Local topology in deformation spaces of hyperbolic 3–manifolds

JEFFREY F BROCK
KENNETH W BROMBERG
RICHARD D CANARY
YAIR N MINSKY

We prove that the deformation space $AH(M)$ of marked hyperbolic 3–manifolds homotopy equivalent to a fixed compact 3–manifold M with incompressible boundary is locally connected at minimally parabolic points. Moreover, spaces of Kleinian surface groups are locally connected at quasiconformally rigid points. Similar results are obtained for deformation spaces of acylindrical 3–manifolds and Bers slices.

30F40; 57M50

1 Introduction

The conjectural picture for the topology of the deformation space $AH(M)$ of all (marked) hyperbolic 3–manifolds homotopy equivalent to a fixed compact 3–manifold M has evolved from one of relative simplicity to one far more complicated in recent years. Indeed, the interior of this space has been well-understood since the late 1970's. Roughly, components of $AH(M)$ are enumerated by (marked) homeomorphism types of compact 3–manifolds homotopy equivalent to M, and each component is a manifold parametrized by natural conformal data. In the last decade, however, a string of results has established that the topology of $AH(M)$ itself is not well-behaved. In particular, $AH(M)$ fails to be locally connected when M is an untwisted I–bundle over a closed surface (see Bromberg [21] and Magid [43]), and a new conjectural picture in which such pathology is prevalent has replaced the old.

The present paper clarifies the role that the geometry and topology of 3–manifolds associated to points in the boundary of $AH(M)$ plays in the local topology at such points. In particular, we show that the topology of $AH(M)$ is well-behaved at many points; if M has incompressible boundary, then $AH(M)$ is locally connected at "generic" points in the boundary. When M is acylindrical or an untwisted I–bundle we obtain finer results.

Published: 26 July 2011 DOI: 10.2140/gt.2011.15.1169

Central to the present discussion are recent fundamental improvements in the understanding of the internal geometry and topology of ends of hyperbolic 3–manifolds. Via the Ending Lamination Theorem of Minsky [51] and Brock, Canary and Minsky [17] and the model manifold developed in its proof, the Tameness Theorem of Agol [1] and Calegari and Gabai [25] and the Density Theorem of [17], Namazi and Souto [53] and Ohshika [58], we develop a more complete picture of the topological complexity at the boundary of deformation spaces.

Our first theorem extracts consequences for the local structure of deformation spaces in terms of the topology of M and the presence of parabolic elements for an element ρ in the boundary of $AH(M)$.

Two components B and C of $\mathrm{int}(AH(M))$ are said to *bump* at $\rho \in \partial AH(M)$ if $\rho \in \overline{B} \cap \overline{C}$. A component B of $\mathrm{int}(AH(M))$ is said to *self-bump* at $\rho \in \partial B$ if there exists a neighborhood W of ρ such that if V is a neighborhood of ρ which is contained in W, then $V \cap B$ is disconnected. A point $\rho \in \partial AH(M)$ is said to be *uniquely approachable* if there is no bumping or self-bumping at ρ. The Density Theorem [17; 53; 58] asserts that $AH(M)$ is the closure of its interior, so $AH(M)$ is locally connected at all uniquely approachable points.

Theorem 1.1 *Let M be a compact 3–manifold with incompressible boundary and $\rho \in \partial AH(M)$. If every parabolic element of $\rho(\pi_1(M))$ lies in a rank-two free abelian subgroup, then ρ is uniquely approachable. In particular, $AH(M)$ is locally connected at ρ.*

Remark Such points ρ are generic in the boundary of $AH(M)$ in the sense of Lemma 4.2 in Canary and Hersonsky [30].

Recall that if $\rho \in AH(M)$, then $N_\rho = \mathbb{H}^3/\rho(\pi_1(M))$ is a hyperbolic 3–manifold homotopy equivalent to M. If $\Omega(\rho)$ is the domain of discontinuity for the action of $\rho(\pi_1(M))$ on $\widehat{\mathbb{C}}$, then $\partial_c N_\rho = \Omega(\rho)/\rho(\pi_1(M))$ is a Riemann surface called the *conformal boundary* of N_ρ. In order to rule out bumping in the presence of parabolics we place the additional restriction on ρ that every component of its conformal boundary is a thrice-punctured sphere. Such a ρ is called *quasiconformally rigid*. Notice that this includes the case that the conformal boundary is empty.

Theorem 1.2 *Let M be a compact 3–manifold. If ρ is a quasiconformally rigid point in $\partial AH(M)$, then there is no bumping at ρ.*

In order to rule out self-bumping, we make additional restrictions on the topology of M.

Theorem 1.3 *Let M be a compact 3–manifold which is either acylindrical or homeomorphic to $S \times I$, for a closed surface S. If ρ is a quasiconformally rigid point in $\partial AH(M)$ then there is no self-bumping at ρ.*

We may combine Theorems 1.2 and 1.3 to establish the following corollary.

Corollary 1.4 *Let M be a compact 3–manifold which is either acylindrical or homeomorphic to $S \times I$, for a closed surface S. If ρ is a quasiconformally rigid point in $\partial AH(M)$ then ρ is uniquely approachable. In particular, $AH(M)$ is locally connected at ρ.*

If $M = S \times I$, then $\mathrm{int}(AH(S \times I))$ is the *quasi-Fuchsian locus*, denoted $QF(S)$, and is naturally identified with $\mathcal{T}(S) \times \mathcal{T}(S)$. Given $Y \in \mathcal{T}(S)$, the *Bers slice* B_Y of $QF(S)$ is the slice $\mathcal{T}(S) \times \{Y\}$ in the product structure. If ρ lies in the boundary of a Bers slice B, then its conformal boundary always has a component homeomorphic to S (see Bers [8, Theorem 8]). In this setting, we say that ρ is *quasiconformally rigid in* ∂B if every other component of its conformal boundary is a thrice-punctured sphere. We say a Bers slice *self-bumps* at a point $\rho \in \partial B$ if there exists a neighborhood W of ρ in the closure $\overline{B}$ of B (within $AH(S \times I)$) such that if V is a neighborhood of ρ in $\overline{B}$ which is contained in W, then $V \cap B$ is disconnected.

Theorem 1.5 *Let B be a Bers slice of $QF(S)$ for some closed surface S. If $\rho \in \partial B$ and ρ is quasiconformally rigid in ∂B, then B does not self-bump at ρ. In particular, its closure $\overline{B}$ is locally connected at ρ.*

An important ingredient in the proofs of these results, which may be of independent interest, is developed in Section 5. Theorem 5.1 in this section provides a tool for controlling the interaction between Fenchel–Nielsen length-twist coordinates on Teichmüller space and the "rough coordinates" associated to curve-complex subsurface projections.

History The Ending Lamination Theorem [51; 17; 18] asserts that hyperbolic 3–manifolds in $AH(M)$ are classified by their (marked) homeomorphism type and ending invariants which encode the asymptotic geometry of their ends. As points in the interior are parametrized by Teichmüller space(s) and ending laminations are associated to points on the boundary, a tenuous analogy between deformation spaces and Thurston's compactification of Teichmüller spaces by the sphere of projective measured laminations clouded the picture of the topological structure of deformation spaces for many years. The noncontinuity of the action of the mapping class group on Bers compactification (see Kerckhoff and Thurston [40]), illustrated some initial failings of this analogy, and elucidated a central example of Jørgenson (see Marden [44]) concerning the disparity between algebraic and geometric convergence that underlies the present discussion.

Anderson and Canary [2] showed that the (marked) homeomorphism type need not vary continuously over $AH(M)$, while Brock [13] showed that ending laminations do not vary continuously in any of the usual topologies, even in the closure of a Bers slice. These results make it clear that the parametrization of $AH(M)$ must be much more complicated than one might naively hope.

Bumping phenomena in deformation spaces were first discovered by Anderson and Canary [2]. Anderson, Canary and McCullough [5] characterized exactly which components of $\mathrm{int}(AH(M))$ bump when M has incompressible boundary. McMullen [49] showed that $QF(S)$ self-bumps, while Bromberg and Holt [22] showed that every component of $\mathrm{int}(AH(M))$ self-bumps whenever M contains a primitive essential annulus. Bromberg [21] and Magid [43] showed that $AH(S\times I)$ is not locally connected. For a more complete overview of recent results on the pathology of the topology of $AH(M)$, see Canary [28].

All known bumping and self-bumping results make use of the "wrapping" construction from [2] which requires the presence of a primitive essential annulus. It is not yet known whether self-bumping can occur in $AH(M)$ when M does not contain primitive essential annuli or in the closure of a Bers slice. However, Bromberg [21] conjectures that if S is a closed surface of genus at least 2, then the closure of every Bers Slice of $QF(S)$ is not locally connected. In the case of Bers slices of the space of punctured torus groups, Minsky [50] showed that the closure of every Bers slice is a disk and hence locally connected. We conjecture, similarly, that $AH(M)$ is not locally connected whenever M has a boundary component of genus at least two.

Theorem 1.5 and the quasifuchsian case of Theorem 1.3 also appear in Ohshika [54].

Outline of the argument

In Section 3, we rule out bumping in the setting of Theorems 1.1 and 1.2. In each case, the point is to rule out change of marked homeomorphism type in a sequence approaching the point in question. The hypotheses allow for the key use of the core embedding results of Anderson, Canary, Culler and Shalen [4].

In Section 4, we rule out self-bumping in the setting of Theorem 1.1. By hypothesis, we consider a point ρ with no extra parabolics and some degenerate ends. To rule out self-bumping at ρ it suffices to consider two sequences $\{\rho_n\}$ and $\{\rho'_n\}$ in $\mathrm{int}(AH(M))$ converging to ρ, and show that they can be connected by a sequence of paths $\{\gamma_n\}$, also in int($AH(M)$), which accumulate only on ρ. Nonbumping implies that ρ_n and ρ'_n are quasiconformally conjugate, so the paths can be chosen as Teichmüller geodesics in the associated quasiconformal deformation space. We can control the behavior of the ending invariants of these sequences, and use the Ending Lamination Theorem to show that any accumulation point of these paths is ρ.

The proof of Theorem 1.5 (the Bers slice case) is given in Section 7, using results from Sections 5 and 6. For clarity, consider first the case of a point $\rho \in \partial B$ which is a *maximal cusp*; that is, where a maximal curve system $\boldsymbol{\alpha}$ on the base surface S is represented by parabolics.

There is a neighborhood basis of ρ in B consisting of sets of the form

$$U(\delta) = \{\rho' \in B : l_{\alpha_j}(\rho') < \delta \ \forall \alpha_j \in \boldsymbol{\alpha}\},$$

where $l_{\alpha_j}(\rho')$ is the translation distance in hyperbolic space of $\rho'(\alpha_j)$, for a component α_j of $\boldsymbol{\alpha}$. To show no self-bumping occurs at ρ, then we must show that for any $\epsilon > 0$ there is a $\delta > 0$ such that any two points in $U(\delta)$ can be joined by a path in $U(\epsilon)$.

To show this would be straightforward if all components of $\boldsymbol{\alpha}$ were already short on the top conformal boundary of our group (the one that varies in the Bers slice): Fenchel–Nielsen coordinates for the Teichmüller space of the top conformal boundary component can be used directly to obtain a path in which the lengths of components of $\boldsymbol{\alpha}$ are controlled.

In general, however, curves in $\boldsymbol{\alpha}$ can have very short geodesic representatives deep inside the convex core of the manifold, while on the boundary they are extremely long. To obtain geometric control over the interior of the convex core via boundary geometry requires tools from the solution of the Ending Lamination Conjecture in [51; 17]. Lemma 6.1 gives the statement needed, namely that when the geodesic representatives of $\boldsymbol{\alpha}$ are very short in the manifold, there is a continuous path in B terminating at a point where $\boldsymbol{\alpha}$ is short in the conformal boundary such that the geodesic representatives of $\boldsymbol{\alpha}$ are short in all the corresponding hyperbolic manifolds along the deformation.

With Lemma 6.1 in hand, starting with two points in $U(\delta)$ we can connect them within $U(\epsilon)$ to a smaller neighborhood in which $\boldsymbol{\alpha}$ is short on the conformal boundary, and there to each other. This is carried out in Section 7.1, completing Theorem 1.5 in the case of maximal cusps.

We develop the necessary machinery for the proof of Lemma 6.1 in Sections 5 and 6. Recall first from [51] that short length for a curve γ in a surface group corresponds to *large projection coefficients* for some subsurface W with $\gamma \subset \partial W$. That is, for each subsurface W we project the ending invariants of the group to the curve complex $\mathcal{C}(W)$

and measure the distance between them. Then a curve $\gamma \in \mathcal{C}(S)$ is short in the hyperbolic 3–manifold if and only if it is either short in the conformal boundary or one of these coefficients is large for a subsurface with γ in its boundary (see Theorem 2.2 for a precise statement).

In Section 5 we examine Fenchel–Nielsen coordinates and their effect on subsurface projections. In particular we prove in Theorem 5.1 that, given a curve system $\boldsymbol{\alpha}$ and a point X in Teichmüller space, we can deform the length and twist parameters of X associated to $\boldsymbol{\alpha}$ as much as we want without changing by more than a bounded amount projections to subsurfaces disjoint from $\boldsymbol{\alpha}$.

In Section 6 we perform the deformation in $AH(S)$. The trickiest issue is that we must adjust the components of the curve system in an order reflecting their arrangement in the manifold, with the curves "closest" to the top boundary being adjusted first. In particular, to each component α_i of $\boldsymbol{\alpha}$ we associate a subsurface W_i with α_i in its boundary, whose projection coefficient is large enough to be responsible for α_i being short. To each W_i is associated a certain geometric region in the manifold, and these regions are partially ordered in terms of their separation properties in the manifold. In order not to disturb the projection coefficients of the other surfaces while adjusting each α_i, we need to start with the highest ones.

In practice we detect this partial order in a combinatorial way, by looking at the projections of the subsurface boundaries to each other's curve complexes. These ideas come from Masur and Minsky [46] and Brock, Canary and Minsky [17], and are also exploited by Behrstock, Kleiner, Minsky and Mosher [6] and elsewhere. The details of this are discussed in Section 2.4, in particular, Lemma 2.3.

In the general case of Theorem 1.5, handled in Section 7.2, we must consider a representation ρ with a mix of parabolics (a nonmaximal system $\boldsymbol{\alpha}$) and degenerate ends. By the Ending Lamination Theorem such representations are uniquely determined by their ending invariants, and we can determine a neighborhood system for ρ by considering constraints not just on the lengths of the curves in $\boldsymbol{\alpha}$ but on the projections of the ending data to the subsurfaces associated to the degenerate ends. The appropriate statement is given in Lemma 7.1, which relies on Theorem 2.7, whose proof will appear in [16].

The acylindrical case of Theorem 1.3 is handled in Section 8. This is quite similar to the Bers slice case, with Thurston's Bounded Image Theorem providing control on the lower conformal boundary of each boundary subgroup.

Finally, the general surface group case of Theorem 1.3 is completed in Section 9. In this case parabolics and degenerate ends can occur on both top and bottom. We deform one end and then the other, taking care to preserve the order of the ends (and, in particular, the order of the curves becoming parabolic).

Acknowledgements The authors gratefully acknowledge the support of the National Science Foundation and the support of their NSF FRG grant in particular. We also thank Francis Bonahon for suggesting an initial form of the argument for Theorem 1.1 and the referee for many useful comments and corrections.

Bromberg was partially supported by NSF grants DMS-0554569 and DMS-0504877, Brock was partially supported by DMS-0553694 and DMS-0505442, Canary was partially supported by DMS-0504791 and DMS-0554239 and Minsky was partially supported by DMS-0504019 and DMS-05504321.

2 Background

In this section, we recall some of the key tools and results which will be used in the paper. (A few new technical lemmas will be derived in Sections 2.4 and 2.5).

In Section 2.1 we survey the Ending Lamination Theorem which provides a classification of hyperbolic 3–manifolds with finitely generated fundamental group in terms of their ending invariants. In Section 2.2, we recall basic facts about deformation spaces of hyperbolic 3–manifolds, for example the parametrization of the interior of $AH(M)$ and Thurston's Bounded Image Theorem. In Section 2.3, we recall results which explain how the internal geometry of hyperbolic 3–manifolds can be detected from its ending invariants, via subsurface projections. In Section 2.4, we introduce the partial order on (certain) subsurfaces discussed in the outline of argument and relate it to the ordering of curves in the hyperbolic 3–manifold. In Section 2.5, we recall basic facts about geometric limits and derive consequences of the core embedding results of [4].

2.1 Ending invariants and the Ending Lamination Theorem

We recall (see Benedetti and Petronio [7] for example) that there exists a Margulis constant $\mu > 0$, such that if $\epsilon < \mu$ and

$$N_{\mathrm{thin}(\epsilon)} = \{x \in N \mid \mathrm{inj}_N(x) < \epsilon\},$$

then every component of $N_{\mathrm{thin}(\epsilon)}$ is either a solid torus, which is a metric neighborhood of a closed geodesic in N or a "cusp", which is a quotient of a horoball in $\mathbf{H}^3$ by a group of parabolic transformations. Each cusp is homeomorphic to $T \times (0,\infty)$ where T is either a torus or an open annulus. We pick a uniform $\epsilon_0 < \mu$ which will be used throughout the paper.

If $\rho \in AH(M)$, let $N_\rho = \mathbf{H}^3/\rho(\pi_1(M))$ and let N_ρ^0 be obtained from N_ρ by removing all the cusps of $(N_\rho)_{\mathrm{thin}(\epsilon_0)}$. A *compact core* for a hyperbolic 3–manifold N is a compact submanifold C such that the inclusion of C into N is a homotopy equivalence. A *relative compact core* M_ρ for N_ρ is a compact core for N_ρ^0 which intersects every component of ∂N_ρ^0 in a compact core for that component. (The existence of a relative compact core is due to Kulkarni and Shalen [42] and McCullough [48].) Let $P_\rho = M_\rho \cap \partial N_\rho^0$. There exists a well-defined, up to homotopy, homotopy equivalence $h_\rho\colon M \to M_\rho$ in the homotopy class determined by ρ, and a well-defined identification of the conformal boundary $\partial_c N_\rho$ with a collection of components of $\partial M_\rho - P_\rho$. The Tameness Theorem of Agol [1] and Calegari and Gabai [25] assures us that we may choose M_ρ so that $N_\rho^0 - M_\rho$ is homeomorphic to $(\partial M_\rho - P_\rho) \times (0, \infty)$.

If a component S of $\partial M_\rho - P_\rho$ is identified with a component of $\partial_c N_\rho$, it is called *geometrically finite* and inherits a natural conformal structure, regarded as a point in $\mathcal{T}(S)$. Otherwise, the component S is called *geometrically infinite* and it bounds a neighborhood of a geometrically infinite end. There exists a collection of simple closed curves $\{\alpha_i\}$ on S, whose geodesic representatives lie in the component of $N_\rho^0 - M_\rho$ bounded by S and leave every compact set. Regarded as a sequence of projective measured laminations, $\{\alpha_i\}$ converges to $\mu \in PL(S)$. The support λ of μ, regarded as a geodesic lamination, is called the *ending lamination* associated to S. The ending lamination λ lies in the set $\mathcal{EL}(S)$ of geodesic laminations admitting measures of full support which *fill the surface*: every component of their complement is a disk or a peripheral annulus. (See Thurston [62], Bonahon [11] and Canary [26] for a discussion of geometrically infinite ends and their ending laminations). The Ending Lamination Theorem (see Minsky [51] and Brock, Canary and Minsky [17; 18]) tells us that this information determines the manifold up to isometry.

Ending Lamination Theorem *Suppose that $\rho_1, \rho_2 \in AH(M)$, then $\rho_1 = \rho_2$ if and only if there exists an orientation-preserving homeomorphism of pairs $g\colon (M_{\rho_1}, P_{\rho_1}) \to (M_{\rho_2}, P_{\rho_2})$ such that*

1. *$g \circ h_{\rho_1}$ is homotopic to h_{ρ_2},*
2. *g is a conformal homeomorphism from the geometrically finite components of $\partial M_{\rho_1} - P_{\rho_1}$ to the geometrically finite components of $\partial M_{\rho_2} - P_{\rho_2}$, and*
3. *g takes the ending lamination of any geometrically infinite component of $\partial M_{\rho_1} - P_{\rho_1}$ to the ending lamination of the image geometrically infinite component of $\partial M_{\rho_2} - P_{\rho_2}$.*

2.2 Deformation spaces of hyperbolic 3–manifolds

We begin by reviewing the classical deformation theory of the interior of $AH(M)$. (See Section 7 of Canary and McCullough [31] for a complete treatment of this theory and its history.) Let $\mathcal{A}(M)$ denote the set of (marked) homeomorphism types of compact, oriented hyperbolizable 3–manifolds homotopy equivalent to M. We recall that $\mathcal{A}(M)$ is the set of pairs (M', h') where M' is an oriented, hyperbolizable compact 3–manifold and $h\colon M \to M'$ is a homotopy equivalence, where (M_1, h_1) and (M_2, h_2) are said to be equivalent if there exists an orientation-preserving homeomorphism $j\colon M_1 \to M_2$ such that $j \circ h_1$ is homotopic to h_2. We get a well-defined map

$$\Theta\colon AH(M) \to \mathcal{A}(M)$$

given by taking ρ to the equivalence class of (M_ρ, h_ρ). This map is surjective and the components of the interior of $AH(M)$ are exactly the preimages of points in $\mathcal{A}(M)$.

If M has incompressible boundary, equivalently if $\pi_1(M)$ is freely indecomposable, then points in $\Theta^{-1}(M', h') \cap \mathrm{int}(AH(M))$ give rise to well-defined conformal structures on $\partial_T M'$, where $\partial_T M'$ is the set of nontoroidal boundary components of $\partial M'$. Moreover, every possible conformal structure arises and the conformal structure determines the manifold. Therefore, we may identify the component $\Theta^{-1}(M', h') \cap \mathrm{int}(AH(M))$ with $\mathcal{T}(\partial_T M')$.

The Density Theorem asserts that $AH(M)$ is the closure of its interior. If M has incompressible boundary, the Density Theorem follows from the Ending Lamination Theorem [51; 17], Bonahon's Tameness Theorem [12] and convergence results of Thurston [65; 66] (see [17]). For the proof of the general case see Namazi and Souto [53] or Ohshika [58]. There is an alternate approach, using cone-manifold deformation theory, pioneered by Bromberg [20] and Brock and Bromberg [14] and completed by Bromberg and Souto [23].

