

ANNALS OF MATHEMATICS

**The classification of Kleinian surface
groups, I: models and bounds**

By YAIR MINSKY



SECOND SERIES, VOL. 171, NO. 1

January, 2010

ANMAAH

The classification of Kleinian surface groups, I: models and bounds

By YAIR MINSKY

Abstract

We give the first part of a proof of Thurston’s Ending Lamination conjecture. In this part we show how to construct from the end invariants of a Kleinian surface group a “Lipschitz model” for the thick part of the corresponding hyperbolic manifold. This enables us to describe the topological structure of the thick part, and to give *a priori* geometric bounds.

1. Introduction
2. End invariants
3. Hyperbolic constructions
4. Complexes of curves and arcs
5. Hierarchies
6. The coarse projection property
7. The projection bound theorem and consequences
8. The model manifold
9. Comparing meridian coefficients
10. The Lipschitz model map
11. Length bounds

References

Partially supported by NSF grant DMS-9971596. The bulk of this work was completed at SUNY at Stony Brook.

1. Introduction

This paper is the first in a three-part series addressing the question: to what extent is a hyperbolic 3-manifold determined by its asymptotic geometry? This question underlies the deformation theory of Kleinian groups, as pioneered by Ahlfors and Bers in the 60's and by Thurston and Bonahon in the 70's and early 80's. Their work provides us with a theory of *end invariants* assigned to the ends of a hyperbolic 3-manifold, and determined by their asymptotic geometric properties. Thurston [61] formulated this conjecture which has been a guiding question in the field:

ENDING LAMINATION CONJECTURE. *A hyperbolic 3-manifold with finitely generated fundamental group is uniquely determined by its topological type and its end invariants.*

When the manifold has finite volume its ends are either empty or cusps, the end invariants are empty, and the conjecture reduces to the well-known rigidity theorems of Mostow and Prasad [52], [54]. When the manifold has infinite volume but is “geometrically finite”, the end invariants are Riemann surfaces arising from the action of $\pi_1(N)$ on the Riemann sphere, and the conjecture follows from the work of Ahlfors-Bers [3], [8], [10] and Marden-Maskit [38], Maskit [40], Kra [36] and others.

The remaining cases are those where the manifold has a “geometrically infinite” end, for which the end invariant is a *lamination*. Here the discussion splits into two, depending on whether the boundary of the compact core is compressible or incompressible. If it is incompressible, then the work of Thurston [60] and Bonahon [13] gives a preliminary geometric and topological description of the end, and allows the ending laminations to be defined (see also Abikoff [1] for a survey). If the core boundary is compressible the situation is more difficult to analyze, and the corresponding question of *tameness of the end* was only recently resolved, by Agol [2] and Calegari-Gabai [21]; see Section 2.2.

In this paper we restrict ourselves to the incompressible boundary case. This case reduces, by restriction to boundary subgroups, to the case of (*marked*) *Kleinian surface groups*, with which we will be concerned for the remainder of the paper.

A marked Kleinian surface group is a discrete, faithful representation

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

where S is a compact surface, and ρ sends elements representing ∂S to parabolic elements. Each ρ determines a set of end invariants $v(\rho)$, which for each end give us Ahlfors-Bers Teichmüller data or an ending lamination, as appropriate.

In broadest outline, our plan for establishing the Ending Lamination Conjecture is to construct a “model manifold” M_ν , depending only on the invariants $\nu(\rho)$, together with a bilipschitz homeomorphism $f : M_\nu \rightarrow N_\rho$ (where $N_\rho = \mathbb{H}^3/\rho(\pi_1(S))$). Then if ρ_1 and ρ_2 are two Kleinian surface groups with the same end invariants ν , we would obtain a bilipschitz homeomorphism between N_{ρ_1} and N_{ρ_2} (in the right homotopy class), and an application of Sullivan’s rigidity theorem [58] would then imply that the map can be deformed to an isometry.

In this paper we will construct the model manifold together with a map satisfying some *Lipschitz* bounds (and some additional geometric properties, including detailed information about the thick-thin decomposition of N_ρ). In the second paper, with Brock and Canary [18], this map will be promoted to a bilipschitz homeomorphism. In a third paper, the case of general finitely-generated Kleinian groups will be treated, where the primary issue is the case of compressible-boundary compact core, and the reduction involves techniques such as Canary’s branched-cover argument.

Structure of the model. For simplicity, let us describe the model manifold M_ν when S is a closed surface, and when ν are invariants of a manifold N_ρ without parabolics, and without geometrically finite ends. (This avoids discussion of parabolic cusps and boundaries of the convex core.) In this case, M_ν is homeomorphic to $S \times \mathbb{R}$, and we fix such an identification.

Within M_ν there is a subset \mathcal{U} , which consists of open solid tori called “tubes” of the form $U = A \times J$, where A is an annulus in S and J is an interval in \mathbb{R} . No two components of \mathcal{U} are homotopic.

M_ν comes equipped with a piecewise-Riemannian metric, with respect to which each tube boundary ∂U is a Euclidean torus. The geometry of ∂U is described by a coefficient we call $\omega_M(U)$, which lies in the upper half-plane $\mathbb{H}^2 = \{z : \text{Im } z > 0\}$, thought of as the Teichmüller space of the torus. U itself is isometric to a tubular neighborhood of a hyperbolic geodesic, whose length goes to 0 as $|\omega_M| \rightarrow \infty$.

Let $\mathcal{U}[k]$ denote the union of components of \mathcal{U} with $|\omega_M| \geq k$, and let $M_\nu[k] = M_\nu \setminus \mathcal{U}[k]$. Then $M_\nu[0] = M_\nu \setminus \mathcal{U}$ is a union of “blocks”, which have a finite number of possible isometry types. This describes a sort of “thick-thin” decomposition of M_ν .

There is a corresponding decomposition of N_ρ , associating a Margulis tube $\mathbb{T}_{\epsilon_1}(\gamma)$ to each sufficiently short geodesic γ . Let $\mathbb{T}[k]$ denote the set of such Margulis tubes (if any) associated to the homotopy classes of components of $\mathcal{U}[k]$ under the homotopy equivalence between M_ν and N_ρ determined by ρ .

Let \hat{C}_{N_ρ} denote the “augmented convex core” of N_ρ (see §3.4), which in our simplified case is equal to N_ρ itself. Our main theorem asserts that M_ν can be

mapped to \widehat{C}_{N_ρ} by a Lipschitz map that respects the thick-thin decompositions of both.

LIPSCHITZ MODEL THEOREM. *Fix a compact oriented surface S . There exist $K, k > 0$ such that, if $\rho : \pi_1(S) \rightarrow \mathrm{PSL}_2(\mathbb{C})$ is a Kleinian surface group with end invariants $\nu(\rho)$, then there is a map*

$$f : M_\nu \rightarrow \widehat{C}_{N_\rho}$$

with the following properties:

- (1) f induces ρ on π_1 , is proper, and has degree 1.
- (2) f is K -Lipschitz on $M_\nu[k]$, with respect to the induced path metric.
- (3) f maps $\mathcal{U}[k]$ to $\mathbb{T}[k]$, and $M_\nu[k]$ to $N_\rho \setminus \mathbb{T}[k]$.
- (4) $f : \partial M_\nu \rightarrow \partial \widehat{C}_{N_\rho}$ is a K -bilipschitz homeomorphism of the boundaries.
- (5) For each tube U in \mathcal{U} with $|\omega_M(U)| < \infty$, $f|_U$ is λ -Lipschitz, where λ depends only on K and $|\omega_M(U)|$.

Remarks. The condition on the degree of f , after appropriate orientation conventions, amounts to the fact that f maps the ends of M_ν to the ends of N_ρ in the “correct order”. In our simplified case condition (4) is vacuous, as is the restriction $|\omega_M(U)| < \infty$ in (5).

The extended model map. In the general case, N_ρ may have parabolic cusps and \widehat{C}_{N_ρ} may not be all of N_ρ . The statement of the Lipschitz Model Theorem is unchanged, but the structure of M_ν is complicated in several ways: Some of the tubes of \mathcal{U} will be “parabolic”, meaning that their boundaries are annuli rather than tori, and the coefficients ω_M may take on the special value $i\infty$. M_ν will have a boundary, and the condition that f is proper is meant to include both senses: it is proper as a map of topological spaces, and it takes ∂M_ν to $\partial \widehat{C}_{N_\rho}$. The blocks of $M_\nu[0]$ will still have a finite number of topological types, but for a finite number of blocks adjacent to the boundary the isometry types will be unbounded, in a controlled way.

In Section 3.4 we will describe the geometry of the exterior of the augmented core, $E_N = N \setminus \widehat{C}_N$, in terms of a model E_ν that depends only on the end invariants of the geometrically finite ends of N , and is a variation of Epstein-Marden’s description of the exterior of the convex hull. M_ν and E_ν attach along their boundaries to yield an extended model manifold ME_ν .

N has a natural conformal boundary at infinity $\partial_\infty N$, and ME_ν has a conformally equivalent boundary $\partial_\infty ME_\nu$. The Lipschitz Model Theorem will then generalize to:

EXTENDED MODEL THEOREM. *The map f obtained in the Lipschitz Model Theorem extends to a proper degree 1 map*

$$f' : ME_v \rightarrow N$$

which restricts to a K -bilipschitz homeomorphism from E_v to E_N , and extends to a conformal map from $\partial_\infty ME_v$ to $\partial_\infty N$.

Length bounds. Note that Part (3) of the Lipschitz Model Theorem implies that for every component of $\mathcal{U}[k]$ there is in fact a corresponding Margulis tube in $\mathbb{T}[k]$, to which it maps properly. On the other hand the bounded isometry types of blocks (ignoring the boundary case) and the Lipschitz bound on f will imply that there is a lower bound $\epsilon > 0$ on the injectivity radius of N_ρ outside the image of \mathcal{U} . In other words, the structure of M_v determines the pattern of short geodesics and their Margulis tubes in N_ρ .

