s)}\right)\right|\end{aligned}$$

Assuming as we may that $s \le t$, we apply Lemma 3.6 to obtain a constant L depending only on $\mathcal{B}$, ρ, δ so that

$$\begin{aligned}\mathrm{len}_s(\lambda^+) &\le L^{t-s}\,\mathrm{len}_t(\lambda^+) \\ a^+(t) &\le L^{t-s}a^+(s) \\ \log\left(\frac{a^+(t)}{a^+(s)}\right) &\le \log(L)(t-s)\end{aligned}$$

and similarly

$$\log\left(\frac{a^-(s)}{a^-(t)}\right) \le \log(L)(t-s)$$

and so

$$d(\Sigma(s),\Sigma(t)) \le \log(L)\,|t-s|\,.$$

For the lower bound we apply Lemma 3.6 again to obtain $(\kappa,n,0)$ flaring of the sequences $a^+(i)$, $a^-(i)$, with κ, n depending only on $\mathcal{B}$, ρ, δ. We also use the fact that $a^+(i)$ achieves its minimum near the right end of I and that $a^-(i)$ achieves its minimum near the left end. To simplify matters, truncate I so that any finite endpoint of I is an integer divisible by n, and if it is a left (resp. right) endpoint np then the minimum of the sequence $(a^-(ni))_i$ (resp. $(a^+(ni))_i$) is achieved uniquely with $i=p$. By Proposition 3.12 we need only chop off an amount of length $\le \eta+2n$ to achieve this effect; at worst the additive quasigeodesic constant for Σ is increased by an amount depending only on $\mathcal{B}$, ρ, δ, and the multiplicative constant is unchanged. For $s = np < t = nq \in I$ we therefore have

$$\begin{aligned}\mathrm{len}_t(\lambda^-) &\ge \kappa^{q-p}\,\mathrm{len}_s(\lambda^-)\\ a^-(s) &\ge \kappa'^{q-p}a^+(s), \qquad \kappa' = \kappa^{1/n}\\ \log\left(\frac{a^-(s)}{a^-(t)}\right) &\ge \kappa'^{t-s}\end{aligned}$$

and similarly

$$\log\left(\frac{a^+(t)}{a^+(s)}\right) \ge \kappa'^{t-s}$$

and so

$$d(\Sigma(s),\Sigma(t)) \ge \log(\kappa')\,|t-s|\,.$$

For general $s \le t \in I$ pick $s' = np \le t' = nq \in I$ so that $|s-s'|, |t-t'| \le n$, and we have:

$$\begin{aligned}d(\Sigma(s),\Sigma(t)) &\ge d(\Sigma(s'),\Sigma(t')) - d(\Sigma(s),\Sigma(s')) - d(\Sigma(t'),\Sigma(s'))\\ &\ge \log(\kappa')\left|t'-s'\right| - 2n\log(L)\\ &\ge \log(\kappa')\,|t-s| - 2n\log(L) - 2n\log(\kappa')\end{aligned}$$

This completes the proof of Theorem 1.1.

4 Model geometries for geometrically infinite ends

Throughout this section a *hyperbolic 3–manifold* will always be complete, with finitely generated, freely indecomposable fundamental group, and with no parabolics.

If N is a hyperbolic 3–manifold, Scott's core theorem [23] produces a *compact core* $K_N \subset N$, a compact, codimension–0 submanifold whose inclusion $K_N \hookrightarrow N$ is a homotopy equivalence. Bonahon proves [4] that N is *geometrically tame*, which by a result of Thurston [25] implies that the inclusion $\mathrm{int}(K_N) \hookrightarrow N$ is homotopic to a homeomorphism. For each end e, let N_e be the closure of the component of $N - K_N$ corresponding to e, and it follows that the inclusion of the closed surface $S_e = K_N \cap N_e$ into N_e extends to a homeomorphism $S_e \times [0,\infty) \to N_e$.

The manifold N is the quotient of a free and properly discontinuous action of $\pi_1 N$ on $\mathbf{H}^3$, with limit set $\Lambda \subset S^2 = \partial \mathbf{H}^3$ and domain of discontinuity $D = S^2 - \Lambda$. The quotient of the convex hull of Λ in $\mathbf{H}^3$ is called the convex hull of N, $\mathrm{Hull}(N) = \mathrm{Hull}(\Lambda)/\pi_1 N \subset N$. Let $\overline{N} = (\mathbf{H}^3 \cup D)/\pi_1 N$.

An end e is *geometrically finite* if N_e is precompact in $\overline{N}$, in which case one may choose the homeomorphism $S_e \times [0,\infty) \to N_e$ so that it extends to a homeomorphism $S_e \times [0,\infty] \to \overline{N}_e \subset \overline{N}$, with $S_e \times \infty$ a component of $\partial\overline{N}$. Under the action of $\pi_1 N$ there is an orbit of components of D, each an open disc δ, such that the projection map $\delta \to S_e \times \infty$ is a universal covering map. Since the stabilizer of δ in $\pi_1 N$ acts conformally on δ, we obtain a conformal structure on $S_e \times \infty$. This conformal structure is independent of the choice of δ in the orbit, and its isotopy class is independent of the choice of the homeomorphism $S_e \times [0,\infty] \to \overline{N}_e$ extending the inclusion map $S_e \hookrightarrow N_e$. We therefore obtain a well-defined point of the Teichmüller space of S_e, called the *conformal structure at* ∞ associated to the end e.

