ISSN 1364-0380

Geometry & Topology
Volume 4 (2000) 117–148
Published: 29 February 2000

Kleinian groups and the complex of curves

Yair N Minsky

Department of Mathematics, SUNY at Stony Brook
Stony Brook, NY 11794, USA

Email: `yair@math.sunysb.edu`

Abstract

We examine the internal geometry of a Kleinian surface group and its relations to the asymptotic geometry of its ends, using the combinatorial structure of the complex of curves on the surface. Our main results give necessary conditions for the Kleinian group to have 'bounded geometry' (lower bounds on injectivity radius) in terms of a sequence of coefficients (subsurface projections) computed using the ending invariants of the group and the complex of curves.

These results are directly analogous to those obtained in the case of punctured-torus surface groups. In that setting the ending invariants are points in the closed unit disk and the coefficients are closely related to classical continued-fraction coefficients. The estimates obtained play an essential role in the solution of Thurston's ending lamination conjecture in that case.

AMS Classification numbers Primary: 30F40

Secondary: 57M50

Keywords: Kleinian group, ending lamination, complex of curves, pleated surface, bounded geometry, injectivity radius

Proposed: David Gabai
Seconded: Jean-Pierre Otal, Walter Neumann

Received: 16 July 1999
Revised: 9 November 1999

1 Introduction

In this paper we examine some new connections between the internal geometry of hyperbolic 3–manifolds and the asymptotic geometry of their ends. Our main theorem is a necessary condition for "bounded geometry" (lower bounds on injectivity radius), in terms of combinatorial properties of the asymptotic geometry, which we think of as analogous to bounds on continued fraction expansions. We conjecture also that this necessary condition is sufficient.

The motivating problem for this analysis is Thurston's Ending Lamination Conjecture, which states that a hyperbolic 3–manifold with finitely-generated fundamental group is determined, up to isometry, by its topological type and a collection of invariants that specify the asymptotic geometry of its ends. Even in the geometrically finite case, where Thurston's conjecture reduces to classical work of Ahlfors–Bers et al, the theory is not geometrically explicit: direct estimates of internal geometry from the end invariants are quite hard to come by.

We will approach this problem by studying the combinatorial/geometric structure of an object known as a *complex of curves on a surface*, a simplicial complex $\mathcal{C}(S)$ whose vertices are homotopy classes of simple closed curves on a surface S (see Section 2 for complete definitions). Adjacency relations in these complexes are related to elementary homotopies between special types of surfaces in hyperbolic 3–manifolds, in a way that gives some control of the internal geometry. Our main methods of proof in this paper depend on some important properties of pleated surfaces in hyperbolic 3–manifolds, originally discovered and applied by Thurston. The results we obtain here should be viewed as the first steps of a project to exploit these connections.

Statements of results

In this paper, a *Kleinian surface group* is a discrete, faithful, type-preserving representation $\rho\colon \pi_1(S) \to \mathrm{PSL}_2(\mathbf{C})$, where $S = S_{g,p}$ is an oriented genus g surface with p punctures, and the type-preserving condition means that the image of a loop around any puncture is parabolic. We exclude $S_{0,p}$ for $p \leq 2$ and $S_{1,0}$, for which there are only elementary representations. Let N_ρ be the quotient $\mathbf{H}^3/\rho(\pi_1(S))$. Such representations arise naturally in the theory when considering hyperbolic 3–manifolds homotopy-equivalent to manifolds with incompressible boundary, and restricting to the boundary groups.

If the group is *quasi-Fuchsian* then its action on the Riemann sphere has two invariant disks whose quotients give two Riemann surface structures on S, or two points in the Teichmüller space $\mathcal{T}(S)$, which we label $\nu_-(\rho)$ and $\nu_+(\rho)$. The work of Ahlfors–Bers [1, 3] shows that this gives a parametrization of all quasi-Fuchsian representations up to isometry by $\mathcal{T}(S)\times\mathcal{T}(S)$, which is in fact a holomorphic isomorphism. And yet, it is hard to answer simple questions such as, given $\nu_\pm$, what is the length of the shortest geodesic (or any given geodesic) in the manifold N_ρ?

For general representations, Thurston and Bonahon [29, 5] showed how to generalize $\nu_\pm$ using *ending laminations* (see Section 5). The same questions are difficult to answer in this case, and in addition Thurston's Ending Lamination Conjecture (the analogue of the Ahlfors–Bers parameterization), which in this setting states that $\nu_\pm$ determine ρ up to isometry, is open in most cases.

Our first theorem gives a necessary condition, in terms of $\nu_\pm(\rho)$, for ρ to have *bounded geometry* in the following sense. Define the *non-cuspidal injectivity radius* $\mathrm{inj}_0(\rho)$ to be half the infimum of the lengths of all closed geodesics in N_ρ. We say ρ has bounded geometry if $\mathrm{inj}_0(\rho)>0$.

For any essential subsurface $Y\subset S$ we will define a "projection" π_Y that takes the invariants $\nu_\pm$ to elements of the complex of curves of Y, and we denote the distance measured in this complex by $d_Y(\nu_-,\nu_+)$.

Theorem A *For any Kleinian surface group ρ with ending invariants $\nu_\pm$, if*

$$\sup_{Y\subset S} d_Y(\nu_+,\nu_-)=\infty$$

then

$$\mathrm{inj}_0(\rho)=0,$$

where the supremum is over proper essential subsurfaces Y in S not all of whose boundaries are mapped to parabolics.

See below for some examples of surface groups to which this theorem applies, and some examples supporting the converse implication.

Theorem A is a consequence of Theorem B, in which we consider the set of geodesics in N_ρ satisfying a certain length bound. If α is a finite set of homotopy classes of simple closed curves in S, let $\ell_\rho(\alpha)$ denote the total length of their geodesic representatives in N_ρ. For $L>0$, let $\mathcal{C}_0(\rho,L)$ denote the set of homotopy classes α of simple closed curves in S with $\ell_\rho(\alpha)\le L$. An understanding of this set is crucial to characterizing the internal geometry of N_ρ. The projection π_Y is defined on this set and the diameter of the image is denoted $\mathrm{diam}_Y(\mathcal{C}_0(\rho,L))$. We will prove:

Theorem B *Given a surface S, $\epsilon > 0$ and $L > 0$, there exists K so that, if $\rho\colon \pi_1(S) \to PSL_2(\mathbf{C})$ is a Kleinian surface group and Y is a proper essential subsurface of S, then*

$$\mathrm{diam}_Y(\mathcal{C}_0(\rho, L)) \geq K \quad \Longrightarrow \quad \ell_\rho(\partial Y) \leq \epsilon.$$

The relation between the invariants $\nu_\pm(\rho)$ and the set of bounded curves $\mathcal{C}_0(\rho, L)$ is, for example in the case where $\nu_\pm$ are laminations rather than points in the Teichmüller space, that the elements of $\mathcal{C}_0(\rho, L)$ for suitable L will accumulate exactly on $\nu_\pm$ in the space of laminations on S. This gives the connection between Theorem B and Theorem A.

Examples

Let us sketch some examples of surface groups where the hypothesis of Theorem A holds. The basic idea is similar to that in Bonahon–Otal [6] and Thurston [27], where surface groups $\rho\colon \pi_1(S) \to \mathrm{PSL}_2(\mathbf{C})$ are constructed with arbitrarily short geodesics. In those examples the projections $\pi_Y(\nu_-, \nu_+)$ which are large correspond to annuli Y whose cores are the short curves. We leave out many details, since they are not needed for the proofs of the main theorems, but the reader is referred to [31, 12, 11] for details about laminations and pseudo-Anosov maps, and [20] for details about quasi-Fuchsian groups and Bers slices.

Let $Y_1, Y_2, \ldots$ be a sequence of essential subsurfaces of S and let $\Psi_i\colon S \to S$ be homeomorphisms so that $\Psi_i = id$ outside Y_i, and $\Psi_i|_{Y_i}$ are pseudo-Anosov. For simplicity it is probably a good idea to take the Y_i and Ψ_i from a finite set, eg, $Y_i = Y_{i+2}$ and $\Psi_i = \Psi_{i+2}$ for each i. Assume also that ∂Y_i and ∂Y_{i+1} intersect each other essentially.

Let $\{m_i\}$ be an increasing sequence of positive integers. Define

$$h_j^k = \Psi_j^{m_j} \circ \cdots \circ \Psi_{k-1}^{m_{k-1}}$$

for integers $0 < j < k$, and extend to all pairs of integers $j, k > 0$ by requiring $h_k^j = (h_j^k)^{-1}$. Note that $h_j^k \circ h_k^m = h_j^m$. Now fix a conformal structure σ on S and consider the sequence of quasi-Fuchsian groups $\rho_k = QF(\sigma, h_1^k(\sigma))$ as $k \to \infty$. (The notation $QF(\sigma, \tau)$ refers to a quasi-Fuchsian group whose Riemann surface structures at infinity are $\nu_- = \sigma$, $\nu_+ = \tau$. This is determined up to conjugacy in $\mathrm{PSL}_2(\mathbf{C})$.)

Let $Z_j = h_1^j(Y_j)$. We claim that, with an *a priori* choice of sufficiently large numbers m_j, the coefficients $c_j(k) = d_{Z_j}(\sigma, h_1^k(\sigma))$ can be made as large as we please (where $j < k$). We refer the reader to sections 2 and 5 where these coefficients are defined, but note that in the case of $d_Z(\sigma,\tau)$ where σ and τ are conformal structures, and Z is not an annulus, we simply consider minimal-length curves α in σ and β in τ which intersect Z, and measure the "elementary move" distance in the complex $\mathcal{C}'(Z)$ between the arc systems $\alpha\cap Z$ and $\beta\cap Z$.

Applying h^1_j to Z_j, σ and $h^k_1(\sigma)$, we obtain $c_j(k) = d_{Y_j}(h^1_j(\sigma), h^k_j(\sigma))$. Let λ^+_i and λ^-_i be the stable and unstable laminations, respectively, of Ψ_i. Then choosing m_{j-1} sufficiently large we can guarantee that $h^1_j(\sigma)$ is as close as we like (in Thurston's compactification of Teichmüller space, say) to λ^-_{j-1}, and choosing m_j sufficiently large $h^k_j(\sigma)$ is as close as we like to λ^+_j. Note that our assumption that Y_j and Ψ_j vary in a finite set makes it easier to verify this in a uniform way. Now since λ^-_{j-1} intersects Y_j in some finite number of arcs up to homotopy, whereas λ^+_j is a lamination that fills Y_j, it is not hard to see (for example see [19]) that the combinatorial distance $d_{Y_j}(h^1_j(\sigma), h^k_j(\sigma))$ can be made as large as we like by increasing m_j — independently of k as long as $k > j$.