The majority of this paper will be concerned with the case where $M = S \times I$ and S is a closed surface. In this case, $\mathcal{A}(S \times I)$ is a single point, and the interior $QF(S)$ of $AH(S \times I)$ (which is often abbreviated to $AH(S)$) is identified with $\mathcal{T}(S) \times \mathcal{T}(S)$. If $\rho \in AH(S)$, then the relative compact core M_ρ is identified with $S \times [0,1]$. (Here we are implicitly identifying $\mathcal{T}(\overline{S})$ with $\mathcal{T}(S)$ where $\overline{S}$ is S with the opposite orientation. Formally, the conformal structure on $S \times \{0\}$ lies in $\mathcal{T}(\overline{S})$.) The orientation on S allows us to identify one component $\partial_1 M_\rho$ as the top, or *upward pointing* component and the other component $\partial_0 M_\rho$ as the bottom or *downward pointing* component. If $\rho \in QF(S)$ has conformal structure X on $\partial_1 M_\rho$ and Y on $\partial_0 M_\rho$, we will use the notation $\rho = Q(X, Y)$. In general, $P \cap \partial_1 M_\rho$ may be identified with the regular neighborhood of a collection $\boldsymbol{\alpha}$ of simple closed curves on S and $P \cap \partial_0 M_\rho$ may also be identified with the regular neighborhood of a collection $\boldsymbol{\beta}$ of simple closed curves on S. We say that the components of $\boldsymbol{\alpha}$ are associated to upward-pointing cusps, while the components of $\boldsymbol{\beta}$ are associated to downward-pointing cusps. Similarly the components of $\partial_1 M_\rho \setminus P$ are said to bound upward-pointing ends, and the components of $\partial_0 M_\rho \setminus P$ are said to bound downward-pointing ends. If $\rho \in AH(S)$ is quasiconformally rigid, a component of $\partial M_\rho - P_\rho$ is geometrically finite if and only if it is a thrice-punctured sphere, while the remaining components each bound neighborhoods of degenerate ends and inherit an ending lamination.

We recall that a Bers slice B_Y of $QF(S)$ is a set of the form $\mathcal{T}(S) \times \{Y\}$ where $Y \in \mathcal{T}(S)$. If B_Y is a Bers slice and $\rho \in \overline{B_Y}$ (the closure of B_Y in $AH(S)$), then the bottom boundary component of M_ρ is geometrically finite and has conformal structure Y (see Bers [8, Theorem 8]). If ρ is quasiconformally rigid in B_Y, one then obtains a collection $\boldsymbol{\alpha}$ of curves on the top boundary component whose regular neighborhood is P_ρ, and an ending lamination on every upward-pointing component of $\partial M_\rho - P_\rho$ which is not a thrice-punctured sphere.

The other special case we will consider is when M is acylindrical. Johannson [36] showed that any homotopy equivalence from an acylindrical manifold to a compact 3–manifold is homotopic to a homeomorphism, so $\mathcal{A}(M)$ has two components (one associated to each possible orientation on M). So, int($AH(M)$) has two components and it follows from [5] that Θ is locally constant. Thurston [64] showed that $AH(M)$ is compact if M is acylindrical.

Our proof of Theorem 1.3 in the acylindrical case will make crucial use of Thurston's Bounded Image Theorem (see Kent [38] for a proof.) If B is a component of int($AH(M)$) then B is identified with $\mathcal{T}(\partial_T M)$. If S is a component of $\partial_T M$, then there is a natural map $r_S\colon B \to AH(S)$ given by restriction, whose image lies in $QF(S)$. If $\tau \in \mathcal{T}(\partial M)$, then $r_S(\tau)$ is a well-defined point $(\tau|_S, \sigma_S(\tau))$ where $\sigma_S(\tau) \in \mathcal{T}(S)$. Letting S vary over all components of $\partial_T M$, we get a well-defined map

$$\sigma\colon \mathcal{T}(\partial_T M) \to \mathcal{T}(\partial_T M)$$

called the skinning map. Thurston's Bounded Image Theorem simply asserts that σ has bounded image in $\mathcal{T}(\partial_T M)$.

2.3 The conformal boundary of a hyperbolic 3–manifold and its internal geometry

In this section, we review a variety of results which relate the geometry of the conformal boundary to the geometry of the hyperbolic 3–manifold. Most classically, a result of Bers [8] shows that lengths of curves in the conformal boundary provide upper bounds for lengths in the manifold. To set notation, if $\rho \in AH(M)$ and α is a (homotopically nontrivial) closed curve in M, then $l_\rho(\alpha)$ is the length of the geodesic representative α^*

of $h_\rho(\alpha)$ in N_ρ (with $l_\rho(\alpha) = 0$ if $h_\rho(\alpha)$ is homotopic into a cusp of N_ρ). Similarly, if $X \in \mathcal{T}(S)$ and α is a closed curve on X, then $l_X(\alpha)$ is the length of the geodesic representative of α on X.

Lemma 2.1 (Bers [8, Theorem 3]) *If $\rho = Q(X, Y) \in QF(S)$, then*

$$l_\rho(\alpha) \le 2 l_X(\alpha)$$

for any closed curve α on X.

Subsurface projections and the curve complex The proof of the Ending Lamination Theorem develops more sophisticated information about the relationship between the geometry of a hyperbolic 3–manifold and its ending invariants. This information is typically expressed in terms of projections onto curve complexes of subsurfaces of the boundary.

Recall from [47] the curve complexes $\mathcal{C}(W)$ where $W \subseteq S$ is an essential subsurface. When W is not an annulus, the vertices of $\mathcal{C}(W)$ are homotopy classes of simple closed nonperipheral curves in W. When W is an annulus, vertices are homotopy classes rel endpoints of arcs connecting the boundaries of the compactified annulus cover $\widehat{W} \to S$ associated to W. Edges in these complexes correspond to pairs of vertices with representatives that intersect in the minimal possible number of points allowed by W. $\mathcal{C}(W)$ is endowed with the path metric $d_{\mathcal{C}(W)}$ assigning length 1 to each edge. If W is a three-holed sphere then $\mathcal{C}(W)$ is empty, and from now on we implicitly ignore this case.

If $\mathcal{C}(S, W)$ denotes the set of curves in S which intersect W essentially, we have, also as in [47], *subsurface projection maps*

$$\pi_W \colon \mathcal{C}(S, W) \to \mathcal{C}(W).$$

If W is not an annulus then $\pi_W(\alpha)$ is obtained by selecting (any) arc of the essential intersection of α with W, and doing surgery with ∂W to obtain a closed curve. When W is an annulus we take more care: we consider the annular cover $\widehat{W}$ of S associated to W and lift α to an arc connecting the two boundaries. All the choices involved in these constructions differ by bounded distance in the image, and in our applications this ambiguity will not matter. Define, for $\alpha, \beta \in \mathcal{C}(S, W)$,

$$d_W(\alpha, \beta) = d_{\mathcal{C}(W)}(\pi_W(\alpha), \pi_W(\beta)).$$

All of these notions can be applied to points in $\mathcal{T}(S)$ as well, giving a map

$$\pi_W \colon \mathcal{T}(S) \to \mathcal{C}(W)$$

defined as follows: Given $X \in \mathcal{T}(S)$ let α be a curve of minimal length in X intersecting W essentially and let $\pi_W(X) = \pi_W(\alpha)$. Except when W is an annulus (or a three-holed sphere, which we always exclude), the length of α has a uniform upper bound known as the Bers constant of S (see [10]). Indeed the shortest maximal curve system has a uniform upper length bound, and one of those curves must intersect W. Any nonuniqueness in the choice of α leads to values for $\pi_W(X)$ that differ by a uniformly bounded amount.

If W is an annulus whose core γ has extremely short length in X, then the shortest curve crossing γ will be long; however, the ambiguity in the definition of π_W will still be uniformly bounded. To see this, note that if two curves β_1 and β_2 crossing γ have projections with distance greater than 2 in $\mathcal{C}(\gamma)$, then there exists a pair of arcs b_1 and b_2 in β_1 and β_2 respectively with common endpoints whose concatenation is homotopic into γ. Exchange of these arcs, and smoothing, will strictly shorten at least one of β_1 or β_2, so they cannot both have minimal length in X. (The same argument actually works for nonannular W as well).

Lengths in Kleinian surface groups In the case of a quasifuchsian hyperbolic manifold $Q(X,Y)$, a curve is short if and only if it is either short in the conformal boundary or there is a subsurface with the curve in its boundary such that $d_W(X,Y)$ is large. To be more explicit, given a simple closed curve γ in S and $X, Y \in \mathcal{T}(S)$, we define

$$\mathbf{m}_\gamma(X,Y) = \max\left(\frac{1}{l_\gamma(X)}, \frac{1}{l_\gamma(Y)}, \sup_{\gamma\subset\partial W} d_W(X,Y)\right).$$

The supremum is over all essential subsurfaces in S whose boundary contains a curve parallel to γ. The following theorem is a restatement (and special case) of the Length Bound Theorem from Brock, Canary and Minsky [17].

Theorem 2.2 *Given $\epsilon > 0$ there exists M such that, for any $Q(X,Y) \in QF(S)$, and simple closed curve γ in S,*

$$\mathbf{m}_\gamma(X,Y) > M \implies l_\gamma(Q(X,Y)) < \epsilon.$$

Conversely, given M' there exists $\epsilon' > 0$ such that

$$l_\gamma(Q(X,Y)) < \epsilon' \implies \mathbf{m}_\gamma(X,Y) > M'.$$

2.4 Partial orders

In view of Theorem 2.2, those subsurfaces W where $d_W(X,Y)$ is large are important because their boundaries correspond to short curves in $Q(X,Y)$. If the curves are sufficiently short then Otal [60] shows that their associated Margulis tubes are *unlinked*, meaning they are isotopic to level curves in a product structure on $Q(X,Y)$, and hence admit a partial order.

If $\alpha,\beta\in\mathcal{C}(S)$ and $i(\alpha,\beta)\neq 0$, then we say that α *lies above* β, and that β *lies below* α, in $N_\rho\in AH(S)$ if their geodesic representatives α^* and β^* are disjoint and α^* may be homotoped to $+\infty$ in the complement of β^* (that is, there is a proper map $F\colon S^1\times[0,\infty)\to N_\rho$ such that $F|_{S^1\times\{0\}}=\alpha^*$, $\beta^*\cap F(S^1\times[0,\infty))=\varnothing$, and $F(S^1\times\{t\})$ is a family of curves exiting the upward-pointing end of N_ρ). If $l_\rho(\alpha)=0$, then α lies above β if α is associated to an upward-pointing cusp. See Section 3.1 of [17] for further discussion of this topological partial order.

There is a closely related combinatorial partial order, which originates in the "hierarchy path" construction of [47].

For (X,Y) an ordered pair of points in Teichmüller space and $c>0$, define the following collection of (isotopy classes of) essential subsurfaces of S:

$$\mathcal{L}_c(X,Y)=\{W\subset S : d_W(X,Y)>c\}.$$

We say two subsurfaces or curves in S *overlap* if they intersect essentially and neither is contained in the other.

The following lemma can be extracted from Lemmas 4.18, 6.1 and 6.2 of [47] (see also [6, Section 4.1]).

Lemma 2.3 *There is a constant m_1 such that, if $c>m_1$ then $\mathcal{L}_c\equiv\mathcal{L}_c(X,Y)$ admits a partial order $\prec$, such that any $U,V\in\mathcal{L}_c$ which overlap are ordered, and $U\prec V$ implies that*

(1) $d_U(\partial V,X)\le m_1$,

(2) $d_U(\partial V,Y)>c-m_1$,

(3) $d_V(Y,\partial U)\le m_1$, *and*

(4) $d_V(\partial U,X)>c-m_1$.

Moreover, if $U\in\mathcal{L}_c(X,Y)$, $c>2m_1$, V overlaps U, and $d_V(X,\partial U)>m_1$, then

(5) $d_V(X,Y)>c-m_1$ *and*

(6) $U\prec V$

with respect to the order on $\mathcal{L}_{c-m_1}(X,Y)$.

One way to make sense of these inequalities is to interpret a large value for $d_U(\partial V, X)$ to mean that U is "between" V and X. In [47] this had a literal meaning, because a large value for $d_U(\partial V, X)$ meant that any hierarchy path connecting a bounded-length marking on X to a marking containing ∂V would have to pass through markings containing ∂U.

Thus, informally (2) says that U is between V and Y, but (1) says that U is *not* between V and X, and so on. Together these inequalities say that, in "traveling" from Y to X, we must first pass through U and then through V.

Theorem 2.2 implies that subsurfaces in $\mathcal{L}_c(X,Y)$, for suitable c, have short boundary curves in $Q(X,Y)$, and therefore are topologically ordered as above. Lemmas 2.2 and 4.10 and the Bilipschitz Model Theorem from [17] combine to show that, indeed, the partial order $\prec$ determines the topological ordering of the boundary components of the subsurfaces when c is large. In particular the combinatorial notion of "betweenness" translates to a topological statement, that in a suitable product structure on the manifold, one level surface lies at a height between two others. The following statement will suffice for us:

Lemma 2.4 *There exists $c_0 > m_1$ such that if $c > c_0$, $U, V \in \mathcal{L}_c(X,Y)$, and $U \prec V$, then if a boundary component α of U overlaps a boundary component β of V, then α lies below β in $Q(X,Y)$.*

It is a simple observation that a curve α which is short in the top conformal boundary lies above any curve β which is short in the manifold, if $i(\alpha,\beta) > 0$.

Lemma 2.5 *If $l_\alpha(X) < \epsilon_0$ and $l_\beta(Q(X,Y)) < \epsilon_0$ and α and β overlap then α lies above β in $Q(X,Y)$. Similarly, if $l_\beta(Y) < \epsilon_0$ and $l_\alpha(Q(X,Y)) < \epsilon_0$ and α and β intersect, then α lies above β in $Q(X,Y)$.*

Proof We give the proof in the case that $l_\alpha(X) < \epsilon_0$. A result of Epstein, Marden and Markovic [33, Theorem 3.1] implies that α has length at most $2\epsilon_0$ in the top boundary component of the convex core of $Q(X,Y)$. Therefore, one may isotope the geodesic representative of α onto the top boundary component of the convex core entirely within the Margulis tube of α. One may then isotope it to $+\infty$ in the complement of the convex core. The geodesic representative β^* of β is contained in the convex core, and since it has length less than ϵ_0 it is contained in its own Margulis tube which is disjoint from that of α. It follows that the homotopy does not intersect β^*. □

The lemma below will be used in the proof of Theorem 1.3 in the surface group case to control the impact of changing the top conformal structure on the ordering of the short curves and on related features. It is really just a repackaging of the preceding sequence of lemmas. It says that, if α is known to be short in $Q(X,Y)$, and Z is "between" α and the top conformal structure X in the combinatorial sense discussed above, then indeed ∂Z is also short in $Q(X,Y)$, and each of its components that overlap α are topologically ordered above it.

Lemma 2.6 *There exists $d_0 > 0$ and $\delta_0 \in (0,\epsilon_0)$ such that, if $l_\alpha(Q(X,Y)) < \delta_0$, $\alpha \in \mathcal{C}(S)$ overlaps Z and $d_Z(X,\alpha) > d_0$, then $l_{\partial Z}(Q(X,Y)) < \delta_0$ and each component of ∂Z which overlaps α lies above α.*

Proof Applying Theorem 2.2 we may choose $\delta_0 \in (0,\epsilon_0)$ so that

$$l_\alpha(Q(X,Y)) < \delta_0 \implies \mathbf{m}_\alpha(X,Y) > \max\{c_0, 2m_1\}.$$

Applying the other direction of Theorem 2.2, we choose $d_0 > c_0 + m_1 + 2$ so that if $W \subset S$, then

$$d_W(X,Y) > d_0 - m_1 - 2 \implies l_{\partial W}(Q(X,Y)) < \delta_0.$$

We note that if $l_\alpha(X) < \delta_0$, then $d_Z(X,\alpha) \le 2 < d_0$, so we may assume that $l_\alpha(X) \ge \delta_0$.

If $l_\alpha(Y) < \delta_0$, then Lemma 2.5 implies that each component of ∂Z which overlaps α lies above α. Moreover, $d_Z(Y,\alpha) \le 2$, so

$$d_Z(X,Y) \ge d_Z(X,\alpha) - d_Z(\alpha,Y) > d_0 - 2$$

so $l_{\partial Z}(Q(X,Y)) < \delta_0$. This completes the proof in this case.

Hence we can now assume $l_\alpha(Y) \ge \delta_0$. Now $\mathbf{m}_\alpha(X,Y) > 2m_1$ implies that there exists an essential subsurface $W \subset S$ with $\alpha \subset \partial W$ such that $d_W(X,Y) > 2m_1$. Since $d_Z(X,\partial W) > d_0 - 1 > m_1$ and $d_W(X,Y) > 2m_1$, Lemma 2.3(6) implies that $W \prec Z$ in $\mathcal{L}_{c-m_1}(X,Y)$. Lemma 2.3(3) implies that $d_Z(Y,\partial W) \le m_1$. Therefore,

$$d_Z(X,Y) \ge d_Z(X,\partial W) - d_Z(\partial W,Y) > d_0 - 1 - m_1,$$

so $l_{\partial Z}(Q(X,Y)) < \delta_0$. Lemma 2.4 then implies that each component of ∂Z which overlaps α lies above α. □

Predicting geometrically infinite ends in an algebraic limit Geometrically infinite surfaces in the algebraic limit can be detected by looking at the limiting behavior of the ending invariants. Recall that Masur and Minsky [46] proved that if W is an essential subsurface of S, then $\mathcal{C}(W)$ is Gromov hyperbolic and Klarreich [41] (see also Hamenstadt [35]) proved that if W is not an annulus or pair of pants, then its Gromov boundary $\partial_\infty\mathcal{C}(W)$ is identified with $\mathcal{EL}(W)$.

Theorem 2.7 [16] *Let $\{\rho_n\}$ be a sequence in $AH(S)$ converging to ρ such that the top ending invariant of ρ_n is $X_n \in \mathcal{T}(S)$. If W is an essential subsurface of S, the following statements are equivalent:*

(1) *N_ρ^0 has an upward-pointing end bounded by W with ending lamination $\lambda \in \mathcal{EL}(W)$.*

(2) *$\{\pi_W(X_n)\}$ converges to λ.*

Moreover, if $\{\rho_n = Q(X_n, Y_n)\}$, then $\pi_W(Y_n)$ does not accumulate at λ if $\{\pi_W(X_n)\}$ converges to λ.

Similarly, we obtain an equivalence if upward is replaced by downward and the roles of X_n and Y_n are interchanged.

A key tool in the proof of Theorem 2.7 is the fact that for any nonannular subsurface W the set of bounded length curves in $Q(X,Y)$ project to a set of curves in $\mathcal{C}(W)$ which are a bounded Hausdorff distance from any geodesic in $\mathcal{C}(W)$ joining $\pi_W(X)$ to $\pi_W(Y)$. This result will be itself used in the proof of Lemma 9.2. We state the results in the special case of quasifuchsian groups.

Theorem 2.8 ([16]) *Given S, there exists $L_0 > 0$ such that for all $L \geq L_0$, there exists D_0, such that, if $X, Y \in \mathcal{T}(S)$, $\rho = Q(X,Y)$, $W \subset S$ is an essential subsurface and*

$$C(\rho, L) = \{\alpha \in \mathcal{C}(S) : l_\alpha(\rho) < L\},$$

then $\pi_W(C(\rho,L) \cap \mathcal{C}(S,W))$ has Hausdorff distance at most D_0 from any geodesic in $\mathcal{C}(W)$ joining $\pi_W(X)$ to $\pi_W(Y)$. Moreover if $d_W(X,Y) > D_0$ then

$$C(W, \rho, L) = \{\alpha \in \mathcal{C}(W) : l_\alpha(\rho) < L\}$$

is nonempty and also has Hausdorff distance at most D_0 from any geodesic in $\mathcal{C}(W)$ joining $\pi_W(X)$ to $\pi_W(Y)$.

2.5 Geometric limits

A sequence $\{\Gamma_n\}$ of torsion-free Kleinian groups *converges geometrically* to a torsion-free Kleinian group Γ if Γ is the set of all accumulation points of sequences of elements $\{\gamma_n \in \Gamma_n\}$ and every $\gamma \in \Gamma$ is a limit of a sequence of elements $\{\gamma_n \in \Gamma_n\}$; or in other words if $\{\Gamma_n\}$ converges to Γ in the Chaubaty topology on closed subsets of $\mathrm{Isom}_+(\mathbf{H}^3)$. One may equivalently express this in terms of Gromov convergence of the quotient hyperbolic 3–manifolds (see [29; 7]). If $N_n = \mathbf{H}^3/\Gamma_n$ and $N = \mathbf{H}^3/\Gamma$ and v_n and v_0 denote the projections of a fixed orthonormal base frame based at the origin for $\mathbf{H}^3$, then $\{\Gamma_n\}$ converges geometrically to Γ if and only if there exists a nested sequence of compact submanifolds $\{X_n\}$ of N which exhaust N and K_n–bilipschitz diffeomorphisms $f_n\colon X_n \to Y_n$ onto submanifolds of N_n such that $df_n(v_0) = v_n$, $\lim K_n = 1$ and f_n converges uniformly on compact subsets of N to an isometry (in the C^∞–topology).

Lemma 3.6 and Proposition 3.8 of Jørgensen and Marden [37] guarantee that if $\{\rho_n\}$ is a sequence in $AH(M)$ converging to ρ, then there is a subsequence of $\{\rho_n(\pi_1(M))\}$ which converges geometrically to a torsion-free Kleinian group $\widehat{\Gamma}$ such that $\rho(\pi_1(M)) \subset \widehat{\Gamma}$.

We say that a sequence $\{\rho_n\}$ in $AH(M)$ converges *strongly* to $\rho \in AH(M)$ if it converges in $AH(M)$ and $\{\rho_n(\pi_1(M))\}$ converges geometrically to $\rho(\pi_1(M))$. One may combine work of Anderson and Canary with the recent resolution of Marden's Tameness Conjecture to show that in the absence of unnecessary parabolics, algebraic convergence implies strong convergence (see also Theorem 1.2 of Brock and Souto [19]).

Theorem 2.9 *Let M be a compact 3–manifold and let $\{\rho_n\}$ be a sequence in $AH(M)$ converging to ρ in $AH(M)$. If every parabolic element of $\rho(\pi_1(M))$ lies in a rank two free abelian subgroup, then $\{\rho_n\}$ converges strongly to ρ.*

Proof Theorem 3.1 of Anderson and Canary [3] and Theorem 9.2 of [27] together imply that if ρ is topologically tame, then $\{\rho_n\}$ converges strongly to ρ. The Tameness Theorem of Agol [1] and Calegari and Gabai [25] assures that ρ is topologically tame, so our convergence is indeed strong. □

Proposition 3.2 of Anderson, Canary, Culler and Shalen [4] shows that whenever the algebraic limit is a maximal cusp (ie geometrically finite and quasiconformally rigid), then the convex core of the algebraic limit embeds in the geometric limit. Remark 3.3 points out that the same argument applies whenever the algebraic limit is topologically tame and its convex core has totally geodesic boundary. In particular, the result holds when the limit is quasiconformally rigid.