The following theorem makes this connection more precise. If γ is a homotopy class of curves in S , then let $\lambda_\rho(\gamma)$ denote the complex translation length of the corresponding conjugacy class $\rho(\gamma)$ in $\rho(\pi_1(S))$. Its real part $\ell_\rho(\gamma)$, which we may assume positive if $\rho(\gamma)$ is not parabolic, is the length of the geodesic representative of this homotopy class in N_ρ . If γ is homotopic to the core of some tube U in \mathcal{U} , then we define $\omega_M(\gamma) \equiv \omega_M(U)$.

SHORT CURVE THEOREM. *There exist $\bar{\epsilon} > 0$ and $c > 0$ depending only on S , and for each $\epsilon > 0$ there exists $K > 0$, such that the following holds: Let $\rho : \pi_1(S) \rightarrow \mathrm{PSL}_2(\mathbb{C})$ be a Kleinian surface group and γ a simple closed curve in S .*

- (1) *If $\ell_\rho(\gamma) < \bar{\epsilon}$, then γ is homotopic to a core of some component U in \mathcal{U} .*
- (2) *(Upper length bounds) If γ is homotopic to the core of a tube in \mathcal{U} , then*

$$|\omega_M(\gamma)| \geq K \implies \ell_\rho(\gamma) \leq \epsilon.$$

- (3) *(Lower length bounds) If γ is homotopic to the core of a tube in \mathcal{U} , then*

$$|\lambda_\rho(\gamma)| \geq \frac{c}{|\omega_M(\gamma)|}$$

and

$$\ell_\rho(\gamma) \geq \frac{c}{|\omega_M(\gamma)|^2}.$$

Part (2) of this theorem is actually a restatement of the main theorem of [48]; part (3) is the main new ingredient.

1.1. Outline of the proofs. In the following summary of the argument, we will continue making the assumptions that S is closed and N_ρ has no geometrically

finite ends or cusps. This greatly simplifies the logic of the discussion, while retaining all the essential elements of the proof. The reader is encouraged to continue making this assumption on a first reading of the proof itself (for example, §3.4 can be skipped, as can all mention of boundary blocks in §§8 and 10).

In this case, N_ρ has two ends, which we label with $+$ and $-$ (see §2.2 for the orientation conventions), and the end invariants $\nu(\rho)$ become two filling laminations ν_+ and ν_- on S (§2.1).

Quasiconvexity and the complex of curves. The central idea is to use the geometry of the *complex of curves* $\mathcal{C}(S)$ to obtain *a priori* bounds on lengths of curves in N_ρ . The vertices of $\mathcal{C}(S)$ are the essential homotopy classes of simple loops in S (see §4 for details), and we will study the sublevel sets

$$\mathcal{C}(\rho, L) = \{v \in \mathcal{C}_0(S) : \ell_\rho(v) \leq L\}$$

where $\ell_\rho(v)$ for a vertex $v \in \mathcal{C}_0(S)$ denotes the length of the corresponding closed geodesic in N_ρ .

In [49] we showed that $\mathcal{C}(\rho, L)$ is *quasiconvex* in the natural metric on $\mathcal{C}(S)$. The main tool for the proof of this is the “short curve projection” $\Pi_{\rho, L}$, which maps $\mathcal{C}(S)$ to $\mathcal{C}(\rho, L)$ by constructing for any vertex v in $\mathcal{C}(S)$ the set of pleated surfaces in N with v in their pleating locus, and finding the curves of length at most L in these surfaces. This map satisfies certain contraction properties which make it coarsely like a projection to a convex set, and this yields the quasiconvexity of $\mathcal{C}(\rho, L)$.

As a metric space $\mathcal{C}(S)$ is δ -hyperbolic, and the ending laminations $\nu_\pm(\rho)$ describe two points on its Gromov boundary $\partial\mathcal{C}(S)$ (see Masur-Minsky [43], Klarreich [34] and §4). In fact $\nu_\pm(\rho)$ are the accumulation points of $\mathcal{C}(\rho, L)$ on $\partial\mathcal{C}(S)$, and this together with quasiconvexity of $\mathcal{C}(\rho, L)$ appears to give a coarse type of control on $\mathcal{C}(\rho, L)$; in particular an infinite geodesic in $\mathcal{C}(S)$ joining ν_- to ν_+ must lie in a bounded neighborhood of $\mathcal{C}(\rho, L)$. However, since $\mathcal{C}(S)$ is locally infinite this estimate is not sufficient for us.

Subsurfaces and hierarchies. In Section 6 we generalize the quasiconvexity theorem of [49] to a *relative* result which incorporates the structure of subsurface complexes in $\mathcal{C}(S)$. In order to do this we recall in Sections 4 and 5 some of the structure of the subsurface projections and hierarchies in $\mathcal{C}(S)$ which were developed in Masur-Minsky [42]. To an essential subsurface $W \subset S$ we associate a “projection”

$$\pi_W : \mathcal{A}(S) \rightarrow \mathcal{A}(W)$$

(where $\mathcal{A}(W)$ is the arc complex of W , containing and quasi-isometric to $\mathcal{C}(W)$). This, roughly speaking, is a map that associates to a curve (or arc) system in S its

essential intersection with W . This map has properties analogous to the orthogonal projection of \mathbb{H}^3 to a horoball; see particularly Lemma 4.1.

A *hierarchy* is a way of enlarging a geodesic in $\mathcal{C}(S)$ to a system of geodesics in subsurface complexes $\mathcal{C}(W)$ that together produces families of markings of S . Such a hierarchy, called H_v , is constructed in Section 5 so that its base geodesic g connects v_- to v_+ (the construction is nearly the same as in Masur-Minsky [42], except for the need to treat infinite geodesics). The vertices which appear in H_v are all within distance 1 of g in $\mathcal{C}(S)$. The structure of H_v is strongly controlled by the maps π_W , as in Lemma 5.9.

Projections and length bounds. Once this structure is in place, we revisit the map $\Pi_{\rho,L}$. We prove, in Theorem 6.1, that the composition $\pi_Y \circ \Pi_{\rho,L}$ for a subsurface Y has contraction properties generalizing those shown in [49]. We then prove Theorem 7.1, which states in particular that

$$d_Y(v, \Pi_{\rho,L}(v))$$

is uniformly bounded for any subsurface Y and all vertices v appearing in H_v , provided v intersects Y essentially. Here $d_Y(x, y)$ denotes distance in $\mathcal{A}(Y)$ between $\pi_Y(x)$ and $\pi_Y(y)$.

This bound implies that v and the bounded-length curves $\Pi_{\rho,L}(v)$ are not too different in some appropriate combinatorial sense, and indeed we go on to apply this to obtain, in Lemma 7.9, an *a priori* upper bound on $\ell_\rho(v)$ for all vertices v that appear in H_v . Another crucial result we prove along the way is Lemma 7.7, which limits the ways in which pleated surfaces constructed from vertices of the hierarchy can penetrate Margulis tubes.

Model manifold construction. At this point we are ready to build the model manifold. In Section 8 we construct M_v out of the combinatorial data in H_v . The blocks of $M_v[0]$ are constructed from edges of geodesics in H_v associated to one-holed torus and 4-holed sphere subsurfaces, and glued together using the “subordinacy” relations in H_v . The structure of H_v is also used to embed $M_v[0]$ in $S \times \mathbb{R}$ (after which we identify it with its embedded image), and the tubes \mathcal{U} are the solid-torus components of $S \times \mathbb{R} \setminus M_v[0]$, and are in one-to-one correspondence with the vertices of H_v . In Section 8.3 we introduce the meridian coefficients $\omega_M(v)$, which encode for each vertex v the geometry of the associated tube boundary. The metric of M_v is described in this section too.

In Section 9 we define alternative meridian coefficients ω_H and ω_v , which are computed, respectively, directly from the data of H_v and directly from v itself. It is useful later in the proof to compare all three of these and in Theorem 9.1 we show that they are essentially equivalent. The proof requires a somewhat careful

analysis of the geometry of the model, and some counting arguments using the structure of the hierarchy.

Lipschitz bounds. In Section 10 we finally build the Lipschitz map from M_ν to N_ρ , establishing the Lipschitz Model Theorem. This is done in several steps, starting with the “gluing boundaries” of blocks, where the *a priori* bound on vertex lengths from Lemma 7.9 provides the Lipschitz control. Extension to the “middle surfaces” of blocks (Step 2) requires another application of Thurston’s Uniform Injectivity Theorem, via Lemma 3.1. Control of the extension to the rest of the blocks requires a reprise of the “figure-8 argument” from [47] to bound homotopies between Lipschitz maps of surfaces (Step 4). The map can be extended to tubes, and the last subtle point comes in Step 7, where we need Lemma 10.1 to relate large meridian coefficients to short curves (this is the point where we apply the results of Section 9, as well as the main theorem of [48]).

The proof of the Short Curve Theorem, carried out in Section 11, is now a simple consequence of the Lipschitz Model Theorem together with the properties of Margulis tubes. Roughly speaking, an upper bound on $|\omega_M(\gamma)|$ gives an upper bound on the meridian disk of $U(\gamma)$, and hence a lower bound on the length of its geodesic core.

