BOUNDED GEOMETRY FOR KLEINIAN GROUPS

YAIR N. MINSKY

Abstract. We show that a Kleinian surface group, or hyperbolic 3-manifold with a cusp-preserving homotopy-equivalence to a surface, has bounded geometry if and only if there is an upper bound on an associated collection of coefficients that depend only on its end invariants. Bounded geometry is a positive lower bound on the lengths of closed geodesics. When the surface is a once-punctured torus, the coefficients coincide with the continued fraction coefficients associated to the ending laminations.

Contents

1. Introduction

Let N be a hyperbolic 3-manifold homeomorphic to the interior of a compact manifold. We say that N has *bounded geometry* if, outside the cusps of N, there is a positive lower bound on injectivity radii. Equivalently, there is a lower bound on the length of all closed geodesics in N.

The condition of bounded geometry is very helpful in understanding some basic questions, such as classification by end invariants (Thurston's Ending Lamination Conjecture), and description of topological structure of the limit set. For example, if N is homeomorphic to the interior of a manifold with incompressible boundary and has no cusps, then bounded geometry provides an explicit quasi-isometric model for N, which is known to imply that N is uniquely determined by its end invariants, and gives a topological model for

Date: May 10, 2001.
Partially supported by NSF grant DMS-9971596.

the limit set of $\pi_1(N)$ acting on the Riemann sphere (see [44, 43] and Klarreich [34]; see also McMullen [38] for progress in the unbounded geometry case). Bounded geometry also has implications for the spectral theory of a hyperbolic 3-manifold and its L^2-cohomology (Canary [14], Lott [35]).

The *end invariants* of N are points in a certain parameter space associated to each end, which describe the asymptotic geometry of N. Conjecturally, N is determined uniquely by these invariants (see Thurston [53]). In this paper we address the question of whether at least the condition of bounded geometry can be detected from the end invariants, and our main theorem is a first application of some tools that were developed with the full conjecture in mind. As a corollary we will obtain a small extension of the setting in which the conjecture itself can be established.

Let us restrict now to the case where N is homeomorphic to $S \times \mathbb{R}$, where S is a surface of finite type. More specifically, we will describe our manifolds as quotients of $\mathbb{H}^3$ by injective representations

$$\rho : \pi_1(S) \to \mathrm{PSL}_2(\mathbb{C})$$

with discrete images, which are also *type preserving*, that is they map elements representing punctures of S to parabolics. Call such a representation a (marked) *Kleinian surface group.* The theory of Ahlfors-Bers, Thurston and Bonahon attaches to ρ two invariants (ν_+, ν_-) lying in a combination of Teichmüller spaces and lamination spaces of S and its subsurfaces.

We will associate to the pair (ν_+, ν_-) a collection of positive integers $\{d_Y(\nu_+, \nu_-)\}$, where Y runs over all isotopy classes of essential subsurfaces in S (see Sections 2.2 and 2.6). These are analogues of the continued fraction coefficients, considered in the setting where S is a once-punctured torus, in [46]. We will establish:

Bounded Geometry Theorem *Let $\rho : \pi_1(S) \to PSL_2(\mathbb{C})$ be a Kleinian surface group with no accidental parabolics, and end invariants (ν_+, ν_-). Then ρ has bounded geometry if and only if the coefficients $\{\pi_Y(\nu_+, \nu_-)\}$ are bounded above.*

Moreover, for any $K > 0$ there exists $\epsilon > 0$, depending only on K and the topological type of S, so that

$$\sup_Y d_Y(\nu_+, \nu_-) < K \quad \Longrightarrow \quad \inf_\gamma \ell_\rho(\gamma) > \epsilon$$

where the infimum is over elements of $\pi_1(S)$ that are not externally short, *and the supremum is over essential subsurfaces for which $d_Y(\nu_+, \nu_-)$ is defined. Similarly given ϵ there exists K for which the implication is reversed.*

As a corollary we obtain the following improvement of the main theorem of [44]:

Corollary 1 (Ending lamination theorem for bounded geometry) *Let $\rho_1, \rho_2 : \pi_1(S) \to PSL_2(\mathbb{C})$ be two Kleinian surface groups where S is a* *closed surface. Suppose that ρ_1, ρ_2 have the same end invariants, and that ρ_1 admits a lower bound*

$$\inf_{\gamma \in \pi_1(S)} \ell_{\rho_1}(\gamma) > 0. \tag{$*$}$$

Then ρ_1 and ρ_2 are conjugate in $PSL_2(\mathbb{C})$.

The main theorem of [44] gives this conclusion if *both* ρ_1 and ρ_2 are known to satisfy the bound $(*)$. The Bounded Geometry theorem implies that $(*)$ is equivalent to a condition depending only on the end invariants, which are common to both representations. Thus the improvement closes a nagging loophole in the earlier result. As a special case we obtain the following, which was not previously known: If the end invariants of N are the stable and unstable laminations of a pseudo-Anosov map $\varphi : S \to S$, then N is the infinite cyclic cover of the hyperbolic mapping torus of φ.

The condition that S be closed, rather than finite type, is not essential, but an extension to the case with cusps would require a reworking of [44].

We can also improve a little on what is known about the Bers/Thurston conjecture that every Kleinian surface group is a limit of quasifuchsian groups:

Corollary 2 (Bers Density for bounded geometry) *Any Kleinian surface group with bounded geometry and no parabolics is a limit of quasifuchsian groups.*

Given ρ with bounded geometry, Thurston's Double Limit Theorem (see Thurston [51] and Ohshika [49]) and a continuity property for the lengths of measured laminations (see Brock [10, 8]) can be used to find a sequence ρ_i of quasifuchsian representations converging to a representation ρ_∞ which has the same ending invariants as ρ. Corollary 1 then implies that ρ and ρ_∞ are conjugate, yielding Corollary 2.

Recent work of Rafi [50] applies the results of this paper to give a relationship between bounded geometry for hyperbolic manifolds and Teichmüller geodesics. A complete Teichmüller geodesic g in the Teichmüller space $\mathcal{T}(S)$ is uniquely determined by its endpoints $g_\pm$ in the projective measured lamination space $\mathcal{PML}(S)$. We say that g has bounded geometry if its projection to the Moduli space of S has compact closure.

Theorem (Rafi) *If ρ is a Kleinian surface group and g is a Teichmüller geodesic such that $g_\pm$ and $\nu_\pm(\rho)$ are the same up to forgetting the measures, then ρ has bounded geometry if and only if g does.*

The "only if" direction was proved in [44]. Rafi shows that the bounded geometry condition for a geodesic g is equivalent to a bound on $\sup_Y d_Y(g_+, g_-)$, so that this result then follows from the Bounded Geometry Theorem.

Another consequence of Rafi's work and the work in this paper is the following (see [50] for details, and Farb-Mosher [20] for related results). A Teichmüller geodesic g determines a canonical metric on $S \times \mathbb{R}$ as follows:

For $t \in \mathbb{R}$ the metric on $S \times \{t\}$ is the metric determined by the quadratic differential of g at the point $g(t)$. The distance between $S \times \{t\}$ and $S \times \{s\}$ is $|s - t|$, and the connection between slices is given by Teichmüller maps. This metric lifts to a metric on $\widetilde{S} \times \mathbb{R}$.

Theorem (Rafi) *Let S be a closed surface and g a Teichmüller geodesic in $\mathcal{T}(S)$. The lifted induced metric of g on $\widetilde{S} \times \mathbb{R}$ is quasi-isometric to hyperbolic 3-space if and only if g has bounded geometry.*

The Bounded Geometry Theorem also gives us a richer class of manifolds known to have bounded geometry. Previously, aside from geometrically finite manifolds for which bounded geometry is automatic, all examples came from iteration of pseudo-Anosov or partial pseudo-Anosov mapping classes acting on quasi-Fuchsian space or a Bers slice (see Thurston [51], McMullen [40] and Brock [9]). The class of Teichmüller geodesics satisfying bounded geometry is larger than this – in particular uncountable (this is a consequence of [20]). Thurston's Double Limit Theorem [51] yields manifolds whose end invariants correspond to the endpoints of these geodesics, and Rafi's theorem guarantees that in fact they have bounded geometry.

In spite of this one should note that bounded geometry is a rare condition. In the boundary of a Bers slice, for example, there is a topologically generic (dense G_δ) set of representations each of which has arbitrarily short elements (see McMullen [39, Cor. 1.6], and Canary-Culler-Hersonsky-Shalen [13] for generalizations). We hope that some of the techniques introduced in this paper, when used more carefully, will yield information about general geometrically infinite hyperbolic 3-manifolds as well.

Some remarks on the technical conditions of the theorem: The condition of no accidental parabolics means that all cusps of N correspond to cusps of S. Although accidental parabolics can be allowed, the statement and the proof become more awkward, and we prefer to defer this case to a later paper.

The restriction of the second conclusion of the theorem to curves that are not externally short rules out those curves that are already short in the domain of discontinuity, if any (see §2.1). There are a bounded number of such curves for any ρ, and we remark that at any rate they are detected by the end invariants.

The theorem can be generalized in the standard ways to more complicated manifolds, by considering their boundary subgroups. In the interests of brevity we omit this discussion as well.

Outline of the argument

See also the research announcement [48] for a more informal, and perhaps more readable, account of the argument. See Section 2 for notation and definitions.

One direction of the theorem,

$$\inf_{\gamma} \ell_\rho(\gamma) > \epsilon \quad \Longrightarrow \quad \sup_Y d_Y(\nu_+, \nu_-) < K$$

for K depending only on ϵ and S and γ varying over the non-parabolic elements of $\pi_1(S)$, has already been established in [47]. In fact a somewhat stronger statement is proved, that for each individual subsurface Y a lower bound on $\ell_\rho(\partial Y)$ implies an upper bound on $d_Y(\nu_-, \nu_+)$.

We will discuss the proof of the opposite direction,

$$\sup_Y d_Y(\nu_+, \nu_-) < K \quad \Longrightarrow \quad \inf_{\gamma} \ell_\rho(\gamma) > \epsilon$$

for ϵ depending only on K and S, with the infimum over γ that are not externally short.

For each $\gamma \in \pi_1(S)$ which is not externally short, we will find a lower bound on $\ell_\rho(\gamma)$ by finding an upper bound on the radius of its Margulis tube $\mathbf{T}_\gamma(\epsilon_0)$. The main idea of the proof involves an interplay between combinatorial structure on surfaces, namely elementary moves between pants decompositions, and geometry of pleated surfaces, particularly homotopies between them. The first of these is controlled by hyperbolicity of the complex of curves and related results in [37, 36], and the second is controlled by Thurston's Uniform Injectivity Theorem [51] and its consequences.

We will construct a map $H : S \times [0, n] \to N_\rho$, whose structure is determined by the end invariants $\nu_\pm$, and which has the following properties.

1. H covers the Margulis tube $\mathbf{T}_\gamma(\epsilon_0)$ with degree 1.
2. For each integer $i \in [0, n]$, $H_i = H|_{S\times\{i\}}$ is a pleated surface, mapping a pants decomposition P_i geodesically.
3. P_i and P_{i+1} are related by an *elementary move*, and the block $H|_{S\times[i,i+1]}$ has uniformly bounded tracks $H(\{x\} \times [i, i+1])$, except in the special case that γ is a component of P_i or P_{i+1}.
4. There is a uniform M, depending only on $\sup_Y d_Y(\nu_+, \nu_-)$, so that H is disjoint from $\mathbf{T}_\gamma(\epsilon_0)$ except on a segment $S \times [j, j + M]$.

Clearly, except for the special case of property (3), these properties suffice to bound the diameter of $\mathbf{T}_\gamma(\epsilon_0)$. If the special case occurs and γ is a component of one of the P_i, then there is a solid torus in $S \times [0, n]$ that is mapped over $\mathbf{T}_\gamma(\epsilon_0)$, and an additional argument must be given, using the projection coefficient $d_\gamma(\nu_-, \nu_+)$ and a lemma on shearing (Section 6) to bound the size of this solid torus.

The initial and final pants decompositions P_0, P_n are chosen using the end invariants; for example in the quasifuchsian case they are minimal-length pants decompositions on the convex hull boundary, and in the degenerate case they are chosen sufficiently close to the ending laminations to guarantee the covering property (1). We connect P_0 to P_n with a special sequence $\{P_i\}$ called a *resolution sequence*, whose properties follow from the work in Masur-Minsky [37], and are described in Section 5.

Property (3) is established in Section 4, where we make use of the consequences of Thurston's Uniform Injectivity theorem (discussed in §2.4).

Property (4) is established using a quasiconvexity property of the subset of the complex of curves spanned by curves of bounded ρ-length (Section 3). This together with the fact that the resolution sequence $\{P_i\}$ follows a geodesic in the complex of curves serves to bound the part of the sequence that can come near the Margulis tube of γ. The dependence of the bound on $\sup_Y d_Y(\nu_-, \nu_+)$ is one of the properties of the resolution sequence.

In addition to the notation and background material discussed in Section 2, some well-known and/or straightforward definitions and constructions in hyperbolic geometry are also included in the appendix §8.

2. Preliminaries and notation

The following notation and constants will be used throughout the paper.

- S: a surface of genus g, with p punctures, admitting a finite-area complete hyperbolic metric. (We may fix such a metric for reference, but its choice is not important).
- $\mathcal{D}(S)$: The space of discrete, faithful representations of $\pi_1(S)$ into $\mathrm{PSL}_2(\mathbb{C})$ that are *type-preserving*, meaning that the image of any element representing a puncture is parabolic. We call these "Kleinian surface groups" for short.
- $\mathcal{D}_{\mathrm{np}}(S)$: The subset of $\mathcal{D}(S)$ consisting of representations without *accidental parabolics*, meaning that *only* elements representing punctures have parabolic images.
- N_ρ: The quotient manifold $\mathbb{H}^3/\rho(\pi_1(S))$ for a representation $\rho \in \mathcal{D}(S)$. N_ρ is homeomorphic to $S \times \mathbb{R}$ (Bonahon [5] and Thurston [52]), and comes equipped with a homotopy equivalence $S \to N_\rho$ determined by ρ.
- $l(\alpha)$: The length of a curve or arc α. A subscript l_σ usually denotes a metric, or sometimes an ambient space as in l_N.
- $\ell(\alpha)$: The minimum of l over the free homtopy class of a closed curve α, or a homotopy class rel endpoints of an arc. Again a subscript denotes an ambient metric or space, and in addition ℓ_ρ for $\rho \in \mathcal{D}(S)$ denotes the length of the shortest representative of α in the manifold N_ρ, or equivalently the translation length of any isometry in the conjugacy class determined by $\rho(\alpha)$.
- ϵ_0: A Margulis constant for $\mathbb{H}^2$ and $\mathbb{H}^3$, chosen as in §2.5.
- ϵ_1: A number in $(0, \epsilon_0)$ satisfying the conditions in §2.5.
- L_1: A number such that every hyperbolic structure on S admits a pants decomposition of total length at most L_1, and furthermore that

for any geodesic α in S of length at least ϵ_1 such a pants decomposition exists that intersects α. See §2.6.
- K_0: the bilipschitz constant in Sullivan's theorem (§2.1).

2.1. **End invariants and convex hulls.** For additional discussions of end invariants in the setting of this paper see [45, 47, 48], as well as Ohshika [49] and of course Thurston [52] and Bonahon [5]. We will recall briefly their relevant properties in the case of no accidental parabolics, i.e. when $\rho \in \mathcal{D}_{\mathrm{np}}(S)$.

Let $\mathcal{T}(S)$ denote the Teichmüller space of S (see e.g. Abikoff [1] or Gardiner [21]). Let $\mathcal{GL}(S)$ denote the space of geodesic laminations on S and let $\mathcal{ML}(S)$ be the space of measured geodesic laminations, i.e. geodesic laminations equipped with transverse invariant measures of full support (see Bonahon [6] or Casson-Bleiler [18]). On $\mathcal{GL}$ we have the topology of Hausdorff convergence of closed subsets of S, and on $\mathcal{ML}$ there is a topology coming from weak-* convergence of the measures on transversals. We will need the fact that if $\lambda_i \to \lambda$ in $\mathcal{ML}$ then the supports of λ_i converge in $\mathcal{GL}$, after restriction to a subsequence, to a lamination containing the support of λ.

Let $\mathcal{UML}(S)$ denote the "unmeasured laminations," or the quotient space of $\mathcal{ML}(S)$ by the equivalence relation that forgets measures. This is a non-Hausdorff topological space, but it has a Hausdorff subset $\mathcal{EL}(S)$, which is the image of the "filling" measured laminations: those laminations μ with the property that μ intersects λ nontrivially for any $\lambda \in \mathcal{ML}(S)$ whose support is not equal to the support of μ. (See Klarreich [33, §7] for a proof. Note that [33] uses the equivalent language of measured foliations rather than laminations).

The invariants ν_+ and ν_- for $\rho \in \mathcal{D}_{\mathrm{np}}(S)$ lie either in $\mathcal{T}(S)$ or in $\mathcal{EL}(S)$. Let Q_ρ denote the union of standard cusp neighborhoods for the cusps of N_ρ, so that $N_\rho \setminus Q_\rho$ has two ends, which we call e_+ and e_-. (There is an orientation convention for deciding which is which, that will not concern us here).

Let $C(N_\rho)$ denote the convex hull of N_ρ. If there is a neighborhood of the end e_s (where s denotes $+$ or $-$) that is disjoint from $C(N_\rho)$ we say e_s is geometrically finite, and there is a boundary component $\partial_s C(N_\rho)$ which is in fact the image of a pleated surface $S \to N$, that bounds a neighborhood of e_s. In the compactification $\bar{N}_\rho = \mathbb{H}^3 \cup \Omega_\rho/\rho(\pi_1(S))$ (where Ω_ρ is the domain of discontinuity in the Riemann sphere), there is a component isotopic to $\partial_s C(N_\rho)$, which inherits a conformal structure from Ω_ρ. This conformal structure, seen as a point in $\mathcal{T}(S)$, is ν_s.

The hyperbolic structure on $\partial_s C(N_\rho)$ yields a point ν'_s in $\mathcal{T}(S)$. A theorem of Sullivan (proof in Epstein-Marden [19]) states that ν'_s and ν_s differ by a uniformly bilipschitz distortion. Let K_0 denote this bilipschitz constant.

If e_s is not geometrically finite it is geometrically infinite, and ν_s is a lamination in $\mathcal{EL}(S)$, with the following properties:

1. There exists a sequence of simple closed curves α_i in S, whose geodesic representatives α_i^* in N_ρ are eventually contained in any neighborhood of e_s (we say "α_i^* exit the end e_s"), and whose lengths are bounded.
2. For any sequence of simple closed curves β_i whose geodesics β_i^* exit the end e_s, $\beta_i \to \nu_s$ in $\mathcal{UML}(S)$.

