SHORT GEODESICS AND END INVARIANTS

YAIR N. MINSKY

Even topologically simple hyperbolic 3-manifolds can have very intricate geometry. Consider in particular a closed surface S of genus 2 or more, and the product $N = S \times \mathbf{R}$. This 3-manifold admits a large family of complete, infinite-volume hyperbolic metrics, corresponding to faithful representations $\rho : \pi_1(S) \to \mathrm{PSL}_2(\mathbf{C})$ with discrete image.

The geometries of N are very different from the product structure that its topology would suggest. Typically, N contains a complicated pattern of "thin" and "thick" parts. The thin parts are collar neighborhoods of very short geodesics, typically infinitely many. Each one, called a "Margulis tube", has a well-understood shape, but the way in which these are arranged in N, and in particular the identities of the short geodesics as elements of the fundamental group, are still something of a mystery.

This issue is closely related to the basic classification conjecture associated with these manifolds, Thurston's "ending lamination conjecture". This conjecture states that certain asymptotic invariants of the geometry of N, called ending invariants, in fact determine N completely. (Actually the classification of hyperbolic structures for any manifold with incompressible boundary reduces to this case, by restriction to boundary subgroups.)

In this expository paper we will focus on the following question: What information do the ending invariants give about the presence of very short geodesics in the manifold? We will summarize and discuss the theorem below, part of whose proof appears in [40] and part of which will be in [33], as well as a few conjectures.

Bounded Geometry Theorem. *Let S be a closed surface, and consider a Kleinian surface group $\rho : \pi_1(S) \to PSL_2(\mathbf{C})$ with no externally short curves, and ending invariants ν_+ and ν_-. Then*

$$\inf_{\gamma\in\pi_1(S)} \ell_\rho(\gamma) > 0 \iff \sup_{Y\subset S} d_Y(\nu_+,\nu_-) < \infty.$$

Here the supremum is over proper essential isotopy classes of subsurfaces in S, and the quantities $d_Y(\nu_+,\nu_-)$, called "projection coefficients", are defined in Section 1.4. The quantity $\ell_\rho(\gamma)$ is the translation distance of $\rho(\gamma)$ in $\mathbf{H}^3$, or the length of the closed geodesic associated to γ in the 3-manifold. (The condition on externally short curves is not really necessary – it is added to simplify the other definitions and discussions – see §1.1 below).

Date: March 27, 2000.

Part of our goal is to advertise a combinatorial object known as the *complex of curves on a surface*, as a tool for studying the geometry of hyperbolic 3-manifolds. This object is used for definining the coefficients d_Y, and in general it encodes something about the structure of the set of simple loops on a surface. In particular, face transitions between simplices in this complex correspond to elementary moves on pants decompositions of S, and these in turn correspond to homotopies between elementary pleated surfaces in a hyperbolic 3-manifold. The interaction between the combinatorial and geometric aspects of these moves is our main object of study, and seems to be worthy of further consideration.

1. Definitions

1.1. **Surface groups and ending laminations.** Let S be a closed surface of genus $g \geq 2$. A *Kleinian surface group* will be a representation $\rho : \pi_1(S) \to \mathrm{PSL}_2(\mathbf{C})$, discrete and faithful. The quotient $\mathbf{H}^3/\rho(\pi_1(S))$ is denoted N_ρ, and comes equipped with a homotopy class of homotopy equivalences $S \to N_\rho$, determined by ρ. In fact N_ρ is homeomorphic to $S \times \mathbf{R}$, by Thurston's theory of tame ends [45] and Bonahon's Tameness theorem [7].

We can associate to ρ two *ending invariants* ν_- and ν_+, which we will describe in the special case that ρ has *no parabolics* (see also [36] and Ohshika [42]).

Let $C(N_\rho)$ be the *convex core* of N_ρ, the smallest convex submanifold whose inclusion is a homotopy equivalence. Fixing an orientation on S and N_ρ, there is an orientation-preserving homeomorphism of N_ρ to $S\times\mathbf{R}$ taking $C(N_\rho)$ onto exactly one of $S \times \mathbf{R}$, $S \times [0, \infty)$, $S \times (-\infty, 1]$ or $S \times [0, 1]$.

The end of N defined by neighborhoods $S \times (a, \infty)$ is called e_+, and the one defined by $S \times (-\infty, a)$ is called e_-. If an end's neighborhoods all meet the convex hull it is called *geometrically infinite*, and otherwise it is *geometrically finite*. Suppose e_+ is geometrically finite. Then the component $\partial_+(C(N_\rho))$ corresponding to $S \times \{1\}$ is a convex surface, and its exterior $S \times (1, \infty)$ develops out to a "conformal structure at infinity" on S, which we call ν_+. (This surface is obtained from the action of $\rho(\pi_1(S)$ on the Riemann sphere). We define ν_- in the same way when e_- is geometrically finite.

Thurston pointed out that boundary $\partial_+(C(N_\rho))$ is itself a hyperbolic surface; let us call its structure ν'_+. A theorem of Sullivan (proof in Epstein-Marden [15]) states that ν'_+ and ν_+ differ by a uniformly bilipschitz distortion.

To describe the invariant of a geometrically infinite end we need to briefly recall the notion of a *geodesic lamination*. Fixing a hyperbolic metric on S, a geodesic lamination is a closed subset of S foliated by geodesics. Let $\mathcal{GL}(S)$ denote the set of all of these. A *measured lamination* is a geodesic lamination equipped with a Borel measure on transverse arcs, invariant under transverse isotopy. The space $\mathcal{ML}(S)$ of measured laminations admits a natural topology coming from weak-* convergence of the measures. On the supporting geodesic laminations, this is related to but not quite the same as the topology of Hausdorff convergence. However the difference will not be important to us here. Simple closed geodesics with positive weights are dense in $\mathcal{ML}(S)$, and we will consider geodesic laminations obtained as supports of limits in $\mathcal{ML}(S)$ of sequences of simple closed curves. Finally we remark that the choice of metric on S is irrelevant, as any other choice yields naturally isomorphic spaces of laminations. For more details on this topic see Bonahon [5, 6], Canary-Epstein-Green [13], or Casson-Bleiler [14].

If e_+ is geometrically infinite then the convex hull contains an infinite sequence of closed geodesics γ_n, all homotopic to *simple* closed loops on S, and eventually contained in $S \times (a, \infty)$ for any a. This is a theorem of Bonahon, and Thurston (previously) showed that for such a sequence the curves on S must converge in the sense of the previous paragraph to a unique geodesic lamination on S. We call this lamination ν_+, the ending lamination of e_+. The corresponding discussion for e_- gives ν_-.