Proposition 2.10 *If ρ is a quasiconformally rigid point in $\partial AH(M)$ and $\{\rho_n\}$ converges algebraically to ρ and $\{\rho_n(\pi_1(M))\}$ converges geometrically to $\widehat{\Gamma}$, then the convex core of N_ρ embeds in $\widehat{N} = \mathbb{H}^3/\widehat{\Gamma}$ under the obvious covering map.*

Proposition 2.10 will be used in Section 3 to rule out bumping at quasiconformally rigid points. We will also use it to control the relative placement of closed curves in manifolds algebraically near to a quasiconformally rigid manifold. Lemma 2.11 will only be needed in the quasifuchsian case discussed in Section 9.

Lemma 2.11 *If $\rho \in AH(S)$ is quasiconformally rigid, α is an upward-pointing cusp in N_ρ and β is a downward-pointing cusp in N_ρ, and α and β intersect in S, then there exists a neighborhood U of ρ in $AH(S)$ such that if $\rho' \in U$, then α lies above $\beta \in N_{\rho'}$.*

Proof Find an embedded surface F in $C(N_\rho)$ which is a compact core for $C(N_\rho)$. Let $\epsilon < \epsilon_0$ be a lower bound for the injectivity radius of N on F. Let A be an embedded annulus in $C(N_\rho)$, intersecting F only in one boundary component and whose other boundary component is curve in the homotopy class of α with length at most $\epsilon/4$. Let B be an embedded annulus in $C(N_\rho)$, intersecting F only in one boundary component and whose other boundary component is curve in the homotopy class of α with length at most $\epsilon/4$.

If the lemma fails we may produce a sequence $\{\rho_n\}$ converging to ρ such that α does not lie above β in any N_{ρ_n}. We may again pass to a subsequence such that $\{\rho_n(\pi_1(M))\}$ converges geometrically to $\widehat{\Gamma}$ and $\rho(\pi_1(M)) \subset \widehat{\Gamma}$. Let $\widehat{N} = \mathbf{H}^3/\widehat{\Gamma}$ and let $\pi\colon N_\rho \to \widehat{N}$ be the natural covering map.

By Proposition 2.10, π embeds $C = F \cup A \cup B$ in $\widehat{N}$. Then, for all large enough n, one can pull C back to $C_n = F_n \cup A_n \cup B_n$ by an orientation-preserving 2–bilipschitz map and F_n is a compact core for N_{ρ_n} (as in the proof of Proposition 3.3 in Canary and Minsky [32]). One may join the geodesic representative of α in N_{ρ_n} to $\partial A_n - F_n$ by an annulus contained entirely within the $\epsilon/2$–Margulis tube associated to α. It follows that this annulus cannot intersect F_n (since F_n is contained entirely in the $\epsilon/2$–thick part of N_{ρ_n}) so we see that the geodesic representative of α in N_{ρ_n} lies above F_n. Similarly, the geodesic representative of β in N_{ρ_n} lies below F_n. Therefore, for sufficiently large n, α lies above β in N_{ρ_n}. This contradiction establishes the result. □

3 Ruling out bumping

In this section, we will show that there is no bumping at points with no unnecessary parabolics or at quasiconformally rigid points. The first case gives the nonbumping portion of Theorem 1.1, while the second case is Theorem 1.2. In each case, we do so by showing that the (marked) homeomorphism type is locally constant at ρ, which immediately implies that there is no bumping at ρ. Note that in this section it will never be necessary to assume that M has incompressible boundary.

The case where ρ contains no unnecessary parabolics is especially easy, since any sequence converging algebraically to ρ converges strongly.

Proposition 3.1 *Let M be a compact 3–manifold and $\rho \in \partial AH(M)$. If every parabolic element of $\rho(\pi_1(M))$ lies in a rank two free abelian subgroup, then Θ is locally constant at ρ. In particular, there is no bumping at ρ.*

Proof Let $\{\rho_n\}$ be a sequence in $AH(M)$ which converges to ρ. Theorem 2.9 implies that $\{\rho_n\}$ converges strongly to ρ. Results of Canary and Minsky [32] and Ohshika [57], then imply that for all large enough n there exists a homeomorphism $h_n\colon N_\rho \to N_{\rho_n}$ in the homotopy class determined by $\rho_n \circ \rho^{-1}$. It follows that $\Theta(\rho_n) = \Theta(\rho)$ for all large enough n, which completes the proof. □

Remark If we assume that $\{\rho_n\} \subset \mathrm{int}(AH(M))$, then strong convergence follows immediately from Theorem 1.2 of Brock and Souto [19]. Consideration of this case would suffice to establish that there is no bumping at ρ.

If ρ is a quasiconformally rigid point in $\partial AH(M)$, then sequences of representations converging to ρ need not converge strongly. However, by Proposition 2.10, the convex core of N_ρ embeds in the geometric limit of any sequence in $AH(M)$ converging to ρ, which will suffice to complete the proof. Proposition 3.2 immediately implies Theorem 1.2

Proposition 3.2 *If M is a compact 3–manifold and $\rho \in \partial AH(M)$ is quasiconformally rigid, then Θ is locally constant at ρ. In particular, there is no bumping at ρ.*

Proof If Θ is not locally constant, then there is a sequence $\{\rho_n\}$ such that $\Theta(\rho_n) \neq \Theta(\rho)$ for all n. We may pass to a subsequence, still called $\{\rho_n\}$, such that $\{\rho_n(\pi_1(M))\}$ converges geometrically to $\widehat{\Gamma}$ and $\rho(\pi_1(M)) \subset \widehat{\Gamma}$. Let $\widehat{N} = \mathbf{H}^3/\widehat{\Gamma}$ and let $\pi\colon N_\rho \to \widehat{N}$ be the natural covering map. Proposition 2.10 implies that π embeds the convex core $C(N_\rho)$ into $\widehat{N}$.

Let C be a compact core for $C(N_\rho)$. We recall that for all sufficiently large n, there exists a K_n–bilipschitz diffeomorphism $f_n\colon X_n \to N_{\rho_n}$ from a compact submanifold X_n of $\widehat{N}$ which contains $\pi(C)$ onto a compact submanifold of N_{ρ_n}. The arguments of Proposition 3.3 of Canary and Minsky [32] go through directly to show that, again for large enough n, $C_n = f_n(\pi(C))$ is a compact core for N_{ρ_n}. Moreover, $(f_n \circ \pi)_*\colon \pi_1(C) \to \pi_1(C_n)$ is the same isomorphism, up to conjugacy, induced by $\rho_n \circ \rho^{-1}$. It follows that $\Theta(\rho_n) = [(C_n, h_{\rho_n})] = [(C, h_\rho)] = \Theta(\rho)$. □

4 Ruling out self-bumping in the absence of parabolics

In this section we rule out self-bumping at points in $\partial AH(M)$ with no unnecessary parabolics when M has incompressible boundary. Propositions 3.1 and 4.1 combine to establish Theorem 1.1.

Proposition 4.1 *Let M be a compact 3–manifold with incompressible boundary and $\rho \in \partial AH(M)$. If every parabolic element of $\rho(\pi_1(M))$ lies in a rank two free abelian subgroup, then there is no self-bumping at ρ.*

Proof Let M_ρ be a relative compact core for N_ρ^0 and let $\{S_1, \dots, S_r\}$ denote the nontoroidal components of ∂M_ρ. We may order the boundary components so that $\{S_1, \dots, S_k\}$ correspond to geometrically finite ends of N_ρ^0 while $\{S_{k+1}, \dots, S_r\}$ correspond to geometrically infinite ends of N_ρ^0. Let $\{\tau_1, \dots, \tau_k, \lambda_{k+1}, \dots \lambda_r\}$ be the end invariants of ρ where $\tau_i \in \mathcal{T}(S_i)$ for all $i \le k$ and $\lambda_i \in \mathcal{EL}(S_i)$ for all $i > k$.

Let B be the component of int$(AH(M))$ corresponding to $[(M_\rho, h_\rho)]$. Since Θ is locally constant at ρ, by Proposition 3.1, B is the only component of int$(AH(M))$ containing ρ in its closure. We may identify B with $\mathcal{T}(S_1) \times \cdots \times \mathcal{T}(S_r)$. Let $\{\rho_n = (\tau_1^n, \dots, \tau_r^n)\}$ be a sequence in B converging to ρ. Theorem 2.7 implies that $\{\pi_{S_i}(\tau_i^n)\} \subset \mathcal{C}(S_i)$ converges to $\lambda_k \in \partial_\infty \mathcal{C}(S_i)$ for all $i > k$. Theorem 2.9 implies that $\{\rho_n\}$ converges strongly to ρ. Then, a result of Ohshika [56] (see also Kerckhoff and Thurston [40, Corollary 2.2]) implies that $\{\tau_i^n\}$ converges to τ_i for all $i \le k$.

Let $\{\rho_n = (\tau_1^n, \dots, \tau_r^n)\}$ and $\{\rho_n' = ((\tau_1^n)', \dots, (\tau_r^n)')\}$ be two sequences in B converging to ρ. In order to rule out self-bumping at ρ, it suffices to construct paths γ_n in B joining ρ_n to ρ_n' such that if $\nu_n \in \gamma_n$, then $\{\nu_n\}$ converges to ρ. We choose γ_n to be the Teichmüller geodesic in $\mathcal{T}(S_1) \times \cdots \times \mathcal{T}(S_r)$ joining ρ_n to ρ_n'. If $\{\nu_n = (\mu_1^n, \dots, \mu_r^n) \in \gamma_n\}$ is a sequence, then, for all $i \le k$, since both $\{\tau_i^n\}$ and $\{(\tau_i^n)'\}$ converge to τ_i, $\{\mu_i^n\}$ also converges to τ_i. In [46] (see Theorems 2.3 and 2.6), it is shown that a Teichmüller geodesic in $\mathcal{T}(S_i)$ projects into a c_2–neighborhood of a geodesic in $\mathcal{C}(S_i)$ (for some uniform choice of c_2). Therefore, since $\{\pi_{S_i}(\tau_i^n)\}$ and $\{\pi_{S_i}((\tau_i^n)')\}$ both converge to $\lambda_i \in \partial_\infty \mathcal{C}(S_i)$ for all $i > k$, we see that $\{\pi_{S_i}(\mu_i^n)\}$ converges to λ_i for all $i > k$.

If $M = S \times I$ for a closed surface S, then Thurston's Double Limit Theorem [65] implies that every subsequence of $\{\nu_n\}$ has a convergent subsequence. If M is not homeomorphic to $S \times I$, then the main result of Ohshika [55] (which is itself derived by combining results of Thurston [65; 66]) implies that every subsequence of $\{\nu_n\}$ has a convergent subsequence.

Let ν be a limit of a subsequence of $\{\nu_n\}$, still denoted $\{\nu_n\}$, in $AH(M)$. In order to complete the proof, it suffices to show that $\nu = \rho$. We do so by invoking the Ending Lamination Theorem. The main difficulty here is that we do not know that ν does not contain any unnecessary parabolics, so we cannot immediately conclude that $\{\nu_n\}$ converges strongly to ν.

Let $h\colon M_\rho \to M_\nu$ be a homotopy equivalence such that $h \circ h_\rho$ is homotopic to h_ν. Consider the sequence $\{\nu_n' = (\tau_1, \dots, \tau_k, \mu_{k+1}^n, \dots, \mu_r^n)\}$. There exists a sequence of K_n–quasiconformal map conjugating ν_n to ν'_n with $K_n \to 1$. It follows that $\{\nu'_n\}$ also converges to ν. Theorem 5 in Bers [8] implies that, for all $i \le k$, the sequence of components $\{\Omega_i^n\}$ of $\Omega(\nu'_n)$ associated to $\nu_n(\pi_1(S_i))$ (where we have chosen a fixed subgroup in the conjugacy class of subgroups associated to $\pi_1(S_i)$) converges in the sense of Caratheodory to a component Ω_i of $\Omega(\nu(\pi_1(S_i)))$ such that $\Omega_i/\nu(\pi_1(S_i))$ is homeomorphic to S_i with conformal structure τ_i. It then follows (again from Ohshika [56]) that Ω_i is a component of the domain of discontinuity of any geometric limit of $\{\nu'_n(\pi_1(M))\}$. Therefore, Ω_i is a component of $\Omega(\nu)$ and the stabilizer of Ω_i in $\nu(\pi_1(M))$ contains $\nu(\pi_1(S_i))$ as a finite index subgroup. Therefore, we may homotope h so that, for all $i \le k$, $h|_{S_i}$ is an orientation-preserving covering map of a component of ∂M_ν which is locally conformal.

If $i > k$, Theorem 2.7 implies that the cover $(N_\nu)_i$ of N_ν associated to $\pi_1(S_i)$ has a geometrically infinite end $\widetilde{E}_i$ with ending lamination λ_i. Moreover, if the orientation on S_i is chosen so that the geometrically infinite end in M_ρ is upward-pointing, then $\widetilde{E}_i$ is also upward-pointing in $(N_\nu)_i$. The Covering Theorem (see [62; 27]) then implies that the covering map $p_i\colon (N_\nu)_i \to N_\nu$ is finite-to-one on a neighborhood of $\widetilde{E}_i$. Therefore, we may homotope h so that $h|_{S_i}$ is an orientation-preserving covering map with image a component of ∂M_ν. If T_j is a toroidal component of ∂M_ρ, then, since all incompressible tori are peripheral in M_ν, $h|_{T_j}$ can again be homotoped to a covering map onto a toroidal component of ∂M_ν. Therefore, we may assume that h is a covering map on each component of ∂M_ρ and is orientation-preserving on each nontoroidal component.

Waldhausen's Theorem [67, Theorem 6.1] now implies that h is homotopic to an orientation-preserving covering map $h'\colon M_\rho \to M_\nu$, by a homotopy keeping $h|_{\partial M_\rho}$ constant. Since h is a homotopy equivalence, h' is a homeomorphism. It follows that (M_ν, h_ν) is equivalent to (M_ρ, h_ρ) and that the ending invariants are identified. The Ending Lamination Theorem then implies that $\nu = \rho$. It follows that $\{\nu_n\}$ converges to ρ as desired. □

5 Fenchel–Nielsen coordinates and projection coefficients

In this section we discuss and compare *length-twist parameters* for $\mathcal{T}(S)$. For traditional Fenchel–Nielsen twist parameters based on a maximal curve system $\boldsymbol{\alpha}$ (also known as a pants decomposition), we will see how the twist parameters compare with coarse twist parameters coming from projections to the annulus complexes associated to each curve in $\boldsymbol{\alpha}$. More generally for a curve system $\boldsymbol{\alpha}$ that may not be maximal, Theorem 5.1 allows us to vary arbitrarily the length and twist parameters of a curve system $\boldsymbol{\alpha}$, while (coarsely) fixing all subsurface projections in the complement of $\boldsymbol{\alpha}$.

To state the main theorem of this section we fix notation for the parameter spaces as follows. Given a curve system $\boldsymbol{\alpha} = \alpha_1 \cup \cdots \cup \alpha_m$, define $T_{\boldsymbol{\alpha}} = \mathbb{R}^m$, $L_{\boldsymbol{\alpha}} = \mathbb{R}^m_+$, and $V_{\boldsymbol{\alpha}} = T_{\boldsymbol{\alpha}} \times L_{\boldsymbol{\alpha}}$. For each component α_j of $\boldsymbol{\alpha}$ we have a geodesic length function $l_{\alpha_j}\colon \mathcal{T}(S) \to \mathbb{R}_+$, and we let

$$l_{\boldsymbol{\alpha}}\colon \mathcal{T}(S) \to L_{\boldsymbol{\alpha}}$$

denote $(l_{\alpha_1}, \dots, l_{\alpha_m})$.

Theorem 5.1 *Let $\boldsymbol{\alpha}$ be a curve system in S. For any $X \in \mathcal{T}(S)$ there is a continuous map*

$$\Phi\colon V_{\boldsymbol{\alpha}} \to \mathcal{T}(S)$$

such that $X \in \Phi(V_{\boldsymbol{\alpha}})$, and such that

(1) $l_{\boldsymbol{\alpha}} \circ \Phi(\mathbf{t}, \boldsymbol{\lambda}) = \boldsymbol{\lambda}$,

(2) $|\mathrm{tw}_{\boldsymbol{\alpha}}(X, \Phi(\mathbf{t}, \boldsymbol{\lambda})) - \mathbf{t}| < m_2$, *and*

(3) *for any essential subsurface $W \subset S$ disjoint from $\boldsymbol{\alpha}$ (except annuli parallel to components of $\boldsymbol{\alpha}$),*

$$\mathrm{diam}_{\mathcal{C}(W)}(\Phi(V_{\boldsymbol{\alpha}})) < m_2$$

where m_2 depends only on S.

We will precisely define $\mathrm{tw}_{\boldsymbol{\alpha}}$ below but roughly speaking it is an m–tuple of signed distances between the projections to the annular complexes associated to the curve system α.

Throughout this section an inequality of the form $|\mathbf{t}| < K$ for an m–tuple $\mathbf{t}$ refers to the sup norm on $\mathbf{t}$, so that we are just bounding each component individually.

Theorem 5.7 will state the special case of Theorem 5.1 when $\boldsymbol{\alpha}$ is a maximal curve system, namely that Fenchel–Nielsen coordinates can be chosen so that their twist parameters agree roughly with the parameters given by $\mathrm{tw}_{\boldsymbol{\alpha}}$.

At the end of the section we will prove Lemma 5.11, which is a connectivity result for a region in $\mathcal{T}(S)$ given by bounding the lengths of a curve system and restricting the structures in the complementary subsurfaces to certain neighborhoods of points at infinity. This lemma will be used in the last steps of the proofs of Theorems 1.5, 8.1 and 9.1.

5.1 Coarse twist parameters

An annulus complex is quasi-isometric to $\mathbb{Z}$. This allows us to define a signed version of distance. If α is the core curve of an annulus W we denote $\mathcal{C}(\alpha)=\mathcal{C}(W)$, $\pi_\alpha=\pi_W$, $d_\alpha=d_W$, and $\mathcal{C}(S,\alpha)=\mathcal{C}(S,W)$.

Given two elements a and b in $\mathcal{C}(\alpha)$ we let $i_\alpha(a,b)$ be the algebraic intersection of a and b. We then define

$$\mathrm{tw}_\alpha\colon \mathcal{C}(S,\alpha)\times\mathcal{C}(S,\alpha)\to\mathbb{Z}$$

by $\mathrm{tw}_\alpha(\gamma,\beta)=i_\alpha(\pi_\alpha(\gamma),\pi_\alpha(\beta))$. (If a and b have endpoints in common then algebraic intersection number is not well-defined. We correct this, in this special case, by taking the algebraic intersection of arcs in the homotopy class of a and b with minimal *geometric* intersection.)

There are two important properties of tw_α that we will use repeatedly:

(1) $d_\alpha(\gamma,\beta)=|\mathrm{tw}_\alpha(\gamma,\beta)|+1$ if $\gamma\neq\beta$.

(2) $|\mathrm{tw}_\alpha(\gamma,\beta)+\mathrm{tw}_\alpha(\beta,\zeta)-\mathrm{tw}_\alpha(\gamma,\zeta)|\leq 1$.

(See [51, Section 4] for closely related properties).

Recall, that in Section 2.3, we defined $\pi_\alpha(X)$, for $X\in\mathcal{T}(S)$, by setting $\pi_\alpha(X)=\pi_\alpha(\beta)$ where β is a shortest curve in X that intersects α. Abusing notation, we define

$$\mathrm{tw}_\alpha\colon \mathcal{T}(S)\times\mathcal{T}(S)\to\mathbb{Z}$$

by letting $\mathrm{tw}_\alpha(X,Y)=\mathrm{tw}_\alpha(\pi_\alpha(X),\pi_\alpha(Y))$. As we saw in Section 2.3 if β and β' are both shortest length curves in X that cross α then

$$|\mathrm{tw}_\alpha(\beta,\beta')|+1=d_\alpha(\beta,\beta')\leq 2.$$

Therefore tw_α is well-defined up to a uniform bound.

Recall that the length spectrum on a hyperbolic surface X, the values of lengths of curves on X, is discrete. Since length functions are continuous on $\mathcal{T}(S)$, this implies that the function on $\mathcal{T}(S)$ which gives back the length of the shortest curve that crosses α is continuous. These two facts allow us to find a neighborhood U of X in $\mathcal{T}(S)$ such that for every $Y\in U$ any shortest length curve in Y that crosses α is also a shortest length curve in X which crosses α. It follows that tw_α is *coarsely continuous*: there is a constant C such that every pair $(X,Y)\in\mathcal{T}(S)\times\mathcal{T}(S)$ has a neighborhood U such that $\mathrm{diam}(\mathrm{tw}_\alpha(U))<C$.

If $\boldsymbol{\alpha}=\alpha_1\cup\cdots\cup\alpha_m$ is a curve system then $\mathrm{tw}_{\boldsymbol{\alpha}}$ takes values in $\mathbb{Z}^m$.

5.2 Earthquakes and Fenchel–Nielsen coordinates

For a curve α, $s \in \mathbb{R}$ and $X \in \mathcal{T}(S)$, a *right earthquake* of magnitude s along α is obtained by cutting X along the geodesic representative of α and shearing to the right by signed distance s before regluing (so negative s corresponds to left shearing). See Thurston [63] and Kerckhoff [39]. Let $e_{\alpha,t}(X)$ denote the result of a right earthquake of magnitude $tl_\alpha(X)$, so that in particular

$$e_{\alpha,1}(X) = \theta_\alpha(X)$$

where θ_α is a left Dehn-twist on X. The equivalence of left twists with right shears corresponds to the fact that a mapping class f acts on $\mathcal{T}(S)$ by precomposing the marking with f^{-1}.

For a curve system $\boldsymbol{\alpha}$ and $\mathbf{t} \in T_{\boldsymbol{\alpha}}$, with components $t_{\alpha_j} = t_j$, note that the shears e_{α_j,t_j} commute and define

$$e_{\boldsymbol{\alpha},\mathbf{t}} = e_{\alpha_1,t_1} \circ \cdots \circ e_{\alpha_m,t_m}.$$

This earthquake map defines a free action of $T_{\boldsymbol{\alpha}}$ on $\mathcal{T}(S)$ which fixes the fibers of the length map $l_{\boldsymbol{\alpha}}$.

Now suppose that $\boldsymbol{\alpha}$ is a maximal curve system. Then the action on the fibers is also transitive and gives $\mathcal{T}(S)$ the structure of principal $\mathbb{R}^m$–bundle over $L_{\boldsymbol{\alpha}}$. A choice of section of this bundle determines *Fenchel–Nielsen coordinates* for $\mathcal{T}(S)$. More explicitly if

$$\sigma\colon L_{\boldsymbol{\alpha}} \to \mathcal{T}(S)$$

is a section then we can define a *Fenchel–Nielsen map*

$$F\colon V_{\boldsymbol{\alpha}} \to \mathcal{T}(S)$$

by

$$F(\mathbf{t}, \boldsymbol{\lambda}) = e_{\boldsymbol{\alpha},\mathbf{t}}(\sigma(\boldsymbol{\lambda})).$$

This map will be a homeomorphism and give *Fenchel–Nielsen coordinates* for $\mathcal{T}(S)$. There are a number of concrete constructions for sections and Fenchel–Nielsen coordinates, but none are particularly canonical.