Preliminaries. Sections 2 through 5 provide some background and notation before the proof itself starts in Section 6. Section 2 introduces compact cores, ends and laminations. Section 3 introduces pleated surfaces, Margulis tubes and collars in surfaces, and the augmented convex core. There is only a little bit of new material here: the augmented convex core and particularly the geometric structure of its exterior, via Lemma 3.4, and a slightly technical variation (Lemma 3.3) on the standard collar of a short geodesic or cusp in a hyperbolic surface. Sections 4 and 5 introduce the complexes of curves and arcs, and hierarchies. Most of this is review of material from [42], with certain generalizations to the case involving infinite geodesics. In particular the existence of an infinite hierarchy connecting the invariants ν_+ and ν_- (and in fact any pair of generalized markings) is shown in Lemma 5.13, and the existence of a “resolution” of a hierarchy, which is something like a sequence of markings separated by elementary moves sweeping through all the data in the hierarchy, is shown in Lemmas 5.7 and 5.8.

An informal but extensive summary of the argument, focusing on the case where S is a five-holed sphere, can be found in the lecture notes [35].

Bilipschitz control of the model map. In [18] we will show that the map $f : M_\nu \rightarrow \hat{C}_N$ can be made a bilipschitz homeomorphism, and thereby establish the Ending Lamination Conjecture. The argument begins with a decomposition of M_ν along quasi-horizontal slices into pieces of bounded size. The surface embedding machinery of Anderson-Canary [6] is used to show that the slices in the boundaries

of the pieces can be deformed to embeddings in a uniform way. The resulting map preserves a certain topological partial order among these slices – this is established using an argument by contradiction and passage to geometric limits – and this can be used to make the map an orientation-preserving embedding on each piece separately. Uniform bilipschitz bounds on these embeddings are obtained again by contradiction and geometric limit, and a global bilipschitz bound follows.

We note that the structure of the model manifold provides new information even in the geometrically finite case, for which the Ending Lamination Conjecture itself reduces to the quasiconformal deformation theory of Ahlfors and Bers. In particular it describes the thick-thin decomposition of the manifold, its volume and other geometric properties explicitly in terms of the Riemann surfaces at infinity.

In the rather long period between the initial and the final versions of this article, a number of other proofs (including the general compressible case) have appeared, due to Bowditch [16], Rees [55] and Soma [57]. Rees' approach works directly in Teichmüller space without the intermediate use of complexes of curves. Bowditch and Soma do use complexes of curves but simplify various aspects of the proof.

Acknowledgements. The author is grateful to Howard Masur, Dick Canary and Jeff Brock, without whose collaborations and encouragement this project would never have been completed. He is also grateful to the referee for detailed reading and suggestions, and for pointing out an error in the proof of Lemma 3.4, and to the editors for their patience.

2. End invariants

Before we discuss the end invariants of a hyperbolic 3-manifold, we set some notation which will be used throughout this paper. Let $S_{g,b}$ denote an oriented, connected and compact surface of genus g with b boundary components. Given a surface R of this type, an *essential subsurface* $Y \subseteq R$ is a compact, connected subsurface all of whose boundary components are homotopically nontrivial, and so that Y is not homotopic into a boundary component of R . (Unless otherwise mentioned we assume throughout the paper that any subsurface is essential.)

Define the *complexity* of a surface to be

$$\xi(S_{g,b}) \equiv 3g + b.$$

Note that for an essential subsurface Y of R , $\xi(Y) < \xi(R)$ unless Y is homeomorphic to R .

A *hyperbolic structure on* $\text{int}(R)$ will be a hyperbolic metric whose completion is a hyperbolic surface with cusps and/or geodesic boundary components, the latter of which we may identify with components of ∂R . If the completion is

compact then its boundary is identified with all of ∂R , and we call this a *hyperbolic structure on R* .

The Teichmüller space $\mathcal{T}(R)$ is the space of (marked) hyperbolic structures on $\text{int}(R)$ which are complete – that is, all ends are cusps. Alternatively $\mathcal{T}(R)$ is the space of marked conformal structures for which the ends are punctures.

2.1. Geodesic laminations. We will assume that the reader is familiar with the basics of geodesic laminations on hyperbolic surfaces; see Casson-Bleiler [25] or the recently written Bonahon [14] for an introduction to this subject. Given a surface R with a complete hyperbolic structure on $\text{int}(R)$ (in the above sense, with all ends cusps), we will be using the following spaces:

$\mathcal{GL}(R)$ is the space of geodesic laminations on R . $\mathcal{ML}(R)$ is the space of *transversely measured* laminations on R with compact support (if $\lambda \in \mathcal{ML}(R)$ its *support*, $|\lambda|$, lies in $\mathcal{GL}(R)$). $\mathcal{GL}(R)$ is usually topologized using the Hausdorff topology on closed subsets of R , whereas $\mathcal{ML}(R)$ admits a topology (due to Thurston) coming from the weak-* topology of the measures induced on transverse arcs.

$\mathcal{UML}(R)$ is the quotient space of $\mathcal{ML}(R)$ obtained by forgetting the measures (the “unmeasured laminations”). Its topology is non-Hausdorff and it is rarely used as an object on its own. Note that $\mathcal{UML}(R)$ is set-theoretically contained in $\mathcal{GL}(R)$, but the topologies are different.

It is part of the basic structure theory of laminations that every element of $\mathcal{UML}(R)$ decomposes into a finite union of disjoint connected components, each of which is a minimal lamination.

Let $\mathcal{EL}(R)$ denote the image in $\mathcal{UML}(R)$ of the *filling laminations* in $\mathcal{ML}(R)$. $\mu \in \mathcal{ML}(R)$ is called *filling* if it has transverse intersection with any $\mu' \in \mathcal{ML}(S)$, unless μ and μ' have the same support. An equivalent condition is that μ is both minimal and maximal as a measured lamination, and another is that the complementary components of μ are ideal polygons or once-punctured ideal polygons. (see [60]). In the topology inherited from $\mathcal{UML}(S)$, $\mathcal{EL}(S)$ is Hausdorff (see [34] for a proof).

Remarks. 1. All of these spaces do not really depend on the hyperbolic structure of $\text{int}(R)$; that is, the spaces obtained from two different choices of structure are canonically homeomorphic. 2. The laminations in $\mathcal{EL}(R)$ are exactly those that appear as ending laminations of Kleinian surface groups $\rho \in \mathcal{D}(R)$ without accidental parabolics – this is the reason for the notation \mathcal{EL} .

2.2. Cores and ends. Let N be an oriented complete hyperbolic 3-manifold with finitely generated fundamental group. N may be expressed as the quotient \mathbb{H}^3 / Γ by a Kleinian group $\Gamma \cong \pi_1(N)$. Let Λ be the limit set of Γ in the Riemann

sphere $\widehat{\mathbb{C}}$, and let $\Omega = \widehat{\mathbb{C}} \setminus \Lambda$ be the domain of discontinuity of Γ . (See e.g. [41], [22], [7] for background.) Then

$$\bar{N} \equiv (\mathbb{H}^3 \cup \Omega) / \Gamma$$

is a 3-manifold with boundary $\partial\bar{N} = \Omega / \Gamma$ and interior N . This boundary, which is also denoted $\partial_\infty N$, inherits a conformal structure from Ω .

Let C_N denote the *convex core* of N , which is the quotient by Γ of the convex hull of Λ in \mathbb{H}^3 . C_N is homeomorphic to \bar{N} by a map homotopic to the inclusion (except when Γ is Fuchsian or elementary and C_N has dimension 2 or less; we will assume this is not the case). Thurston showed that ∂C_N is a union of hyperbolic surfaces in the induced path metric from N , and by a theorem of Sullivan, if ∂C_N is incompressible these metrics are within universal bilipschitz distortion from the Poincaré metric on $\partial\bar{N}$ (see Epstein-Marden [26]). Ahlfors' finiteness theorem [4] states that $\partial\bar{N}$ and ∂C_N have finite hyperbolic area.

Let Q denote the union of (open) ϵ_0 -Margulis tubes of cusps of N (see §3.2.2 for a discussion of Margulis tubes), and let $N_0 = N \setminus Q$. By Scott's compact core theorem [56] there is a compact 3-manifold $K \subset N$ whose inclusion is a homotopy equivalence. The relative core theorem of McCullough [44] and Kulkarni-Shalen [37] tells us that K can be chosen in N_0 so that ∂K meets the boundary of each rank-1 cusp of Q in an essential annulus, and contains the entire torus boundary of each rank-2 cusp. Let $P = \partial K \cap \partial Q$. Note that no two components of P can be homotopic, since no two Margulis tubes can have homotopic core curves. The topological ends of N_0 are in one to one correspondence with the components of $N_0 \setminus K$ (which are *neighborhoods* of the ends), and hence with the components of $\partial K \setminus P$ (see Bonahon [13]). If R is the closure of a component of $\partial K \setminus P$, let E_R be the component of $N_0 \setminus K$ adjacent to R . We say that the end *faces* R .

Geometrically finite ends. An end of N_0 is *geometrically finite* if its associated neighborhood E_R meets C_N in a bounded subset. This implies that the boundary of E_R in \bar{N} consists of R , some annuli in ∂Q , and a component X of $\partial\bar{N}$ which is homotopic to R . Indeed, K may be chosen so that $E_R \cong \text{int}(R) \times (0, 1)$. The conformal structure of X gives rise to a point in $\mathcal{T}(R)$, and this is the end invariant associated to R , which we name ν_R .

Geometrically infinite ends. An end of N_0 is *geometrically infinite* if its associated E_R intersects C_N in an unbounded set. Ahlfors' finiteness theorem [4] implies that $\partial C_N \cap N_0$ is compact, and hence cannot separate E_R into two unbounded sets. Thus it follows that in fact there is a (possibly smaller) neighborhood of the end which is contained in C_N .

In order to describe the end invariant for a geometrically infinite end, we must consider the following definition from Thurston [60]:

Definition 2.1. An end E_R of N_0 , when R is incompressible, is called *simply degenerate* if there exists a sequence of essential simple closed curves α_i in R whose geodesic representatives α_i^* exit the end.