Externally short curves. Call a curve γ in S *externally short*, with respect to a representation ρ, if it is either parabolic or if at least one of $\nu_\pm$ is a hyperbolic structure with respect to which γ has length less than ϵ_0.

2.2. **Complexes of arcs and curves:** Let Z be a compact surface, possibly with boundary. If Z is not an annulus, define $\mathcal{A}_0(Z)$ to be the set of essential homotopy classes of simple closed curves or properly embedded arcs in Z. Here "homotopy class" means free homotopy for closed curves, and homotopy rel ∂Z for arcs. "Essential" means the homotopy class does not contain a constant map or a map into the boundary.

If Z is an annulus, we make the same definition except that homotopy for arcs is rel endpoints.

If Z is a non-annular surface with punctures as well as boundaries, we make a similar definition, in which arcs are allowed to terminate in punctures.

We can extend $\mathcal{A}_0$ to a simplicial complex $\mathcal{A}(Z)$ by letting a k-simplex be any $(k+1)$-tuple $[v_0, \ldots, v_k]$ with $v_i \in \mathcal{A}_0(Z)$ distinct and having pairwise disjoint representatives.

Let $\mathcal{A}_i(Z)$ denote the i-skeleton of $\mathcal{A}(Z)$, and let $\mathcal{C}(Z)$ denote the subcomplex spanned by vertices corresponding to simple closed curves. This is the "complex of curves of Z", originally introduced by Harvey [26] (see Harer [24, 25] and Ivanov [28, 30, 32] for subsequent developments).

If we put a path metric on $\mathcal{A}(Z)$ making every simplex regular Euclidean of sidelength 1, then it is clearly quasi-isometric to its 1-skeleton. It is also quasi-isometric to $\mathcal{C}(Z)$ except in a few simple cases when $\mathcal{C}(Z)$ has no edges. When Z has no boundaries or punctures, of course $\mathcal{A}(Z) = \mathcal{C}(Z)$. Note that if Z is a torus with one hole or a sphere with four holes then this definition does not agree with the one we used in previous papers (e.g. [46, 47]) – in particular $\mathcal{C}(Z)$ is 0-dimensional, whereas it was 1-dimensional before. However the 0-skeletons are the same in both definitions, and the distance function restricted from $\mathcal{A}(Z)$ agrees up to bilipschitz distortion with the distance function of the earlier version.

Note that $\mathcal{A}_0(S)$ can identified with a subset of the geodesic lamination space $\mathcal{GL}(S)$. Let $Y \subset S$ be a proper essential closed subsurface (all boundary curves are homotopically nontrivial, and Y is not deformable into a cusp). We have a "projection map"

$$\pi_Y : \mathcal{GL}(S) \to \mathcal{A}(\widehat{Y}) \cup \{\emptyset\}$$

defined as follows: there is a unique cover of S corresponding to the inclusion $\pi_1(Y) \subset \pi_1(S)$, to which we can append a boundary using the circle at infinity of the universal cover of S to yield a surface $\widehat{Y}$ homeomorphic to Y (take the quotient of the compactified hyperbolic plane minus the limit set of $\pi_1(Y)$). Any lamination $\lambda \in \mathcal{GL}(S)$ lifts to this cover. If this lift has leaves that are either non-peripheral closed curves or essential arcs that terminate in the boundary components or the cusps of $\widehat{Y}$, these components determine a simplex of $\mathcal{A}(\widehat{Y})$ and we can take $\pi_Y(\lambda)$ to be its barycenter. If there are no such components then $\pi_Y(\lambda) = \emptyset$. Note that $\pi_Y(\lambda) \neq \emptyset$ whenever λ contains a leaf that is either a closed non-peripheral curve in Y, or intersects ∂Y essentially.

A version of this projection also appears in Ivanov [31, 29].

If $\beta, \gamma \in \mathcal{GL}(S)$ have non-empty projections π_Y, we denote their "Y-distance" by:

$$d_Y(\beta, \gamma) \equiv d_{\mathcal{A}(\widehat{Y})}(\pi_Y(\beta), \pi_Y(\gamma)).$$

Note that $\mathcal{A}(\widehat{Y})$ can be identified naturally with $\mathcal{A}(Y)$, except when Y is an annulus, in which case the pointwise correspondence of the boundaries matters. In the annulus case d_Y measures relative twisting of arcs determined rel endpoints, and in all other cases we ignore twisting on the boundary of $\widehat{Y}$. If α is the core curve of an annulus Y we will also write

$$d_\alpha = d_Y.$$

Note that, if Y is a three-holed sphere (pair of pants), $\mathcal{A}(Y)$ is a finite complex with diameter 1, and there is not much information to be had from π_Y. We will usually exclude three-holed spheres when considering the projection π_Y.

We make a final observation that one can bound $d_{\mathcal{A}(S)}(\beta, \gamma)$, as well as $d_Y(\beta, \gamma)$ when defined, in terms of the number of intersections of β and γ (although there is no bound in the opposite direction). See e.g. Hempel [27].

Elementary moves on pants decompositions. An elementary move on a maximal curve system P is a replacement of a component α of P by α', disjoint from the rest of P, so that α and α' are in one of the two configurations shown in Figure 1.

We indicate this by $P \to P'$ where $P' = P \setminus \{\alpha\} \cup \{\alpha'\}$ is the new curve system. Note that there are infinitely many choices for α', naturally indexed by $\mathbb{Z}$.

In Section 4 we will show how to relate these moves to controlled homotopies between pleated surfaces in N_ρ. In Section 5 we will describe the combinatorial aspects of connecting any two pants decompositions by an efficient sequence of elementary moves.

2.3. **Unmeasured laminations and the boundary of $\mathcal{C}(S)$.** In addition to being the space of ending laminations for $\rho \in \mathcal{D}_{\rm np}(S)$, $\mathcal{EL}(S)$ also has an

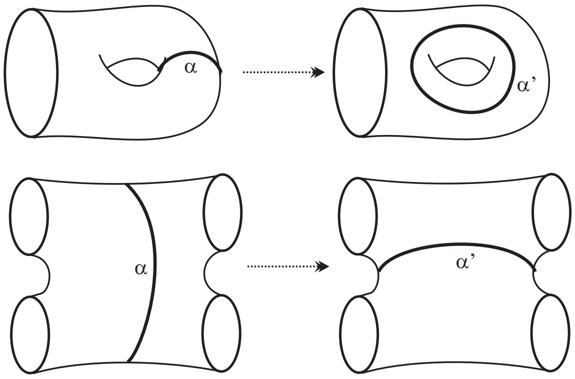

FIGURE 1. The two types of elementary moves.

interpretation as the Gromov boundary $\partial\mathcal{C}(S)$ of the δ-hyperbolic space $\mathcal{C}(S)$ (see e.g. [17, 23, 22, 2, 7] for material on δ-hyperbolicity).

In the following theorem note that $\mathcal{C}_0(S)$ can be considered as a subset of $\mathcal{UML}(S)$.

Theorem 2.1. (Klarreich [33]) *There is a homeomorphism*

$$k : \partial\mathcal{C}(S) \to \mathcal{EL}(S),$$

which is natural in the sense that a sequence $\{\beta_i \in \mathcal{C}_0(S)\}$ converges to $\beta \in \partial\mathcal{C}(S)$ if and only if it converges to $k(\beta)$ in $\mathcal{UML}(S)$.

2.4. **Pleated surfaces and uniform injectivity.** A *pleated surface* is a map $f : S \to N$ together with a hyperbolic metric on S, written σ_f and called the *induced metric*, and a σ_f-geodesic lamination λ on S, so that the following holds: f is length-preserving on paths, maps leaves of λ to geodesics, and is totally geodesic on the complement of λ. Pleated surfaces were introduced by Thurston [52]. See Canary-Epstein-Green [16] for more details.

The set of all pleated surfaces (in fact all maps $S \to N_\rho$) admits a standard equivalence relation, in which $f \sim f\circ h$ if h is a homeomorphism of S isotopic to the identity. Let us refer to this as equivalence *up to domain isotopy.*

If P is a curve or arc system, i.e. a simplex in $\mathcal{A}(S)$, let

$$\mathbf{pleat}_\rho(P)$$

denote the set of pleated surfaces $f : S \to N_\rho$, in the homotopy class determined by ρ, which map representatives of P to geodesics. Thurston observed that such maps always exist provided P has no closed component which is parabolic in ρ. Let $\mathbf{pleat}_\rho$ denote the set of *all* pleated surfaces in the homotopy class of ρ.

Note, if P contains an arc terminating in punctures, the corresponding leaf in the lamination will be infinite and properly embedded in S, its ends exiting the cusps.

In particular, if P is a maximal curve system, or "pants decomposition", $\mathbf{pleat}_\rho(P)$ consists of finitely many equivalence classes, all constructed as follows: Extend P to a triangulation of S with one vertex on each component of P (and a vertex in each puncture, if any) and "spin" this triangulation around P, arriving at a lamination λ whose closed leaves are P and whose other leaves spiral onto P, as in Figure 2, or go out the cusps.

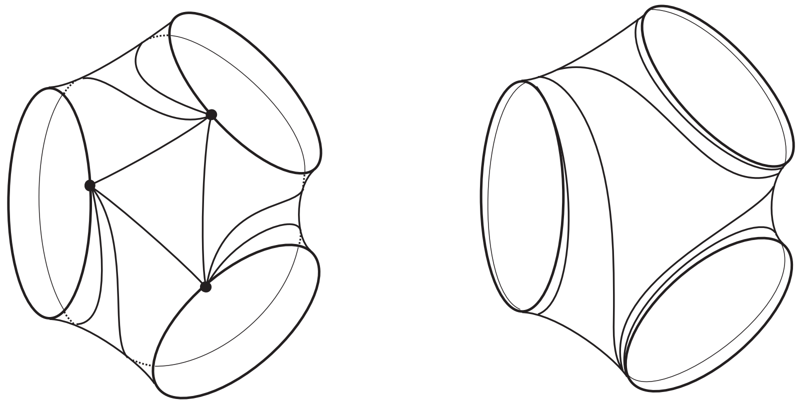

FIGURE 2. The lamination obtained by spinning a triangulation around a curve system. The picture shows one pair of pants in a decomposition.

Uniform Injectivity

Thurston's Uniform Injectivity theorem for pleated surfaces [54] has two corollaries that we will use here. For further discussion and proofs see [42], [47] and also Brock [10].

Bridge arcs. If α is a lamination in S, a *bridge arc* for α is an arc in S with endpoints on α, which is not deformable rel endpoints into α. A *primitive bridge arc* is a bridge arc whose interior is disjoint from α. If σ is a hyperbolic metric on S and τ is a bridge arc for α, let $[\tau]$ denote the homotopy class of τ with endpoints fixed, and for a metric σ let $\ell_\sigma([\tau])$ denote the length of the minimal representative of $[\tau]$.

For a lamination μ and two maps $g, g' \in \mathbf{pleat}_\rho(\mu)$, we say that g and g' are *homotopic relative to* μ if there is a homotopy between them fixing μ pointwise. Lemma 3.3 in [47] guarantees that we can always precompose g' by a homeomorphism isotopic to the identity to obtain a map that is homotopic to g relative to μ.

Let $\mathbf{P}\mathbb{H}^3$ denote the tangent line bundle over $\mathbb{H}^3$. For $g \in \mathbf{pleat}_\rho$ mapping a lamination μ geodesically, and a bridge arc τ of μ, define $d_{\mathbf{p}}(g(\tau))$ as follows: Lift $g(\tau)$ to an arc $\widetilde{g}(\tau)$ in $\mathbb{H}^3$ connecting two leaves of the lift of $g(\mu)$. The endpoints of $\widetilde{g}(\tau)$ and the leaves on which they lie determine two points in $\mathbf{P}\mathbb{H}^3$, and we let $d_{\mathbf{p}}(g(\tau))$ be their distance. In other words, $d_{\mathbf{p}}$ is small when the two leaves are both close together and nearly tangent.

The following strengthening of Thurston's Uniform Injectivity theorem is essentially Lemma 3.4 in [47], which follows from Lemma 2.3 in [42].

Lemma 2.2. (Short bridge arcs) *Fix the surface S. Given $\delta_1 > 0$ there exists $\delta_0 \in (0, \delta_1)$ such that the following holds.*

Let $\rho \in \mathcal{D}(S)$ and let $g \in \mathbf{pleat}_\rho$, mapping a lamination μ geodesically. Suppose that τ a bridge arc for μ which is either primitive, or contained in the ϵ_1-thick part of σ_g. Then

$$d_{\mathbf{p}}(g(\tau)) \leq \delta_0 \implies \ell_{\sigma_g}([\tau]) \leq \delta_1.$$

Moreover if g' is another map in $\mathbf{pleat}_\rho(\mu)$, chosen so it is homotopic to g relative to μ, then

$$\ell_{\sigma_g}([\tau]) \leq \delta_0 \implies \ell_{\sigma_{g'}}([\tau]) \leq \delta_1.$$

The second statement follows from the first, since the short bridge arc in σ_g connects two leaves that contain two segments that remain close to each other for roughly $|\log \delta_0|$, thus the same is true for their images by g (hence by g'). We get a bound on $d_{\mathbf{p}}(g'(\tau))$ and up to revising the constants we have the desired statement.

Efficiency of pleated surfaces. Thurston also used Uniform Injectivity to establish an estimate relating lengths of curves in a pleated surface to their lengths in the 3-manifold. In order to state it we need the *alternation number*

$$a(\lambda, \gamma)$$

where λ is a lamination with finitely many leaves and γ is a simple closed curve (more generally a measured lamination. This quantity, a sort of refined intersection number, is defined carefully in Thurston [51] and Canary [15]. For our purposes we need only the following observations: a finite-leaved lamination consists of finitely many closed leaves, and finitely many infinite leaves whose ends spiral around the closed leaves. If γ crosses only infinite leaves of λ, then $a(\lambda, \gamma)$ is bounded by the number of intersection points with λ.

The following statement is a slight generalization of the theorem proved in [51, Thm 3.3]. Thurston sketches the argument for this generalization, and it also follows from a relative version of the theorem proved in [47].

Theorem 2.3. (Efficiency of pleated surfaces) *Given S and any $\epsilon > 0$, there is a constant $C > 0$ for which the following holds.*

Let $\rho \in \mathcal{D}(S)$ and suppose $g \in \mathbf{pleat}_\rho$ maps geodesically a maximal finite-leaved lamination λ Suppose γ is a measured geodesic lamination in (S, σ_g) which does not intersect any closed curve of λ whose length is less than ϵ. Then

$$\ell_\rho(\gamma) \leq \ell_{\sigma_g}(\gamma) \leq \ell_\rho(\gamma) + Ca(\lambda, \gamma). \tag{2.1}$$

2.5. **Margulis tubes.** We shall denote by $\mathbf{T}_\alpha(\epsilon)$ the ϵ-Margulis tube in N_ρ for a closed curve α in hyperbolic manifold, that is the region where α can be represented with length at most ϵ. We let ϵ_0 be a Margulis constant for 2 and 3 dimensions, meaning that $\mathbf{T}_\alpha(\epsilon_0)$ is always a solid torus neighborhood of a closed geodesic, or a horoball neighborhood of a cusp, and any two such tubes are disjoint. We also choose ϵ_0 sufficiently small that, on a hyperbolic surface, any simple closed geodesic is disjoint from any ϵ_0-Margulis tube but its own.

If $\rho \in \mathcal{D}(S)$ is fixed, we will usually use $\mathbf{T}_\alpha$ to denote $\mathbf{T}_{\rho(\alpha)}$, where α is a conjugacy class in $\pi_1(S)$.

The radius of an ϵ-Margulis tube grows as the length of the core curve shrinks (see Brooks-Matelski [11] and Meyerhoff [41]). We shall need to use the following facts. First, if $\ell_\rho(\alpha) < \epsilon < \epsilon_0$ then the radius of $\mathbf{T}_\alpha(\epsilon_0)$ is at least $\frac{1}{2}\log(\epsilon_0/\epsilon) - c$ for a universal c. Second, the distance from $\mathbf{T}_\alpha(\epsilon/2)$ to $\partial\mathbf{T}_\alpha(\epsilon)$, if the former is non-empty, is uniformly bounded away from 0 and ∞.

Margulis tubes in surface groups. Thurston observed that a constant ϵ_1 exists, depending only on S, so that for any $\rho \in \mathcal{D}(S)$, if a pleated surface $g \in \mathbf{pleat}_\rho$ meets $\mathbf{T}_\alpha(\epsilon_1)$ it can only do so in its own ϵ_0-Margulis tube. Thus α (if it is a primitive element) must be homotopic to the core of this tube in S, and in particular simple. Together with Bonahon's tameness theorem [5], which implies that every point in $C(N_\rho)$ is within a bounded distance of a pleated surface, we have that ϵ_1 can be chosen so that $\ell_\rho(\alpha) < \epsilon_1$ implies that α is simple.

In the remainder of the paper, we fix ϵ_1 so that it has these properties, and in addition $\epsilon_1 < \epsilon_0/K_0$, where K_0 is the constant in Sullivan's theorem (see §2.1) relating $\nu_\pm$ to $\nu'_\pm$. This has the effect that a curve which is not externally short has length at least ϵ_1 in $\nu'_\pm$ as well as $\nu_\pm$.

2.6. **The projection coefficients.** Let us now see how to define the coefficients

$$d_Y(\nu_+, \nu_-)$$

which appear in the main theorem, where $\nu_\pm$ are end invariants for some $\rho \in \mathcal{D}_{\mathrm{np}}(S)$, and Y is any essential subsurface of S.

Using π_Y as above, we can already define this whenever $\nu_\pm$ are laminations. In the case of a geometrically finite end when ν_+ or ν_- are hyperbolic metrics, we can extend this definition as follows:

A theorem of Bers (see [3, 4] and Buser [12]) says that a constant L_1 exists, depending only on the topological type of S, so that for any hyperbolic metric on S there is a maximal curve system (a pants decomposition) with total length bounded by L_1. Moreover, L_1 can be chosen so that, if α is a geodesic of length at least ϵ_1 (the constant defined in §2.5), a pants decomposition can be chosen with length bounded by L_1, and intersecting α essentially. Fix this constant for the remainder of the paper.

Now we define

$$\mathbf{short}(\sigma)$$

to be the set of curve systems of S with total σ-length at most L_1.

Thus e.g. if both ν_+ and ν_- are hyperbolic structures, we may consider distances

$$d_Y(P_+, P_-)$$

for any $P_\pm \in \mathbf{short}(\nu_\pm)$ that both intersect Y essentially, and notice that the numbers obtained cannot vary by more than a uniformly bounded constant (because two different curves in $\mathbf{short}(\sigma)$ have a bounded intersection number, depending only on L_1). We let $d_Y(\nu_+, \nu_-)$ be, say, the minimum over all choices. The case when one of $\nu_\pm$ is a lamination and the other is a hyperbolic metric is handled similarly.

