

WINDOWS, CORES AND SKINNING MAPS

JEFF BROCK, KEN BROMBERG, DICK CANARY, AND YAIR MINSKY

ABSTRACT. We give a generalization of Thurston's Bounded Image Theorem for skinning maps, which applies to pared 3-manifolds with incompressible boundary that are not necessarily acylindrical. Along the way we study properties of divergent sequences in the deformation space of such a manifold, establishing the existence of compact cores satisfying a certain notion of uniform geometry.

1. Introduction

Critical to Thurston's Geometrization Theorem for Haken 3-manifolds was a fixed-point problem, phrased for a self-mapping of the deformation space of a hyperbolic 3-manifold with boundary. This *skinning map* implicitly describes how to enhance a topological gluing of a 3-manifold along its boundary with geometric information; a fixed point for the skinning map realizes a geometric solution to the gluing problem, resulting in a hyperbolic structure on the gluing.

Beyond its utility in geometrization, the map itself reveals more quantitatively the relationship between topological and geometric features of a hyperbolic 3-manifold. Indeed, Thurston's Bounded Image Theorem (see Thurston [45], Morgan [41], Kent [30]) which provides the desired fixed-point, guarantees that for an acylindrical 3-manifold, the image of the skinning map is a bounded subset of Teichmüller space. In this paper, we investigate the skinning map in the more general case when the 3-manifold is only assumed to have incompressible boundary. Here, the image of the skinning map need not be bounded, but the restriction onto any essential subsurface of the boundary that is homotopic off of the characteristic submanifold is bounded. One may think of this as a strong form of Thurston's Relative Compactness Theorem. Along the way, we refine our understanding of the interior geometry of a hyperbolic 3-manifold, establishing existence of a uniformly controlled family of compact cores for each deformation space of such manifolds.

Let M be a compact orientable hyperbolizable 3-manifold whose boundary is incompressible and contains no tori and let $CC_0(M)$ denote the set of convex cocompact hyperbolic structures on the interior of M . A convex cocompact hyperbolic structure N_ρ on the interior of M gives rise to a holonomy representation $\rho : \pi_1(M) \rightarrow \mathrm{PSL}(2, \mathbb{C})$ and induces a well-defined conformal structure on its

Date: January 20, 2016.

Brock was partially supported by NSF grant DMS-1207572, Bromberg was partially supported by NSF grants DMS-1207873 and DMS-1509171, Canary was partially supported by NSF grant DMS -1306992, and Minsky was partially supported by NSF grant DMS-1311844.

boundary ∂M . Bers [8] shows that one obtains an identification of $CC_0(M)$ with $\mathcal{T}(\partial M)$, the Teichmüller space of conformal structures on ∂M . The skinning map

$$\sigma_M : CC_0(M) \rightarrow \mathcal{T}(\overline{\partial M})$$

records the asymptotic geometry of the “inward-pointing” end of the cover associated to each boundary component. If M has connected boundary S and N_ρ is in $CC_0(M)$, then the cover N_S of N_ρ associated to $\pi_1(S)$ is quasifuchsian, i.e. a point in $CC_0(S \times [0, 1])$, so may be identified to a point $(X, Y) \in \mathcal{T}(S) \times \mathcal{T}(\overline{S})$. Then, $\sigma_M(\rho) = Y$ and X is the point in $\mathcal{T}(S)$ associated to ρ by the Bers parametrization. (The skinning map will be defined more carefully, and in greater generality, in Section 2.1.)

A compact, orientable, hyperbolizable 3-manifold M is said to be *acylindrical* if it contains no essential annuli, or, equivalently, if $\pi_1(M)$ does not admit a non-trivial splitting over a cyclic subgroup. In this setting, Thurston’s Bounded Image Theorem has the following form.

Thurston’s Bounded Image Theorem: *If M is a compact, orientable, acylindrical, hyperbolizable 3-manifold with no torus boundary components, then the skinning map $\sigma_M : CC_0(M) \rightarrow \mathcal{T}(\overline{\partial M})$ has bounded image.*

The skinning map has been studied extensively when M is acylindrical. This study has focussed on obtaining bounds on the diameter of the skinning map in terms of the topology of M and the geometry of its unique hyperbolic metric with totally geodesic boundary, see Kent [30] and Kent-Minsky [31]. In the case that M is not required to be acylindrical, it is known that the skinning map is non-constant (Dumas-Kent [24]) and that its image is Zariski dense in the appropriate character variety (Dumas-Kent [25]). Recently, Dumas [23] has shown that the skinning map is finite-to-one.

In order to state our generalization of Thurston’s Bounded Image Theorem to the setting where M is only assumed to have incompressible boundary, we recall that the *characteristic submanifold* $\Sigma(M)$ is a minimal collection of solid tori and interval bundles in M whose frontier is a collection of essential annuli such that every (embedded) essential annulus in M is isotopic into $\Sigma(M)$ (see Johannson [28] or Jaco-Shalen [27]). Thurston [47] defines the *window* of M to be the union of the interval bundles in $\Sigma(M)$, together with a regular neighborhood of each component of the frontier of $\Sigma(M)$ which is not homotopic into an interval bundle. Let $\partial_{nw}M$ denote the intersection of ∂M with the complement of the window. Components of $\partial_{nw}M$ are either components of the intersection of ∂M with the (relatively) acylindrical pieces of the decomposition, or annuli in the boundaries of the solid torus pieces.

Our main theorem asserts that any curve in $\partial_{nw}M$, equivalently any curve in ∂M which may be homotoped off the characteristic submanifold, has uniformly bounded length in the hyperbolic structures which arise in the skinning image.

Theorem 1.1. *Let M be a compact, orientable, hyperbolizable 3-manifold whose boundary is incompressible. For each curve α in $\partial_{nw}M$, its length $\ell_Y(\alpha)$ is bounded as Y varies over the image of σ_M .*

If W is an essential subsurface of ∂M , then σ_M induces a map

$$\sigma_M^W : CC_0(M) \rightarrow \mathcal{F}(W)$$

where $\mathcal{F}(W)$ is the Fricke space of all hyperbolic structures on the interior of W . Notice that these hyperbolic structures are allowed to have either finite or infinite area. Our main theorem immediately translates to the fact that σ_M^W has bounded image:

Corollary 1.2. *Let M be a compact, orientable, hyperbolizable 3-manifold whose boundary is incompressible. If W is a component of $\partial_{nw}M$, then the image of σ_M^W is bounded in $\mathcal{F}(W)$.*

Note that these statements allow ∂M to have torus components, and moreover the theorem applies more generally, when we consider the space $AH_0(M)$ of all hyperbolic structures on the interior of M . Given an element of $AH_0(M)$ with holonomy representation $\rho : \pi_1(M) \rightarrow PSL(2, \mathbb{C})$ and quotient manifold N_ρ , we obtain an end invariant $\sigma_M(\rho)$ on the non-toroidal part of the boundary, which records the asymptotic geometry of the inward-pointing end of the covers of N_ρ associated to those boundary components of M . This ending invariant consists of a multicurve, known as the parabolic locus, and either a finite area hyperbolic structure or a filling lamination on each component of the complement. If a curve α on ∂M is homotopic into a component of the complement of the parabolic locus which has a hyperbolic structure, then $\ell_{\sigma_M(\rho)}(\alpha)$ denotes the length of the geodesic representative of α in this hyperbolic structure; if α is homotopic into the parabolic locus then $\ell_{\sigma_M(\rho)}(\alpha) = 0$; and otherwise $\ell_{\sigma_M(\rho)}(\alpha) = +\infty$.

With these definitions, Theorem 1.1 continues to hold. In addition, σ_M^W can still be defined on $AH_0(M)$ when W is a component of $\partial_{nw}M$, and Corollary 1.2 holds as well. Notice that Corollary 1.2 contains Thurston's original Bounded Image Theorem as a special case. Both results also have natural generalizations to the pared setting, which we will state and prove in Section 5.

Along the way to proving Theorem 1.1, we will develop some tools and structural results on the geometry of hyperbolic 3-manifolds that may be of independent interest. One is a new tool for studying the geometry of surface groups, and the other is a uniformity statement for compact cores of manifolds in a deformation space $AH_0(M)$.

Length bounds in surface groups. In [40, 15], Brock-Canary-Minsky give an explicit connection between the geometry at infinity of a surface group $[\rho] \in AH(S)$ and its internal geometry, by means of a “model manifold” combinatorially built up from the ending data $v(\rho)$. In Section 3 we use the structure of this model, together with some refinements proved in [14], to “reverse engineer” a usable criterion for bounding the length at infinity of a curve in S , based on the situation in the interior.

Specifically, let α be a simple curve in S , and $\hat{\alpha}$ a representative of α in the 3-manifold N_ρ , which is contained in a “level surface” in a product structure on N_ρ . Let $\mathcal{C}(\hat{\alpha}, L)$ denote the set of curves on S which intersect α essentially and whose geodesic representatives in N_ρ have length at most L and do not “lie above”

$\widehat{\alpha}$ (in a sense to be made precise in Section 2.4). We will show that, given a length bound on $\widehat{\alpha}$ and constraints on $\mathcal{C}(\widehat{\alpha}, L)$, we can obtain bounds on the length of α in the “bottom” conformal structure $v^-(\rho)$.

The idea here is that, in order for α to be considerably longer in $v^-(\rho)$ than $\widehat{\alpha}$ is, there must be some kind of geometric complexity between $\widehat{\alpha}$ and the bottom end of N_ρ , and this is what $\mathcal{C}(\widehat{\alpha}, L)$ captures. We give a loose statement of the theorem here, in the closed case; see section 3 for a more carefully quantified general version.

Theorem 3.1. *Let S be a closed surface and α an essential simple closed curve in S . Let $\widehat{\alpha}$ be a representative of α in N_ρ for $[\rho] \in AH(S)$ which is contained in a level surface that avoids the thin part of N_ρ . Given a length bound on $\widehat{\alpha}$, an upper bound on the number of elements of $\mathcal{C}(\widehat{\alpha}, L)$ (for suitable L) and a lower bound on the length of all elements of $\mathcal{C}(\widehat{\alpha}, L)$, we obtain an upper bound on $\ell_{v^-(\rho)}(\alpha)$.*

Uniform Core Models. If M is acylindrical, then it follows from Thurston’s Compactness Theorem [46] and work of Anderson-Canary [4] that there is a fixed metric on M such that for each $[\rho] \in AH_0(M)$ there is a uniformly bilipschitz embedding of M into N_ρ in the homotopy class of ρ . The image of this embedding is a compact core with uniformly bounded geometry. In our setting this is no longer possible, since a sequence of hyperbolic manifolds in $AH_0(M)$ may “pull apart” along Margulis tubes. Moreover, one must take into account the action of the outer automorphism group, which will be infinite if M is not acylindrical.

In Section 4 we define a notion of a *model core* for an element $[\rho] \in AH_0(M)$. Roughly, a model core is a metric m on the complement $M_{\mathcal{A}}$ of a collection of solid and thickened tori in M . We say that a model core *controls* N_ρ if there is an embedding of M into N_ρ which is 2-bilipschitz with respect to m on $M_{\mathcal{A}}$, takes the components of $M \setminus M_{\mathcal{A}}$ into the thin part of N_ρ , and lies in the homotopy class of $\rho \circ \phi_*$ where $\phi : M \rightarrow M$ is a homeomorphism which is the identity on the complement of the characteristic submanifold. The image of the embedding is a compact core for N_ρ with uniformly bounded geometry on $M_{\mathcal{A}}$ (and, in particular, on the complement of the characteristic submanifold).

Theorem 4.1. *If M is a hyperbolizable 3-manifold with incompressible boundary then there is a finite collection of model cores so that each element of $AH_0(M)$ is controlled by one of them.*

The proof of Theorem 4.1 utilizes a version of Thurston’s Relative Compactness Theorem [47] from Canary-Minsky-Taylor [22] and the analysis of the relationship between algebraic and geometric limits of sequences of isomorphic Kleinian groups carried out by Anderson, Canary, and McCullough [4, 6].

Summary of Proof of Theorem 1.1: It is instructive to first think about our argument in the case where M is acylindrical, which is the setting of Thurston’s original Bounded Image Theorem. We will also simplify the discussion by assuming that $\partial M = S$ is connected. Let α be a fixed closed curve in S .

Theorem 4.1 gives us, in this case, a finite collection $\{C_1, \dots, C_r\}$ of compact Riemannian manifolds with boundary so that any $[\rho] \in AH(M)$ has a compact core

which is 2-bilipschitz to some C_i . (In the acylindrical case this can be directly obtained from Thurston's Compactness Theorem [46] and the work of Anderson-Canary [4] on cores and limits). For such a compact core C , its boundary F lifts to a level surface \widehat{F} for the cover N_S associated to $\rho(\pi_1(S))$, which contains a representative $\widehat{\alpha}$ of α whose length is uniformly bounded. Let \widehat{C} be the component of the pre-image of C in N_ρ^S which contains \widehat{F} in its boundary. Since M is acylindrical, each component of $N_S \setminus \widehat{C}$ which lies below $\widehat{\alpha}$ is simply connected. Therefore, every bounded length closed geodesic which is not above $\widehat{\alpha}$ must intersect \widehat{C} . Moreover, again since M is acylindrical, distinct closed geodesics in N_S project to distinct geodesics in N_ρ . Since C has uniformly bounded geometry we obtain an upper bound on the number of such geodesics and a lower bound on their length. Theorem 3.1 then implies that $\ell_{v^-(\rho^S)}(\alpha)$, which is exactly α 's length in the skinning structure, is uniformly bounded.

In order to generalize this proof to the setting where M is only assumed to have incompressible boundary, we need the more general statement of Theorem 4.1 about model cores. This theorem gives us, for $[\rho] \in AH_0(M)$, a compact core C in N_ρ whose geometry is uniformly controlled on the ‘‘non-window’’ part (as well as on parts of the window). Picking a component F of ∂C and a curve α on the non-window part of F , we can again lift to the corresponding cover N_ρ^F . Now we learn that a bounded length geodesic lying below \widehat{F} will either intersect \widehat{C} or lie in a complementary region of \widehat{C} which retracts to the window part. In the latter case this geodesic cannot represent a curve that intersects α essentially, and hence does not count as a member of $\mathcal{C}(\widehat{\alpha}, L)$. An application of Theorem 3.1, with a bit of additional care, again gives us the desired bound.

Relationship to proofs of Thurston's Bounded Image Theorem: Our proof is new even in the acylindrical setting. The only known proof of Thurston's original Bounded Image Theorem, due to Kent [30], uses the work of Anderson-Canary [4] to extend the skinning map continuously to all of $AH(M)$ and then applies Thurston's Compactness Theorem, which gives that $AH(M)$ is compact, to conclude that the image is bounded. In our more general setting, the skinning map does not extend continuously to $AH(M)$, in any reasonable sense, and $AH(M)$ fails to be compact, so it was necessary to develop a new approach to the proof.

Potential applications: The model manifolds constructed in [40, 15] are bilipschitz equivalent to N_ρ for any $[\rho] \in AH(M)$, but the bilipschitz constant is not in general uniform over all of $AH(M)$. For the case of surface groups, we do have uniformity once the genus is fixed, but a uniform model construction that would apply in general would be helpful for a better global understanding of $AH(M)$.

When M is acylindrical, Thurston's Bounded Image Theorem suggests a natural candidate for this uniform model: In this case, there is a unique hyperbolic 3-manifold C_M with totally geodesic boundary homeomorphic to M . Let $Y \in \mathcal{T}(\partial M)$ denote the conformal structure of this boundary. A model for the hyperbolic

3-manifold in $CC_0(M)$ with conformal structure $X \in \mathcal{T}(\partial M)$ can then be assembled from C_M and the quasifuchsian hyperbolic 3-manifold with conformal structure (X, Y) . One expects Thurston's Bounded Image Theorem and the Bilipschitz Model Manifold Theorem from [15] to be the main tools for showing that this model is uniform over $CC_0(M)$.

Our relative bounded image theorem should play a similar role in constructing uniform models for hyperbolic 3-manifolds with finitely generated, freely indecomposable fundamental group. The complement of the characteristic submanifold admits a cusped hyperbolic metric with totally geodesic boundary, and this should play the role of the manifold C_M .

The relative bounded image theorem should also be a tool in the study of the local topology of $AH(M)$. We recall that Bromberg [16] and Magid [33] showed that $AH(S)$ fails to be locally connected. In earlier work [13], we used Thurston's Bounded Image Theorem to show that if M is acylindrical and $[\rho] \in AH(M)$ is quasiconformally rigid (i.e. every component of the conformal boundary of N_ρ is a thrice-punctured sphere), then $AH(M)$ is locally connected at ρ . We hope our relative bounded image theorem will allow us to attack the following conjecture.

Conjecture: *Let M be a compact, orientable, hyperbolizable 3-manifold whose boundary is incompressible. If $[\rho] \in AH(M)$ is quasiconformally rigid, then $AH(M)$ is locally connected at ρ .*

The relative bounded image theorem may also be useful in establishing a particularly mysterious step in Thurston's original proof of his Geometrization Theorem (see Morgan [41, Section 10]). If $\tau : M \rightarrow M$ is an orientation-reversing homeomorphism so that the 3-manifold M_τ is atoroidal, then τ induces a map $\tau_\# : \mathcal{T}(\partial M) \rightarrow \mathcal{T}(\overline{\partial M})$. (Here we allow M to have two components and make the natural adjustments to the notation.) Thurston claimed the following result, which, along with Maskit's combination theorems [34], completed the final step of the proof of his Geometrization Theorem.

Theorem: (Thurston) *If M_τ is atoroidal, then there exists $n \in \mathbb{N}$ so that $(\tau_\# \circ \sigma_M)^n$ has bounded image.*

As far as we know, no one currently knows a proof of Thurston's result. All written versions of Thurston's proof of hyperbolization only produce a fixed point of $\tau_\# \circ \sigma_M$ (which suffices to hyperbolize M_τ), see for example Kapovich [29], McMullen [38], Morgan [41] and Otal [42].

2. Background

In this section we recall some standard ideas and notation, with a few new notions particular to this paper. Section 2.1 is a brief review of deformation spaces, end invariants and the definition of skinning maps. Section 2.2 reviews pared manifolds and the Jaco-Shalen Johannson theory of characteristic submanifolds, and introduces the notion of *robust systems of annuli* which give useful ways to cut up 3-manifolds that correspond to the way in which sequences of hyperbolic structures diverge along cusps. Section 2.3 discusses algebraic and geometric convergence for

representations, and introduces language for describing divergent sequences that converge on subgroups. Section 2.4 recalls the notion of markings and hierarchies for curve complexes of surfaces, which are the combinatorial ingredients for the model manifold construction of [40] and are needed here in the proof of Theorem 3.1. Section 2.5 shows how to use Thurston's Relative Compactness Theorem to extract convergent behavior from a diverging sequence of elements of $AH_0(M, P)$. Corollary 2.9 is the main result we will need, adapting ideas from Canary-Minsky-Taylor [22] to provide a robust system of annuli for such a sequence so that a subsequence converges in a suitable sense on its complementary pieces.