Finally let us define the technical simplifying condition in the statement of the Bounded Geometry Theorem. Call a curve γ in S *externally short*, with respect to a representation ρ, if it is either parabolic or has length less than ϵ_1 with respect to the structures ν_- and ν_+ (if these are not laminations), where ϵ_1 is some fixed constant small enough that there exist hyperbolic structures on S with no curves of length less than ϵ_1. Note in particular that if ρ has two degenerate ends then it automatically has no externally short curves.

1.2. **Pleated surfaces.** 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, called the pleating locus, 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. These were introduced by Thurston [45], and we will see some explicit examples in §4.1.

It is a consequence of the work of Thurston and Bonahon that a geometrically infinite end of a surface group ρ admits pleated surfaces in the homotopy class of ρ contained in any neighborhood of the end. The pleating loci of these surfaces must converge to the ending lamination, and their hyperbolic structures converge to this lamination in Thurston's compactification of the Teichmüller space.

1.3. **Complexes of arcs and curves:** Let Z be a compact finite genus 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 the 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.

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".

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 $\mathcal{C}(S)$ except in a few simple cases when $\mathcal{C}(S)$ has no edges. When $\partial Z = \emptyset$, of course $\mathcal{A} = \mathcal{C}$.

It is a nice exercise to compute $\mathcal{A}(Z)$ exactly for Z a one-holed torus, and we leave this to the reader. The answer is closely related to the Farey graph in the plane – see [37].

Fix our closed surface S and let $\mathcal{GL}(S)$ denote the set of geodesic laminations on S (note that $\mathcal{A}_0(S) = \mathcal{C}_0(S)$ can identified with a subset of $\mathcal{GL}(S)$). Let $Y \subset S$ be a proper essential closed subsurface (all boundary curves homotopically nontrivial). 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 $\pi_1(Y) \subset \pi_1(S)$, which can be naturally compactified using the circle at infinity of the universal cover of S to yield a surface $\widehat{Y}$ homeomorphic to Y (remove the limit set of $\pi_1(Y)$ and take the quotient of the rest). Any lamination $\lambda \in \mathcal{GL}(S)$ lifts to this cover as a collection of closed curves or arcs that have well-defined endpoints in $\partial\widehat{Y}$. Removing the trivial components, we have a simplex of $\mathcal{A}(\widehat{Y})$ and we can take, say, its barycenter (we can also get the empty set if there are no essential components). A version of this projection also appears in Ivanov [26, 24].

If $\beta, \gamma \in \mathcal{GL}(S)$ (in particular in $\mathcal{C}(S)$ have non-trivial intersection with 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 naturally identified 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

$$d_\alpha = d_Y.$$

See [16] for an application of this construction in the annulus case.

The complex of curves $\mathcal{C}(S)$ was first introduced by Harvey [20]. It was applied by Harer [18, 19] and Ivanov [23, 27, 25] to study the mapping class group of S. Similar complexes were introduced by Hatcher-Thurston [22]. Masur-Minsky [30] proved that $\mathcal{C}(S)$ is δ-hyperbolic in the sense of

Gromov, and then applied this in [29] to prove the structural theorems on pants decompositions that we will use in Section 4.

1.4. **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 ending invariants for a surface group. 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:

If σ is a hyperbolic metric on S, and L_1 a fixed constant, define

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

to be the set of pants decompositions of S with total σ-length at most L_1. A theorem of Bers (see [3, 4] and Buser [10]) says that L_1 can be chosen, depending only on genus of S, so that $\mathbf{short}(\sigma)$ is always non-empty. Let us also choose L_1 sufficiently large that, if σ has no curves of length less than ϵ_1 (the constant from the end of §1.1), then *every* curve in S intersects some $P \in \mathbf{short}(\sigma)$.

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. 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. Note that the condition that ρ has no externally short curves implies that $d_Y(\nu_+, \nu_-)$ is well-defined for all Y.

2. Margulis tubes

Let γ be a loxodromic element of a Kleinian group Γ. We denote its complex translation length by $\lambda(\gamma) = \ell + i\theta$ (determined mod $2\pi i$). Let $\mathcal{T}_\epsilon$ be the γ-invariant set $\{x \in \mathbf{H}^3 : \inf_n d(x, \gamma^n(x)) \leq \epsilon\}$. If $\ell(\gamma) < \epsilon$ This is a tube of some radius r around the axis of γ, and The Margulis Lemma and Thick-Thin decomposition tell us (see e.g. [28, 46, 1]) that there is a universal constant ϵ_0 such that if $\ell(\gamma) < \epsilon_0$ then $\mathcal{T}_{\epsilon_0}/\langle\gamma\rangle$ embeds as a solid torus $\mathbf{T}_\gamma$ in $N = \mathbf{H}^3/\Gamma$, called a *Margulis tube*, and furthermore that all Margulis tubes in N are disjoint.

The radius r of the tube goes to ∞ as the length of the core goes to 0. See Brooks-Matelski [9] and Meyerhoff [32] for more precise bounds.

Thus in some sense the geometry around a very short curve in N is very well understood. It is more difficult to determine the pattern in which these tubes are arranged in the manifold, and in particular which curves γ have length less than a given ϵ.

2.1. **Margulis tubes in surface groups.** When Γ is the image $\rho(\pi_1(S))$ of a Kleinian surface group, there is a little more we can say. An observation of Thurston [44], together with Bonahon's tameness theorem [7], imply that *only simple curves can be short:* that is, ϵ_0 may be chosen so that, if $\ell_\rho(\gamma) < \epsilon_0$ and γ is a primitive element of $\pi_1(S)$ then γ is represented by a simple loop in S. This is because, by Bonahon's theorem, every point in N_ρ is uniformly near the image of a pleated surface. Thurston pointed out using a simple area bound that if ϵ_0 is sufficiently short a π_1-injective pleated surface can only meet $\mathbf{T}_\gamma$ in the image of its own 2-dimensional Margulis tube. The core of this tube must therefore be γ.

2.2. **Bounds.** An *upper bound* on the length of a curve in a surface group can be obtained in terms of the conformal boundary at infinity. Bers showed [2] for a Quasi-Fuchsian representation ρ, that

$$\frac{1}{\ell_\rho(\gamma)} \geq \frac{1}{2}\left(\frac{1}{\ell_+(\gamma)} + \frac{1}{\ell_-(\gamma)}\right)$$

where $\ell_\pm$ denote lengths in the hyperbolic structures on S coming from the two conformal structures $\nu_\pm$ at infinity. The argument uses a monotonicity property for conformal moduli and the action of γ on the Riemann sphere. When S is a once-punctured torus this upper bound can be generalized to an estimate in both directions (see [39]). In general we have no such result, but in Section 5 we will state a conjectural estimate.