This defines $d_Y(\nu_+, \nu_-)$ for all Y, with the exception of an annulus whose core curve has length less than ϵ_1 in one of ν_+ or ν_-, and a three-holed sphere all of whose boundary curves have this property. The case of three-holed spheres will not make any difference, since at any rate their curve complexes are finite. The case of annuli will require a bit of attention at the end of the proof of the main theorem.

3. Quasiconvexity of the bounded curve set

Let $\mathcal{C}(\rho, L)$ denote the subcomplex of $\mathcal{C}(S)$ spanned by the vertices with ρ-length at most L. In this section we will show that this set has certain quasiconvexity properties. Before stating them we will need to define a map

$$\Pi_\rho : \mathcal{A}(S) \to \mathcal{P}(\mathcal{C}(\rho, L_1))$$

where $\mathcal{P}(X)$ denotes the set of subsets of X. Given $x \in \mathcal{A}(S)$ let P_x be the curve/arc system associated to the smallest simplex containing x. We define

$$\Pi_\rho(x) = \bigcup_{f \in \mathbf{pleat}_\rho(P_x)} \mathbf{short}(\sigma_f).$$

For convenience we will often write $\Pi_\rho(P)$ where P is a curve/arc system.

In the following theorem, a set A in a geodesic metric space is *b-quasiconvex* if every geodesic with endpoints in A is contained in the b-neighborhood of A.

Theorem 3.1. (Quasiconvexity) *For any $\rho \in \mathcal{D}(S)$ and $L \geq L_1$, $\mathcal{C}(\rho, L)$ is B-quasiconvex, where B depends only on L and the topology of S.*

Moreover, if β is a geodesic in $\mathcal{C}(S)$ with endpoints in $\mathcal{C}(\rho, L)$ then

$$d_{\mathcal{C}}(x, \Pi_\rho(x)) \leq B$$

for each $x \in \beta$.

This will follow from the Coarse Projection lemma below, together with the hyperbolicity of $\mathcal{C}(S)$, via the argument in Lemma 3.3.

The coarse projection property

Lemma 3.2. (Coarse Projection) *For any $\rho \in \mathcal{D}(S)$ the map Π_ρ satisfies the following:*

1. *(Coarse Lipschitz) If $d_{\mathcal{A}}(x, y) \leq 1$ then*
$$\operatorname{diam}_{\mathcal{C}}(\Pi_\rho(x) \cup \Pi_\rho(y)) \leq b.$$
2. *(Coarse idempotence) If $x \in \mathcal{C}(\rho, L_1)$ then*
$$d_{\mathcal{C}}(x, \Pi_\rho(x)) = 0$$

where b depends only on S.

Note that here distance between sets (as in property (2)) is the minimal distance, whereas $\operatorname{diam}_{\mathcal{C}}$ controls maximal distances.

Proof. Property (2) (Coarse idempotence) is immediate from the definition, since for any $f \in \mathbf{pleat}_\rho(x)$, x is realized in σ_f with its minimal length, and so is included in $\mathbf{short}(\sigma_f)$.

Note next that there exists $B = B(L_1)$ such that, for any hyperbolic structure σ,
$$\operatorname{diam}_{\mathcal{C}}(\mathbf{short}(\sigma)) < B \tag{3.1}$$
since there is a uniform bound on the intersection number of any two curves of length at most L_1 in the same metric on S.

Now to prove part (1), it clearly suffices to consider $x, y \in \mathcal{A}_0(S)$. The condition $d(x, y) = 1$ means x and y are disjoint, so $x \cup y$ is a curve/arc system, and since
$$\Pi_\rho(x \cup y) \subset \Pi_\rho(x) \cap \Pi_\rho(y),$$
the intersection is non-empty. Thus it will suffice to obtain a bound of the form
$$\operatorname{diam}_{\mathcal{C}}(\Pi_\rho(\gamma)) \leq B \tag{3.2}$$
for any curve/arc system γ. To do this we must compare the short curves in any two different surfaces pleated along γ. Let $f, g \in \mathbf{pleat}_\rho(\gamma)$, and assume (see discussion in §2.4) that they are homotopic relative to γ.

Suppose first that γ meets the (non-cuspidal) ϵ_1-thin part of σ_f. Then its f-image meets the ϵ_1-thin part of N_ρ, and hence so does its g-image (since they agree). Let α be the core curve of this component of the thin part. The length of α in σ_g must also be at most ϵ_0 (by the choice of ϵ_1 in §2.5), and in particular
$$\alpha \in \mathbf{short}(\sigma_f) \cap \mathbf{short}(\sigma_g).$$
Together with (3.1), this implies a bound on
$$\operatorname{diam}_{\mathcal{C}}(\mathbf{short}(\sigma_f) \cup \mathbf{short}(\sigma_g)).$$
The bound (3.2) follows.

Now suppose that γ stays in the ϵ_1-thick part of σ_f, except possibly for cusps (γ may contain an infinite leaf terminating in a cusp).

By the second part of Lemma 2.2 (Short bridge arcs), there exists an $\epsilon > 0$ so that, if τ is a bridge arc for γ in the ϵ_1-thick part of σ_f and whose σ_f-length is at most ϵ, then τ is homotopic rel endpoints to an arc of σ_g-length ϵ_0.

Given this ϵ, we may construct a homotopically non-trivial curve γ_ϵ in the ϵ_1-thick part of σ_f, whose σ_f-length is at most a constant L_2, and which is composed of at most two arcs of γ and at most two primitive bridge arcs of length ϵ or less. (The proof is a standard argument, which we sketch in Lemma 8.5).

The bridge arcs can be homotoped to have σ_g length at most ϵ_0, and hence γ_ϵ can be realized in σ_g with length at most $L_2 + 2\epsilon_0$. In each surface this bounds its $\mathcal{C}$-distance from the curves of length L_1, and together with (3.1) we again obtain a bound on

$$\mathrm{diam}_{\mathcal{C}}(\mathbf{short}(\sigma_f) \cup \mathbf{short}(\sigma_g)),$$

and the desired bound (3.2) follows. ☐

Proof of Quasiconvexity

The Quasiconvexity theorem follows from the Coarse Projection lemma via the following standard argument, which has its roots in the proof of Mostow's rigidity theorem (the difference between this and the standard argument is the need to consider two projections, to a geodesic and to the candidate quasiconvex set, whereas the standard argument involves only projection to a geodesic).

Lemma 3.3. *Let X be a δ-hyperbolic geodesic metric space and $Y \subset X$ a subset admitting a map $\Pi : X \to Y$ which is coarse-Lipschitz and coarse-idempotent. That is, there exists $C > 0$ such that*

- *If $d(x, x') \leq 1$ then $d(\Pi(x), \Pi(x')) \leq C$, and*
- *If $y \in Y$ then $d(y, \Pi(y)) \leq C$.*

Then Y is quasi-convex, and furthermore if g is a geodesic in X whose endpoints are within distance a of Y then

$$d(x, \Pi(x)) \leq b$$

for some $b = b(a, \delta, C)$, and every $x \in g$.

Proof. The condition of δ-hyperbolicity for X implies that for any geodesic g (finite or infinite) the closest-points projection $\pi_g : X \to g$ is coarsely contracting in this sense: If $x \in X$ and $r = d(x, g)$ then the ball $B_r(x)$ has π_g-image whose diameter is bounded by a constant b_0 depending only on δ. (This is an easy exercise in the definitions – see e.g. [23, 22, 2, 7]. Indeed this condition for all geodesics g implies δ-hyperbolicity [36]).

In the rest of the proof, let a "dotted path" be a sequence $p = \{x_i \in X\}$ with $d(x_i, x_{i+1}) \leq C$, and let its "length" be $l(p) = \sum_i d(x_i, x_{i+1})$. Let

$r > C$. If p is a dotted path in X outside an r-neighborhood of g then the contraction property of the previous paragraph implies

$$l(p) \geq (r - C)\left(\frac{1}{b_0}\operatorname{diam}(\pi_g(p)) - 1\right). \tag{3.3}$$

Now suppose that g has endpoints within a of Y and assume also that $r > a$. Let J be a segment of g whose endpoints ξ, η are within r of $x, y \in Y$, respectively, but whose interior is outside the r-neighborhood of Y. On the concatenation of geodesics from x to ξ, across J to η and back to y, select a sequence of points spaced at most 1 apart and apply Π, obtaining (via the coarse-Lipschitz property) a dotted path p in Y satisfying

$$l(p) \leq C(2r + l(J) + 1)$$

by the coarse-Lipschitz property. Since x and y are distance C from the respective endpoints $\Pi(x)$ and $\Pi(y)$ of p (by coarse idempotence) and $r > C$, the contraction property of π_J implies that $\pi_J(\Pi(x))$ and $\pi_J(\Pi(y))$ are within b_0 of $\pi_J(x) = \xi$ and $\pi_J(y) = \eta$ respectively. It follows that $\operatorname{diam}(\pi_J(p)) \geq l(J) - 2b_0$. Thus, together with (3.3) we obtain

$$(r - C)\left(\frac{1}{b_0}(l(J) - 2b_0) - 1\right) \leq l(p) \leq C(2r + l(J) + 1),$$

and hence

$$\left(\frac{r - C}{b_0} - C\right) l(J) \leq 3(r - C) + C(2r + 1).$$

Now if we choose r so that $\frac{r-C}{b_0} - C \geq 1$, we obtain an upper bound

$$l(J) \leq 3(r - C) + C(2r + 1).$$

This bounds by $b_1 = r + \frac{1}{2}(3(r - C) + C(2r + 1))$ the maximum distance from a point in J to Y. Since this applies to every excursion of g from the r-neighborhood of Y, we conclude that Y is b_1-quasi-convex.

Now let $x \in g$ be any point. We have the bound $d(x, Y) \leq b_1$. Let $y \in Y$ be a nearest point to x. We have $d(y, \Pi(y)) \leq C$ by coarse idempotence. Now applying coarse Lipschitz to the path from y to x, whose length is at most b_1, we find that $d(\Pi(x), \Pi(y))$ is at most $C(b_1 + 1)$. Finally by the triangle inequality we obtain a bound on $d(x, \Pi(x))$. □

To apply this lemma to our setting, we recall first that in [36] we proved that $\mathcal{C}(S)$, and hence $\mathcal{A}(S)$, is δ-hyperbolic. Our map Π_ρ has images that are subsets of $\mathcal{C}(S)$ rather than single points, but this can easily be remedied by choosing any method at all to select a single point from each set $\Pi_\rho(x)$. Lemma 3.2 implies that the resulting map has the properties required in Lemma 3.3. □

4. Elementary moves on pleated surfaces

In this section we will show how to realize an elementary move between two pants decompositions P_0 and P_1 of S as a controlled homotopy between pleated surfaces in $\mathbf{pleat}_\rho(P_0)$ and $\mathbf{pleat}_\rho(P_1)$. Lemma 4.1 (Homotopy bound) will show that two surfaces that are "good" with respect to a single pants decomposition admit a controlled homotopy. Lemma 4.2 (Halfway surfaces) shows that a pleated surface exists which is "good" for both P_0 and P_1 simultaneously. Thus we can concatenate a controlled homotopy from a surface in $\mathbf{pleat}_\rho(P_0)$ to the halfway surface, with one from the halfway surface to a surface in $\mathbf{pleat}_\rho(P_1)$.

We begin with some definitions. Let $\mathbf{collar}(\gamma, \sigma)$ denote the standard collar for γ in the surface S with metric σ, as defined in Section 8. Similarly define $\mathbf{collar}(P, \sigma)$ for a curve system P.

Good maps: If P is a curve system in S, we let

$$\mathbf{good}_\rho(P, C)$$

denote the set of pleated maps $g : S \to N_\rho$ in the homotopy class determined by ρ, such that

$$\ell_{\sigma_g}(\gamma) \leq \ell_\rho(\gamma) + C$$

for all components γ of P. Note that

$$\mathbf{pleat}_\rho(P) \subset \mathbf{good}_\rho(P, C)$$

for any $C \geq 0$.

Good homotopies: Let $f, g \in \mathbf{pleat}_\rho$ and let P be a curve system. We say that f and g admit a K-*good homotopy* with respect to P if there exists a homotopy $H : S \times [0, 1] \to N_\rho$ such that the following holds:

1. $H_0 \sim f$ and $H_1 \sim g$ up to domain isotopy (§2.4).
2. Denoting by σ_i the induced metric by H_i for $i = 0, 1$,
$$\mathbf{collar}(P, \sigma_0) = \mathbf{collar}(P, \sigma_1).$$
We henceforth omit the metric when referring to these collars.
3. The metrics σ_0 and σ_1 are locally K-bilipschitz outside $\mathbf{collar}(P)$.
4. Let P_0 denote the subset of P consisting of curves γ with $\ell_\rho(\gamma) < \epsilon_0$. The tracks $H(p \times [0, 1])$ are bounded in length by K when $p \notin \mathbf{collar}(P_0)$.
5. For each $\alpha \in P_0$, the image $H(\mathbf{collar}(\alpha) \times [0, 1])$ is contained in a K-neighborhood of the Margulis tube $\mathbf{T}_\alpha(\epsilon_0)$.

The homotopy bound lemma

Lemma 4.1. (Homotopy bound) *Given C there exists K so that for any $\rho \in \mathcal{D}_{\mathrm{np}}(S)$ and maximal curve system P, if*

$$f, g \in \mathbf{good}_\rho(P, C)$$

then f and g admit a K-good homotopy with respect to P.

Proof. Let us first give the proof in the case that S is a closed surface. At the end we will remark on the changes necessary to allow cusps.

Since the σ_f and σ_g lengths of the components of P differ by at most an additive constant C, Lemma 8.2 (applied to each component of $S \setminus P$) gives us a homeomorphism $\varphi : S \to S$ isotopic to the identity, which takes $\mathbf{collar}(P, \sigma_f)$ to $\mathbf{collar}(P, \sigma_g)$ and is locally K-bilipschitz in its complement, with K depending only on C. Moreover arclengths on $\partial\,\mathbf{collar}(P, \sigma_f)$ are *additively* distorted in a bounded way: if $\alpha \subset \partial\,\mathbf{collar}(P, \sigma_f)$ is any arc then $|l_{\sigma_f}(\alpha) - l_{\sigma_g}(\varphi(\alpha))| \leq K$. After replacing g with $g \circ h$, we may assume $\mathbf{collar}(P, \sigma_f) = \mathbf{collar}(P, \sigma_g)$ (and henceforth denote it just $\mathbf{collar}(P)$), and that σ_f and σ_g are locally K-bilipschitz off $\mathbf{collar}(P)$, and have bounded additive length distortion on $\partial\,\mathbf{collar}(P)$.

Now let $H : S \times [0,1] \to N_\rho$ be the homotopy between f and g whose tracks $H|_{\{x\}\times[0,1]}$ are geodesics parameterized at constant speed. (H exists and is unique as a consequence of negative curvature and the fact that $\pi_1(S)$ is non-elementary, see e.g. [42]).

We will bound the tracks of H on succesively larger parts of the surface.

Let Y denote a component of $S \setminus P$, and let $Y_0 = Y \setminus int(\mathbf{collar}(\partial Y))$. We first remark that, in either σ_f or σ_g, the length of any boundary component γ of Y_0 is at most 2 more than its corresponding geodesic in S, by (8.1), and this in turn is bounded by the assumption that $f, g \in \mathbf{good}_\rho(P, C)$. Thus we have

$$l_\sigma(\gamma) \leq \ell_\rho(\gamma) + C + 2 \tag{4.1}$$

for $\sigma = \sigma_f$ or σ_g.

Bounds on the tripods. Let Δ be an essential tripod in Y_0 with legs of σ_f-length bounded by δ, as given by Lemma 8.1 in the Appendix. Let X be a component of the preimage of Δ in the universal cover $\widetilde{S}$. Lift H to $\widetilde{H} : \widetilde{S} \times [0,1] \to \mathbb{H}^3$.

We claim there is a uniform bound on the lengths of the tracks $\widetilde{H}(\{x\} \times [0,1])$ for $x \in X$.

The image $\widetilde{H}(X \times \{0\})$ connects three lifts of boundary curves of Y_0, each invariant by a primitive deck translation γ_i in $\rho(\pi_1(S))$ ($i = 1,2,3$). Let L_i be the axis of γ_i (see Figure 3). If γ_i has translation length less than $\epsilon_1/2$, define N_i to be the lift of the corresponding ϵ_0-Margulis tube in N_ρ. Otherwise define $N_i = L_i$.

Each image (under f or g) of a boundary curve of Y_0 admits a homotopy to its geodesic representative in N, and because of the length bound (4.1), Lemma 8.4 tells us that this homotopy can, in uniformly bounded distance, be made to reach either the geodesic or its ϵ_0-Margulis tube if it is short. Lifting to the universal cover, we conclude that the endpoints of $\widetilde{H}(X \times \{j\})$ ($j = 0,1$) are within bounded distance of the corresponding N_i.

Since we have a uniform diameter bound $d_0 = 2K\delta$ on $\widetilde{H}(X \times \{0\})$ and $\widetilde{H}(X \times \{1\})$, we find that

$$\widetilde{H}(X \times \{j\}) \subset \mathcal{N}_{d_1}(N_1) \cap \mathcal{N}_{d_1}(N_2) \cap \mathcal{N}_{d_1}(N_3) \tag{4.2}$$

for $j = 0, 1$, with a uniform d_1. Here $\mathcal{N}_r$ denotes an r-neighborhood in $\mathbb{H}^3$.

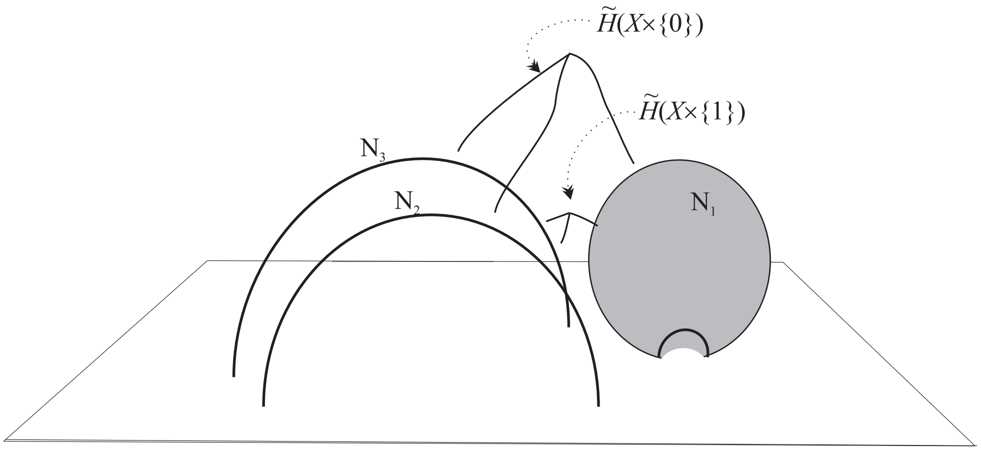

FIGURE 3. The images of $X \times \{0\}$ and $X \times \{1\}$ are within uniform distance of the three invariant sets N_1, N_2, N_3.