Let $c_j = \lim_{k\to\infty} c_j(k)$. We can restrict to a subsequence so that these limits exist, and by choosing $m_j \to \infty$ as above, we guarantee that $c_j \to \infty$. Now considering the limit (in the Bers slice) of ρ_k, we obtain a limiting representation ρ with invariants $\nu_- = \sigma$, $\nu_+ = \lim_{k\to\infty} h^k_1(\sigma)$, and it is not hard to see that $d_{Z_j}(\nu_-,\nu_+)$ are estimated (to within a bounded distance) by c_j. Theorem A implies that there are arbitrarily short geodesics in this manifold. In fact Theorem B tells us that the boundary curves ∂Z_j have lengths going to 0 as $j \to \infty$. The estimates of Theorem B hold for all the groups ρ_k in the sequence.

In the opposite direction (supporting the conjectural converse of Theorem A), consider the representations $\rho\colon \pi_1(S) \to \mathrm{PSL}_2(\mathbf{C})$ that arise as fibres of hyperbolic manifolds that fibre over the circle (Thurston [27]). In this case the end-invariants are the stable and unstable laminations $\lambda^\pm$ of the pseudo-Anosov monodromy Ψ of the fibration. Such a manifold certainly has bounded geometry since it covers a compact manifold, and we can show that the distances $d_Y(\nu_-,\nu_+)$ are bounded over all subsurfaces Y: Let $\mathcal{PML}(S)$ denote the projective measured lamination space [26], and let K be any compact subset of $\mathcal{PML}(S)\setminus\{\nu_-,\nu_+\}$. Fix a hyperbolic metric on S, and realize all laminations geodesically in this metric. For any $\lambda\in K$, $(S-\lambda)\cap\nu_+$ is a collection of arcs with length bounded by some finite $L(\lambda)$ (since ν_+ fills S). Furthermore L is bounded on K: if not, then by compactness of K there is some convergent sequence $\lambda_i \to \lambda \in K$ for which $L(\lambda_i) \to \infty$, and hence there is a sequence of arcs α_i of ν_+ of length going to infinity with λ_i disjoint from α_i. Since ν_+ is minimal these arcs converge to ν_+ and we conclude that ν_+ and λ have no transverse intersection, and hence are the same in $\mathcal{PML}(S)$ since ν_+ fills and is uniquely ergodic. This contradicts the assumption that K misses ν_+. The same argument can be made for ν_-. Thus, for any subsurface Y whose boundary ∂Y, viewed as a lamination, determines an element of K, there is a uniform upper bound L_K on the lengths of components of $\nu_+\cap Y$ and $\nu_-\cap Y$. This gives, as in Section 4, an upper bound D_K on $d_Y(\nu_-,\nu_+)$ over all Y with ∂Y in K. Now the action of the pseudo-Anosov Ψ on $\mathcal{PML}(S)-\{\nu_-,\nu_+\}$ has some compact fundamental domain K, and since $\nu_\pm$ are invariant under Ψ, the bound $d_Y(\nu_-,\nu_+)\le D_K$ holds for ∂Y in each translate $\Psi^n(K)$. Thus the bound holds for all Y.

A more complete class of examples is obtained when S is a once-punctured torus. In this case, the set $\mathcal{C}_0(S)$ of simple closed curves in S can be identified with $\hat{\mathbf{Q}} = \mathbf{Q}\cup\{\infty\}$ by taking each curve to its slope in the homology $H_1(S) = \mathbf{Z}^2$ after fixing a basis. The natural completion of $\hat{\mathbf{Q}}$ is the circle $\hat{\mathbf{R}} = \mathbf{R}\cup\{\infty\}$. The Teichmüller space $\mathcal{T}(S)$ can be identified with the hyperbolic plane $\mathbf{H}^2$, and the natural compactification of $\mathbf{H}^2$ by $\hat{\mathbf{R}}$ is exactly Thurston's compactification of $\mathcal{T}(S)$ by the space of projective measured laminations on S. Thus ν_+ and ν_- are points in the closed 2–disk $\mathbf{H}^2\cup\hat{\mathbf{R}}$. The essential subsurfaces Y in S are all annuli, so they can again be parametrized by $\hat{\mathbf{Q}}$, by associating to each Y the slope of its core curve. For simplicity, consider the special case where $\nu_- = \infty$ and $\nu_+\in\mathbf{R}\setminus\mathbf{Q}$. Then in fact $d_Y(\nu_-,\nu_+)$ is uniformly bounded except when Y's slope is a continued-fraction approximant of ν_+, and in that case it is equal (up to bounded error) to the corresponding coefficient $w(Y)$ in the continued-fraction expansion of ν_+.

Theorem B implies that $\ell_\rho(\partial Y)$ is small if $w(Y)$ is large, and in fact there is an explicit estimate for the *complex* translation length $\lambda_\rho = \ell_\rho + i\theta_\rho$ (where θ_ρ gives the rotation of the associated isometry), of the form $\lambda_\rho(\partial Y)\asymp(1/w(Y)^2 + 2\pi i/w(Y))$. See [24] for more details.

Outline of the paper

Section 2 outlines the basic definitions and properties of the complexes of curves $\mathcal{C}(S)$. Section 3 discusses some properties of pleated surfaces in hyperbolic 3–manifolds, concentrating on Thurston's Uniform Injectivity Theorem and some consequences, old and new. A number of the proofs in these two sections are sketched because they are reasonably well known, although perhaps hard to find in print.

In Section 4 we give the proof of Theorem B. This is the main construction of the paper, and the main tools for the proof are our extension, Theorem 3.5, of Thurston's theorem on Efficiency of pleated surfaces, and Lemma 3.4 on comparing pleated surfaces that share part of their pleating locus. The proof of Theorem A appears next in Section 5, after a short discussion of the end invariants $\nu_\pm$.

2 Complexes of curves and subsurface projections

We recall here the definitions of Harvey's complexes of curves [13] and related complexes from [18].

Let $S = S_{g,p}$ be an orientable surface of genus g with p punctures. We will always assume that S is not a sphere with 0,1 or 3 holes, nor a torus with 0 holes. The case of the once-punctured torus $S_{1,1}$ and quadruply-punctured sphere, $S_{0,4}$, are special and will be treated below, as will the case of the annulus $S_{0,2}$, which will only occur as a subsurface of larger surfaces. We call all other cases "generic".

If S is generic, we define $\mathcal{C}(S)$ as the following simplicial complex: vertices of $\mathcal{C}(S)$ are non-trivial, non-peripheral homotopy classes of simple curves, and k–simplices are sets $\{v_0, \ldots, v_k\}$ of distinct vertices with disjoint representatives. For $k \geq 0$ let $\mathcal{C}_k(S)$ denote the k–skeleton of $\mathcal{C}(S)$.

We define a metric on $\mathcal{C}(S)$ by making each simplex regular Euclidean with sidelength 1, and taking the shortest-path distance. Distance in $\mathcal{C}(S)$ will be denoted $d_{\mathcal{C}(S)}$ or d_S. Note that $\mathcal{C}_k(S)$ is quasi-isometric to $\mathcal{C}_1(S)$ (with the induced path metrics) for $k \geq 1$. These conventions also apply to the non-generic cases below.

Once-punctured tori and 4–punctured spheres

If S is $S_{0,4}$ or $S_{1,1}$, we define the vertices $\mathcal{C}_0(S)$ as before, but let edges denote pairs $\{v_0, v_1\}$ which have the minimal possible geometric intersection number (2 for $S_{0,4}$ and 1 for $S_{1,1}$).

We remark that in both these cases $\mathcal{C}(S)$ is isomorphic to the classical Farey graph in the plane (see eg Hatcher–Thurston [16] or [23]).

Arc complexes

If Y is a non-annular surface with punctures, let us also define the larger *arc complex* $\mathcal{C}'(Y)$ whose vertices are either properly embedded arcs in Y, up to homotopy keeping the endpoints in the punctures, or essential closed curves up to homotopy. Simplices correspond to sets of vertices with representatives that have disjoint interiors.

Subsurface projections

By an *essential subsurface* of S we shall always mean an open subset Y which is incompressible (π_1 injects), non-peripheral (not homotopic into a puncture) and not homeomorphic to $S_{0,3}$. We usually consider isotopic subsurfaces to be equivalent.

Let Y be a non-annular essential subsurface of S. We can define a "projection"

$$\pi_Y\colon \mathcal{C}'_0(S) \to \mathcal{C}'(Y) \cup \{\emptyset\}$$

as follows (where in the definition of $\mathcal{C}'(Y)$ we consider each end of Y to be a puncture):

If $\alpha \in \mathcal{C}'_0(S)$ has no essential intersections with Y (including the case that α is homotopic to ∂Y) then define $\pi_Y(\alpha) = \emptyset$. Otherwise, α intersects Y in a collection of disjoint embedded arcs. Keeping only the essential ones (for example by taking geodesic representatives of α and ∂Y in some hyperbolic metric), and taking their homotopy classes modulo the ends of Y, we obtain a simplex in $\mathcal{C}'(Y)$. Let $\pi_Y(\alpha)$ be the barycenter of this simplex.

For convenience we also extend the definition of π_Y to $\mathcal{C}'_0(Y)$, where it is the identity map.

We let $d_Y(\alpha,\beta)$ denote $d_{\mathcal{C}'(Y)}(\pi_Y(\alpha),\pi_Y(\beta))$, when these projections are non-empty. Similarly $\operatorname{diam}_Y(A)$ denotes $\operatorname{diam}_{\mathcal{C}'(Y)}(\cup_{a\in A}\pi_Y(a))$, where $A \subset \mathcal{C}'_0(S)$.

Annuli

Now let Y be an essential annulus in S. We will define $\mathcal{C}'(Y)$ a little differently: let $\widetilde{Y}$ denote the annular cover of S to which Y lifts homeomorphically. We can compactify $\widetilde{Y}$ to a closed annulus $\widehat{Y}$ in a natural way: the universal cover of S can be identified with $\mathbf{H}^2$, which has a natural compactification as the closed disk, and the covering $\mathbf{H}^2 \to \widetilde{Y}$ has deck group $\mathbf{Z}$ which acts with two fixed points on the boundary. The quotient of the closed disk minus these two points is the closed annulus $\widehat{Y}$. Now define $\mathcal{C}'_0(Y)$ to be the set of all homotopy classes *rel endpoints* of arcs connecting the two boundaries of $\widehat{Y}$. Define $\mathcal{C}'(Y)$ by putting a simplex between any finite set of arcs that have representatives with pairwise disjoint interiors.