2.1. Deformation spaces and skinning maps

We begin with a quick review of the terminology for deformation spaces of hyperbolic 3-manifolds.

Hyperbolic 3-manifolds. A hyperbolic 3-manifold is the quotient $N = \mathbb{H}^3/\Gamma$ of hyperbolic 3-space by a group Γ of isometries acting freely and properly discontinuously. We will always assume that N is orientable, so

$$\Gamma \subset \text{Isom}_+(\mathbb{H}^3) \cong \text{PSL}(2, \mathbb{C}).$$

The *convex core* $C(N)$ of N is the smallest convex submanifold of N .

There exists a constant $\mu_3 > 0$, called the *Margulis constant*, such that if $\varepsilon < \mu_3$ and

$$N_{(0,\varepsilon)} = \{x \in N \mid \text{inj}_N(x) < \varepsilon\}$$

is the set of points in N with injectivity radius less than ε , then every component of $N_{(0,\varepsilon)}$ is a *Margulis region*, i.e. either a *Margulis tube*, (a metric neighborhood of a closed geodesic) or a *cuspidal region* (a quotient of a horoball by a parabolic subgroup). If γ is a nontrivial element of $\pi_1(N)$, we denote by $\mathbb{T}_\varepsilon(\gamma)$ the component of $N_{(0,\varepsilon)}$, possibly empty, associated to the maximal abelian subgroup containing γ , and similarly define $\mathbb{T}_\varepsilon(a)$ (or $\mathbb{T}_\varepsilon(A)$) if a is a curve in N (or A is an incompressible annulus in N). We fix an explicit value $\varepsilon_0 \in (0, \mu_3)$ which we use as a default value for Margulis regions. (For convenience, we choose the value ε_0 used in the Bilipschitz Model Theorem from [15]). In particular, we will use the shorthand $\mathbb{T}(a) = \mathbb{T}_{\varepsilon_0}(a)$.

Let N^0 be obtained from N by removing all cuspidal regions in $N_{(0,\varepsilon_0)}$. If N has finitely generated fundamental group, then N^0 contains a *relative compact core* $(M, P) \subset (N^0, \partial N^0)$, where M is a codimension zero submanifold such that the inclusion into N is a homotopy equivalence, each component X of ∂N^0 contains exactly one component Q of P , and the inclusion of Q into X is a homotopy equivalence (see Kulkarni-Shalen [32] or McCullough [37]). More generally, if $\varepsilon \in (0, \mu_3)$, we define N^ε to be the complement of the non-compact components of $N_{(0,\varepsilon)}$ and note that one may similarly define relative compact cores for N^ε .

Pared manifolds. A pared manifold is a pair (M, P) where M is a compact oriented irreducible 3-manifold and P is a closed subsurface of ∂M whose components are incompressible annuli and tori, such that

- (1) every abelian, non-cyclic subgroup of $\pi_1(M)$ is conjugate into a subgroup of $\pi_1(Q)$ where Q is a component of P , and
- (2) any π_1 -injective map $f : (S^1 \times [0, 1], S^1 \times \{0, 1\}) \rightarrow (M, P)$ of an annulus into M is homotopic, as a map of pairs, to a map with image in P .

Thurston's Geometrization Theorem, see Morgan [41, Theorem B'], implies that (M, P) is pared if and only if it is the relative compact core of a hyperbolic 3-manifold with finitely generated fundamental group.

The closure $\overline{\partial M - P}$ of the complement of P in ∂M is called the *pared boundary* and denoted $\partial_0(M, P)$. We say that (M, P) has *pared incompressible boundary* if each component of $\partial_0(M, P)$ is incompressible.

Bonahon's Tameness Theorem [10] implies that if N^0 has finitely generated fundamental group and its relative compact core (M, P) has pared incompressible boundary, then $N^0 - \text{int}(M)$ is homeomorphic to $\partial_0(M, P) \times [0, \infty)$. It follows that the relative compact core is well-defined up to isotopy in this case.

Spaces of representations. If G is a group let $D(G)$ denote the subset of $\text{Hom}(G, \text{PSL}_2(\mathbb{C}))$ consisting of discrete, faithful representations. For a 3-manifold M , let $D(M) = D(\pi_1(M))$, and for a pared manifold (M, P) , let $D(M, P)$ denote the subset of $D(M)$ consisting of representations such that $\pi_1(Q)$ maps to a parabolic subgroup for each component Q of P . Let $AH(M, P)$ be the set of conjugacy classes in $D(M, P)$. If P is empty or consists only of tori, then we use $AH(M)$ as shorthand for $AH(M, P)$. We give $\text{Hom}(G, \text{PSL}_2(\mathbb{C}))$ the compact-open topology, $D(M)$ and $D(M, P)$ inherit the subspace topology and we give $AH(M, P)$ the quotient topology.

If $\rho \in D(M, P)$, let $N_\rho = \mathbb{H}^3 / \rho(\pi_1(M))$ be the quotient manifold, and let (M_ρ, P_ρ) be a relative compact core for N_ρ^0 . Note that the conjugacy class $[\rho] \in AH(M)$ determines a homotopy class of homotopy equivalences $h_\rho : M \rightarrow M_\rho$ and associated maps of pairs $h_\rho : (M, P) \rightarrow (M_\rho, P_\rho)$. Let $AH_0(M, P)$ comprise those elements of $AH(M, P)$ for which the homotopy class of h_ρ contains an orientation-preserving pared homeomorphism.

Let $GF(M, P) \subset AH(M, P)$ denote the geometrically finite representations (i.e. those for which the intersection of the convex core $C(N_\rho)$ with N_ρ^0 is compact) whose *only* parabolic subgroups correspond to components of P . (In this case, $(C(N_\rho), C(N_\rho) \cap \partial N_\rho^0)$ is a relative compact core unless ρ is Fuchsian, in which case $C(N_\rho)$ is two-dimensional.) Let $GF_0(M, P) \subset GF(M, P)$ denote the geometrically finite elements which also lie in $AH_0(M, P)$. If P is empty, then $GF_0(M, P)$ agrees with the deformation space $CC_0(M)$ introduced in the introduction.

End invariants. If (M, P) has pared incompressible boundary, the Bers parameterization (see Bers [7, 8] or [20, Chapter 7]) gives rise to a homeomorphism

$$B : GF_0(M, P) \rightarrow \mathcal{T}(\partial_0(M, P)).$$

Here $\mathcal{T}(S)$, for a compact surface S , denotes the Teichmüller space of marked, finite area hyperbolic structures on the interior of S . A special case of this is the product case $M = S \times [0, 1]$ for a compact surface S , and $P = \partial S \times [0, 1]$. In this case $GF_0(S \times [0, 1], \partial S \times [0, 1]) \cong \mathcal{T}(S) \times \mathcal{T}(\bar{S})$ denotes the space of *quasifuchsian representations* of S , where \bar{S} is S with the opposite orientation. We use the notation

$QF(S)$ for $GF_0(S \times [0, 1], \partial S \times [0, 1])$, $D(S)$ for $D(S \times [0, 1], \partial S \times [0, 1])$, and $AH(S)$ for $AH(S \times [0, 1], \partial S \times [0, 1])$.

Thurston's *end invariants* generalize Bers' parameters. An end invariant for a point $[\rho]$ in $AH_0(M, P)$ is the following data: A multicurve p in $\partial_0(M, P)$, called the parabolic locus, and for each component X of $\partial_0(M, P) \setminus p$ that is not a 3-holed sphere, either a point of $\mathcal{T}(X)$, in which case X is called a geometrically finite component, or a filling geodesic lamination on X , in which case X is called a geometrically infinite component. The parabolic locus p is a collection of core curves of components of $h_\rho^{-1}(P_\rho)$. Each geometrically finite component corresponds to a component of the conformal boundary at infinity which bounds a component of the complement of the relative compact core of $(N_\rho)^0$, and each geometrically infinite component corresponds to a component which is "simply degenerate" in the sense of Thurston, in which case the lamination is the ending lamination of that component. See [40, 15] for a detailed discussion.

We write the end invariant of $\rho \in AH_0(M, P)$ as $v(\rho)$. If S is a component of $\partial_0(M, P)$ and α is a curve on S , then we define $\ell_{v(\rho)}(\gamma) \in [0, +\infty]$ to be 0 if γ is homotopic into the parabolic locus, to be the length of γ in the associated hyperbolic metric if γ is contained in a geometrically finite component of the parabolic locus, and to be $+\infty$ otherwise.

In the quasifuchsian space $QF(S)$ we write $v(\rho)$ as a pair $(v^+(\rho), v^-(\rho))$ associated to the two copies of the surface. The ordering of this pair is determined by keeping track of the orientation of S and of the quotient 3-manifold. In particular given an orientation-preserving embedding $S \times [0, 1]$ as a relative compact core and identifying $N_\rho \setminus \mathbb{T}(\partial S)$ with $S \times \mathbb{R}$, v^+ corresponds to the component $S \times (1, \infty)$, and v^- to the $S \times (-\infty, 0)$, which we call the *top* and *bottom ends*, respectively.

Let $\mathcal{E}(S)$ denote the set of all possible end invariants on a surface S . The Ending Lamination Theorem [15] tells us that the map

$$v_{(M,P)} : AH_0(M, P) \rightarrow \mathcal{E}(\partial_0(M, P))$$

is an injection. In the special case of a compact surface S we obtain

$$v_S : AH_0(S) \rightarrow \mathcal{E}(S) \times \mathcal{E}(\bar{S})$$

which has the form (v_S^+, v_S^-) . We note that, while the Bers parameterization is a holomorphic isomorphism, there is no natural topology on $\mathcal{E}(\partial_0(M, P))$ for which v is even continuous (see Brock [11]).

Skinning maps. If (M, P) has pared incompressible boundary, then for each component S of $\partial_0(M, P)$, we have a restriction map $D(M, P) \rightarrow D(S)$ which induces a map

$$r_{(M,P)}^S : AH_0(M, P) \rightarrow AH_0(S).$$

Thurston proved that every cover of a geometrically finite Kleinian group of infinite volume which has finitely generated fundamental group is itself geometrically finite (see Morgan [41, Prop. 7.1]). If P is empty, or if no non-peripheral curve in S is homotopic into P , then $r_{(M,P)}^S(GF_0(M, P)) \subset QF(S)$. In these cases, using the

Bers parametrization, we obtain a map

$$\hat{r}_{(M,P)}^S : \mathcal{T}(\partial_0(M,P)) \rightarrow \mathcal{T}(S) \times \mathcal{T}(\bar{S})$$

and, putting this together over all components we get

$$\hat{r}_{(M,P)} : \mathcal{T}(\partial_0(M,P)) \rightarrow \mathcal{T}(\partial_0(M,P)) \times \mathcal{T}(\overline{\partial_0(M,P)})$$

which has the form $\hat{r}_{(M,P)} = (Id, \sigma_{(M,P)})$, where

$$\sigma_{(M,P)} : \mathcal{T}(\partial_0(M,P)) \rightarrow \mathcal{T}(\overline{\partial_0(M,P)}).$$

is Thurston's skinning map.

To define a skinning map in general, for each component S of $\partial_0(M,P)$ we consider the map

$$\sigma_{(M,P)}^S = v_S^- \circ r_{(M,P)}^S : AH_0(M,P) \rightarrow \mathcal{E}(\bar{S}).$$

The skinning map is the product

$$\sigma_{(M,P)} = \prod_{S \subset \partial_0(M,P)} \sigma_{(M,P)}^S$$

where the product is taken over all components of $\partial_0(M,P)$. Notice that

$$\prod_{S \subset \partial_0(M,P)} v_S \circ r_{(M,P)}^S : AH_0(M,P) \rightarrow \mathcal{E}(\partial_0(M,P)) \times \mathcal{E}(\overline{\partial_0(M,P)})$$

has the form $(v_{(M,P)}, \sigma_{(M,P)})$.

Fricke spaces. If S is a compact, orientable surface with negative Euler characteristic, one may consider the *Fricke space* $\mathcal{F}(S)$ of all complete marked hyperbolic structures on the interior of S . Notice that the Teichmüller space $\mathcal{T}(S)$ of all finite area complete hyperbolic structures on S is naturally the boundary of the Fricke space. If S has genus g and n boundary components, then

$$\mathcal{F}(S) \cong \mathbb{R}^{6g-6+2n} \times [0, \infty)^n$$

and one may define Fenchel-Nielsen coordinates just as for $\mathcal{T}(S)$, the only difference being that one has an additional coordinate which records the lengths of the geodesics in the homotopy class of ∂S , see Bers-Gardiner [9] for details.

For an annulus A we declare $\mathcal{F}(A)$ to be $[0, \infty)$, where the parameter denotes the translation length of the core of A in a complete hyperbolic structure on the interior of A .

2.2. Characteristic submanifold theory

Jaco-Shalen [27] and Johannson [28] introduced the characteristic submanifold $\Sigma(M)$ of a compact, orientable, irreducible 3-manifold M with incompressible boundary, which is a canonical collection of Seifert-fibered and interval bundle submanifolds which contain all the incompressible annuli and tori in M . Johannson [28] further proved that every homotopy equivalence of M can be homotoped to preserve $\Sigma(M)$ and its complement $M - \Sigma(M)$. In this section, we recall the key properties of the characteristic submanifold in the setting of pared manifolds. (For

a more detailed discussion of the characteristic submanifold in the pared setting, see either Morgan [41, Sec. 11] or Canary-McCullough [20, Chap. 5].)

We must introduce some additional notation to handle the pared case. If B is a compact surface with boundary, we say that $f : B \rightarrow M$ is *admissible in (M, P)* if whenever C is a component of ∂B , then either $f(C) \subset \text{int}(P)$ or $f(C) \subset \partial_0(M, P)$. We say that two admissible maps (or admissible embeddings) $f : B \rightarrow M$ and $g : B \rightarrow M$ are *admissibly homotopic* (or *admissibly isotopic*) if there exists a homotopy (isotopy) $H : B \times [0, 1] \rightarrow M$ from f to g such that if C is a component of ∂B , then either $H(C \times [0, 1]) \subset P$ or $H(C \times [0, 1]) \subset \partial_0(M, P)$. We recall that if A is an annulus or Möbius band, then an admissible map $f : A \rightarrow M$ is an *essential immersed annulus or Möbius band in (M, P)* , if $f_* : \pi_1(A) \rightarrow \pi_1(M)$ is injective, and f is not admissibly homotopic to a map with image in ∂M . If f is an embedding, we will simply call f an *essential annulus or Möbius band in (M, P)* . In this case we will often abuse notation and also refer to the image $f(A)$ as an essential annulus or Möbius band in (M, P) . Notice that with this definition annuli which are parallel to components of P are inessential.

The following result encodes the basic properties of the characteristic submanifold of a pared manifold with pared incompressible boundary.

Theorem 2.1. (Jaco-Shalen [27], Johannson [28]) *Let (M, P) be a pared 3-manifold with incompressible pared boundary. There exists a codimension zero submanifold $\Sigma(M, P)$ satisfying the following properties:*

- (1) *Each component Σ_i of $\Sigma(M, P)$ is either*
 - (i) *an interval bundle over a compact surface F_i with negative Euler characteristic so that $\Sigma_i \cap \partial_0(M, P)$ is its associated ∂I -bundle,*
 - (ii) *a solid torus V , or*
 - (iii) *a thickened torus X such that ∂X contains a toroidal component of P .*
- (2) *The frontier $\text{Fr}(\Sigma(M, P))$ of $\Sigma(M, P)$ is a collection of essential annuli with boundary in $\partial M - P$.*
- (3) *If an annular component of P is homotopic into a solid torus component V , of $\Sigma(M, P)$, then it is contained in the interior of a component of $\partial V \cap \partial M$.*
- (4) *Any immersed essential annulus or Möbius band in (M, P) is admissibly homotopic into $\Sigma(M, P)$. If the annulus or Möbius band is embedded, then it is admissibly isotopic into $\Sigma(M, P)$.*
- (5) *If Y is a component of $M - \Sigma(M, P)$, then either $\pi_1(Y)$ is non-abelian or $(\bar{Y}, \text{Fr}(Y)) \cong (S^1 \times [0, 1] \times [0, 1], S^1 \times [0, 1] \times \{0, 1\})$ and Y lies between an interval bundle component of $\Sigma(M, P)$ and a thickened or solid torus component of $\Sigma(M, P)$.*

Moreover, the component of $\Sigma(M, P) \cup Y$ which contains Y is not a submanifold of type (1).

A submanifold with these properties is unique up to isotopy, and is called the characteristic submanifold of M .

Remark: For statements of Jaco, Shalen and Johannson's results in a form quite similar to the above statement, see Canary-McCullough [20, Thm. 2.9.1 and 5.3.4].




FIGURE 1. Schematic of a typical JSJ decomposition. A and C represent relatively acylindrical pieces. $\Sigma(M, P)$ consists of an I -bundle B and a solid torus U . The window surface has six components: w_1, w_2 in $\partial_0 B$, and four annuli w_3-w_6 .

In Johansson [28] (and in Canary-McCullough [20]), every toroidal component of P is contained in a component of the characteristic submanifold. In this paper, we use the convention of Jaco-Shalen [27], that a toroidal component T of P is only contained in a component of $\Sigma(M, P)$ if there is an essential annulus with one boundary component in T .

Thurston defines the *window* of (M, P) , denoted $\text{window}(M, P)$, to be the union of the interval bundle components of $\Sigma(M, P)$ with a regular neighborhood of every component of $\text{Fr}(\Sigma(M, P))$ which is not homotopic into an interval bundle component. The window is an interval bundle over a surface $\text{wb}(M, P)$, called the *window base*, which contains a copy of the base surface of each interval bundle component of $\Sigma(M, P)$ and an annulus for each component of $\text{Fr}(\Sigma(M, P))$ which is not homotopic into an interval bundle component. The *window surface* $\text{ws}(M, P) = \text{window}(M, P) \cap \partial_0(M, P)$ is a two-fold cover of the window base. We then define

$$\partial_{nw}(M, P) = \partial_0(M, P) - \text{ws}(M, P)$$

to be the non-window portion of the boundary. Each component of $\partial_{nw}(M, P)$ is either an annulus homotopic to a component of $\text{Fr}(\Sigma(M, P)) \cap \partial M$ or a subsurface with negative Euler characteristic contained in a component of $\partial M - \Sigma(M, P)$.