The idea now is that, by virtue of discreteness, the triple intersection of the $\mathcal{N}_{d_1}(N_i)$ cannot have very large diameter. Roughly, if the N_i are Margulis tubes their strict convexity implies this, and if they are axes L_i an application of Lemma 2.2 (Short bridge arcs) can be used to forbid the existence of long parallel sections. Once we prove this we will obtain a diameter bound on $\widetilde{H}(X \times [0, 1])$, which will give the desired homotopy bound.

If one of the N_i corresponds to a γ_i of translation length less than $\epsilon_1/2$, let $N_i(\epsilon_1) \subset N_i$ denote a lift of the corresponding ϵ_1-Margulis tube. Note that $\operatorname{dist}(N_i(\epsilon_1), \partial N_i)$ is uniformly bounded away from 0 and ∞. If two of the N_i, say N_1 and N_2, are in this case, then they are disjoint, and hence their convex subsets $N_i(\epsilon_1)$ are a definite distance apart. Applying Lemma 8.6 of the Appendix to $N_1(\epsilon_1)$ and $N_2(\epsilon_1)$, we can deduce that the intersection $\mathcal{N}_{d_1}(N_1) \cap \mathcal{N}_{d_1}(N_2)$ has diameter bounded by some d_2.

For the remainder of the argument, pick a pleated surface $h \in \mathbf{pleat}_\rho(P)$.

Suppose that just N_1 corresponds to a curve of length less than $\epsilon_1/2$. We claim that L_2 is disjoint from $N_1(\epsilon_1)$. Let γ_2^h denote the σ_h-geodesic boundary component of Y corresponding to to γ_2. If L_2 meets $N_1(\epsilon_1)$ then γ_2^h has h-image that meets an ϵ_1-Margulis tube, and hence (by our choice of ϵ_1) γ_2^h meets the ϵ_0-collar of another boundary component of Y (in S with the metric σ_h). By our choice of ϵ_0, a simple geodesic cannot meet an ϵ_0-Margulis tube in a surface unless it is the core of that tube, so this is a contradiction. Now $N_1(\epsilon_1)$ contains the convex subset $N_1(\epsilon_1/2)$ which is definite distance from its boundary, and hence we can apply Lemma 8.6 to $N_1(\epsilon_1/2)$ and L_2, deducing a bound on $\operatorname{diam}(\mathcal{N}_{d_1}(N_1) \cap \mathcal{N}_{d_1}(N_2))$.

Finally suppose all three γ_i have length at least $\epsilon_1/2$. Then in fact the three axes $\{L_i\}$ themselves come within bounded distance $2d_3$. If the intersection of all three $\mathcal{N}_{d_3}(L_i)$ has diameter D then L_1, L_2 and L_3 contain segments of length at least $D - 2d_3$ that remain distance d_3 apart. There are therefore two a-priori constants $d_4, b > 0$ so that there exists a point $p \in \mathbb{H}^3$ which is at most $\epsilon(D) = d_4 e^{-bD}$ from all three L_i, and so that the tangent directions to L_i at the points x_i closest to p are at most $2\epsilon(D)$ apart in $\mathbf{P}\mathbb{H}^3$.

Extend the tripod Δ to a tripod Δ' with endpoints in ∂Y, and let X' be the component of its lift to $\widetilde{S}$ that contains X. The endpoints of X' are mapped by $\widetilde{h}$ to points $y_i \in L_i$ $(i = 1, 2, 3)$. The arcs $[y_i, x_i]$ on L_i pull back to arcs on ∂Y which, if we append them to Δ' and perturb slightly to an embedding in Y, give us a new tripod Δ'' whose endpoints (after lifting and applying $\widetilde{h}$) map to x_i.

Let δ' be the constant given in Lemma 8.1, and let δ_0 be the constant given by Lemma 2.2 after setting $\delta_1 = \delta'$. Note that each pair of legs of Δ'' is a primitive bridge arc for ∂Y, whose h-image is homotopic rel endpoints to an arc of length at most $2\epsilon(D)$. If D is sufficiently large that $2\epsilon(D) < \delta_0$, then Lemma 2.2 tells us that each pair of legs of Δ'' is homotopic to an arc of length at most δ_1 in σ_h. This gives us a triangle in Y with the same vertices as Δ'' whose sidelengths are at most δ_1. Joining its barycenter to its vertices we obtain a new tripod Δ''' all of whose legs are bounded by $\delta_1 = \delta'$. This contradicts Lemma 8.1. We conclude that the triple intersection of (4.2) has diameter at most D.

This diameter bound and (4.2) now imply a uniform bound on the track lengths of $\widetilde{H}$ restricted to X, and hence H restricted to Δ.

Bounds outside the collars. To bound the tracks of H on the rest of $S \setminus \mathbf{collar}(P)$, we first bound them on $\partial\,\mathbf{collar}(P)$.

Let γ be a component of $\partial\,\mathbf{collar}(P)$. Let Y be the component of $S \setminus P$ containing γ, so that γ is in the boundary of $Y_0 \equiv Y \setminus int(\mathbf{collar}(\partial Y))$. Lemma 8.1 gives us a bounded essential tripod Δ in Y_0 with endpoint $x \in \gamma$. By the forgoing discussion we already have a bound, say t_0, on the track length $H(\{x\} \times [0,1])$.

If $\ell_\rho(\gamma) \le 1$, we have the bound (4.1) on the $l_{\sigma_f}(\gamma)$ and $l_{\sigma_g}(\gamma)$. Via the triangle inequality we immediately get a bound on $H(\{y\} \times [0,1])$ for any $y \in \gamma$.

Suppose that $\ell_\rho(\gamma) > 1$. Let $G = H|\gamma \times [0,1]$, and let $\widetilde{G} : \widetilde{\gamma} \times [0,1] \to \mathbb{H}^3$ be a lift to the universal covers. Lift x to $\{x_i | i \in \mathbb{Z}\}$ in $\widetilde{\gamma}$. Let Γ be the geodesic lift of γ^* to $\mathbb{H}^3$ which is invariant by the holonomy of $\widetilde{G}$.

Because of the bound (4.1) on $l_\sigma(\gamma)$ (for $\sigma = \sigma_f$ or σ_g), we may apply Lemma 8.4 to either $\widetilde{G}_0 = \widetilde{G}(\cdot, 0)$ or $\widetilde{G}_1 = \widetilde{G}(\cdot, 1)$. For any $y \in \widetilde{\gamma}$ let $\eta_i = \pi(\widetilde{G}_i(y))$, where $\pi : \mathbb{H}^3 \to \Gamma$ is the orthogonal projection. Then the lemma gives us, first, a bound

$$d(\widetilde{G}_i(y), \eta_i) \leq a_1$$

and $i = 0, 1$.

Let $\xi_{ij} = \pi(\widetilde{G}_i(x_j))$. The main part of Lemma 8.4 tells us that for any $y \in [x_0, x_1]$, its projection η_0 lies a bounded distance a_2 from the point in $[\xi_{00}, \xi_{01}]$ that is at distance $|[x_0, y]|_f$ from ξ_{00}. Here $|\cdot|_f$ and $|\cdot|_g$ denote the lifts of σ_f and σ_g arclength to $\widetilde{\gamma}$. Similarly we have a bound for η_1, $[\xi_{10}, \xi_{11}]$ and $|[x_0, y]|_g$.

The length bound on $G(x \times [0,1])$ tells us that $d(\xi_{0j}, \xi_{1j}) \leq 2a_1 + t_0$. Since $|[x_0, y]|_f$ and $|[x_0, y]|_g$ differ by at most K (via Lemma 8.2, as discussed in the beginning of the proof), we obtain a uniform upper bound on $d(\eta_0, \eta_1)$ and hence on the track length $\widetilde{G}(\{y\} \times [0, 1])$.

This proves our uniform bound on the track lengths for $\partial\,\mathbf{collar}(P)$.

Now for any point z in $S \setminus \mathbf{collar}(P)$ there is an arc β connecting z to some point z' in $\partial\,\mathbf{collar}(P)$ whose length in either σ_f or σ_g is uniformly bounded. The length of $H(z' \times [0, 1])$ is bounded by the above, so we may bound $H(z \times [0, 1])$ using the triangle inequality.

Bounds in the collars. It remains to control H on $\mathbf{collar}(P) \times [0, 1]$. For each component γ of P, $\mathbf{collar}(\gamma) \times [0, 1]$ is a solid torus, and we will control H by considering a meridian curve

$$m = \partial(a \times [0, 1])$$

where a is an arc connecting the boundary components of $\mathbf{collar}(\gamma)$.

Suppose first that $\ell_\rho(\gamma) \geq \epsilon_0$. Then the radius of $\mathbf{collar}(\gamma)$ is bounded above by $s \equiv w(\epsilon_0)$ in either σ_f or σ_g. Let a denote a minimal σ_f-length arc crossing $\mathbf{collar}(\gamma)$, and let us find a bound on $L \equiv l_{\sigma_g}(a)$. (Note that up to the usual equivalence of maps we may assume a is geodesic in both metrics).

Let γ_g be the σ_g-geodesic representative of γ. We first replace a by a homotopic path (with fixed endpoints) $q_0 * a' * q_1$ where a' is the σ_g-orthogonal projection of a to γ_g, and q_i are orthogonal segments to γ_g of length at most s. Thus $l(a') \geq L - 2s$. Now apply Lemma 8.4 to deform $g(a')$, fixing endpoints, to a path $p_0 * a'' * p_1$ in N with p_i of length bounded by a_1, a'' traveling along the geodesic representative of γ,

$$l_N(a'') \geq c_1 l(a') - c_2$$

where c_1 and c_2 depend on the constants in the lemma. It follows that the geodesic representative of the path $g(a)$ in N has length at least $c_3 L - c_4$ for suitable constants c_3, c_4.

Now consider a lift $\widetilde{H}(m)$ of the meridian to $\mathbb{H}^3$. The endpoints of $\widetilde{H}(a \times \{1\})$ are at least $c_3 L - c_4$ apart, and on the other hand the other three legs of m give us an upper bound of $2s+2t_1$, where t_1 is the track bound we have already obtained on H outside the collars. This gives us an upper bound on L, and hence on the length of $H(m)$.

A bound on the tracks of H follows immediately for any point of a, and since we can foliate $\mathbf{collar}(\gamma)$ by arcs such as a, we obtain a bound in the entire collar.

Finally suppose that $\ell_\rho(\gamma) < \epsilon_0$, that is $\gamma \in P_0$. Since $\ell_{\sigma_f}(\gamma)$ and $\ell_{\sigma_g}(\gamma)$ are bounded by $\epsilon_0 + C$, we may foliate $\mathbf{collar}(\gamma)$ by closed curves with this length bound in both metrics (again, up to precomposing the maps by homeomorphisms homotopic to the identity, we may assume the same curves are bounded in both metrics). For each such curve β, the geodesic homotopy H must be contained in $\mathbf{T}_\gamma(\epsilon_0)$ for all but a bounded portion. This is a standard area argument, for the area of the ruled annulus $H(\beta \times [0,1])$ is bounded by the length of its boundary, and on the other hand a long section of the annulus outside of $\mathbf{T}_\gamma(\epsilon_0)$ would have area at least ϵ_0 times its length. (See [52, 5] for similar area arguments).

Thus $H(\mathbf{collar}(\gamma) \times [0,1])$ is contained in a uniformly bounded neighborhood of $\mathbf{T}_\gamma(\epsilon_0)$, which is what we needed to prove.

Surfaces with cusps. It remains to discuss the case when S has cusps. The main difference is that the components Y of $S \setminus P$ may have cusps rather than boundary components. All the arguments go through in the same way, with $\mathbf{collar}(P)$ replaced by the union of $\mathbf{collar}(P)$ with the collars of the cusps. In the complement of these collars we still obtain a bilipschitz relation between σ_g and σ_f. In bounding the tracks on tripods, the neighborhoods N_i must be allowed to be Margulis tubes of parabolic elements when the corresponding boundary component of Y is a cusp. ☐

Remark: The bilipschitz relation between σ_f and σ_g generally breaks down in the collars of P, but a careful consideration of the proof will show that, if $\gamma \in P$ is a component with $\ell_\rho(\gamma)$ bounded both above and below, then we may extend the bilipschitz relation to $\mathbf{collar}(\gamma)$ as well. This is a consequence of the simultaneous bound on the arcs a crossing the collar in both metrics.

Halfway surfaces

Lemma 4.2. (Halfway surfaces) *There exists a $C_1 > 0$ depending only on S, so that if $\rho \in \mathcal{D}(S)$ and $P_0 \to P_1$ is an elementary move then*

$$\mathbf{good}_\rho(P_0, C_1) \cap \mathbf{good}_\rho(P_1, C_1) \neq \emptyset.$$

A map in this intersection is called a *halfway surface* for P_0 and P_1.

Proof. We will construct a pleated surface $g \in \mathbf{pleat}_\rho(P_0 \cap P_1)$ to which we can apply Thurston's Efficiency of Pleated Surfaces (Theorem 2.3).

Let $\alpha_0 \in P_0$ and $\alpha_1 \in P_1$ be the curves exchanged by the elementary move. Let Y be the component of $S - (P_0 \cap P_1)$ containing α_0 and α_1. Note that Y is a 4-holed sphere or 1-holed torus.

Let us describe a lamination λ on Y by first considering its lift to a planar cover. In Figure 4 we indicate the plane $\mathbb{R}^2$ with a small disk removed around every point in the lattice $\mathbb{Z}^2$. When Y is a one-holed torus it is obtained as the quotient of this by the action of $\mathbb{Z}^2$, and when Y is a four-holed sphere it is the quotient by the group generated by $(2\mathbb{Z})^2$ and $-I$. Normalize the picture so that α_0 lifts to lines parallel to the x-axis and α_1 lifts to lines parallel to the y-axis.

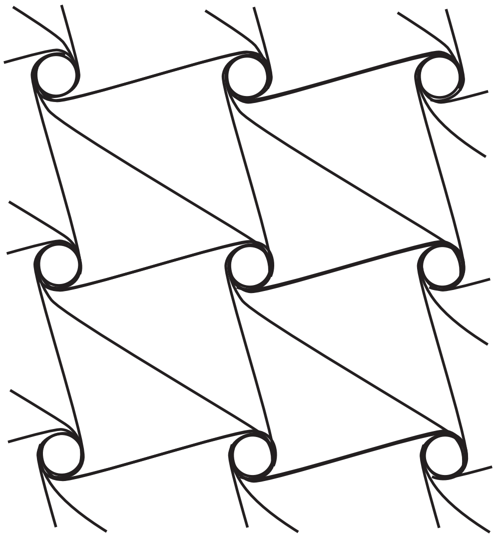

FIGURE 4. The lamination $\widetilde{\lambda}$ in the planar cover of Y.

We have indicated the lamination $\widetilde{\lambda}$ in this cover, which is obtained from a standard triangulation by spinning leftward around the boundary components. The projection of this to Y is λ. This discussion applies equally well when Y has ends that are cusps of S. In this case the leaves that we drew as winding around ∂Y simply go out the corresponding cusp.

Let $g \in \mathbf{pleat}_\rho(P_0 \cap P_1)$ be the pleated surface mapping λ geodesically. An inspection of the diagram gives this bound on the alternation numbers:

$$a(\lambda, \alpha_i) \leq 4$$

for $i = 0, 1$ (in fact it is 2 when Y is a one-holed torus and 4 when Y is a four-holed sphere). And furthermore that α_0 and α_1 cross no closed leaves of λ. Efficiency of Pleated Surfaces then gives us the inequality

$$\ell_{\sigma_g}(\alpha_i) \leq \ell_\rho(\alpha_i) + C$$

for $i = 0, 1$ and a uniform C. Thus $g \in \mathbf{good}_\rho(P_i, C)$ for $i = 0, 1$, and the lemma is proved. $\square$

The following lemma controls the geometry of a halfway surface. It will be instrumental in the last part of the proof of the main theorem.

Lemma 4.3. *Let $P_0 \to P_1$ be an elementary move exchanging α and α', and let $g \in \mathbf{good}_\rho(P_0, C_1) \cap \mathbf{good}_\rho(P_1, C_1)$. Assume that α and α' are realized as geodesics in σ_g, and suppose that*

$$\ell_\rho(\alpha) < \epsilon_1/2$$

but

$$\ell_{\sigma_g}(\alpha) \geq \epsilon_1.$$

Then there is an upper bound C_2, depending only on ϵ_1, C_1 and the topological type of S, so that

$$l_{\sigma_g}(a) \leq C_2$$

for each arc a of $\alpha' \cap \mathbf{collar}(\alpha, \sigma_g)$.

Proof. Let a be an arc of $\alpha' \cap \mathbf{collar}(\alpha, \sigma_g)$ – there may be one or two such arcs. Let $L = l_{\sigma_g}(a) = l_N(g(a))$. Deform a fixing endpoints to a path $q_0 * a' * q_1$ with q_i orthgonal paths from ∂a to α, and a' running along α. The lengths of q_i are at most $w(\epsilon_0)$, and $l(a') \leq L$. Lemma 8.4 allows us to deform $g(a')$, fixing endpoints, to $p_0 * a'' * p_1$ with $l_N(p_i) \leq a_1$ and a'' in the $\epsilon_1/2$-Margulis tube of α, so that

$$l_N(a'') = n\epsilon_1/2 + r'$$

where we write $l(a') = n\ell_{\sigma_g}(\alpha) + r$, $n \in \mathbb{Z}_{\geq 0}$, $r \in [0, \ell_{\sigma_g}(\alpha)]$, and $|r - r'| \leq a_2$. Since $\ell_{\sigma_g}(\alpha) \geq \epsilon_1$, we conclude that $g(a)$ can be deformed to an arc $a''' = g(q_0) * p_0 * a'' * p_1 * g(q_1)$, so that

$$l_N(a''') \leq \tfrac{1}{2}L + c \tag{4.3}$$

with c depending on the previous constants. This is shorter than $g(a)$ by at least $L/2 - c$, so we can shorten the curve $g(\alpha')$ by at least this much, concluding that

$$\ell_\rho(\alpha') \leq \ell_{\sigma_g}(\alpha') - \tfrac{1}{2}L + c.$$

On the other hand $\ell_{\sigma_g}(\alpha') - C_1 \leq \ell_\rho(\alpha')$ since $g \in \mathbf{good}_\rho(\alpha', C_1)$, so we obtain an upper bound on L. ☐

5. Resolution sequences

In Masur-Minsky [37], we show the existence of special sequences of elementary moves that are controlled in terms of the geometry of the complex of curves, and particularly the projections π_Y. First some terminology: if $P_0 \to P_1 \to \cdots \to P_n$ is an elementary-move sequence and β is any vertex of $\mathcal{C}(S)$, we denote

$$J_\beta = \{j \in [0, n] : \beta \in P_j\}.$$

Note that if J_β is an interval $[k, l]$, then the elementary move $P_{k-1} \to P_k$ exchanges some α for β, and $P_l \to P_{l+1}$ exchanges β for some α'. Both α and α' intersect β, and we call them the *predecessor* and *successor* of β, respectively.