We can make the simplices of $\mathcal{C}'(Y)$ regular Euclidean as before. Although now $\mathcal{C}'(Y)$ is infinite-dimensional it is in fact quasi-isometric to its 1–skeleton $\mathcal{C}'_1(Y)$, which in turn is quasi-isometric to the integers $\mathbf{Z}$ with the usual distance (see [18]). In particular it is easy to see that distance in $\mathcal{C}'_1(Y)$ is determined by intersection number: If α and β are two distinct vertices in $\mathcal{C}'_0(Y)$ with geometric intersection number $i(\alpha,\beta)$ in $\widehat{Y}$ then

$$d_Y(\alpha,\beta) = 1 + i(\alpha,\beta). \tag{2.1}$$

We can now define π_Y in this case as follows: For $\alpha \in \mathcal{C}'_0(S)$, *lift* a representative of α to the annular cover $\widetilde{Y}$. Each component of the lift extends naturally to an arc in the closed annulus $\widehat{Y}$, and a finite number of these connect the two boundary components, and have disjoint interiors. Thus they determine a simplex of $\mathcal{C}'(Y)$ and again we let $\pi_Y(\alpha)$ be the barycenter of this simplex.

Properties of $\mathcal{C}(S)$

We note first that $\mathcal{C}(S)$ is connected and has infinite diameter in all cases we consider, and is a δ–hyperbolic metric space for some $\delta(S) > 0$ — see [19].

We will also need an elementary lemma relating $\mathcal{C}$–distance to intersection number, analogous to (2.1):

Lemma 2.1 *Given a surface Y and $D > 0$ there exists D' such that, if α,β are vertices of $\mathcal{C}'(Y)$ and $i(\alpha,\beta) \le D$ then*

$$d_Y(\alpha,\beta) \le D'.$$

In fact a more precise bound can be given — see eg [19] and Hempel [17]. For us it will suffice to note that *some* bound exists. The proof is an easy surgery argument: One can inductively replace β with β' which is disjoint from β and intersects α fewer times.

3 Pleated surfaces

Thurston introduced pleated surfaces in [29, 30] as a powerful tool for studying hyperbolic 3–manifolds. In this section we recall the basic definitions of pleated surfaces and related notions (for further details see Canary–Epstein–Green [10]), and prove a collection of results, some of them technical and all of them directed towards some extensions and applications of Thurston's theorems on Uniform Injectivity and Efficiency of pleated surfaces.

Definitions

See Thurston [29], Canary–Epstein–Green [10] and Penner–Harer [26] for definitions and examples of geodesic laminations and measured geodesic laminations on hyperbolic surfaces. We note here that a geodesic lamination with respect to one hyperbolic metric on a surface can be isotoped to a geodesic lamination with regard to any other hyperbolic metric — see Hatcher [15]. If we use the term "geodesic lamination" in the absence of a specific metric we will mean this isotopy class.

A pleated surface is a map $f\colon S \to N$, where S is a surface of finite type and N is a hyperbolic 3–manifold, together with a complete, finite-area hyperbolic metric σ on S, satisfying the following two conditions.

First, f is *path-isometric* with respect to σ. That is, any σ–rectifiable path in S is mapped by f to a path of N–length equal to its σ–length. With this in mind we call σ the metric *induced by* f, and note that in fact it is uniquely determined by f.

Second, there is a σ–geodesic lamination λ on S whose leaves are mapped to geodesics by f. In the complement of λ, f is totally geodesic.

Double incompressibility A map $f\colon S \to N$ is *incompressible* if $\pi_1(f)$ is injective. Following Thurston [30], we say that a map $f\colon S \to N$ is *doubly incompressible* if in addition the following conditions hold:

(1) *Arcs modulo cusps are mapped injetively:*
Homotopy classes of maps $(I, \partial I) \to (S, cusps(S))$ map injectively to homotopy classes of maps $(I, \partial I) \to (N, cusps(N))$.

(2) *No essential annuli except at parabolics:*
For any cylinder $c\colon S^1 \times I \to N$ whose boundary $\partial c\colon S^1 \times \{0,1\} \to N$ factors as $f \circ c_0$ where $c_0\colon S^1 \times \{0,1\} \to S$, if $\pi_1(c)$ is injective then either the image of $\pi_1(c_0)$ consists of parabolic elements of $\pi_1(S)$, or c_0 extends to a map of $S^1 \times I$ into S.

(3) *Primitive elements are preserved and no rank–2 cusps in image:*
Each maximal abelian subgroup of $\pi_1(S)$ is mapped to a maximal abelian subgroup of $\pi_1(N)$.

Here $I = [0,1]$ and $cusps(S)$, $cusps(N)$ are the union of ϵ_0–cusp neighborhoods of the parabolic cusps of S and N respectively.

In particular, a map which induces an isomorphism on π_1 and sends cusps to cusps is doubly incompressible.

Noded pleated surfaces

Let S' be an essential subsurface of S whose complement is a disjoint union of simple curves. Given a homotopy class $[f]$ of maps from S to N, we say that $g\colon S' \to N$ is a *noded pleated surface* in the class $[f]$ if g is pleated with respect to a hyperbolic metric on S' (in which the ends are cusps), and g is homotopic to the restriction to S' of an element of $[f]$.

We note that if $f\colon S \to N$ is a doubly-incompressible map and $g\colon S' \to N$ is a noded pleated surface in the homotopy class of f, then g is also doubly incompressible.

Finite-leaved laminations

In this paper, pleated surfaces and noded pleated surfaces will only arise with *finite-leaved* laminations, which always consist of a finite number (possibly zero) of closed geodesics and a finite number of infinite leaves whose ends either spiral around the closed geodesics or exit a cusp. Thurston observed that, if $f\colon S \to N$ is a doubly-incompressible map and λ is a finite leaved lamination all of whose closed curves are taken by f to non-parabolic loops, then there is a pleated surface homotopic to f mapping λ geodesically (λ is "realizable").

Similarly if some subset C of the closed loops are taken by f to parabolic loops and $S' = S \setminus C$ then there is a *noded* pleated surface $g\colon S' \to N$ in the homotopy class of f. To simplify notation we will still refer to this noded pleated surface as a "pleated surface mapping λ geodesically" (see in particular Lemma 3.5 (Efficiency of pleated surfaces)).

Basic properties

We refer the reader to Thurston [32], Canary–Epstein–Green [10] or Benedetti–Petronio [2] for a discussion of the Margulis lemma and the thick–thin decomposition. In what follows ϵ_0 will always denote a constant no greater than the Margulis constant for $\mathbf{H}^3$. We may also assume ϵ_0 has the property that the intersection of any simple closed geodesic in a hyperbolic surface with an ϵ_0–Margulis tube, is either the core of the tube or a union of arcs which connect the two boundaries (see eg [10]).

It is obvious that a π_1–injective pleated surface $g\colon S \to N$ takes the ϵ_0–thin part of S to the ϵ_0–thin part of N. In the other direction we have this observation (see Thurston [30]):

Lemma 3.1 *Given ϵ_0 there exists ϵ_1 (depending only on ϵ_0 and the topological type of S) such that, if $g\colon S \to N$ is a π_1–injective pleated surface then only the ϵ_0–thin part of the surface can be mapped into the ϵ_1–thin part of N.*

Thurston's uniform injectivity theorem

One of the most important properties of pleated surfaces is that, under appropriate topological assumptions, there is some control over the ways in which they can fold in the target manifold. Let $g\colon S \to N$ be a pleated surface mapping a lamination λ geodesically, and let $\mathbf{p}_g\colon \lambda \to \mathbf{P}(N)$ be the natural lift of $g|_\lambda$ to the projectivized tangent bundle $\mathbf{P}(N)$ of N, taking $x \in \lambda$ to $(g(x), g_*(l))$ where l is the tangent line to the leaf of λ through x. In [27], Thurston established the following theorem:

Theorem 3.2 (Uniform Injectivity Theorem) *Fix a surface S of finite type, and a constant $\bar\epsilon > 0$. Given $\epsilon > 0$ there exists $\delta > 0$ such that, given any type-preserving doubly incompressible pleated surface $g\colon S \to N$, with pleating locus λ and induced metric σ, if $x, y \in \lambda$ are in the $\bar\epsilon$–thick part of (S, σ) then*

$$d_{\mathbf{P}(N)}(\mathbf{p}_g(x), \mathbf{p}_g(y)) \le \delta \quad \Longrightarrow \quad d_\sigma(x, y) \le \epsilon.$$

For more general versions see also Thurston [28], and Canary [8].)

As a consequence of this theorem, Lemma 3.4 below allows us to compare pleated surfaces which have a subset of their pleating locus in common. We first need to address the following somewhat technical point, arising from the fact that, in the absence of a single fixed hyperbolic metric, laminations on a surface are usually considered only up to isotopy.

Let us say that two pleated surfaces $f, g\colon S \to N$ are *homotopic relative to a common pleating lamination* μ if μ is a lamination on S which is mapped geodesically by f, if $f|_\mu = g|_\mu$, and finally if f and g are homotopic by a family of maps that fixes μ pointwise. The next lemma tells us that in our setting, if two pleated surfaces share a pleating sublamination only up to isotopy, then after re-parameterizing the domain the stronger pointwise condition will hold:

Lemma 3.3 *Let $f, g\colon S \to N$ be homotopic pleated surfaces that are injective on π_1. Suppose that μ is a sub-lamination of the pleating locus of f and μ' is a sub-lamination of the pleating locus of g, such that μ and μ' are isotopic to each other. Then there is a homeomorphism $h\colon S \to S$, isotopic to the identity, such that $h(\mu) = \mu'$, and f and $g \circ h$ are homotopic relative to the common pleating lamination μ.*

Proof By assumption there is some homeomorphism $k\colon S \to S$ isotopic to the identity such that $k(\mu) = \mu'$. Thus possibly replacing g by $g \circ k$ we may assume $\mu = \mu'$.

Next, after possibly lifting to a suitable cover of N we may assume that f and g are isomorphisms on π_1. The image of μ by f is a geodesic lamination μ^* in N — this is also the image of μ by g since they are homotopic, as below.