Here “exiting the end” means that the geodesics are eventually contained in E_R minus any bounded subset. Note that a geometrically finite end cannot be simply degenerate, since all closed geodesics are contained in the convex hull.

Thurston established this theorem (see also Canary [23]):

THEOREM 2.2 (Thurston [60]). *Let e be an end of N_0 facing $R \subset \partial K$, and suppose that R is incompressible in K . If e is simply degenerate, then there exists a unique lamination $\nu_R \in \mathcal{ML}(R)$ such that for any sequence of simple closed curves α_i in R ,*

$$\alpha_i \rightarrow \nu_e \iff \alpha_i^* \text{ exit the end } e.$$

A sequence $\alpha_i \rightarrow \nu_e$ can be chosen so that the lengths $\ell_N(\alpha_i^*) \leq L_0$, where L_0 depends only on S .

Furthermore, $\nu_e \in \mathcal{EL}(R)$ – that is, it fills R .

Thurston also proved that an incompressible simply degenerate end is *topologically tame*, meaning that it has a neighborhood homeomorphic to $R \times (0, \infty)$, and that manifolds obtained as limits of quasifuchsian manifolds have ends that are geometrically finite or simply degenerate. Bonahon completed the picture, in the incompressible boundary case, with his “tameness theorem”,

THEOREM 2.3 ([13]). *Suppose that each component of $\partial K \setminus P$ is incompressible in K . Then the ends of N_0 are either geometrically finite or simply degenerate.*

An end that is geometrically finite or simply degenerate is known as *geometrically tame*.

The case of an end facing a compressible boundary component is considerably harder to understand. Canary [23] showed that the analogue of Bonahon’s theorem holds for such ends (with a suitably strengthened notion of simple degeneracy) if the end is known to be topologically tame. Marden had conjectured in [39] that all hyperbolic 3-manifolds with finitely generated fundamental groups have topologically tame ends. When this article was originally written the question of tameness was still open, but it has since been resolved by Agol [2] and Calegari-Gabai [21].

Ends for Kleinian surface groups. From now on restrict to the case of a Kleinian surface group $\rho : \pi_1(S) \rightarrow \mathrm{PSL}_2(\mathbb{C})$, and let $N = N_\rho$ be the quotient manifold $\mathbb{H}^3 / \rho(\pi_1(S))$. Let Q and K be defined as above.

Let $Q_0 \subseteq Q$ be the set of cusp tubes associated to the boundary of S (if any), and let P_0 be the union of annuli in $\partial Q_0 \cap K$. Then K is homeomorphic to $S \times [-1, 1]$ with P_0 identified with $\partial S \times [-1, 1]$.

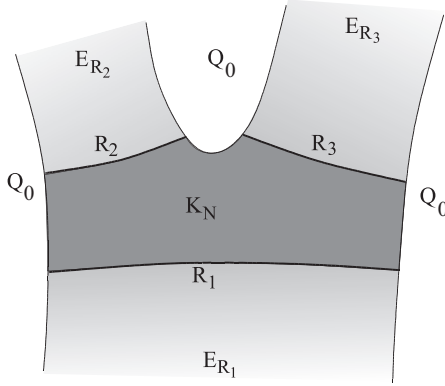


Figure 1. A schematic of the division of N into cusps, compact core and ends.

Divide the remaining components of Q into Q_+ and Q_- according to whether they meet K on $S \times \{1\}$ or $S \times \{-1\}$. Let P denote the union of annuli $\partial Q \cap K$, divided similarly into P_0 , P_+ and P_- .

The closure R of a component of $\partial K \setminus P$ then has an associated end invariant ν_R as above, in either $\mathcal{T}(R)$ or $\mathcal{EL}(R)$. We group these according to whether they come from the “+” or “-” side. More specifically, let p_+ denote the set of core curves of P_+ . Let R_+^L denote the union of (closures of) components of $S \times \{1\} \setminus P_+$, for which the invariant ν_R is a lamination in $\mathcal{EL}(R)$. Let R_+^T denote the remaining components. Define ν_+ as a pair (ν_+^L, ν_+^T) , where $\nu_+^L \in \mathcal{ULM}(S)$ is the union of p_+ and the laminations of components of R_+^L , and ν_+^T is the set of Teichmüller end invariants of components of R_+^T , which can be seen as an element of $\mathcal{T}(R_+^T)$. We also have to allow either ν_+^T or ν_+^L to be empty, in the case that there are no lamination or Teichmüller invariants, respectively. Define ν_- in the analogous way.

We remark that, because distinct cusps in a hyperbolic manifold cannot have homotopic curves, the parabolics p_+ and p_- have no elements in common. This implies, when p_+ or p_- are nonempty, that ν_+^L and ν_-^L have no infinite-leaf components in common either. In fact, this is true if $p_+ = p_- = \emptyset$ as well, as shown by Thurston [60].

It may be helpful to discuss special cases for clarity: If ρ is quasifuchsian, then p_\pm are both empty (there are no parabolics aside from those associated to ∂S) and R_+^T and R_-^T are both copies of S . Thus $\nu_\pm = (\emptyset, \nu_\pm^T)$, where $\nu_\pm^T \in \mathcal{T}(S)$ are the classical Ahlfors-Bers parameters for ρ . In this case of course the Ending Lamination Conjecture is well established, but our construction will still provide new information about the geometry of N_ρ .

If ρ has no parabolics aside from ∂S and no geometrically finite ends, then $\nu_\pm = (\nu_\pm^L, \emptyset)$ with $\nu_\pm^L \in \mathcal{EL}(S)$. This is called the *doubly degenerate case*.

In Section 7.1, we will replace v_{\pm} with a pair of *generalized markings* which will be used in the rest of the construction.

Orientation convention. In order to have a natural ordering of the ends, let us fix an orientation convention throughout the paper. An orientation on a manifold X will induce an orientation on ∂X by the convention that, if e is a baseframe for $T_p(\partial X)$ and e' is an inward-pointing vector in $T_p X$, then e is positively oriented if and only if (e, e') is. We orient hyperbolic space \mathbb{H}^3 so that it induces the standard orientation on its boundary $\widehat{\mathbb{C}}$, and note that the induced orientation on the upper half-plane \mathbb{H}^2 in turn induces the standard orientation on its boundary $\widehat{\mathbb{R}}$.

The manifold N_{ρ} inherits an orientation from \mathbb{H}^3 , and the compact core K inherits one as well. We are given a fixed orientation on S , and this determines (up to proper isotopy) an identification of K with $S \times [-1, 1]$ by the condition that the induced boundary orientation on $S \times \{-1\}$ agrees with the given orientation on S . Thus we know which end is up.

3. Hyperbolic constructions

3.1. Pleated surfaces. A *pleated surface* is a map $f : \text{int}(S) \rightarrow N$ together with a hyperbolic structure on $\text{int}(S)$, written σ_f and called the *induced metric*, and a σ_f -geodesic lamination λ on S , so that the following holds: f is length-preserving on paths, maps leaves of λ to geodesics, and is totally geodesic on the complement of λ . Pleated surfaces were introduced by Thurston [60]. See Canary-Epstein-Green [22] for more details. We include the case that the hyperbolic structure on $\text{int}(S)$ is incomplete and the boundary of the completed surface is mapped geodesically. In almost every case in this paper, however, boundary components of S are mapped to cusps, so that σ_f will be a complete metric with all ends cusps, and leaves of λ will go straight out this cusp in the σ_f metric.

As in [48] we extend the definition slightly to include *noded pleated surfaces*: Suppose S' is an essential subsurface of S whose complement is a disjoint union of open collar neighborhoods of simple curves Δ . Let $[f]$ be a homotopy class of maps from S to N , such that f takes Δ to cusps of N . We say that $g : S' \rightarrow N$ is a *noded pleated surface* in the class $[f]$ if g is pleated with respect to a hyperbolic metric on S' (in which the ends are cusps), and g is homotopic to the restriction to S' of an element of $[f]$. We say that g is “noded on Δ ”. By convention, we represent the “pleating locus” λ of f as a lamination in $\mathcal{GL}(S)$ that contains Δ as components, and leaves of λ that spiral onto Δ will be taken to leaves that go out the corresponding cusp.

Now fixing a Kleinian surface group $\rho : \pi_1(S) \rightarrow \text{PSL}_2(\mathbb{C})$, we have a natural homotopy class of maps $S \rightarrow N_{\rho}$ inducing ρ on fundamental groups. If $\lambda \in \mathcal{GL}(S)$,




Figure 2. A triangulation of a pair of pants and the lamination obtained by spinning it – moving the vertices leftward along the boundaries and taking a limit.

define

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

to be the set of pleated surfaces $g : \text{int}(S) \rightarrow N_\rho$ in the homotopy class of ρ , which map the leaves of λ to geodesics (if some closed leaves of λ correspond to parabolic elements in $\pi_1(N_\rho)$, then we allow the surface to be noded on these curves).

When v is a pants decomposition there is only a finite number of laminations containing the curves of v , so $\mathbf{pleat}_\rho(v)$ is finite up to isotopy equivalence (the natural equivalence of precomposition by homeomorphisms homotopic to the identity). Each of these laminations consists of the curves of v , together with finitely many arcs that spiral around them. We call $g \in \mathbf{pleat}_\rho(v)$ *leftward* if this spiraling is to the left for each leaf. This determines a unique isotopy equivalence class. Figure 2 shows a leftward lamination, restricted to a pair of pants, and also indicates how it can be obtained from a triangulation by “spinning”.

In particular we find that $\mathbf{pleat}_\rho(v)$ is nonempty for any lamination v consisting of simple closed curves. The same is true for any *finite-leaved* laminations, i.e. one in which we also allow infinite arcs that either spiral on closed curves, or go out the cusps of S . See [22].