In the following theorem, where $\beta_0, \ldots, \beta_m$ is a sequence of vertices in $\mathcal{C}(S)$ we use the notation

$$J_{[s,t]} \equiv \bigcup_{i=s}^{t} J_{\beta_i}.$$

Theorem 5.1. (Controlled Resolution Sequences) *Let P and Q be maximal curve systems in S. There exists a geodesic in $\mathcal{C}_1(S)$ with vertex sequence $\beta_0, \ldots, \beta_m$, and an elementary move sequence $P_0 \to \ldots \to P_n$, with the following properties:*

1. *$\beta_0 \in P_0 = P$ and $\beta_m \in P_n = Q$.*
2. *Each P_j contains some β_i.*
3. *J_β, if nonempty, is always an interval, and if $[s,t] \subset [0,m]$ then*
$$|J_{[s,t]}| \leq K(t-s) \sup_Y d_Y(P,Q)^a,$$
where the supremum is over only those non-annular subsurfaces Y whose boundary curves are components of some P_k with $k \in J_{[s,t]}$.
4. *If β is a curve with non-empty J_β, then its predecessor and successor curves α and α' satisfy*
$$|d_\beta(\alpha, \alpha') - d_\beta(P,Q)| \leq \delta.$$

The constants K, a, δ depend only on the topological type of S. The expression $|J|$ for an interval J denotes its diameter.

The sequence $\{P_i\}$ in this theorem is called a resolution sequence. Such sequences are constructed in [37], using what we call "hierarchies of geodesics" in $\mathcal{C}(S)$. The machinery of [37] is cumbersome to describe fully, so we will give just a rough description of the construction and an indication of how the results of [37] imply Theorem 5.1.

The construction is by an inductive procedure: we begin with what we call a "tight geodesic", which is a sequence m_i of simplices in $\mathcal{C}(S)$ with the first in P and the last in Q, so that any sequence $\beta_i \in m_i$ of vertices yields a geodesic in $\mathcal{C}_1(S)$ joining P to Q. (The m_i satisfy an additional condition which we need not use here). The link of each m_i is itself a curve complex for a subsurface, in which m_{i-1} and m_{i+1} represent two simplices, and we construct a geodesic connecting them. We then repeat, with the complexity of the subsurfaces decreasing at each step. (The actual construction is considerably complicated by the need to take care of endpoints of geodesics correctly, and by the fact that a typical simplex cuts up the surface into several components, in each of which the construction continues independently). The final structure is a collection of geodesics in subsurfaces related by inclusion, and the pants decompositions P_k are obtained by taking "slices": picking a vertex at one level and then inductively adding vertices from the geodesics supported in its complementary subsurfaces. These slices can be "resolved" into the sequence described in Theorem 5.1, where the elementary moves $P_k \to P_{k+1}$ correspond to steps along geodesics in highest-level subsurfaces, which are always one-holed tori or four-holed spheres.

The fact that J_β is always an interval follows from the proof of Proposition 5.4 in [37], in the course of which we establish a monotonicity property of the way a resolution steps through the geodesics in a hierarchy, that implies no curve is ever repeated once it has been traversed.

To obtain the inequality in part (3), note first that the length of $J_{[s,t]}$ is just the sum of the lengths of the geodesics in the highest level subsurfaces meeting the slices based at $\beta_s, \ldots, \beta_t$. Each of these subsurfaces arises from vertices in a lower-level (higher complexity) subsurface, so their number is bounded by the sum of the lengths of the geodesics at the lower level. Continuing inductively, if we have a length bound of B on all the geodesics encountered (and a length $(t-s)$ at the bottom level), we obtain a bound of the form $K(t-s)B^a$, where a bounds the number of levels, which only depends on the topological complexity of S.

Finally, the length of a geodesic supported in a subsurface Y is bounded by a multiple of the projection distance $d_Y(P,Q)$: this is the substance of Lemma 6.2 of [37], which involves a crucial use of the hyperbolicity property of $\mathcal{C}(S)$, from [36]. The bound of part (3) follows.

Part (4) follows from the same construction, which in fact includes annuli and their arc complexes as part of the discussion. The bound is just a restatement of Lemma 6.2 of [37] applied to annuli.

6. The bounded shear lemma

In this section we will develop some estimates of shearing in annuli that will be used near the end of the proof of the main theorem, in Section 7.

We begin with an observation. Let σ be a hyperbolic metric on S, γ a simple geodesic in S, and $B = \mathbf{collar}(\gamma,\sigma)$. Let $\hat{Y}$ be the compactified annular lift associated to γ and $\hat{B}$ the annular lift of B to $\hat{Y}$. Let E denote one of the components of $\hat{Y} \setminus \hat{B}$.

There is only a bounded amount of twisting that a geodesic crossing $\hat{Y}$ can do in E. In fact, let β_1 and β_2 be any two geodesic lines connecting the two boundaries of $\hat{Y}$. We have:

$$d_{\mathcal{A}(E)}(\beta_1 \cap E, \beta_2 \cap E) \leq 4. \tag{6.1}$$

Proof. The width of the collar of γ is the function $w(\ell(\gamma))$ described in the Appendix.

Lift $\hat{B}$ to a w-neighborhood $\widetilde{B}$ of a geodesic lift Γ of γ in $\mathbb{H}^2$. For $i = 1, 2$, β_i lifts to a geodesic arc $\widetilde{\beta}_i$ connecting $\partial\widetilde{B}$ to the circle at infinity. Let p denote the length of its orthogonal projection to Γ. This is largest when $\widetilde{\beta}_i$ is tangent to $\partial\widetilde{B}$, so let us consider this case. The arcs $\widetilde{\beta}_i$ and its projection to Γ form opposite sides of a quadrilateral with three right angles and an ideal vertex, whose two finite sides have lengths w and p (see Figure 6). Hyperbolic trigonometry (Buser [12, 2.3.1(i)]) gives us

$$\sinh(w)\sinh(p) = 1.$$

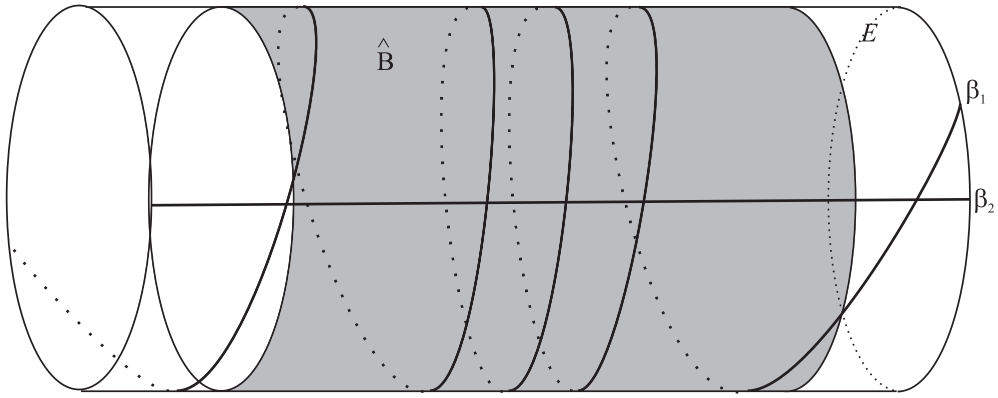

FIGURE 5. The relative twisting $d_{\mathcal{A}(E)}$ of $\beta_1 \cap E$ and $\beta_2 \cap E$ is bounded.

On the other hand we have from the definitions in §8 that $w = \max(w_0/2, w_0 - 1)$ where w_0 satisfies

$$\sinh(w_0)\sinh(\ell/2) = 1.$$

From this we obtain an expression for p/ℓ as a function of w_0, and one can deduce that this quantity is bounded. Indeed the maximum is obtained at $w_0 = 2$, and we have

$$\frac{p}{\ell} < 1.5$$

In other words, $\beta_i \cap E$ travels less than 1.5 times around the annulus, as measured by its orthogonal projection. We deduce that β_1 and β_2 intersect at most 3 times in E. The bound (6.1) follows. □

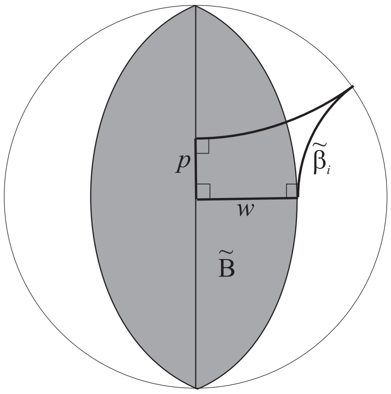

FIGURE 6. The quadrilateral formed by $\widetilde{\beta}_i$ and its projection to the lift of γ

This observation prompts us to define the following measure of shearing of two different metrics outside a collar. Let σ, τ be two hyperbolic metrics on S and let B be an annulus which is equal to both $\mathbf{collar}(\gamma, \sigma)$ and

collar(γ,τ). Lift to $\hat{B}$ in $\hat{Y}$ as above. We let the *shear outside* B of the two metrics be the quantity

$$\sup_{E,\alpha_\sigma,\alpha_\tau} d_{\mathcal{A}(E)}(\alpha_\sigma\cap E,\alpha_\tau\cap E).$$

Here E varies over the two complementary annuli of $\hat{B}$ in $\hat{Y}$, α_σ varies over all σ-geodesics that connect the two boundaries, and α_τ varies over all τ-geodesics that connect the two boundaries. Note that the shear depends on the *pointwise* metrics, and not just their isotopy classes.

The point of having a bound on this shear is that it allows us to measure twisting by restricting to a collar. That is, suppose in the setting above the shear outside B is bounded by D. Then if α_σ and α_τ are any two σ and τ geodesics, respectively, which cross γ, we immediately have

$$|d_{\mathcal{A}(B)}(\alpha_\sigma\cap B,\alpha_\tau\cap B)-d_\gamma(\alpha_\sigma,\alpha_\tau)|\le 2D. \tag{6.2}$$

Note that $d_{\mathcal{A}(B)}$ is a measure of twisting inside the collar which depends on the particular curves we chose, whereas d_γ depends only on homotopy classes.

The main lemma of this section bounds the shear outside a collar for pairs of metrics satisfying a special condition.

Lemma 6.1. (Shear bound) *Suppose R is a subsurface of S which is convex in two hyperbolic metrics σ and τ, and that σ and τ are locally K-bilipschitz in the complement of R. Suppose that one component of R is an annulus B which is equal to both* **collar**(γ,σ) *and* **collar**(γ,τ) *for a certain curve γ.*

Then the shear outside B of σ and τ is bounded above by $\delta_0 K$, where δ_0 depends only on the topological type of S.

Proof. Let $\hat{Y}$ be the compactified annular lift, and let $\hat\sigma$ and $\hat\tau$ be the lifts of σ and τ to $\hat{Y}$. Let $\hat{R}$ denote the lift of R to $\hat{Y}$, and let $\hat{B}$ be the component of $\hat{R}$ which is a homeomorphic lift of B. Let E be one of the components of $\hat{Y}\setminus\hat{B}$. There is a uniform upper bound b to the $\hat\sigma$-length of the shortest arc α connecting $\hat{B}$ to any component D of $\hat{R}$ contained in E. This comes from a standard area bound: Since $\ell_\sigma(\partial B)$ is bounded below by b_0 (see §8), an embedded collar of radius r around B has area at least $b_0 r$. Since the area of S is fixed by the Gauss-Bonnet theorem, r is uniformly bounded. The first self-tangency of this collar yields a bound for b.

Since by (6.1) the choice of α_σ cannot change the twisting number in the lemma by more than 4, we may assume that $\alpha_\sigma\cap E$ has an initial segment which is this shortest arc α. Similarly we may assume that α_τ is orthogonal to ∂B in the metric τ. The Lipschitz condition in the complement of R implies that the $\hat\tau$-length of α is bounded by Kb. Thus the number of essential intersections that α_τ can have with α_σ in the interval α is bounded by Kb/b_0, since between any two such there is a segment of α that projects orthogonally (in $\hat\tau$) to an entire boundary component of B. The remainder of α_σ is contained in a convex set whose boundary ξ is the lift of a boundary component of R. Since R is a convex subsurface both in σ and τ, ξ still

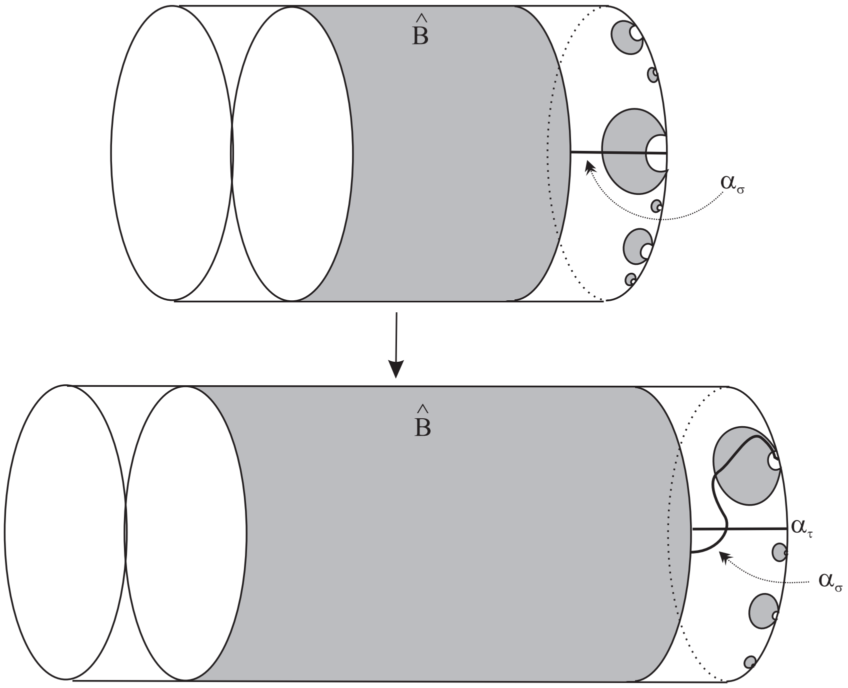

FIGURE 7. $\hat{Y}$, seen in the metric $\hat{\sigma}$ on top, and the metric $\hat{\tau}$ on the bottom. Some components of $\hat{R}$ are shown, shaded.

bounds a convex set in the metric $\hat{\tau}$. Thus α_τ can have at most one essential intersection with α_σ in this convex set. □

7. THE PROOF OF THE MAIN THEOREM

As discussed in the outline, after [47] it suffices to prove this direction of the theorem:

$$\sup_Y d_Y(\nu_+,\nu_-) < K \quad\implies\quad \inf_\gamma \ell_\rho(\gamma) > \epsilon$$

for ϵ depending only on K and S, with the infimum over γ that are not externally short.

For simplicity of exposition, let us first prove the theorem in the case when both ends of N_ρ are degenerate. At the end of the section we will indicate the changes in the argument necessary if one or both of the ends are geometrically finite.

Let K_1 be a constant to be determined shortly, and let ϵ_2 be such that a K_1-neighborhood of any ϵ_2-Margulis tube is still contained in an ϵ_1-Margulis tube (see §2.5).

Fix a closed curve γ in S, and assume $\ell_\rho(\gamma) < \epsilon_2$. In particular γ must be simple (§2.5), and represents a vertex of $\mathcal{C}_0(S)$. Our goal will be to bound the radius of the Margulis tube $\mathbf{T}_\gamma(\epsilon_0)$.

Initial pants

Our first step is to obtain two pants decompositions P_+ and P_- of moderate ρ-length, and pleated surfaces $f_\pm \in \mathbf{pleat}_\rho(P_\pm)$ which homologically encase $\mathbf{T}_\gamma(\epsilon_1)$.

Lemma 7.1. *Suppose that $\rho \in \mathcal{D}_{\rm np}(S)$ has two degnerate ends, so that $\nu_\pm \in \mathcal{EL}(S)$.*

Let $\mathcal{U}_+$ and $\mathcal{U}_-$ denote neighborhoods in $\mathcal{UML}(S)$ of ν_+ and ν_-, respectively. There exist pants decompositions P_+ and P_- lying in $\mathcal{U}_+$ and $\mathcal{U}_-$, respectively, with the following properties:

1. $\ell_\rho(P_\pm) \le L_1$
2. *Given $f_+ \in \mathbf{pleat}_\rho(P_+)$ and $f_- \in \mathbf{pleat}_\rho(P_-)$, $\mathbf{T}_\gamma(\epsilon_1)$ is homologically encased by f_+ and f_-.*

The last statement in the lemma means that $\mathbf{T}_\gamma(\epsilon_1)$ is covered with degree 1 by any 3-chain whose boundary relative the cusps is $f_+ - f_-$. In particular, any proper homotopy from f_+ to f_- must cover all of $\mathbf{T}_\gamma(\epsilon_1)$.

Proof. Since the e_+ end is degenerate, there is a sequence of pleated surfaces f_i^+ eventually contained in every neighborhood of the end. Choose $P_i^+ \in \mathbf{short}(\sigma_{f_i^+})$. Then the geodesic representatives of P_i^+ must also eventually exit e_+. In particular they must converge to ν_+ in $\mathcal{UML}$. The same discussion applies with $+$ replaced by $-$.

Since $(N_\rho, cusps)$ is homeomorphic to $(S \times \mathbb{R}, cusps \times \mathbb{R})$ and $\mathbf{T}_\gamma(\epsilon_1)$ is compact, any pleated surface in the proper homotopy class of ρ which is sufficiently far in a neighborhood of e_+ (or e_-) can be deformed to infinity, through cusp-preserving maps, without meeting $\mathbf{T}_\gamma(\epsilon_1)$.

Choosing i high enough, then, insures that f_i^+ and f_i^- homologically encase $\mathbf{T}_\gamma(\epsilon_1)$ and that $P_i^\pm \in \mathcal{U}_\pm$. Furthermore $P_i^\pm \in \mathbf{short}(\sigma_{f_i^+})$ implies that $f_i^\pm \in \mathbf{good}_\rho(P_i^\pm, L_1)$, so that by Lemma 4.1 there is a K-good homotopy (with K depending on L_1) between $f_i^\pm$ and any map in $\mathbf{pleat}(P_i^\pm)$. This homotopy has bounded tracks outside the Margulis tubes of $P_i^\pm$, if any. Thus we may choose i high enough that this homotopy also avoids $\mathbf{T}_\gamma(\epsilon_1)$, so that $P_\pm = P_i^\pm$ satisfies the conclusions of the lemma. □

Resolution sequence and block map

Join P_+ to P_- with a resolution sequence $P_- = P_0 \to \cdots \to P_n = P_+$, as in Theorem 5.1. Let $\{\beta_i\}_{i=0}^m$ be the vertex sequence of the associated geodesic in $\mathcal{C}(S)$.