An end e is *geometrically infinite* if it is not geometrically finite. Bonahon proves [4] that if e is geometrically infinite then e is *geometrically tame*, which means that there is a sequence of hyperbolic structures σ_i on S_e and pleated surfaces $g_i \colon (S_e, \sigma_i) \to N_e$, each homotopic to the inclusion $S_e \hookrightarrow N_e$, such that the sequence of sets $g_i(S_e)$ leaves every compact subset of N_e. In this situation, the σ_i form a sequence in the Teichmüller space $\mathcal{T}(S_e)$ which accumulates on the boundary $\mathcal{PMF}(S_e)$. The unmeasured lamination on S_e corresponding to any accumulation point gives a unique point in the space $\mathcal{GL}(S_e)$, called the *ending lamination* of e.

If N' is another 3–hyperbolic manifold and $f\colon N \to N'$ is a homeomorphism, then f induces a bijection between the ends of N and of N', an isometry between the corresponding Teichmüller spaces, and a homeomorphism between the corresponding geodesic lamination spaces. Thurston's ending lamination conjecture says that if the end invariants of N and N' agree under this correspondence, then f is homotopic to an isometry. Define the *injectivity radius* of a hyperbolic manifold N at a point x, denoted $\mathrm{inj}_x(N)$, to be the smallest $\epsilon \geq 0$ such that the ball of radius ϵ about x is isometric to the ball of radius ϵ in hyperbolic space. The injectivity radius of N itself is $\mathrm{inj}(N) = \inf\limits_{x \in N} \mathrm{inj}_x N$. The manifold N has *bounded geometry* if $\mathrm{inj}(N) > 0$.

Theorem 4.1 (Minsky [16]) *The ending lamination conjecture holds for complete hyperbolic 3–manifolds N, N' with finitely generated, freely indecomposable fundamental group and with bounded geometry.*

In the introduction we recounted briefly how Minsky reduced this theorem to the construction of model manifolds for geometrically infinite ends, Theorems 1.3 and 1.4. The proofs of these theorems occupy sections 4.1–4.3. The construction of a model manifold for an end of bounded geometry, even when the ambient manifold does not have bounded geometry, is given in Section 4.4.

4.1 Pleated surfaces

We review here several facts about pleated surfaces, their geometry, and their homotopies. For fuller coverage the reader is referred to [25], [8], and in particular results of [16] which are crucial to our proofs of Theorems 1.3 and 1.4.

Given a hyperbolic 3–manifold N, a *pleated surface* in N is a π_1–injective, continuous map $\theta\colon F \to N$ where F is a closed surface, together with a hyperbolic structure σ on F and a geodesic lamination λ on σ, such that any rectifiable path in F is taken to a rectifiable path in N of the same length, θ is totally geodesic on each leaf of λ, and θ is totally geodesic on each component of $F - \lambda$. The minimal such lamination λ is called the *pleating locus* of θ. We incorporate the hyperbolic structure σ into the notation by writing $\theta\colon (F, \sigma) \to N$.

When e is a geometrically infinite end with neighborhood $N_e \approx S_e \times [0, \infty)$, we shall assume implicitly that any pleated surface with image in N_e is of the form $\theta\colon (S_e, \sigma) \to N_e$ where θ is homotopic to the inclusion $S_e \hookrightarrow N_e$. We may therefore drop S_e from the notation and write $\theta\colon \sigma \to N_e$. A similar notational convention will be used when $N \approx S \times (-\infty, +\infty)$.

An end e is *geometrically tame* if for each compact subset of N_e there is a pleated surface $\theta\colon \sigma \to N_e$ which misses that compact set. Bonahon proved [4] (using free indecomposability of $\pi_1 N$) that each geometrically infinite end is geometrically tame.

The first lemma controls the geometry of a pleated surface in the large, as long as the ambient manifold has bounded geometry. Note that any pleated surface is distance nonincreasing, that is, $(1,0)$–coarse lipschitz.

Lemma 4.2 *For each $\epsilon > 0$, $g \ge 2$ there exists $d > 0$, and a properness gauge $\rho\colon [0,\infty) \to [0,\infty)$, such that if $\theta\colon (F,\sigma) \to N$ is a pleated surface with* $\mathrm{genus}(F) \le g$ *and* $\mathrm{inj}(N) \ge \epsilon$*, then*

(1) *The diameter of $\theta(F)$ in N is $\le d$.*

(2) ([16], Lemma 4.4) *The map $\widetilde{\theta}\colon \widetilde{F} \to \widetilde{N} = \mathbf{H}^3$ is ρ uniformly proper.*

The next fact relates the geometry of a nearby pair of pleated surfaces:

Lemma 4.3 ([16], Lemma 4.5) *For each $\epsilon > 0$, $g \ge 2$, and $a \ge 0$ there exists $r \ge 0$ such that the following holds. Let $\theta_i\colon (F,\sigma_i) \to N$, $i = 0,1$, be homotopic pleated surfaces, and suppose that* $\mathrm{inj}(N) \ge \epsilon$, $\mathrm{genus}(F) \le g$, *and* $d(\mathrm{image}(\theta_0), \mathrm{image}(\theta_1)) \le a$. *Then the distance from σ_0 to σ_1 in the Teichmüller space of F is at most r.*

Next we need a result controlling the geometry of a homotopy between nearby pleated surfaces. Recall that if $\mathcal{T}$ is the Teichmüller space of a surface F of genus g, then for each $\epsilon > 0$, $r \ge 0$ there exists $\kappa \ge 1$ depending only on g, ϵ, r such that if $\sigma_0, \sigma_1 \in \mathcal{T}$ both have injectivity radius $\ge \epsilon$, and if σ_0, σ_1 have distance $\le r$ in $\mathcal{T}$, then there is a κ–bilipschitz map $\phi\colon \sigma_0 \to \sigma_1$ isotopic to the identity.