Halfway surfaces. Let v and v' be pants decompositions that agree except on one curve: Thus $v = v_0 \cup w$ and $v' = v_0 \cup w'$ where v_0 divides S into three-holed spheres and one component W of type $\xi(W) = 4$. Assume also that v and v' intersect a minimal number of times in W (once if $W = S_{1,1}$ and twice if $W = S_{0,4}$). We say that v and v' differ by an *elementary move on pants decompositions*.

As in [49], define a lamination $\lambda_{v,v'}$ as follows: $\lambda_{v,v'}$ contains v_0 as a sublamination; in each 3-holed sphere complementary component of v_0 it is the same as the lamination in Figure 2; and in W , $\lambda_{v,v'}$ is given by the following diagram in the \mathbb{R}^2 -cover of W : The lift of ∂W is the lattice of small circles at integer points. W is




Figure 3. The lamination $\tilde{\lambda}_{v,v'}$ in the planar cover of W .

the quotient under the action of \mathbb{Z}^2 if $W = S_{1,1}$, and under the group generated by $(2\mathbb{Z})^2$ and $-I$ if $W = S_{0,4}$. The standard generators of \mathbb{Z}^2 in the figure correspond to the curves w and w' .

There is then a unique isotopy equivalence class $g_{v,v'} \in \mathbf{pleat}_\rho(\lambda_{v,v'})$, which we call the *halfway surface* associated to the pair (v, v') .

Let $g_v \in \mathbf{pleat}_\rho(v)$ and $g_{v'} \in \mathbf{pleat}_\rho(v')$ be leftward pleated surfaces. Figure 4 indicates the laminations λ_v and $\lambda_{v'}$ associated to these surfaces, restricted to W and lifted to the planar cover as above. In the complement of W , all three laminations agree.

Consider the set of leaves labeled \tilde{l} in $\tilde{\lambda}_v$ in Figure 4. These are disjoint from the lifts of w and project to either one or two leaves l in W (depending on whether W is a 1-holed torus or 4-holed sphere) that spiral on its boundary. The leaves l are common to both λ_v and $\lambda_{v,v'}$. Furthermore it is easy to see that any closed curve α that has an essential intersection with γ_w must also have an essential intersection with l . If α is contained in W this is evident from the fact that l cuts $\text{int}(W)$ into an annulus with core γ_w ; if not then α must cross ∂W , and we use the fact that l spirals around every component of W . The corresponding facts hold for the leaves \tilde{l}' which project to l' in $\lambda_{v'}$.

The following lemma, a restatement of Lemma 4.2 in [49], serves to control the geometry of a halfway surface.

LEMMA 3.1. *Let $v = v_0 \cup w$ and $v' = v_0 \cup w'$ be pants decompositions that differ by an elementary move. Let $\sigma_v, \sigma_{v'}$ and $\sigma_{v,v'}$ be the metrics induced on S by the pleated surfaces $g_v, g_{v'}$ and $g_{v,v'}$, respectively. Then for a constant C depending only on the topology of S ,*

$$\ell_{\sigma_v}(w) \leq \ell_{\sigma_{v,v'}}(w) \leq \ell_{\sigma_v}(w) + C,$$

and

$$\ell_{\sigma_{v'}}(w') \leq \ell_{\sigma_{v,v'}}(w') \leq \ell_{\sigma_{v'}}(w') + C.$$




Figure 4. The laminations associated to g_v and $g_{v'}$. The representative of w lifts to a horizontal line in the left figure, and the representative of w' lifts to a vertical line in the right figure.

Note that $\ell_{\sigma_v}(w) = \ell_\rho(w)$ since the curve representing w is mapped to its geodesic representative by g_v ; and similarly for $\ell_{\sigma_{v'}}(w')$. Hence the left-hand inequality on each line is immediate. The right-hand inequalities follow from an application of Thurston's Efficiency of Pleated Surfaces [62].

In [49] this result is proved and used without any assumption about the lengths $\ell_\rho(v)$ or $\ell_\rho(v')$. In this paper we will only use it in the case where these lengths are already bounded both above and below.

3.2. Tubes and constants.

3.2.1. Hyperbolic tubes. A *hyperbolic tube* is the quotient of an r -neighborhood of a geodesic in \mathbb{H}^3 by a translation or screw motion.

Given $\lambda \in \mathbb{C}$ with $\operatorname{Re} \lambda > 0$, and $r > 0$, we define $\mathbb{T}(\lambda, r)$ to be the quotient of the open r -neighborhood of the vertical line above $0 \in \mathbb{C}$ in the upper half-space model of \mathbb{H}^3 by the loxodromic $\gamma : z \mapsto e^\lambda z$. Let $\mathbb{T}(\lambda, \infty)$ denote the quotient of $\mathbb{H}^3 \cup \mathbb{C} \setminus \{0\}$ by γ . Any hyperbolic tube is isometric to some $\mathbb{T}(\lambda, r)$, but we note that the imaginary part of λ is, so far, only determined modulo 2π .

Marked boundaries. We discuss now how to describe the geometry of a hyperbolic tube \mathbb{T} in terms of the structure of its boundary torus.

If T is an oriented Euclidean torus, a *marking* of it is an ordered pair (α, β) of homotopy classes of unoriented simple closed curves with intersection number 1. There is a unique $t > 0$ and $\omega \in \mathbb{H}^2$ such that T can be identified with $\mathbb{C}/t(\mathbb{Z} + \omega\mathbb{Z})$ by an orientation-preserving isometry, so that the images of \mathbb{R} and $\omega\mathbb{R}$ are in the classes α and β , respectively. The parameter ω describes the conformal structure of T as a point in the Teichmüller space $\mathcal{T}(T) \cong \mathbb{H}^2$.

The boundary torus of a hyperbolic tube \mathbb{T} inherits a Euclidean metric and an orientation from \mathbb{T} (see §2.2 for orientation conventions), and it admits an almost

uniquely defined marking: Let μ denote the homotopy class of a meridian of the torus (i.e. the boundary of an essential disk in \mathbb{T}) and let α denote a homotopy class in $\partial\mathbb{T}$ of simple curves homotopic to the core curve of \mathbb{T} . While μ is unique, α is only defined up to multiples of μ . Fixing such a choice (α, μ) , we obtain *boundary parameters* (ω, t) as above.

Given $\lambda \in \mathbb{C}$ we can determine a marking (α, μ) of $\partial\mathbb{T}(\lambda, \infty)$, in a way that relates ambiguities in the choice of α to the freedom of adding $2\pi i$ to λ . The boundary at infinity $\partial\mathbb{T}(\lambda, \infty)$ is the quotient of $\mathbb{C} \setminus \{0\}$ by $z \mapsto e^\lambda z$. Using e^z as the universal covering $\mathbb{C} \rightarrow \mathbb{C} \setminus \{0\}$, we obtain $\partial\mathbb{T}(\lambda, \infty)$ as the quotient $\mathbb{C}/(\lambda\mathbb{Z} + 2\pi i\mathbb{Z})$. The line $i\mathbb{R}$ maps to the meridian, and we let α be the image of $\lambda\mathbb{R}$. Thus adding $2\pi i$ to λ corresponds to twisting α once around μ . We note that the corresponding boundary parameters for $\partial\mathbb{T}(\lambda, \infty)$ (with the standard metric and orientation inherited from \mathbb{C}) are $\omega = 2\pi i/\lambda$ and $t = |\lambda|$.

The marking (α, μ) at $\partial\mathbb{T}(\lambda, \infty)$ determines a unique marking via orthogonal projection, denoted also (α, μ) , on $\partial\mathbb{T}(\lambda, r)$, and for this marked torus we also have boundary data, (ω_r, t_r) .

The next lemma allows us to recover the length-radius parameters (λ, r) from (ω_r, t_r) , and indeed to construct a tube realizing any desired boundary data.

LEMMA 3.2. *Given $\omega \in \mathbb{H}^2$ and $t > 0$, there is a unique pair (λ, r) with $\operatorname{Re} \lambda > 0$, $r > 0$ such that $\mathbb{T}(\lambda, r)$ has boundary data (ω, t) in the marking determined by λ .*

Proof. Let X_r be the boundary of the r -neighborhood of the vertical geodesic over $0 \in \mathbb{C}$. This is a cone in the upper half-space, and in the induced metric is a Euclidean cylinder with circumference $2\pi \sinh r$. We may identify the universal cover \tilde{X}_r isometrically with \mathbb{C} , with deck translation $z \mapsto z + 2\pi i \sinh r$. Let $X_\infty = \mathbb{C} \setminus \{0\}$ with the Euclidean metric of circumference 2π , and identify the universal cover \tilde{X}_∞ isometrically with \mathbb{C} , with deck translation $z \mapsto z + 2\pi i$.

Let $\Pi_r : X_\infty \rightarrow X_r$ be the orthogonal projection map. It lifts to a map $\Phi_r : \tilde{X}_\infty \rightarrow \tilde{X}_r$ which, in our coordinates, can be written

$$(3.1) \quad \Phi_r(x + iy) = x \cosh r + iy \sinh r$$

(up to translation).