Let $H\colon S \times I \to N$ be the homotopy from f to g, where $I = [0,1]$ and $H(x,0) = f(x)$, $H(x,1) = g(x)$. Now, each leaf l of μ is taken by f to a leaf of μ^* in N, and similarly by g. If we lift the homotopy H to a map of the universal covers $\widetilde{H}\colon \widetilde{S} \times I \to \widetilde{N}$, and let $\widetilde{l}$ be a lift of l, then we find that the paths $\widetilde{H}(\widetilde{l} \times \{t\})$ for $t \in I$ are all homotopic with bounded trajectories, and hence the geodesics $\widetilde{H}(\widetilde{l} \times \{0\})$ and $\widetilde{H}(\widetilde{l} \times \{1\})$ are the same. Projected back down to the manifolds, this means that f and g take l to the *same* leaf l^*. Furthermore it means that each trajectory $H(\{x\} \times I)$, for $x \in \mu$, is homotopic rel endpoints into the image l^* of the leaf l containing x.

Now define a new homotopy $\hat{H}\colon S \times I \to N$, by letting $\hat{H}|_{\{x\}\times I}$, for each $x \in S$, be the unique constant speed geodesic homotopic to $H_{\{x\}\times I}$ with endpoints fixed. This is clearly a continuous map (here we are using the negative curvature of the target), and by the previous paragraph, when x is in a leaf l of μ, $\hat{H}(\{x\} \times I)$ lies in l^*.

For each $t \in I$, let $g_t(\cdot) = \hat{H}(\cdot, t)$, so that $g_1 = g$ and $g_0 = f$. Then g_t maps μ to μ^*, and in fact the lift $\mathbf{p}_{g_t}\colon \mu \to \mathbf{P}(N)$ takes μ homeomorphically to the lift $\mathbf{P}\mu^*$ in the projective tangent bundle. Let $h_t\colon \mu \to \mu$ be the homeomorphism $\mathbf{p}_{g_t}^{-1} \circ \mathbf{p}_f$. Note that h_0 is the identity, and that h_t takes each leaf of μ to itself. Furthermore h_t has the property that $g_t(h_t(x)) = f(x)$ for all $x \in \mu$.

It remains to extend h_t to a homeomorphism of S, and then $g \circ h_1$ is the desired reparameterization of g, with the desired homotopy being $G(x,t) = g_t(h_t(x))$.

Each complementary region R of μ is an open hyperbolic surface with finite area. Taken with the induced path metric, it has a metric completion $\overline{R}$, which is a hyperbolic surface with geodesic boundary possibly with “spikes” — that is, each end of R is completed either by a closed geodesic or by a chain of geodesic lines with each two successive lines asymptotic along rays. The region between two such rays where they are less than ϵ apart (for some small $\epsilon > 0$) is called an ϵ–spike. All these lines naturally immerse in S as leaves of μ. Each h_t induces a homeomorphism of the boundary of $\overline{R}$ which is uniformly continuous (since h_t is continuous on a compact set), and is easy to extend continuously to an embedding of each ϵ–spike into $\overline{R}$ for suitable ϵ, which remains uniformly continuous and continuous in t (for example use horocyclic coordinates in each spike and extend the map by a linear interpolation of the boundary values). $\overline{R}$ minus its spikes is a compact surface and the standard methods apply to extend h_t to a homeomorphism of $\overline{R}$, continuously varying with t. These extensions piece together continuously when all the components of $\overline{R}$ are immersed in S, because of the uniform continuity of each. $\square$

Applications of uniform injectivity

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, we let $\ell_\sigma([\tau])$ denote the σ–length of the shortest arc homotopic to τ, with endpoints fixed.

Lemma 3.4 (Short bridge arcs) *Fix a surface S of finite type. Fix a positive $\epsilon_1 \leq \epsilon_0$. Given $\delta_1 > 0$ there exists $\delta_0 \in (0, \delta_1)$ such that the following holds.*

Let $g_0\colon S \to N$ and $g_1\colon S \to N$ be type-preserving doubly incompressible pleated surfaces in a hyperbolic 3–manifold N, which are homotopic relative to a common pleating lamination μ. Let σ_0, σ_1 be the induced metrics on S. If τ is a bridge arc for μ, and either

(1) *τ is in the ϵ_1–thick part of (S, σ_0), or*

(2) *τ is a primitive bridge arc,*

then

$$\ell_{\sigma_0}([\tau]) \leq \delta_0 \implies \ell_{\sigma_1}([\tau]) \leq \delta_1.$$

This lemma is a direct consequence of Lemma 2.3 in [21], and the discussion of Uniform Injectivity there. Let us sketch the proof.

The bound $\ell_{\sigma_0}([\tau]) \leq \delta_0$ implies a proportional bound $d_{\mathbf{P}(N)}(\mathbf{p}_{g_0}(x), \mathbf{p}_{g_0}(y)) \leq C\delta_0$, where x and y are the endpoints of τ and C is a universal constant. This is because, for a uniform c, there are segments of leaves of μ centered on x and y of radius at least $\log 1/\delta_0 - c$, with lifts to the universal cover of (S,σ_0) that are bounded Hausdorff distance ϵ_0 — these map to segments in N with similar bounds since g_0 is Lipschitz, and hyperbolic trigonometry implies the bound on $d_{\mathbf{P}(N)}$.

Note also that $d_{\mathbf{P}(N)}(\mathbf{p}_{g_0}(x),\mathbf{p}_{g_0}(y)) = d_{\mathbf{P}(N)}(\mathbf{p}_{g_1}(x),\mathbf{p}_{g_1}(y))$, since $g_0|_\mu = g_1|_\mu$ by definition.

Now consider case (1) first. Because τ is contained in the ϵ_1–thick part of (S,σ_0), by Lemma 3.1 there is some $\epsilon_2(\epsilon_1)$ so that $g_0(\tau)$ lies in the ϵ_2–thick part of N. Choose $\epsilon > 0$ so that $\epsilon + \delta_0 < \epsilon_2$. Now x,y must be in the ϵ_2–thick part of (S,σ_1), so the Uniform Injectivity Theorem applied to g_1, with $\bar\epsilon = \epsilon_2$, gives a choice of δ_0 that guarantees $d_{\sigma_1}(x,y) \le \epsilon$. Let τ' be an arc in S joining x and y of length ϵ. Then the concatenation $g_0|_\tau * g_1|_{\tau'}$ has length bounded by $\delta_0 + \epsilon < \epsilon_2$ and is hence null-homotopic. By assumption, g_0 and g_1 are homotopic on τ' rel endpoints and we conclude that τ and τ' are homotopic rel endpoints. Thus the bound on τ' yields the desired bound on $\ell_{\sigma_1}([\tau])$.

In case (2), although τ can be in the thin part, the extra topological restriction that τ is a primitive bridge arc is used in Lemma 2.3 of [21] to obtain the desired bound. Essentially, this condition prevents x and y from being on opposite sides of the core of a Margulis tube that is very badly folded in the 3–manifold — this is the basic example of failure of Uniform Injectivity. See also Brock [7] for a stronger version of this result.

A further consequence of Uniform Injectivity allows us to estimate the lengths of curves in a hyperbolic 3–manifold based on their representatives in a certain kind of pleated surface. Let λ be a finite-leaved geodesic lamination in a hyperbolic surface S which is maximal, in the sense that its complementary regions are ideal triangles (when S has cusps there will be leaves that enter the cusps). If c is any closed geodesic in S we define its *alternation number* with λ, $a(\lambda,c)$, following Thurston [27]:

If c is a leaf of λ, then $a(\lambda,c)=0$. If not then c meets λ in a countable set of transverse intersection points. If x and y are two points of $\lambda\cap c$ bounding an interval on c with no intersection points in its interior, then x and y are on leaves of λ that are legs of an ideal triangle, and hence are asymptotic, either on one side of c or the other. If x,y and z are successive intersection points with no intervening intersections then we call y a *boundary intersection* if the leaves through x and y are asymptotic on the opposite side of c from the leaves through y and z. If $y \in \lambda\cap c$ is an accumulation point of leaves of λ, then since λ is finite-leaved the leaf through y is closed, and there are accumulations

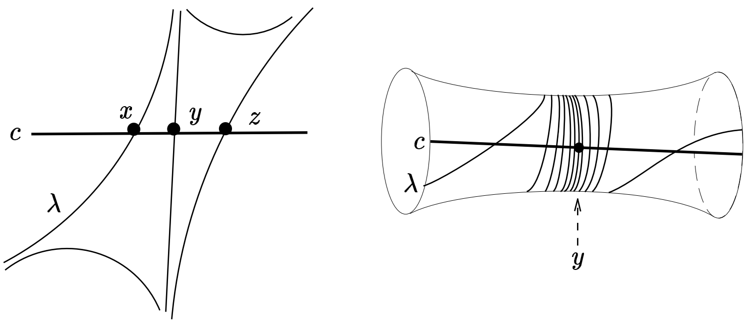

Figure 1: The two kinds of boundary intersections used in defining $a(\lambda, c)$

from each side, by spiraling leaves. Again we call y a boundary intersection if the spiraling is in opposite directions on the two sides (see Figure 1).

Each boundary intersection is isolated, so there are finitely many. We let $a(\lambda, c)$ denote the number of boundary intersections.

The following theorem will be used in the proof of Theorem B. The first part (3.1) was proved by Thurston in [27, Theorem 3.3]. A generalization was proved by Canary in [8, Proposition 5.4], to which we refer the reader for a detailed discussion. The second conclusion (3.2) in our statement comes from a straightforward extension of Thurston's argument, which we sketch below.

Theorem 3.5 (Efficiency of pleated surfaces) *Given S and any $\epsilon > 0$, there is a constant $C > 0$ such that, if $g\colon S \to N$ is a type-preserving doubly incompressble pleated surface with induced metric σ mapping geodesically a maximal finite-leaved lamination λ, and each closed leaf of λ has image length at least ϵ, then for any measured geodesic lamination γ in S,*

$$\ell_N(g(\gamma)^*) \leq \ell_\sigma(\gamma) \leq \ell_N(g(\gamma)^*) + Ca(\lambda, \gamma). \tag{3.1}$$

If we remove the length condition on the closed leaves of λ and let R_λ be the complement in S of the ϵ_0–thin parts (in the metric σ) whose cores are closed leaves of λ, we have

$$\ell_\sigma(\gamma \cap R_\lambda) \leq \ell_N(g(\gamma)^*) + Ca(\lambda, \gamma). \tag{3.2}$$

This estimate applies also if some closed leaves of λ have zero length, in which case g is a noded pleated surface defined in the complement of those leaves.