5.3 Proof of Theorem 5.1

In this section we reduce the proof of Theorem 5.1 to three lemmas. We will prove these lemmas in the sections that follow.

It is not hard to measure how the twist parameter changes under powers of Dehn twists. In particular,

$$|\mathrm{tw}_\alpha(X, \theta_\alpha^n(X)) - n|$$

is uniformly bounded. Rather than prove this directly we replace the Dehn twist with the earthquake map which allows us to replace the integer n with a real number t. The first lemma generalizes the above bound for Dehn twists and is considerably more subtle to prove.

Lemma 5.2 *There exists a constant m_3 such that*

$$|\mathrm{tw}_{\boldsymbol{\alpha}}(X, e_{\boldsymbol{\alpha},\mathbf{t}}(X)) - \mathbf{t}| \leq m_3.$$

Next we see that projections to subsurfaces disjoint from $\boldsymbol{\alpha}$ remain coarsely constant when we earthquake along $\boldsymbol{\alpha}$.

Lemma 5.3 *There exists an m_4 such that for any essential subsurface $W \subset S$ disjoint from $\boldsymbol{\alpha}$ (except annuli parallel to components of $\boldsymbol{\alpha}$), and any $\mathbf{t} \in T_{\boldsymbol{\alpha}}$,*

$$d_W(X, e_{\boldsymbol{\alpha},\mathbf{t}}(X)) < m_4$$

where m_4 only depends on S.

Finally we will construct a section of the bundle $l_{\boldsymbol{\alpha}}\colon \mathcal{T}(S) \to L_{\boldsymbol{\alpha}}$ such the projection of all subsurfaces disjoint from $\boldsymbol{\alpha}$ is coarsely constant.

Lemma 5.4 *There exists an m_5 depending only on S such that the following holds. For any $X \in \mathcal{T}(S)$ there exists a section*

$$\sigma\colon L_{\boldsymbol{\alpha}} \to \mathcal{T}(S)$$

such that $X \in \sigma(L_{\boldsymbol{\alpha}})$ and if $W \subset S$ is an essential subsurface disjoint from $\boldsymbol{\alpha}$, then

$$\mathrm{diam}_{C(W)}(\sigma(L_{\boldsymbol{\alpha}})) < m_5.$$

Assuming these three lemmas it is easy to prove Theorem 5.1.

Proof of Theorem 5.1 We define the map Φ by

$$\Phi(\mathbf{t}, \boldsymbol{\lambda}) = e_{\boldsymbol{\alpha},\mathbf{t}}(\sigma(\boldsymbol{\lambda}))$$

where σ is the section given by Lemma 5.4. In particular $l_{\boldsymbol{\alpha}} \circ \sigma(\boldsymbol{\lambda}) = \boldsymbol{\lambda}$. Since the earthquake maps fix the lengths of $\boldsymbol{\alpha}$ we also have $l_{\boldsymbol{\alpha}} \circ \Phi(\mathbf{t}, \boldsymbol{\lambda}) = \boldsymbol{\lambda}$ and (1) holds.

Let $m_2 = \max\{m_3 + m_5, m_4 + m_5\}$. Note that

$$|\mathrm{tw}_{\boldsymbol{\alpha}}(X, \Phi(\mathbf{t}, \boldsymbol{\lambda})) - \mathrm{tw}_{\boldsymbol{\alpha}}(X, \sigma(\boldsymbol{\lambda})) - \mathrm{tw}_{\boldsymbol{\alpha}}(\sigma(\boldsymbol{\lambda}), \Phi(\mathbf{t}, \boldsymbol{\lambda}))| \leq 1.$$

Lemma 5.4 implies that $|\mathrm{tw}_{\boldsymbol{\alpha}}(X,\sigma(\boldsymbol{\lambda}))|+1 < m_5$ and Lemma 5.2 implies that $|\mathrm{tw}_{\boldsymbol{\alpha}}(\sigma(\boldsymbol{\lambda}),\Phi(\mathbf{t},\boldsymbol{\lambda}))-\mathbf{t}| \leq m_3$. Therefore

$$|\mathrm{tw}_{\boldsymbol{\alpha}}(X,\Phi(\mathbf{t},\boldsymbol{\lambda}))-\mathbf{t}| < m_3+m_5 \leq m_2$$

proving (2).

Let $W \subset S$ be an essential subsurface in S disjoint from $\boldsymbol{\alpha}$ which is not an annulus parallel to a component of $\boldsymbol{\alpha}$. By Lemma 5.4,

$$d_W(X,\sigma(\boldsymbol{\lambda})) < m_5,$$

and Lemma 5.3 implies that

$$d_W(\sigma(\boldsymbol{\lambda}),\Phi(\mathbf{t},\boldsymbol{\lambda})) < m_4.$$

Therefore

$$d_W(X,\Phi(\mathbf{t},\boldsymbol{\lambda})) < m_4+m_5 \leq m_2,$$

proving (3). □

5.4 Comparing twist coefficients

To prove Lemma 5.2 we need an effective method of calculating $\mathrm{tw}_{\boldsymbol{\alpha}}$. The map $\mathrm{tw}_{\boldsymbol{\alpha}}$ can be difficult to compute because, unlike other subsurface projections, it is defined by lifting curves to a cover rather than restricting them to a subsurface. We now describe a method for approximating $\mathrm{tw}_{\boldsymbol{\alpha}}$ by restricting the curves to an annular neighborhood of α (See Minsky [51] for a similar discussion.)

First, recall there is a uniform way to choose a regular neighborhood of a geodesic in a hyperbolic surface. Namely there is a function $w\colon \mathbb{R}^+ \to \mathbb{R}^+$ such that, for a simple closed geodesic γ of length l in any hyperbolic surface, the neighborhood of radius $w(l)$, which we call $\mathbf{collar}(\gamma)$, is an embedded annulus, and moreover

(1) $\mathbf{collar}(\gamma)\cap\mathbf{collar}(\beta)=\varnothing$ whenever $\gamma\cap\beta=\varnothing$, and

(2) the length l' of each component of $\partial\,\mathbf{collar}(\gamma)$ satisfies

$$\max(a_0,l(\gamma)) < l' < l(\gamma)+a_1$$

where a_0, a_1 are universal positive constants.

See eg [24, Theorem 4.4.6]. We can also define $\mathbf{collar}(\gamma)$ for a boundary component of a surface, and extend the definition to give horocyclic neighborhoods of cusps (here $l=0$ and $w=\infty$) by requiring that the boundary length of the neighborhood be fixed. If $\boldsymbol{\alpha}$ is a curve system then $\mathbf{collar}(\boldsymbol{\alpha})=\bigcup_{\alpha_j\in\alpha}\mathbf{collar}(\alpha_j)$.

If α is a single curve and a and b are properly embedded arcs in $\mathbf{collar}(\alpha)$ let $i^{\mathbf{c}}_{\alpha}(a,b)$ be their algebraic intersection. (When a and b have common endpoints we modify the definition just as we did for $i_{\boldsymbol\alpha}(a,b)$.) If γ and β are simple closed curves on S that intersect $\mathbf{collar}(\alpha)$ essentially and minimally in their homotopy class define

$$\mathrm{tw}^{\mathbf{c}}_{\boldsymbol\alpha}(\gamma,\beta) = i^{\mathbf{c}}_{\alpha}(a,b)$$

where a and b are components of $\gamma \cap \mathbf{collar}(\alpha)$ and $\beta \cap \mathbf{collar}(\alpha)$, respectively. As usual this definition depends on the choice of component but only up to a bounded amount. Note that while $\mathrm{tw}_{\boldsymbol\alpha}(\gamma,\beta)$ only depends on the homotopy classes of γ and β, $\mathrm{tw}^{\mathbf{c}}_{\alpha}(\gamma,\beta)$ depends strongly on the choice of curves. However, as we will see in the next lemma if γ and β satisfy certain geometric conditions then $\mathrm{tw}^{\mathbf{c}}_{\alpha}(\gamma,\beta)$ is a good approximation for $\mathrm{tw}_{\boldsymbol\alpha}(\gamma,\beta)$.

Notation To prevent a proliferation of constants throughout the remainder of this section we will use the following notation. The expression $x \sim y$ means that $|x-y|<c$ for some constant c that depends only on S. We write $x \overset{K}{\sim} y$ if the constant depends on S and some other constant K. For example, if $f \sim 0$ then the quantity $|f|$ is uniformly bounded.

Lemma 5.5 *Let α be a curve in a curve system $\boldsymbol\alpha$ on S and $X \in \mathcal{T}(S)$. Let γ and β be simple closed curves which intersect $\mathbf{collar}(\alpha)$ nontrivially, so that all components of their intersections with $\mathbf{collar}(\boldsymbol\alpha)$ and with $S \setminus \mathbf{collar}(\boldsymbol\alpha)$ are essential.*

Further assume that every component of $\gamma \cap (S \setminus \mathbf{collar}(\boldsymbol\alpha))$ that is adjacent to $\mathbf{collar}(\alpha)$ has length $< L$, and similarly for β. Then

$$\mathrm{tw}_{\alpha}(\gamma,\beta) \overset{L}{\sim} \mathrm{tw}^{\mathbf{c}}_{\alpha}(\gamma,\beta).$$

Proof We consider another measure of twisting. For two intersecting simple closed curves α and β and a hyperbolic structure X, we define a geometric shear of β about α in X, $s_{\alpha,X}(\beta) \in \mathbb{R}$, as follows. Let A be a lift of the geodesic representative of α to $\mathbb{H}^2$, let B be a lift of β which crosses A, and let $s_{\alpha,X}(\beta)$ denote $1/l_{\alpha}(X)$ times the signed distance along A between the orthogonal projections to A of the endpoints of B. The sign is chosen so that a left-earthquake of X along α will increase $s_{\alpha,X}(\beta)$.

Since any two lifts of β are disjoint, the values they give for $s_{\alpha,X}$ differ by at most 1 (see Farb, Lubotzky and Minsky [34] for a discussion along these lines). Moreover, $s_{\alpha,X}$ measures roughly the (signed) number of fundamental domains of α crossed by the lift of β, and this means that a difference of shears $s_{\alpha,X}(\gamma) - s_{\alpha,X}(\beta)$ coarsely measures the algebraic intersection numbers of lifts of γ and β to the annulus cover associated to α. In other words, comparing this with the definition of $\mathrm{tw}_{\boldsymbol{\alpha}}$ we can see that, for any X, α, and γ, β both crossing α,

$$\mathrm{tw}_\alpha(\gamma, \beta) \sim s_{\alpha,X}(\beta) - s_{\alpha,X}(\gamma). \tag{5-1}$$

We now make a similar definition using only $\mathbf{collar}(\alpha)$. Let $\mathbf{C}$ be a neighborhood of A in $\mathbb{H}^2$ that is a lift of $\mathbf{collar}(\alpha)$ and consider the arc $B \cap \mathbf{C}$. Let $s^{\mathbf{c}}_{\alpha,X}(\beta)$ denote $1/l_\alpha(X)$ times the signed distance along A between the orthogonal projections to A of the endpoints of $B \cap \mathbf{C}$. As for $s_{\alpha,X}(\beta)$ the signs are chosen so that a left-earthquake of X along α will increase $s^{\mathbf{c}}_{\alpha,X}(\beta)$. Using the same reasoning as above we see that

$$\mathrm{tw}^{\mathbf{c}}_\alpha(\gamma, \beta) \sim s^{\mathbf{c}}_{\alpha,X}(\beta) - s^{\mathbf{c}}_{\alpha,X}(\gamma).$$

Note that $s_{\alpha,X}(\beta)$ only depends on the homotopy class of β and the choice of lift. On the other hand, $s^{\mathbf{c}}_{\alpha,X}(\beta)$ depends strongly on the curve β. However, given the restrictions we have put on β we claim

$$s_{\alpha,X}(\beta) \overset{L}{\sim} s^{\mathbf{c}}_{\alpha,X}(\beta). \tag{5-2}$$

The lemma follows from this estimate.

To establish (5-2) we further examine the lift B of β. Let $x^{\mathbf{c}}$ be an endpoint of $B \cap \mathbf{C}$. After leaving $\mathbf{C}$ at $x^{\mathbf{c}}$, B must continue to another lift $\mathbf{D}$ of a component of $\mathbf{collar}(\boldsymbol{\alpha})$, and terminate at infinity at a point x on the other side of $\mathbf{D}$. The distance in $\partial\mathbf{C}$ between $x^{\mathbf{c}}$ and the orthogonal projection of x to $\partial\mathbf{C}$ will be bounded by L plus the diameter of the projection of $\mathbf{D}$. The latter projects to at most one fundamental domain of $\mathbf{C}$ because the collars of $\boldsymbol{\alpha}$ are embedded. The arc of length L projects, on the boundary of $\mathbf{C}$, to at most L/a_0 fundamental domains because the length of each of them is at least a_0. The bound of $1 + L/a_0$ fundamental domains therefore applies to the projection to the axis A as well. Applying the same estimate to the other endpoints, (5-2) follows. □

We can now prove Lemma 5.2.

Proof of Lemma 5.2 We first assume that $\boldsymbol{\alpha}$ is a maximal curve system. Let α_j be a curve in $\boldsymbol{\alpha}$ and let β be a shortest curve in X that crosses α_j, chosen so that $\pi_{\alpha_j}(\beta) = \pi_{\alpha_j}(X)$. Note that $\mathbf{collar}(\boldsymbol{\alpha})$ has a canonical affine structure given by the orthogonal foliations consisting of vertical geodesics orthogonal to core geodesics and horizontal curves equidistant to the core curve. There is then a canonical map from X to $e_{\boldsymbol{\alpha},\mathbf{t}}(X)$ that is an isometry on $X \setminus \mathbf{collar}(\boldsymbol{\alpha})$ and is an affine shear on each component of $\mathbf{collar}(\boldsymbol{\alpha})$. Let β' be the image of β under this map and let $\gamma = \pi_{\alpha_j}(X)$ be a shortest curve in $e_{\boldsymbol{\alpha},\mathbf{t}}(X)$ that crosses α_j. Then $\mathrm{tw}_{\alpha_j}(X, e_{\boldsymbol{\alpha},\mathbf{t}}(X)) = \mathrm{tw}_{\alpha_j}(\beta, \gamma)$.

Since $\boldsymbol{\alpha}$ is maximal and β is a shortest curve that crosses α_j, the length of every component of $\beta \setminus \mathbf{collar}(\boldsymbol{\alpha})$ in X is uniformly bounded. It follows that every arc in $\beta' \setminus \mathbf{collar}(\boldsymbol{\alpha})$ is uniformly bounded in $e_{\boldsymbol{\alpha},\mathbf{t}}(X)$. Similarly every component of $\gamma \setminus \mathbf{collar}(\boldsymbol{\alpha})$ has uniformly bounded length in $e_{\boldsymbol{\alpha},\mathbf{t}}(X)$. Therefore we can apply Lemma 5.5 to β' and γ.

Since β is a shortest curve crossing α_j in X there is a vertical arc b in $\mathbf{collar}(\alpha_j)$ that is disjoint from a component of $\beta \cap \mathbf{collar}(\alpha_j)$. Let b' be the image of b under the affine shear determined by $e_{\boldsymbol{\alpha},\mathbf{t}}$. In particular b' will be disjoint from a component of β'. Similarly there is a vertical arc a disjoint from a component of $\gamma \cap \mathbf{collar}(\alpha_j)$. Therefore

$$|\mathrm{tw}^{\mathbf{c}}_{\alpha_j}(\beta', \gamma) - i^{\mathbf{c}}_{\alpha_j}(b', a)| \leq 2.$$

From the construction of the earthquake map we also see that

$$|i^{\mathbf{c}}_{\alpha_j}(b', a) - t_j| \leq 1$$

and it follows that

$$|\mathrm{tw}^{\mathbf{c}}_{\alpha_j}(\beta', \gamma) - t_j| \leq 3.$$

Lemma 5.5 then gives us our desired estimate for $\mathrm{tw}_{\alpha_j}(X, e_{\boldsymbol{\alpha},\mathbf{t}}(X))$ and applying this estimate to each component of $\boldsymbol{\alpha}$ gives us the lemma when $\boldsymbol{\alpha}$ is maximal.

If $\boldsymbol{\alpha}$ is not maximal we extend it to a maximal system $\widehat{\boldsymbol{\alpha}}$. Given $\mathbf{t} \in T_{\boldsymbol{\alpha}}$, we extend it to $\widehat{\mathbf{t}} \in T_{\widehat{\boldsymbol{\alpha}}}$ by letting all the coordinates corresponding to components of $\widehat{\boldsymbol{\alpha}} - \boldsymbol{\alpha}$ be 0. We then have

$$\begin{aligned} |\mathrm{tw}_{\boldsymbol{\alpha}}(X, e_{\boldsymbol{\alpha},\mathbf{t}}(X)) - \mathbf{t}| &= |\,\mathrm{tw}_{\boldsymbol{\alpha}}(X, e_{\widehat{\boldsymbol{\alpha}},\widehat{\mathbf{t}}}(X)) - \mathbf{t}| \\ &\leq |\mathrm{tw}_{\widehat{\boldsymbol{\alpha}}}(X, e_{\widehat{\boldsymbol{\alpha}},\widehat{\mathbf{t}}}(X)) - \widehat{\mathbf{t}}|. \end{aligned}$$

The desired bound then follows from the bound in the maximal case because $e_{\widehat{\boldsymbol{\alpha}},\widehat{\mathbf{t}}}(X) = e_{\boldsymbol{\alpha},\mathbf{t}}(X)$. □

We can now prove a special case of Lemma 5.4 when $\boldsymbol{\alpha}$ is a maximal curve system. This special case is required to prove the more general version of the lemma.

Lemma 5.6 *Let $\boldsymbol{\alpha}$ be a maximal curve system on S and let $X \in \mathcal{T}(S)$. Then there exists a section*

$$\sigma\colon L_{\boldsymbol{\alpha}} \to \mathcal{T}(S)$$

such that $X \in \sigma(L_{\boldsymbol{\alpha}})$ and

$$|\mathrm{tw}_{\boldsymbol{\alpha}}(X, Y)| \sim 0$$

for all $Y \in \sigma(L_{\boldsymbol{\alpha}})$.

Proof Let

$$\hat{\sigma}\colon L_{\boldsymbol{\alpha}} \to \mathcal{T}(S)$$

be an arbitrary choice of section. We will use Lemma 5.2 to "twist" $\hat{\sigma}$ to our desired section σ.

Define a function $g\colon L_{\boldsymbol{\alpha}} \to T_{\boldsymbol{\alpha}}$ by

$$g(\boldsymbol{\lambda}) = \mathrm{tw}_{\boldsymbol{\alpha}}(X, \hat{\sigma}(\boldsymbol{\lambda})).$$

Since $\hat{\sigma}$ is continuous and $\mathrm{tw}_{\boldsymbol{\alpha}}$ is coarsely continuous, the function g is coarsely continuous. Recall this means there exists a constant $C > 0$ such that any $\boldsymbol{\lambda} \in L_{\boldsymbol{\alpha}}$ has a neighborhood U with $\mathrm{diam}(g(U)) < C$.

In particular, there exists a continuous function $\hat{g}\colon L_{\boldsymbol{\alpha}} \to T_{\boldsymbol{\alpha}}$, such that $|g - \hat{g}| < 2C$: Simply triangulate $L_{\boldsymbol{\alpha}}$ sufficiently finely, set $\hat{g} = g$ on the 0–skeleton, and extend by affine maps to each simplex.

We now define σ by setting

$$\sigma(\boldsymbol{\lambda}) = e_{\boldsymbol{\alpha}, -\hat{g}(\boldsymbol{\lambda})}(\hat{\sigma}(\boldsymbol{\lambda})).$$

Lemma 5.2 then implies that

$$|\mathrm{tw}_{\boldsymbol{\alpha}}(\hat{\sigma}(\boldsymbol{\lambda}), \sigma(\boldsymbol{\lambda})) + \hat{g}(\boldsymbol{\lambda})| < m_3.$$

Using the fact that

$$|\mathrm{tw}_{\boldsymbol{\alpha}}(X, \hat{\sigma}(\boldsymbol{\lambda})) + \mathrm{tw}_{\boldsymbol{\alpha}}(\hat{\sigma}(\boldsymbol{\lambda}), \sigma(\boldsymbol{\lambda})) - \mathrm{tw}_{\boldsymbol{\alpha}}(X, \sigma(\boldsymbol{\lambda}))| \leq 1$$

and the bound on the difference between g and $\hat{g}$ we have

$$|\mathrm{tw}_{\boldsymbol{\alpha}}(X, \sigma(\boldsymbol{\lambda}))| < m_3 + 2C + 1. \qquad \square$$

Note that if $\boldsymbol{\alpha}$ is a maximal curve system then Lemma 5.3 is vacuous. In particular we have already proven Theorem 5.1 in this special case. As it may be of independent interest we state it as a theorem here.

Theorem 5.7 *Let $\boldsymbol{\alpha}$ be a maximal curve system for S. For any $X \in \mathcal{T}(S)$ there exist Fenchel–Nielsen coordinates*

$$F\colon V_{\boldsymbol{\alpha}} \to \mathcal{T}(S)$$

such that

$$\mathrm{tw}_{\boldsymbol{\alpha}}(X, F(\mathbf{t}, \boldsymbol{\lambda})) \sim \mathbf{t}$$

for all $\mathbf{t} \in T_{\boldsymbol{\alpha}}$ and all $\boldsymbol{\lambda} \in L_{\boldsymbol{\alpha}}$.

5.5 Proof of Lemma 5.3

To prove Lemmas 5.3 and 5.4 we need to control subsurface projections along subsurfaces in the complement of the curve system $\boldsymbol{\alpha}$ as we twist along $\boldsymbol{\alpha}$ and as we vary the length of $\boldsymbol{\alpha}$. The difficulty is that as we vary the lengths of $\boldsymbol{\alpha}$ we can not hope to control the behavior of the collection of shortest curves, especially when all components of $\boldsymbol{\alpha}$ are very long. What we will do instead is control the lengths of arcs on complementary subsurfaces and we will see that this is sufficient. The following lemma contains a more precise statement. It will be used in the proofs of both Lemmas 5.3 and 5.4.

If $R \subset S$ is an essential nonannular subsurface and X is a given hyperbolic structure on S, let $R^{\mathbf{c}}$ denote the component of $S \setminus \mathbf{collar}(\partial R)$ which is isotopic to R.