Robust systems of annuli. We will see, in Section 2.5, that one may remark a sequence in $AH_0(M, P)$ so that it converges on the complement of a collection of essential annuli, in a sense to be defined in Section 2.3. It will be convenient to always choose this collection of essential annuli so that each solid or thickened torus of (M, P) is either cut out cleanly in one piece or does not participate at all in the collection. This is part of a condition we call *robustness*, which we define and study below.

We say that a collection \mathcal{A} of essential annuli in $\Sigma(M, P)$ is *robust* if

- (1) the components of \mathcal{A} are disjoint and mutually non-parallel,

- (2) \mathcal{A} contains the frontier of any thickened torus component X of $\Sigma(M, P)$, and contains no other annuli homotopic into X , and
- (3) if \mathcal{A} has a component homotopic into a solid torus component V of $\Sigma(M, P)$, then \mathcal{A} contains the frontier of V and no other annuli homotopic into V . Moreover, if V contains a component of P , then the frontier of V is contained in \mathcal{A} .

Any component of \mathcal{A} which is not homotopic into a solid or thickened torus component of $\Sigma(M, P)$, must be isotopic into an interval bundle component Σ_i of $\Sigma(M, P)$ and hence to the interval bundle over a two-sided curve in F_i (see [20, Thm. 2.7.2] for example). Hence, one may always isotope \mathcal{A} so that it is the I -bundle associated to a two-sided multi-curve a in the window base.

A robust collection of essential annuli \mathcal{A} in (M, P) naturally decomposes (M, P) into a collection of solid and thickened tori and submanifolds which inherit the structure of a pared manifold. Let $\mathcal{N}(\mathcal{A})$ be a (pared) regular neighborhood of the union of the components of \mathcal{A} . A solid or thickened torus component of $M - \mathcal{N}(\mathcal{A})$ is either a regular neighborhood of a solid or thickened torus component of $\Sigma(M, P)$ or a regular neighborhood of an essential Möbius band in an interval bundle component of $\Sigma(M, P)$. Let $M_{\mathcal{A}}$ be obtained from $M - \mathcal{N}(\mathcal{A})$ by deleting all components which are solid or thickened tori and let

$$P_{\mathcal{A}} = (M_{\mathcal{A}} \cap P) \cup \text{Fr}(M_{\mathcal{A}}).$$

Lemma 2.2. *If (M, P) is a pared manifold with pared incompressible boundary and \mathcal{A} is robust collection of essential annuli in (M, P) , then each component of $(M_{\mathcal{A}}, P_{\mathcal{A}})$ is a pared manifold with pared incompressible boundary.*

Proof. Let (M', P') be a component of $(M_{\mathcal{A}}, P_{\mathcal{A}})$. Every component of P' is either a component of P or isotopic to a component of \mathcal{A} , hence is either an incompressible annulus or torus.

Any non-cyclic abelian subgroup H of $\pi_1(M')$ is also a non-cyclic abelian subgroup of $\pi_1(M)$, since $\pi_1(M')$ injects into $\pi_1(M)$. Since (M, P) is pared, H is conjugate in $\pi_1(M)$ to a subgroup of $\pi_1(T)$ where T is a toroidal component of P . Since the frontier of M' is a collection of essential annuli, T must be a toroidal component of $\partial M'$ and H is conjugate in $\pi_1(M')$ to a subgroup of $\pi_1(T')$.

We next claim that if P'_0 is a component of P' then $\pi_1(P'_0)$ is a maximal abelian subgroup of $\pi_1(M')$. If P'_0 is a component of P , then it is a maximal abelian subgroup of $\pi_1(M)$, see [20, Lemma 5.1.1], hence a maximal abelian subgroup of $\pi_1(M') \subset \pi_1(M)$. If $\pi_1(P'_0)$ is not a maximal abelian subgroup of $\pi_1(M)$, then it is a component of $\partial \mathcal{N}(A_0)$ where A_0 is either a component of the frontier of a solid or thickened torus or the boundary of a regular neighborhood of an essential Möbius band, see Johannson [28, Lemma 32.1]. In all these cases, P'_0 is a component of the frontier of a component X_0 of $M - M_{\mathcal{A}}$ so that $\pi_1(P'_0)$ is a proper subgroup of the abelian group $\pi_1(X_0)$. If $\pi_1(P'_0)$ is not a maximal abelian subgroup of $\pi_1(M')$, then there exists an element α of $\pi_1(M') \setminus \pi_1(P'_0)$ which centralizes $\pi_1(P'_0)$. If β is an element of $\pi_1(X_0) \setminus \pi_1(P'_0)$, then, by the Seifert-van Kampen Theorem, α and β both centralize $\pi_1(P'_0)$ but do not themselves commute. However, this is impossible in a torsion-free discrete subgroup of $\text{PSL}(2, \mathbb{C})$.

Now consider a π_1 -injective map $f : (A, \partial A) \rightarrow (M', P')$ of an annulus A into (M', P') . Notice that, by construction, no two distinct components of P' contain homotopic curves, so $f(\partial A)$ is contained in a single component B of P' . If B is a torus then, since (M, P) is pared, f can be admissibly homotoped in (M, P) to a map with image in B . Since the frontier of M' consists of essential annuli, one may alter the homotopy to lie entirely within (M', P') . (One may see this by lifting the homotopy to the cover of M associated to $\pi_1(M')$ and noticing that the cover deformation-retracts onto M' .) If B is an annulus, note that both components of ∂A map to homotopic curves in B : if not then the core of B would have two distinct but conjugate powers, but this is impossible in a torsion-free discrete subgroup of $PSL(2, \mathbb{C})$ since the stabilizer of the fixed points of an element must commute with it. A homotopy between these two curves, adjoined to f , gives a map \bar{f} of a torus into M' . Since, $\pi_1(B)$ is a maximal abelian subgroup of $\pi_1(M')$, the map \bar{f} , and hence f itself, lifts to the cover of M' associated to $\pi_1(B)$. Since this cover deformation retracts to the lift of B , the lift of f admissibly homotopes into the lift of B , which implies that f itself is admissibly homotopic into B . This completes the proof that (M', P') is a pared manifold.

Since every component of $\partial_0 M' - P'$ is an incompressible subsurface of a component of $\partial_0 M$ and $\partial_0 M$ is incompressible in M , (M', P') has pared incompressible boundary. \square

2.3. Partial Convergence and Geometric Limits

We now recall standard notation on algebraic and geometric convergence, as well as introducing terminology for sequences of representations which may not converge, but converge (in various senses) on subgroups. We take care, in Lemmas 2.3 and 2.4, to keep track of basepoints, keeping in mind the situation where there may be more than one geometric limit of interest.

Partial convergence and basepoints. Fix a group H . If $\rho \in D(H)$ and $g \in PSL_2(\mathbb{C})$, we let ρ^g be the conjugate representation $h \mapsto g\rho(h)g^{-1}$. For $h \in H$ we let $I_h : H \rightarrow H$ be conjugation by h and note that $\rho \circ I_h = \rho^{\rho(h)}$.

We call a group *nonelementary* if it is not virtually abelian. For a sequence $\{\rho_n\}$ in $D(H)$ and a nonelementary subgroup $J < H$, we need two distinct notions of partial convergence of the sequence on J .

- We say $\{\rho_n\}$ *converges up to conjugacy on J* if there exists a sequence $\{g_n\}$ in $PSL_2(\mathbb{C})$ such that $\{\rho_n^{g_n}|_J\}$ converges in $D(J)$.
- We say $\{\rho_n\}$ *converges up to inner automorphism on J* if there exists a sequence $\{h_n\}$ in H such that $\{\rho_n^{\rho_n(h_n)}|_J\}$ (equivalently $\{\rho_n \circ I_{h_n}|_J\}$) converges in $D(J)$.

Note that both of these definitions depend only on ρ_n and the conjugacy class of J in H . The first definition also depends only on the conjugacy class of ρ_n , i.e. on $[\rho_n] \in AH(H)$.

Note that, if both $\{\rho_n|_J\}$ and $\{\rho_n^{g_n}|_J\}$ converge, then g_n must remain in a compact subset of $PSL_2(\mathbb{C})$. This is a consequence of the assumption that J is nonelementary, and the fact that the action of $PSL_2(\mathbb{C})$ by conjugation on its nonelementary

discrete subgroups is proper. It follows from this that, if $\{\rho_n\}$ converges up to conjugacy on J , then the limit representation is determined up to conjugacy.

A representation $\rho \in D(H)$ determines a basepoint b_ρ in $N_\rho = \mathbb{H}^3/\rho(H)$, namely the image of the orbit $\rho(H)\mathbf{0}$, where $\mathbf{0}$ is a fixed basepoint for \mathbb{H}^3 . We note that this basepoint is invariant under inner automorphisms since ρ and $\rho \circ I_h$ have the same image for any $h \in H$. We can now relate behavior of basepoints to these notions of partial convergence.

Lemma 2.3. *Suppose that a sequence $\{\rho_n\}$ in $D(H)$ converges up to inner automorphism on a nonelementary subgroup J of H . For a sequence $\{g_n\}$ in $PSL_2(\mathbb{C})$,*

- (1) *if the sequence $\{d(b_{\rho_n}, b_{\rho_n^{g_n}})\}$ is bounded, then a subsequence of $\{\rho_n^{g_n}\}$ also converges up to inner automorphism on J .*
- (2) *if $\{\rho_n^{g_n}\}$ converges up to inner automorphism on J , then $\{d(b_{\rho_n}, b_{\rho_n^{g_n}})\}$ is bounded.*

Proof. To compare the basepoints of two conjugate representations ρ and ρ^g , consider the cover $\mathbb{H}^3 \rightarrow N_\rho = N_{\rho^g}$ induced by the action of $\rho(H)$. The lift of b_ρ to this cover is the orbit $\rho(H)\mathbf{0}$, whereas the lift of b_{ρ^g} to this cover is the orbit $\rho(H)g^{-1}\mathbf{0}$ (since it is the preimage by the conjugating map g of the orbit $\rho^g(H)\mathbf{0}$). Thus, $d_{N_\rho}(b_\rho, b_{\rho^g})$ is equal to the distance from $\mathbf{0}$ to $\rho(H)g^{-1}\mathbf{0}$.

Now, after suitable conjugation we may assume that $\{\rho_n\}$ already converges on J . If the distances $d(b_{\rho_n}, b_{\rho_n^{g_n}})$ are bounded, then $d_{\mathbb{H}^3}(\mathbf{0}, \rho_n(H)g_n^{-1}\mathbf{0})$ is bounded, and so there is a sequence $\{h_n\}$ in H such that $\{\rho_n(h_n)g_n^{-1}\}$ remains in a compact subset of $PSL_2(\mathbb{C})$. Hence a subsequence of $\{\rho_n^{g_n \rho_n(h_n^{-1})}|_J\}$ converges.

For the second part, if $\{\rho_n^{g_n \rho_n(k_n)}|_J\}$ converges for some sequence $k_n \in H$, then $\{g_n \rho_n(k_n)\}$ lies in a compact set as we argued above, and hence the distances $\{d(b_{\rho_n}, b_{\rho_n^{g_n}})\}$ are bounded. \square

These observations allow us to analyze the situation where one sequence $\{\rho_n\}$ converges up to inner automorphism on several subgroups, in terms of distances between their basepoints. Given $\{\rho_n \in D(H)\}$ and a collection \mathcal{J} of nonelementary subgroups of H , if $\{\rho_n\}$ converges up to conjugacy on each $J \in \mathcal{J}$, fix a conjugating sequence $\{g_n^J\}$ for which $\{\rho_n^{g_n^J}|_J\}$ converges, and let b_n^J denote the sequence of associated basepoints. Lemma 2.3 implies that the property that $d(b_n^J, b_n^L)$ is bounded for $J, L \in \mathcal{J}$ is independent of the choice of conjugating sequences. We now record the following consequence:

Lemma 2.4. *Let \mathcal{J} be a collection of nonelementary subgroups of H and suppose that a sequence $\{\rho_n\}$ in $D(H)$ converges up to conjugacy on each $J \in \mathcal{J}$. With choice of basepoints as above, if $\{d(b_n^J, b_n^L)\}$ is bounded for each $J, L \in \mathcal{J}$, then there is a single conjugating sequence $\{g_n\}$ such that, after possibly passing to a subsequence, $\{\rho_n^{g_n}\}$ converges up to inner automorphism on each $J \in \mathcal{J}$.*

If b_n is the basepoint associated to $\rho_n^{g_n}$, then the sequence $\{d(b_n, b_n^J)\}$ is bounded for each $J \in \mathcal{J}$.

Proof. Fixing one member $J_0 \in \mathcal{J}$ let $g_n = g_n^{J_0}$ so that $\{\rho_n^{g_n}|_{J_0}\}$ converges. For any other $J \in \mathcal{J}$, we have a bound on $d(b_n^{J_0}, b_n^J)$, so by Lemma 2.3 we find that (a

subsequence of) $\{\rho_n^{g_n}\}$ also converges up to inner automorphism on J . The final statement follows since $b_n = b_n^{J_0}$. \square

Geometric convergence. We say that a sequence $\{\Gamma_n\}$ of torsion-free discrete subgroups of $\mathrm{PSL}_2(\mathbb{C})$ *converges geometrically* to a discrete group Γ , if it converges as a sequence of closed subsets in the sense of Gromov-Hausdorff convergence. It follows that the sequence $\{\mathbb{H}^3/\Gamma_n\}$ of quotient manifolds converges geometrically to $N = \mathbb{H}^3/\Gamma$ (see Canary-Epstein-Green [19] for an extensive discussion). The following lemma recalls standard properties of geometric convergence which will be used throughout the paper.

Lemma 2.5. *Suppose that a sequence $\{\rho_n\}$ in $D(H)$ converges on a nonelementary subgroup J of H to some $\rho \in D(J)$. Then after possibly restricting to a subsequence, $\{\rho_n(H)\}$ converges geometrically to a discrete nonelementary group $\hat{\Gamma}$. Let $\hat{N} = \mathbb{H}^3/\hat{\Gamma}$ and let $\pi : N_\rho \rightarrow \hat{N}$ be the natural covering map. Given $\varepsilon \leq \varepsilon_0$ there exists a nested sequence $\{Z_n\}$ of compact submanifolds exhausting \hat{N} and K_n -bilipschitz smooth embeddings $\psi_n : Z_n \rightarrow N_{\rho_n}$ such that:*

- (1) $K_n \rightarrow 1$.
- (2) $\psi_n \circ \pi$ carries the basepoint b_ρ of N_ρ to the basepoint b_{ρ_n} of N_{ρ_n} .
- (3) If Q is a compact subset of a component of $\partial\hat{N}_{(0,\varepsilon)}$, then, for all large enough n , $\psi_n(Q)$ is contained in $\partial(N_{\rho_n})_{(0,\varepsilon)}$ and $\psi_n(Z_n \cap (\hat{N} - \hat{N}_{(0,\varepsilon)}))$ does not intersect $(N_{\rho_n})_{(0,\varepsilon)}$.
- (4) If X is a finite complex and $h : X \rightarrow N_\rho$ is continuous, then, for all large enough n , $(\psi_n \circ \pi \circ h)_*$ is conjugate to $\rho_n \circ \rho^{-1} \circ h_*$.

Remark: The existence of the exhausting sequence of submanifolds with properties (1) and (2) is discussed in [19, Cor. I.3.2.11]. Property (3) is established in [15, Lemma 2.8]. Property (4) is an immediate consequence of the proof of [21, Prop. 3.3] which it generalizes.

2.4. Hierarchies, model manifolds, and topological order

We briefly review here some notation and definitions that arise in [36, 40, 15]. See also [14] for a summary. This machinery will play a central role in Section 3.

Markings, subsurface projections and hierarchies. A *marking* μ on a surface S , as in [36, Section 2.5], is a multicurve, denoted $\mathrm{base}(\mu)$, together with at most one transversal curve for each component β of $\mathrm{base}(\mu)$ (i.e. a curve intersecting β minimally which is disjoint from $\mu - \beta$). A *generalized marking* μ , as in [40, Section 5.1], is a geodesic lamination $\mathrm{base}(\mu)$ which supports a measure, together with at most one transversal curve for each closed curve component of $\mathrm{base}(\mu)$.

To an end invariant ν on a surface S , as described in Section 2.1, we associate in [40, Section 7.1] a generalized marking μ as follows: $\mathrm{base}(\mu)$ consists of the parabolic locus of ν , the ending laminations of geometrically infinite components of the complement of the parabolic locus and a minimal length pants decomposition of the hyperbolic structure on each geometrically finite component of ν . To each

curve in $\text{base}(\mu)$ which is non-peripheral in a geometrically finite component we choose a minimal length transversal curve.

Recall the *curve complex* $\mathcal{C}(Y)$ of a surface, whose vertices are homotopy classes of essential simple curves (except when Y is an annulus and a special definition is needed), see [35]. If Y is an essential subsurface of S and $\alpha \in \mathcal{C}(S)$ essentially intersects Y , one defines $\pi_Y(\alpha) \in \mathcal{C}(Y)$ by taking a component of $\alpha \cap Y$ and combining with arcs in ∂Y to obtain a (coarsely well-defined) element of $\mathcal{C}(Y)$ (again a special definition is needed when Y is an annulus). The curve complex and *subsurface projections* $\pi_Y : \mathcal{C}(S) \rightarrow \mathcal{C}(Y)$ are studied in [35, 36] and elsewhere. We also note that $\pi_Y(\mu)$ is well defined when μ is a generalized marking, and by extension so is $\pi_Y(\mathbf{v})$ for an end invariant \mathbf{v} .

A *hierarchy*, as in [36] and [40], is a collection of *tight geodesics* in curve complexes of subsurfaces of S , with certain interlocking properties. A tight geodesic in $\mathcal{C}(W)$ is a sequence of simplices (w_n) such that whenever $v_j \in w_j$ are vertices we have $d(v_i, v_j) = |i - j|$ for $i \neq j$, and such that w_i is the boundary of the subsurface filled by $w_{i-1} \cup w_{i+1}$.

To a pair (μ^+, μ^-) of generalized markings (which share no infinite leaves), one can associate a hierarchy $H(\mu^+, \mu^-)$, see [40, Section 5.5]. Given a subsurface $W \subseteq S$, we define

$$d_W(\mu^+, \mu^-) \equiv d_{\mathcal{C}(W)}(\pi_W(\mu^+), \pi_W(\mu^-)).$$

If $d_W(\mu^+, \mu^-)$ is sufficiently large (where the threshold depends only on S) then $\mathcal{C}(W)$ supports a (unique) geodesic $h_W \in H$, whose initial vertex i_W (respectively terminal vertex t_W) is uniformly close in $\mathcal{C}(W)$ to $\pi_W(\mu^-)$ (resp. $\pi_W(\mu^+)$), see [36, Lemma 6.2] and [40, Lemma 5.9]. Let $\mathcal{CH} \subset \mathcal{C}(S)$ denote the set of all curves appearing in geodesics in H .