3. Bounded geometry

We say that ρ has *bounded geometry* if there is a positive lower bound on the translation lengths of all group elements. This condition incidentally disallows parabolic elements (in a more general discussion we would allow them and revise the condition), but the real point is that there is a positive lower bound on the lengths of all closed geodesics in the quotient manifold.

In [34, 35], we showed that bounded geometry implies a positive solution to the ending lamination conjecture. That is, if ρ_1 and ρ_2 both have bounded geometry, and have the same ending invariants, then they are conjugate in $\mathrm{PSL}_2(\mathbf{C})$. This result was accompanied by a fairly explicit bilipschitz model for the metric on N, derived from the Teichmüller geodesic joining the two ending laminations.

The Bounded Geometry theorem gives us a way to strengthen this result, since it implies that bounded geometry is detected by the ending invariants:

Corollary 3.1. *Let ρ_1, ρ_2 be Kleinian surface groups with the same ending invariants, and suppose that ρ_1 has bounded geometry. Then ρ_1 and ρ_2 are conjugate in $PSL_2(\mathbf{C})$.*

It is worth noting 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 [31, Cor. 1.6], and Canary-Culler-Hersonsky-Shalen [11] for generalizations).

4. The proof of the bounded geometry theorem

The proof of the direction ($\Rightarrow$) of the Bounded Geometry Theorem appears in [40]. The essential tool used there is Thurston's "efficiency of pleated surfaces" theorem from [44]. We will outline the proof of ($\Leftarrow$), for which the details will appear in [33].

In roughest form, the argument is this: Let $\gamma \in \pi_1(S)$ be an element with $\ell_\rho(\gamma) < \epsilon_0$, and let $\mathbf{T}_\gamma$ be its Margulis tube. We will use the condition $\sup d_Y(\nu_+, \nu_-) < \infty$ to construct a sequence of pleated surfaces $\{f_i\}_{i=0}^M$ with the following properties:

1. The size M of the sequence is bounded by
$$M \leq K(\sup d_Y(\nu_+, \nu_-))^a$$
where K, a depend only on the genus of S.
2. Any successive f_i, f_{i+1} are connected by a homotopy $H : S \times [i, i+1] \to N_\rho$ which is *uniformly bounded* except in a special case, described below.
3. The total homotopy $H : S \times [0, M] \to N_\rho$ homologically encloses $\mathbf{T}_\gamma$.

Part (3) means that the image of H must cover all of $\mathbf{T}_\gamma$. Thus, if the "special case" of (2) does not occur, then the bounds of (1) and (2) give a uniform diameter bound on $\mathbf{T}_\gamma$, and hence a lower bound on $\ell_\rho(\gamma)$.

The "special case" of (2) corresponds to the curve γ itself appearing in the pleating locus of some subsequence of the f_i. In this case a more delicate argument is needed, using the annulus projection distance $d_\gamma(\nu_+, \nu_-)$ to bound the size of $\mathbf{T}_\gamma$.

Let us now introduce the ingredients needed for this construction. In §4.5 we will return to the main proof.

4.1. Adapted pleated surfaces.

If Q is a collection of disjoint, homotopically distinct curves on S (henceforth a "curve system"), and ρ a fixed Kleinian surface group, we let $\mathbf{pleat}(Q, \rho)$ denote the set of pleated surfaces $f : S \to N_\rho$, in the homotopy class determined by ρ, which map representatives of Q to geodesics. There is the usual equivalence relation on this set, in which $f \sim f \circ h$ if h is a homeomorphism of S homotopic to the identity. Let σ_f denote the hyperbolic metric on S induced by f.

In particular, if Q is a maximal curve system, or "pants decomposition", $\mathbf{pleat}(Q, \rho)$ consists of finitely many equivalence classes, all constructed as follows: Extend Q to a triangulation of S with one vertex on each component of Q, and "spin" this triangulation around Q, arriving at a lamination λ whose closed leaves are Q and whose other leaves spiral onto Q, as in Figure 1.

A unique pleated surface (up to equivalence) exists carrying λ to geodesics, since no element of Q is parabolic (by hypothesis on ρ). This was originally

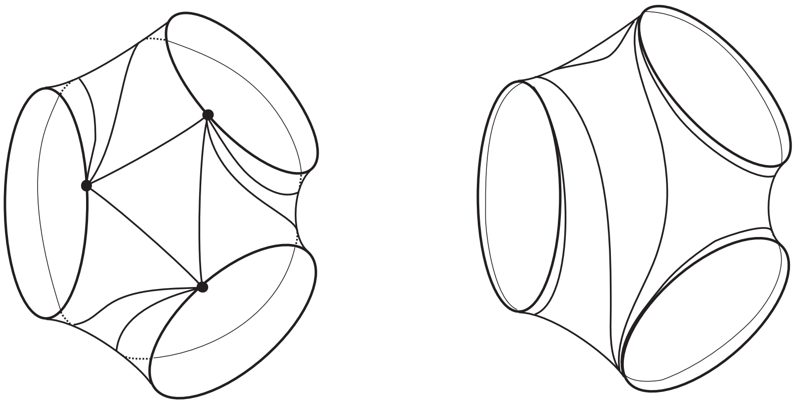

FIGURE 1. The lamination obtained by spinning a triangulation around a curve system. The picture shows one pair of pants in a decomposition.

observed by Thurston (see [45] and Canary-Epstein-Green [13, Thm 5.3.6] for a proof). The choices of λ coming from the finite number of possible triangulations up to isotopy, and the different directions of spiraling, account for all of $\mathbf{pleat}(Q,\rho)$.

4.2. **Elementary moves.** 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 2.

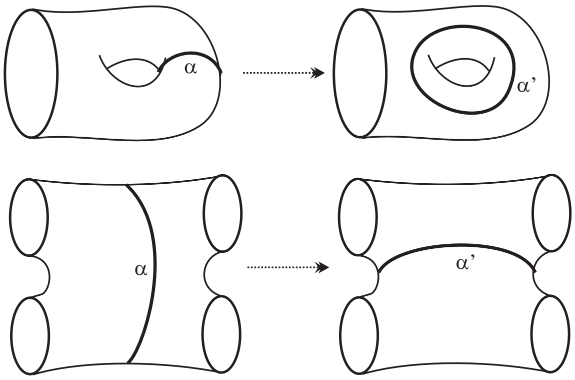

FIGURE 2. The two types of elementary moves.