Lemma 5.8 *Let $R \subseteq S$ be a nonannular essential subsurface and $W \subseteq R$ an essential (possibly annular) subsurface nested in R. Let κ be an essential simple closed curve or properly embedded arc in R that intersects W essentially and let $L > 0$ be a constant. If X and Y are hyperbolic structures in $\mathcal{T}(S)$ such that the length of $\kappa \cap R^{\mathbf{c}}$ is bounded by L in both X and Y then*

$$d_W(X, Y) \overset{L}{\sim} 0.$$

Proof We first extend κ to an essential simple closed γ. If both endpoints of κ lie on components of ∂R that are on the boundary of the same component of $S \setminus R$ then we choose γ such that $\gamma \cap R = \kappa$. If the endpoints are on the boundary of different components then we construct γ such that $\gamma \cap R$ is the union of κ and an arc parallel to κ. If κ is a simple closed curve then $\gamma = \kappa$. In all cases each component of the restriction of γ to $R^{\mathbf{c}}$ has length bounded by L.

We first assume that W is nonannular. Let β be a shortest curve in X that intersects W essentially, such that $\pi_W(\beta) = \pi_W(X)$. The restriction of both γ and β to $W^{\mathbf{c}}$ will have uniformly bounded length and hence uniformly bounded intersection. Therefore $\pi_W(\gamma)$ and $\pi_W(\beta)$ have bounded intersection giving a uniform bound on $d_{\mathcal{C}(W)}(\pi_W(\gamma), \pi_W(\beta))$.

If β' is a shortest curve in Y that intersects W essentially such that $\pi_W(\beta') = \pi_W(Y)$, the same argument shows that $d_{\mathcal{C}(W)}(\pi_W(\gamma), \pi_W(\beta'))$ is uniformly bounded. The triangle inequality then implies that

$$d_{\mathcal{C}(W)}(\pi_W(\beta), \pi_W(\beta')) = d_W(X, Y)$$

is uniformly bounded which completes the proof in the nonannular case.

We now assume that W is an annulus with core curve ζ. Since each arc of $\gamma \cap \mathbf{collar}(\zeta)$ has length at most L, the width of the collar is bounded from above, which gives a bound from below on $l_\zeta(X)$. Together these bounds imply a bound on the number of times a component of $\gamma \cap \mathbf{collar}(\zeta)$ winds around $\mathbf{collar}(\zeta)$. (More concretely, it gives an upper bound on the absolute value of the algebraic intersection number of the component with a geodesic arc in $\mathbf{collar}(\zeta)$ which is orthogonal to ξ.) Let β be a shortest curve in X crossing ζ, such that $\pi_\zeta(\beta) = \pi_\zeta(X)$. Since $l_\zeta(X)$ is bounded below, $l_\beta(X)$ is uniformly bounded above. Since β is a shortest curve each arc in $\beta \cap \mathbf{collar}(\zeta)$ intersects each geodesic arc in $\mathbf{collar}(\zeta)$ which is orthogonal to ζ at most once. Therefore, there is a uniform bound on $|\mathrm{tw}^{\mathbf{c}}_\zeta(\beta,\gamma)|$ (measured with respect to X).

If $\boldsymbol{\alpha} = \partial R \cup \zeta$ then every component of $\gamma \cap \mathbf{collar}(\boldsymbol{\alpha})$ that is adjacent to $\mathbf{collar}(\zeta)$ has length bounded by L so we can apply Lemma 5.5 to conclude that

$$\mathrm{tw}_\zeta(\beta,\gamma) \stackrel{L}{\sim} \mathrm{tw}^{\mathbf{c}}_\zeta(\beta,\gamma) \stackrel{L}{\sim} 0.$$

Repeating the argument with a curve β' that is shortest in Y, such that $\pi_\zeta(Y) = \pi_\zeta(\beta')$, we get a bound on $\mathrm{tw}_\zeta(\beta',\gamma)$, and the desired bound on $\mathrm{tw}_\zeta(\beta,\beta') = \mathrm{tw}_\zeta(X,Y)$ follows. □

Lemma 5.3 now follows easily. The proof of Lemma 5.4 is more involved.

Proof of Lemma 5.3 Let W be a nonannular subsurface in the complement of α. Let κ be a shortest curve on X that intersects W, so that there is a uniform length bound on κ. Since the earthquake map is an isometry on W we have the same length bound on the intersection of κ with $W^{\mathbf{c}}$ in the metric $e_{\boldsymbol{\alpha},\mathbf{t}}(X)$. Therefore by Lemma 5.8, $d_W(X, e_{\boldsymbol{\alpha},\mathbf{t}}(X))$ is uniformly bounded.

Now let W be an annulus with core curve ζ. Add ζ to $\boldsymbol{\alpha}$ to make a new curve system $\hat{\boldsymbol{\alpha}}$ and let $\hat{\mathbf{t}} \in L_{\hat{\boldsymbol{\alpha}}}$ be equal to $\mathbf{t}$ on the original $\boldsymbol{\alpha}$–coordinates and 0 on the ζ–coordinate. Then $e_{\hat{\boldsymbol{\alpha}},\hat{\mathbf{t}}}(X) = e_{\boldsymbol{\alpha},\mathbf{t}}(X)$. The bound on $|\mathrm{tw}_\zeta(X, e_{\boldsymbol{\alpha},\mathbf{t}}(X))|$ now follows from Lemma 5.2. □

5.6 Geometry of pants

Before we begin the proof of Lemma 5.4, we need to make some geometric observations about pairs of pants. These are fairly basic but we will take some care because we need statements that will hold uniformly for curves of all lengths.

Let Y be a hyperbolic pair of pants with geodesic boundary, and let l_1, l_2, l_3 denote its boundary lengths (we allow 0 for a cusp). Recall that $Y^{\mathbf{c}}$ denotes $Y \setminus \mathbf{collar}(\partial Y)$. Now for each permutation (i,j,k) of $(1,2,3)$, call a properly embedded essential arc in $Y^{\mathbf{c}}$ of type *ii* if both its endpoints lie on the i–th boundary component, and of type *jk* if its endpoints lie in the j–th and k–th boundary components. Define

- x_i to be the length of the shortest arc of type ii,
- y_i to be the length of the shortest arc of type jk, and
- $\Delta_i = \frac{1}{2}(l_j + l_k - l_i)$.

The following lemma encodes the fact that y_i is estimated by Δ_i when $\Delta_i > 0$, and x_i is estimated by $-\Delta_i$ when $\Delta_i < 0$ – and that $\min(x_i, y_i)$ is always bounded above. This is because $Y^{\mathbf{c}}$ retracts uniformly to a 1–complex whose combinatorial type and geometry are (approximately) dictated by the numbers Δ_i.

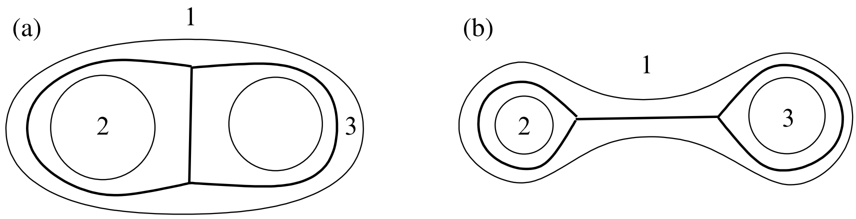

Figure 1: The two types of thick hyperbolic pants $Y^{\mathbf{c}}$. In type (a), the edges of the 1–complex have lengths Δ'_1, Δ'_2 and Δ'_3. In type (b), the edge lengths are $-\Delta'_1, l'_2$ and l'_3.

Lemma 5.9 *There exists $a > 0$ such that, for a hyperbolic pair of pants labeled as above,*

$$\max(\Delta_i, 0) - a \le x_i \le 2\max(\Delta_i, 0) + a$$

and

$$\max(-\Delta_i, 0) - a \le y_i \le \max(-\Delta_i, 0) + a.$$

Proof (Sketch) There is a subdivision (Voronoi diagram) of $Y^{\mathbf{c}}$ into three convex annuli, of width bounded by a uniform w_1, each containing the points closest to one of the boundary components. The annuli meet in a geodesic 1–complex to which $Y^{\mathbf{c}}$ retracts. The i–th annulus is attached to the 1–complex along a curve whose length we denote by l'_i; note that $l_i < l'_i < l_i + a_2$ for a uniform a_2.

Now defining $\Delta'_i = \frac{1}{2}(l'_j + l'_k - l'_i)$, it is easy to see that the signs of the Δ'_i (either all nonnegative or exactly one negative) determine the combinatorial type of this 1–complex: If all Δ'_i are nonnegative then the 1–complex is a "theta", three arcs attached along endpoints so any two make a loop, and each Δ'_i is the length of the arc which, when deleted, leaves a loop homotopic into the i–th annulus (see Figure 1 (a)). If one $\Delta'_i < 0$ then the 1–complex is a "pair of glasses", ie two disjoint loops homotopic into annuli j and k respectively and attached to the endpoints of an arc, whose length is $-\Delta_i'$ (Figure 1 (b)). Consider for example the theta case: each y_i is bounded above by $2w_1$, and each x_i is between Δ_i' and $\Delta_i' + 2w_1$. The pair of glasses case is similar with a bit less symmetry, accounting for the factor of 2 in the inequality. Finally, the fact that $|\Delta_i - \Delta_i'| < 3a_2$ finishes the proof. □

Let $\mathcal{P}$ be a pair of pants and $\mathcal{T}(\mathcal{P})$ the Teichmüller space of hyperbolic structures with geodesic boundary on $\mathcal{P}$. We also allow the possibility that one or more of the boundary components is a cusp. The l_i, x_i, y_i and Δ_i are now functions on $\mathcal{T}(\mathcal{P})$. We also let $l_{ij} = (l_i, l_j)$ be the function which gives back the lengths of the i–th and j–th boundary component.

The following lemma should be thought of as a version of Lemma 5.4 for pairs of pants.

Lemma 5.10 *Given $s > 0$ there exists an s' such that the following holds. Let Y be a hyperbolic structure in $\mathcal{T}(\mathcal{P})$.*

(1) *If $x_1(Y) < s$ then there exists a section $\sigma\colon [0,\infty) \to \mathcal{T}(\mathcal{P})$ such that $l_1 \circ \sigma = \mathrm{id}$, $Y = \sigma(l_1(Y))$ and $x_1(Z) < s'$ for all $Z \in \sigma([0,\infty))$.*

(2) *If $y_1(Y) < s$ there exists a section $\sigma\colon [0,\infty)^2 \to \mathcal{T}(\mathcal{P})$ such that $l_{23} \circ \sigma = \mathrm{id}$, $Y = \sigma(l_{23}(Y))$ and $y_1(Z) < s'$ for all $Z \in \sigma([0,\infty)^2)$.*

Proof We first prove (2). By Lemma 5.9 we need to find a section such that the function $\max(-\Delta_1, 0)$ is bounded on the image of σ. The Teichmüller space $\mathcal{T}(\mathcal{P})$ is parametrized by the lengths of the boundary curves. This gives $\mathcal{T}(\mathcal{P})$ a linear structure on which $\max(-\Delta_1, 0)$ is a convex. Triangulate $[0,\infty)^2$ with linear triangles and such that $l_{23}(Y)$ is a vertex in the triangulation. Define $\sigma(l_{23}(Y)) = Y$ and for any other vertex v in the triangulation we define $\sigma(v)$ such that $\Delta_1(\sigma(v)) = 0$. We then extend σ linearly across each triangle. By Lemma 5.9, $\max(-\Delta_i(Y), 0)$ is bounded by a constant only depending on s. On all other vertices $\max(-\Delta(\sigma(v)), 0) = 0$. Therefore, by convexity, $\max(-\Delta_1, 0) \le \max(-\Delta_1(Y), 0)$ on the image of σ as desired.

We can follow the same strategy to prove (1) except that now the triangulation of $[0,\infty)$ is just a partition into countably many compact segments. □

Proof of Lemma 5.4 We will enlarge $\boldsymbol{\alpha}$ to a suitably chosen maximal curve system $\widehat{\boldsymbol{\alpha}}$ and write $L_{\widehat{\boldsymbol{\alpha}}} = L_{\boldsymbol{\alpha}} \times L_{\widehat{\boldsymbol{\alpha}} \setminus \boldsymbol{\alpha}}$. We will then define the section σ by taking the section

$$\sigma_{\widehat{\boldsymbol{\alpha}}}\colon L_{\widehat{\boldsymbol{\alpha}}} \to \mathcal{T}(S)$$

given by Lemma 5.6 and precomposing it with a suitable section

$$\psi\colon L_{\boldsymbol{\alpha}} \to L_{\widehat{\boldsymbol{\alpha}}}.$$

That is we set $\sigma = \sigma_{\widehat{\boldsymbol{\alpha}}} \circ \psi$.

We will select $\widehat{\boldsymbol{\alpha}}$ satisfying the following geometric properties:

(1) We can write $X \setminus \boldsymbol{\alpha}$ as

$$X \setminus \boldsymbol{\alpha} = Z_1 \supset Z_2 \supset \cdots \supset Z_k$$

such that, for $1 \leq i < k$, Z_{i+1} is obtained from Z_i by cutting along a properly embedded arc κ_i. More precisely, we will let Y_i be a pair of pants component of a regular neighborhood of $\kappa_i \cup \partial Z_i$, and let $Z_{i+1} = \operatorname{int}(Z_i \setminus Y_i)$.

(2) The boundary components of Y_i that are incident to κ_i are exactly those that are parallel to ∂Z_i.

(3) The length of $\kappa_i \cap Y_i^{\mathbf{c}}$ will be bounded by a uniform constant b.

(4) Z_k will be a disjoint union of pairs of pants.

See Figure 2 for an illustration of condition (2).

The maximal curve system $\widehat{\boldsymbol{\alpha}}$ will then be the union of $\boldsymbol{\alpha}$ with representatives of the isotopy classes of the boundaries of the Y_i.

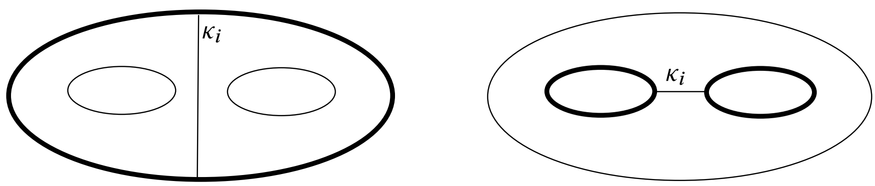

Figure 2: The two configurations of κ_i in Y_i allowed by condition (2). Heavily shaded boundary components are parallel to ∂Z_i.

We proceed by induction. Let $Z_1 = X \setminus \boldsymbol{\alpha}$, and let U be a component of Z_1 which is not a 3–holed sphere. Because each component of $\partial U^{\mathbf{c}}$ has length uniformly bounded below, area considerations give a uniform r_1, such that the neighborhood of $\partial U^{\mathbf{c}}$ of radius r_1 cannot be an embedded collar, and hence there is an essential properly embedded arc $\kappa_1' \subset U^{\mathbf{c}}$ of length bounded by $b = 2r_1$ (so we let κ_1 be properly embedded in U so that its intersection with $U^{\mathbf{c}}$ is κ_1'). Now let Y_1 be the pair of pants obtained from a regular neighborhood of $\kappa_1 \cup \partial Z_1$, and let $Z_2 = \operatorname{int}(Z_1 \setminus Y_1)$. Hence $\kappa_1 \subset Y_1$ satisfies conditions (1) and (3) above.

Since by construction both ends of κ_1 are on ∂U and ∂Y_1 must have at least one boundary component not parallel to ∂U, the only way condition (2) can fail is if (numbering ∂Y_1 appropriately and using the notation of Lemma 5.9), κ_1 is of type 11 in Y_1, while boundary component number 2 is (isotopic to) a boundary component of U. In this case, we can replace the 11 arc by a 12 arc, whose length we can also bound. Indeed, note that the bound on κ_1 gives a bound on Δ_1 by Lemma 5.9, and from the definition of the Δ_i we have that $-\Delta_3 \leq \Delta_1$, so that again by Lemma 5.9 we obtain a bound on the length of the 12 arc in $Y_1^{\mathbf{c}}$. We therefore replace κ_1 by the 12 arc (keeping the name κ_1), noting that conditions (1), (2), and (3) now hold.

Now repeat inductively in Z_i until we have reduced to a disjoint union of pairs of pants.

Having found $\widehat{\boldsymbol{\alpha}}$ using this construction, we construct the section $\psi\colon L_{\boldsymbol{\alpha}} \to L_{\widehat{\boldsymbol{\alpha}}}$ – that is, we build a continuous function $h\colon L_{\boldsymbol{\alpha}} \to L_{\widehat{\boldsymbol{\alpha}} \smallsetminus \boldsymbol{\alpha}}$. We will do this inductively, using Lemma 5.10. For Y_1, a point in $L_{\boldsymbol{\alpha}}$ determines the boundary lengths of the components of ∂Y_1 that are adjacent to κ_1, and the map given by Lemma 5.10 gives lengths for the remaining components. For each successive Y_i, then, the already-defined coordinates of h determine the lengths for the components of ∂Y_i that are adjacent to κ_i (here we use property (2) of the curve system $\widehat{\boldsymbol{\alpha}}$), and the lemma again determines the rest. We then let ψ be the section $\psi(\boldsymbol{\lambda}) = (\boldsymbol{\lambda}, h(\boldsymbol{\lambda}))$.

It remains to verify that $\operatorname{diam}_{\mathcal{C}(W)}(\sigma(L_{\boldsymbol{\alpha}}))$ is uniformly bounded for all subsurfaces $W \subset S$ that are disjoint from $\boldsymbol{\alpha}$.

If W is an annulus whose core curve is a component of $\widehat{\boldsymbol{\alpha}}$ then the bound follows from Lemma 5.6 as the image of σ lies in the image of $\sigma_{\widehat{\boldsymbol{\alpha}}}$.

If W is an annulus whose core curve is not in $\widehat{\boldsymbol{\alpha}}$ or W is nonannular, choose Z_i such that $W \subseteq Z_i$ but $W \not\subseteq Z_{i+1}$. This implies that κ_i intersects W essentially and we can apply Lemma 5.8 to W and a component of $\kappa_i \cap W$ to obtain the bound. □

5.7 Connectivity near infinity

The following lemma will be used in the final steps of the proofs of Theorems 1.5, 8.1 and 9.1. It is a connectivity result for subsets of Teichmüller space of the following type. Given a multicurve $\boldsymbol{\alpha}$, let $\{S_1, \dots, S_l\}$ be the components of $S \smallsetminus \boldsymbol{\alpha}$ that are not 3–holed spheres, select laminations $\lambda_i \in \mathcal{EL}(S_i)$, and let U_i be neighborhoods of λ_i in $\mathcal{C}(S_i)$ (recalling that $\mathcal{EL}(S_i)$ is the Gromov boundary of $\mathcal{C}(S_i)$ by Klarreich's theorem). Let $\mathbf{U}$ denote the tuple (U_i). Then define for $\epsilon < \epsilon_0$

$$\mathcal{T}(\epsilon, \mathbf{U}) = \{X \in \mathcal{T}(S) \mid \pi_{S_i}(X) \in U_i \ \forall i = 1, \dots, l, \\ l_{\alpha_j}(X) < \epsilon \ \forall \alpha_j \in \boldsymbol{\alpha}\}.$$

Lemma 5.11 *Given a multicurve $\boldsymbol{\alpha}$ on S, let $\{S_1, \dots, S_l\}$ be the components of $S \setminus \boldsymbol{\alpha}$ which are not thrice-punctured spheres. Given $\lambda_i \in \mathcal{EL}(S_i)$ and neighborhoods U_i of λ_i for all i and $\epsilon < \epsilon_0$, there exist neighborhoods $U_i' \subset U_i$ of λ_i in $\mathcal{C}(S_i)$, such that any two points in $\mathcal{T}(\epsilon, \mathbf{U}')$ are connected by a path in $\mathcal{T}(\epsilon, \mathbf{U})$.*

Proof Let $\mathcal{T}_\epsilon(\boldsymbol{\alpha})$ denote the region of $\mathcal{T}(S)$ where $l_{\alpha_j} < \epsilon$ for all $\alpha_j \in \boldsymbol{\alpha}$.

Recall that the Deligne–Mumford compactification of the Moduli space of S lifts to an "augmentation" of the Teichmüller space in which a stratum $\mathcal{T}_0(\gamma)$ is added for each curve system γ, corresponding to "noded" Riemann surfaces where exactly the elements of γ are pinched, and parametrized by $\mathcal{T}(S \setminus \gamma)$. The topology of this bordification is the smallest one for which the length functions of simple closed curves, extended to allow the value 0, are continuous (see eg Bers [9]).

Extended Fenchel–Nielsen coordinates give us an explicit description of the local topology at a stratum: enlarge $\boldsymbol{\alpha}$ to a maximal curve system $\widehat{\boldsymbol{\alpha}}$ and let $l_{\widehat{\boldsymbol{\alpha}}}$ and $t_{\widehat{\boldsymbol{\alpha}}}$ be associated Fenchel–Nielsen length and twist parameters as in Section 5.2. Adding $\mathcal{T}_0(\boldsymbol{\alpha})$ to $\mathcal{T}(S)$ corresponds to enlarging the parameter spaces $V_{\widehat{\boldsymbol{\alpha}}} = T_{\widehat{\boldsymbol{\alpha}}} \times L_{\widehat{\boldsymbol{\alpha}}}$ to allow points where $l_{\alpha_j} = 0$ exactly for $\alpha_j \in \boldsymbol{\alpha}$, and then taking a quotient by identifying points which agree on all coordinates except possibly those t_j for which $l_{\alpha_j} = 0$ (in other words shearing around a pinched curve is ignored). Let $\overline{V}_{\widehat{\boldsymbol{\alpha}}}$ denote this augmented parameter space, which gives a homeomorphic model for $\overline{\mathcal{T}}^{\boldsymbol{\alpha}} \equiv \mathcal{T}(S) \cup \mathcal{T}_0(\boldsymbol{\alpha})$. The map $V_{\widehat{\boldsymbol{\alpha}}} \to V_{\widehat{\boldsymbol{\alpha}} \smallsetminus \boldsymbol{\alpha}}$ that forgets the $\boldsymbol{\alpha}$ coordinates extends to a map of $\overline{V}_{\widehat{\boldsymbol{\alpha}}}$, and gives us a retraction $\overline{\mathcal{T}}^{\boldsymbol{\alpha}} \to \mathcal{T}_0(\boldsymbol{\alpha})$, which on $\mathcal{T}(S)$ is a fibration with contractible fibers.

Because length functions are continuous in this topology, there is a neighborhood $\mathcal{V}_{\boldsymbol{\alpha}}$ of $\mathcal{T}_0(\boldsymbol{\alpha})$ in $\mathcal{T}(S)$ for which the fibration, which we write $\psi\colon \mathcal{V}_{\boldsymbol{\alpha}} \to \mathcal{T}_0(\boldsymbol{\alpha}) = \mathcal{T}(S \smallsetminus \boldsymbol{\alpha})$, changes the lengths of, say, the set of shortest curves in the complement of $\boldsymbol{\alpha}$ by a ratio arbitrarily close to 1. Shrinking $\mathcal{V}_{\boldsymbol{\alpha}}$ if necessary we may also assume $\mathcal{V}_{\boldsymbol{\alpha}} \subset \mathcal{T}_\epsilon(\boldsymbol{\alpha})$. The small perturbation of lengths implies a distance bound in $\mathcal{C}(S_i)$, namely,

$$d_{\mathcal{C}(S_i)}(X, \psi(X)) < c_1, \tag{5-3}$$

for a uniform c_1, when $X \in \mathcal{V}_{\boldsymbol{\alpha}}$.