Model manifolds. In [40, Section 8] we associate to the hierarchy $H = H(\mu^+, \mu^-)$ a *model manifold* $M(\mu^+, \mu^-)$ which is a copy of $\text{int}(S) \times \mathbb{R}$ endowed with a certain metric and a subdivision into blocks and tubes. We will mostly be concerned with the tubes. There is a tube $U(\gamma)$ for each curve γ in \mathcal{CH} and each component of ∂S . Each tube $U(\gamma)$ has the form $A \times J$ where A is an open regular neighborhood of γ in S and J is an interval, where $J = \mathbb{R}$ if γ is a component of ∂S . In particular, the complement $\hat{M}(\mu^+, \mu^-)$ in $M(\mu^+, \mu^-)$ of the tubes associated to ∂S is homeomorphic to $S \times \mathbb{R}$.

The following result is the main outcome of [40] and [15].

Bilipschitz Model Manifold Theorem: ([40, 15]) *If S is a compact surface, there exists $L_h(S) > 1$ and $\varepsilon_h(S) > 0$ such that if $\rho \in D(S)$ has end invariants $(\mathbf{v}^+, \mathbf{v}^-)$ and associated generalized markings (μ^+, μ^-) , then there exists a $L_h(S)$ -bilipschitz map*

$$f_\rho : N_\rho \rightarrow M(\mu^+, \mu^-)$$

such that if $\alpha \in \mathcal{C}(S)$ and $\ell_\rho(\alpha) < \varepsilon_h(S)$, then $\alpha \in \mathcal{CH}$ and $f_\rho(U(\alpha)) = \mathbb{T}(\alpha)$.

By construction the core curves of the tubes of M have length at most 1, so L_h also bounds the length of their images in N_ρ , i.e. if $\alpha \in \mathcal{CH}$, then $\ell_\rho(\alpha) \leq L_h(S)$.

Ordering. If A and B are curves and/or subsurfaces of S and $f : A \rightarrow S \times \mathbb{R}$ and $g : B \rightarrow S \times \mathbb{R}$ are maps which are homotopic to the inclusions of A and B into $S \times \{0\}$, then we say that f *lies above* g if f is homotopic to $+\infty$ in the complement of $g(B)$ (i.e. there is a proper map $F : A \times [0, \infty) \rightarrow S \times [r, \infty)$, for some r , such that $f = F(\cdot, 0)$, whose image is disjoint from $g(B)$). We similarly say that f *lies below* g if f is homotopic to $-\infty$ in the complement of $g(B)$. (This topological ordering is discussed extensively in [15, Section 3].)

We will often identify a curve α in $S \times \mathbb{R}$ with a function whose image is α . Notice that a curve α may lie above a curve β without β lying below α . We say that a curve α is *unknotted* if it is isotopic to $\hat{\alpha} \times \{0\}$ for some curve $\hat{\alpha}$ on S .

A *level surface* in $S \times \mathbb{R}$ is an embedding $f : (S, \partial S) \rightarrow (S \times \mathbb{R}, \partial S \times \mathbb{R})$ which is properly isotopic to the identification of S with $S \times \{0\}$. Again, we often identify a level surface with its image in $S \times \mathbb{R}$.

The following result records a special property of homotopy equivalences which will be used in Section 3.

Lemma 2.6. ([14, Lemma 2.7]) *Let α and β be simple closed curves on S that intersect essentially. Let $f : (S, \partial S) \rightarrow (S \times \mathbb{R}, \partial S \times \mathbb{R})$ be a homotopy equivalence with image disjoint from $\beta \times \{0\}$. Then f lies above (below) $\beta \times \{0\}$ if and only if $f|_{\alpha}$ lies above (below) $\beta \times \{0\}$.*

We may apply this ordering to tubes in $\hat{M}(\mu^+, \mu^-)$. If $v, w \in \mathcal{CH}$, we say that $U(v)$ lies above (below) $U(w)$ if a core curve of $U(v)$ lies above (below) a core curve of $U(w)$ in $\hat{M}(\mu^+, \mu^-)$. By construction, if v and w both lie in simplices of a geodesic h_w in $H(\mu^+, \mu^-)$, then $U(v)$ lies above (below) $U(w)$ if and only if the simplex containing v occurs after (before) the simplex containing w on h_w .

2.5. Thurston's Relative Compactness Theorem and its consequences

The goal of this section is Corollary 2.9, which shows that any sequence in $AH(M, P)$ may be remarked by homeomorphisms supported on $\text{window}(M, P)$, so that it converges, up to subsequence, on $M_{\mathcal{A}}$ for some robust collection \mathcal{A} of essential annuli in (M, P) . The core curves of the components of \mathcal{A} have length 0 in the limit representations, and \mathcal{A} can be chosen to be maximal in the sense that the core curve of any essential annulus in a component of $(M_{\mathcal{A}}, P_{\mathcal{A}})$ has positive length in the associated limit representation.

Our major tool is Thurston's Relative Compactness Theorem [47] (see also Morgan [41, Cor. 11.5]):

Theorem 2.7. (Thurston [47]) *Suppose that (M, P) is a pared manifold with pared incompressible boundary and R is a component of $M \setminus \text{window}(M, P)$. The restriction mapping from $AH(M, P)$ to $AH(R)$ has bounded image.*

If X is a submanifold of M whose frontier is a collection of essential annuli in (M, P) , then each component of X determines a conjugacy class of subgroups of $\pi_1(M)$. We say that a sequence $\{\rho_n\} \subset D(M, P)$ *converges up to conjugacy on X* if $\{\rho_n\}$ converges up to conjugacy on $\pi_1(R)$ (in the sense of Section 2.3) for each

component R of X . In this language, Theorem 2.7 implies that any sequence in $D(M, P)$ has a subsequence which converges up to conjugacy on $M \setminus \text{window}(M, P)$.

Canary, Minsky and Taylor [22, Thm. 5.5] observed that, if one allows oneself to remark by pared homeomorphisms supported on $\Sigma(M, P)$, then one can find a subsequence which converges up to conjugacy on the complement of a robust collection of essential annuli.

We will adopt the following convenient notational convention throughout the paper. If $\rho \in AH(M, P)$ and \mathcal{B} is a collection of essential annuli in (M, P) , then $\ell_\rho(\mathcal{B})$ denotes the sum of the lengths of the geodesic representatives in N_ρ of the core curves of annuli in \mathcal{B} .

Theorem 2.8. (Canary-Minsky-Taylor [22]) *Let (M, P) be a pared 3-manifold with pared incompressible boundary. If $\{\rho_n\}$ is a sequence in $D(M, P)$, then, after passing to a subsequence, there is a robust collection \mathcal{B} of essential annuli in (M, P) and a sequence of pared homeomorphisms $\{\phi_n : (M, P) \rightarrow (M, P)\}$, each supported on $\text{window}(M, P)$, such that*

- (1) $\lim \ell_{\rho_n \circ (\phi_n)_*}(\mathcal{B}) = 0$, and
- (2) $\{\rho_n \circ (\phi_n)_*\}$ converges up to conjugacy on $M_{\mathcal{B}}$.

We provide a sketch of the proof of Theorem 2.8, since our statement is slightly different and more general than the one given in [22], although their proof goes through directly to yield our statement.

Sketch of proof of Theorem 2.8: We first apply Theorem 2.7 to pass to a subsequence so that $\{\rho_n\}$ converges up to conjugacy on $M - \text{window}(M, P)$. We then construct \mathcal{B} and $\{\phi_n\}$ piece by piece. We first include the frontier of any thickened torus component of $\Sigma(M, P)$ in \mathcal{B} . If V is a solid torus component of $\Sigma(M, P)$ with core curve v , then we include $\text{Fr}(V)$ in \mathcal{B} if and only if $\lim \ell_{\rho_n}(v) = 0$.

Let F be the collection of components of $\text{wb}(M, P)$ which are base surfaces of interval bundle components of $\Sigma(M, P)$. For each n , let F_n be obtained by removing any boundary component f so that $\ell_{\rho_n}(f) = 0$. Consider a pleated surface $f_n : (F_n, \tau_n) \rightarrow N_{\rho_n}$ in the homotopy class of the inclusion map, where τ_n is a finite area hyperbolic metric on F_n . (A pleated surface is a 1-Lipschitz map which is totally geodesic on a geodesic lamination which includes the boundary, called the pleating locus, and totally geodesic on the complement of the pleating locus.) One may pass to a subsequence so that there exists a two-sided multicurve b_I on F and a sequence $\{\hat{\psi}_n\}$ of homeomorphisms of $\text{int}(F)$, which extend to the identity on ∂F , so that $\lim \ell_{\hat{\psi}_n^*(\tau_n)}(b_I) = 0$ and if y is a curve on F which is disjoint from b_I , then $\{\ell_{\hat{\psi}_n^*(\tau_n)}(y)\}$ is bounded. (See [22, Prop. 5.6] and its use in the proof of [22, Thm. 5.5] for more details.) We complete the construction of \mathcal{B} by adding the interval bundle over b_I to \mathcal{B} .

We may extend each $\hat{\psi}_n$ to a homeomorphism of $\text{window}(M, P)$ which is the identity on $\text{Fr}(\text{window}(M, P))$ and hence to a pared homeomorphism ψ_n of (M, P) which is supported on $\text{window}(M, P)$. One then checks that, up to subsequence, there exists a sequence D_n of Dehn multitwists in $\text{Fr}(\text{window}(M, P))$ so that if $\phi_n = D_n \circ \psi_n^{-1}$, then $\{\rho_n \circ (\phi_n)_*\}$ converges up to conjugacy on $M_{\mathcal{B}}$, which verifies

property (2) (see [22, Lem. 5.7] and its use in the proof of [22, Thm. 5.5] for more details). Properties (1) and (3) hold by construction. \square

It will be useful to be able to choose the robust collection of annuli \mathcal{B} in Theorem 2.8 to be maximal in the sense that the length of any essential annulus in $(M_{\mathcal{B}}, P_{\mathcal{B}})$ has positive length in its associated limit representation.

Corollary 2.9. *Let (M, P) be a pared 3-manifold with pared incompressible boundary. If $\{\rho_n\}$ is a sequence in $D(M, P)$, then, after passing to a subsequence, there is a robust collection \mathcal{A} of essential annuli in (M, P) and a sequence of pared homeomorphisms $\{\phi_n : (M, P) \rightarrow (M, P)\}$, each supported on $\text{window}(M, P)$, such that*

- (1) $\lim \ell_{\rho_n \circ (\phi_n)_*}(\mathcal{A}) = 0$
- (2) $\{\rho_n \circ (\phi_n)_*\}$ converges up to conjugacy on $M_{\mathcal{A}}$, and
- (3) If B is an essential annulus in a component of $(M_{\mathcal{A}}, P_{\mathcal{A}})$, then $\lim \ell_{\rho_n \circ (\phi_n)_*}(B) > 0$.

Proof. Theorem 2.8 guarantees that there exists a subsequence, still called $\{\rho_n\}$, a robust collection \mathcal{B} of essential annuli with base multicurve b and sequence of pared homeomorphisms $\{\phi_n : (M, P) \rightarrow (M, P)\}$ such that $\lim \ell_{\rho_n \circ (\phi_n)_*}(b) = 0$ and $\{\rho_n \circ (\phi_n)_*\}$ converges up to conjugacy on $M_{\mathcal{B}}$. We may further assume that if V is a solid torus component of $\Sigma(M, P)$ with core curve v , then $\text{Fr}(V) \subset \mathcal{B}$ if and only if $\lim \ell_{\rho_n}(v) = 0$

Let \mathcal{C} be a maximal collection of disjoint, non-parallel essential annuli in $(M_{\mathcal{B}}, P_{\mathcal{B}})$ such that $\lim \ell_{\rho_n \circ (\phi_n)_*}(\mathcal{C}) = 0$. (Alternatively, we may choose c to be a maximal two-sided multicurve in $\text{wb}(M, P) - b$ such that $\lim \ell_{\rho_n \circ (\phi_n)_*}(c) = 0$ and let \mathcal{C} be the interval bundle over c .) Let \mathcal{A} be the union of \mathcal{B} and \mathcal{C} . Properties (1) and (2) hold for \mathcal{A} , since they held for the subcollection \mathcal{B} of \mathcal{A} . Property (3) holds by construction. \square

3. Length bounds

Let S be a compact surface and $[\rho] \in AH(S)$ a surface group. The goal of this section is a criterion for bounding the length at infinity of a curve in N_ρ , in terms of the situation of a representative of the curve in the interior. That is, given a simple curve $\alpha \subset S$ with a representative $\hat{\alpha}$ in N_ρ of given length, we wish to know when we can bound $\ell_{v^-(\rho)}(\alpha)$. The key is to study the collection of bounded length geodesics in N_ρ that lie “between” $\hat{\alpha}$ and the bottom end of N_ρ , in the following precise sense.

Let $\mathcal{C}(\hat{\alpha}, L)$ denote the set of curves on S which intersect α essentially and whose geodesic representatives in N_ρ have length at most L and do not lie above $\hat{\alpha}$ (in the sense of Section 2.4). Parabolic curves have no geodesic representatives, but we adopt the convention that β lies above every curve in N_ρ if it is in the parabolic locus of v^+ , and below if it is in the parabolic locus of v^- .

We will see that given an upper bound on the length of α in F , an upper bound on the size of $\mathcal{C}(\hat{\alpha}, L)$ and a lower bound on the length of any curve in $\mathcal{C}(\hat{\alpha}, L)$, one obtains an upper bound on the length of α in the bottom ending invariant. In

the statement, $L_h(S)$ is the constant from the Bilipschitz Model Manifold Theorem in Section 2.4.

Theorem 3.1. *Let S be a compact surface. Given $R, L_0, \varepsilon \in (0, \mu_3)$ and $L \geq L_h(S)$, there exists $L_1 > 0$ such that, if $\rho \in AH(S)$, α is a simple closed curve on S , and $\hat{\alpha}$ is a representative of α in N_ρ such that*

- (1) $\hat{\alpha}$ has length at most L_0 ,
- (2) $\mathcal{C}(\hat{\alpha}, L)$ contains at most R elements,
- (3) $\mathcal{C}(\hat{\alpha}, L)$ contains no curves of length less than ε , and
- (4) $\hat{\alpha}$ lies on a level surface F which does not intersect $(N_\rho)_{(0, \varepsilon)}$,

then

$$l_{v^-(\rho)}(\alpha) < L_1.$$

The idea of the proof is the following: If α has large, or infinite, length in $v^-(\rho)$, then there must be a subsurface W in S so that the projection distance $d_W(\alpha, v^-(\rho))$ is large or infinite. This forces the hierarchy $H(\mu^+(\rho), \mu^-(\rho))$ to have a long geodesic h_W associated to W , which corresponds to a large region in the model manifold of N_ρ containing a large number of bounded-length curves (corresponding to hierarchy curves) isotopic into W . In the model manifold, the topological placement of these curves corresponds to their location along the hierarchy geodesic. Thus, using what we know about the projection of α to h_W , we conclude that a large number of the geodesic representatives of these curves must not lie above $\hat{\alpha}$ in N_ρ – hence must belong to $\mathcal{C}(\alpha, L)$. The hypotheses bound the number of such curves and therefore the size of $d_W(\alpha, v^-(\rho))$. The main technical difficulties in the proof involve converting the relatively nice picture in the model manifold to the slightly messier arrangement of true geodesic representatives in N_ρ .

Proof. We first observe that α lies in the thick part of a Riemann surface component of $v^-(\rho)$.

Lemma 3.2. *With the same assumptions as Theorem 3.1, there exists $\varepsilon' < \varepsilon$ such that the curve α is homotopic into the ε' -thick part Z' of a Riemann surface component Z of $v^-(\rho)$.*

Proof. If a component p of the parabolic locus of $v^-(\rho)$ crosses α essentially, then the cusp associated to p is downward-pointing and hence p lies below α^* . However, this contradicts the fact that $\mathcal{C}(\alpha, L)$ contains no curves of length less than ε . If α essentially intersects a subsurface W which supports an ending lamination in $v^-(\rho)$, then there exists a sequence $\{\beta_n\}$ of hierarchy curves so that $\{\beta_n^*\}$ exits the downward-pointing simply degenerate end with base surface W . Therefore, $\ell_\rho(\beta_n) \leq L_h(S)$ for all n and β_n^* lies below $\hat{\alpha}$ for all large enough n . However, this contradicts the fact that $\mathcal{C}(\hat{\alpha}, L)$ contains only finitely many curves. Therefore, $v^-(\rho)$ has a Riemann-surface component supported on a subsurface Z containing α .

Choose $\delta = \delta(L_0) > 0$ so that, if a homotopically non-trivial curve γ in a hyperbolic 3-manifold is of length at most L_0 and intersects $\mathbb{T}_\delta(\beta)$ for some primitive curve β , then γ is homotopic to a power of β .

Suppose that β is a curve such that $\ell_Z(\beta) < 2\varepsilon'$ where $\varepsilon' = \min(\delta(L_0), \varepsilon)$. Theorem 3.1 in Epstein-Marden-Markovic [26] implies that β has a representative $\hat{\beta}$ on the bottom boundary of the convex core which has length at most $2\varepsilon' \leq \varepsilon$. So, $\hat{\beta}$ lies inside $\mathbb{T}_{\varepsilon'}(\beta)$ and one may then connect β^* to $\hat{\beta}$ by an annulus A_0 in the convex core which is entirely contained within $\mathbb{T}_{\varepsilon'}(\beta)$. We construct a homotopy A from β^* to $-\infty$, by concatenating A_0 with an annulus in $r^{-1}(\hat{\beta})$ where r is the nearest point retraction of N_ρ onto its convex core. Since $r(\hat{\alpha})$ has length at most L_0 , it must be disjoint from $\mathbb{T}_{\varepsilon'}(\beta)$ and hence from A . Therefore, β^* lies below $r(\hat{\alpha})$. Since $r(A) = A_0$, A is also disjoint from $\hat{\alpha}$, so β^* also lies below $\hat{\alpha}$. However, this contradicts our assumptions, so α must be homotopic into the ε' -thick part of Z . \square

Let $\mu^-(\rho)$ be the generalized marking of S which we associate with $v^-(\rho)$, and let μ' be the restriction of $\mu^-(\rho)$ to Z' . Notice that, assuming ε' is small enough, $\partial Z'$ is contained in $\mu^-(\rho)$, so $\text{base}(\mu')$ consists of the essential curves in Z' which are contained in $\text{base}(\mu^-(\rho))$ and the transversals in μ' are simply the transversals in $\mu^-(\rho)$ to the curves in $\text{base}(\mu')$. Since Z' is ε' -thick, both the base curves and the transversals in μ' have uniformly bounded length. Therefore, $l_Z(\alpha)$ can be bounded in terms of the complexity of its intersection with μ' , and this in turn is controlled by the subsurface projections of α and μ' , as in Theorem 6.12 of [36], or Lemma 2.3 [12]. In particular in order to bound $l_Z(\alpha)$ it suffices to bound

$$\sup_W d_W(\alpha, \mu') = \sup_W d_W(\alpha, v^-(\rho)) \quad (3.1)$$

where W varies over all subsurfaces of Z' intersecting α .