Here $g(\gamma)^*$ denotes the geodesic representative of $g(\gamma)$ in N, if it exists. For a measured lamination γ this means the (unique) image of γ by a pleated surface homotopic to g that maps γ geodesically. The length $\ell_N(g(\gamma)^*)$ is then well-defined, and if no geodesic representative exists (γ is mapped to a parabolic element, or is an ending lamination) then we define $\ell_N(g(\gamma)^*) = 0$. This quantity is continuous as a function of the measured lamination γ (see Thurston [27] and also Brock [7]).

Proof (Sketch) Let us first recall Thurston's original argument for (3.1). We may assume that γ is a simple closed curve — the general case can be obtained by taking limits, since multiples of simple closed curves are dense in the measured lamination space and the length function is continuous. The first step is to construct a polygonal curve μ on S, homotopic to γ, which consists of $a(\lambda, \gamma)$ segments on leaves of λ (which meet γ at the boundary intersections), connected by $a(\lambda, \gamma)$ "jumps" of bounded length.

This is best seen by lifting γ to a line $\widetilde{\gamma}$ in $\mathbf{H}^2$ and considering the chain of geodesics $\{g_i\}_{i=-\infty}^{\infty}$ in the lift of λ that cross it in boundary intersections. Thus g_i and g_{i+1} are asymptotic on alternating sides of $\widetilde{\gamma}$. Let X be a 1–neighborhood of $\widetilde{\gamma}$. A path $\widetilde{\mu}$ is constructed as a chain of segments, alternately subsegments of $g_i \cap X$ and segments in X of uniformly bounded length joining $g_i \cap X$ to $g_{i+1} \cap X$. See Figure 2 for an example, and Thurston or Canary for the exact construction.

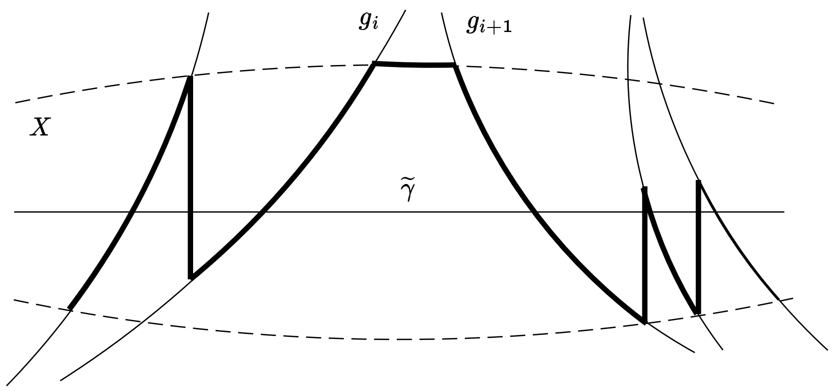

Figure 2: Thurston's polygonal approximation to a geodesic using the leaves of a maximal finite-leaved lamination. The boundary of X is dotted; $\widetilde{\mu}$ is thickened.

We project $\widetilde{\mu}$ to S to obtain a closed polygonal curve μ with $2a(\lambda, \gamma)$ segments. This curve has the property that $\ell_\sigma(\mu) \leq \ell_\sigma(\gamma) + C_0 a(\lambda, \gamma)$. (In the rest of the argument each C_i denotes some constant independent of anything but ϵ and the topological type of S.)

Furthermore, Thurston points out that we can adjust μ, by moving the points where the bounded jumps occur, so that those jumps never occur inside the ϵ_0 Margulis tubes whose cores have length ϵ or less and are not leaves of λ (in Thurston's setting all the Margulis tubes with core ϵ or less are not leaves of λ).

The image of μ by g consists of the images of arcs along λ, which are still geodesic, and of the jumps, which still have bounded lengths. After straightening out the images of the bounded jumps we obtain a polygonal curve ν in N, and we have the estimates

$$\ell_\sigma(\gamma) \leq \ell_\sigma(\mu) \leq \ell_N(\nu) + C_1 a(\lambda,\gamma). \tag{3.3}$$

Thurston next forms a "pleated annulus" connecting ν to the geodesic representative $g(\gamma)^*$. This is a pleated map of a hyperbolic annulus A with one geodesic boundary component that maps to $g(\gamma)^*$, and one polygonal boundary component that maps to ν. The Gauss–Bonnet theorem bounds the area of A in terms of the number of corners of ν, which is $2a(\lambda,\gamma)$. Let us partition ν into several parts: Fixing some $\epsilon' > 0$, ν_0 is the union of segments that admit collar neighborhoods of width ϵ' in A, together with all the bounded length jumps. The length of ν_0 is bounded by $C_2 a(\lambda,\gamma)/\epsilon'$ because of the area bound on A. Any component of $\nu \setminus \nu_0$ must run within distance ϵ' (in A) of some other segment of ∂A. Let ν_1 denote those segments that are distance ϵ' from the boundary component mapping to $g(\gamma)^*$ — the total length of these is at most $\ell_N(g(\gamma)^*)$ plus a constant times the number of segments in ν_1, which again is at most $C_3 a(\lambda,\gamma)$.

Let ν_2 denote the rest of ν — any component of ν_2 is within ϵ' in A of another segment of ν. We then bound the length of ν_2 by $C_4 a(\lambda,\gamma)$, using the Uniform Injectivity Theorem: If two long segments of ν along λ are nearly parallel in A and hence in N, then there are three possibilities:

If the segments are in the ϵ_0–thick part, then the Uniform Injectivity Theorem implies they are nearly parallel in S as well — In fact (see Lemma 3.4 for such an argument), with ϵ' chosen sufficiently small compared to ϵ_0, the short arc connecting them in A is homotopic to a short arc in S, so the curve μ can be shortened significantly by a homotopy in S, reducing its length by at least the sum of the lengths of the two parallel segments. However, since $\ell_\sigma(\gamma) \leq \ell_\sigma(\nu) \leq \ell_\sigma(\gamma) + C_5 a(\lambda,\gamma)$, this is a contradiction unless the segments involved have length bounded by $C_6 a(\lambda,\gamma)$.

If one of the segments is in a Margulis tube of S whose core is not a component of λ, then both must be in the same Margulis tube, since g takes thin parts to thin parts, bijectively by the condition on π_1. Thus one can follow both segments until they exit the tube (since all jumps happen outside this tube) and where they exit one can again use Uniform Injectivity to show the leaves are close in S. Again we obtain too much shortening of ν if the segments are too long.

The last possibility, that the segments are in a Margulis tube whose core has length shorter than ϵ and is a component of λ, is disallowed by Thurston's hypothesis. Thus he obtains

$$\ell_N(\nu) \leq \ell_N(g(\gamma)^*) + C_7 a(\lambda, \gamma)$$

from which (3.1) follows.

Now in order to obtain (3.2), we allow the possibility of Margulis tubes of σ whose cores are closed curves in λ, and let R_λ be the complement of these tubes. Let us first assume that no closed leaf of λ maps to a parabolic curve in N, and return to this case (the noded case) at the end.

Thurston's argument applies to all parts of ν in R_λ, and hence immediately yields

$$\ell_N(\nu \cap R_\lambda) \leq \ell_N(g(\gamma)^*) + C_8 a(\lambda, \gamma). \tag{3.4}$$

We also have

$$\ell_\sigma(\mu \cap R_\lambda) \leq \ell_N(\nu \cap R_\lambda) + C_1 a(\lambda, \gamma) \tag{3.5}$$

by the same argument as for (3.3). It remains to establish a connection between $\ell_\sigma(\mu \cap R_\lambda)$ and $\ell_\sigma(\gamma \cap R_\lambda)$. To do this, we note that the configuration of λ in a Margulis tube outside R_λ is as follows: There is the core curve c of the tube, which is a leaf of λ, and the other leaves of λ that meet the tube must enter from one side or the other and spiral around c. An arc of γ that runs through this tube has one boundary intersection each with two leaves of λ, on opposite sides of the core c, and possibly one boundary intersection with the c itself (if the spiraling on the two sides of c is in opposite directions). Thus the intersection of μ with this Margulis tube consists of at most three arcs of λ, each of which remains in a bounded neighborhood of $\widetilde{\gamma}$ in the lift to $\mathbf{H}^2$, and at most two bounded jumps. It follows that the arc of μ intersected with this tube has length at most a constant plus the length of the corresponding arc of γ. We conclude that

$$\ell_\sigma(\gamma \cap R_\lambda) \leq \ell_\sigma(\mu \cap R_\lambda) + C_9 a(\lambda, \gamma).$$

This together with (3.4) and (3.5) gives the desired inequality.

Finally, we consider the noded case, where some closed leaves $\{c_1, \ldots c_k\}$ of λ are mapped to parabolics in N. Then we let S' be the complement of these leaves and $g\colon S' \to N$ a noded pleated surface mapping geodesically all leaves of λ except the c_i. The g–image of an end of a leaf of μ that spirals around c_i must then terminate in the corresponding cusp of N, and two such ends of leaves have asymptotic images. If γ does not intersect any of the c_i the argument can be repeated as before, but in general $\gamma \cap S'$ may be a union of arcs, each represented by an infinite geodesic in $g(S')$ with its ends in the cusps. The construction of μ in S is the same as before, but where μ crosses one of the c_i it leaves the domain of g. For each such crossing, the image $g(\mu \cap S')$ traverses two leaves of λ that are asymptotic into the corresponding cusp. To build ν, we simply extend these leaves far enough into the cusp that we can jump between them with a bounded arc in the correct homotopy class. As before Thurston's argument applies to all parts of ν in R_λ, which excludes the segments in the parabolic cusps, and the length of $\gamma \cap R_\lambda$ and $\mu \cap R_\lambda$ again differ by a bounded multiple of $a(\lambda, \gamma)$. □

4 The proof of Theorem B

Recall the statement of our second main theorem:

Theorem B *Given a surface S, $\epsilon > 0$ and $L > 0$, there exists K so that, if $\rho\colon \pi_1(S) \to PSL_2(\mathbf{C})$ is a Kleinian surface group and Y is a proper essential subsurface of S, then*

$$\mathrm{diam}_Y(\mathcal{C}_0(\rho, L)) \geq K \quad \Longrightarrow \quad \ell_\rho(\partial Y) \leq \epsilon.$$

In fact, we will prove the equivalent conclusion

$$\ell_\rho(\partial Y) > \epsilon \quad \Longrightarrow \quad \mathrm{diam}_Y(\mathcal{C}_0(\rho, L)) < K.$$

We first give the proof in the case where Y is not an annulus. The annular case is similar, and we describe it at the end of the section.