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 $\mathbf{Z}$.

Pleated surfaces associated to an elementary move are homotopic in a controlled way. Let us first recall (see Buser [10]) that a simple geodesic γ in a hyperbolic surface (S,σ) always admits a "standard collar", which is an annulus of radius depending only on $\ell_\sigma(\gamma)$, such that disjoint geodesics have disjoint collars, and when $\ell_\sigma(\gamma) < \epsilon_0$ the collar covers all but a bounded part of the ϵ_0-Margulis tube. We write this collar as $\mathbf{collar}(\gamma, \sigma)$, or $\mathbf{collar}(P, \sigma)$ for the union of collars over a curve system P.

Lemma 4.1. (Elementary Homotopy) *If $P_0 \to P_1$ is an elementary move exchanging α_0 and α_1, ρ is a Kleinian surface group, and $f_i \in \mathbf{pleat}(P_i, \rho)$ for $i = 0, 1$, then there exists a homotopy $H : S \times [0, 1] \to N_\rho$ with the following properties:*

1. *$H_0 \sim f_0$ and $H_1 \sim f_1$ under the usual equivalence.*
2. *If σ_i is the induced metric of H_i (for $i = 0, 1$) then $\mathbf{collar}(P_j, \sigma_i) = \mathbf{collar}(P_j, \sigma_{1-i})$, for $j = 0, 1$.*
3. *The metrics σ_0 and σ_1 are K-bilipschitz except possibly when $\ell_\rho(\alpha_i) < \epsilon_0$ for $i = 0$ or 1. In that case the metrics are locally K-bilipschitz outside $\mathbf{collar}(\alpha_0) \cup \mathbf{collar}(\alpha_1)$ (or just one collar if only one curve is short in N_ρ).*
4. *The trajectories $H(p \times [0, 1])$ are bounded in length by K except possibly when $p \in \mathbf{collar}(\alpha_i)$ and $\ell_\rho(\alpha_i) < \epsilon_0$, in which case they are bounded outside of $\mathbf{T}_{\rho(\alpha_i)}$.*

The constant K depends only on the genus of S.

(Note that $\mathbf{collar}(\alpha_i)$ in (3) and (4) makes sense without specifying the metric σ_j, since the two are equal by (2).)

It is worth pointing out that this theorem applies without any a-priori bounds on the lengths $\ell_\rho(P_i)$. The proof is an application of Thurston's Uniform Injectivity theorem for pleated surfaces, and the closely related Efficiency of Pleated Surfaces [44] (see also Canary [12]). These theorems control the amount kind of bending that can occur in a pleated surface, and in particular can be used to compare two pleated surfaces that share part of their pleating locus.

We also remark that part (2) is just for convenience – it is easy to arrange by an appropriate isotopy.

4.3. **Resolution sequences.** In [29], 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 simple closed curve, denote

$$J_\beta = \{i \in [0, n] : \beta \in P_i\}.$$

(Here $\beta \in P$ means β is a component of P.) We also denote $J_{\beta_1, \ldots, \beta_k} = \cup J_{\beta_i}$.

Note that if β is a curve and 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.

Theorem 4.2. (Controlled Resolution Sequences) *Let P and Q be maximal curve systems in S. There exists a geodesic $\beta_0, \ldots, \beta_m$ in $\mathcal{C}_1(S)$ 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_i contains some β_j.*
3. *J_β, if nonempty, is always an interval, and if $[i,j] \subset [0,m]$ then*
$$|J_{\beta_i,\ldots,\beta_j}| \leq K(j-i) \sup_Y d_Y(P,Q)^a,$$
where the supremum is over only those subsurfaces Y whose boundary curves are components of some P_k with $k \in J_{\beta_i,\ldots,\beta_j}$.
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 genus 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 [29] by an inductive procedure: beginning with a geodesic $\{\beta_i\}$ in $\mathcal{C}_1(S)$ joining P to Q (we are describing a geodesic here as a sequence of vertices where successive ones are joined by edges), we note that the link of each β_i is itself a curve complex for a subsurface. In each such complex we add a new geodesic, and repeat. The final structure can then be "resolved" into a sequence of maximal curve systems. Control of the size of the construction at each stage is achieved by applying the hyperbolicity theorem of [30].

4.4. **Contraction and quasi-convexity.** Let $\mathcal{C}(S,\rho,L)$ denote the subcomplex of $\mathcal{C}$ spanned by the vertices with ρ-length at most L. We will define a map $\Pi_\rho : \mathcal{C}(S) \to \mathcal{P}(\mathcal{C}(S,\rho,L_1))$, where $\mathcal{P}(X)$ is the power set of X, as follows. For $x \in \mathcal{C}(S)$, let P_x be the curve system associated to the smallest simplex containing x. We define
$$\Pi_\rho(x) = \bigcup_{f \in \mathbf{pleat}(P_x,\rho)} \mathbf{short}(\sigma_f).$$
This map turns out to have coarsely the properties of a closest-point projection to a convex subset of a hyperbolic space.

Lemma 4.3. (Contraction Properties) *There are constants $b, c > 0$, depending only on the genus of S, such that for any ρ the map Π_ρ has the following properties:*

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

3. *(Contraction) If $r = d_{\mathcal{C}}(x, \Pi_\rho(x))$ then*

$$\operatorname{diam}_{\mathcal{C}} \Pi_\rho(B(x, cr)) \le b.$$

Here $d_{\mathcal{C}}$ and $\operatorname{diam}_{\mathcal{C}}$ refer to distance and diameter measured in $\mathcal{C}(S)$, and $B(x,s)$ is a ball of $d_{\mathcal{C}}$-radius s around x. By $\Pi_\rho(X)$ for a set X we mean $\cup_{x\in X}\Pi_\rho(x)$.

Compare this with the contraction property in [30], which was used to prove hyperbolicity of $\mathcal{C}(S)$, and the property in [38], which was used to prove stability properties for certain geodesics in Teichmüller space.

An easy consequence of this theorem is the following quasiconvexity property for $\mathcal{C}(S, \rho, L_1)$:

Lemma 4.4. (Quasiconvexity) *If $\beta_0, \dots, \beta_m$ is a geodesic in $\mathcal{C}_1(S)$ and $\beta_0, \beta_m \in \mathcal{C}(S, \rho, L_1)$, then*

$$d_{\mathcal{C}}(\beta_i, \Pi_\rho(\beta_i)) \le C$$

for all $i \in [0, m]$ and a constant C depending only on the genus of S.