Theorem B of [39] tells us that when the diameter in $\mathcal{C}(W)$ of the projections of curves in S with a given length bound is sufficiently large, ∂W must be short. In particular there exists A_0 , depending only on ε' and on the length bounds on α and μ' , such that, if

$$d_W(\alpha, \mu') > A_0$$

then $\ell_\rho(\partial W) < \varepsilon'$. Since if $d_W(\alpha, \mu') \leq A_0$ we are done, we may assume from now on that $\ell_\rho(\partial W) < \varepsilon' \leq \varepsilon$.

Suppose now that α essentially intersects a component β of ∂W (in particular this must be the case if W is an annulus). By assumption, since $\ell_\rho(\beta) < \varepsilon$, β^* lies above $\hat{\alpha}$. Since $\hat{\alpha}$ lies on a level surface F disjoint from β^* , it follows that $\hat{\alpha}$ lies below β^* . Notice that one may homotope $\hat{\alpha}$ to α^* through curves of length at most L_0 . None of the curves in the homotopy can intersect $\mathbb{T}_{\varepsilon'}(\beta)$, since we can choose $\varepsilon' < \delta(L_0)$ (see the proof of Lemma 3.2), so α^* lies below β^* as well.

Now Theorem 1.3 of [14] tells us that a geodesic of bounded length that lies below β^* (and overlaps it on S) must have projection in W close to the bottom end invariant. That is, there is a constant A_1 (depending on L_0 and S) such that, if α^* lies below β^* then

$$d_W(\alpha, v^-) \leq A_1.$$

Thus, we are done in this case.

The only remaining possibility is that α is a non-peripheral curve contained in W . This case follows from the following lemma:

Lemma 3.3. *Let S be a compact surface. Given $L_0, L \geq L_h(S)$, and K , there exists $D > 0$ such that if α is a curve on S and $\hat{\alpha}$ is a representative of α in N_ρ with length at most L_0 , α is non-peripheral in an essential subsurface W of S , and $d_W(\alpha, \nu^-(\rho)) > D$, then $\mathcal{C}(\hat{\alpha}, L)$ contains at least K curves.*

Proof. Let $g : X \rightarrow N_\rho$ be a 1-Lipschitz surface that realizes α and ∂W , i.e. X is a finite area hyperbolic surface homeomorphic to S , g is a 1-Lipschitz map in the homotopy class of ρ and g takes the geodesic representatives of α and ∂W in X to their geodesic representatives in N_ρ by an isometry. (For example, one may choose a pleated surface, see [19, Chapter I.5.3]). By Lemma 3.1 in [14] there exists a multicurve Γ of hierarchy curves, each of which has length at most $L_2 = L_2(S, L_0)$ on X such that there are no hierarchy curves supported on the complement of Γ .

Theorem 1.2 in [14] tells us that the projection to W of the set of curves of length at most L_0 in N_ρ lies in a uniformly bounded Hausdorff distance of a $\mathcal{C}(W)$ -geodesic joining $\pi_W(\nu^+(\rho))$ to $\pi_W(\nu^-(\rho))$. So, there exists a constant $a = a(S, L_0)$ so that

$$d_W(\nu^+(\rho), \nu^-(\rho)) > d_W(\alpha, \nu^-(\rho)) - a$$

Choose $\varepsilon_2 > 0$ small enough that every curve on S whose geodesic representative in N_ρ has length at most ε_2 lies in the hierarchy and such that any geodesic of length less than ε_2 on a hyperbolic surface cannot intersect another geodesic of length at most L_2 . If $d_W(\nu^+(\rho), \nu^-(\rho))$ is sufficiently large, then W is a domain in the hierarchy, see [40, Lemma 5.9]. So, combining this fact with Theorem B from [39], we see that there exists $D_1 = D_1(S, \varepsilon_2)$ so that if $d_W(\nu^+(\rho), \nu^-(\rho)) > D_1$, then $\ell_\rho(\partial W) < \varepsilon_2$ and W is a domain in the hierarchy. We may assume that $D > D_1 + a$, so that this is the case.

Since g is 1-Lipschitz and realizes ∂W , no curve in Γ can intersect ∂W and ∂W must be contained in Γ . Moreover, since W is a domain in the hierarchy, there exists $\gamma \in \Gamma \cap \mathcal{C}(W)$. Let γ^X be a representative of γ on $g(X)$ of length at most L_2 .

Let $M_\rho = M(\mu^+(\rho), \mu^-(\rho))$ and let $f_\rho : M_\rho \rightarrow N_\rho$ be the model map. If β is a hierarchy curve, let β^M be the image in N_ρ of the core curve of $U(\beta)$. Since $\ell_X(\alpha^*) \leq L_0$ and $\ell_X(\gamma^X) \leq L_2$, $d_W(\alpha, \gamma) < c = c(S, L_0, L_2)$. Therefore, see [14, Lemma 2.6], there exists $b = b(S, L_0, L_2)$, so that there are at least $d_W(\alpha, \nu^-(\rho)) - b$ curves β on h_W so that $U(\beta)$ lies below $U(\gamma)$ in M_ρ . (In the simple situation when γ lies on h_W , we may take $b = c$ and the curves are simply the curves on h_W coming before γ .) If $U(\beta)$ lies below $U(\gamma)$ in M_ρ , then β^M lies below γ^M in N_ρ . In order




FIGURE 2. When β lies in the indicated range, on the part of h_W to the left of γ , $U(\beta)$ lies below $U(\gamma)$ in the model M_ρ .

to complete the proof, it suffices to bound from above the number of such curves whose geodesic representatives lie above $\hat{\alpha}$.

Let Y be a ruled homotopy from $\alpha^* = \alpha^X$ to $\hat{\alpha}$, and let Y' be a ruled homotopy from γ^X to γ^M (see Figure 3). We claim that if β is a curve in h_W so that β overlaps




FIGURE 3. If H avoids $g(X) \cup Y \cup Y'$ then the ordering of $\hat{\alpha}$ and β^* is consistent with that of γ^M and β^M .

α , β^M lies below γ^M , and there is a homotopy H of β^M to β^* that is disjoint from $g(X) \cup Y \cup Y'$, then β^* does not lie above $\hat{\alpha}$. We can prove this in a few steps.

- (1) Since β^M is disjoint from Y' , γ^X lies above β^M . One concatenates Y' with the homotopy of γ^M to $+\infty$, to obtain a homotopy of γ^X to $+\infty$ disjoint from β^M .
- (2) Since β^M is disjoint from $g(X)$, α^X lies above β^M . Using Lemma 2.6, the fact that γ^X lies above β^M implies that g lies above β^M , and therefore that α^X does as well.
- (3) Since β^M is disjoint from Y , $\hat{\alpha}$ lies above β^M . As in (1), one concatenates Y with the homotopy of α^X to $+\infty$, to obtain a homotopy of α^X to $+\infty$ disjoint from β^M .
- (4) Since $\hat{\alpha}$ lies above β^M and α overlaps β , β^M does *not* lie above $\hat{\alpha}$. Since H is disjoint from $\hat{\alpha}$, β^* also does not lie above $\hat{\alpha}$.

We next observe that there is an upper bound $m(L_h)$ on the number of homotopy classes of primitive curves of length at most L_h that intersect $g(X) \cup Y \cup Y'$. Let $\varepsilon_3 > 0$ be chosen so that an ε_3 -Margulis tube cannot meet a curve of length L_h which is not homotopic to a power of its core curve. So, if a curve of length L_h intersects $g(X)$ it must do so in the image of the ε_3 -thick part of X , which consists of a bounded number of pieces of bounded diameter. Similarly, each of Y and Y' is a union of a bounded diameter neighborhood of its boundary and possibly an annulus inside an ε_3 -Margulis tube. Since a uniformly bounded number of

homotopy classes of primitive curves of length at most L_h can intersect a set of uniformly bounded diameter, we obtain the desired bound $m(L_h)$.

Now, if β is a curve such that β^M lies below γ^M , but β^* lies above $\hat{\alpha}$, then either α does not overlap β or there exists a homotopy H_β from β^M to β^* through curves of length at most L_h which intersects $g(X) \cup Y \cup Y'$. Since at most three simplices of h_W contain curves which do not overlap α , and we have the above bound $m(L_h)$, we conclude that there are at most $d = d(S, L_0)$ such curves. Therefore, $\mathcal{C}(\hat{\alpha}, L)$ contains at least $d_W(\alpha, \nu^-(\rho)) - b - d$ curves. If we also choose $D > K + b + d$ the proof is complete. \square

\square

4. Uniform cores

In this section we show that all of the marked hyperbolic manifolds in $AH_0(M, P)$ can be given compact cores whose geometry is controlled by a finite family of “models”, in a suitable sense. These are quite different from the bilipschitz models for entire manifolds that are constructed in [40, 15], in that they give a model only for a compact subset of the manifold and not for its ends. Because a sequence of elements of $AH_0(M, P)$ can degenerate along annuli in $\Sigma(M, P)$, our notion of control must allow this kind of degeneration. Similarly the possibility of changes of marking inside the characteristic submanifold must be taken into account. We make this precise with the following definition.

If (M, P) is a pared manifold with pared incompressible boundary, a *model core* is given by a triple $(\mathcal{A}, m, \varepsilon)$ where \mathcal{A} is a robust collection of annuli in M , m is a metric on $M_{\mathcal{A}}$, and $\varepsilon \in (0, \mu_3/2)$. If $\rho \in AH_0(M, P)$ we say that a model core $(m, \mathcal{A}, \varepsilon)$ *controls* ρ if there exists an embedding

$$f : (M, P) \rightarrow (N_\rho \setminus \mathbb{T}_\varepsilon(P), \partial \mathbb{T}_\varepsilon(P))$$

such that

- (1) $(f(M), f(P))$ is a relative compact core for $N_\rho \setminus \mathbb{T}_\varepsilon(P)$,
- (2) there is a homeomorphism $\phi : (M, P) \rightarrow (M, P)$ which is the identity on the complement of $\Sigma(M, P)$, such that $f \circ \phi$ is in the homotopy class determined by ρ ,
- (3) the restriction of f to $M_{\mathcal{A}}$ is 2-bilipschitz with respect to m ,
- (4) $f(M_{\mathcal{A}}) \subset N_\rho \setminus \mathbb{T}_{2\varepsilon}(f(\mathcal{A}))$,
- (5) $f(\partial M) \subset N_\rho \setminus (N_\rho)_{(0, \varepsilon)}$.

Here $\mathbb{T}_\varepsilon(f(\mathcal{A}))$ denotes the union of Margulis tubes $\mathbb{T}_\varepsilon(f(A))$ for annuli $A \in \mathcal{A}$.

We call f a *model core map* for $(\mathcal{A}, m, \varepsilon)$.

Theorem 4.1. *If (M, P) is a pared 3-manifold with pared incompressible boundary then, for any $\varepsilon < \mu_3/2$, there exists a finite collection of model cores $(\mathcal{A}, m, \varepsilon)$ such that every $[\rho] \in AH_0(M, P)$ is controlled by one of them.*

Theorem 4.1 will follow directly from this statement:




FIGURE 4. A model core and its controlling map. Here \mathcal{A} has 5 components (indicated by thick segments) whose associated solid tori map to three Margulis tubes.

Theorem 4.2. *If (M, P) is a pared 3-manifold with pared incompressible boundary and $\varepsilon \in (0, \mu_3/2)$, then any sequence in $AH_0(M, P)$ has a subsequence which is controlled by a single model core $(\mathcal{A}, m, \varepsilon)$.*

Indeed, if Theorem 4.1 were to fail, there would be a sequence $\{[\rho_n]\}$ in $AH_0(M, P)$ such that no two elements can be controlled by the same model core. This would contradict Theorem 4.2.

The proof of Theorem 4.2 can be briefly sketched as follows. Starting with a sequence $\{[\rho_n]\}$ in $AH_0(M, P)$, we will identify a robust system of annuli \mathcal{A} which breaks M into pieces on which a subsequence converges, embed these pieces into corresponding geometric limits of the subsequence, and then push those embedded pieces into the approximating manifolds N_{ρ_n} and join them along Margulis tubes to obtain the desired compact cores.

Corollary 2.9 gives the basic convergence result, identifying a robust collection of annuli \mathcal{A} such that some subsequence of $\{[\rho_n]\}$ converges on the pieces of the cut-up manifold $M_{\mathcal{A}}$ after appropriate remarking on the characteristic submanifold. Our basic embedding result is Proposition 4.4 in Section 4.2, which shows that, if we consider a collection of pieces of $M_{\mathcal{A}}$ which live in the same geometric limit, then they admit disjoint embeddings into that geometric limit. The machinery for these embedding theorems is a variation on results of Anderson-Canary [4] and Anderson-Canary-McCullough [6], and Section 4.1 contains some background material and notation for these results.

In Section 4.3 we put these ingredients together to finish the proof. Pieces of $M_{\mathcal{A}}$ which live in one geometric limit can be mapped, via the approximating maps of the geometric limiting process, into the manifolds N_{ρ_n} . Pieces that live in different geometric limits are easy to embed disjointly since their distance in N_{ρ_n} goes to ∞ . Thus we can eventually disjointly embed all the pieces of $M_{\mathcal{A}}$ in N_{ρ_n} . (Something similar was done in the surface-group case in Brock-Canary-Minsky [15].) Note that these embeddings in the geometric limit are pared manifolds, where the frontier annuli $\text{Fr}(M_{\mathcal{A}})$, as well as the annuli of the original pared locus P , map to the boundaries of their associated cusps.

The next step will be to attach the pieces along the Margulis tubes in N_{ρ_n} that approximate the cusps, using solid or thickened tori (whose diameter can be growing as n grows) and obtain compact cores C_n . The identification of M with C_n is at this point on the level of homotopy, and it will require an additional step to obtain a homeomorphism that respects the subdivision into pieces. The geometry of the geometric limits gives us the metric on the pieces which defines the core model m for the subsequence.

4.1. Relative compact carriers

In this section, we recall a criterion from Anderson-Canary-McCullough [6] which guarantees that a collection of subgroups of a Kleinian group is associated to a disjoint collection of submanifolds of the quotient manifold.

We first recall some terminology from Kleinian groups. Let Γ be a finitely generated, torsion-free, non-elementary Kleinian group. We say that Γ is *quasifuchsian* if its domain of discontinuity $\Omega(\Gamma)$ has 2 components, each of which is invariant under the entire group, and is *degenerate* if $\Omega(\Gamma)$ is connected and simply connected. A *component subgroup* of Γ is the stabilizer of a component of its domain of discontinuity. We say that Γ is a *generalized web group* if all of its component subgroups are quasifuchsian. (This includes the case where the domain of discontinuity is empty.) An *accidental parabolic* for Γ is a non-peripheral curve c in $\Omega(\Gamma)/\Gamma$ which is associated to a parabolic element of Γ .

We say that a Γ -invariant collection \mathcal{H} of disjoint horoballs in \mathbb{H}^3 is an *invariant system of horoballs* for a torsion-free Kleinian group Γ if every element of \mathcal{H} is invariant under a non-trivial parabolic subgroup of Γ and every non-trivial parabolic subgroup of Γ fixes some element of \mathcal{H} . If $\varepsilon \in (0, \mu_3)$, then the set of pre-images of the non-compact components of $N_{(0,\varepsilon)}$, where $N = \mathbb{H}^3/\Gamma$, is an invariant system of horoballs for Γ . If $\widehat{\Gamma} \subset \Gamma$ and \mathcal{H} is an invariant system of horoballs for Γ , then the collection $\widehat{\mathcal{H}}$ of elements of \mathcal{H} based at fixed points of non-trivial parabolic subgroups of $\widehat{\Gamma}$ is an invariant system of horoballs for $\widehat{\Gamma}$, which we call the *induced sub-collection of invariant horoballs* for Γ . Notice that if \mathcal{H} is the pre-image of the non-compact components of $\widehat{N}_{(0,\varepsilon)}$, then $\widehat{\mathcal{H}}$ need not be the pre-image of the non-compact components of $N_{(0,\varepsilon)}$. It is this unpleasant fact which necessitates the introduction of this cumbersome terminology. However, for most purposes one can simply imagine that our invariant system of horoballs is associated to the non-compact components of the ε -thin part.

If \mathcal{H} is an invariant system of horoballs for a finitely generated, torsion-free Kleinian group Γ , let

$$N^{\mathcal{H}} = \left(\mathbb{H}^3 - \bigcup_{H \in \mathcal{H}} H \right) / \Gamma.$$

For example, if \mathcal{H} is the pre-image of the non-compact components of $N_{(0, \varepsilon_0)}$, then $N^0 = N^{\mathcal{H}}$. If (M, P) is a relative compact core for $N^{\mathcal{H}}$, then Γ has an accidental parabolic if and only if there is an essential annulus in (M, P) joining a geometrically finite component of $\partial_0(M, P)$ to a component of P . If every geometrically finite component of $\partial_0(M, P)$ is incompressible and Γ has no accidental parabolics, then it is either a generalized web group or degenerate (see [6, Lemma 3.2]).

We will make crucial use of a criterion which guarantees that a collection of subgroups of a Kleinian group is associated to a collection of disjoint submanifolds. A finitely generated subgroup Θ of a Kleinian group Γ is *precisely embedded* if it is the stabilizer, in Γ , of its limit set and if $\gamma \in \Gamma - \Theta$, then there is a component of $\Omega(\Theta)$ whose closure contains $\gamma(\Lambda(\Theta))$ (here Λ denotes the limit set of a Kleinian group). More generally, a collection $\{\Gamma_1, \dots, \Gamma_n\}$ of precisely embedded subgroups of Γ is a *precisely embedded system* if whenever $\gamma \in \Gamma$ and $i \neq j$, then there is a component of $\Omega(\Gamma_i)$ whose closure contains $\gamma(\Lambda(\Gamma_j))$.