The proof in the non-annular case

The proof reduces to two main lemmas. First, given a hyperbolic surface (Z, σ) with geodesic boundary, we define a "minimal proper arc" to be an embedded arc τ with endpoints on ∂Z, not homotopic into ∂Z, which is minimal in σ–length among all such arcs.

Lemma 4.1 *For $L > 0$ there exists $D > 0$, depending only on L and the topology of S, such that the following holds:*

Given a Kleinian surface group $\rho\colon \pi_1(S) \to PSL_2(\mathbf{C})$, a non-annular proper essential subsurface $Y \subset S$, and any simple closed curve γ in S intersecting Y essentially, such that $\ell_\rho(\gamma) \le L$, there is a pleated surface g_γ in the homotopy class of ρ mapping ∂Y geodesically, with induced metric σ_γ, such that for any minimal proper arc τ in (Y, σ_γ) we have

$$d_Y(\gamma, \tau) \le D. \tag{4.1}$$

We remark that in this lemma we do not use the assumption on $\ell_\rho(\partial Y)$.

The second lemma will show that, over all pleated surfaces $g\colon S \to N$ mapping ∂Y geodesically the set of minimal proper arcs in Y with respect to the induced metrics has uniformly bounded diameter in $\mathcal{C}'(Y)$. More precisely:

Lemma 4.2 *For any $\epsilon > 0$ there exists D (depending on ϵ and the topological type of S) such that the following holds.*

Let $\rho\colon \pi_1(S) \to PSL_2(\mathbf{C})$ be a Kleinian surface group, and let Y be a non-annular proper essential subsurface of S and g_0, g_1 a pair of pleated surfaces in the homotopy class $[\rho]$ mapping ∂Y to geodesics. Let σ_0 and σ_1 be the induced metrics on S. Suppose that $\ell_\rho(\partial Y) \ge \epsilon > 0$. If τ_0 and τ_1 are minimal proper arcs in (Y, σ_0) and (Y, σ_1), respectively, then

$$d_Y(\tau_0, \tau_1) \le D.$$

Thus, assuming $\ell_\rho(\partial Y) \ge \epsilon$, for any two curves γ_0, γ_1 in $\mathcal{C}_0(\rho, L)$ whose projections $\pi_Y(\gamma_i)$ are non-empty, we apply Lemma 4.1 to obtain two pleated surfaces mapping ∂Y geodesically and minimal proper arcs τ_0, τ_1 in Y with respect to the two induced metrics, with a bound on $d_Y(\gamma_i, \tau_i)$ from (4.1). Lemma 4.2 then implies the final bound on $d_Y(\gamma_0, \gamma_1)$. This completes the proof modulo the two lemmas.

Before giving the proofs let us recall the following geometric fact: Let Z be a complete hyperbolic or simplicial hyperbolic surface with boundary curved outward (curvature vector pointing out of the surface) — for example a surface with geodesic boundary or a cusped surface minus a horoball neighborhood of the cusp. Then the length r of a minimal proper arc in Z satisfies

$$2\pi|\chi(Z)| \ge \ell(\partial Z)\sinh(r/2), \tag{4.2}$$

by the Gauss–Bonnet theorem and an elementary formula for the area of an embedded collar around ∂Z.

Proof of Lemma 4.1 Let γ denote any simple closed curve with $\ell_\rho(\gamma) \le L$, which has non-trivial intersection with Y. If $\gamma \subset Y$ then we can simultaneously pleat along ∂Y and γ yielding a pleated surface g_γ in which γ has length bound L.

Let R_γ denote the complement in S of the ϵ_0–Margulis tubes (with respect to the induced metric of g_γ) whose cores are components of ∂Y.

Note also that we allow g_γ to be a *noded* pleated surface, in case any components of $\partial Y \cup \gamma$ are parabolic in N (see Section 3). In this case g_γ is defined on the complement S' of the parabolic curves, and τ is a geodesic arc in Y with endpoints in the cusps and whose length within R_γ is minimal.

Inequality (4.2) applied to $R_\gamma \cap Y$ gives a uniform upper bound on the length of $\tau \cap R_\gamma$. This gives us a uniform upper bound on the intersection number $i(\gamma, \tau)$, and hence on $d_Y(\gamma, \tau)$ by Lemma 2.1.

From now on assume that γ has nontrivial intersection with ∂Y.

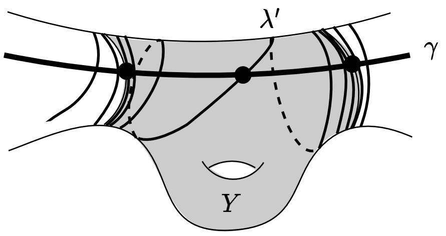

Figure 3: An intersection of γ with Y and the resulting leaves of the spun lamination λ'. Y is shaded and the resulting boundary intersections are indicated by bullets.

Let $\mathcal{D}$ denote the mapping class that performs one positive Dehn twist on each component of ∂Y. The sequence of curves $\mathcal{D}^n(\gamma)$ converge, as $n \to \infty$, to a finite-leaved lamination λ' whose non-compact leaves spiral around ∂Y, in such a way that the spiraling on opposite sides of any component is in opposite directions. Its closed leaves are exactly ∂Y. (This construction is often called "spinning" γ around ∂Y — see Figure 3 for an example).

After adding a finite number of other infinite leaves spiraling around ∂Y, we can obtain a *maximal* finite-leaved lamination λ containing λ'. We immediately observe that

$$a(\lambda, \gamma) = a(\lambda', \gamma) = 2i(\partial Y, \gamma) \tag{4.3}$$

since γ can be represented by a chain of segments of infinite leaves of λ, each asymptotic to a component of ∂Y at each end, with successive segments joined by paths along ∂Y. After perturbation to a curve transverse to λ we obtain a curve which has one boundary intersection for each intersection of γ with ∂Y, and one boundary intersection with each of the leaves in the chain.

Let g_λ be the pleated surface mapping λ geodesically, and let σ_λ be its induced metric. Again if some components of ∂Y are parabolic in N then g_λ will be a noded pleated surface defined on the complement S' of these components, and σ_λ will have cusps corresponding to the ends of S'. Otherwise let $S' = S$.

Again let R_λ be the complement in (S', σ_λ) of the ϵ_0–Margulis tubes whose cores are components of λ (hence of ∂Y). Realize γ by its geodesic representative in (S', σ_λ) (which we continue to call γ; in the noded case γ is replaced by the geodesic representative rel cusps of $\gamma \cap S'$). Theorem 3.5 part (3.2) then implies

$$\ell_{\sigma_\lambda}(\gamma \cap R_\lambda) \le \ell_\rho(\gamma) + Ca(\lambda, \gamma). \tag{4.4}$$

Using (4.3) and the fact that $\ell_\rho(\gamma) \le L$, we conclude that

$$\ell_{\sigma_\lambda}(\gamma \cap R_\lambda) \le L + 2Ci(\gamma, \partial Y). \tag{4.5}$$

We remark that, in case all components of ∂Y have length at least ϵ_0, we may use Thurston's original version of Theorem 3.5 (3.1), or equivalently set $R_\lambda = S$.

Let $\gamma_1, \dots, \gamma_m$ be the component arcs of $\gamma \cap (R_\lambda \setminus \partial Y)$, where $m = i(\gamma, \partial Y)$. At least half of these arcs are contained in $Y \cap R_\lambda$. A bit of algebra shows that, choosing $K = 4C + 2L$, less than $m/2$ of the arcs can have length greater than K. Thus we are guaranteed at least one arc γ_i in $Y \cap R_\lambda$ with σ_λ–length bounded by K.

The length of $\tau \cap R_\lambda$ is bounded above as before by (4.2), and hence again we obtain a bound on the intersection number $i(\gamma_i, \tau)$. Lemma 2.1 then bounds $d_Y(\gamma_i, \tau)$, which is within 1 of $d_Y(\gamma, \tau)$ since γ_i is in the same simplex of $\mathcal{C}'(Y)$ as $\pi_Y(\gamma)$. □

Proof of Lemma 4.2 We need to recall an additional geometric fact: There is a constant $\delta_1 > 0$ for which, if Z is any hyperbolic surface with geodesic boundary and τ and τ' are essential properly embedded arcs in Z whose lengths are at most δ_1, then τ and τ' are either homotopic keeping endpoints in ∂Z, or they are disjoint. (This follows directly from the fact that boundary components of a hyperbolic surface which are close at one point must be nearly parallel for long stretches.)

Given this constant δ_1, let $\delta_0 \leq \delta_1$ be the constant given by Lemma 3.4 (Short bridge arcs). Note that, by Lemma 3.3, we may assume, possibly precomposing g_1 by a domain homeomorphism isotopic to the identity, that g_0 and g_1 are homotopic relative to the common pleating lamination ∂Y. This allows us to apply Lemma 3.4.

There are now two cases.

Case a: Suppose $\ell_{\sigma_0}(\tau_0) \leq \delta_0$. Since τ_0 is a primitive bridge arc for ∂Y, Lemma 3.4 part (2) guarantees that $\ell_{\sigma_1}(\tau_0) \leq \delta_1$ as well. Then $\ell_{\sigma_1}(\tau_1) \leq \delta_1$ since τ_1 is minimal in the metric σ_1, and we may then conclude, by choice of δ_1, that the arcs τ_0 and τ_1 are either homotopic or disjoint, and hence $d_{\mathcal{C}'(Y)}(\tau_0, \tau_1) \leq 1$.

Case b: Suppose $\ell_{\sigma_0}(\tau_0) > \delta_0$. Then (4.2) implies that there is some $A(\delta_0)$ such that $\ell_N(\partial Y) = \ell_{\sigma_0}(\partial Y) \leq A$.

The rest of the proof resembles that of the Connectivity Lemma 8.1 in [24]:

Let $g_t, t \in [0,1]$ be a continuous family of maps connecting g_0 to g_1 so that for each $t \in (0,1)$, g_t is a simplicial hyperbolic map, and maps ∂Y geodesically. Recall that this means that g_t may be noded on parabolic components of ∂Y, but not on all of them since we assume $\ell(\partial Y) \geq \epsilon > 0$.

The existence of such a family follows from the techniques of Thurston in [29] and Canary [9], for example. The pleated maps g_0, g_1 may be approximated by simplicial hyperbolic surfaces, in which lamination leaves that spiral an infinite number of times along the geodesics ∂Y are replaced by triangulation edges that terminate on ∂Y. Any two such triangulations may be connected by a sequence of elementary moves in which the edges on ∂Y are fixed and the other edges are replaced one-by-one, by a theorem of Hatcher [14]. Each such triangulation gives rise to a simplicial hyperbolic surface in which the components of ∂Y are mapped to their geodesic representatives (or the surface is noded on the ones that are parabolic), and Canary shows in [9] that each elementary move between two such surfaces may be realized by a continuous family of simplicial hyperbolic surfaces.