In particular a geodesic with endpoints in $\mathcal{C}(S, \rho, L_1)$ never strays too far from $\mathcal{C}(S, \rho, L_1)$. This can be compared to the "Connectivity" lemma in [39].

The argument for this lemma is very simple, and has its origins in the stability of quasi-geodesics argument in the proof of Mostow's Rigidity Theorem [41]: We compare the path $\{\beta_i\}$ to its image "quasi-path" $\{\Pi_\rho(\beta_i)\}$. If the distance between these grows too much then the images slow down because of the Contraction property (3). Since $\{\beta_i\}$ is a shortest path and the two paths have nearly the same endpoints (Coarse idempotence (2)), there is a bound on how far apart they can get.

The proof of Lemma 4.3 is another application of Thurston's Uniform Injectivity theorem, as well as some of the tools developed in [30]. For example, to prove part (1), we note that if two vertices of $\mathcal{C}(S)$ are at distance 1 then they correspond to disjoint curves, and hence a pleated surface exists that maps both geodesically. Thus the argument reduces to bounding $\Pi_\rho(x)$ for one x. Suppose two pleated surfaces share a curve x. If x is short then their short curve sets intersect, and we finish by noting that $\operatorname{diam}_{\mathcal{C}}(\mathbf{short}(\sigma))$ is uniformly bounded for any σ. If the curve x is long, then in one of the pleated surfaces we can find a curve x' of bounded length that runs along x and then makes a very small jump in its complement (a long curve in a hyperbolic surface must run very close to itself). The Uniform Injectivity theorem is then applied to show that x' can be realized with bounded length on the second pleated surface as well.

Part (3) is the main point of the lemma. Its proof depends on the analysis in [30], which shows roughly that if $x \in \mathcal{C}(S)$ is far in $\mathcal{C}(S)$ from the short curves of a given hyperbolic metric σ on S, then sets of the form $B(x, R)$ for large R can be carried in a long nested chain of "train tracks" (see [43]) whose branches mostly run nearly parallel to x. These train tracks are then used to control $\Pi_\rho(B(x,R))$, via a Uniform Injectivity argument similar to the previous paragraph.

4.5. **Building a resolution sequence for ρ.** We can now use Theorem 4.2 (Controlled Resolution Sequences) and Lemma 4.4 (Quasiconvexity) to produce a resolution sequence adapted to the geometry of our representation ρ.

As a starting point we need an initial and terminal curve system:

Lemma 4.5. *Given ρ with no externally short curves, and a Margulis tube $\mathbf{T}_\gamma$ in N_ρ, there exist maximal curve systems P_+ and P_-, and pleated surfaces f_+, f_- (in the homotopy class of ρ) with the following properties:*

1. *$P_\pm \in \mathbf{short}(\sigma_{f_\pm})$,*
2. *f_+ and f_- homologically encase $\mathbf{T}_\gamma$.*

This is done roughly as follows. If ν_+ is a lamination then there exists a sequence g_i of pleated surfaces exiting the end of N_ρ corresponding to ν_+. The curves in $\mathbf{short}(\sigma_{g_i})$ converge to ν_+ in the space of laminations (modulo measure), and for large enough i, g_i can be deformed to e_+ without meeting $\mathbf{T}_\gamma$. We can pick $f_+ = g_i$ and let $P_+ \in \mathbf{short}(\sigma_{g_i})$. The same goes for f_-, P_-, so if both invariants are laminations we have the conclusion that f_+ and f_- must encase $\mathbf{T}_\gamma$.

If the end e_+ is geometrically finite we can let f_+ be the pleated map to the convex hull boundary itself, and similarly for e_-. Again choose $P_\pm \in \mathbf{short}(\sigma_{f_\pm}) = \mathbf{short}(\nu'_\pm)$.

Note that, if the pleated surfaces g_i are chosen far enough out the end (in the geometrically infinite case) then the homotopy from g_i to a map in $\mathbf{pleat}(P_+,\rho)$ does not pass through $\mathbf{T}_\gamma$, and so we may assume $f_\pm \in \mathbf{pleat}(P_\pm,\rho)$ and still have the encasing condition. When there are geometrically finite ends this is trickier because $\mathbf{T}_\gamma$ may be close to the convex hull boundary. Slightly more care is needed in the rest of the construction in that case. Let us from now on assume that $f_\pm \in \mathbf{pleat}(P_\pm,\rho)$, and the encasing condition holds.

Join P_+ to P_- with a resolution sequence $P_- = P_0 \to \cdots \to P_n = P_+$, as in Theorem 4.2. Let $\{\beta_i\}_{i=0}^m$ be the associated geodesic. This sequence may be much longer than we need, so we will use Lemma 4.4 to find a suitable subsequence. Recall that we would like our sequence to have the property of homologically encasing $\mathbf{T}_\gamma$, so let us try to throw away those surfaces that we are sure cannot meet $\mathbf{T}_\gamma$. In particular, let $f \in \mathbf{pleat}(P_i,\rho)$ for some $i \in [0,n]$, and let P_i contain a curve β_j. If $f(S) \cap \mathbf{T}_\gamma \neq \emptyset$, then γ itself is short in σ_f (as in §2.1) and so γ is distance 1 from $\Pi_\rho(\beta_j)$. It follows from Lemma 4.4 that

$$d_{\mathcal{C}}(\beta_j,\gamma) \le C \qquad (*)$$

where C is a new constant depending only on the genus of S. Thus we conclude that there is a subinterval I_γ of $[0,m]$ of diameter at most $2C$, such that f can only meet $\mathbf{T}_\gamma$ when β_j satisfies $j \in I_\gamma$. Let us therefore restrict our elementary move sequence to

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

where $[s,t] = \cup_{j\in I_\gamma} J_{\beta_j}$, and renumber it as $P_0 \to \cdots \to P_M$. This subsequence must also encase $\mathbf{T}_\gamma$, since none of the pieces we have thrown away can meet $\mathbf{T}_\gamma$. Part (3) of Theorem 4.2 tells us that

$$M \le K(2C) \sup_Y d_Y(P_+, P_-)^a,$$

where the supremum is over subsurfaces Y whose boundaries appear among the P_i in our subsequence. This means by $(*)$ that the $\mathcal{C}(S)$-distance $d_{\mathcal{C}}(\partial Y, \gamma)$ is bounded by $C + 1$ for all such Y. The analysis of [29] shows that, for a fixed such bound,

$$d_Y(P_+, P_-) \le d_Y(\nu_+, \nu_-) + \delta$$

with δ depending only on the genus of S, provided, when e_+ or e_- are geometrically infinite, that the surface $f_\pm$ are taken sufficiently far out in the ends (for geometrically finite ends this is an easier consequence of Sullivan's theorem comparing $\nu_\pm$ with $\nu'_\pm$, though here we must take a bit more care with the constants to make sure that ∂Y intersects $P_\pm$). Since the right side is a priori bounded by hypothesis, we obtain our desired uniform bound on M.