We now give a strong pared version of what it means for a subgroup to be associated to a submanifold of the quotient manifold. If $\widehat{\Gamma}$ is a subgroup of a torsion-free, finitely generated Kleinian group Γ with invariant system of horoballs \mathcal{H} , then (Y, Z) is a *relative compact carrier* for Γ if $Y \subset N^{\mathcal{H}}$, $Z = Y \cap \partial N^{\mathcal{H}}$, and (Y, Z) lifts to a relative compact core for $\widehat{N}^{\mathcal{H}}$, where $\widehat{N} = \mathbb{H}^3 / \widehat{\Gamma}$, $N = \mathbb{H}^3 / \Gamma$ and $\widehat{\mathcal{H}}$ is the induced subcollection of invariant horoballs for Γ .

The following result combines Lemma 4.1 and Proposition 4.2 in [6].

Proposition 4.3. *Let Γ be a torsion-free Kleinian group with invariant system of horoballs \mathcal{H} and let $\{\Gamma_1, \dots, \Gamma_n\}$ be a collection of non-conjugate generalized web subgroups of Γ . Then $\{\Gamma_1, \dots, \Gamma_n\}$ is a precisely embedded system if and only if there exists a disjoint collection $\{R_1, \dots, R_n\}$ of compact submanifolds of $N^{\mathcal{H}}$ such that, for all j , R_j is a relative compact carrier for Γ_j and no component of $N^{\mathcal{H}} - R_j$ is a compact twisted I -bundle whose associated ∂I -bundles lies in ∂R_j .*

4.2. Partial Cores

In this section we establish our basic embedding result for partially convergent sequences in $D(M, P)$. In our setting we have a robust collection \mathcal{A} of essential annuli and a sequence $\{\rho_n \in D(M, P)\}$ which is convergent up to inner automorphism on some subset of components of $M_{\mathcal{A}}$ (this situation will arise, in Section 4.3, using Corollary 2.9). Assuming also that the sequence has a geometric limit, and with additional assumptions on parabolics, we will obtain a collection of disjoint relative compact carriers in the geometric limit for the convergent components.

The following convention will be used from now on: Given a component R of $M_{\mathcal{A}}$ on which $\{\rho_n\}$ converges up to inner automorphism, there exists a sequence $\{g_n\} \subset \pi_1(M)$ such that $\{\rho_n^{\rho_n(g_n)}|_{\pi_1(R)}\}$ converges in $D(\pi_1(R))$. We take $\rho^R \in D(R)$

to be the limit of such a sequence. Notice that ρ^R is well-defined up to conjugacy in the geometric limit group.

Proposition 4.4. *Let (M, P) be a pared 3-manifold with pared incompressible boundary and let \mathcal{A} be a robust collection of essential annuli in (M, P) . Suppose that $\{\rho_n\}$ is a sequence in $D(M, P)$ which is convergent up to inner automorphism on a collection \mathcal{R} of components of $(M_{\mathcal{A}}, P_{\mathcal{A}})$ so that*

- (1) $\ell_{\rho_n}(\mathcal{A}) \rightarrow 0$,
- (2) $\{\rho_n(\pi_1(M))\}$ converges geometrically to Γ , and
- (3) if B is an essential annulus in a component of $(M_{\mathcal{A}}, P_{\mathcal{A}})$, then $\lim \ell_{\rho_n}(B) > 0$.

If $\varepsilon \in (0, \varepsilon_0]$, then there exists a disjoint collection $\{(Y_R, Z_R)\}_{(R, Q) \in \mathcal{R}}$ of relative compact carriers for $\{\rho^R(\pi_1(R))\}_{(R, Q) \in \mathcal{R}}$ in N^ε where $N = \mathbb{H}^3/\Gamma$ and N^ε is obtained from N by removing the non-compact components of $N_{(0, \varepsilon)}$.

We begin by showing that if $(R, Q), (S, T) \in \mathcal{R}$, then any conjugate of the limit set of $\rho^R(\pi_1(R))$ intersects $\rho^S(\pi_1(S))$ in at most one point. We next show that each $\rho^R(\pi_1(R))$ is either a generalized web group or a degenerate group without accidental parabolics. We can then use Proposition 4.3 and the Covering Theorem [18, 44] to complete the proof of Proposition 4.4, as in the proof of [15, Prop 6.4].

Our first lemma is a common generalization of Proposition 2.7 from [5] and Proposition 6.7 from [15].

Lemma 4.5. *Let (M, P) be a pared 3-manifold with pared incompressible boundary and let \mathcal{A} be a robust collection of essential annuli in (M, P) . Suppose that $\{\rho_n\}$ is a sequence in $D(M, P)$ which is convergent up to inner automorphism on a collection \mathcal{R} of components of $(M_{\mathcal{A}}, P_{\mathcal{A}})$ so that*

- (1) $\ell_{\rho_n}(\mathcal{A}) \rightarrow 0$, and
- (2) $\{\rho_n(\pi_1(M))\}$ converges geometrically to Γ .

If (i) R and S are distinct elements of \mathcal{R} and $\gamma \in \Gamma$ or (ii) $R = S$ and $\gamma \notin \rho^R(\pi_1(R))$, then the intersection of limit sets

$$\Lambda(\gamma \rho^R(\pi_1(R)) \gamma^{-1}) \cap \Lambda(\rho^S(\pi_1(S)))$$

contains at most one point.

Proof. A result of Anderson [2] and Soma [43] implies that if Φ_1 and Φ_2 are non-elementary, finitely generated subgroups of a torsion-free Kleinian group Γ , then

$$\Lambda(\Phi_1) \cap \Lambda(\Phi_2) = \Lambda(\Phi_1 \cap \Phi_2) \cup P(\Phi_1, \Phi_2)$$

where $P(\Phi_1, \Phi_2)$ is the set of points $x \in \Lambda(\Gamma)$ such that the stabilizers of x in Φ_1 and Φ_2 are rank one parabolic subgroups which generate a rank two parabolic subgroup of Γ . (Anderson and Soma's results require that the subgroups are topologically tame, but it is now known, by work of Agol [1] and Calegari-Gabai [17], that all finitely generated, torsion-free Kleinian groups are topologically tame.)

Without loss of generality we can assume that $\{\rho_n|_{\pi_1(R)}\}$ converges to ρ^R and there exists $g_n \in \pi_1(M)$ so that $\{\rho_n^{\rho_n(g_n)}|_{\pi_1(S)}\}$ converges to ρ^S .

One may show that $\gamma\rho^R(\pi_1(R))\gamma^{-1} \cap \rho^S(\pi_1(S))$ is purely parabolic, i.e. consists entirely of parabolic and trivial elements, exactly as in the proofs of [5, Lemma 2.4] and [15, Lemma 6.7]. Therefore, $\Lambda(\gamma\rho^R(\pi_1(R))\gamma^{-1} \cap \rho^S(\pi_1(S)))$ contains at most one point. Thus, it only remains to prove that

$$P(\gamma\rho^R(\pi_1(R))\gamma^{-1}, \rho^S(\pi_1(S))) = \emptyset.$$

Suppose that $x \in P(\gamma\rho^R(\pi_1(R))\gamma^{-1}, \rho^S(\pi_1(S)))$. The stabilizer $\text{stab}_{\rho^S(\pi_1(S))}(x)$ is generated by $\rho^S(s)$, and $\text{stab}_{\gamma\rho^R(\pi_1(R))\gamma^{-1}}(x)$ is generated by $\gamma\rho^R(r)\gamma^{-1}$ where r and s are primitive elements of $\pi_1(R)$ and $\pi_1(S)$, respectively. Since these two elements must commute, $h_n r h_n^{-1}$ commutes with $g_n s g_n^{-1}$ for sufficiently large n , see [15, Prop. 2.7].

If $h_n r h_n^{-1}$ and $g_n s g_n^{-1}$ lie in a cyclic subgroup of $\pi_1(M)$, then, since they are primitive, they agree, up to taking inverses, so their limits cannot generate a rank two abelian group. Otherwise, $h_n r h_n^{-1}$ and $g_n s g_n^{-1}$ are primitive elements generating a rank two abelian subgroup Δ_n of $\pi_1(M)$ for all sufficiently large n . This can only occur if $R = S$, and the rank two abelian subgroup Δ_n is the maximal rank two abelian subgroup Δ of $\pi_1(R)$ containing s . So, γ lies in the geometric limit of $\{\rho_n(\Delta)\}$, which is simply $\rho^R(\Delta)$ (since any element of the geometric limit of $\{\rho_n(\Delta)\}$ is a parabolic element which shares a fixed point with $\rho^R(\Delta)$, and hence, since $\rho^R(\Delta)$ has rank two and the geometric limit is discrete, has a power which lies in $\rho^R(\Delta)$, which is again ruled out by an application of [15, Prop. 2.7]). We have achieved a contradiction, so $P(\gamma\rho^R(\pi_1(R))\gamma^{-1}, \rho^S(\pi_1(S)))$ is empty, which completes the proof. \square

The following lemma is the main new ingredient in our proof.

Lemma 4.6. *If we are in the setting of Proposition 4.4 and (R, Q) is a component of \mathcal{R} , then $\rho^R(\pi_1(R))$ is either a generalized web group or a degenerate group without accidental parabolics.*

Proof. Let \mathcal{H} be the invariant system of horoballs for Γ obtained by considering the pre-images of the non-compact components of $N_{(0, \varepsilon)}$ where $N = \mathbb{H}^3/\Gamma$. If $(R, Q) \in \mathcal{R}$, let \mathcal{H}_R be the induced subcollection of \mathcal{H} associated to the subgroup $\rho^R(\pi_1(R))$ and let $N_R^{\mathcal{H}} = N_{\rho^R}^{\mathcal{H}_R}$.

We first show that a relative compact core (M_R, P_R) of $N_R^{\mathcal{H}}$ has pared incompressible boundary. Recall that Lemma 2.2 implies that (R, Q) is a pared manifold with pared incompressible boundary. Let P_R^Q denote the collection of components of P_R whose fundamental groups are in the conjugacy class of $\rho^R(\pi_1(Q_0))$ for some component Q_0 of Q . Since there is a pared homotopy equivalence

$$j : (R, Q) \rightarrow (M_R, P_R^Q),$$

every component of $\partial_0(M_R, P_R^Q)$ is incompressible (see Bonahon [10, Prop. 1.2] or [20, Lemma 5.2.1]). Since $\partial_0(M_R, P_R)$ is obtained from $\partial_0(M_R, P_R^Q)$ by removing incompressible annuli, $\partial_0(M_R, P_R)$ is also incompressible, so (M_R, P_R) also has pared incompressible boundary.

We next show that $\rho^R(\pi_1(R))$ has no accidental parabolics. If $\rho^R(\pi_1(R))$ has an accidental parabolic, there exists an essential annulus E in (M_R, P_R) so that one component of ∂E is contained in P_R and the other is contained in $\partial_0(M_R, P_R)$. Notice that E is also an essential annulus in (M_R, P_R^Q) . Thus, there is a component Σ_0 of the characteristic submanifold $\Sigma(X_R, P_R^Q)$, so that E is isotopic into Σ_0 . Johannson's Classification Theorem [28, Thm. 24.2] implies that we may assume that the pared homotopy equivalence j between (R, Q) and (M_R, P_R^Q) has the property that $j(\Sigma(R, Q)) = \Sigma(M_R, P_R^Q)$ and j is a homeomorphism from the (closure of the) complement of $\Sigma(R, Q)$ to the (closure of the) complement of $\Sigma(M_R, P_R^Q)$. If Σ_0 is a solid or thickened torus component of $\Sigma(M_R, P_R^Q)$, then $\Sigma_1 = j^{-1}(\Sigma_0)$ is a solid or thickened torus component of $\Sigma(R, Q)$. It follows that any annulus B in the frontier of Σ_1 would be an essential annulus in (R, Q) with the property that $\ell_{\rho^R}(B) = 0$, which is disallowed by our assumptions. If Σ_0 is an interval bundle component of $\Sigma(M_R, P_R^Q)$, then $\Sigma_1 = j^{-1}(\Sigma_0)$ is an interval bundle component of $\Sigma(R, Q)$. However, by [20, Lemma 2.11.3], the restriction of j to $(\Sigma_1, \text{Fr}(\Sigma_1))$ is pared homotopic to a pared homeomorphism $h : (\Sigma_1, \text{Fr}(\Sigma_1)) \rightarrow (\Sigma_0, \text{Fr}(\Sigma_0))$. Then, $B = h^{-1}(E)$ would again be an essential annulus in (R, Q) with the property that $\ell_{\rho^R}(B) = 0$, which is again disallowed by our assumptions. Therefore, $\rho^R(\pi_1(R))$ has no accidental parabolics.

Since the relative compact core of $N_R^{\mathcal{H}}$ has pared incompressible boundary and $\rho^R(\pi_1(R))$ has no accidental parabolics, Lemma 3.2 in [6] implies that $\rho^R(\pi_1(R))$ is either a generalized web group or degenerate group without accidental parabolics. \square

Proof of Proposition 4.4: Let $\{\Gamma^1, \dots, \Gamma^s\} = \{\rho^R(\pi_1(R))\}_{(R, Q) \in \mathcal{R}}$. We re-order so that $\{\Gamma^1, \dots, \Gamma^r\}$ are generalized web groups and $\{\Gamma^{r+1}, \dots, \Gamma^s\}$ are degenerate groups without accidental parabolics. Let (R_i, Q_i) be the component of \mathcal{R} , so that $\Gamma_i = \rho^{R_i}(\pi_1(R_i))$.

For $i, j \in \{1, \dots, r\}$, Lemma 4.5 implies that if $\gamma \in \Gamma$ and $i \neq j$ or if $i = j$ and $\gamma \in \Gamma - \Gamma^i$, then the limit sets of $\gamma\Gamma^i\gamma^{-1}$ and Γ^j intersect in at most one point. Since each component of $\Omega(\Gamma^i)$ is bounded by a quasi-circle, we immediately conclude that $\gamma(\Lambda(\Gamma^j))$ is contained in the closure of a component of $\Omega(\Gamma^i)$. Therefore, $\{\Gamma^1, \dots, \Gamma^r\}$ is a precisely embedded collection of generalized web subgroups of Γ . Proposition 4.3 then implies that there is an associated collection $\{(Y_1, Z_1), \dots, (Y_r, Z_r)\}$ of disjoint relative compact carriers for $\{\Gamma^1, \dots, \Gamma^r\}$ in $N^{\mathcal{H}}$.

Since Γ^{r+1} is a degenerate group without accidental parabolics, the Covering Theorem [18, 44] may be used, exactly as in the proof of [15, Prop. 6.10], to show that there exists a neighborhood U of an end of $N_{R_{r+1}}^{\mathcal{H}} \cong F_{r+1} \times \mathbb{R}$, for some compact surface F_{r+1} , which is identified with $F_{r+1} \times (k_{r+1}, \infty)$ and $\pi_{R_{r+1}}(U)$ embeds in N , under the obvious covering map $p_{r+1} : N_{R_{r+1}}^{\mathcal{H}} \rightarrow N$. One may then choose a relative compact carrier (Y_{r+1}, Z_{r+1}) for Γ_{r+1} of the form

$$(Y_{r+1}, Z_{r+1}) = (p_{r+1}(F_{r+1} \times [s, s+1]), p_{r+1}(\partial F_r \times [s, s+1]))$$

for some $s > k_{r+1}$ so that Y_{r+1} is disjoint from $Y_1 \cup \dots \cup Y_r$. One similarly uses the Covering Theorem to iteratively choose a relative compact carrier (Y_{r+i}, Z_{r+1}) for Γ_{r+i} which is disjoint from $Y_1 \cup \dots \cup Y_{r+i-1}$. One finally arrives at a disjoint collection $\{(Y_R, Z_R)\}_{(R,Q) \in \mathcal{R}}$ of relative compact carriers for $\{\Gamma_i^s\}_{i=1}^s$. \square

4.3. Proof of the uniform cores theorem

We now complete the proof of Theorem 4.2 (and hence Theorem 4.1). The proof breaks into several steps which we briefly summarize.

The first step is to use the convergence theorems we have been developing to identify a robust system of annuli \mathcal{A} so that (a subsequence of) our given sequence $\{\rho_n\}$ converges on the pieces of $M_{\mathcal{A}}$. We then find an embedded submanifold Y in the union of geometric limits of $\{N_{\rho_n}\}$, and a homotopy equivalence $h : M_{\mathcal{A}} \rightarrow Y$. In Step 1 we do this for each geometric limit separately. In Step 2 we combine the geometric limits and produce an embedding $\psi_n : Y \rightarrow N_{\rho_n}$ for each large enough n . Attaching the pieces of $Y_n = \psi_n(Y)$ along Margulis tubes, we obtain the submanifold $C_n \subset N_{\rho_n}$ which will be our compact core. In Step 3, we combine h and ψ_n to get a map $s_n : M \rightarrow C_n$ in the correct homotopy class.

Note that, at this point, s_n is not known to be a homeomorphism even on the components of $M_{\mathcal{A}}$. Indeed, even in the case of a convergent sequence of representations the homeomorphism type of the limit may differ from that of the approximates, as discovered by Anderson-Canary [3], and this phenomenon accounts for many of the difficulties in this proof.

In Step 4 we show that s_n is a homotopy equivalence and conclude that C_n is in fact a (relative) compact core of N_{ρ_n} . In Step 5 we show that s_n is indeed homotopic to a homeomorphism \widehat{s}_n . Note that here the assumption that $[\rho_n] \in AH_0(M, P)$ is crucial: it means that an embedding from M to N_{ρ_n} in the right homotopy class actually exists. It requires some extra work to ensure that \widehat{s}_n respects the decomposition induced by the annuli \mathcal{A} . Finally in Step 6 we use the geometry of the geometric limits to choose a metric m on $M_{\mathcal{A}}$ which makes our final maps 2-bilipschitz.

Step 1: Embedding the partial cores. Consider a sequence $\{\rho_n\}$ in $D(M, P)$ of representatives of the sequence $\{[\rho_n]\}$. We may apply Corollary 2.9 to remark, subdivide by annuli and obtain a subsequence that converges on the pieces. That is, let $\{\phi_n : (M, P) \rightarrow (M, P)\}$ be the sequence of homeomorphisms which are the identity on the complement of $\Sigma(M, P)$, and \mathcal{A} the robust collection of annuli for which the conclusions of Corollary 2.9 hold. To lighten the notation we can assume, without loss of generality, that each ϕ_n is the identity and that $\{\rho_n\}$ is the subsequence. Hence we have

- (1) $\lim \ell_{\rho_n}(\mathcal{A}) = 0$,
- (2) $\{\rho_n\}$ converges on $M_{\mathcal{A}}$ up to conjugacy, and
- (3) if B is an essential annulus in $(M_{\mathcal{A}}, P_{\mathcal{A}})$, then

$$\lim \ell_{\rho_n}(B) > 0.$$

For each component (R, Q) of $M_{\mathcal{A}}$ we have a sequence $\{\rho_n^{g_n}\}$ of conjugates which converge on $\pi_1(R)$ to ρ^R , and a corresponding sequence $\{b_n^R \in N_{\rho_n}\}$ of base-points, which are the projections of a fixed basepoint $\mathbf{0} \in \mathbb{H}^3$ as in the setup of Lemma 2.4. We may pass to a subsequence so that if (R, Q) and (S, T) are components of $M_{\mathcal{A}}$, then $d(R, S) = \lim d(b_n^R, b_n^S)$ exists and lives in $[0, \infty]$. We say two components (R, Q) and (S, T) are *nearby* if $d(R, S)$ is finite. This defines an equivalence relation on the set of components of $M_{\mathcal{A}}$. For an equivalence class \mathcal{R} of components, Lemma 2.4 implies that there is a single sequence of conjugates $\{\rho_n^{g_n}\}$ which converges up to inner automorphisms on \mathcal{R} , and (again passing to a subsequence) shares a geometric limit $\Gamma^{\mathcal{R}}$ with quotient $N_{\mathcal{R}}$.