Lemma 4.4 ([16], Lemma 4.2) *For each $\epsilon > 0$, $g \ge 2$, $\kappa \ge 1$ there exists $B \ge 0$ such that the following hold. If N has injectivity radius $\ge \epsilon$ and F has genus $\le g$, if $\theta_i\colon (F,\sigma_i) \to N$ are homotopic pleated surfaces, and if $\phi\colon \sigma_0 \to \sigma_1$ is a κ–bilipschitz map isotopic to the identity, then there is a straight line homotopy between θ_0 and $\theta_1 \circ \phi$ whose tracks have length $\le B$.*

The next fact gives us the raw material for constructing nearby pleated surfaces as needed. It is an almost immediate consequence of section 9.5 of [25]; we provide extra details for convenience.

Lemma 4.5 *There exists a constant $\eta > 0$ such that for any homotopy equivalence $S \to N$ from a closed surface S to a complete hyperbolic 3–manifold N without parabolics, every point in the convex hull of N comes within distance η of the image of some pleated surface $\theta \colon S \to N$.*

Proof We prove this with a bound η equal to the smallest positive real number such that for any ideal hyperbolic tetrahedron $\tau \subset \mathbf{H}^3$ and any two faces σ_1, σ_2 of τ, any point of τ is within distance η of $\sigma_1 \cup \sigma_2$. By slicing τ with totally geodesic planes passing through the cusps of τ, it follows that η is equal to the thinness constant for the hyperbolic plane $\mathbf{H}^2$, namely $\eta = \log(1+\sqrt{2})$ (as computed in [6], Theorem 11.8).

For each point x in the convex hull of N, there exists a disjoint pair of pleated surfaces $\theta_0, \theta_1 \colon S \to N$ such that if C is the unique component of $N - (\theta_0(S) \cup \theta_1(S))$ whose closure $\overline{C}$ intersects both $\theta_0(S)$ and $\theta_1(S)$, then $x \in \overline{C}$. If the ends of N are e_0, e_1, we may choose θ_i to separate x from e_i, either the convex hull boundary when e_i is geometrically finite, or a pleated surface very far out when e_i is geometrically tame; we use here Bonahon's theorem that N, and each end of N, is geometrically tame [4]. It follows that x is in the image of any homotopy from θ_0 to θ_1.

Now we use the construction of section 9.5 of [25]. Given a homotopic pair of pleated surfaces $\theta_0, \theta_1 \colon S \to N$, this construction produces a homotopy θ_t, $t \in [0,1]$, such that each map $\theta_t \colon S \to N$ in the homotopy is either a pleated surface, or there exist $t' < t < t''$ in $[0,1]$ such that $\theta_{[t',t'']}$ is an arbitrarily short straight line homotopy, or $\theta_{[t',t'']}$ is supported on the image of a locally isometric map $\tau \to N$ where τ is an ideal tetrahedron in $\mathbf{H}^3$. In the latter situation, there are two cases. In the first case, on page 9.47 of [25], the homotopy $\theta_{[t',t'']}$ moves two faces of τ to the opposite two faces through the image of τ in N. In the second case, on the bottom of page 9.48 of [25], there are actually two tetrahedra involved in the homotopy but we may homotop through them one at a time; the map $\tau \to N$ identifes two faces of the tetrahedron, wrapping their common edge infinitely around a closed curve in N, and the homotopy $\theta_{[t',t'']}$ moves one of the remaining two faces to last remaining face through the image of τ in N. The upshot in either case is that there are at least two faces of τ whose images in N lie on the image of one or the other of the pleated surfaces $\theta_{t'}$, $\theta_{t''}$, and each point swept out by the homotopy is within a bounded distance of the union of these two faces. □

Given a metric space with metric d and two subsets α, β, let

$$d(\alpha,\beta) = \inf\{d(x,y) \bigm| x \in \alpha, y \in \beta\}.$$

A sequence of homotopic pleated surface $\theta_i\colon \sigma_i \to N_e$, defined for $i \in I$ where I is a subinterval of the integers, is said to be *uniformly distributed* if there are constants $B > A > 0$ such that for each $i \geq 1$,

$$A \leq d(\mathrm{image}(\theta_{i-1}), \mathrm{image}(\theta_i)) \leq B$$

and $\mathrm{image}(\theta_i)$ separates $\mathrm{image}(\theta_{i-1})$ from $\mathrm{image}(\theta_{i+1})$ in N.