Define a map $H : S \times [0, n] \to N_\rho$ as follows. For each $j = 0, \ldots, n$ choose $g_j \in \mathbf{pleat}_\rho(P_j)$, and for $j = 0, \ldots, n-1$ let $g_{j+\frac{1}{2}}$ be a map in $\mathbf{good}_\rho(P_j, C_1) \cap \mathbf{good}_\rho(P_{j+1}, C_1)$, promised to exist by Lemma 4.2 (Halfway surfaces).

Lemma 4.1 (Homotopy bound) gives us a constant K_1 so that g_j and $g_{j+\frac{1}{2}}$ admit a K_1-good homotopy, and so do $g_{j+\frac{1}{2}}$ and g_{j+1}. Let $H|_{S\times[j,j+\frac{1}{2}]}$ and $H|_{S\times[j+\frac{1}{2},j+1]}$ be these two homotopies. Up to suitable precomposition by homeomorphisms of S isotopic to the identity, we can arrange for the definitions to agree on each integer and half-integer level, so they may be concatenated to yield a single map H. The constant K_1 described here gives us our choice of ϵ_2, which we note depends only on the topological type of S.

Let $\sigma(u) = \sigma_{H_u}$ be the induced metric on $S \times \{u\}$, which we note is hyperbolic for u an integer or half-integer. We endow $S \times [0,n]$ with the metric that restricts to $\sigma(u)$ in the horizontal direction and to the induced metric from $\mathbb{R}$ in the vertical direction, and such that the two directions are orthogonal.

Restricting the block map

Assuming the neighborhoods $\mathcal{U}_\pm$ have been chosen sufficiently small, we can throw away all but a bounded number of the blocks and still have the encasing condition. To see this, begin with the following claim:

Claim 7.2. *There is a constant C_2 depending only on the topological type of S, and a subinterval $I_\gamma \subseteq [0,m]$ of diameter at most C_2, so that $H(S \times [j,j+1])$ can meet $\mathbf{T}_\gamma(\epsilon_2)$ only if j or $j+1$ are in J_{I_γ}.*

(The notation J_I is defined in §5.)

Proof. Let β_i be a vertex that is in P_j. If $H(S \times \{j\})$ meets $\mathbf{T}_\gamma(\epsilon_1)$ then (see §2.5) $\ell_{\sigma(j)}(\gamma) \leq \epsilon_0$, and in particular

$$\gamma \in \mathbf{short}(\sigma(j)) \subset \Pi_\rho(\beta_i).$$

It follows from Lemma 3.1 (Quasiconvexity) that

$$d_{\mathcal{C}}(\beta_i, \gamma) \leq C$$

where C depends only on the topological type of S. Thus, since $\{\beta_0, \dots, \beta_m\}$ are the vertices of a geodesic, the possible values of i lie in an interval of diameter at most $2C$, which we call I_γ. In other words, $j \in J_{I_\gamma}$.

It remains to notice that, by the choice of ϵ_2 and the K_1-good homotopy property, if any part of a block $H(S \times [j,j+1])$ meets $\mathbf{T}_\gamma(\epsilon_2)$ then one of the boundaries must meet $\mathbf{T}_\gamma(\epsilon_1)$, and hence j or $j+1$ are in J_{I_γ}. ☐

Let us therefore restrict our elementary move sequence to

$$P_{s-1} \to \cdots \to P_{t+1}$$

where $[s,t] = J_{I_\gamma}$. This subsequence must still encase $\mathbf{T}_\gamma(\epsilon_2)$, since we have thrown away only blocks that avoid $\mathbf{T}_\gamma(\epsilon_2)$. Let $M = t - s = |J_{I_\gamma}|$.

In order to deduce a bound on M from the diameter bound on I_γ, we consider part (3) of Theorem 5.1, which tells us that

$$M \leq K(2C) \sup_Y d_Y(P_+, P_-)^a, \tag{7.1}$$

where the supremum is over subsurfaces Y whose boundaries appear among the P_j in our subsequence. Such P_j must lie in a $C+1$ neighborhood of γ, in the $d_{\mathcal{C}}$ metric. In order to compare $d_Y(P_+, P_-)$ to $d_Y(\nu_+, \nu_-)$ for such Y, we will need the following lemma:

Lemma 7.3. *There exists a constant b depending only on S, so that given $\gamma \in \mathcal{C}_0(S)$, $R > 0$ and $\nu \in \mathcal{EL}(S)$, there is a neighborhood $\mathcal{U}$ of ν in $\mathcal{UML}(S)$ for which the following holds:*

If $Y \subset S$ is a subsurface with $d_{\mathcal{C}}(\gamma, \partial Y) \leq R$ and $\beta \in \mathcal{C}_0(S) \cap \mathcal{U}$, then

$$d_Y(\beta, \nu) \leq b.$$

Proof. By Klarreich's Theorem 2.1, $\mathcal{EL}(S)$ is naturally identified with the Gromov boundary of $\mathcal{C}(S)$. By the definition of the Gromov boundary, there is a neighborhood $\mathcal{U}$ of ν such that, if $\beta, \beta' \in \mathcal{C}_0(S)$ are in $\mathcal{U}$ then any geodesic g in $\mathcal{C}_1(S)$ connecting them must lie outside an $R+2$-ball of γ. In particular every vertex of g must be at least distance 2 from ∂Y, and hence must intersect ∂Y essentially. Theorem 3.1 of [37] states that in such a situation

$$\operatorname{diam}_Y(\pi_Y(g)) \leq A$$

for a constant A depending only on the topological type of S, and in particular $d_Y(\beta, \beta') \leq A$.

Now consider a sequence β_i converging to ν in $\mathcal{UML}(S)$, with $\beta_0 = \beta$. After restricting to a subsequence if necessary, the β_i converge in the Hausdorff topology to a lamination that includes the support of ν, which means that eventually $\operatorname{diam}_Y(\pi_Y(\beta_i) \cup \pi_Y(\nu)) \leq 2$. Thus $d_Y(\beta, \nu) \leq A + 2$. □

Thus if we choose the original neighborhoods $\mathcal{U}_+$ and $\mathcal{U}_-$ sufficiently small, this lemma gives us

$$d_Y(P_+, P_-) \leq d_Y(\nu_+, \nu_-) + 2b. \tag{7.2}$$

The hypothesis of the main theorem bounds the right side, so together with (7.1) we obtain our desired uniform bound on M.

Suppose that γ is not a component of any P_j. Then according to Lemma 4.1, each block $H|_{S\times[j,j+1]}$ has track lengths of at most $2K_1$ within $\mathbf{T}_\gamma(\epsilon_1)$. There are only $M+2$ blocks in our restricted sequence, and they cover all of $\mathbf{T}_\gamma(\epsilon_1)$. The beginning and the end of the sequence are outside of $\mathbf{T}_\gamma(\epsilon_1)$, so that any point in $\mathbf{T}_\gamma(\epsilon_1)$ is at most

$$2K_1(M+2)$$

from its boundary. This bounds the radius of $\mathbf{T}_\gamma(\epsilon_1)$ by $2K_1(M+2)$, and a corresponding lower bound for $\ell_\rho(\gamma)$ follows as in §2.5.

Bounding twists

Now suppose that γ does appear among the $\{P_j\}$. Then J_γ is some nonempty subinterval of J_{I_γ} by Theorem 5.1, and we let α and α' be the predecessor and successor curves to γ in the sequence. Both of them cross γ, and we have by part (4) of Theorem 5.1 that $d_\gamma(\alpha,\alpha')$ is uniformly approximated by $d_\gamma(P_+,P_-)$, which by (7.2) and the hypothesis of the main theorem, is uniformly bounded. Let D denote this bound.

Write $J_\gamma = [k,l]$. By the normalization used in Lemma 4.1, for all integer and half-integer $u \in [k,l]$ the annuli $\mathbf{collar}(\gamma,\sigma(u))$ coincide. Name this common annulus B. consider the solid torus

$$U = B \times [k - \tfrac{1}{2}, l + \tfrac{1}{2}].$$

The map H can take the complement of U at most $2K_1(M+2)$ into $\mathbf{T}_\gamma(\epsilon_1)$, by the same argument as above, and hence there is a uniform $\epsilon_3 > 0$ so that $H(U)$ must cover $\mathbf{T}_\gamma(\epsilon_3)$. Assume $\ell_\rho(\gamma) < \epsilon_3$, for otherwise we are done.

Consider the geometry of ∂U, in the metric we have placed on $S \times [0,n]$. The top annulus $B \times \{k - \frac{1}{2}\}$ is $\mathbf{collar}(\gamma, \sigma(k-\frac{1}{2}))$. Since $H_{k-\frac{1}{2}} \in \mathbf{good}(\gamma, C_1)$ by construction, we have an upper bound of $\epsilon_3 + C_1$ on the circumference of this collar. We also have a lower bound of ϵ_3 since the image of $H_{k-\frac{1}{2}}$ avoids the interior of $\mathbf{T}_\gamma(\epsilon_3)$. Thus, this annulus has uniformly bounded geometry. The same holds for $B \times \{l + \frac{1}{2}\}$ in the metric $\sigma(l+\frac{1}{2})$

The vertical annuli $\partial B \times [k-\frac{1}{2}, l+\frac{1}{2}]$ have height bounded by $M+1$. Thus ∂U is a concatenation of four annuli of bounded geometry, and hence is itself up to uniformly bounded distortion a Euclidean square torus, and H maps it with a uniform Lipschitz constant into N. Let us now try to control a meridian curve for U.

Realize α as a geodesic in $\sigma(k-\frac{1}{2})$ and let a be an arc of $\alpha \cap B$ (there may be two). Similarly assume α' is a geodesic in $\sigma(l+\frac{1}{2})$ and let a' be an arc of $\alpha' \cap B$. Lemma 4.3 gives an upper bound for the length of a in $\sigma(k-\frac{1}{2})$, and for the length of a' in $\sigma(l+\frac{1}{2})$. The curve

$$m = \partial(a \times [k-\tfrac{1}{2}, l+\tfrac{1}{2}]),$$

is a meridian of U, and we claim that its length is uniformly bounded.

The arc a may a priori be long in $\sigma(l+\frac{1}{2})$, but its length is estimated by the number of times it twists around the annulus, which in turn is estimated by $d_{\mathcal{A}(B)}(a,a')$ since a' has bounded length in this metric.

We can bound this twisting using the results in Section 6. Note first that, for each $j = k-\frac{1}{2}, k, \ldots, l-\frac{1}{2}, l$, the pair of metrics $\sigma(j), \sigma(j+\frac{1}{2})$ satisfy the hypotheses of Lemma 6.1 (Shear bound). That is, they are K_1-bilipschitz outside a union of collars including B. Thus, Lemma 6.1 bounds the shearing outside B at each step, and the number of steps is bounded by $2M+2$. We conclude that there is a bound on the shearing outside B of the two metrics $\sigma(k-\frac{1}{2})$ and $\sigma(l+\frac{1}{2})$. It follows as in (6.2) that we have a bound of the form

$$|d_{\mathcal{A}(B)}(a, a') - d_\gamma(\alpha, \alpha')| \leq M' \tag{7.3}$$

With this estimate and the bound $d_\gamma(\alpha, \alpha') \leq D$, we find that a and a' intersect a bounded number of times, so that the length of a is uniformly bounded in $S \times \{l + \frac{1}{2}\}$. It follows that $H(m)$ is uniformly bounded. It therefore spans a disk of bounded diameter, and now by a coning argument we can homotope H on all of U to a new map of bounded diameter. This bounds the radius of $\mathbf{T}_\gamma(\epsilon_3)$ from above, and again we are done.

The case of geometrically finite ends

The main change in the argument comes at the beginning, in our choice of P_+ and P_-.

Suppose that e_+ is geometrically finite. It suffices to consider γ which is not externally short, i.e. $\ell_{\nu_+}(\gamma) > \epsilon_0$, because the second statement in the main theorem considers only such curves, and the first statement is insensitive to the removal of a finite number of curves from consideration. The choice of ϵ_1 and Sullivan's theorem (see §2.5) imply that we also have $\ell_{\nu'_+}(\gamma) > \epsilon_1$.

We may therefore choose a pants decomposition $P_+ \in \mathbf{short}(\nu'_+)$ so that γ is not one of its components (our choice of L_1 was made so that this would be possible, see §2.6). Let $f'_+ : S \to \partial_+(C(N_\rho))$ be the pleated surface in the homotopy class of ρ that parametrizes the e_+-boundary of the convex hull. In particular $f'_+ \in \mathbf{good}_\rho(P_+, L_1)$. Thus, choosing $f_+ \in \mathbf{pleat}_\rho(P_+)$, Lemma 4.1 gives us a K_2-good homotopy G_+ between f'_+ and f_+, with K_2 depending on L_1. Since γ is not a component of P_+, G_+ can only penetrate a distance K_2 into $\mathbf{T}_\gamma(\epsilon_1)$, and hence does not meet $\mathbf{T}_\gamma(\epsilon_4)$, with ϵ_4 depending only on K_2.

We do the same thing with f_-, P_- if e_- is geometrically finite, or we repeat the original discussion if e_- is degenerate. Thus we have f_+, f_- encasing $\mathbf{T}_\gamma(\epsilon_4)$.

We can therefore continue as before, constructing a resolution sequence from P_- to P_+ and then restricting it. The argument goes through with slightly altered constants (since we have replaced ϵ_1 with ϵ_4), and the one step that needs attention is the comparison (7.2) between $d_Y(P_+, P_-)$ and $d_Y(\nu_+, \nu_-)$, for appropriate Y. In other words we must bound

$$d_Y(P_+, \nu_+)$$

in the geometrically finite case (and similarly for e_-). This quantity by definition is just $\min_{Q_+} d_Y(P_+, Q_+)$ where Q_+ varies over pants decompositions in $\mathbf{short}(\nu_+)$, provided P_+ meets Y nontrivially and Q_+ can be found which does the same.

Since ν_+ and ν'_+ admit a K_0-bilipschitz map by Sullivan's theorem (§2.1), the ν'_+-length of Q_+ is bounded by K_0L_1, which gives a bound on its intersection number with P_+. This bounds $d_Y(P_+, Q_+)$ provided both pants decompositions intersect Y.

If Y is not an annulus (and recall that we never consider three-holed spheres), then it automatically intersects any pants decomposition. Thus the only problematic case is if Y is an annulus whose core is a component of P_+ or Q_+. However, we recall that the only annulus that actually comes into the argument is the annulus with core γ. Since γ is not externally short, we have already chosen P_+ so that γ is not a component of it, and we can do the same for Q_+.

The rest of the proof goes through in the same way. □

8. Appendix: Hyperbolic geometry

In this appendix we write out statements, and sketch some proofs, of a few facts and constructions in hyperbolic geometry. These are "well-known" in the sense that those working in this field are familiar at least with some variation of them, or would find it straightforward to derive them. Still it seems advisable to include some discussion.

Throughout, a *hyperbolic surface* always means a finite area hyperbolic surface which could have closed geodesic boundary components and/or cusps. By abuse of notation we usually think of a cusp as a boundary component of length 0.

Collars

Let w_0 be the function

$$w_0(t) = \sinh^{-1}\left(\frac{1}{\sinh(t/2)}\right).$$

For a simple closed geodesic γ of length ℓ in a hyperbolic surface X, we define

$$\mathbf{collar}_0(\gamma, X) = \{p \in X : \mathrm{dist}(p, \gamma) \le w_0(\ell)\}.$$

When the ambient surface X is understood we omit it from the notation. This set is always an embedded annulus, and in fact if $\gamma_1 \ldots, \gamma_k$ are disjoint and homotopically distinct then $\mathbf{collar}_0(\gamma_i)$ are pairwise disjoint. (See e.g. Buser [12, Chapter 4].)

We will need a slightly smaller collar, so as to guarantee a definite amount of space in its complement. Define

$$w = \max(w_0/2, w_0 - 1)$$

and let

$$\mathbf{collar}(\gamma, X) = \{p \in X : \mathrm{dist}(p, \gamma) \le w(\ell)\}.$$

Let γ'_0 be a boundary component of $\mathbf{collar}_0(\gamma)$ (assume $\gamma'_0 \ne \gamma$ if γ itself is in ∂X), and similarly let $\gamma' \ne \gamma$ be a boundary component of $\mathbf{collar}(\gamma)$.

The length of γ_0' (resp. γ') is given by $\ell\cosh(w_0(\ell))$ (resp. $\ell\cosh(w(\ell))$). A bit of arithmetic shows that

$$l(\gamma') \le l(\gamma_0') \le \ell(\gamma) + 2 \tag{8.1}$$

We can define $\mathbf{collar}_0(\gamma)$ and $\mathbf{collar}(\gamma)$ also when γ represents a cusp of X, describing them either explicitly or as a limit as $\ell \to 0$. The boundary of $\mathbf{collar}_0(\gamma)$ for a cusp is horocyclic, and its length is 2 (the limiting value of $\ell\cosh(w_0(\ell))$ as $\ell \to 0$). The slightly smaller $\mathbf{collar}(\gamma)$ has horocylic boundary a distance 1 inside the boundary of $\mathbf{collar}_0(\gamma)$, so its length is $2/e$.

We remark in fact that the boundary length $\ell\cosh w_0(\ell)$ of $\mathbf{collar}_0(\gamma)$ is increasing with ℓ, and hence always at least 2. There is a similar lower bound b_0 for the boundary of $\mathbf{collar}(\gamma)$, which we will not compute explicitly.

If $P = \{\gamma_1, \dots, \gamma_k\}$ is a curve system we let $\mathbf{collar}(P) = \cup_i \mathbf{collar}(\gamma_i)$.

Hyperbolic pairs of pants

If Y is a hyperbolic surface as above, with genus 0 and three boundary components, we call it a hyperbolic pair of pants. The three boundary lengths determine the metric on Y completely (up to isotopy). Let $Y_0 = Y \setminus int(\mathbf{collar}(\partial Y))$.

A *tripod* is a copy of the 1-complex Δ obtained from three disjoint copies of $[0,1]$ (called "legs") by identifying the three copies of $\{0\}$. The three copies of $\{1\}$ are called the boundary of Δ. An *essential tripod* in Y_0 is an embedding of Δ (also called Δ) taking $\partial\Delta$ to ∂Y_0, such that each subarc of Δ obtained by deleting one copy of $(0,1]$ is not homotopic rel endpoints into ∂Y_0 (Figure 8).