Using Theorem 5.1, for each $X \in \mathcal{T}_\epsilon(\boldsymbol{\alpha})$ one can find a path $\{X_t\}$ in $\mathcal{T}_\epsilon(\boldsymbol{\alpha})$ connecting X to a point $X' \in \mathcal{V}_{\boldsymbol{\alpha}}$, such that projections to each $\mathcal{C}(S_i)$ remain uniformly bounded. (In fact this is just a pinching deformation and the full power of Theorem 5.1 is not needed.) Hence we can and do choose c_1 sufficiently large that

$$\operatorname{diam}_{S_i}(\{X_t\}) < c_1. \tag{5-4}$$

In [46], it is shown that a Teichmüller geodesic in $\mathcal{T}(S_i)$ projects to a c_2–neighborhood of a $\mathcal{C}(S_i)$–geodesic, with c_2 uniform.

Now by the definition of the Gromov boundary, there is a neighborhood U_i^0 of each $\lambda_i \in \partial_\infty \mathcal{C}(S_i)$ in $\mathcal{C}(S_i)$ such that any $\mathcal{C}(S_i)$–geodesic with endpoints in U_i^0 has the property that its c_2–neighborhood is in U_i.

Let U_i' be a neighborhood of λ_i in $\mathcal{C}(S_i)$ whose $2c_1$–neighborhood is in U_i^0. Now suppose $X_1, X_2 \in \mathcal{T}_\epsilon(\boldsymbol{\alpha})$, and $\pi_{S_i}(X_j) \in U_i'$ $(j = 1, 2)$. Using (5-4) we can deform X_j to X_j' $(j = 1, 2)$ within $\mathcal{T}_\epsilon(S_i)$ so that $X_j' \in \mathcal{V}_{\boldsymbol{\alpha}}$ and the π_{S_i}–image of the path stays in a c_1–neighborhood of U_i'. Then, by (5-3), $\pi_{S_i}(\psi(X_j')) \in U_i^0$. Let G be a Teichmüller geodesic in $\mathcal{T}(S \setminus \boldsymbol{\alpha})$ connecting $\psi(X_1')$ to $\psi(X_2')$. Then $\pi_{S_i}(G) \subset U_i$, so a lift of G back to $\mathcal{V}_{\boldsymbol{\alpha}}$ with endpoints X_1' and X_2' will, again by (5-3), give us the desired continuous family. □

6 Deformations with controlled projections

In this section, we establish Lemma 6.1 which is a key technical tool in the paper. We begin with a system of curves on the top conformal boundary which are short in the manifold. Lemma 6.1 allows us to shrink the lengths of the curves on the top conformal boundary, without disrupting the subsurface projections on complementary subsurfaces and keeping the curves short in the manifold throughout the process.

Lemma 6.1 *Given S and $K > 1/\epsilon_0$, there exists $c = c(S)$, depending only on S, and $h = h(K, S)$, which depends on both K and S, such that if $X, Y \in \mathcal{T}(S)$ and $\boldsymbol{\alpha}$ is a curve system on S, such that*

$$\mathbf{m}_{\alpha_i}(X, Y) > h$$

or

$$l_{\alpha_i}(X) < 1/K$$

for each component α_i of $\boldsymbol{\alpha}$, then there exists a path $\{X_t : t \in [0, T]\}$ in $\mathcal{T}(S)$ with $X_0 = X$ such that

(1) $l_{\alpha_i}(X_T) < 1/K$ *for each α_i,*

(2) $\mathbf{m}_{\alpha_i}(X_t, Y) > K$ *for each α_i and each $t \in [0, T]$, and*

(3) $\mathrm{diam}(\pi_W(\{X_t : t \in [0, T]\})) < c$, *for any W disjoint from $\boldsymbol{\alpha}$.*

Recall that ϵ_0 is a specific Margulis constant chosen in Section 2.1 and that

$$\mathbf{m}_\gamma(X, Y) = \max\left(\frac{1}{l_\gamma(X)}, \frac{1}{l_\gamma(Y)}, \sup_{\gamma \subset \partial W} d_W(X, Y)\right)$$

is defined in Section 2.3. Theorem 2.2 asserts that $\mathbf{m}_\gamma(X, Y)$ is large if and only if γ is short in $Q(X, Y)$. Lemma 6.1 will follow from Theorem 5.1, which allows us to change lengths and maintain control on subsurface projections, and Lemma 2.3, which records key estimates concerning the partial order on subsurfaces of S.

Proof Let $\epsilon = 1/K$. Write $\boldsymbol{\alpha} = \boldsymbol{\alpha}^X \cup \boldsymbol{\alpha}^Y \cup \boldsymbol{\alpha}^0$, where $\boldsymbol{\alpha}^X$ consists of those components α_i with $l_{\alpha_i}(X) < \epsilon$, $\boldsymbol{\alpha}^Y$ consists of the components α_i of $\boldsymbol{\alpha} - \boldsymbol{\alpha}^X$ such that $l_{\alpha_i}(Y) < \epsilon$, and $\boldsymbol{\alpha}^0$ consists of the remaining components.

We argue by induction on the cardinality n of $\boldsymbol{\alpha}^0$ (which we note is bounded from above in terms of S). We will iteratively construct h_n (which implicitly depends on K and S) and show that the lemma holds if $\boldsymbol{\alpha}^0$ has n components and $\mathbf{m}_{\alpha_i}(X,Y) > h_n$ for all α_i in $\boldsymbol{\alpha}^0$ with a constant c_n in (3), which only depends on S.

If $n = 0$, we let $h_0 = K$ and $c_0 = m_2$ (the constant from Theorem 5.1). If $\boldsymbol{\alpha}^Y = \varnothing$ then the deformation is trivial, ie $T = 0$.

If $\boldsymbol{\alpha}^Y \neq \varnothing$, let $\Phi\colon V_{\boldsymbol{\alpha}} \to \mathcal{T}(S)$ be the map given by Theorem 5.1, such that $X \in \Phi(V_{\boldsymbol{\alpha}})$ – in fact we must have $X = \Phi((\mathbf{t}, l_{\boldsymbol{\alpha}}(X)))$ for some $\mathbf{t} \in T_{\boldsymbol{\alpha}}$. Let $\{X_t : t \in [0, T_1]\}$ be the Φ–image of the path in $V_{\boldsymbol{\alpha}}$ that begins at $(\mathbf{t}, l_{\boldsymbol{\alpha}}(X))$, shrinks the length of each α_i in $\boldsymbol{\alpha}^Y$ monotonically to $\epsilon/2$, fixes the length of every component of $\boldsymbol{\alpha}^X$ and fixes each twist coordinate. In particular, $l_{X_T}(\alpha_i) < \epsilon$ for all α_i in $\boldsymbol{\alpha}$. Since $l_{X_t}(\alpha_i) < \epsilon$ for all α_i in $\boldsymbol{\alpha}^X$ and $l_Y(\alpha_i) < \epsilon$ if α_i lies in $\boldsymbol{\alpha}^Y$, we see immediately $\mathbf{m}_{\alpha_i}(X_t, Y) > K$ for all i and t. Theorem 5.1 also implies that if W is a subsurface disjoint from $\boldsymbol{\alpha}$, then

$$\operatorname{diam}(\pi_W(\{X_t : t \in [0,T]\})) < m_2 = c_0.$$

The base case follows.

For $n > 0$, set $h_n = h_{n-1} + 2m_1 + m_2$ and $c_n = c_{n-1} + m_2$, where m_1 and m_2 are the constants from Lemma 2.3 and Theorem 5.1.

For each α_i in $\boldsymbol{\alpha}^0$, there must be some subsurface W_i with $\alpha_i \subset \partial W_i$, such that

$$d_{W_i}(X,Y) > h_n, \tag{6-1}$$

since $\mathbf{m}_{\alpha_i}(X,Y) > h_n > K$ but $l_{\alpha_i}(X) \geq \epsilon$ and $l_{\alpha_i}(Y) \geq \epsilon$. (Note that possibly $W_i = \mathbf{collar}(\alpha_i)$). Fix one such W_i for each $\alpha_i \in \boldsymbol{\alpha}^0$.

Since $h_n > m_1$, Lemma 2.3 implies that the set of domains $\mathcal{L}_{h_n}(X,Y)$, which contains all the W_i, is partially ordered by the relation $\prec$.

Reordering $\boldsymbol{\alpha}$ if necessary, we may assume that $\alpha_1 \in \boldsymbol{\alpha}^0$, and W_1 is $\prec$–maximal among the W_i, as well as maximal with respect to inclusion among the $\prec$–maximal elements (Lemma 2.3 implies that any two maximal elements are either disjoint or nested). In particular, the curves in ∂W_1 all lie above any curves they intersect in ∂W_i, so, intuitively, W_1 is the closest surface, among all the W_i, to the top of the manifold. Let $\boldsymbol{\beta} = \boldsymbol{\alpha}^X \cup \{\alpha_1\}$.

Now let $\Phi\colon V_{\boldsymbol{\beta}} \to \mathcal{T}(S)$ be the map given by Theorem 5.1 such that $X = \Phi(\mathbf{t}, l_{\boldsymbol{\beta}}(X))$ for some $\mathbf{t} \in T_{\boldsymbol{\beta}}$. Let $\{X_t : t \in [0, T_1]\}$ then be the Φ–image of the path in $V_{\boldsymbol{\beta}}$ that begins at $(\mathbf{t}, l_{\boldsymbol{\beta}}(X))$, shrinks the length of α_1 monotonically to ϵ, keeps the lengths of each element of $\boldsymbol{\alpha}^X$ fixed, and fixes each twist coordinate of an element of $\boldsymbol{\beta}$. Theorem 5.1 guarantees that

$$\operatorname{diam}(\pi_W(\{X_t : t \in [0, T_1]\})) \le m_2 \tag{6-2}$$

if W is disjoint from $\boldsymbol{\beta}$ (including the case when W is an annulus with core a component of $\boldsymbol{\beta}$). If W intersects one of the curves of $\boldsymbol{\alpha}^X$, then since their lengths are bounded by ϵ over the family X_t, we again have a bound on $\operatorname{diam}_W(\{X_t\})$, by a constant which we may assume is smaller than m_2. It follows, for any W_i disjoint from α_1, that for all $t \in [0, T_1]$,

$$d_{W_i}(X_t, Y) \ge h_n - m_2 = h_{n-1} + 2m_1. \tag{6-3}$$

In particular, $d_{W_1}(X_t, Y) \ge h_n - m_2$ for all t. More generally, we see that $\mathbf{m}_{\alpha_i}(X_t, Y) > h_{n-1}$ for all t, whenever W_i is disjoint from α_1.

If W_i intersects α_1 then, by the choice of W_1, we see that W_i and W_1 overlap and $W_i \prec W_1$, with respect to the order $\prec$ on $\mathcal{L}_{h_n}(X_0, Y)$. Lemma 2.3(2) implies that

$$d_{W_i}(Y, \partial W_1) > h_n - m_1 > m_1.$$

Then, since $d_{W_1}(X_t, Y) \ge h_n - m_2 > 2m_1$ for all t, the subsurface W_1 overlaps W_i, and $d_{W_i}(Y, \partial W_1) > m_1$, Lemma 2.3(5) implies that

$$d_{W_i}(X_t, Y) \ge h_n - m_2 - m_1.$$

In particular this implies that $\mathbf{m}_{\alpha_i}(X_t, Y) > h_{n-1}$ for all t and all α_i in $\boldsymbol{\alpha}^0$.

We now have a family $\{X_t : t \in [0, T_1]\}$ such that the number of components α_i of $\boldsymbol{\alpha}$ with $l_{\alpha_i}(X_{T_1}) \ge \epsilon$ and $l_{\alpha_i}(Y) \ge \epsilon$ is at most $n-1$. Moreover, for each α_i either, $\mathbf{m}_{\alpha_i}(X_{T_1}, Y) > h_{n-1}$, $l_{\alpha_i}(X_{T_1}) < \epsilon$ or $l_{\alpha_i}(Y) < \epsilon$ and if W is disjoint from $\boldsymbol{\alpha}$, then

$$\operatorname{diam}(\pi_W(\{X_t : t \in [0, T_1]\})) \le m_2.$$

Now applying the inductive hypothesis to X_{T_1}, we can concatenate this family with one that shrinks the remaining components of $\boldsymbol{\alpha}$ to have length at most ϵ, so that $\mathbf{m}_{\alpha_i}(X_t, Y) > K$ for each α_i and each t, and

$$\operatorname{diam}(\pi_W(\{X_t : t \in [0, T]\})) < c_n,$$

for any W disjoint from $\boldsymbol{\alpha}$. □

7 Bers slices

In this section, we prove Theorem 1.5 which we restate here for the reader's convenience.

Theorem 1.5 *Let B be a Bers slice of $QF(S)$ for some closed surface S. If $\rho \in \partial B$ and ρ is quasiconformally rigid in ∂B, then B does not self-bump at ρ. In particular, its closure $\overline{B}$ is locally connected at ρ.*

We will begin by proving that there is no self-bumping at a maximal cusp in the boundary of a Bers slice. The proof in this case is much simpler but follows the same outline as the proof of the general case.

7.1 The maximal cusp case

We first assume that ρ is a maximal cusp in the boundary of a Bers slice $B = B_Y$ in $AH(S)$. Let $\boldsymbol{\alpha}$ be the maximal curve system on S which is cusped in N_ρ.

If $\{\rho_n\}$ is a sequence in B_Y, then $\{\rho_n\}$ converges to ρ if and only if $\lim l_{\rho_n}(\alpha_j) = 0$ for all $\alpha_j \in \boldsymbol{\alpha}$. (Theorem 5 of Bers [8] implies that B_Y has compact closure in $AH(S \times I)$ while Theorem 1 in Maskit [45] implies that a maximal cusp in ∂B_Y is determined by its parabolic elements.) Therefore the sets

$$U(\delta) = \{\rho' \in B_Y : l_{\alpha_j}(\rho') < \delta \ \forall \alpha_j \in \boldsymbol{\alpha}\}.$$

for $\delta > 0$ give a neighborhood system for ρ in B_Y.

We will show that for each $\delta > 0$ there exists a neighborhood V of ρ such that any two points in $V \cap B_Y$ are connected by a path in $U_\delta \cap B_Y$. It then follows that there is no self-bumping at ρ.

First, let
$$W(\epsilon) = \{Q(X,Y) \in B_Y : l_{\alpha_j}(X) < \epsilon \ \forall \alpha_j \in \boldsymbol{\alpha}\}.$$

The set $W(\epsilon)$ is path-connected for any $\epsilon > 0$, because it is parametrized by a convex set in the Fenchel–Nielsen coordinates for $\mathcal{T}(S)$. Bers' Lemma 2.1 implies that $W(\epsilon) \subset U(2\epsilon)$. Hence it suffices to choose V so that any point in V can be connected to $W(\delta/2)$ by a path in $U(\delta)$.

Given $\delta > 0$, Theorem 2.2 allows us to choose K such that, for any $X, Y \in \mathcal{T}(S)$,

$$\mathbf{m}_\gamma(X,Y) > K \implies l_\gamma(Q(X,Y)) < \delta.$$

We may moreover require that $K > 2/\delta$. Let $h = h(K, S)$ be the constant given by Lemma 6.1. Theorem 2.2 then gives $\delta' > 0$ such that

$$l_\gamma(Q(X,Y)) < \delta' \implies \mathbf{m}_\gamma(X,Y) > h.$$

Now consider $V = U(\delta')$. If $Q(X,Y) \in V$, then $\mathbf{m}_{\alpha_i}(X,Y) > h$ for all $\alpha_j \in \boldsymbol{\alpha}$, so Lemma 6.1 gives a family $\{X_t \mid t \in [0,T]\} \subset \mathcal{T}(S)$ with $X_0 = X$ such that, for each $\alpha_j \in \boldsymbol{\alpha}$,

(1) $\mathbf{m}_{\alpha_j}(X_t, Y) > K$ for all $t \in [0,T]$, and

(2) $l_{\alpha_j}(X_T) \leq 1/K < \delta/2$.

It follows, from (1), that $Q(X_t, Y) \in U(\delta)$ for all t and, from (2), that $Q(X_T, Y)$ is contained in $W(\delta/2)$.

This completes the proof of Theorem 1.5 for maximal cusps.

7.2 General quasiconformally rigid points on the Bers boundary

In order to prove that there is no self-bumping at quasiconformally rigid points we must also allow for geometrically infinite ends. Theorem 2.7 and the Ending Lamination Theorem allow us to use subsurface projections to construct a neighborhood system about a general quasiconformally rigid point. Once we have constructed this neighborhood system the control we obtained on subsurface projections in Lemma 6.1 allows us to proceed much as in the proof of the maximal cusp case.

If $\rho \in \partial B_Y$ is quasiconformally rigid, then its geometrically infinite ends are associated with a disjoint collection of subsurfaces $\{S_1, \ldots, S_l\}$ of S and the cusps are associated with a collection $\boldsymbol{\alpha}$ of disjoint simple closed curves such that the components of $S \setminus \boldsymbol{\alpha}$ are precisely the S_i together with a (possibly empty) collection of thrice-punctured spheres. Let $\{\lambda_1, \ldots, \lambda_l\}$ be the ending laminations supported on $\{S_1, \ldots, S_l\}$.

Let U_i be a neighborhood of $\lambda_i \in \partial_\infty \mathcal{C}(S_i)$ in $\mathcal{C}(S_i)$ for each $i = 1, \ldots, l$. We denote by $\mathbf{U}$ the tuple $(U_1, \ldots, U_l)$, and for $\delta > 0$ we let $\mathcal{U}(\delta, \mathbf{U})$ be the set

$$\begin{aligned} \mathcal{U}(\delta, \mathbf{U}) = \{Q(X,Y) : \pi_{S_i}(X) \in U_i \ \forall i = 1, \ldots, l, \\ l_{\alpha_j}(Q(X,Y)) < \delta \ \forall \alpha_j \in \boldsymbol{\alpha}\}. \end{aligned}$$

Theorem 2.7 and the Ending Lamination Theorem allow us to show that the $\mathcal{U}(\delta, \mathbf{U})$ give a neighborhood system for ρ in B_Y. However, we should note that the sets $\mathcal{U}(\delta, \mathbf{U})$ need not be open in B_Y, since π_{S_i} is not continuous.

Lemma 7.1 *The sets $\mathcal{U}(\delta, \mathbf{U})$, where δ varies in $(0, \epsilon_0)$ and the U_i vary over neighborhoods of λ_i in $\overline{\mathcal{C}}(S_i)$, are the intersections with B_Y of a neighborhood system for ρ.*

Proof It suffices to show that a sequence $\{\rho_n = Q(X_n, Y)\}$ converges to ρ if and only if it is eventually contained in any $\mathcal{U}(\delta, \mathbf{U})$.

Let $\{\rho_n\}$ be a sequence eventually contained in any $\mathcal{U}(\delta, \mathbf{U})$. Since $\overline{B}_Y$ is compact, it suffices to show that any accumulation point of $\{\rho_n\}$ is ρ. Therefore we may assume $\{\rho_n\}$ converges to ρ'. By hypothesis, for each S_i, $\{\pi_{S_i}(X_n)\}$ converges to λ_i. Theorem 2.7 now implies that S_i faces an upward-pointing end of ρ' with ending lamination λ_i, for each i. Since $\lim l_{\alpha_j}(\rho_n) = 0$, each α_j corresponds to a cusp of ρ'. Since $\rho' \in \overline{B}_Y$, it has a downward pointing end associated to the full surface S, with conformal structure Y (see Bers [8, Theorem 8]). Thus, each cusp of ρ' is upward-pointing. Therefore, the end invariants of ρ' are the same as those of ρ. By the Ending Lamination Theorem, $\rho' = \rho$.

In the other direction, suppose $\{\rho_n\}$ converges to ρ. Then $\lim l_{\alpha_j}(\rho_n) = 0$ for all $\alpha_j \in \boldsymbol{\alpha}$, by continuity of length, and $\{\pi_{S_i}(X_n)\}$ converges to λ_i for all i, by Theorem 2.7. Hence $\{\rho_n\}$ is eventually contained in any $\mathcal{U}(\delta, \mathbf{U})$. □

Let $\mathcal{W}(\epsilon, \mathbf{U})$ denote a similarly defined set where the length bounds on $\boldsymbol{\alpha}$ take place in the boundary structure X, ie

$$\mathcal{W}(\epsilon, \mathbf{U}) = \{Q(X, Y) : \pi_{S_i}(X) \in U_i \ \forall i = 1, \dots, l, \quad l_{\alpha_j}(X) < \epsilon \ \forall \alpha_j \in \boldsymbol{\alpha}\}.$$

Notice that $\mathcal{W}(\epsilon, \mathbf{U}) = \{Q(X, Y) : X \in \mathcal{T}(\epsilon, \mathbf{U})\}$, where $\mathcal{T}(\epsilon, \mathbf{U})$ is as in Section 5.7. By Bers' Lemma 2.1, $\mathcal{W}(\delta/2, \mathbf{U}) \subset \mathcal{U}(\delta, \mathbf{U})$.

Theorem 1.5 follows from the following lemma:

Lemma 7.2 *Given $\delta > 0$ and neighborhoods U_i of λ_i, there exists $\epsilon > 0$ and neighborhoods V_i of λ_i such that any two points in $\mathcal{U}(\epsilon, \mathbf{V})$ can be connected by a path that remains in $\mathcal{U}(\delta, \mathbf{U})$.*

Proof By Theorem 2.2, choose K such that

$$\mathbf{m}_\gamma(X, Y) > K \implies l_\gamma(Q(X, Y)) < \delta$$

and also suppose $K > 2/\delta$. Let $h = h(K, S)$ be the constant given by Lemma 6.1, and let $c = c(S)$ be the constant in part (3) of Lemma 6.1. Lemma 5.11 allows us to choose neighborhoods W_i of λ_i such that any two points in $\mathcal{W}(\delta/2, \mathbf{W})$ are connected by a path in $\mathcal{W}(\delta/2, \mathbf{U})$.

Choose $\epsilon > 0$ small enough that (again by Theorem 2.2)

$$l_\gamma(Q(X,Y)) < \epsilon \implies \mathbf{m}_\gamma(X,Y) > h.$$

Finally, choose neighborhoods V_i of λ_i such that a c–neighborhood of V_i in $\mathcal{C}(S_i)$ is contained in W_i.