Combining Lemma 4.2(1) with Lemma 4.5 we immediately obtain:

Lemma 4.6 *Suppose N has bounded geometry. Each geometrically infinite end N_e has a uniformly distributed sequence of pleated surfaces $\theta_i\colon \sigma_i \to N_e$, $i \geq 1$, escaping the end. If $N \approx S \times (-\infty, +\infty)$ has two geometrically infinite ends then there is a uniformly distributed sequence of pleated surfaces $\theta_i\colon \sigma_i \to N_e$, $i \in \mathbf{Z}$, escaping both ends.*

In the first case of this lemma, note that if $\widetilde{N} \to N$ is the covering space corresponding to the injection $\pi_1 S_e \hookrightarrow \pi_1 N$, then the sequence $\theta_i\colon \sigma_i \to N_e$ must be contained in the projection of $\mathrm{Hull}(\widetilde{N})$ to N.

4.2 The doubly degenerate case: Theorem 1.3

Let N, as above, be a complete hyperbolic 3–manifold with finitely generated, freely indecomposable fundamental group and with bounded geometry. Let e be a geometrically infinite end of N with corresponding surface $S = S_e$ and neighborhood $N_e \approx S \times [1, +\infty)$, and we assume that the injection $\pi_1(S) \hookrightarrow \pi_1(N)$ is doubly degenerate. Let $\widetilde{N} \to N$ be the covering map corresponding to the subgroup $\pi_1(S)$, so we have a homeomorphism $\widetilde{N} \approx S \times (-\infty, +\infty)$. The manifold $\widetilde{N}$ has an end $\tilde{e}$ with neighborhood $\widetilde{N}_{\tilde{e}}$ such that the covering map restricts to an isometry $\widetilde{N}_{\tilde{e}} \to N_e$, and we may assume that the notation is chosen so that this isometry is expressed by the identity map on $S \times [1, +\infty)$.

By Theorem 9.2.2 of [25], each end of $\widetilde{N}$ has a neighborhood mapping properly to a neighborhood of a geometrically infinite end of N, and hence N has at most two ends, each geometrically infinite.

Consider first the case that N has exactly two ends, and so the covering map $\widetilde{N} \to N$ is bijective on ends. It follows clearly that the covering map $\widetilde{N} \to N$ has degree 1 on some neighborhood of e. Since degree is locally constant, the covering map is a homeomorphism, and we have $N \approx \widetilde{N} \approx S \times (-\infty, +\infty)$.

Consider next the case that N has just the one end e. The manifold $\widetilde{N}$ has two ends, $\tilde{e}$ with neighborhood $\widetilde{N}_{\tilde{e}} \approx S \times [1, +\infty)$, and $\tilde{e}'$ with neighborhood

$\widetilde{N}_{\tilde{e}'} \approx S \times (-\infty, -1]$. Each of these maps by a finite degree covering map to $N_e \approx S \times [1, +\infty)$, in each case induced by a covering map $S \to S$ which must be degree 1, and hence $\widetilde{N} \to N$ is a degree 2 map. It follows that N has a compact core which is doubly covered by $S \times [-1, +1]$, and which has an orbifold fibration with generic fiber S and base orbifold $[0, 1]$ with a $\mathbf{Z}/2$ mirror point at 0. This implies that N itself has an orbifold fibration with generic fiber S and base orbifold $[0, \infty)$ with a $\mathbf{Z}/2$ mirror point at 0.

Here for convenience is a restatement of theorem 1.3:

Theorem (Doubly degenerate model manifold) *Under the above conditions, there exists a unique cobounded geodesic line g in $\mathcal{T}(S)$ such that $\Sigma(S \to N)$ is Hausdorff equivalent to g in $\mathcal{T}(S)$. Moreover: the homeomorphism $S \times (-\infty, +\infty) \approx \widetilde{N}$ is properly homotopic to a map which lifts to a quasi-isometry $\mathcal{H}_g^{\text{SOLV}} \to \mathbf{H}^3$; and the ideal endpoints of g in $\mathcal{PMF}(S)$ are the respective ending laminations of the two ends of $\widetilde{N}$. In the degree 2 case, the order 2 covering transformation group on $\widetilde{N}$ acts isometrically on $\mathcal{S}_g^{\text{SOLV}}$ so as to commute with the homeomorphism $\mathcal{S}_g^{\text{SOLV}} \approx \widetilde{N}$.*

We'll focus on the two-ended case $N \approx S \times (-\infty, +\infty)$, mentioning later the changes needed for the one-ended case.

Let $\mathcal{T}$ be the Teichmüller space of S.

Applying Lemma 4.6, let $\theta_n \colon (S, \sigma_n) \to N$, $n \in \mathbf{Z}$, be a uniformly distributed sequence of pleated surfaces. By Lemma 4.3 there is a constant ρ such that for any $n \in \mathbf{Z}$, the distance in $\mathcal{T}$ between σ_n and σ_{n+1} is at most ρ. It follows that there is a $\mathbf{Z}$–piecewise affine, ρ–lipschitz path $\gamma \colon \mathbf{R} \to \mathcal{T}$ with $\gamma(n) = \sigma_n$. Since $\text{inj}(\sigma_n)$ is uniformly bounded away from zero it follows that the path γ is cobounded.

Consider now the canonical hyperbolic surface bundle $\mathcal{S}_\gamma \to \mathbf{R}$ and its universal cover, the canonical hyperbolic plane bundle $\mathcal{H}_\gamma \to \mathbf{R}$. We have an identification $\mathcal{S}_n \approx \sigma_n \in \mathcal{T}$ for every $n \in \mathbf{Z}$.