We can then apply Proposition 4.4 to obtain, for each equivalence class \mathcal{R} , a collection $\{(Y_R, Z_R)\}_{(R, Q) \in \mathcal{R}}$ of disjoint relative compact carriers for $\{\rho^R(\pi_1(R))\}_{(R, Q) \in \mathcal{R}}$ in $N_{\mathcal{R}}^{2\epsilon}$. Let $Y_{\mathcal{R}} = \bigcup_{(R, Q) \in \mathcal{R}} Y_R$ and $Z_{\mathcal{R}} = \bigcup_{(R, Q) \in \mathcal{R}} Z_R$.

For each $(R, Q) \in \mathcal{R}$, let Z_R^Q be the collection of components of Z_R which lie in the homotopy class of $\rho^R(\pi_1(Q))$. There exists a map of pairs

$$h_R : (R, Q) \rightarrow (Y_R, Z_R^Q)$$

which is a homotopy equivalence from R to Y_R and restricts to an orientation-preserving embedding on Q .

The following elementary lemma from [20] implies that h_R is a homotopy equivalence of pairs.

Lemma 4.7. (Canary-McCullough [20, Prop. 5.2.3]) *If (M_1, P_1) and (M_2, P_2) are pared manifolds and $h : (M_1, P_1) \rightarrow (M_2, P_2)$ is a map of pairs, so that h is a homotopy equivalence from M_1 to M_2 , then h is a pared homotopy equivalence.*

We may combine the maps on each component to obtain a map

$$h_{\mathcal{R}} = \bigcup h_R : \bigcup_{(R, Q) \in \mathcal{R}} (R, Q) \rightarrow (Y_{\mathcal{R}}, Z_{\mathcal{R}}).$$

Step 2: Building the core C_n . A crucial step in the proof is the construction of a compact core C_n in the approximate N_{ρ_n} for all large enough n . Using elementary facts about geometric limits we will pull back the carriers $Y_{\mathcal{R}}$ to submanifolds of N_{ρ_n} , and then attach them along the thin parts associated to the annuli \mathcal{A} to obtain C_n together with maps from M to C_n which we will then show are homotopy equivalences, and eventually promote to homeomorphisms.

We first make a small adjustment to $(Y_{\mathcal{R}}, Z_{\mathcal{R}})$, caused by a technical need to assure that C_n may be pared so that it is a relative compact core for $N_{\rho_n} \setminus \mathbb{T}_{\epsilon, n}(P)$, whereas the annuli of $Z_{\mathcal{R}}$ have been chosen to lie in $\partial \mathbb{T}_{2\epsilon}(Z_{\mathcal{R}})$. Thus, for each annulus P_0 of P which lies in a component (R, Q) of $M_{\mathcal{A}}$, let $Z_0 = h_R(P_0)$ be the associated component of Z_R . We “push” Z_0 into $\partial \mathbb{T}_{\epsilon}(Z_0)$ by adding to Y_R a radial collar of Z_0 in $\mathbb{T}_{2\epsilon}(Z_0) \setminus \mathbb{T}_{\epsilon}(Z_0)$. To simplify notation we continue to call the resulting pared manifolds (Y_R, Z_R) , and the homotopy equivalences h_R .

For all large enough n , Lemma 2.5 provides a 2-bilipschitz map

$$\psi_n^{\mathcal{R}} : Y_{\mathcal{R}} \rightarrow N_{\rho_n}$$

such that for each $(R, Q) \in \mathcal{R}$,

- (1) $\psi_n|_{Y_R}$ is in the homotopy class determined by $\rho_n \circ (h_R)_*^{-1}$,
- (2) If $Z_0 = h_R(P_0)$ and P_0 is a component of $Q \cap P$, then

$$\psi_n(Y_R) \subset N_{\rho_n} \setminus \mathbb{T}_{\varepsilon, n}(P_0) \quad \text{and} \quad \psi_n(Z_0) \subset \partial \mathbb{T}_{\varepsilon, n}(P_0).$$

- (3) If $Z_0 = h_R(Q_0)$ and Q_0 is a component of $Q \setminus P$, then

$$\psi_n(Y_R) \subset N_{\rho_n} \setminus \mathbb{T}_{2\varepsilon, n}(Q_0) \quad \text{and} \quad \psi_n(Z_0) \subset \partial \mathbb{T}_{2\varepsilon, n}(Q_0).$$

(Here, $\mathbb{T}_{\varepsilon, n}(P_0)$ denotes the component of $(N_{\rho_n})_{(0, \varepsilon)}$ in the homotopy class of P_0 .)

If \mathcal{R} and \mathcal{R}' are distinct equivalence classes of components of $M_{\mathcal{A}}$, the fact that $d(R, R') = \infty$ for all $(R, Q) \in \mathcal{R}$ and $(R', Q') \in \mathcal{R}'$ implies that $\psi_n^{\mathcal{R}}(Y_{\mathcal{R}})$ and $\psi_n^{\mathcal{R}'}(Y_{\mathcal{R}'})$ are disjoint for all large enough values of n . To see this, note first that Lemma 2.4 tells us that the basepoints b_n^R for $(R, Q) \in \mathcal{R}$ remain a bounded distance from the basepoints $b_n^{\mathcal{R}}$ for the common sequence of conjugates that converges for all $R \in \mathcal{R}$. Now the maps $\psi_n^{\mathcal{R}}$ take the limiting basepoint $b^{\mathcal{R}}$ to $b_n^{\mathcal{R}}$ (see Lemma 2.5) so we may conclude that $\psi_n^{\mathcal{R}}(Y_{\mathcal{R}})$ remain at bounded distance from $b_n^{\mathcal{R}}$ for each $(R, Q) \in \mathcal{R}$, and similarly for \mathcal{R}' . Hence the distance between $\psi_n^{\mathcal{R}}(Y_{\mathcal{R}})$ and $\psi_n^{\mathcal{R}'}(Y_{\mathcal{R}'})$ goes to ∞ , so they are eventually disjoint. Now let

$$Y = \bigcup Y_{\mathcal{R}}, \quad h = \bigcup h_{\mathcal{R}}, \quad \psi_n = \bigcup \psi_n^{\mathcal{R}}, \quad \text{and} \quad Y_n = \psi_n(Y),$$

noting that $\psi_n : Y \rightarrow Y_n$ is an embedding for large n .

We now construct a submanifold C_n as the union of Y_n with solid or thickened tori in $\mathbb{T}_{2\varepsilon, n}(\mathcal{A}) - \mathbb{T}_{\varepsilon, n}(\mathcal{A})$, each of which is attached to one or more annuli in $\psi_n(h(\text{Fr}(M_{\mathcal{A}})))$. Let $\mathbb{T}_{2\varepsilon, n}(A_i)$ be a Margulis tube for a component (possibly several) A_i of \mathcal{A} . If $\ell_{\rho_n}(A_i) > 0$, we append $\mathbb{T}_{2\varepsilon, n}(A_i)$ to Y_n . If not, we choose an annulus (if $\mathbb{T}_{2\varepsilon, n}(A_i)$ is a rank one cusp) or torus (if $\mathbb{T}_{2\varepsilon, n}(A_i)$ is a rank two cusp) B_i in $\partial \mathbb{T}_{2\varepsilon, n}(A_i)$ which contains $Y_n \cap \partial \mathbb{T}_{\varepsilon, n}(A_i)$ in its interior, and append to Y_n the set of points in $\mathbb{T}_{2\varepsilon, n}(A_i) \setminus \mathbb{T}_{\varepsilon, n}(A_i)$ which project radially to B_i .

To give C_n a pared structure, note first that each annulus $P_0 \subset P$ in a component (R, Q) of $M_{\mathcal{A}}$ is already taken by h_R to an annulus in $\partial \mathbb{T}_{\varepsilon, n}(P_0)$. If P_0 is a component of P lying in a component U of $M \setminus M_{\mathcal{A}}$, then U corresponds to a cusp $\mathbb{T}_{2\varepsilon, n}(U)$ and the intersection of C_n with $\partial \mathbb{T}_{\varepsilon, n}(U)$ is an annulus or torus. Putting all these components together we obtain a pared locus $P_n = C_n \cap \partial \mathbb{T}_{\varepsilon, n}(P)$.

Step 3: Mapping M to C_n . Next we define a map $s_n : (M, P) \rightarrow (C_n, P_n)$, such that

- s_n is in the homotopy class of ρ_n , viewed as a map from M to N_{ρ_n} ,
- $s_n|_{M_{\mathcal{A}}} = \psi_n \circ h$,
- s_n maps P to P_n homeomorphically.

Since $[\rho_n] \in AH_0(M, P)$, there exists an orientation-preserving embedding

$$\kappa_n : (M, P) \rightarrow (N_{\rho_n}^{\varepsilon}, \partial N_{\rho_n}^{\varepsilon})$$

in the homotopy class of ρ_n . After an isotopy along the boundary, we may assume that $\kappa_n(P)$ is the pared locus P_n of C_n .

Since $\psi_n \circ h$ is homotopic to $\kappa_n|_{M_{\mathcal{A}}}$, and each element of $\mathcal{A} \cup P$ is collared in M , there exists an extension s'_n of $\psi_n \circ h$ to all of M , which is homotopic to κ_n as a

map of pairs and equals κ_n on P . If U is a component of $M \setminus M_{\mathcal{A}}$ then it is a solid or thickened torus, and $\text{Fr}(U)$ is mapped by s'_n into the boundary of the associated Margulis tube $\mathbb{T}_{2\varepsilon,n}(U)$. Since $\pi_1(U)$ is a maximal abelian subgroup of $\pi_1(M)$ and $\pi_1(U)$ is its own centralizer in $\pi_1(M)$, the restriction $s'_n|_U$ is homotopic, relative to $\text{Fr}(U) \cup (\partial U \cap P)$, into $C_n \cap \mathbb{T}_{2\varepsilon,n}(U)$. After performing such a homotopy for each U we obtain the desired map s_n .

Step 4: Showing C_n is a core. We now claim that (C_n, P_n) is a relative compact core for $N_n - \mathbb{T}_{\varepsilon,n}(P)$. To prove this we first note the following lemma proved in [6].

Lemma 4.8. ([6, Lemma 5.2]) *For $i = 1, 2$, let M_i be a compact, orientable, irreducible 3-manifold with incompressible boundary and let V_i be a 3-dimensional submanifold whose frontier $\text{Fr}(V_i)$ is non-empty and incompressible. If $g : M_1 \rightarrow M_2$ is a continuous map such that*

- (1) $g^{-1}(V_2) = V_1$,
- (2) g restricts to a homeomorphism from V_1 to V_2 , and
- (3) g restricts to a homotopy equivalence from $M_1 \setminus V_1$ to $M_2 \setminus V_2$,

then g is a homotopy equivalence.

We can apply this lemma with $M_1 = M$ and $M_2 = C_n$, where V_1 is a collar neighborhood $\mathcal{N}(\mathcal{A})$. Since s_n by construction is an embedding on \mathcal{A} and takes $M \setminus \mathcal{A}$ to the complement of $s_n(\mathcal{A})$, we can assume after a small homotopy that s_n is an embedding on V_1 , and setting $V_2 = s_n(V_1)$, that $s_n^{-1}(V_2) = V_1$ as well. The components of $M \setminus \mathcal{A}$ are isotopic to components of $M_{\mathcal{A}}$ or to the solid or thickened tori of $M \setminus M_{\mathcal{A}}$, and s_n is a homotopy equivalence from each of these to the corresponding component of $C_n \setminus s_n(\mathcal{A})$, again by construction. Moreover we note that s_n is bijective on the components of $M \setminus \mathcal{A}$ since no two components have conjugate fundamental groups, and that it is surjective to the components of $C_n \setminus s_n(\mathcal{A})$ by definition of C_n . We conclude that s_n is a homotopy equivalence of $M_1 \setminus V_1$ to $M_2 \setminus V_2$. Lemma 4.8 therefore implies that $s_n : M \rightarrow C_n$ is a homotopy equivalence.

Since s_n is in the homotopy class of ρ_n , as a map into N_{ρ_n} , the inclusion of $\pi_1(C_n)$ into $\pi_1(N_{\rho_n})$ is an isomorphism, so C_n is a compact core for N_{ρ_n} . Then, by construction, (C_n, P_n) is a relative compact core for $N_{\rho_n}^{\varepsilon}$. Lemma 4.7 implies that

$$s_n : (M, P) \rightarrow (C_n, P_n)$$

is a homotopy equivalence of pairs. It follows from this that (C_n, P_n) is a relative compact core for $N_{\rho_n} \setminus \mathbb{T}_{\varepsilon,n}(P)$.

Step 5: Homotopy to the final embedding. We next show that $s_n : (M, P) \rightarrow (C_n, P_n)$ is pared homotopic to a homeomorphism. Recall that relative compact cores are unique up to admissible isotopy and their complements have product structures. (See section 2.1.) Therefore, there exists a pared homeomorphism $g_n : (C_n, P_n) \rightarrow (\kappa_n(M), P_n)$ in the isotopy class of the inclusion map. Then $\widehat{s}_n = g_n^{-1} \circ \kappa_n$ is a pared homeomorphism in the homotopy class of ρ_n and hence homotopic to s_n . Since $N_{\rho_n}^{\varepsilon}$ deformation retracts onto C_n , the homotopy between s_n and \widehat{s}_n can be presumed to lie in C_n .

We wish to show that \widehat{s}_n can be admissibly isotoped so that it preserves the decomposition of \mathcal{A} . That is, it takes any component R of $M_{\mathcal{A}}$ to $\psi_n(Y_R)$ and any component U of $M - M_{\mathcal{A}}$ to $\mathbb{T}_{2\varepsilon,n}(U) \cap C_n$. To do this, we first show that \widehat{s}_n can be admissibly isotoped so that $\widehat{s}_n(\text{Fr}(M_{\mathcal{A}})) = s_n(\text{Fr}(M_{\mathcal{A}}))$.

Let \mathcal{B} denote the collection of components of $\text{Fr}(M_{\mathcal{A}})$. If $B \in \mathcal{B}$, let $\widehat{B} = \widehat{s}_n^{-1}(s_n(B))$ and

$$\widehat{\mathcal{B}} = \{\widehat{B}\}_{B \in \mathcal{B}}.$$

Notice that if $B \in \mathcal{B}$, then $s_n(B)$ is an essential annulus in (C_n, P_n) , since otherwise it would bound a solid torus V in C_n such that $\partial V - B \subset \partial C_n$. So, since \widehat{s}_n is a pared homeomorphism, \widehat{B} is an essential annulus in (M, P) . The enclosing property for characteristic submanifolds, see Johannson [28, Prop. 10.7], implies that $\widehat{\mathcal{B}}$ is admissibly isotopic into $\Sigma(M, P)$, so we may assume that $\widehat{\mathcal{B}}$ is contained in the interior of $\Sigma(M, P)$.

If a component B of $\widehat{\mathcal{B}}$ is homotopic into a solid or thickened torus component X of $\Sigma(M, P)$ but is not contained in X , it must be contained in an adjacent interval bundle component, and cobound a solid torus with a component of $\text{Fr}(X)$. Since $\widehat{\mathcal{B}}$ is embedded, all these solid tori are disjoint or nested. Thus, after adjusting \widehat{s}_n by an isotopy supported on a regular neighborhood of the union of such solid tori, we may assume that each component of $\widehat{\mathcal{B}}$ that is homotopic to a solid or thickened torus component of $\Sigma(M, P)$ is already contained in it.

Solid/thickened torus components: Consider a component V of $M \setminus M_{\mathcal{A}}$ which is a regular neighborhood of a solid or thickened torus component X of $\Sigma(M, P)$. Let $\mathcal{B}_V \subset \mathcal{B}$ be the components of $\text{Fr}(V)$ and let $\widehat{\mathcal{B}}_V$ denote the corresponding subset of $\widehat{\mathcal{B}}$. By the previous paragraph we know that the components of \mathcal{B}_V lie in the interior of X .

The pair $(X, \text{Fr}(X))$ admits a Seifert-fibration over (E, d) , where E is either a disk with d a collection of arcs in ∂E , or an annulus with d a collection of arcs in one component of ∂E . There is at most one singular point in E if it is a disk, and none if it is an annulus. Since $\widehat{\mathcal{B}}_V$ is a collection of essential annuli in X , it may be isotoped, by an isotopy supported on X , to the S^1 -bundle over a collection e of arcs in E with the end points of each arc in distinct components of the complement of d (see Johannson [28, Prop. 5.6]). Since $s_n(\mathcal{B}_V)$ bounds a solid or thickened torus in C_n whose fundamental group is a maximal abelian subgroup, the same is true for $\widehat{\mathcal{B}}_V$. It follows that the arcs e are the boundary in E of a connected region W which is either an essential subannulus when E is an annulus, or a disk when E is a disk. Moreover if E has a singular point then it must lie in W , since otherwise the preimage of W would be a solid torus whose core is homotopic to a proper power of the core of X , and hence generates a non-maximal subgroup of $\pi_1(M)$. Since the cardinalities of e and d are equal, it follows that each component of e must be parallel to a component of d across a region that does not contain a singular point. It follows that $\widehat{\mathcal{B}}_V$ is isotopic to the frontier of X , and hence to \mathcal{B}_V , by an isotopy supported on V .