Let $Q(X,Y)$ be in $\mathcal{U}(\epsilon,\mathbf{V})$. Then, by our choice of ϵ, $\mathbf{m}_{\alpha_i}(X,Y) > h$ for each component of $\boldsymbol{\alpha}$, and so Lemma 6.1 can be applied to give a path $\{X_t : t \in [0,T]\}$ such that

(1) $l_{\alpha_j}(X_T) < 1/K < \delta/2$ for all $\alpha_j \in \boldsymbol{\alpha}$,

(2) $\mathbf{m}_{\alpha_j}(X_t,Y) > K$ for all $\alpha_j \in \boldsymbol{\alpha}$ and all $t \in [0,T]$, and

(3) $\mathrm{diam}_{\mathcal{C}(S_i)}(\pi_{S_i}(\{X_t\})) < c$ for all i.

It follows immediately from (1) and (3), that $Q(X_T,Y) \in \mathcal{W}(\delta/2,\mathbf{W})$. Moreover, (2) implies that $l_{\alpha_j}(Q(X_t,Y)) < \delta$ for all t and all $\alpha_j \in \boldsymbol{\alpha}$, so, again applying (3), we see that the entire path $\{Q(X_t,Y)\}$ lies in $\mathcal{U}(\delta,\mathbf{W})$.

This shows that any point in $\mathcal{U}(\epsilon,\mathbf{V})$ can be connected to $\mathcal{W}(\delta/2,\mathbf{W})$ by a path in $\mathcal{U}(\delta,\mathbf{W})$. Now since any two points in $\mathcal{W}(\delta/2,\mathbf{W})$ can be connected by a path in $\mathcal{W}(\delta/2,\mathbf{U})$, and since $\mathcal{W}(\delta/2,\mathbf{U}) \subset \mathcal{U}(\delta,\mathbf{U})$, we conclude that any two points in $\mathcal{U}(\epsilon,\mathbf{V})$ can be connected by a path in $\mathcal{U}(\delta,\mathbf{U})$. □

8 Acylindrical manifolds

In this section, we rule out self-bumping at quasiconformally rigid points in boundaries of deformation spaces of acylindrical 3–manifolds. Thurston's Bounded Image Theorem allows us to use essentially the same argument as in the Bers Slice case. Theorem 8.1 is the special case of Theorem 1.3 where M is acylindrical.

Theorem 8.1 *Let M be an acylindrical compact 3–manifold. If ρ is a quasiconformally rigid point in $\partial AH(M)$, then there is no self-bumping at ρ.*

Proof If B is a component of int$(AH(M))$ then there is an identification of B with $\mathcal{T}(S)$ where $S = \partial_T M$ is the nontoroidal portion of ∂M. Explicitly, we identify $\nu \in B$ with $\partial_c N_\nu$, regarded as a point in $\mathcal{T}(S)$. Thurston's Bounded Image Theorem asserts that the skinning map $\sigma\colon \mathcal{T}(S) \to \mathcal{T}(S)$ has bounded image. Let L be the diameter of $\sigma(\mathcal{T}(S))$.

We again begin by constructing a neighborhood system about ρ. Suppose that B is the component of int$(AH(M))$ such that $\rho \in \partial B$. Let (M_ρ, P_ρ) be a relative compact core for N_ρ. Let $\{S_1, \dots, S_l\}$ be the components of $\partial M_\rho - P_\rho$ which are not thrice-punctured spheres. Each S_i may be thought of as a subsurface of S and comes equipped with an ending lamination λ_i. The annular components of P_ρ are associated with a disjoint collection $\boldsymbol{\alpha}$ of simple closed curves on S. Since M is acylindrical, Θ is locally constant (see [5]), so our identification of B with $\mathcal{T}(S)$ is consistent with our identification of ∂M_ρ with S.

Let U_i be a neighborhood of $\lambda_i \in \partial_\infty \mathcal{C}(S_i)$ in $\mathcal{C}(S_i)$ for each $i = 1, \dots, i$. We denote by $\mathbf{U}$ the tuple $(U_1, \dots, U_l)$, and for $\delta > 0$ we let $\mathcal{U}(\delta, \mathbf{U})$ be the set

$$\mathcal{U}(\delta, \mathbf{U}) = \{\nu \in B : \pi_{S_i}(\partial_c N_\nu) \in U_i \ \forall i = 1, \dots, l, \\ l_{\alpha_j}(\nu) < \delta \ \forall \alpha_j \in \boldsymbol{\alpha}\}.$$

Since $AH(M)$ is compact (see [64]) if M is acylindrical, the proof of Lemma 7.1 generalizes directly to give:

Lemma 8.2 *The sets $\mathcal{U}(\delta, \mathbf{U})$, where δ varies in $(0, \epsilon_0)$ and the U_i vary over neighborhoods of λ_i in $\mathcal{C}(S_i)$, are the intersections with B of a neighborhood system for ρ.*

We again define a related set $\mathcal{W}(\epsilon, \mathbf{U})$ where the length bounds on $\boldsymbol{\alpha}$ take place in the conformal boundary:

$$\mathcal{W}(\epsilon, \mathbf{U}) = \{\nu \in B : \pi_{S_i}(\partial_c N_\nu) \in U_i \ \forall i = 1, \dots, l, \\ l_{\alpha_j}(\partial_c N_\nu) < \epsilon \ \forall \alpha_j \in \boldsymbol{\alpha}\}.$$

Again Bers' Lemma 2.1 implies that $\mathcal{W}(\delta/2, \mathbf{U}) \subset \mathcal{U}(\delta, \mathbf{U})$.

The proof of Theorem 8.1 is completed by Lemma 8.3 whose proof mimics that of Lemma 7.2 but must be adapted to account for the fact that σ is bounded rather than constant.

Lemma 8.3 *Given $\delta > 0$ and neighborhoods U_i of λ_i, there exists $\epsilon > 0$ and neighborhoods V_i of λ_i such that any two points in $\mathcal{U}(\epsilon, \mathbf{V})$ can be connected by a path that remains in $\mathcal{U}(\delta, \mathbf{U})$.*

Proof We will assume that S is connected for simplicity, but the general case is handled easily one component at a time.

Notice that if $\gamma \in \mathcal{C}(S)$ and $X = \partial_c N_\nu \in \mathcal{T}(S)$, then $l_\gamma(\nu) = l_\gamma(Q(X, \sigma(X)))$, since $Q(X, \sigma(X))$ is the cover of N_ν associated to $\pi_1(S)$. Since $\sigma(\mathcal{T}(S))$ is bounded, by Thurston's Bounded Image Theorem, there exists R such that for all $W \subset S$,

$$\operatorname{diam} \pi_W(\sigma(\mathcal{T}(S))) < R.$$

By Theorem 2.2, we may choose K such that

$$\mathbf{m}_\gamma(X,Y) > K \implies l_\gamma(\nu) < \delta$$

and also suppose that $K > 2/\delta$, and $K > \max\{1/l_{\alpha_j}(Y) : \alpha_j \in \boldsymbol{\alpha}, Y \in \sigma(\mathcal{T}(S))\}$. Let $h = h(K+R,S)$ be the constant given by Lemma 6.1, and let $c = c(S)$ be the constant in part (3) of Lemma 6.1.

Lemma 5.11 allows us to choose neighborhoods W_i of λ_i such that any two points in $\mathcal{W}(\delta/2,\mathbf{W})$ are connected by a path in $\mathcal{W}(\delta/2,\mathbf{U})$.

Choose $\epsilon > 0$ small enough that (again by Theorem 2.2)

$$l_\gamma(Q(X,Y)) < \epsilon \implies \mathbf{m}_\gamma(X,Y) > h,$$

and choose neighborhoods V_i of λ_i such that a c–neighborhood of V_i in $\mathcal{C}(S_i)$ is contained in W_i.

If $\nu \in \mathcal{U}(\epsilon,\mathbf{V})$ and $X = \partial_c N_\nu \in \mathcal{T}(S)$, then Lemma 6.1 gives a path $\{X_t : t \in [0,T]\}$ beginning at $X = X_0$, such that

(1) $l_{\alpha_j}(X_T) < 1/(K+R) < \delta/2$ for all $\alpha_j \in \boldsymbol{\alpha}$,

(2) $\mathbf{m}_{\alpha_j}(X_t,\sigma(X)) > K+R$ for all $\alpha_j \in \boldsymbol{\alpha}$ and all $t \in [0,T]$, and

(3) $\operatorname{diam}_{\mathcal{C}(S_i)}(\pi_{S_i}(\{X_t\})) < c$.

Let $\{\nu_t \mid t \in [0,T]\}$ be the associated path in B, where $\partial_c N_{\nu_t} = X_t$. Then, (1) and (3) imply that $\nu_T \in \mathcal{W}(\delta/2,\mathbf{W})$. By choice of K, the term of $\mathbf{m}_{\alpha_j}(X_t,\sigma(X))$ that contributes to (2) does not involve $l_{\alpha_j}(\sigma(X))$, and by choice of R all the other terms cannot change by more than R if $\sigma(X)$ is replaced by $\sigma(X_t)$. Hence we have $\mathbf{m}_{\alpha_j}(X_t,\sigma(X_t)) > K$ for all t, so $l_{\alpha_j}(\nu_t) < \delta$ for all t and all $\alpha_j \in \boldsymbol{\alpha}$. Combining this again with (3), we see that the entire path $\{\nu_t\}$ lies in $\mathcal{U}(\delta,\mathbf{W})$.

We can now complete the argument exactly as in the proof of Lemma 7.2. □

9 Surface groups

In this section we prove that quasifuchsian space doesn't self-bump at quasiconformally rigid points in its boundary. The proof is closely modeled on the Bers slice case (Section 7), with the main complication being that we need to keep track of the ordering of the ends, and of the relevant Margulis tubes, during the deformation. Theorem 2.7 allows us to keep track of the ordering of the ends, while Lemma 2.11 will be used to control the ordering of the Margulis tubes.

Theorem 9.1 *If S is a closed surface and ρ is a quasiconformally rigid point in $\partial AH(S\times I)$, then there is no self-bumping at ρ.*

Theorem 1.3 follows from Theorems 8.1 and 9.1.

Proof We begin by constructing a neighborhood system for ρ in $QF(S)$. Let the upward-pointing end invariants of ρ be denoted by a collection $\boldsymbol{\alpha}$ of simple closed curves on S associated to upward-pointing cusps and subsurfaces $\{S_i\}$ with laminations $\{\lambda_i\}$, and let its downward-pointing end invariants be denoted by a collection $\boldsymbol{\beta}$ of simple closed curves on S associated to downward-pointing cusps, and subsurfaces $\{T_k\}$ with laminations $\{\mu_k\}$. For all i and k, let U_i be a neighborhood of $\lambda_i\in\partial_\infty\mathcal{C}(S_i)$ of $\mathcal{C}(S_i)$ and let V_k be a neighborhood of μ_k in $\mathcal{C}(T_k)$. Let $\mathbf{U}$ and $\mathbf{V}$ denote the corresponding tuples of neighborhoods. Define $\mathcal{U}(\delta,\mathbf{U},\mathbf{V})$ to be the set of all quasifuchsian groups $Q(X,Y)$ such that

1. $\pi_{S_i}(X)\in U_i$ for all i,
2. $l_{\alpha_j}(Q(X,Y))<\delta$ for all $\alpha_j\in\boldsymbol{\alpha}$,
3. $\pi_{T_k}(Y)\in V_k$ for all k,
4. $l_{\beta_l}(Q(X,Y))<\delta$ for all $\beta_l\in\boldsymbol{\beta}$, and
5. if $\alpha_j\in\boldsymbol{\alpha}$ and $\beta_l\in\boldsymbol{\beta}$ intersect on S, then α_j lies above β_l in $Q(X,Y)$.

Lemma 9.2 *The sets $\mathcal{U}(\delta,\mathbf{U},\mathbf{V})$ are the intersections with $QF(S)$ of a neighborhood system for ρ.*

Proof As in the proof of Lemma 7.1, it suffices to show that a sequence $\{\rho_n=Q(X_n,Y_n)\}$ converges to ρ if and only if it is eventually contained in any $\mathcal{U}(\delta,\mathbf{U},\mathbf{V})$.

Suppose $\{\rho_n=Q(X_n,Y_n)\}\subset QF(S)$ converges to ρ. Then, by continuity of length, $\lim l_{\alpha_j}(Q(X_n,Y_n))=0$ for all $\alpha_j\in\boldsymbol{\alpha}$ and $\lim l_{\beta_l}(Q(X_n,Y_n))=0$ for all $\beta_l\in\boldsymbol{\beta}$. Theorem 2.7 implies that $\{\pi_{S_i}(X_n)\}$ converges to λ_i for all i and $\{\pi_{T_k}(Y_n)\}$ converges to μ_k for all k. If $\alpha_j\in\boldsymbol{\alpha}$ and $\beta_l\in\boldsymbol{\beta}$ intersect, then Lemma 2.11 ensures that, for all large n, α_j lies above β_l in N_{ρ_n}. Therefore, $\{\rho_n\}$ is eventually contained in any $\mathcal{U}(\delta,\mathbf{U},\mathbf{V})$.

Now suppose that $\{\rho_n\}$ is eventually contained in any $\mathcal{U}(\delta,\mathbf{U},\mathbf{V})$. We must first show that any such $\{\rho_n\}$ has a convergent subsequence in $AH(S)$. If not, then some subsequence, still denoted $\{\rho_n\}$, converges to a small action, by isometries, of $\pi_1(S)$ on an $\mathbb{R}$–tree T, ie there exists $\{\epsilon_n\}$ converging to 0, such that $\{\epsilon_n l_\gamma(\rho_n)\}$ converges to the translation distance $l_T(\gamma)$ of the action of γ on T for any closed curve γ on S (see Morgan and Shalen [52]). Skora's theorem [61] implies that there exists a measured lamination ν on S dual to the tree such that $l_T(\gamma) = i(\nu,\gamma)$ for all γ. If any $\alpha_j \in \boldsymbol{\alpha}$ or $\beta_l \in \boldsymbol{\beta}$ intersects ν, then we obtain an immediate contradiction since $\lim l_{\alpha_j}(\rho_n) = 0$ and $\lim l_{\beta_l}(\rho_n) = 0$. Therefore, ν must be contained both in some S_i and in some T_k. The support of ν cannot agree with both λ_i and μ_k, since λ_i and μ_k do not agree, so ν must intersect either λ_i or μ_k transversely.

Suppose without loss of generality that ν intersects λ_i transversely. We will now show that the geodesics $[\pi_{S_i}(X_n), \pi_{S_i}(Y_n)]$ come uniformly close to a fixed point in $\mathcal{C}(S_i)$ as $n \to \infty$. (Recall that $[a,b]$ refers to *any* geodesic connecting the points a and b). Suppose first that some β_l intersects S_i essentially. Then since $l_{\beta_l}(\rho_n)$ is bounded (in fact goes to 0), Theorem 2.8 gives a D_0 such that $\pi_{S_i}(\beta_l)$ stays within D_0 of $[\pi_{S_i}(X_n), \pi_{S_i}(Y_n)]$. Now suppose that S_i is disjoint from $\boldsymbol{\beta}$, and hence is contained in T_k. Since μ_k fills T_k, it intersects S_i essentially. Let τ_n be a shortest curve on Y_n intersecting T_k essentially, such that $\pi_{T_k}(\tau_n) = \pi_{T_k}(Y_n)$. Since $\pi_{T_k}(Y_n) \to \mu_k$, the Hausdorff limit of $\tau_n \cap T_k$ must contain μ_k. Since μ_k intersects S_i, so must τ_n for high enough n, and moreover eventually $d_{S_i}(\tau_n, \mu_k) \le 1$. Since $l_{\tau_n}(\rho_n)$ is bounded, Theorem 2.8 again tells us that $\pi_{S_i}(\tau_n)$, and hence the fixed point $\pi_{S_i}(\mu_k)$, lie within bounded distance of $[\pi_{S_i}(X_n), \pi_{S_i}(Y_n)]$.

Now, since $\pi_{S_i}(X_n)$ converges to $\lambda_i \in \partial_\infty \mathcal{C}(S_i)$, we see that $d_{S_i}(X_n, Y_n) \to \infty$. Thus for large enough n Theorem 2.8 tells us that $\mathcal{C}(S_i, \rho_n, L_0)$ is nonempty and within bounded Hausdorff distance of $[\pi_{S_i}(X_n), \pi_{S_i}(Y_n)]$. In particular there exists a sequence $\{\gamma_n\} \subset \mathcal{C}(S_i)$ with $\{l_{\gamma_n}(\rho_n)\}$ bounded, and $d_{S_i}(\gamma_n, X_n)$ bounded. The last bound implies that $\gamma_n \to \lambda_i$.

However, the fact that λ_i intersects ν essentially implies, by Corollary 3.1.3 in Otal [59], that λ_i is realizable in the tree T. Since $\gamma_n \to \lambda_i$, Theorem 4.0.1 in Otal [59] then implies that $l_{\gamma_n}(\rho_n) \to \infty$, so we have achieved a contradiction. We conclude that in fact $\{\rho_n\}$ has a convergent subsequence.

Consider any accumulation point ρ' of $\{\rho_n\}$. Each $\alpha_j \in \boldsymbol{\alpha}$ and $\beta_l \in \boldsymbol{\beta}$ is associated to a cusp of $N_{\rho'}$. Theorem 2.7 implies that each S_i is associated to an upward pointing geometrically infinite end with ending lamination λ_i and each T_k is associated to a downward pointing end with ending lamination μ_k. So, there exists a pared homotopy equivalence $h\colon (M_\rho, P_\rho) \to (M_{\rho'}, P_{\rho'})$ which can be taken to be an orientation-preserving homeomorphism on each S_i and T_k. Proposition 8.1 in Canary and Hersonsky [30] implies that there exists a pared homeomorphism $h'\colon (M_\rho, P_\rho) \to (M_{\rho'}, P_{\rho'})$ which agrees with h on each S_i and T_j. In particular, this implies that ρ' is quasiconformally rigid.

In order to apply the Ending Lamination Theorem it remains to check that our pared homeomorphism h' is orientation-preserving. If N_ρ has a geometrically infinite end associated to some S_i or T_k, then h' is orientation-preserving on that surface, so it is orientation-preserving. If N_ρ has no geometrically infinite ends, then it is a maximal cusp. So, each $\alpha_j \in \boldsymbol{\alpha}$ intersects some β_l. As ρ' is quasiconformally rigid and α_j lies above β_l in N_{ρ_n} for all large enough n, Lemma 2.11 implies that α_j is associated to an upward-pointing cusp of $N_{\rho'}$. Similarly, each $\beta_l \in \boldsymbol{\beta}$ is associated to a downward-pointing cusp in $N_{\rho'}$, so h' must be orientation-preserving. The Ending Lamination Theorem then allows us to conclude that $\rho' = \rho$. □

Remark The convergence portion of the above argument can also be derived from the main result of Brock, Bromberg, Canary and Lecuire [15] or by using efficiency of pleated surfaces as in Thurston's proof of the Double Limit Theorem [65].

If $\delta > 0$, $\mathbf{U}$ and $\mathbf{V}$ are as above, then we define $\mathcal{W}(\delta, \mathbf{U}, \mathbf{V})$ to be the set of all quasifuchsian groups $Q(X, Y)$ such that

(1) $\pi_{S_i}(X) \in U_i$ for all i,

(2) $l_{\alpha_j}(X) < \delta$ for all $\alpha_j \in \boldsymbol{\alpha}$,

(3) $\pi_{T_k}(Y) \in V_k$ for all k, and

(4) $l_{\beta_l}(Y) < \delta$ for all $\beta_l \in \boldsymbol{\beta}$.

Lemma 2.5 and Bers' Lemma 2.1 give:

Lemma 9.3 *If $\delta < \epsilon_0$, then $\mathcal{W}(\delta/2, \mathbf{U}, \mathbf{V}) \subset \mathcal{U}(\delta, \mathbf{U}, \mathbf{V})$.*

Lemma 2.5 also allows us to restrict to neighborhoods where the $\alpha_j \in \boldsymbol{\alpha}$ are not short on the bottom conformal boundary component and the $\beta_l \in \boldsymbol{\beta}$ are not short on the top conformal boundary component.

Lemma 9.4 *There exist neighborhoods $(U_i)_0$ of λ_i in $\mathcal{C}(S_i)$ and $(V_k)_0$ of μ_k in $\mathcal{C}(T_i)$ such that if $Q(X, Y) \in \mathcal{U}(\epsilon_0, \mathbf{U}_0, \mathbf{V}_0)$, then $l_{\beta_l}(X) > \epsilon_0$ and $l_{\alpha_j}(Y) > \epsilon_0$ for all $\beta_l \in \boldsymbol{\beta}$ and $\alpha_j \in \boldsymbol{\alpha}$.*

Proof Suppose that $l_{\beta_l}(X) \le \epsilon_0$ for some $\beta_l \in \boldsymbol{\beta}$. If β_l intersects some $\alpha_j \in \boldsymbol{\alpha}$, then Lemma 2.5 would imply that β_l lies above α_j which is a contradiction. If β_l does not intersect any α_j, then it lies in some S_i. Then $d_{S_i}(X, \beta_l) \le 2$. So, if we choose the neighborhood $(U_i)_0$ to have the property that $\pi_{S_i}(\beta_l)$ does not lie in the 2–neighborhood of $(U_i)_0$, we again have a contradiction.

The proof that $l_{\alpha_j}(Y) > \epsilon_0$ for all $\alpha_j \in \boldsymbol{\alpha}$ is similar. □

Theorem 9.1 now follows from:

Lemma 9.5 *Given $\delta > 0$ and neighborhoods U_i of λ_i and V_i of μ_i, there exists $\epsilon > 0$ and neighborhoods U_i'' of λ_i in $\mathcal{C}(S_i)$ and V_k'' of μ_k in $\mathcal{C}(T_k)$ such that any two points in $\mathcal{U}(\epsilon, \mathbf{U}'', \mathbf{V}'')$ can be connected by a path that remains in $\mathcal{U}(\delta, \mathbf{U}, \mathbf{V})$.*

Proof Without loss of generality, we may assume $\delta < \delta_0$ (from Lemma 2.6) and $\mathbf{U} \subset \mathbf{U}_0$, $\mathbf{V} \subset \mathbf{V}_0$ (from Lemma 9.4). By Theorem 2.8, we may further assume that if W is an essential subsurface of S, $\gamma \in \mathcal{C}(S, W)$ and $l_\gamma(Q(X,Y)) < \delta$, then $\pi_W(\gamma)$ lies within D_0 of any geodesic joining $\pi_W(X)$ to $\pi_W(Y)$.

By Theorem 2.2, we may choose K such that

$$\mathbf{m}_\gamma(X,Y) > K \implies l_\gamma(Q(X,Y)) < \delta$$

and also suppose $K > 2/\delta$. Let $h = h(K, S)$ be the constant given by Lemma 6.1, and let $c = c(S)$ be the constant in part (3) of Lemma 6.1. Let d_0 be the constant from Lemma 2.6.