Claim 4.7 *There exists a map $\Phi \colon \mathcal{S}_\gamma \to N$ in the correct proper homotopy class which lifts to a quasi-isometry of universal covers $\widetilde{\Phi} \colon \mathcal{H}_\gamma \to \mathbf{H}^3$.*

Proof To construct the map Φ, restricting the bundle $\mathcal{S}_\gamma$ to the affine subpath $\gamma \mid [n, n+1]$ we obtain the bundle $\mathcal{S}_{[n,n+1]} \to [n, n+1]$, on which there is a connection with bilipschitz constant κ depending only the coboundedness and lipschitz constant of γ. In particular, we obtain a κ–bilipschitz map $\phi_n\colon \sigma_n \to \sigma_{n+1}$, which may be regarded as a map on S isotopic to the identity; but in fact, under the topological identification $\sigma_n \approx \mathcal{S}_n \approx S$, the map $\phi_n\colon S \to S$ must be the identity. Applying Lemma 4.4 there is a constant B and, for each n, a straight line homotopy in N with tracks of length $\le B$ from θ_n to θ_{n+1}. The domain of this homotopy may be taken to be $\mathcal{S}_{[n,n+1]} \approx S \times [n, n+1]$, and in particular the homotopy is a B–lipschitz map when restricted to any $x \times [n, n+1]$. Piecing these homotopies together for each n we obtain the desired map $\Phi\colon \mathcal{S}_\gamma \to N$.

To prove that any lift $\widetilde\Phi\colon \mathcal{H}_\gamma \to \mathbf{H}^3$ is a quasi-isometry, obviously $\widetilde\Phi$ is surjective, and so it suffices by Lemma 2.1(2) to prove that $\widetilde\Phi$ is coarse lipschitz and uniformly proper. Coarse lipschitz is immediate from the fact that $\widetilde\Phi \mid \mathcal{H}_n$ is distance nonincreasing for each n, and that $\widetilde\Phi$ is B–lipschitz along connection lines (indeed this implies that $\widetilde\Phi$ is lipschitz).

Uniform properness of $\widetilde\Phi$ will follow from two facts. Lemma 4.2(2) tells us that the maps $\widetilde\theta_n\colon \widetilde\sigma_n \to \mathbf{H}^3$, which are identified with the maps $\widetilde\Phi \mid \mathcal{H}_n$, are uniformly proper with a properness gauge independent of n. Also, the images $\widetilde\Phi(\mathcal{H}_n)$ are uniformly distributed in $\mathcal{H}_\gamma$, by Lemma 4.6. We put these together as follows.

Consider a number $r \ge 0$ and points $x, y \in \mathcal{H}_\gamma$. We must show that there is a number s independent of x, y, depending only on r, such that

$$\text{if} \quad d(x,y) \ge s \quad \text{then} \quad d(\xi,\eta) \ge r, \quad \text{with} \quad \xi = \widetilde\Phi(x), \eta = \widetilde\Phi(y) \tag{4.1}$$

We claim there are constants $L \ge 1, C \ge 0$ such that if $x \in \mathcal{H}_t$ and $y \in \mathcal{H}_u$, then $d(\xi,\eta) \ge \frac{1}{L}|t-u| - C$. To see why, note that we may assume that $t, u \in \mathbf{Z}$ and $t < u$. Consider the geodesic $\overline{\xi\eta}$ in $\mathbf{H}^3$. By Lemma 4.6, along this geodesic there is a monotonic sequence of points $\xi = \zeta_t, \zeta_{t+1}, \ldots, \zeta_u = \eta$ such that $\zeta_n \in \widetilde\Phi(\mathcal{H}_n)$, and such that $\operatorname{len}(\overline{\zeta_n\zeta_{n+1}})$ is bounded away from zero, establishing the claim.

In order to prove 4.1 it therefore suffices to consider the case $x \in \mathcal{H}_t$, $y \in \mathcal{H}_u$ with $|u-t| \le R_0 = Lr + LC$. Let y' be the point of $\mathcal{H}_t$ obtained by moving from y into $\mathcal{H}_t$ along a connection line, and so $d(y,y') \le R_0$. Setting $\eta' = \widetilde\Phi(y')$ we have $d(\eta,\eta') \le BR_0$. To prove that $d(\xi,\eta) \ge r$ it therefore suffices to prove

$$d(\xi,\eta') \ge R_1 = r + BR_0.$$

The restriction of $\widetilde\Phi$ to $\mathcal{H}_n$ is just a lift of the pleated surface $\theta_n\colon \sigma_n \to N$. Applying Lemma 4.2(2) it follows that $\widetilde\Phi \mid \mathcal{H}_n$ is ρ-uniformly proper with ρ independent of n, and so

$$d(\xi,\eta') \ge \rho(d_n(x,y'))$$

where d_n is the distance function on $\mathcal{H}_n$ (an isometric copy of $\mathbf{H}^2$). By properness of ρ there exists therefore a number $R_2 \ge 0$ depending only on R_1 such that

$$\text{if} \quad d_n(x,y') \ge R_2 \quad \text{then} \quad \rho(d_n(x,y')) \ge R_1$$

and so it suffices to prove $d_n(x,y') \ge R_2$. Since the inclusion map $\mathcal{H}_n \hookrightarrow \mathcal{H}_\gamma$ is distance nonincreasing it suffices to prove $d(x,y') \ge R_2$. Setting $s = R_2 + R_0$ and using the fact that $d(y,y') \le R_0$, we have:

$$\text{if} \quad d(x,y) \ge s \quad \text{then} \quad d(x,y') \ge R_2$$

Putting it altogether, if $d(x,y) \ge s$ then $d(\xi,\eta) \ge r$. This completes the proof that $\widetilde{\Phi}$ is a quasi-isometry. □