Now let $H : S \times [i, i+1] \to N_\rho$ be the homotopy provided by Lemma 4.1 (Elementary Homotopy), where $H_i \in \mathbf{pleat}(P_i, \rho)$. After possibly adjusting by homeomorphisms of S homotopic to the identity, we can piece these together to a map $H : S \times [0, M] \to N_\rho$.

Assume first that γ is not a component of any P_i. Then according to Lemma 4.1, H can make only uniformly bounded progress through the Margulis tube $\mathbf{T}_\gamma$. Thus diam $\mathbf{T}_\gamma$ is bounded above, and $\ell_\rho(\gamma)$ is bounded below, and we are done.

Now suppose that γ does appear in the $\{P_i\}$. Then J_γ is some subinterval of $[0, M]$ by Theorem 4.2, 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 4.2 that $d_\gamma(\alpha, \alpha')$ is uniformly approximated by $d_\gamma(P_+, P_-)$ and hence uniformly bounded.

For simplicity, let us consider now the case that both $\ell_\rho(\alpha)$ and $\ell_\rho(\alpha')$ are uniformly bounded above and below. (There is in fact a uniform upper bound on their lengths; if they become too short a small additional argument is needed).

Let $\sigma_i \equiv \sigma_{H_i}$ and note that, by Lemma 4.1, for all $i \in J_\gamma$ the annuli $\mathbf{collar}(\gamma, \sigma_i)$ coincide. Name this common annulus B. Write $J_\gamma = [k, l]$, and consider in $S \times [0, M]$ the solid torus

$$U = B \times [k-1, l+1].$$

By Lemma 4.1, this is the only part of $S \times [0, M]$ that H can map more than a bounded distance into $\mathbf{T}_\gamma$. The height of this torus, $|k - l + 2|$, is at most M and this is uniformly bounded. The top and bottom annuli $B\times\{k-1\}$ and $B\times\{l+1\}$ have uniformly bounded geometry (in σ_{k-1} and σ_{l+1}, respectively), by the length bounds we've assumed on α and α'. We will control the size of the *meridian* of U, and this will in turn bound the size of $\mathbf{T}_\gamma$.

Assume α is a geodesic in σ_{k-1} (where we note its length is bounded above), and let $a = \alpha\cap B$. Similarly assume α' is a geodesic in σ_{l+1} and let $a' = \alpha'\cap B$. The arc a may a priori be long in σ_{l+1}, but its length is estimated by the number of times it twists around a', or $d_{\mathcal{A}(B)}(a,a')$.

A lemma in 2 dimensional hyperbolic geometry establishes

$$|d_{\mathcal{A}(B)}(a,a') - d_\gamma(\alpha,\alpha')| \le C$$

where this C depends only on M, which we have already bounded uniformly. The idea of this is that, in each elementary move, the metric σ_i changes in a bilipchitz way outside the collars of the curves involved in the elementary move. From this it follows that, starting with a geodesic passing through a collar, we obtain a curve which does only a bounded amount of additional twisting outside the collar. After M such moves the relative twisting of α and α' can still be estimated by their twisting inside the collar, up to an additive bound proportional to M.

With this estimate and the bound on $d_\gamma(\alpha,\alpha')$ in terms of $d_\gamma(P_+,P_-)$, we find that a and a' intersect a bounded number of times, so that the length of a is uniformly bounded in $S\times\{l+1\}$. It follows that the meridian of U

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

has uniformly bounded length in the induced metric. Thus its image is bounded in N_ρ. It therefore spans a disk of bounded diameter, and in fact we can homotope H on all of U to a new map of bounded diameter. This bounds the diameter of $\mathbf{T}_\gamma$ from above, and again we are done.

5. Conjectures

5.1. **Length estimates.** The reader may have noticed that in fact the argument outlined in the previous section shows that the infimum $\epsilon = \inf_\gamma \ell_\rho(\gamma)$ and the supremum $D = \sup_Y d_Y(\nu_+,\nu_-)$ can be bounded one in terms of the other. That is, any positive lower bound for ϵ implies some upper bound for D independent of ρ, and vice versa. Thus there is a version of the theorem which yields non-empty information for quasi-Fuchsian groups (where $\epsilon > 0$ and $D < \infty$ automatically) as well. However it would be nice to have bounds that are more specific and more explicit.

"More specific" means that we would like to know an estimate on $\ell_\rho(\gamma)$ for *a particular* γ. In [40] we actually show that for any subsurface Y, a large lower bound on $d_Y(\nu_+,\nu_-)$ implies a small upper bound for $\ell_\rho(\partial Y)$. In the other direction something more complicated would need to be stated, since any curve γ can be a boundary curve for many different subsurfaces.

"More explicit" means we would like to know the estimate itself more explicitly. Furthermore it would be nice to estimate the *complex translation length* λ and not just its real part ℓ. In [39] this was done for the punctured-torus case. Here is a possible generalization, stated again in the case of ρ with no externally short curves.

Conjecture 5.1. *Let ρ be a Kleinian surface group with no externally short curves. There exist $K, \epsilon > 0$ depending only on the genus of S such that*

$$\ell_\rho(\gamma) > \epsilon \implies \sup_{\gamma \subset Y} d_Y(\nu_+, \nu_-) < K.$$

Conversely, if $\sup_Y d_Y(\nu_+, \nu_-) \geq K$ *then*

$$\frac{2\pi i}{\lambda_\rho(\gamma)} \asymp d_\gamma(\nu_+, \nu_-) + i \widetilde{\sum_{\substack{Y \subset S \\ \gamma \subset \partial Y \\ Y \not\sim \gamma}}} d_Y(\nu_+, \nu_-)$$

Let us explain the notation used here. The expression $\widetilde{\sum_{x \in X}} f(x)$ denotes

$$1 + \sum_{\substack{x \in X \\ f(x) \geq K}} f(x)$$

where K is our a-priori "threshold" constant. Our sum then is over all subsurfaces whose boundary contains the isotopy class of γ, except for the annulus homotopic to γ, excluding those where $d_Y(\nu_+, \nu_-)$ is below K. Both sides of the "$\asymp$" symbol are points in the upper half plane of $\mathbf{C}$, and we take "$\asymp$" to mean that the hyperbolic distance between them is bounded by an a-priori constant D_0. Implicit in the statement is that it holds for some D_0 which depends only on the genus of S.