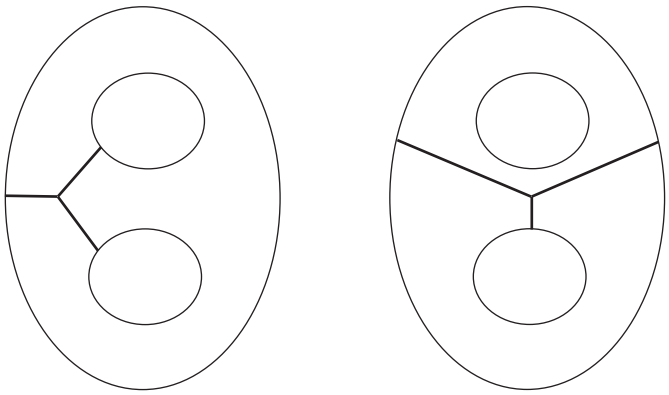

FIGURE 8. The two types of essential tripods

Lemma 8.1. *There is a constant $\delta > 0$ such that for any hyperbolic pair of pants Y and each boundary component γ of Y_0, there is an essential tripod $\Delta \subset Y_0$ whose three legs have length at most δ, and which meets γ.*

On the other hand, there exists $\delta' > 0$ such no essential tripod in any Y_0 has all three legs of length less than δ'.

The next lemma describes how an estimate on the boundary lengths of Y yields an estimate on its isometry type.

Lemma 8.2. *Given $C \geq 0$ there exists $K \geq 1$ such that the following holds.*
Let Y and Y' be two hyperbolic pairs of pants with boundaries $\{\gamma_i\}_{i=1}^3$ and $\{\gamma_i'\}_{i=1}^3$, respectively. Suppose that

$$|\ell(\gamma_i) - \ell(\gamma_i')| \leq C.$$

Then there is a homeomorphism $h : int(Y) \to int(Y')$ taking Y_0 to Y_0' and $\mathbf{collar}(\gamma_i)$ to $\mathbf{collar}(\gamma_i')$, which is K-bilipschitz on Y_0. Furthermore, for any arc a on ∂Y_0 we have $|l(\alpha) - l(h(\alpha))| \leq K$.

Let us prove Lemma 8.2 first. Any hyperbolic pair of pants Y admits a canonical decomposition into two congruent right-angled hexagons, by cutting along the shortest arcs connecting each pair of boundary components (see Buser [12]. In the limiting case of cusps we obtain degenerate hexagons with ideal vertices). Thus it will suffice to find appropriate bilipschitz maps between such hexagons.

Comparison of hexagons

Let Ξ be a hyperbolic right-angled hexagon, with three alternating sides A_1, A_2, A_3 of lengths a_1, a_2, a_3. (we will adopt the convention here of denoting edges by capital letters and their lengths by the corresponding lower case letters). Each side A_i admits an embedded "collar rectangle", which we call $\mathbf{collar}(A_i, \Xi)$, and is just the $w(2a_i)$-neighborhood of A_i in Ξ. If Y is the hyperbolic pair of pants obtained from Ξ by doubling along its other three boundaries, then $\mathbf{collar}(A_i, \Xi)$ is clearly just $\mathbf{collar}(\gamma_i, Y) \cap \Xi$, where γ_i is the double of A_i.

Lemma 8.2 clearly follows from the following property of hexagons:

Lemma 8.3. *Given $C > 0$ there exists $K \geq 1$ so that the following holds:*
Let Ξ, Ξ' be two hyperbolic right-angled hexagons with alternating sides $\{A_i\}$ and $\{A_i'\}$, respectively. Suppose that

$$|a_i - a_i'| \leq C, \quad i = 1, 2, 3.$$

Then there is a label-preserving homeomorphism $f : \Xi \to \Xi'$ which takes $\mathbf{collar}(A_i, \Xi)$ to $\mathbf{collar}(A_i', \Xi')$ and is K-bilipschitz on the complement of the collars. Furthermore for any arc α in $\cup A_i$, we have $|l(\alpha) - l(f(\alpha))| \leq K$.

Proof. We begin by describing a decomposition of Ξ into controlled pieces in a way that is determined by the $\{a_i\}$. In the following, (i, j, k) always denotes a permutation of $(1, 2, 3)$.

Case 1: Suppose that the three "triangle inequalities"

$$a_i \leq a_j + a_k \tag{8.2}$$

all hold. Then there is a unique triple $r_{12}, r_{23}, r_{13} \geq 0$ satisfying

$$r_{ij} + r_{ik} = a_i. \tag{8.3}$$

(where we use the convention $r_{ij} = r_{ji}$). In fact we simply set $r_{ij} = \frac{1}{2}(a_i + a_j - a_k)$.

Now define three "bands" B_{ij} in Ξ as follows: B_{ij} is the (closed) r_{ij}-neighborhood of the edge C_{ij} of Ξ which is the common perpendicular of A_i and A_j. See Figure 9.

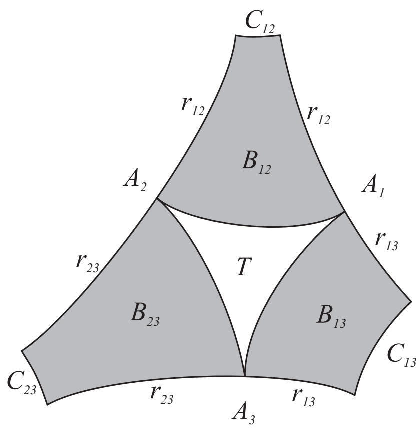

FIGURE 9. The bands decomposition of Ξ in case 1.

The two segments $B_{ij} \cap A_i$ and $B_{ik} \cap A_i$ cover all of A_i and meet in their common boundary point. The interiors of B_{ij} are disjoint as a consequence of the triangle inequality and the fact that $\operatorname{dist}(C_{ij}, C_{ik}) = a_i$. The closure T of the complement $\Xi \setminus \bigcup_{ij} B_{ij}$ is a "triangle" with the following properties:

- Each edge $E_{ij} = T \cap B_{ij}$ has curvature $\kappa_{ij} \in [0, 1]$, which is in fact given by $\tanh r_{ij}$. The curvature vector points outward.
- Adjacent edges meet at angle 0.

From this and the Gauss-Bonnet theorem we have

$$\sum_{ij=12,13,23} e_{ij}\kappa_{ij} = \pi - \operatorname{area}(T)$$

which implies the upper bound

$$e_{ij} \leq \frac{\pi}{\tanh r_{ij}}. \tag{8.4}$$

There is also a lower bound

$$e_{ij} \geq 2 \log \frac{2}{\sqrt{3}} \tag{8.5}$$

obtained as follows: Extend each E_{ij} to a line of constant curvature in $\mathbb{H}^2$. Each bounds a region R_{ij} that contains B_{ij}, so that the three regions have disjoint interiors. Since $\kappa_{ij} \leq 1$, through any point $x \in \partial R_{ij}$ passes the boundary of a horoball contained in R_{ij}. One can check that for any three horoballs in $\mathbb{H}^2$ with disjoint interiors, and any point p, the distance from p to at least one of the horoballs is at least $\log \frac{2}{\sqrt{3}}$ (the extreme case is where all three are tangent). The midpoint of E_{ij} meets a horoball in R_{ij} and is a distance at most $e_{ij}/2$ from horoballs in R_{ij} and R_{jk}. Thus we have (8.5).

The geometry of the band B_{ij} is easy to describe: it can parametrized by the rectangle $[0, r_{ij}] \times [0, c_{ij}]$ with the metric

$$dx^2 + \cosh^2(x)dy^2, \tag{8.6}$$

Where C_{ij} is identified with $\{0\} \times [0, c_{ij}]$.

Case 2: Suppose that one of the opposite triangle inequalities holds, e.g.

$$a_1 \geq a_2 + a_3. \tag{8.7}$$

We then let $r_{12} = a_2$, $r_{13} = a_3$, and

$$r_{11} = a_1 - a_2 - a_3,$$

So that $a_1 = r_{12} + r_{11} + r_{13}$. We then have bands B_{12} and B_{13} as before, and B_{11} is defined as follows (see Figure 10): Let H_1 be the common perpendicular of A_1 and its opposite edge C_{23}. In A_1 let J_{11} be the closure of the complement of $B_{12} \cap A_1$ and $B_{13} \cap A_1$. This has length r_{11}. Join each $x \in J_{11}$ to a point $y \in C_{23}$ with a curve equidistant from H_1. We obtain a foliated rectangle, B_{11} (if $r_{11} = 0$ then B_{11} is a single segment). Similarly to the other bands, we can describe the metric in B_{11} by the formula (8.6), but on a rectangle of the form $[u_2, u_3] \times [0, h_1]$, where $u_3 - u_2 = r_{11}$.

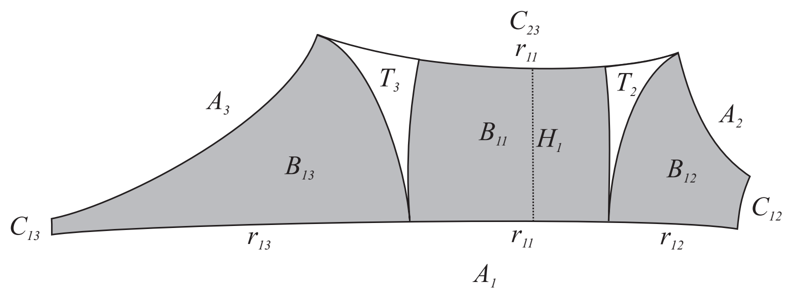

FIGURE 10. The bands decomposition of Ξ in case 2.

The closure of $\Xi \setminus \bigcup_i B_{1i}$ is now two triangles T_2, T_3, with angles 0, 0 and $\pi/2$. Let $E_{1m} = T_m \cap B_{1m}$ and let $E_{11,m} = T_m \cap B_{11}$, for $m = 2, 3$. Note that the curvature of $E_{11,m}$ is $\kappa_{11,m} = \tanh |u_m|$. Upper and lower bounds on e_{1m} and $e_{11,m}$ can be derived from (8.4) and (8.5), by observing that the union of T_m with its reflection across C_{23} is a triangle of the same type as

T in case 1. For e_{1m} we obtain (8.4) and (8.5) exactly, whereas for $e_{11,m}$ we get

$$e_{11,m} \le \frac{\pi}{2\tanh|u_m|} \tag{8.8}$$

and

$$e_{11,m} \ge \log\frac{2}{\sqrt{3}}. \tag{8.9}$$

Let $|u_m|$ be the larger of $|u_2|, |u_3|$. Then $|u_m| \ge r_{11}/2$, and since $e_{11,m}$ must then be the larger of $|e_{11,2}|, |e_{11,3}|$, we have a bound for $j = 2, 3$:

$$e_{11,j} \le \frac{\pi}{2\tanh(r_{11}/2)} \tag{8.10}$$

The comparison: Now consider two hexagons Ξ, Ξ' satisfying the bound $|a_i - a_i'| \le C$.

Suppose first that both Ξ and Ξ' are in case 1. Note immediately that we also have $|r_{ij} - r_{ij}'| \le \frac{3C}{2}$.

If $r_{ij} > 3C$ for each i, j, then $r_{ij}' > \frac{3C}{2}$. In this case we have upper bounds of $\pi/\tanh(3C/2)$ on both e_{ij} and e_{ij}'. The bilipschitz map $\Xi \to \Xi'$ can be constructed separately on each piece of the decomposition.

Map each band B_{ij} to B_{ij}' using an affine stretch on the parameter rectangles. The metric described in (8.6) gives a uniform bilipschitz bound on this map. Note that the assumption that both r_{ij} and r_{ij}' are bounded away from 0 bounds the bilipschitz constant in the x direction, and the fact that their difference is bounded gives a bound in the other direction (since the cosh factor gives an exponential scaling).

The triangles T and T' are also in uniform bilipschitz correspondence, since in fact they vary in a compact family of possible figures (due to the length and curvature bounds).

If at least one $r_{ij} < C$ we must take a bit more care. Let us consider a limiting case where at least one $r_{ij} = 0$ and the rest are no smaller than C. These cases are illustrated in Figure 11.

If $r_{12} = r_{23} = r_{13} = 0$ then $a_1 = a_2 = a_3 = 0$ and Ξ is an ideal triangle. Recall that we are interested only in bounds on the complements of the collars, and this is in this case a right-angled "hexagon" with three of the sides equal to horocyclic arcs of definite length. If $r_{12} = r_{13} = 0$ but $r_{23} > C$, then T has two geodesic edges meeting in an ideal vertex, and a third leg E_{23} satsifying the bounds (8.4,8.5). To this leg is attached the band B_{23}.

If $r_{12} = 0$ and $r_{13}, r_{23} > C$ then we obtain a triangle T with one geodesic leg and two curved legs to which are attached the bands. Note that the length bounds on the curved legs imply a length bound on the straight leg.

If we now take a general Ξ with some $r_{ij} < C$, consider the family of shapes obtained by taking those r_{ij} to 0, and fixing the rest. These must converge to one of the three cases above, and the bands corresponding to

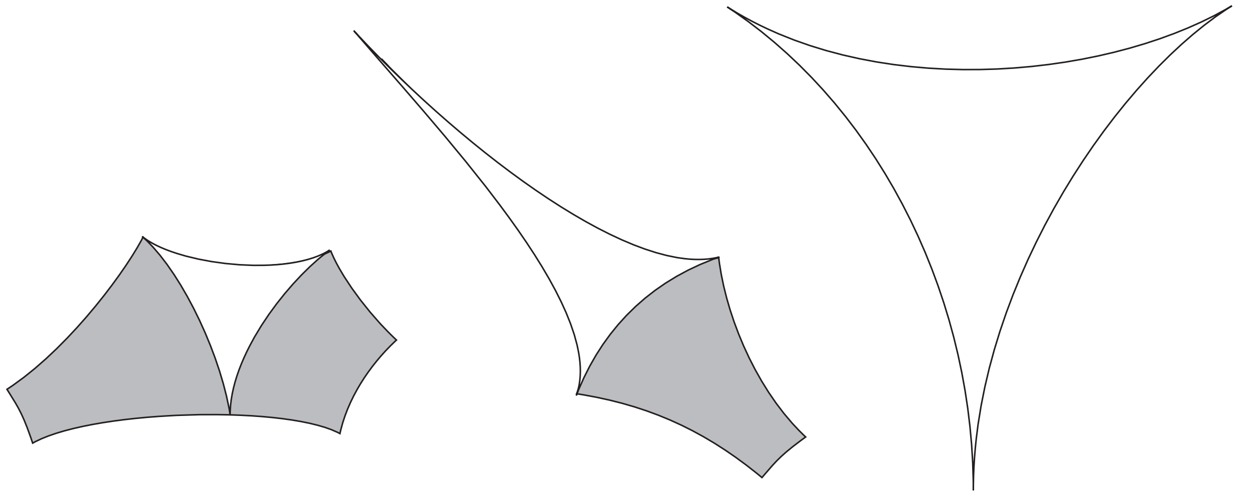

FIGURE 11. The boundary cases where some $r_{ij} = 0$. From left to right, one, two and all three are 0.

$r_{jk} \geq C$ remain fixed. Thus we can find some uniformly bilipschitz mapping taking Ξ to the limiting case.

Two hexagons Ξ and Ξ' both near one of the boundary cases are therefore near to each other, via the composition of the two maps.

Now suppose our two hexagons are in case 2. Again we obtain bounds $|r_{1j} - r'_{1j}| \leq C$ for $j = 2, 3$ and $|r_{11} - r'_{11}| \leq 3C$. Assume first $r_{1j} > 6C$ for $j = 1, 2, 3$. The bands B_{12}, B_{13} again have uniformly bilipschitz maps to B'_{12} and B'_{13}, respectively, coming from the affine stretches of the parameter rectangles. The same is true for $B_{11} \to B'_{11}$, using the parametrization $[u_2, u_3] \times [0, h_1]$ with metric (8.6), together with the bounds (8.8–8.10).

The cases where some $r_{1j} \leq 6C$ are handled as before, by comparing to the boundary cases where some $r_{1j} = 0$. We leave the details to the reader.

It is also possible for Ξ to be in case 1 and Ξ' to be in case 2. However, this can only happen if both of them are near their respective boundary cases, and one can check that these boundary cases are in fact in the overlap of case 1 and case 2, and hence can be compared with each other.

What we have shown is that Ξ and Ξ' admit a homeomorphism which is bilipschitz outside the collars of those A_i, A'_i that are sufficiently short. For the long A_i we observe that, as in (8.1), the length of the boundary of $\mathbf{collar}(A_i)$ in $int(\Xi)$ is at most $a_i + 1$, and the collar width w has been chosen so that it is at most half the width of the largest embedded collar. It follows from this that the map can be adjusted to take all collars of Ξ to collars of Ξ', with a bounded change in the bilipschitz constant. The additive distortion in length of all subarcs of A_i follows from the construction. $\square$

The proof of Lemma 8.1 is fairly simple from the geometric picture of hexagons we have just produced. Let Y be decomposed as the double of a hexagon Ξ. If Ξ is in Case 1, the curved triangle T, minus the collars of the A_i, is a bounded-diameter hexagon, and an essential tripod of bounded size can be embedded in it. This tripod meets all three boundary collars.

In Case 2, we consider the double of T_2 (or T_3) across C_{23}. This is a curved triangle of the same type as T, and after cutting away the intersections with the collars we can again obtain a bounded essential tripod. Note that this tripod has two legs terminating on the same boundary component of Y. It only meets two boundary collars, but since we can choose both T_2 and T_3, for each collar there is a tripod meeting it.

To obtain the lower bound on any essential tripod in Y, we argue as in the lower bound on e_{ij} in the previous proof. Suppose the legs of the tripod have length less than l. Lifting the tripod to $\mathbb{H}^2$ we see that the intersection point of the legs has distance $\leq l$ from three distinct lifts of the collars, each of which is bounded by by a curve of constant curvature $\kappa \in [0,1]$. The smallest possible l is again obtained if $\kappa = 1$ for each curve, namely $l \geq 2\log\frac{2}{\sqrt{3}}$. □

A curve shortening lemma

Let γ be a circle of length $l(\gamma)$, and let $h : \gamma \to N$ be a homotopically nontrivial, arclength-preserving map into a hyperbolic 3-manifold N. Lift h to a map $\widetilde{h} : \widetilde{\gamma} \to \mathbb{H}^3$. After appropriate choice of basepoints, $h_*(\pi_1(\gamma))$ is a nontrivial cyclic subgroup of $\pi_1(N)$ preserving $\widetilde{h}(\widetilde{\gamma})$. Let $\ell_N(h) \geq 0$ denote the translation length of its generator. If it is loxodromic let Γ be its axis and let $h(\gamma)^*$ denote the quotient of Γ in N, and if it is parabolic or has translation length less than ϵ_0, let $\widetilde{\mathbf{T}}_h(\epsilon)$ (for $0 < \epsilon \leq \epsilon_0$) be the invariant lift of its corresponding ϵ-Margulis tube in N.