Applying Claim 4.7, since $\mathbf{H}^3$ is a hyperbolic metric space, it follows that $\mathcal{H}_\gamma$ is also a hyperbolic metric space. Theorem 1.1 now applies, and we conclude that there is a bi-infinite geodesic g in $\mathcal{T}$ such that γ and g are asynchronous fellow travellers, that is, there is a quasi-isometry $s\colon \mathbf{R} \to \mathbf{R}$ such that $d(\gamma(t), g(s(t)))$ is bounded independent of t. Note that since γ is cobounded and γ and g are asynchronous fellow travellers, it follows that g is cobounded. Applying Proposition 2.3 it follows that the map $t \mapsto s(t)$ lifts to a map $\mathcal{S}_g^{\text{SOLV}} \to \mathcal{S}_\gamma$ any of whose lifts $\mathcal{H}_g^{\text{SOLV}} \to \mathcal{H}_\gamma$ is a quasi-isometry. By composition we therefore obtain a map $\mathcal{S}_g^{\text{SOLV}} \to N$ any of whose lifts to universal covers $\mathcal{H}_g^{\text{SOLV}} \to \mathbf{H}^3$ is a quasi-isometry, as required for proving Theorem 1.3.

To complete the proof in the case where $N \approx S \times (-\infty, +\infty)$, there are a few loose ends to clean up.

From what we have proved it is evident that g is contained in a bounded neighborhood of the subset $\{\sigma_i\}$ in $\mathcal{T}$, and hence g is contained in a bounded neighborhood of the subset $\Sigma(S \to N)$. The opposite containment is an immediate consequence of the fact that the pleated surface sequence $\theta_i\colon \sigma_i \to N$ is uniformly spaced, combined with Lemma 4.2(1) and Lemma 4.3.

Since γ is evidently cobounded, it follows that g is cobounded.

Uniqueness of g is a consequence of the general fact that if g, g' are two cobounded Teichmüller geodesics whose Hausdorff distance is finite then $g = g'$; for a proof see [10] Lemma 2.4.

The fact that the two ends of the geodesic g are the two ending laminations of N is a consequence of the following fact: if $\theta_i\colon \sigma_i \to N$, $i \ge 1$, is a sequence of pleated surfaces going out an end e of N, then for any sequence $\sigma_i' \in \mathcal{T}(S)$ such that $d(\sigma_i, \sigma_i')$ is bounded, any accumulation point of σ_i' in $\overline{\mathcal{T}}(S)$ is an element of $\mathcal{PMF}(S)$ whose image in $\mathcal{GL}(S)$ is the ending lamination of e. For the proof see Lemmas 9.2 and 9.3 of [16].

Now we turn to the case where N fibers over the ray orbifold with generic fiber S and one geometrically infinite end. The above analysis applies to the degree 2 covering manifold $\widetilde{N} \approx S \times (-\infty, +\infty)$, for which we obtain a cobounded Teichmüller geodesic g and a map $\mathcal{S}_g^{\text{SOLV}} \to \widetilde{N}$. There is an order 2 isometric covering transformation $\tau\colon \widetilde{N} \to \widetilde{N}$ which exchanges the two ends, inducing an order 2 mapping class on S and an order 2 isometry of the Teichmüller space $\mathcal{T}$. We may then easily tailor the proof of Lemma 4.5 to produce a uniformly distributed sequence of pleated surfaces $\theta_i\colon \sigma_i \to N$ such that in $\mathcal{T}$ we have $\tau(\sigma_i) = \sigma_{-i}$ for all $i \in \mathbf{Z}$. By the uniqueness clause of Theorem 1.3 for $\widetilde{N} \approx S \times (-\infty, +\infty)$ it follows that τ preserves the Teichmüller geodesic g, acting on it by reflection across some fixed point. This implies that the homeomorphism $\mathcal{S}_g^{\text{SOLV}} \to \widetilde{N}$ may be chosen in its proper homotopy class so that the action of τ on $\widetilde{N}$ commutes with an action on $\mathcal{S}_g^{\text{SOLV}}$.

4.3 The singly degenerate case: Theorem 1.4

Again let N be a complete hyperbolic 3–manifold with finitely generated, freely indecomposable fundamental group and with bounded geometry, and let e an end of N with corresponding surface $S = S_e \hookrightarrow N$ bounding a neighborhood $S \times [0, \infty) \approx N_e \subset N$ of the end e. We assume that the injection $\pi_1(S) \hookrightarrow \pi_1(N)$ is singly degenerate. By lifting to the covering space of N corresponding to the subgroup $\pi_1(S) \hookrightarrow N$ we may assume that we have $N \approx S \times (-\infty, +\infty)$. Since we are in the singly degenerate case, the convex hull of N is bounded by a surface homotopic to S, and so we may assume that S is equal to this surface, and hence N_e is the convex hull of N.