The significance of the hyperbolic distance estimate on $2\pi i/\lambda(\gamma)$ is that we can interpret $2\pi i/\lambda(\gamma)$ as a Teichmüller parameter for the Margulis tube $\mathbf{T}_\gamma$, as follows (cf. [39] and McMullen [31]). Normalize $\rho(\gamma)$ so that it acts on $\hat{\mathbf{C}}$ by $z \mapsto e^\lambda z$. The quotient $(\mathbf{C} \setminus \{0\})/\rho(\gamma)$ is then a torus, and there is a preferred marking of this torus by the pair $(\hat{\gamma}, \mu)$, where μ is the meridian of the torus, or the image of the unit circle in $\mathbf{C}$, and $\hat{\gamma}$ is the image of the curve $\{e^{t\lambda} : t \in [0, 1]\}$. Note, this curve depends on the choice of λ mod $2\pi i$. In [39] we point out that if $\ell_\rho(\gamma)$ is sufficiently short then we can choose $\hat{\gamma}$ to be a minimal representative of γ on the torus just by choosing $\theta = \operatorname{Im} \lambda \in [0, 2\pi)$.

The quantity $2\pi i/\lambda$ turns out to be the point in the upper half-plane representation of the Teichmüller space of the torus which represents the marked quotient torus. Estimating this quantity up to bounded hyperbolic distance is then equivalent to estimating the torus structure up to bounded Teichmüller distance, which corresponds to knowing the action of $\rho(\gamma)$ up to bounded quasi-conformal conjugacy. This finally is equivalent to bounded bilipchitz conjugacy of the action on $\mathbf{H}^3$, and thus is the "right" kind of estimate if we are interested in knowing the quotient geometry up to bilipschitz equivalence.

The imaginary part of the conjectural estimate is supposed to estimate the "height" of the margulis tube boundary for γ, and its real part is supposed to measure the "twist" of the meridian around $\hat{\gamma}$. In our discussion of the Bounded Geometry Theorem, we essentially showed that the height was bounded by the number of elementary moves it took to pass $\mathbf{T}_\gamma$, and the twisting was bounded by the relative twisting of the predecessor and successor curves α and α'. In general we expect that large values of $d_Y(\nu_+, \nu_-)$ with $\gamma \subset \partial Y$ will contribute to parts of the elementary move sequence that make progress along the sides of $\mathbf{T}_\gamma$, and thus give a good estimate for its height.

In [39] we obtained a similar estimate for the case where S is a once-punctured torus. (In this case we are not requiring S to be closed, and our representations must satisfy the added condition that the conjugacy class corresponding to loops around the puncture is mapped to parabolics.) Let us state this just in the case that $\nu_\pm$ are both laminations. For the torus, a lamination are determined by its slope in $H_1(S, \mathbf{R}) = \mathbf{R}^2$, which takes values in $\widehat{\mathbf{R}} = \mathbf{R} \cup \{\infty\}$. Simple closed curves correspond to rational points. For any simple closed curve α we defined a quantity analogous to $d_\alpha(\nu_-, \nu_+)$ as follows: after an appropriate basis change for S (or equivalently action by an element of $\mathrm{SL}_2(\mathbf{Z})$, we may assume that α is represented by ∞, and let $\nu_-(\alpha), \nu_+(\alpha)$ be the irrational numbers representing the ending laminations. Then define

$$w(\alpha) = \nu_+(\alpha) - \nu_-(\alpha).$$

We showed that $\ell_\rho(\alpha)$ can only be short if $w(\alpha)$ is above a uniform threshold, and in this case we estimated

$$\frac{2\pi i}{\lambda_\rho(\alpha)} \asymp w(\alpha) + i.$$

In fact $w(\alpha)$ is just a measure of relative twisting of ν_- and ν_+ around α, and it is not hard to see that $|w(\alpha)|$ is estimated by our $d_\alpha(\nu_-, \nu_+)$, up to a uniform additive error. Thus, this is really the same estimate as in Conjecture 5.1, since there are no essential subsurfaces in S other than annuli.

5.2. **General representations.** All the methods that we have presented here depend heavily on the assumption that ρ is both faithful and discrete. It can be argued, however, that a full understanding of the deformation space of hyperbolic structures on a manifold would require some better geometric description of the whole representation variety, including indiscrete or non-faithful points, and it is tempting to try to enlist the complex of curves for this purpose.

The only results I know that offer any hope are in a paper of Bowditch [8], in which he studies general representations for the once-punctured torus (again with the parabolicity condition for the puncture). Such a representation determines a trace (closely related to complex translation length) for every conjugacy class, and in particular for the simple closed curves, which in this case correspond to $\mathbf{Q} \cup \{\infty\}$, viewed as the vertices of the Farey tesselation of the disk. To every triangle and adjacent pair of triangles is associated a relation among the traces of the vertices, coming from the standard trace identities in $\mathrm{SL}_2(\mathbf{C})$. Bowditch uses these relations alone, without discreteness, to analyze the global properties of the trace function, in particular obtaining a connectedness property for sublevel sets closely analogous to the quasi-convexity property of Lemma 4.4. Using this he is able to define an invariant of the representation that generalizes the ending lamination for discrete representations; but it is hard to know how to extract more information from this invariant.

In the higher genus case, no such analysis has been done, and it would be very interesting to try it. Elementary moves between pants decompositions still give rise to trace identities among the curves involved, although they are a bit more complicated. One wonders at least whether a result like Lemma 4.4 can be generalized to all representations.

Bowditch is led to the following question: Consider the quantity

$$\frac{\ell_\rho(\gamma)}{\ell_{\rho_0}(\gamma)}$$

where ρ_0 is some fixed Fuchsian representation, ρ is a general representation, and γ is a non-trivial element of $\pi_1(S)$. The infimum of this ratio is positive for quasi-Fuchsian representations. For a non-quasi-Fuchsian discrete, faithful representation, the infimum is 0, and can be achieved by considering only γ with *simple* representatives. The limit points of minimizing sequences in the space of laminations give the ending laminations for ρ.