Interval bundle components: If Σ is an interval bundle component of $\Sigma(M, P)$ with base surface F , let $\widehat{\Sigma}$ be obtained from Σ by appending the closure of $\mathcal{N}(A)$ whenever A is a component of $\text{Fr}(\Sigma)$ contained in \mathcal{A} . Then $\widehat{\Sigma}$ is an interval bundle with base surface \widehat{F} (which is obtained from F by appending collar neighborhoods of certain components of ∂F). Let \mathcal{B}_Σ denote the components of \mathcal{B} which are contained in $\widehat{\Sigma}$ and not isotopic into a solid or thickened torus component of $\Sigma(M, P)$. By construction, the corresponding subset $\widehat{\mathcal{B}}_\Sigma$ of $\widehat{\mathcal{B}}$ is contained in the interior of Σ . In an interval bundle, every system of disjoint essential annuli is admissibly isotopic to the subbundle over a multicurve in the base (again see Johannson [28, Prop. 5.6]), and two such multicurves are isotopic if and only if they are homotopic. Since \mathcal{B}_Σ and $\widehat{\mathcal{B}}_\Sigma$ are homotopic, they are isotopic to sub-bundles over homotopic collections of curves, so they are isotopic. Therefore, we may admissibly isotope \widehat{s}_n , by an isotopy supported on a regular neighborhood of $\widehat{\Sigma}$, so that $\widehat{\mathcal{B}}_\Sigma = \mathcal{B}_\Sigma$.

Since the supports of all the resulting isotopies may be chosen to be disjoint, they can be performed simultaneously. Therefore, we have isotoped \widehat{s}_n so that $\widehat{s}_n(\mathcal{B}) = s_n(\mathcal{B})$.

Since \widehat{s}_n is a homeomorphism, $\widehat{s}_n(\text{Fr}(M_{\mathcal{A}})) = s_n(\text{Fr}(M_{\mathcal{A}}))$ and s_n takes components of $M \setminus \text{Fr}(M_{\mathcal{A}})$ to components of $C_n \setminus s_n(\text{Fr}(M_{\mathcal{A}}))$, we see that \widehat{s}_n also takes components of $M \setminus \text{Fr}(M_{\mathcal{A}})$ homeomorphically to components of $C_n \setminus s_n(\text{Fr}(M_{\mathcal{A}}))$. Notice that every component of $M_{\mathcal{A}}$ has non-abelian fundamental group, while every component of $M - M_{\mathcal{A}}$ has abelian fundamental group. Since \widehat{s}_n is homotopic to s_n , and no two components of $M_{\mathcal{A}}$ are homotopic, we see that $\widehat{s}_n(R) = s_n(R) = \psi_n(Y_R)$ if R is a component of $M_{\mathcal{A}}$. Similarly, if U is a component of $M \setminus M_{\mathcal{A}}$, then $\widehat{s}_n(U) = s_n(U) = C_n \cap \mathbb{T}_{2\varepsilon, n}(U)$. In short, our decomposition is preserved by \widehat{s}_n .

Step 6: The metric. Let n_0 be large enough that our constructions work and let

$$g : (M_{\mathcal{A}}, P_{\mathcal{A}}) \rightarrow (Y, \partial\mathbb{T}_\varepsilon(P_{\mathcal{A}} \cap P) \cup \partial\mathbb{T}_{2\varepsilon}(P_{\mathcal{A}} \setminus P))$$

be the diffeomorphism given by the restriction of $\psi_{n_0}^{-1} \circ \widehat{s}_{n_0}$ to $M_{\mathcal{A}}$. Let m be the metric given by pulling back, by g , the metric on Y to a metric on $M_{\mathcal{A}}$. For all large enough n , define

$$f_n : (M_{\mathcal{A}}, P_{\mathcal{A}}) \rightarrow (Y_n, \partial\mathbb{T}_{\varepsilon, n}(P_{\mathcal{A}} \cap P) \cup \partial\mathbb{T}_{2\varepsilon, n}(P_{\mathcal{A}} \setminus P))$$

to be given by $\psi_n \circ g$. Notice that f_n is then 2-bilipschitz with respect to m on $M_{\mathcal{A}}$. Since f_n is isotopic to \widehat{s}_n on $\text{Fr}(\mathcal{A})$ (and pared homotopic on $(M_{\mathcal{A}}, P_{\mathcal{A}})$) and \widehat{s}_n extends to a pared homeomorphism between (M, P) and (C_n, P_n) , f_n also extends to a pared homeomorphism from (M, P) to (C_n, P_n) which is homotopic to \widehat{s}_n and hence in the homotopy class of ρ_n . The remaining properties hold by construction and are easily checked. Thus f_n is a model core map for $(\mathcal{A}, m, \varepsilon)$. This concludes the proof of Theorem 4.2. \square

5. Proof of the main theorem

We are now ready to prove the main theorem, which we state here in the pared setting:

Theorem 5.1. *Let (M, P) be a compact, orientable, pared 3-manifold with incompressible pared boundary. For each curve α in $\partial_{nw}(M, P)$, there exists $K = K(\alpha)$ such that*

$$\ell_{\sigma_{(M,P)}(\rho)}(\alpha) \leq K$$

for all $\rho \in AH_0(M, P)$.

Proof of Theorem 5.1: Let α be a curve in $\partial_{nw}(M, P)$, and let S be the component of $\partial_0(M, P)$ containing α . Fix $r < \mu_3/2$ such that any curve of length L_h intersecting $\mathbb{T}_r(\beta)$ is homotopic to a power of β . Applying Theorem 4.1 with $\varepsilon = r$ we see that, given $\rho \in AH_0(M, P)$, there is a model core (\mathcal{A}, m, r) , chosen from a finite list, with an associated model core map f .

Noting that α can be made disjoint from \mathcal{A} , choose a minimal length representative of α in $\partial_0(M, P) \cap M_{\mathcal{A}}$ and let α_f be its f -image in N_ρ . Its length is bounded by some L_0 depending on α and the finitely many possibilities for m . α_f is contained in $F = f(S)$ which, by the properties of a model map, is contained in the r -thick part of N_ρ .

Let $\pi : N_S \rightarrow N_\rho$ be the cover associated to $\pi_1(S)$ and let \widehat{F} be the homeomorphic lift of F to this cover. Let $\widehat{\alpha}$ be the lift of α_f to \widehat{F} . By definition, $\sigma_{(M,P)}^S(\rho)$ is the bottom ending invariant of N_S .

Let \mathcal{H} be the invariant system of horoballs for $\rho(\pi_1(M))$ which is the pre-image of the non-compact components of $N_{(0,r)}$ and let $\widehat{\mathcal{H}}$ denote the induced subcollection of invariant horoballs for $\rho(\pi_1(S))$. If $C = f(M)$ and $D = f(P)$, then (C, D) is a relative compact core for $N_\rho^{\mathcal{H}}$. Moreover, \widehat{F} is a level surface for $N_S^{\widehat{\mathcal{H}}}$.

Let \widehat{C} be the component of $\pi^{-1}(C)$ which contains \widehat{F} in its boundary. Then \widehat{C} lies below \widehat{F} . Suppose that β is a curve in S which intersects α essentially with length $\ell_\rho(\beta) \leq L_h$. By our choice of r , β^* lies in $N_S^{\widehat{\mathcal{H}}}$. We claim that either β^* lies above \widehat{F} or β^* intersects \widehat{C} .

Suppose β^* does not lie above \widehat{F} . If β^* does not intersect \widehat{C} it must be contained in a component U of $N_S \setminus \widehat{C}$ which lies below \widehat{F} . The component U shares a boundary component E with \widehat{C} . A homotopy of β^* to \widehat{F} , intersected with \widehat{C} and surgered, gives rise to an immersed essential annulus in \widehat{C} joining the lift of $f(\beta)$ to a curve in E , and hence to an immersed essential annulus in C joining $f(\beta)$ to a curve in ∂C . Therefore, by Theorem 2.1(4), β is homotopic into the window of M , so cannot have essential intersection with α . This contradiction implies that β^* must intersect \widehat{C} .

Any curve that does not lie above $\widehat{\alpha}$ also does not lie above \widehat{F} . We conclude that the geodesic representative of every curve in $\mathcal{C}(\widehat{\alpha}, L_h)$ intersects \widehat{C} .

Notice that if two (homotopically) distinct curves in S are in the same homotopy class in M , then there is an immersed essential annulus in M joining them, so both




FIGURE 5. In N_S , β_1 intersects α so β_1^* must meet \widehat{C} if it doesn't lie above \widehat{F} . But β_2 is in the window so β_2^* is not constrained.

curves are homotopic into the window (again by Theorem 2.1(4)). It follows that neither curve could intersect α . Therefore, any two distinct curves in $\mathcal{C}(\widehat{\alpha}, L_h)$ project to (homotopically) distinct curves in N_ρ .

So, the geodesic representatives of curves in $\mathcal{C}(\widehat{\alpha}, L_h)$ project to distinct curves which intersect C . Our choice of r guarantees that the geodesic representative of any curve in $\mathcal{C}(\widehat{\alpha}, L_h)$ cannot intersect $\mathbb{T}_r(f(\mathcal{A}))$ (if it did it would be homotopic into $\Sigma(M, P)$, again by Theorem 2.1(4), so would not intersect α essentially), so it must intersect $f(M_{\mathcal{A}})$. Since each component of $f(M_{\mathcal{A}})$ has uniformly bounded diameter, there exists a uniform bound on the number of geodesics of length at most L_h which can intersect $f(M_{\mathcal{A}})$. This bounds the number of elements of $\mathcal{C}(\widehat{\alpha}, L_h)$. Moreover, there is a lower bound on the length of the geodesic representative of any curve in $\mathcal{C}(\widehat{\alpha}, L_h)$, since if the closed geodesic is very short, its Margulis tube would have very large radius and would thus be forced to contain an entire component of $f(M_{\mathcal{A}})$, which is impossible. Finally, note that since F is in the r -thick part, so is \widehat{F} . This establishes all the hypotheses of Theorem 3.1, which therefore gives a uniform upper bound on $\ell_{\sigma_{M,S(\rho)}}(\alpha)$ and completes the proof. \square

The generalization of Corollary 1.2 to the pared setting is the following:

Corollary 5.2. *Let (M, P) be a compact, orientable, pared 3-manifold with incompressible pared boundary. If W is a component of $\partial_{nw}(M, P)$, then the image of $\sigma_{(M,P)}^W$ is bounded in $\mathcal{F}(W)$.*

Proof. If W is an annulus then $\mathcal{F}(W) = [0, \infty)$ and $\sigma_{(M,P)}^W$ is the map that records the length of the core of W in the skinning image. Hence the bound follows immediately from Theorem 5.1.

For any non-annular W we note that, since $\ell_{\sigma_{(M,P)}(\rho)}(\alpha) < \infty$ for each α in W , W must be contained in a geometrically finite component of $\sigma_{(M,P)}(\rho)$ (in fact this statement is explicitly demonstrated in the first step of the proof of Theorem 3.1). Thus, lifting to the cover associated to W we are able to define $\sigma_{(M,P)}^W(\rho) \in \mathcal{F}(W)$. Now for any $X \in \mathcal{F}(W)$ an upper bound on the lengths in X of a filling system of non-peripheral curves in W restricts X to a compact subset of $\mathcal{F}(W)$. Thus again Theorem 5.1 implies our statement. \square

References

- [1] I. Agol, “Tameness of hyperbolic 3-manifolds,” preprint (2004), ArXiv:math/0405568.
- [2] J.W. Anderson, “Intersections of topologically tame subgroups of Kleinian groups,” *J. Anal. Math.* **65**(1995), 77–94.
- [3] J.W. Anderson and R. Canary, “Algebraic limits of Kleinian groups which rearrange the pages of a book,” *Invent. Math.* **126** (1996), 205–214.
- [4] J.W. Anderson and R. Canary, “Cores of hyperbolic 3-manifolds and limits of Kleinian groups,” *Amer. J. Math.* **118** (1996), 745–779.
- [5] J.W. Anderson, R.D. Canary, M. Culler, and P.B. Shalen, “Free Kleinian groups and volumes of hyperbolic 3-manifolds,” *J. Diff. Geom.* **43**(1996), 738–782.
- [6] J.W. Anderson, R.D. Canary, and D. McCullough, “The topology of deformation spaces of Kleinian groups,” *Annals of Math.* **152** (2000), 693–741.
- [7] L. Bers, “Spaces of Kleinian groups,” in *Maryland Conference in Several Complex Variables I*. Lecture Notes in Math, Springer-Verlag, Vol. 155(1970), 9–34.
- [8] L. Bers, “On moduli of Kleinian groups,” *Russian Math. Surveys* **29** (1974), 88–102.
- [9] L. Bers and F. Gardiner, “Fricke spaces,” *Adv. Math.* **62**(1986), 249–284.
- [10] F. Bonahon, “Bouts des variétés hyperboliques de dimension 3,” *Annals of Math.* **124** (1986), 71–158.
- [11] J. Brock, “Boundaries of Teichmüller spaces and end-invariants for hyperbolic 3-manifolds,” *Duke Math. J.* **106**(2001), 527–552.
- [12] J.F. Brock, K.W. Bromberg, R.D. Canary and C. Lecuire, “Convergence and divergence of Kleinian surface groups,” *J. Topology* **8**(2015), 811–841,
- [13] J.F. Brock, K.W. Bromberg, R.D. Canary and Y.N. Minsky, “Local topology in deformation space of hyperbolic 3-manifolds,” *Geom. Top.* **15** (2011), 1169–1224.
- [14] J.F. Brock, K.W. Bromberg, R.D. Canary and Y.N. Minsky, “Convergence properties of end invariants,” *Geom. Top.* **17** (2013), 2877–2922.
- [15] J.F. Brock, R.D. Canary and Y.N. Minsky, “The Classification of Kleinian Surface Groups II: The Ending Lamination Conjecture,” *Annals of Math.*, **176** (2012), 1–149.
- [16] K. Bromberg, “The space of Kleinian punctured torus groups is not locally connected,” *Duke Math. J.* **156**(2011), 387–427.
- [17] D. Calegari and D. Gabai, “Shrinkwrapping and the taming of hyperbolic 3-manifolds,” *J. Amer. Math. Soc.* **19**(2006), 385–446.
- [18] R.D. Canary, “A covering theorem for hyperbolic 3-manifolds and its applications,” *Topology* **35**(1996), 751–778.
- [19] R.D. Canary, D.B.A. Epstein and P. Green, “Notes on notes of Thurston,” in *Analytical and Geometrical Aspects of Hyperbolic Spaces*, Cambridge University Press, 1987, 3–92.
- [20] R.D. Canary and D. McCullough, “Homotopy equivalences of 3-manifolds and deformation theory of Kleinian groups,” *Mem. Amer. Math. Soc.*, vol. **172**(2004), no. 812.
- [21] R.D. Canary and Y.N. Minsky “On limits of tame hyperbolic 3-manifolds,” *J. Diff. Geom.* **43**(1996), 1–41.
- [22] R.D. Canary, Y.N. Minsky and E.C. Taylor, “Spectral theory, Hausdorff dimension and the topology of hyperbolic 3-manifolds,” *J. Geom. Anal.* **9**(1999), 17–40.
- [23] D. Dumas, “Skinning maps are finite-to-one,” *Acta Math.* **215**(2015), 55–126.

- [24] D. Dumas and R. Kent, “Slicing, skinning, and grafting,” *Amer. J. Math.* **131**(2009), 1419–1429.
- [25] D. Dumas and R. Kent, “Bers slices are Zariski dense,” *J. Top.* **2**(2009), 373–379.
- [26] D.B.A. Epstein, A. Marden and V. Markovic, “Quasiconformal homeomorphisms and the convex hull boundary,” *Annals of Math.* **159** (2004), 305–336.
- [27] W. Jaco and P.B. Shalen, *Seifert fibered spaces in 3-manifolds*, Mem. Amer. Math. Soc. **21** (1979), no. 220.
- [28] K. Johannson, *Homotopy Equivalences of 3-manifolds with Boundary*, Lecture Notes in Mathematics, vol. 761, Springer-Verlag, 1979.
- [29] M. Kapovich, *Hyperbolic manifolds and discrete groups*, Progr. Math. **183**, Birkhäuser, 2001.
- [30] R. Kent, “Skinning maps,” *Duke Math. J.* **151**(2010), 279–336.
- [31] R. Kent and Y. Minsky, “Thick-skinned 3-manifolds,” *G.A.F.A.* **24**(2014), 1981–2001.
- [32] R. Kulkarni and P. Shalen, “On Ahlfors’ finiteness theorem,” *Adv. Math.* **111**(1991), 155–169.
- [33] A. Magid, “Deformation spaces of Kleinian surface groups are not locally connected,” *Geometry and Topology* **16**(2012) 1247–1320.
- [34] B. Maskit, “On Klein’s combination theorem. III,” in *Advances in the Theory of Riemann Surfaces* Ann. of Math. Studies, No. 66, Princeton Univ. Press, 1971, 297–316.
- [35] H.A. Masur and Y.N. Minsky, “Geometry of the complex of curves I: Hyperbolicity,” *Invent. Math.* **138** (1999), 103–149.
- [36] H.A. Masur and Y.N. Minsky, “Geometry of the complex of curves II: Hierarchical structure,” *Geom. Funct. Anal.* **10** (2000), 902–974.
- [37] D. McCullough, “Compact submanifolds of 3-manifolds with boundary,” *Quart. J. Math. Oxford* **37**(1986), 299–306.
- [38] C. McMullen, “Riemann surfaces and the geometrization of 3-manifolds,” *Bull. A.M.S.* **27**(1992), 207–216.
- [39] Y. Minsky, “Kleinian groups and the complex of curves,” *Geom. and Top.* **4** (2000), 117–148.
- [40] Y. Minsky, “The classification of Kleinian surface groups I: models and bounds,” *Annals of Math.* **171** (2010), 1–107.
- [41] J. Morgan, “On Thurston’s uniformization theorem for three-dimensional manifolds” in *The Smith Conjecture*, Pure Appl. Math. **112**, Academic Press, 1984, 37–125.
- [42] J.P. Otal, “Thurston’s hyperbolization of Haken manifolds,” in *Surveys in differential geometry, Vol. III* Int. Press, 1996, 77–194.
- [43] T. Soma, “Function groups in Kleinian groups,” *Math. Ann.* **292**(1992), 181–190.
- [44] W. P. Thurston, *The Geometry and Topology of Three-Manifolds*, Princeton University course notes, 1980, available at <http://www.msri.org/publications/books/gt3m/>
- [45] W.P. Thurston, “Hyperbolic geometry and 3-manifolds,” in *Low-dimensional topology (Bangor, 1979)*, Cambridge University Press, 1982, 9–25.
- [46] W.P. Thurston, “Hyperbolic structures on 3-manifolds, I: deformation of acylindrical manifolds,” *Annals of Math.* **124**(1986), 203–246.
- [47] W.P. Thurston, “Hyperbolic structures on 3-manifolds, III: Deformations of 3-manifolds with incompressible boundary,” preprint, available at: arxiv.org/abs/math/9801058

BROWN UNIVERSITY

UNIVERSITY OF UTAH

UNIVERSITY OF MICHIGAN

YALE UNIVERSITY