Next we proceed as in the doubly degenerate case to get a uniformly distributed sequence of pleated surfaces $\theta_n\colon (S, \sigma_n) \to N_e$, $n \geq 0$, where $\theta_0\colon S \to N_e$ is just the inclusion. We obtain a $\mathbf{Z}$–piecewise affine, cobounded, lipschitz ray $\gamma\colon [0, \infty) \to \mathcal{T}$ with $\gamma(n) = \sigma_n$. The proof of Claim 4.7 goes through as before, and we obtain a map $\Phi\colon \mathcal{S}_\gamma \to N_e$ in the correct proper homotopy class any of whose lifts $\widetilde{\Phi}\colon \mathcal{H}_\gamma \to \widetilde{N}_e$ is a quasi-isometry, where $\widetilde{N}_e$ is the universal cover of N_e.

Now there is a trick: the convex hull of the limit set of $\pi_1 S$ is a convex subset of $\mathbf{H}^3$. This convex hull is precisely the space $\widetilde{N}_e$, and hence the restriction of the geodesic metric on $\mathbf{H}^3$ is a hyperbolic geodesic metric on $\widetilde{N}_e$. The metric space $\mathcal{H}_\gamma$ is therefore also hyperbolic, and so we may apply Theorem 1.1 to γ, concluding that there is a geodesic ray $g\colon [0,\infty) \to \mathcal{T}$ such that γ and g are asynchronous fellow travellers, and such that $g(0) = \gamma(0)$.

The proof that $\Sigma(S \to N)$ is Hausdorff equivalent to g in $\mathcal{T}$ goes through as in the previous case. Given that g has base point $\gamma(0)$, uniqueness of g is proved as before using coboundedness of g. For any other base point $\tau \in \mathcal{T}$, there is a unique ray g_τ in $\mathcal{T}$ with base point τ asymptotic to g; uniqueness of g_τ is proved as before, and existence is a standard fact.

Since γ and g_τ are Hausdorff equivalent and γ is a quasigeodesic, it follows that γ and g_τ are asynchronous fellow travellers. The same argument as before produces a proper homotopy equivalence $N_e \to \mathcal{S}_{g_\tau}$ in the correct proper homotopy class which lifts to a quasi-isometry $\widetilde{N}_e \to \mathcal{H}_{g_\tau}$.

The proof that the endpoint of g in $\mathcal{PMF}(S)$ gives the ending lamination of e is as before.

4.4 Ends of bounded geometry

Suppose that N is a hyperbolic 3–manifold as above, but without bounded geometry. Given an end e of N, we say that e has bounded geometry if there exists $\epsilon > 0$ such that $\mathrm{inj}(N_e) \ge \epsilon$, meaning that for each $x \in N_e$ we have $\mathrm{inj}_x(N) \ge \epsilon$. In this case we can push the construction of the model manifold through; we are grateful to Jeff Brock for suggesting this possibility. We shall merely sketch the outline of how to do this.

The first tricky part is that the pleated surface results of Minsky quoted in Section 4.1 must be adapted to hold without the assumption that N has bounded geometry. Instead, using only bounded geometry of the end e, one must prove that the results hold in N_e.

In the statement of Lemma 4.2, the hypothesis that $\mathrm{inj}(N) > \epsilon$ is replaced with $\mathrm{inj}(N_e) > \epsilon$, and the conclusion that $\tilde\theta\colon \tilde F \to \mathbf{H}^3$ is ρ–uniformly proper is replaced by the conclusion that $\tilde\theta\colon \tilde F \to \tilde N_e$ is ρ–uniformly proper. With these changes, Minsky's proof taken from [16] Lemma 4.4 (which goes back to [15] Lemma 4.5) goes through, using basic compactness results for pleated surfaces such as the theorem of Thurston quoted in [15] as Theorem 4.1.

The statement of Lemma 4.3 undergoes a similar change, and again Minsky's proof taken from [16] Lemma 4.5 (which goes back to [15] Lemma 4.6) goes through.

Lemma 4.4 also undergoes a similar change.

With these results suitably adapted, it follows that there is an evenly spaced sequence of pleated surfaces $\theta_i \colon \sigma_i \to N_e$ escaping the end e, and there is a $\mathbf{Z}$–piecewise affine, cobounded, Lipschitz ray $\gamma \colon [0,\infty) \to \mathcal{T}$ with $\gamma(i) = \sigma_i$. The proof of Claim 4.7 goes through, producing a map $\Phi \colon \mathcal{S}_\gamma \to N_e$ in the correct proper homotopy class, which lifts to a quasi-isometry $\tilde\Phi \colon \mathcal{H}_\gamma \to \widetilde{N}_e$.

Now we need to check that $\mathcal{H}_\gamma$, or equivalently $\widetilde{N}_e$, is a hyperbolic metric space. The convexity trick used in Section 4.3 is not available; instead, we appeal to the methods of [10]. As remarked earlier, hyperbolicity of $\mathcal{H}_\gamma$ is equivalent to the horizontal flaring property, with a uniform relation between the constants. It therefore suffices to show directly that the horizontal flaring property holds in $\mathcal{H}_\gamma$. Suppose that $\alpha, \beta \colon I \to \mathcal{H}_\gamma$, $I \subset [0,\infty)$, are λ–quasihorizontal paths in $\mathcal{H}_\gamma$. We must prove that the sequence $d_i(\alpha(i), \beta(i))$ satisfies $\kappa, n, A(\lambda)$ flaring with appropriate constants, and we do this by applying the proof of Lemma 5.2 of [10]. All that is needed to apply that proof is to show that a λ–quasihorizontal path in $\mathcal{H}_\gamma$ maps via Φ to a quasigeodesic in $\mathbf{H}^3$, with quasigeodesic constants depending only on λ; but this follows immediately from even spacing of the pleated surface sequence $\theta_i \colon \sigma_i \to N_e$.