For our torus with parameters (ω, t) the meridian length is $t|\omega|$. Hence in order to build the right hyperbolic tube we must choose r so that

$$(3.2) \quad 2\pi \sinh r = t|\omega|.$$

It remains to determine λ . In \tilde{X}_∞ , we see λ as the translation that, together with the meridian $2\pi i$, produces the quotient torus $\partial\mathbb{T}(\lambda, \infty)$, with marking. The map Φ_r takes λ to a translation λ' that, together with the meridian $2\pi i \sinh r$, yields

the marking for $\partial\mathbb{T}(\lambda, r)$. Our boundary parameter ω is given by the ratio

$$\omega = \frac{2\pi i \sinh r}{\lambda'}$$

so that $\lambda' = 2\pi i \sinh r / \omega$, and hence we can define

$$(3.3) \quad \lambda = \Phi_r^{-1} \left(\frac{2\pi i \sinh r}{\omega} \right)$$

$$(3.4) \quad = h_r \left(\frac{2\pi i}{\omega} \right)$$

where we define

$$h_r(z) \equiv \Phi_r^{-1}(z \sinh r) = \operatorname{Re} z \tanh r + i \operatorname{Im} z.$$

These parameters therefore yield the desired torus. Notice that for large r , λ and $2\pi i / \omega$ are nearly the same. \square

Parabolic tubes. The parabolic transformation $z \mapsto z + t$, acting on the region of height > 1 in the upper half-space model of \mathbb{H}^3 , gives a quotient that we call a (rank 1) parabolic tube, with boundary parameters $(i\infty, t)$. This can be obtained as the geometric limit of hyperbolic tubes with parameters (ω_n, t) with $\operatorname{Im} \omega_n \rightarrow \infty$ and $\operatorname{Re} \omega_n$ bounded.

3.2.2. Margulis tubes. Let N_J for $J \subset \mathbb{R}$ denote the region in N where the injectivity radius times 2 is in J . Thus $N_{(0,\epsilon)}$ denotes the (open) ϵ -thin part of a hyperbolic manifold N , and $N_{[\epsilon,\infty)}$ denotes the ϵ -thick part.

By the Margulis Lemma (see Kazhdan-Margulis [33]) or Jørgensen's inequality (see Jørgensen [32], and Hersonsky [31] or Waterman [64] for the higher-dimensional case), there is a constant $\epsilon_M(n)$ known as the Margulis constant for \mathbb{H}^n , such that every component of an ϵ -thin part when $\epsilon \leq \epsilon_M(n)$ is of a standard shape. In particular, in dimensions 2 and 3 in the orientable case, such a component is either an open tubular neighborhood of a simple closed geodesic, or the quotient of an open horoball by a parabolic group of isometries of rank 1 or 2. We call these components " ϵ -Margulis tubes". In dimension 3, the former type are hyperbolic tubes as in the previous subsection, and the rank 1 parabolic (or "cusp") tubes were mentioned above as well. See Thurston [63] for a discussion of the thick-thin decomposition.

If γ is a nontrivial homotopy class of closed curves in a hyperbolic manifold with length less than ϵ , then it corresponds to an ϵ -Margulis tube, which we shall denote by $\mathbb{T}_\epsilon(\gamma)$. If Γ is a union of several curves we let $\mathbb{T}_\epsilon(\Gamma)$ denote the union of their Margulis tubes.

The radius $r(\gamma)$ of a hyperbolic ϵ -Margulis tube depends on its complex translation length, $\lambda(\gamma)$. Brooks-Matelski [19] apply Jørgensen's inequality to show that

$$(3.5) \quad r(\gamma) \geq \log \frac{1}{|\lambda(\gamma)|} - c_1$$

where the constant c_1 depends only on ϵ . Writing $\lambda = \ell + i\theta$, we note that when ℓ is small $|\lambda|$ can still be large. However, an additional pigeonhole principle argument in Meyerhoff [45] implies

$$(3.6) \quad r(\gamma) \geq \frac{1}{2} \log \frac{1}{\ell(\gamma)} - c_2$$

where again c_2 depends only on ϵ .

We also have some definite separation between Margulis tubes for different values of ϵ . Given $\epsilon \leq \epsilon_M(3)$ we have, for any $\epsilon' < \epsilon$,

$$(3.7) \quad \text{dist}(\partial\mathbb{T}_\epsilon(\gamma), \partial\mathbb{T}_{\epsilon'}(\gamma)) \geq \frac{1}{2} \log \frac{\epsilon}{\epsilon'} - c_3$$

(see [47] for a discussion).

Margulis constants. For the remainder of the paper we fix $\epsilon_0 \leq \epsilon_M(3)$.

Thurston [59] pointed out that for a π_1 -injective pleated surface $f : S \rightarrow N$, the thick part of S maps to the thick part of N . More precisely, there is a function $\epsilon_T : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, depending on S , such that

$$f(S_{[\epsilon, \infty)}) \subset N_{[\epsilon_T(\epsilon), \infty)}.$$

(See also Minsky [47] for a brief discussion.) Let ϵ_1 be a fixed constant smaller than $\min\{1, \epsilon_T(\epsilon_0)\}$.

Bers constant. There is a constant L_0 , depending only on S (see Bers [11], [12] and Buser [20]), such that any hyperbolic metric on S admits a pants decomposition of total length at most L_0 . In fact we can and will choose L_0 so that such a pants decomposition includes all simple geodesics of length bounded by ϵ_0 . This L_0 will also do as the constant in Theorem 2.2, since the sequence of bounded curves there is obtained by taking shortest curves in a sequence of pleated surfaces.

We can refine this slightly as follows: There is a function $\mathcal{L}(\epsilon)$ such that, if γ is a curve of σ -length at least ϵ , then there is a pants decomposition of total length at most $\mathcal{L}(\epsilon)$ which intersects γ (equivalently, does not contain γ as a component). To see this, start with a pants decomposition w of length bounded by L_0 . If γ is a component of w consider the shortest replacement for γ that produces an elementary move on w . The upper bound on w and lower bound on γ gives an upper bound on this replacement curve.

3.2.3. *Collars in surfaces.* Components of the thin part in a hyperbolic surface are annuli, and the radius estimate (3.5) can be made much more explicit. We note also that as a result of this estimate there is a positive function $\varepsilon(L)$ such

that, if γ is a curve in a hyperbolic surface which essentially intersects a curve t of σ -length at most L , then $\ell_\sigma(\gamma) > \varepsilon(L)$. This will be used in Section 6.2.

It will also be convenient to fix a standard construction of a *collar* for each nontrivial homotopy class of simple curves in a hyperbolic surface:

LEMMA 3.3. *Fix a constant $0 < c \leq 1$. For every hyperbolic structure σ on $\text{int}(R)$ and every homotopically nontrivial simple closed curve γ in R there is an open annulus $\mathbf{collar}(\gamma, \sigma) \subset R$ whose core is homotopic to γ , so that the following holds:*

- (1) *If $\ell_\sigma(\gamma) = 0$, then $\mathbf{collar}(\gamma, \sigma)$ is a horospherical neighborhood of the cusp associated to γ . Otherwise it is an open embedded tubular neighborhood of radius $w = w(\ell_\sigma(\gamma))$ of the geodesic representative of γ .*
- (2) *If $\ell_\sigma(\gamma) < \frac{c}{2}$ (including the cusp case), then the length of each boundary component of $\mathbf{collar}(\gamma, \sigma)$ is exactly c .*
- (3) *If β and γ are homotopically distinct, disjoint curves, then $\mathbf{collar}(\gamma, \sigma)$ and $\mathbf{collar}(\beta, \sigma)$ have disjoint closures. Indeed the collars are at least d apart, for a constant $d > 0$, if $\ell_\sigma(\gamma) \leq c$.*

Proof. This is done by a small variation of a construction of Buser [20] (in [49] we used a slightly different variation).

Let

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

and

$$w_c(t) = \cosh^{-1} \left(\frac{c}{t} \right)$$

(the latter is defined only for $t \in (0, c]$). Assuming for the moment that $\ell_\sigma(\gamma) > 0$, define $\mathbf{collar}_i(\gamma, \sigma)$ to be the $w_i(\ell_\sigma(\gamma))$ neighborhood of the geodesic representative of γ , where $i = 0$ or c . (If γ is isotopic to a boundary component of R , we take this neighborhood in the metric completion of $\text{int}(R)$, and then intersect with $\text{int}(R)$.)

Buser shows that $\mathbf{collar}_0(\gamma, \sigma)$ is always an embedded open annulus, and such collars are disjoint if nonhomotopic. On the other hand when $\ell_\sigma(\gamma) < c$, $\mathbf{collar}_c(\gamma, \sigma)$ has boundaries in $\text{int}(R)$ of length exactly c , since the boundary length of a collar of width w and core length ℓ is $\ell \cosh(w)$. We will obtain our desired collars by interpolating between these two.

Let $\delta(t) = w_0(t) - w_c(t)$. One can check that $\delta(t)$ is a positive function for $t \leq c/2$, and bounded away from 0 and ∞ . Thus, for $\ell_\sigma(\gamma) < c/2$, $\mathbf{collar}_c(\gamma, \sigma)$ is a subannulus of $\mathbf{collar}_0(\gamma, \sigma)$, and their boundaries are separated by a definite but bounded distance.

Thus we define

$$w(t) = \begin{cases} w_c(t) & t \leq c/2 \\ \alpha w_0(t) & t \geq c/2 \end{cases}$$

where $\alpha = w_c(c/2)/w_0(c/2) < 1$. Define $\mathbf{collar}(\gamma, \sigma)$ to be the $w(\ell_\sigma(\gamma))$ -neighborhood of the geodesic representative. This collar has all the desired properties. In particular the definite separation in part (3) follows from the lower bound on δ when $\ell_\sigma(\gamma) < c/2$, and from the fact that $\alpha < 1$ when $\ell_\sigma(\gamma) \in [c/2, c]$.