Lemma 5.11 implies that we may choose neighborhoods U_i' of λ_i in $\mathcal{C}(S_i)$ and neighborhoods V_k' of μ_k in $\mathcal{C}(T_k)$ such that any two points in $\mathcal{W}(\delta/2, \mathbf{U}', \mathbf{V}')$ are connected by a path in $\mathcal{W}(\delta/2, \mathbf{U}, \mathbf{V})$. Moreover, we may further assume that if $\beta_l \in \boldsymbol{\beta}$ is contained in S_i, then for all $\gamma \in U_i'$,

$$d_{S_i}(\beta_l, \gamma) > R = m_1 + D_0 + 1.$$

Choose ϵ small enough that (again by Theorem 2.2)

$$l_\gamma(Q(X,Y)) < \epsilon \implies \mathbf{m}_\gamma(X,Y) > h' = h + 2d_0 + D_0 + m_1 + c.$$

Finally, choose neighborhoods U_i'' of λ_i in $\mathcal{C}(S_i)$ and V_k'' of μ_k in $\mathcal{C}(T_k)$ such that a c–neighborhood of U_i'' in $\mathcal{C}(S_i)$ is contained in U_i', and a c–neighborhood of V_k'' in $\mathcal{C}(T_k)$ is contained in V_k'.

Suppose that $Q(X,Y) \in \mathcal{U}(\epsilon, \mathbf{U}'', \mathbf{V}'')$. Then, by our choice of ϵ, $\mathbf{m}_{\alpha_j}(X,Y) > h$ for each $\alpha_j \in \boldsymbol{\alpha}$, and so Lemma 6.1 can be applied to give a path $\{X_t \mid t \in [0,T]\}$ beginning at $X = X_0$ such that

(1) $l_{\alpha_j}(X_T) < 1/K < \delta/2$ for all $\alpha_j \in \boldsymbol{\alpha}$,

(2) $\mathbf{m}_{\alpha_j}(X_t, Y) > K$ for all $\alpha_j \in \boldsymbol{\alpha}$ and $t \in [0,T]$, and

(3) $\operatorname{diam}_{\mathcal{C}(S_i)}(\{\pi_{S_i}(X_t) \mid t \in [0,T]\}) < c$ for all i.

Condition (2), Bers' Lemma 2.1 and our choice of K, give that $l_{\alpha_j}(Q(X_t,Y)) < \delta$ for all $\alpha_j \in \boldsymbol{\alpha}$ and all $t \in [0,T]$. Condition (3) and our choice of U_i'' give that $\pi_{S_i}(X_t) \in U_i'$ for all i and all $t \in [0,T]$.

In order to guarantee that $Q(X_t,Y) \in \mathcal{U}(\delta,\mathbf{U}',\mathbf{V}')$ for all t, it remains to check that $l_{\beta_l}(Q(X_t,Y)) < \delta$ for all $\beta_l \in \boldsymbol{\beta}$ and that each β_l remains correctly ordered with respect to relevant $\alpha_j \in \boldsymbol{\alpha}$. Recall that, again by our choice of ϵ,

$$\mathbf{m}_{\beta_l}(X,Y) > h'$$

for all $\beta_l \in \boldsymbol{\beta}$. We will additionally need to establish that

$$\mathbf{m}_{\beta_l}(X_T,Y) > h. \tag{9-1}$$

Condition (9-1) is necessary to invoke Lemma 6.1 to construct the deformation of the bottom conformal structure Y.

If $l_{\beta_l}(Y) \le 1/h' < \delta/2$ then $\mathbf{m}_{\beta_l}(X_t,Y) > h$ and $l_{\beta_l}(Q(X_t,Y)) < \delta$ for all t by Bers' Lemma 2.1. Lemma 2.5 then implies that if β_l intersects $\alpha_j \in \boldsymbol{\alpha}$ in S, then β_l lies below α_j in $Q(X_t,Y)$ for all t.

If $l_{\beta_l}(Y) > 1/h'$, then, since $l_{\beta_l}(X) > \epsilon_0$ (by Lemma 9.4) and $\mathbf{m}_{\beta_l}(X,Y) \ge h'$, there must be a subsurface Z_l with $\beta_l \subset \partial Z_l$ such that

$$d_{Z_l}(X,Y) > h'.$$

If Z_l does not intersect $\boldsymbol{\alpha}$, then, by Lemma 6.1(3),

$$\operatorname{diam}_{\mathcal{C}(Z_l)}(\{\pi_{Z_l}(X_t) \mid t \in [0,T]\}) < c,$$

so $d_{Z_l}(X_t,Y) > h' - c > h$ for all t and β_l does not intersect $\boldsymbol{\alpha}$. Therefore, $l_{\beta_l}(Q(X_t,Y)) < \delta$ for all t and condition (9-1) holds.

Suppose β_l intersects $\alpha_j \in \boldsymbol{\alpha}$ on S. For each $t \le T$, we know that $l_{\alpha_j}(Q(X_t,Y)) < \delta_0$. Lemma 2.6 asserts that if $d_{Z_l}(X_t,\alpha_j) \ge d_0$, then β_l lies above α_j in $Q(X_t,Y)$. Since β_l lies below α_j in $Q(X,Y)$, we have that $d_{Z_l}(X,\alpha_j) < d_0$, so

$$d_{Z_l}(\alpha_j,Y) \ge d_{Z_l}(X,Y) - d_{Z_l}(X,\alpha_j) > h' - d_0 > h + d_0.$$

It then follows, again from Lemma 2.6 (this time with the roles of X and Y reversed), that α_j lies above β_l in $Q(X_t,Y)$ for all t. So, one must have $d_{Z_l}(X_t,\alpha_j) < d_0$ for all t and hence

$$d_{Z_l}(X_t,Y) \ge d_{Z_l}(\alpha_j,Y) - d_{Z_l}(X_t,\alpha_j) > h$$

which in turn implies that $l_{\beta_j}(Q(X_t,Y)) < \delta$ for all t. In particular, we have established condition (9-1).

It remains to consider the case where Z_l intersects some $\alpha_j \in \boldsymbol{\alpha}$, but β_l does not intersect $\boldsymbol{\alpha}$. In this case, we do not need to worry about the ordering of β_l, but only need to check that $l_{\beta_l}(Q(X_t, Y)) < \delta$ for all $t \in [0, T]$ and verify condition (9-1). Notice that β_l is contained in some S_i. We see that $d_{S_i}(\beta_l, X) > R$, since $\pi_{S_i}(X) \in U_i'$. Since $l_{\beta_l}(Q(X, Y)) < \delta$, β_l lies within D_0 of any geodesic joining $\pi_{S_i}(X)$ to $\pi_{S_i}(Y)$. Therefore,

$$d_{S_i}(X, Y) \geq R - D_0 > m_1.$$

Since $d_{Z_l}(X, Y) > h' > m_1$, Lemma 2.3 implies that S_i and Z_l are $\prec$–ordered in $\mathcal{L}_b(X, Y)$ where $b = \min\{R - D_0, h'\} > m_1$. Since

$$d_{S_i}(X, \partial Z_l) \geq d_{S_i}(X, \beta_l) - 1 > R - 1 > m_1,$$

Lemma 2.3(3) implies that $Z_l \prec S_i$. Thus, Lemma 2.3(2) shows that $d_{Z_l}(\partial S_i, X) \leq m_1$, which implies that $d_{Z_l}(\partial S_i, Y) \geq h' - m_1$. But, since $l_{\alpha_j}(Q(X_t, Y)) < \delta$ if α_j is a component of ∂S_i, we conclude, as above, that

$$d_{Z_l}(X_t, Y) \geq d_{Z_l}(\partial S_i, Y) - D_0 \geq h' - m_1 - D_0 > h$$

for all $t \in [0, T]$. Therefore, $l_{\beta_l}(Q(X_t, Y)) < \delta$ for all t and condition (9-1) holds.

We have considered all cases, so have completed the proof that $Q(X_t, Y) \in \mathcal{U}(\delta, \mathbf{U}', \mathbf{V}')$ for all $t \in [0, T]$.

Now we can fix X_T and apply Lemma 6.1 to the bottom side, obtaining a path $\{Y_t \mid t \in [0, T']\}$ beginning at $Y = Y_0$ such that

(1) $l_{\beta_l}(Q(X_T, Y_t)) < \delta$ for all $t \leq T'$ and $\beta_l \in \boldsymbol{\beta}$,

(2) $\pi_{T_k}(Y_t) \in V_k'$ for all k and $t \leq T'$, and

(3) $l_{\beta_l}(Y_{T'}) < \delta/2$ for all $\beta_l \in \boldsymbol{\beta}$.

Recall that $\pi_{S_i}(X_T) \in U_i$ for all i and $l_{X_T}(\alpha_j) < \delta/2 < \epsilon_0$ for all $\alpha_j \in \boldsymbol{\alpha}$. Therefore, Lemma 2.5 implies that α_j lies above β_l in $Q(X_T, Y_t)$ for all $t \in [0, T']$ whenever α_j and β_l intersect on S. Therefore, the path $\{Q(X_T, Y_t) \mid t \in [0, T']\}$ lies entirely in $\mathcal{U}(\delta, \mathbf{U}', \mathbf{V}')$. The concatenation of the paths $\{Q(X_t, Y) \mid t \in [0, T]\}$ and $\{Q(X_T, Y_t) \mid t \in [0, T']\}$ remains in $\mathcal{U}(\delta, \mathbf{U}', \mathbf{V}')$, and joins $Q(X, Y)$ to a point $Q(X_T, Y_{T'}) \in \mathcal{W}(\delta/2, \mathbf{U}', \mathbf{V}')$.

Since any two points in $\mathcal{W}(\delta/2, \mathbf{U}', \mathbf{V}')$ can be connected by a path in $\mathcal{W}(\delta/2, \mathbf{U}, \mathbf{V})$, and since $\mathcal{W}(\delta/2, \mathbf{U}, \mathbf{V}) \subset \mathcal{U}(\delta, \mathbf{U}, \mathbf{V})$, by Lemma 9.3, we conclude that any two points in $\mathcal{U}(\epsilon, \mathbf{U}'', \mathbf{V}'')$ can be connected by a path in $\mathcal{U}(\delta, \mathbf{U}, \mathbf{V})$. □

References

[1] **I Agol**, *Tameness of hyperbolic 3–manifolds* `arXiv:math.GT/0405568`

[2] **J W Anderson**, **R D Canary**, *Algebraic limits of Kleinian groups which rearrange the pages of a book*, Invent. Math. 126 (1996) 205–214 MR1411128

[3] **J W Anderson**, **R D Canary**, *Cores of hyperbolic 3–manifolds and limits of Kleinian groups. II*, J. London Math. Soc. (2) 61 (2000) 489–505 MR1760675

[4] **J W Anderson**, **R D Canary**, **M Culler**, **P B Shalen**, *Free Kleinian groups and volumes of hyperbolic 3–manifolds*, J. Differential Geom. 43 (1996) 738–782 MR1412683

[5] **J W Anderson**, **R D Canary**, **D McCullough**, *The topology of deformation spaces of Kleinian groups*, Ann. of Math. (2) 152 (2000) 693–741 MR1815699

[6] **J Behrstock**, **B Kleiner**, **Y N Minsky**, **L Mosher**, *Geometry and rigidity of mapping class groups* `arXiv:0801.2006`

[7] **R Benedetti**, **C Petronio**, *Lectures on hyperbolic geometry*, Universitext, Springer, Berlin (1992) MR1219310

[8] **L Bers**, *On boundaries of Teichmüller spaces and on Kleinian groups. I*, Ann. of Math. (2) 91 (1970) 570–600 MR0297992

[9] **L Bers**, *On spaces of Riemann surfaces with nodes*, Bull. Amer. Math. Soc. 80 (1974) 1219–1222 MR0361165

[10] **L Bers**, *An inequality for Riemann surfaces*, from: "Differential geometry and complex analysis", (I Chavel, H M Farkas, editors), Springer, Berlin (1985) 87–93 MR780038

[11] **F Bonahon**, *Bouts des variétés hyperboliques de dimension 3*, prepublicationes de Orsay (1985)

[12] **F Bonahon**, *Bouts des variétés hyperboliques de dimension 3*, Ann. of Math. (2) 124 (1986) 71–158 MR847953

[13] **J F Brock**, *Boundaries of Teichmüller spaces and end-invariants for hyperbolic 3–manifolds*, Duke Math. J. 106 (2001) 527–552 MR1813235

[14] **J F Brock**, **K W Bromberg**, *On the density of geometrically finite Kleinian groups*, Acta Math. 192 (2004) 33–93 MR2079598

[15] **J F Brock**, **K W Bromberg**, **R D Canary**, **C Lecuire**, *Convergence and divergence of Kleinian surface groups*, in preparation

[16] **J F Brock**, **K W Bromberg**, **R D Canary**, **Y N Minsky**, *Convergence properties of ending invariants*, in preparation (2011)

[17] **J F Brock**, **R D Canary**, **Y N Minsky**, *The classification of Kleinian surface groups II: The Ending Lamination Conjecture* `arXiv:math.GT/0412006`

[18] **J F Brock**, **R D Canary**, **Y N Minsky**, *The classification of finitely generated Kleinian groups*, in preparation (2011)

[19] **J F Brock**, **J Souto**, *Algebraic limits of geometrically finite manifolds are tame*, Geom. Funct. Anal. 16 (2006) 1–39 MR2221251

[20] **K W Bromberg**, *Projective structures with degenerate holonomy and the Bers density conjecture*, Ann. of Math. (2) 166 (2007) 77–93 MR2342691

[21] **K W Bromberg**, *The space of Kleinian punctured torus groups is not locally connected*, Duke Math. J. 156 (2011) 387–427 MR2772066

[22] **K W Bromberg**, **J Holt**, *Self-bumping of deformation spaces of hyperbolic* 3*–manifolds*, J. Differential Geom. 57 (2001) 47–65 MR1871491

[23] **K W Bromberg**, **J Souto**, *The density conjecture: a prehistoric approach*, in preparation (2011)

[24] **P Buser**, *Geometry and spectra of compact Riemann surfaces*, Progress in Math. 106, Birkhäuser, Boston (1992) MR1183224

[25] **D Calegari**, **D Gabai**, *Shrinkwrapping and the taming of hyperbolic* 3*–manifolds*, J. Amer. Math. Soc. 19 (2006) 385–446 MR2188131

[26] **R D Canary**, *Ends of hyperbolic* 3*–manifolds*, J. Amer. Math. Soc. 6 (1993) 1–35 MR1166330

[27] **R D Canary**, *A covering theorem for hyperbolic* 3*–manifolds and its applications*, Topology 35 (1996) 751–778 MR1396777

[28] **R D Canary**, *Introductory bumponomics: the topology of deformation spaces of hyperbolic 3-manifolds*, from: "Teichmüller theory and moduli problem", (I Biswas, R S Kulkarni, S Mitra, editors), Ramanujan Math. Soc. Lect. Notes Ser. 10, Ramanujan Math. Soc., Mysore (2010) 131–150 MR2667553

[29] **R D Canary**, **D B A Epstein**, **P Green**, *Notes on notes of Thurston*, from: "Analytical and geometric aspects of hyperbolic space (Coventry/Durham, 1984)", (D B A Epstein, editor), London Math. Soc. Lecture Note Ser. 111, Cambridge Univ. Press (1987) 3–92 MR903850

[30] **R D Canary**, **S Hersonsky**, *Ubiquity of geometric finiteness in boundaries of deformation spaces of hyperbolic* 3*–manifolds*, Amer. J. Math. 126 (2004) 1193–1220 MR2102392

[31] **R D Canary**, **D McCullough**, *Homotopy equivalences of* 3*–manifolds and deformation theory of Kleinian groups*, Mem. Amer. Math. Soc. 172, no. 812, Amer. Math. Soc. (2004) MR2096234

[32] **R D Canary**, **Y N Minsky**, *On limits of tame hyperbolic* 3*–manifolds*, J. Differential Geom. 43 (1996) 1–41 MR1424418

[33] **D B A Epstein**, **A Marden**, **V Markovic**, *Quasiconformal homeomorphisms and the convex hull boundary*, Ann. of Math. (2) 159 (2004) 305–336 MR2052356

[34] **B Farb**, **A Lubotzky**, **Y Minsky**, *Rank–*1 *phenomena for mapping class groups*, Duke Math. J. 106 (2001) 581–597 MR1813237

[35] **U Hamenstädt**, *Train tracks and the Gromov boundary of the complex of curves*, from: "Spaces of Kleinian groups", (Y N Minsky, M Sakuma, C Series, editors), London Math. Soc. Lecture Note Ser. 329, Cambridge Univ. Press (2006) 187–207 MR2258749

[36] **K Johannson**, *Homotopy equivalences of* 3*–manifolds with boundaries*, Lecture Notes in Math. 761, Springer, Berlin (1979) MR551744

[37] **T Jørgensen**, **A Marden**, *Algebraic and geometric convergence of Kleinian groups*, Math. Scand. 66 (1990) 47–72 MR1060898

[38] **R P Kent, IV**, *Skinning maps*, Duke Math. J. 151 (2010) 279–336 MR2598379

[39] **S P Kerckhoff**, *The Nielsen realization problem*, Ann. of Math. (2) 117 (1983) 235–265 MR690845

[40] **S P Kerckhoff**, **W P Thurston**, *Noncontinuity of the action of the modular group at Bers' boundary of Teichmüller space*, Invent. Math. 100 (1990) 25–47 MR1037141

[41] **E Klarreich**, *The boundary at infinity of the curve complex and the relative Teichmüller space*, preprint (1999)

[42] **R S Kulkarni**, **P B Shalen**, *On Ahlfors' finiteness theorem*, Adv. Math. 76 (1989) 155–169 MR1013665

[43] **A D Magid**, *Deformation spaces of Kleinian surface groups are not locally connected*, PhD thesis, University of Michigan (2009) `arXiv:1003.4541`

[44] **A Marden**, *Geometric relations between homeomorphic Riemann surfaces*, Bull. Amer. Math. Soc. (N.S.) 3 (1980) 1001–1017 MR585180

[45] **B Maskit**, *On the classification of Kleinian groups. I: Koebe groups*, Acta Math. 135 (1975) 249–270 MR0444942

[46] **H A Masur**, **Y N Minsky**, *Geometry of the complex of curves. I: Hyperbolicity*, Invent. Math. 138 (1999) 103–149 MR1714338

[47] **H A Masur**, **Y N Minsky**, *Geometry of the complex of curves. II: Hierarchical structure*, Geom. Funct. Anal. 10 (2000) 902–974 MR1791145

[48] **D McCullough**, *Compact submanifolds of 3–manifolds with boundary*, Quart. J. Math. Oxford Ser. (2) 37 (1986) 299–307 MR854628

[49] **C T McMullen**, *Complex earthquakes and Teichmüller theory*, J. Amer. Math. Soc. 11 (1998) 283–320 MR1478844

[50] **Y N Minsky**, *The classification of punctured-torus groups*, Ann. of Math. (2) 149 (1999) 559–626 MR1689341

[51] **Y N Minsky**, *The classification of Kleinian surface groups. I: Models and bounds*, Ann. of Math. (2) 171 (2010) 1–107 MR2630036

[52] **J W Morgan**, **P B Shalen**, *Valuations, trees, and degenerations of hyperbolic structures. I*, Ann. of Math. (2) 120 (1984) 401–476 MR769158

[53] **H Namazi**, **J Souto**, *Non-realizability and ending laminations: Proof of the Density Conjecture*, preprint (2010) Available at `http://www.ma.utexas.edu/users/hossein`

[54] **K Ohshika**, *Divergence, exotic convergence and self-bumping in quasi-Fuchsian spaces* `arXiv:1010.0070`

[55] **K Ohshika**, *On limits of quasi-conformal deformations of Kleinian groups*, Math. Z. 201 (1989) 167–176 MR997219

[56] **K Ohshika**, *Strong convergence of Kleinian groups and Carathéodory convergence of domains of discontinuity*, Math. Proc. Cambridge Philos. Soc. 112 (1992) 297–307 MR1171166

[57] **K Ohshika**, *Kleinian groups which are limits of geometrically finite groups*, Mem. Amer. Math. Soc. 177, no. 834, Amer. Math. Soc. (2005) MR2154090

[58] **K Ohshika**, *Realising end invariants by limits of minimally parabolic, geometrically finite groups*, Geom. Topol. 15 (2011) 827–890

[59] **J-P Otal**, *Le théorème d'hyperbolisation pour les variétés fibrées de dimension 3*, Astérisque 235, Soc. Math. France (1996) MR1402300

[60] **J-P Otal**, *Les géodésiques fermées d'une variété hyperbolique en tant que nœuds*, from: "Kleinian groups and hyperbolic 3–manifolds (Warwick, 2001)", (Y Komori, V Markovic, C Series, editors), London Math. Soc. Lecture Note Ser. 299, Cambridge Univ. Press (2003) 95–104 MR2044546

[61] **R K Skora**, *Splittings of surfaces*, J. Amer. Math. Soc. 9 (1996) 605–616 MR1339846

[62] **W P Thurston**, *The geometry and topology of three-manifolds*, Princeton Univ. Math. Dept. Lecture Notes (1979) Available at `http://msri.org/publications/books/gt3m/`

[63] **W P Thurston**, *Earthquakes in two-dimensional hyperbolic geometry*, from: "Low-dimensional topology and Kleinian groups (Coventry/Durham, 1984)", (D B A Epstein, editor), London Math. Soc. Lecture Note Ser. 112, Cambridge Univ. Press (1986) 91–112 MR903860

[64] **W P Thurston**, *Hyperbolic structures on 3–manifolds. I: Deformation of acylindrical manifolds*, Ann. of Math. (2) 124 (1986) 203–246 MR855294

[65] **W P Thurston**, *Hyperbolic structures on 3-manifolds, II: Surface groups and 3–manifolds which fiber over the circle* `arXiv:math.GT/9801045`

[66] **W P Thurston**, *Hyperbolic structures on 3–manifolds, III: Deformations of 3–manifolds with incompressible boundary* `arXiv:math.GT/9801045`

[67] **F Waldhausen**, *On irreducible 3–manifolds which are sufficiently large*, Ann. of Math. (2) 87 (1968) 56–88 MR0224099

Department of Mathematics, Brown University
Box 1917, Providence RI 02912, USA

Department of Mathematics, University of Utah
Salt Lake City UT 84112, USA

Department of Mathematics, University of Michigan, Ann Arbor
2074 East Hall, 530 Church St, Ann Arbor MI 48109-1043, USA

Department of Mathematics, Yale University
10 Hillhouse Ave, New Haven CT 06520-8283, USA

`brock@math.brown.edu, bromberg@math.utah.edu, canary@umich.edu, yair.minsky@yale.edu`

`http://www.math.brown.edu/~brock, http://www.math.utah.edu/~bromberg, http://www.math.lsa.umich.edu/~canary, http://www.math.yale.edu/users/yair`

Proposed: Jean-Pierre Otal — Received: 12 November 2009
Seconded: Benson Farb, Danny Calegari — Accepted: 14 April 2011