Since $\mathcal{H}_\gamma$ is a hyperbolic metric space, Theorem 1.1 applies as above to conclude that γ fellow travels a unique cobounded geodesic ray g in $\mathcal{T}$ with $\gamma(0) = g(0)$, and from Proposition 2.3 we obtain the desired quasi-isometry $\mathcal{H}_g^{\text{SOLV}} \to \mathcal{H}_\gamma \xrightarrow{\tilde\Phi} \widetilde{N}_e$.

References

[1] **L Bers**, *Fiber spaces over Teichmüller spaces*, Acta Math. 130 (1973) 89–126

[2] **M Bestvina**, **M Feighn**, *A combination theorem for negatively curved groups*, J. Diff. Geom. 35 (1992) 85–101

[3] **J Brock**, **B Farb**, *Curvature and rank of Teichmüller space*, preprint (2001) `arXiv:math.GT/0109045`

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

[5] **B Bowditch**, *Stacks of hyperbolic spaces and ends of 3–manifolds*, preprint (2002)

[6] **J Cannon**, *The theory of negatively curved spaces and groups*, from: "Ergodic theory, symbolic dynamics, and hyperbolic spaces" (C Series, T Bedford, M Keane, editors) Oxford Univ. Press (1991)

[7] **A Casson**, **S Bleiler**, *Automorphisms of surfaces after Nielsen and Thurston*, London Mathematical Society Student Texts, vol. 9, Cambridge University Press (1988)

[8] **R D Canary**, **D B A Epstein**, **P Green**, *Notes on notes of Thurston*, from: "Analytical and geometric aspects of hyperbolic space", London Math. Soc. Lecture Notes Series, vol. 111, Cambridge University Press (1987) 3–92

[9] **D B A Epstein**, **A Marden**, *Convex hulls in hyperbolic space, a theorem of Sullivan, and measured pleated surfaces*, from: "Analytical and geometric aspects of hyperbolic space", London Math. Soc. Lecture Notes Series, vol. 111, Cambridge University Press (1987) 113–254

[10] **B Farb**, **L Mosher**, *Convex cocompact subgroups of mapping class groups*, Geometry and Topology 6 (2002) 91–152

[11] **S Gersten**, *Cohomological lower bounds for isoperimetric functions on groups*, Topology 37 (1998) 1031–1072

[12] **M Gromov**, *Hyperbolic groups*, from "Essays in group theory" (S Gersten, editor) MSRI Publications, vol. 8, Springer (1987)

[13] **J Hubbard**, **H Masur**, *Quadratic differentials and foliations*, Acta Math. 142 (1979) 221–274

[14] **M Hirsch**, **C Pugh**, **M Shub**, *Invariant manifolds*, Lecture Notes in Mathematics, no. 583, Springer (1977)

[15] **Y Minsky**, *Teichmüller geodesics and ends of hyperbolic 3–manifolds*, Topology 32 (1993) 625–647

[16] **Y Minsky**, *On rigidity, limit sets, and end invariants of hyperbolic 3–manifolds*, Jour. AMS 7 (1994) 539–588

[17] **Y Minsky**, *Quasi-projections in Teichmüller space*, J. Reine Angew. Math. 473 (1996) 121–136

[18] **Y Minsky**, *Bounded geometry for Kleinian groups*, Invent. Math. 146 (2001) 143–192

[19] **H Masur**, **Y Minsky**, *Geometry of the complex of curves, I. Hyperbolicity*, Invent. Math. 138 (1999) no. 1, 103–149

[20] **H Masur**, **Y Minsky**, *Unstable quasigeodesics in Teichmüller space*, from: "In the tradition of Ahlfors and Bers (Stony Brook, NY, 1998)", Contemp. Math., vol. 256, Amer. Math. Soc. (2000) 239–241

[21] **H M Morse**, *A fundamental class of geodesics on any closed surface of genus greater than one*, Trans. AMS 26 (1924) 25–60

[22] **H Masur**, **M Wolf**, *Teichmüller space is not Gromov hyperbolic*, Ann. Acad. Sci. Fenn. Ser. A I Math. 20 (1995) 259–267

[23] **P Scott**, *Compact submanifolds of 3–manifolds*, J. London Math. Sco. 7 (1973) 246–250

[24] **D Sullivan**, *On the ergodic theory at infinity of an arbitrary discrete group of hyperbolic motions*, from; "Riemann surfaces and related topics, Proceedings of the 1978 Stony Brook Conference", Ann. Math. Studies, vol. 97, Princeton University Press (1981) 465–496

[25] **W P Thurston**, *Geometry and topology of 3–manifolds*, Princeton University Lecture Notes (1987)

[26] **W P Thurston**, *Three-manifolds, foliations and circles, I*, preliminary version (Dec. 29 1997)