Collars in the cusp case are easily seen to be obtained from these in the limit as $\ell(\gamma) \rightarrow 0$. In particular one may check that $\delta(t)$ converges to $\log(2/c)$ as $t \rightarrow 0$, and \mathbf{collar}_0 has boundary length 2 in the limit. \square

Collar normalization. In the remainder of the paper we will assume that the constant c is equal to the constant $\epsilon_1 < 1$. Note that this means that $\mathbf{collar}(\gamma, \sigma)$ is contained in the component of the ϵ_1 -thin part associated with γ . If σ is implicitly understood; then we may write $\mathbf{collar}(\gamma)$. For a system Γ of disjoint, homotopically distinct simple closed curves we let $\mathbf{collar}(\Gamma, \sigma)$ or $\mathbf{collar}(\Gamma)$ be the union of the collars of the components.

3.3. Convention on isotopy representatives. Although we usually think of curves and subsurfaces in terms of their isotopy classes, it will be useful for our constructions and arguments to fix explicit representatives. Thus we will adopt the following convention for the remainder of the paper. We fix once and for all an oriented surface $S = S_{g,b}$. Let \hat{S} denote a separate copy of $\text{int}(S)$ and fix a complete, finite-area hyperbolic metric σ_0 on \hat{S} . Embed S inside \hat{S} as the complement of $\mathbf{collar}(\partial S, \sigma_0)$.

Now if v is an essential homotopy class of simple closed curves or simple properly embedded arcs in S , we let γ_v denote its geodesic representative with respect to σ_0 (occasionally we conflate v and γ_v when this is convenient). If v is a class of arcs, then γ_v is an infinite geodesic whose ends exit the cusps of \hat{S} .

We let $\mathbf{collar}(v)$ denote $\mathbf{collar}(v, \sigma_0)$ in \hat{S} , and we assume from now on that every open annulus is of the form $\mathbf{collar}(v)$ (and every closed annulus is the closure of such a collar). Similarly every other subsurface (including S itself) is a component of

$$S \setminus \mathbf{collar}(\Gamma)$$

for a system of curves Γ .

This has the property that if two isotopy classes of subsurfaces have disjoint representatives then these chosen representatives are already disjoint (except for a closed collar that may share a boundary component with an adjacent nonannular surface).

For a nonannular surface Y it is also true that all its intersections with any closed geodesic are essential, since its boundary is concave.

3.4. The augmented convex core and its exterior. Let C_N^r be the closed r -neighborhood of the convex core of any hyperbolic 3-manifold N , where $r \geq 0$. Define the *augmented core of N* to be

$$\widehat{C}_N = C_N^1 \cup N_{(0, \epsilon_0]}.$$

Assume from now on that $\partial\bar{N}$ is incompressible.

In this section we will develop a model for the exterior of the augmented core, analogous to the description given by Epstein-Marden in [26] for the exterior of the convex core.

First let us define a modified metric σ_m on the boundary at infinity, $\partial\bar{N}$, as follows:

Let σ_∞ denote the Poincaré metric on $\partial\bar{N}$. Let Γ denote the set of homotopy classes of curves in $\partial\bar{N}$ whose σ_∞ -length is at most ϵ_1 (including curves homotopic to cusps), and let $\mathbf{collar}(\Gamma, \sigma_\infty)$ be their standard collars in $\partial\bar{N}$. We will let σ_m be conformally equivalent to σ_∞ , with $d\sigma_m/d\sigma_\infty$ a continuous function which is equal to 1 on $\partial\bar{N} \setminus \mathbf{collar}(\Gamma, \sigma_\infty)$, and such that each annulus of $\mathbf{collar}(\Gamma, \sigma_\infty)$ is flat in the σ_m metric (a *flat annulus* for us is an annulus isometric to the product of a circle with an interval). This defines σ_m uniquely. The flat annuli have circumference exactly ϵ_1 , by the definition of the collars.

Now let E_N denote $N \setminus \text{int}(\widehat{C}_N)$, and let $\bar{E}_N = E_N \cup \partial\bar{N}$ be the closure of E_N in \bar{N} . Let E_ν denote a copy of $\partial\bar{N} \times [0, \infty)$, endowed with the metric

$$(3.8) \quad e^{2r} d\sigma_m^2 + dr^2$$

where r is a coordinate for the second factor. We can also let $\bar{E}_\nu = \partial\bar{N} \times [0, \infty]$ where the “boundary at infinity” $\partial_\infty E_\nu \equiv \partial\bar{N} \times \{\infty\}$ is endowed with the conformal structure of $\partial\bar{N}$. Note that the metric σ_m is determined completely by the end invariants ν , justifying the notation E_ν and \bar{E}_ν .

The rest of this subsection is devoted to proving the following lemma, which taken together with the Lipschitz Model Theorem will give us the proof of the Extended Model Theorem (see §10).

LEMMA 3.4. *When $\partial\bar{N}$ is incompressible, there is a homeomorphism*

$$\varphi : E_\nu \rightarrow E_N$$

which is locally bilipschitz with constant depending only on the topological type of $\partial\bar{N}$. Furthermore φ extends to a homeomorphism $\bar{E}_\nu \rightarrow \bar{E}_N$ where the map on $\partial_\infty E_\nu$ is given by $\text{id}_{\partial\bar{N}}$.

Proof. Recall first the following facts from [26]. The function $\delta : N \setminus C_N \rightarrow \mathbb{R}$ measuring distance from C_N is a C^1 function with $|\nabla\delta| = 1$ and so its level sets ∂C_N^r (for $r > 0$) are nonsingular C^1 surfaces. Convexity gives an orthogonal projection

$$\Pi_r : \bar{N} \setminus \text{int}(C_N^r) \rightarrow \partial C_N^r$$

(projection along gradient lines of δ) such that, by Theorem 2.3.1 of [26], the restriction to $\partial\bar{N}$ is a $K(r)$ -bilipschitz homeomorphism if $\partial\bar{N}$ is endowed with $\cosh r$ times its Poincaré metric, and the range is given the induced Riemannian metric from N , denoted $d\sigma_{\partial C_N^r}^2$. If r is bounded below, then $K(r)$ is bounded above. We also note that the estimates in [26] are stated for $r < \log 2$, but the proof applies to all r . Hence we will assume $r \geq 1$ and $K(r) \leq K_0$, a universal constant.

This gives us a homeomorphism

$$(3.9) \quad \Phi : N \setminus \text{int}(C_N^1) \rightarrow E_1 \equiv \partial C_N^1 \times [1, \infty)$$

defined by $\Phi(x) = (\Pi_1(x), \delta(x))$. The estimates on Π_1 imply that Φ is locally K_1 -bilipschitz, with K_1 a universal constant, where we take the hyperbolic metric on the domain and the metric

$$(3.10) \quad e^{2t} d\sigma_{\partial C_N^1}^2 + dt^2$$

on E_1 .

The first thing to note about our augmented core is that its boundary is a graph over the boundary of C_N^1 . To see this we work in the universal cover; let X^r be the lift to \mathbb{H}^3 of C_N^r , and continue to denote by Π_r and δ the lifts of those maps to \mathbb{H}^3 and X^r . Let $\tilde{\mathbb{T}}$ be a component of the preimage $\tilde{N}_{(0, \epsilon_0)}$ of $N_{(0, \epsilon_0)}$ in \mathbb{H}^3 . If the stabilizer of $\tilde{\mathbb{T}}$ is hyperbolic, then $\tilde{\mathbb{T}}$ is an R -neighborhood (for some R) of its geodesic axis L . Let $\beta(x) = d(x, L)$ be the distance function to L , so that $\tilde{\mathbb{T}} = \beta^{-1}([0, R])$. If $\tilde{\mathbb{T}}$ has a parabolic stabilizer, then it has a Busemann function, namely a function β constant on concentric horospheres in $\tilde{\mathbb{T}}$ and measuring signed distance between them, so that $\tilde{\mathbb{T}} = \beta^{-1}((-\infty, R])$.

We claim now that, for any $x \in \partial\tilde{\mathbb{T}}$ outside X^0 ,

$$(3.11) \quad \nabla\delta(x) \cdot \nabla\beta(x) \geq \tanh \delta(x)/2.$$

To see this, let $y = \Pi_0(x)$. Then $\nabla\delta(x)$ is the outward tangent vector to the geodesic arc $[y, x]$ at x . $\nabla\beta(x)$ is the outward tangent vector to a geodesic $[z, x]$ at x , where z is either a point on the axis L of $\tilde{\mathbb{T}}$ or the parabolic fixed point of $\tilde{\mathbb{T}}$. In either case z is in the closure of X^0 in \mathbb{H}^3 . Let P be the plane through y orthogonal to $[y, x]$. By convexity, all of X^0 , and in particular z , lies on the closure of the side of P opposite from x . The lowest possible value for $\nabla\delta \cdot \nabla\beta$ is therefore obtained when $z \in \partial P$, and for this configuration the value $\tanh \delta(x)/2$ is obtained easily from hyperbolic trigonometry. This proves (3.11).

In particular $\nabla\delta \cdot \nabla\beta$ is positive, and has a positive lower bound outside X^1 . Hence the gradient lines of δ can only exit $\tilde{N}_{(0, \epsilon_0)}$, not enter it. Thus the gradient

line meets $\partial\tilde{N}_{(0,\epsilon_0)}$ at most once. This tells us that $\partial\hat{C}_N$ is a graph over ∂C_N^1 with respect to the product structure of (3.9), and moreover that there is an upper bound to the angle between the tangent planes of $\partial\hat{C}_N$ and those of ∂C_N^r , the level sets of δ . (This applies at smooth points of the boundary, i.e. all points except the (transverse) intersection locus of $\partial N_{(0,\epsilon_0)}$ with ∂C_N^1 .)