The induced metrics σ_t vary continuously, in the sense that for any fixed homotopy class β of curves or arcs rel boundary, $\ell_{\sigma_t}(\beta)$ is continuous in t.

Now given any essential properly embedded arc τ in Y, let $E_\tau \subset [0,1]$ denote the set of t–values for which τ is (homotopic rel ∂Y to) a minimal proper arc with respect to σ_t. Since we are assuming $\ell_\rho(\partial Y) \geq \epsilon$, (4.2) gives an upper bound $L_3 = L_3(\epsilon)$ on $\ell_{\sigma_s}([\tau])$ for $s \in E_\tau$. Continuity of the metrics implies that E_τ is closed, and clearly the family $\{E_\tau\}$ covers $[0,1]$.

We now observe two facts.

First, if $E_\tau \cap E_{\tau'} \neq \emptyset$ then τ and τ' are simultaneously (homotopic to) shortest arcs in the same metric σ_s, $s \in E_\tau \cap E_{\tau'}$. We may assume, possibly after homotopy, that τ and τ' are shortest-length representatives of their classes. We claim that $i(\tau,\tau') \leq 1$. If not, let x and y be two intersection points of τ and τ', cutting each into three successive arcs τ_1, τ_2, τ_3 and $\tau_1', \tau_2', \tau_3'$, respectively. The concatenations $\xi_1 = \tau_1 * \tau_1'$ and $\xi_3 = \tau_3 * \tau_3'$ are arcs with endpoints in ∂Y, and are not homotopic into the boundary: if, say, ξ_1 were, then since it meets ∂Y at right angles by the minimality of τ and τ', we would obtain a disk that violates the Gauss–Bonnet theorem. (Actually if the metric σ_t is singular exactly at a point where ξ_1 meets ∂Y then there is an angle of *at least* $\pi/2$ on each side, and the same argument holds.) Let a_i and a_i' be the lengths of τ_i and τ_i', respectively. Let $L = a_1 + a_2 + a_3 = a_1' + a_2' + a_3'$. Then since $a_2, a_2' > 0$ we have $a_1 + a_3 + a_1' + a_3' < 2L$, so either $a_1 + a_1' < L$ or $a_3 + a_3' < L$. This means that either ξ_1 or ξ_3 is strictly shorter than τ and τ', contradicting their minimality.

Thus $i(\tau,\tau') \leq 1$, and this easily gives $d_{\mathcal{C}'(Y)}(\tau,\tau') \leq 2$.

The second fact is that, since $\ell_\rho(\partial Y) \leq A$, all the arcs τ for which $E_\tau \neq \emptyset$, together with the boundary components of Y on which their endpoints lie, can be realized in the quotient hyperbolic 3–manifold N_ρ by a 1–complex with at most k components (where k is the number of components of ∂Y), each of which has diameter at most $A + L_3(k-1)$. Each τ together with one or two segments on ∂Y gives rise to a loop in this 1–complex of length at most $A + 2L_3$, and the loops for homotopically distinct τ's are homotopically distinct. A standard application of the Margulis lemma gives an upper bound $M = M(R)$ for any non-elementary Kleinian group Γ on the number of distinct elements of Γ that can translate any one point a distance R or less. This gives a bound $M' = M'(A, L_3)$ on the number of (homotopically distinct) arcs τ with $E_\tau \neq \emptyset$.

Now consider the graph T whose vertices are those homotopy classes of τ with $E_\tau \neq \emptyset$, and whose edges are those $([\tau],[\tau'])$ for which $E_\tau \cap E_\tau' \neq \emptyset$. The $\mathcal{C}'(Y)$–distance between two endpoints of an edge is at most 2 by the first observation, and the number of vertices is at most M', by the second. The fact that the E_τ cover $[0,1]$ and are closed and finite in number means that T is connected. The diameter of the set of vertices of T is therefore bounded by $2(M'-1)$.

For our original arcs τ_0 and τ_1 we have $i \in E_{\tau_i}$ and therefore $d_Y(\tau_0, \tau_1) \leq 2(M' - 1)$.

This concludes the proof of Lemma 4.2 and hence of Theorem B in the non-annular case. $\square$

The annular case

The proof again reduces to two lemmas analogous to Lemmas 4.1 and 4.2, and their proofs are similar. We introduce a new bit of notation and discuss the differences in the proofs:

Let Y be an annulus and α its core curve. If α is geodesic in some hyperbolic surface S we can consider a *minimal curve crossing* α to be a curve β constructed as follows: Pick one side of α in S and let τ be a minimal length primitive bridge arc for α that is incident to it on this side. Let τ' be a minimal length primitive bridge arc for α that is incident to it on the other side. If α is non-separating and τ meets it on both sides then we let $\tau' = \tau$. Let β be a minimal length shortest simple curve that can be represented as a concatenation of τ, τ' (if they are different), and arcs on α. In particular β crosses α once or twice.

The analogue of Lemma 4.1 is now the following:

Lemma 4.3 *Given $\epsilon > 0$ and S there exists $M > 0$ such that, for any Kleinian surface group $\rho\colon \pi_1(S) \to PSL_2(\mathbf{C})$, the following holds:*

Let Y be an essential annulus with core α such that $\ell_\rho(\alpha) \geq \epsilon$, and let $\gamma \in \mathcal{C}_0(\rho, L)$ intersect α essentially. Then there exists a pleated surface g_γ in whose induced metric σ_γ a minimal curve β crossing α satisfies

$$d_Y(\gamma, \beta) \leq M \tag{4.6}$$

Proof The construction of g_γ is the same as before (starting by spinning γ around α to obtain a finite-leaved lamination λ'), but now we find that if α is very short the conclusion is false: γ could wind arbitrarily inside the Margulis tube of α without violating Theorem 3.5 (Efficiency of pleated surfaces). Note that in the non-annular case γ could also have wound around ∂Y but this had no effect on the distance d_Y. At any rate we now must use the hypothesis of Theorem B that $\ell_\rho(\alpha) \geq \epsilon$.

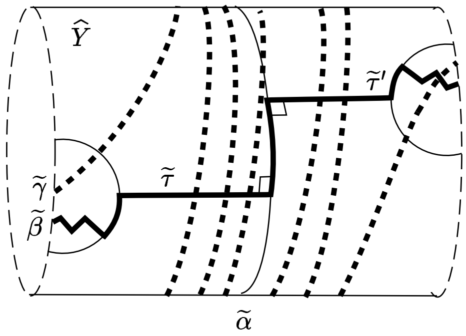

Figure 4: A component $\widetilde{\gamma}$ of the lift of γ to the annulus $\widehat{Y}$ is indicated by the broken line. A component $\widetilde{\beta}$ of the lift of β is indicated in heavy lines, using a representative that traces along the lifts of α, τ and τ'. Note that $\widetilde{\tau}$ and $\widetilde{\tau}'$ terminate on translates of $\widetilde{\alpha}$, pictured as semicircles.

In this case, part (3.1) of Theorem 3.5 (Efficiency of pleated surfaces), together with (4.3), give us

$$\ell_{\sigma_\gamma}(\gamma) \leq L + 2Ci(\alpha, \gamma). \tag{4.7}$$

We next relate $d_Y(\beta, \gamma)$ to the number of times γ intersects the arcs τ and τ': lift α to the core $\widetilde{\alpha}$ of the annular cover $\widehat{Y}$ corresponding to Y, and consider a component $\widetilde{\gamma}$ of the lift of γ that crosses $\widetilde{\alpha}$. In Figure 4 we indicate $\widetilde{\gamma}$, together with a lift $\widetilde{\beta}$ of a representative of β that travels along $\widetilde{\alpha}$ and the lifts of τ and τ'. The essential intersections of $\widetilde{\beta}$ with $\widetilde{\gamma}$ can be read off from this diagram (using the fact that all the segments shown are geodesics) to obtain

$$i(\widetilde{\gamma}, \widetilde{\beta}) - 2 \leq \#(\widetilde{\gamma} \cap (\widetilde{\tau} \cup \widetilde{\tau}')) \leq i(\widetilde{\gamma}, \widetilde{\beta}) + 2. \tag{4.8}$$

(The ± 2 comes from various possibilities for the endpoints and the intersection with $\widetilde{\alpha}$). Thus by (2.1),

$$d_Y(\beta, \gamma) \leq \#(\widetilde{\gamma} \cap (\widetilde{\tau} \cup \widetilde{\tau}')) + 3. \tag{4.9}$$

Now we see from the situation in $\widehat{Y}$ that any two successive intersections of $\widetilde{\gamma}$ with $\widetilde{\tau}$ (or $\widetilde{\tau}'$) bound a segment of $\widetilde{\gamma}$ of length at least $\ell(\alpha) \geq \epsilon$. All these segments project to disjoint segments in γ, and summing over all lifts of γ that cross $\widetilde{\alpha}$, we obtain

$$\begin{aligned} \ell_{\sigma_\gamma}(\gamma) &\geq \epsilon i(\alpha, \gamma)(\#(\widetilde{\gamma} \cap (\widetilde{\tau} \cup \widetilde{\tau}')) - 2) \\ &\geq \epsilon i(\alpha, \gamma)(d_Y(\beta, \gamma) - 5) \end{aligned} \tag{4.10}$$

where the last line follows from (4.9). Putting (4.7) and (4.10) together, we find that

$$d_Y(\beta,\gamma) \le 5 + (L+2C)/\epsilon. \tag{4.11}$$

This is the desired bound. □

The analogue of Lemma 4.2 is the following. A closely related result appears in Brock [7].

Lemma 4.4 *For any $\epsilon > 0$ there exists D (depending on ϵ and the topological type of S) such that for any Kleinian surface group $\rho\colon \pi_1(S) \to PSL_2(\mathbf{C})$ the following holds.*

Let Y be an essential annulus and g_0, g_1 a pair of pleated surfaces in the homotopy class $[\rho]$ mapping the core α of Y to a geodesic. Let σ_0 and σ_1 be the induced metrics on S. Suppose that $\ell_N(\alpha) \ge \epsilon > 0$. If β_0 and β_1 are minimal curves crossing α relative to σ_0 and σ_1, respectively, then

$$d_Y(\beta_0,\beta_1) \le D.$$

As in the proof of Lemma 4.2, if $\ell_{\sigma_0}(\tau_0) \le \delta_0$ then τ_0 and τ_1 are either equal or disjoint. This bounds $i(\beta_0,\beta_1)$ and, in turn, $d_Y(\beta_0,\beta_1)$.