Lemma 8.4. *Suppose $h : \gamma \to N$ is as above and satisfies the condition*

$$l(\gamma) \leq \ell_N(h) + C.$$

Let $\alpha \subset \widetilde{\gamma}$ be an arc whose length is $l(\alpha) = nl(\gamma) + r$, with $n \in \mathbb{Z}_{\geq 0}$ and $r \in [0, l(\gamma))$. Let $\epsilon \in (0, \epsilon_0]$ be given. Then $h(\alpha)$ can be deformed, rel endpoints, to a path $p_0 * \alpha' * p_1$ with the following properties:

1. *p_i are geodesics orthogonal to Γ or $\widetilde{\mathbf{T}}_h(\epsilon)$, and $l(p_i) \leq a_1$, for $i = 0, 1$.*
2. *α' lies on the geodesic Γ if $\ell_N(h) \geq \epsilon$, and in $\widetilde{\mathbf{T}}_h(\epsilon)$ otherwise.*
3. *$l_N(\alpha') = nl' + r'$, where*

$$l' = \max(\ell_N(h), \epsilon) \tag{8.11}$$

and

$$|r - r'| \leq a_2. \tag{8.12}$$

The constants a_i depend only on ϵ and C.

Proof. Fix a generator $\bar{\gamma}$ of $\pi_1(\gamma)$, and let $A = h_*(\bar{\gamma})$. Let $\ell = \ell_N(h)$ be the translation distance of A. Suppose first that A is loxodromic. If $x \in \mathbb{H}^3$ with $d(x, \Gamma) = r$ and $s = d(x, Ax)$ then one can obtain (see e.g. Buser [12, 2.3.1]

$$\sinh \frac{s}{2} \geq \sinh \frac{\ell}{2} \cosh r \tag{8.13}$$

Suppose that $s \leq \ell + C$ and $\ell \geq \epsilon$. Then, since the function $\sinh(t+C)/\sinh t$ is decreasing, we have $\cosh r \leq \sinh(\epsilon + C)/\sinh(\epsilon)$, yielding an upper bound on r.

Let $\xi \in \Gamma$ be the closest point to x (so $[x, \xi] \perp \Gamma$), and let x_t be the point on $[x, \xi]$ with $d(x_t, \xi) = t$. Suppose that $\ell \leq \epsilon$ and hence $s \leq \epsilon + C$. We claim that

$$d(x_t, A(x_t)) \leq a \frac{\cosh t}{\cosh r} \tag{8.14}$$

where a depends only on $\epsilon + C$. To see this, note that the distance from x to Ax along the r-equidistant surface to Γ is at most a constant c', since we can project the geodesic $[x, Ax]$ to this surface with bounded distortion. On the other hand the distance from x_t to Ax_t along the t-equidistant surface is $\sqrt{\ell^2 \cosh^2 t + \theta^2 \sinh^2 t}$, where $\ell + i\theta$ is the complex translation length of A. A little algebra yields (8.14).

It follows that there exists a_1 depending only on ϵ such that $x_t \in \widetilde{T}(\epsilon)$ for all $t \in [0, r - a_1]$.

When A is parabolic the same discussion works using coordinates adapted to its parabolic fixed point, replacing equidistant surfaces to Γ with horospheres.

Now for any arc $\alpha \subset \widetilde{\gamma}$, with endpoints x, y, let p_0 and p_1 be the orthogonal geodesics from $\widetilde{h}(x)$ and $\widetilde{h}(y)$, respectively, to Γ when $\ell \geq \epsilon$, and to $\widetilde{T}(\epsilon)$ when $\ell < \epsilon$. The above discussion bounds the lengths of p_0 and p_1 uniformly. Let ξ, η be the other endpoints of p_0 and p_1.

Consider first the case that $\ell \geq \epsilon$. Let α' be the geodesic segment $[\xi, \eta]$ on Γ. Write $l(\alpha) = nl(\gamma) + r$ as in the statement of the lemma. If $r = 0$ then $\eta = A^{\pm n}\xi$, and it is immediate that $l(\alpha') = n\ell$, which is what we wished to prove.

Suppose $n = 0$, so that $l(\alpha) = r < l(\gamma)$, and let us prove (8.12). Identify $\widetilde{\gamma}$ with $\mathbb{R}$ so that the generator $\bar{\gamma}$ of $\pi_1(\gamma)$ acts by $t \mapsto t + l(\gamma)$. Let $x' = x + l(\gamma)$. We may assume $y \in [x, x']$. Let $\xi' \in \Gamma$ be the closest point to $\widetilde{h}(x')$. Then $\widetilde{h}(x), \widetilde{h}(x')$ and $\widetilde{h}(y)$ are distance at most a_1 from ξ, ξ' and η, respectively.

We have

$$d(\xi, \eta) \leq 2a_1 + |x - y| \tag{8.15}$$

and

$$d(\xi', \eta) \leq 2a_1 + |x' - y| \tag{8.16}$$

by the triangle inequality and the fact that $\widetilde{h}$ is arclength-preserving. We have by hypothesis

$$|x - x'| - d(\xi, \xi') \leq C. \tag{8.17}$$

We have $|x' - y| = |x' - x| - |y - x|$ since $y \in [x, x']$, and we have $d(\xi', \eta) \geq d(\xi, \xi') - d(\xi, \eta)$ by the triangle inequality, so (8.16) yields:

$$d(\xi, \xi') - d(\xi, \eta) \leq 2a_1 + |x' - x| - |y - x|. \tag{8.18}$$

Now rearranging and applying (8.17) we have

$$d(\xi, \eta) \geq |y - x| - 2a_1 - C. \tag{8.19}$$

Now $d(\xi, \eta) = r'$ and $|y - x| = r$, so (8.15) and (8.19) together bound $|r - r'|$.

Together with the upper bound (8.16) on $d(\xi', \eta)$, we find that η is in fact constrained to within bounded distance of the point in $[\xi, \xi']$ at distance r from ξ (as opposed to the point on the other side of ξ at the same distance). This together with the case of $n > 0, r = 0$ yields the general case, by concatentation.

Now consider the case that $\ell < \epsilon$. Writing $l(\alpha) = nl(\gamma) + r$, let α' be the concatenated chain of geodesic segments connecting the sequence of $\xi = A^0\xi, \ldots, A^n\xi, \eta$. The length estimate on α' is immediate.

□

Truncated curves

Lemma 8.5. *Let γ be a simple closed geodesic or a simple properly embedded geodesic line on a hyperbolic surface S. Suppose that γ avoids the ϵ_1-thin part of S, except possibly for its ends leaving the ϵ_1-cusps. Given ϵ, there exists a homotopically non-trivial, non-peripheral curve γ_ϵ, homotopic to a simple curve, with the following properties:*

1. *γ_ϵ has length at most L_2.*
2. *γ_ϵ is composed of at most 2 segments of γ and at most 2 bridge arcs for γ of length at most ϵ.*

The constant L_2 depends only on ϵ, ϵ_1 and the topology of S.

Proof. (Sketch) We may assume $\epsilon << \epsilon_1$. Let $L = 3\pi|\chi(S)|/\epsilon$, and let $L_2 = 2(L + \epsilon)$.

First, in case γ is already a simple closed geodesic with length less than L_2, take $\gamma_\epsilon = \gamma$.

If γ is a properly embedded geodesic whose length outside the ϵ-cusps is at most L, γ_ϵ can be constructed as a boundary component of a slight thickening of γ union the cusps (see Figure 12).

Finally if γ has length larger than L outside the ϵ-cusps, let γ_0 denote any arc of γ outside the cusps of length L, and consider the immersion of $\gamma_0 \times [-\epsilon, \epsilon]$ taking each $\{p\} \times [-\epsilon, \epsilon]$ to the geodesic segment of length 2ϵ orthogonal to γ at its midpoint p. The choice of L_2 and the bound of

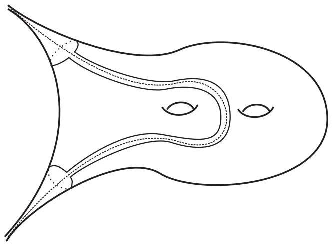

FIGURE 12. Constructing γ_ϵ when γ is a bounded arc modulo cusps.

$2\pi|\chi(S)|$ on the area of S imply that there must be a point z covered with multiplicity 3 by this map. Thus γ passes within ϵ of z three times with (unoriented) directions varying by $O(\epsilon)$ in the projectivized tangent bundle. At least two of those times the orientations match, and we may build γ_ϵ as shown in Figure 13. The resulting curve is homotopically nontrivial and non-peripheral because it can be smoothed out to have curvature near 0 and length bounded away from 0. □

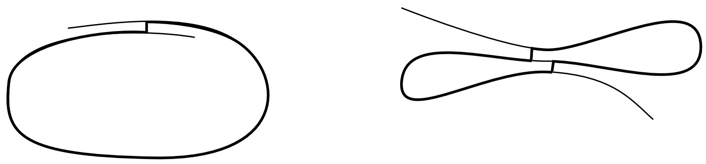

FIGURE 13. Constructing γ_ϵ when γ is long

Junctures of convex sets

Lemma 8.6. *Fix $r_0 > 0$. For any $b > 0$ there is a $d > 0$ such that the following holds: if A and B are convex sets in $\mathbb{H}^3$, and suppose that* $\operatorname{dist}(A,B) \geq r_0$. *Then*

$$\operatorname{diam}(\mathcal{N}_b(A) \cap \mathcal{N}_b(B)) \leq d.$$

Here $\mathcal{N}_b$ denotes b-neighborhood in $\mathbb{H}^3$.

Proof. Convexity and the lower bound on $\operatorname{dist}(A,B)$ imply that A and B are contained in disjoint half-spaces H_A and A_B that are at least r_0 apart. The intersection $\mathcal{N}_b(H_A) \cap \mathcal{N}_b(H_B)$ is compact so it has some diameter bound d. A sharp value is not hard to compute but we will not need it. □

References

[1] W. Abikoff, *The real-analytic theory of Teichmüller space*, Springer-Verlag, 1980, Lecture Notes in Mathematics no. 820.

[2] J. M. Alonso, T. Brady, D. Cooper, V. Ferlini, M. Lustig, M. Mihalik, M. Shapiro, and H. Short, *Notes on word hyperbolic groups*, Group Theory from a Geometrical Viewpoint, ICTP Trieste 1990 (E. Ghys, A. Haefliger, and A. Verjovsky, eds.), World Scientific, 1991, pp. 3–63.

[3] L. Bers, *Spaces of degenerating Riemann surfaces*, Discontinuous groups and Riemann surfaces, Ann. of Math. Stud. 79, Princeton Univ. Press, 1974, pp. 43–59.

[4] ______, *An inequality for Riemann surfaces*, Differential geometry and complex analysis, Springer, Berlin, 1985, pp. 87–93.

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

[6] ______, *Geodesic laminations on surfaces*, Laminations and Foliations in Dynamics, Geometry and Topology (Stony Brook, NY) (M. Lyubich, J. Milnor, and Y. Minsky, eds.), Contemporary Mathematics, vol. 269, AMS, 2001, pp. 1–37.

[7] B. Bowditch, *Notes on Gromov's hyperbolicity criterion for path-metric spaces*, Group theory from a geometrical viewpoint (Trieste, 1990), World Scientific Publishing, 1991, pp. 64–167.

[8] J. Brock, *Boundaries of Teichmüller spaces and end-invariants for hyperbolic 3-manifolds*, Preprint (1999). To appear, Duke Math. J.

[9] ______, *Iteration of mapping classes and limits of hyperbolic 3-manifolds*, Preprint (2000), to appear in Invent. Math.

[10] ______, *Continuity of Thurston's length function*, Geom. Funct. Anal. **10** (2000), no. 4, 741–797.

[11] R. Brooks and J. P. Matelski, *Collars in Kleinian groups*, Duke Math. J. **49** (1982), no. 1, 163–182.

[12] P. Buser, *Geometry and Spectra of Compact Riemann Surfaces*, Birkhäuser, 1992.

[13] R. Canary, M. Culler, S. Hersonsky, and P. Shalen, *Density of cusps in boundaries of quasiconformal deformation spaces*, in preparation.

[14] R. D. Canary, *On the laplacian and the geometry of hyperbolic 3-manifolds*, J. Differential Geom. **36** (1992), 349–367.

[15] ______, *Algebraic limits of schottky groups*, Trans. Amer. Math. Soc. **337** (1993), 235–258.

[16] R. D. Canary, D. B. A. Epstein, and P. Green, *Notes on notes of Thurston*, Analytical and Geometric Aspects of Hyperbolic Space, Cambridge University Press, 1987, London Math. Soc. Lecture Notes Series no. 111, pp. 3–92.

[17] J. Cannon, *The theory of negatively curved spaces and groups*, Ergodic theory, symbolic dynamics, and hyperbolic spaces (Trieste, 1989), Oxford Univ. Press, 1991, pp. 315–369.

[18] A. J. Casson and S. A. Bleiler, *Automorphisms of surfaces after Nielsen and Thurston*, Cambridge University Press, 1988.

[19] D. B. A. Epstein and A. Marden, *Convex hulls in hyperbolic space, a theorem of Sullivan, and measured pleated surfaces*, Analytical and Geometric Aspects of Hyperbolic Space, Cambridge University Press, 1987, London Math. Soc. Lecture Notes Series no. 111, pp. 113–254.

[20] B. Farb and L. Mosher, *Schottky subgroups of mapping class groups*, In preparation.

[21] F. Gardiner, *Teichmüller theory and quadratic differentials*, Wiley Interscience, 1987.

[22] E. Ghys and P. de la Harpe, *Sur les groupes hyperboliques d'aprés Mikhael Gromov*, Birkhäuser, 1990.

[23] M. Gromov, *Hyperbolic groups*, Essays in Group Theory (S. M. Gersten, editor), MSRI Publications no. 8, Springer-Verlag, 1987.

[24] J. Harer, *Stability of the homology of the mapping class group of an orientable surface*, Ann. of Math. **121** (1985), 215–249.

[25] ______, *The virtual cohomological dimension of the mapping class group of an orientable surface*, Invent. Math. **84** (1986), 157–176.

[26] W. J. Harvey, *Boundary structure of the modular group*, Riemann Surfaces and Related Topics: Proceedings of the 1978 Stony Brook Conference (I. Kra and B. Maskit, eds.), Ann. of Math. Stud. 97, Princeton, 1981.

[27] J. Hempel, *3-manifolds as viewed from the curve complex*, MSRI Preprint #1997-095 and E-print math.GT/9712220 at http://front.math.ucdavis.edu.

[28] N. V. Ivanov, *Complexes of curves and the Teichmüller modular group*, Uspekhi Mat. Nauk **42** (1987), 55–107.

[29] ______, *The rank of Teichmüller modular groups*, Mat. Zametki **44** (1988), no. 5, 636–644, 701, translation in Math. Notes 44 (1988), no. 5-6, 829–832.

[30] ______, *Complexes of curves and Teichmüller spaces*, Math. Notes **49** (1991), 479–484.

[31] ______, *Subgroups of Teichmüller modular groups*, American Mathematical Society, Providence, RI, 1992, translated from the Russian by E. J. F. Primrose and revised by the author.

[32] ______, *Automorphisms of complexes of curves and of Teichmüller spaces*, Internat. Math. Res. Notices (1997), no. 14, 651–666.

[33] E. Klarreich, *The boundary at infinity of the curve complex and the relative Teichmüller space*, preprint.

[34] ______, *Semiconjugacies between Kleinian group actions on the Riemann sphere*, Amer. J. Math. **121** (1999), no. 5, 1031–1078.

[35] J. Lott, *L^2-cohomology of geometrically infinite hyperbolic 3-manifolds*, Geom. Funct. Anal. **7** (1997), no. 1, 81–119.

[36] H. A. Masur and Y. Minsky, *Geometry of the complex of curves I: Hyperbolicity*, Invent. Math. **138** (1999), 103–149.

[37] ______, *Geometry of the complex of curves II: Hierarchical structure*, Geom. Funct. Anal. **10** (2000), 902–974.

[38] C. McMullen, *Local connectivity, Kleinian groups and geodesics on the blowup of the torus*, Preprint: http://www.math.harvard.edu/~ctm/papers.

[39] ______, *Cusps are dense*, Ann. of Math. **133** (1991), 217–247.

[40] ______, *Renormalization and 3-manifolds which fiber over the circle*, Princeton University Press, Princeton, NJ, 1996.

[41] R. Meyerhoff, *A lower bound for the volume of hyperbolic 3-manifolds.*, Canad. J. Math. **39** (1987), 1038–1056.

[42] Y. Minsky, *Harmonic maps into hyperbolic 3-manifolds*, Trans. Amer. Math. Soc. **332** (1992), 607–632.

[43] ______, *Teichmüller geodesics and ends of hyperbolic 3-manifolds*, Topology **32** (1993), 625–647.

[44] ______, *On rigidity, limit sets and end invariants of hyperbolic 3-manifolds*, J. Amer. Math. Soc. **7** (1994), 539–588.

[45] ______, *On Thurston's ending lamination conjecture*, Proceedings of Low-Dimensional Topology, May 18-23, 1992, International Press, 1994.

[46] ______, *The classification of punctured-torus groups*, Annals of Math. **149** (1999), 559–626.

[47] ______, *Kleinian groups and the complex of curves*, Geometry and Topology **4** (2000), 117–148.

[48] ______, *Short geodesics and end invariants*, Comprehensive Research on Complex Dynamical Systems and Related Fields (M. Kisaka and S. Morosawa, eds.), vol. 1153, RIMS Kokyuroku, 2000, E-print: arXiv:math.GT/0006002, pp. 1–20.

[49] K. Ohshika, *Ending laminations and boundaries for deformation spaces of Kleinian groups*, J. London Math. Soc. **42** (1990), 111–121.

[50] K. Rafi, *Hyperbolic 3-manifolds and geodesics in Teichmüller space*, PhD Thesis, SUNY at Stony Brook, 2001.

[51] W. Thurston, *Hyperbolic structures on 3-manifolds, II: surface groups and manifolds which fiber over the circle*, E-print: math.GT/9801045 at http://front.math.ucdavis.edu. Original 1986.

[52] ———, *The geometry and topology of 3-manifolds*, Princeton University Lecture Notes, online at http://www.msri.org/publications/books/gt3m, 1982.

[53] ———, *Three dimensional manifolds, Kleinian groups and hyperbolic geometry*, Bull. Amer. Math. Soc. **6** (1982), 357–381.

[54] ———, *Hyperbolic structures on 3-manifolds, I: deformation of acylindrical manifolds*, Ann. of Math. **124** (1986), 203–246.

SUNY Stony Brook