If ρ is indiscrete or non-faithful the infimum is again 0 (indeed $\inf \ell_\rho$ is 0 as well), but the question is, is the infimum also 0 for the simple elements. In other words:

Question 5.2. *Let S be a closed surface of genus at least 2, and let let $\rho : \pi_1(S) \to PSL_2(\mathbf{C})$ be a representation. If*

$$\inf \frac{\ell_\rho(\gamma)}{\ell_{\rho_0}(\gamma)} > 0$$

where γ varies over all simple loops in S, must ρ be quasi-Fuchsian?

This question appears to be difficult, and a positive answer would be a good starting point in using the complex of curves to analyze general representations. To indicate its difficulty, note that it is closely related to the following:

Question 5.3. *If $\rho : \pi_1(S) \to PSL_2(\mathbf{C})$ is any representation with non-trivial kernel, does the kernel contain elements represented by simple loops?*

A positive answer to this question is at least as hard to prove as the simple loop conjecture for hyperbolic 3-manifolds; see Gabai [17] and Hass [21].

References

1. R. Benedetti and C. Petronio, *Lectures on hyperbolic geometry*, Springer-Verlag Universitext, 1992.
2. L. Bers, *On boundaries of Teichmüller spaces and on Kleinian groups I*, Ann. of Math. **91** (1970), 570–600.
3. ______, *Spaces of degenerating Riemann surfaces*, Discontinuous groups and Riemann surfaces, Ann. of Math. Stud. 79, Princeton Univ. Press, 1974, pp. 43–59.
4. Lipman Bers, *An inequality for Riemann surfaces*, Differential geometry and complex analysis, Springer, Berlin, 1985, pp. 87–93.
5. F. Bonahon, *Closed curves on surfaces*, monograph in preparation.
6. ______, *Geodesic laminations on surfaces*, to appear in Proceedings of the Stony Brook 1998 Workshop on Foliations and Laminations.
7. ______, *Bouts des variétés hyperboliques de dimension 3*, Ann. of Math. **124** (1986), 71–158.
8. B. H. Bowditch, *Markoff triples and quasi-Fuchsian groups*, Proc. London Math. Soc. (3) **77** (1998), no. 3, 697–736.
9. R. Brooks and J. P. Matelski, *Collars in Kleinian groups*, Duke Math. J. **49** (1982), no. 1, 163–182.
10. P. Buser, *Geometry and Spectra of Compact Riemann Surfaces*, Birkhäuser, 1992.
11. R. Canary, M. Culler, S. Hersonsky, and P. Shalen, *Density of cusps in boundaries of quasiconformal deformation spaces*, in preparation.
12. R. D. Canary, *Algebraic limits of schottky groups*, Trans. Amer. Math. Soc. **337** (1993), 235–258.
13. 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.
14. A. J. Casson and S. A. Bleiler, *Automorphisms of surfaces after Nielsen and Thurston*, Cambridge University Press, 1988.
15. 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.
16. B. Farb, A. Lubotzky, and Y. Minsky, *Rank one phenomena for mapping class groups*, preprint.
17. David Gabai, *The simple loop conjecture*, J. Differential Geom. **21** (1985), no. 1, 143–149.
18. J. Harer, *Stability of the homology of the mapping class group of an orientable surface*, Ann. of Math. **121** (1985), 215–249.
19. ______, *The virtual cohomological dimension of the mapping class group of an orientable surface*, Invent. Math. **84** (1986), 157–176.
20. 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.
21. Joel Hass, *Minimal surfaces in manifolds with S^1 actions and the simple loop conjecture for Seifert fibered spaces*, Proc. Amer. Math. Soc. **99** (1987), no. 2, 383–388.
22. A. E. Hatcher and W. P. Thurston, *A presentation for the mapping class group*, Topology **19** (1980), 221–237.
23. N. V. Ivanov, *Complexes of curves and the Teichmüller modular group*, Uspekhi Mat. Nauk **42** (1987), 55–107.

24. ———, *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.
25. ———, *Complexes of curves and Teichmüller spaces*, Math. Notes **49** (1991), 479–484.
26. ———, *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.
27. ———, *Automorphisms of complexes of curves and of Teichmüller spaces*, Internat. Math. Res. Notices (1997), no. 14, 651–666.
28. D. Kazhdan and G. Margulis, *A proof of Selberg's conjecture*, Math. USSR Sb. **4** (1968), 147–152.
29. H. A. Masur and Y. Minsky, *Geometry of the complex of curves II: Hierarchical structure*, E-print math.GT/9807150 at http://front.math.ucdavis.edu. To appear in Geom. Funct. Anal.
30. ———, *Geometry of the complex of curves I: Hyperbolicity*, Invent. Math. **138** (1999), 103–149.
31. C. McMullen, *Cusps are dense*, Ann. of Math. **133** (1991), 217–247.
32. R. Meyerhoff, *A lower bound for the volume of hyperbolic* 3*-manifolds.*, Canad. J. Math. **39** (1987), 1038–1056.
33. Y. Minsky, *Bounded geometry in Kleinian groups*, in preparation.
34. ———, *Teichmüller geodesics and ends of hyperbolic 3-manifolds*, Topology **32** (1993), 625–647.
35. ———, *On rigidity, limit sets and end invariants of hyperbolic 3-manifolds*, J. Amer. Math. Soc. **7** (1994), 539–588.
36. ———, *On Thurston's ending lamination conjecture*, Proceedings of Low-Dimensional Topology, May 18-23, 1992, International Press, 1994.
37. ———, *A geometric approach to the complex of curves*, Proceedings of the 37th Taniguchi Symposium on Topology and Teichmüller Spaces (S. Kojima et. al., ed.), World Scientific, 1996, pp. 149–158.
38. ———, *Quasi-projections in Teichmüller space*, J. Reine Angew. Math. **473** (1996), 121–136.
39. ———, *The classification of punctured-torus groups*, Annals of Math. **149** (1999), 559–626.
40. ———, *Kleinian groups and the complex of curves*, Geometry and Topology **4** (2000), 117–148.
41. G. D. Mostow, *Quasiconformal mappings in n-space and the rigidity of hyperbolic space forms*, Publ. I.H.E.S. **34** (1968), 53–104.
42. K. Ohshika, *Ending laminations and boundaries for deformation spaces of Kleinian groups*, J. London Math. Soc. **42** (1990), 111–121.
43. R. Penner and J. Harer, *Combinatorics of train tracks*, Annals of Math. Studies no. 125, Princeton University Press, 1992.
44. 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.
45. ———, *The geometry and topology of 3-manifolds*, Princeton University Lecture Notes, online at http://www.msri.org/publications/books/gt3m, 1982.
46. ———, *Three-Dimensional Geometry and Topology*, Princeton University Press, 1997, (S. Levy, ed.).

SUNY Stony Brook