If $\ell_{\sigma_0}(\tau_0) > \delta_0$ we deduce a length upper bound on α as in (4.2), and consider the homotopy argument of case (b) in Lemma 4.2, in which we join g_0 to g_1 by a family g_t. In each metric σ_t we build β_t as before, and the same argument shows that only a bounded number of curves can be built in this way. The intersection number of two such curves that occur for the same value of t is again bounded, and we obtain the bound on $d_Y(\beta_0,\beta_1)$ in the same way.

5 The proof of Theorem A

Recall again the statement:

Theorem A *For any Kleinian surface group ρ with ending invariants $\nu_\pm$, if*

$$\sup_{Y\subset S} d_Y(\nu_+,\nu_-) = \infty$$

then

$$\operatorname{inj}_0(\rho) = 0,$$

where the supremum is over proper essential subsurfaces Y in S not all of whose boundaries map to parabolics.

As in the proof of Theorem B we will work with the contrapositive statement, that for any Kleinian surface group ρ with ending invariants $\nu_\pm$, if $\mathrm{inj}_0(\rho) > 0$ then

$$\sup_{Y \subset S} d_Y(\nu_+, \nu_-) < \infty,$$

with the supremum taken over essential subsurfaces Y not all of whose boundary components map to parabolics.

To make sense of this statement we must first give a general description of the ending invariants $\nu_\pm$ of a Kleinian surface group $\rho\colon \pi_1(S) \to \mathrm{PSL}_2(\mathbf{C})$. For further discussion see Thurston [29], Bonahon [5], and Ohshika [25], as well as [22].

One convenient way to think of an end invariant ν_+ is as a pair (ν_+^G, ν_+^T) with the following structure. The first component ν_+^G is either a geodesic lamination that admits a transverse measure of full support, or the "empty" lamination $\emptyset$. The second component ν_+^T is either a conformal structure of finite type on a (possibly disconnected) essential subsurface $R_+ \subset S$ (ie $\nu_+^T \in \mathcal{T}(R_+)$), or "$\emptyset$" (in which case R_+ is the empty set as well).

Furthermore, ν_+^G is supported in the complement of R_+, and ν_+ *fills* in the sense that any nontrivial curve in S either intersects a component of ν_+^G, or of R_+, or it is isotopic to a closed curve component of ν_+^G. The other end invariant ν_- is described in the same way.

(One should keep in mind the special case $\nu_+^G = \emptyset$, in which case $R_+ = S$ and ν_+^T is the conformal structure on a geometrically finite end of N, as well as the case $\nu_+^T = \emptyset$, when ν_+^G is a geodesic lamination that fills S and the manifold has a simply degenerate end. The other cases are hybrids of these two in which ends of different types are separated by parabolics.)

Let Y be any proper essential subsurface whose boundary components are not all homotopic to components of ν_+^G. We can define $\pi_Y(\nu_+)$ as the union of $\pi_Y(\nu_+^G)$ and $\pi_Y(\nu_+^T)$, which are defined as follows.

$\pi_Y(\nu_+^G)$ is defined similarly to the definition of π_Y for closed curves: it is the barycenter in $\mathcal{C}'(Y)$ of the simplex whose vertices are the equivalence classes of essential closed curves and properly embedded arcs in $\nu_+^G \cap Y$ (if there are infinite leaves of ν_+^G that are wholly contained in Y, we ignore them). Note that by choice of Y this is empty only if Y is contained in R_+.

To define $\pi_Y(\nu^T_+)$, let α_+ be the union of shortest curves in the conformal structure ν^T_+ that intersect Y (this could be the boundary of R_+, whose ν^T_+-length is 0), and let $\pi_Y(\nu^T_+) = \pi_Y(\alpha_+)$. This is now a subset of $\mathcal{C}'(Y)$ rather than a point, but its diameter is uniformly bounded under one more assumption: Let $\epsilon > 0$ be fixed and suppose that Y is anything other than an annulus whose core is homotopic to a curve in R_+ with ν^T_+-length less than ϵ. Then (assuming $\nu^T_+ \neq \emptyset$) the curves α_+ have ν^T_+–length bounded above by some $L(\epsilon)$, and hence their intersection number is bounded above. Thus their distance in $\mathcal{C}'(Y)$ is bounded above, by Lemma 2.1.

The union $\pi_Y(\nu_+) = \pi_Y(\nu^T_+) \cup \pi_Y(\nu^G_+)$ is always nonempty under our assumptions on Y, and also has bounded diameter: Every arc in the definition of $\pi_Y(\nu^G_+)$ is disjoint from the arcs in $\pi_Y(\nu^T_+)$ since the supports are disjoint.

Without giving the complete definition, let us just list the properties of the ending invariants which we will be using. First, the closed components of ν^G_+ (which include the boundary curves of the support of ν^T_+), are all taken to parabolics by ρ. The same is true for ν_-, and this accounts for all parabolics in ρ other than those coming from punctures of S. The Riemann surface (R_+, ν^T_+) is part of the quotient of the domain of discontinuity of $\rho(\pi_1(S))$, and so by an inequality of Bers [4], for each curve γ in R_+, its length $\ell_\rho(\gamma)$ is bounded above by $2\ell_{\nu^T_+}(\gamma)$. Finally, for each component λ of ν^G_+ there is a sequence of simple closed curves α_i in S, with $\ell_\rho(\alpha_i) \leq L_0$, which converge as projective measured laminations to a limit whose support is λ. In particular the Hausdorff limit of $\{\alpha_i\}$ contains λ, so that $\pi_Y(\alpha_i)$ for any fixed Y is eventually in the same simplex as $\pi_Y(\lambda)$.

Let us now also assume that $\mathrm{inj}_0(\rho) = \epsilon > 0$. Then in particular (applying Bers' inequality from above) the conformal structures $\nu^T_\pm$ contain no non-peripheral curves of length less than ϵ, and there is a bound $L_4 \geq L_0$ on the ρ–lengths of the curves used to construct $\pi_Y(\nu^T_\pm)$. Putting these facts together we then have the following statement:

$\pi_Y(\nu_-)$ and $\pi_Y(\nu_+)$ are contained in a 1–neighborhood of $\pi_Y(\mathcal{C}(\rho, L_4))$.

Now the proof of the theorem is almost immediate. Let Y be any essential subsurface of S not all of whose boundaries map to parabolics. Then $\pi_Y(\nu_+)$ and $\pi_Y(\nu_-)$ are both non-empty, and $\ell_\rho(\partial Y) \geq \mathrm{inj}_0(\rho) = \epsilon$, and hence by Theorem B, $\mathrm{diam}_Y(\mathcal{C}_0(\rho, L_4)) \leq K(\epsilon, L_4)$. The above observation then yields $d_Y(\nu_-, \nu_+) \leq K + 2$. This bound, independent of Y, gives the desired result.

References

[1] **L Ahlfors**, **L Bers**, *Riemann's mapping theorem for variable metrics*, Ann. of Math. 72 (1960) 385–404

[2] **R Benedetti**, **C Petronio**, *Lectures on hyperbolic geometry*, Springer–Verlag Universitext (1992)

[3] **L Bers**, *Simultaneous uniformization*, Bull. Amer. Math. Soc. 66 (1960) 94–97

[4] **L Bers**, *On boundaries of Teichmüller spaces and on Kleinian groups I*, Ann. of Math. 91 (1970) 570–600

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

[6] **F Bonahon**, **J P Otal**, *Variétés hyperboliques a geodesiques arbitrairement courtes*, Bull. London Math. Soc. 20 (1988) 255–261

[7] **J Brock**, *Continuity of Thurston's length function*, preprint (1998), to appear in Geom. Funct. Anal.

[8] **R D Canary**, *Algebraic limits of schottky groups*, Trans. Amer. Math. Soc. 337 (1993) 235–258

[9] **R D Canary**, *A covering theorem for hyperbolic 3–manifolds and its applications*, Topology, 35 (1996) 751–778

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

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

[12] **A Fathi**, **F Laudenbach**, **V Poenaru**, *Travaux de Thurston sur les surfaces*, Asterisque, vol. 66–67 (1979)

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

[14] **A Hatcher**, *On triangulations of surfaces*, Topology Appl. 40 (1991) 189–194

[15] **A E Hatcher**, *Measured lamination spaces for surfaces, from the topological viewpoint*, Topology Appl. 30 (1988) 63–88

[16] **A E Hatcher**, **W P Thurston**, *A presentation for the mapping class group*, Topology 19 (1980) 221–237

[17] **John Hempel**, *3–manifolds as viewed from the curve complex*, MSRI Preprint #1997-095 and E-print `math.GT/9712220`

[18] **H A Masur**, **Y Minsky**, *Geometry of the complex of curves II: Hierarchical structure*, E-print `math.GT/9807150`, to appear in Geom. Funct. Anal.

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

[20] **C McMullen**, *Renormalization and 3–manifolds which fiber over the circle*, Princeton University Press, Princeton, NJ (1996)

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

[22] **Y Minsky**, *On Thurston's ending lamination conjecture*, from: "Proceedings of Low-Dimensional Topology, May 18–23, 1992", International Press (1994)

[23] **Y Minsky**, *A geometric approach to the complex of curves*, from: "Proceedings of the 37th Taniguchi Symposium on Topology and Teichmüller Spaces", (S Kojima et al, editors), World Scientific (1996) 149–158

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

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

[26] **R Penner**, **J Harer**, *Combinatorics of train tracks*, Annals of Math. Studies no. 125, Princeton University Press (1992)

[27] **W Thurston**, *Hyperbolic structures on 3–manifolds, II: surface groups and manifolds which fiber over the circle*, E-print: `math.GT/9801045`, (original 1986)

[28] **W Thurston**, *Hyperbolic structures on 3–manifolds, III: deformations of 3–manifolds with incompressible boundary*, E-print: `math.GT/9801058`, (original 1986)

[29] **W Thurston**, *The geometry and topology of 3–manifolds*, Princeton University Lecture Notes, online at `http://www.msri.org/publications/books/gt3m` (1982)

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

[31] **W Thurston**, *On the geometry and dynamics of diffeomorphisms of surfaces*, Bull. Amer. Math. Soc. 19 (1988) 417–431

[32] **W Thurston**, *Three-Dimensional Geometry and Topology*, (S Levy, editor), Princeton University Press (1997)