From this we conclude that $\partial\hat{C}_N$ is a *locally uniformly Lipschitz graph* with respect to the product structure induced by δ , in the following sense: For any $p \in \partial\hat{C}_N$ there is a neighborhood W in N parametrized as $U \times J$, where $U \subset \partial C_N^{\delta(p)}$ and $J \subset \mathbb{R}$, projection to J is δ and vertical lines are gradient lines for δ , such that $\partial\hat{C}_N \cap W$ is the graph of a K_2 -Lipschitz function $U \rightarrow J$, where K_2 is a universal constant. Note that this is true for *all* points of $\partial\hat{C}_N$, not just the smooth ones, because at the corners where $\partial N_{(0,\epsilon_0)}$ meets ∂C_N^1 , both tangent planes are controlled. Since Φ preserves our product structures and is bilipschitz for the level sets, it follows that $F_0 \equiv \Phi(\partial\hat{C}_N)$ is also a locally uniformly Lipschitz graph with respect to the product structure of E_1 , with different constant.

Now consider for $s \geq 0$ the map $\lambda_s : E_1 \rightarrow E_1$ given by $\lambda_s(p, t) = (p, t + s)$. This map preserves length in the vertical direction and expands by e^s in the horizontal direction, and it follows that $F_s = \lambda_s(F_0)$ is a locally uniformly Lipschitz with the same (actually better) constant as F_0 . Let $\hat{E} = \cup_{s \geq 0} F_s$, i.e. the region above F_0 . Note that $\hat{E} = \Phi(N \setminus \text{int}(\hat{C}_N))$.

We may no longer have differentiability of the surfaces F_s (because the maps Π_r may not be C^1), but the Lipschitz graph property implies that each F_s is rectifiable. At this point it is convenient to consider a substitute for the Riemannian product in the smooth setting: If \mathcal{F} and \mathcal{G} are two transverse foliations in a manifold M and the leaves of each are endowed with intrinsic path metrics, we can consider the *product path metric* which is defined by considering only paths composed of segments that travel either along leaves of \mathcal{F} or of \mathcal{G} . If the metrics on \mathcal{F} and \mathcal{G} are Riemannian this recovers up to uniform bilipschitz equivalence the Riemannian product.

Now it is easy to check that the uniform Lipschitz graph property of the leaves F_t implies that the metric on \hat{E} is (uniformly) bilipschitz equivalent to the product path metric associated to the vertical foliation of E_1 and the transverse foliation by the F_t . This allows us to “straighten” \hat{E} : We consider the product $F_0 \times [0, \infty)$, endow $F_0 \times \{0\}$ with the metric of F_0 and each $F_0 \times \{t\}$ with that metric expanded by e^t . We give each vertical line the metric dt , and endow the product with the product path metric. Our discussion so far implies that the map from \hat{E} to $F_0 \times [0, \infty)$ which preserves the vertical lines and takes F_t to $F_0 \times \{t\}$ is uniformly bilipschitz.

To complete the picture, we have to relate the metric on F_0 (or equivalently $\partial\hat{C}_N$) to the metric σ_m on $\partial\tilde{N}$.

Note that F_0 is uniformly bilipschitz equivalent to a conformal multiple of ∂C_N^1 . This is because, near every point of F_0 , the local uniform Lipschitz graph property gives a local bilipschitz map to a neighborhood in a level surface $\partial C_N^1 \times \{r\}$ in E_1 , and this in turn is a conformal rescaling (by e^r) of ∂C_N^1 . In fact we claim that this conformal factor is estimated by the ratio of injectivity radii $\text{inj}_{F_0}/\text{inj}_{\partial C_N^1}$ at the point in question. This is because, at all points of $\partial \hat{C}_N$, the injectivity radius is bounded above and below by uniform constants (these depend on the topological type of the surface, unlike the previous constants). The log of the injectivity radius is itself a Lipschitz function on a hyperbolic surface, and it follows that the conformal factor can be recovered up to bounded multiple from the injectivity radius ratio.

Now the metric σ_m is a conformal rescaling of σ_∞ , which also has the property that its injectivity radius is uniformly bounded above and below. Since Π_1 is a uniformly bilipschitz map from σ_∞ to $\sigma_{\partial C_N^1}$, it follows that σ_m and the metric on F_0 are uniformly bilipschitz. Thus letting E_ν be defined as above, its metric is bilipschitz equivalent to the metric we've described for \hat{E} , and this gives the statement of Lemma 3.4 for the interior.

For the boundary at infinity, it suffices to note that Π_1 extends continuously there, so that our bilipschitz map from $E_\nu \rightarrow E_N$ extends to the desired map from \bar{E}_ν to \bar{E}_N . \square

4. Complexes of curves and arcs

The definitions below are originally due to Harvey [29], with some modifications in [43], [42] and [49].

The case of the one-holed torus $S_{1,1}$ and four-holed sphere, $S_{0,4}$, are special and will be treated below, as will the case of the annulus $S_{0,2}$, which will only occur as a subsurface of larger surfaces. We call all other cases “generic”.

If S is generic, we define $\mathcal{C}(S)$ to be the simplicial complex whose vertices are nontrivial, nonperipheral homotopy classes of simple curves, and whose k -simplices are sets $\{v_0, \dots, v_k\}$ of distinct vertices with disjoint representatives. For $k \geq 0$ let $\mathcal{C}_k(S)$ denote the k -skeleton of $\mathcal{C}(S)$.

We define a metric on $\mathcal{C}(S)$ by making each simplex regular Euclidean with sidelength 1, and taking the shortest-path distance. We will more often use the shortest-path distance in the 1-skeleton, which we denote $d_{\mathcal{C}_1(S)}$ – but note that the path metrics on $\mathcal{C}(S)$ and $\mathcal{C}_1(S)$ are quasi isometric. These conventions also apply to the nongeneric cases below.

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

In both these cases, $\mathcal{C}(S)$ is isomorphic to the classical Farey graph in the plane (see e.g. Hatcher-Thurston [30] or [46]).

Arc complexes. If Y is a nonannular surface with nonempty boundary, let us also define the larger *arc complex* $\mathcal{A}(Y)$ whose vertices are either properly embedded essential arcs in Y , up to homotopy keeping the endpoints in ∂Y , or essential closed curves up to homotopy. Simplices again correspond to sets of vertices with disjoint representatives.

Subsurface projections. Note that the vertices $\mathcal{A}_0(S)$ can be identified with a subset of the geodesic lamination space $\mathcal{GL}(S)$ – a geodesic leaf with ends in the cusps of a complete finite area hyperbolic structure on $\text{int}(S)$ determines a homotopy class of properly embedded arcs in S . Let Y be a nonannular essential subsurface of S . We can define a “projection”

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

as follows:

If $\alpha \in \mathcal{GL}(S)$ has no essential intersections with Y (including the case that α is homotopic to ∂Y), then define $\pi_Y(\alpha) = \emptyset$. Otherwise, $\alpha \cap Y$ is a collection of disjoint essential curves and/or properly embedded arcs (this follows from our use of geodesic representatives of isotopy classes), which therefore span a simplex in $\mathcal{A}(Y)$. Let $\pi_Y(\alpha)$ be an arbitrary choice of vertex of this simplex.

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

Let us denote

$$d_Y(\alpha, \beta) \equiv d_{\mathcal{A}_1(Y)}(\pi_Y(\alpha), \pi_Y(\beta)),$$

provided $\pi_Y(\alpha)$ and $\pi_Y(\beta)$ are nonempty. Similarly $\text{diam}_Y(A)$ denotes

$$\text{diam}_{\mathcal{A}_1(Y)}(\cup_{a \in A} \pi_Y(a)),$$

where $A \subset \mathcal{GL}(S)$.

Annuli. If Y is a closed annulus, let $\mathcal{A}(Y)$ be the complex whose vertices are essential homotopy classes, rel endpoints, of properly embedded arcs, and whose simplices are sets of vertices with representatives with disjoint interiors. Here it is important that endpoints are *not* allowed to move in the boundary.

It is again easy to see that $\mathcal{A}(Y)$ is quasi-isometric to its 1-skeleton $\mathcal{A}_1(Y)$, and we will mostly consider this.

Twist numbers. Fix an orientation of Y . Given vertices a, b in $\mathcal{A}(Y)$, we will define their *twist number* $\text{tw}_Y(a, b)$, which is a sort of rough signed distance function satisfying

$$(4.1) \quad |\text{tw}_Y(a, b)| \leq d_{\mathcal{A}_1(Y)}(a, b) \leq |\text{tw}_Y(a, b)| + 1.$$

First, let $\tilde{\mathcal{R}} \subset \mathbb{C}$ be the strip $\{\text{Im } z \in [0, 1]\}$, \mathcal{R} its quotient by $z \mapsto z + 1$, and identify Y with \mathcal{R} by an orientation-preserving homeomorphism. Choosing a lift \tilde{a} in $\tilde{\mathcal{R}}$ of a representative of a , denote its endpoints as x_a and $y_a + i$.

Next define a function $\xi \mapsto \xi'$ on \mathbb{R} by letting $\xi' = \xi$ if $\xi \in \mathbb{Z}$ and $\xi' = n + 1/2$ if $\xi \in (n, n + 1)$ for $n \in \mathbb{Z}$. We can now define

$$(4.2) \quad \text{tw}_Y(a, b) \equiv (y_b - y_a)' - (x_b - x_a)'.$$

This definition does not depend on any of the choices made. An isotopy of ∂Y along $\partial \mathcal{R}$ corresponds to changing the endpoints by a homeomorphism $h : \mathbb{R} \rightarrow \mathbb{R}$ satisfying $h(y + 1) = h(y) + 1$, and one can check that $(h(y_2) - h(y_1))' = (y_2 - y_1)'$. Interchanging boundary components (by an orientation-preserving map of Y) is taken care of by the identity $(-\xi)' = -\xi'$, and choosing different lifts of a and b corresponds to the identity $(\xi + 1)' = \xi' + 1$. It is also evident that $\text{tw}_Y(b, a) = -\text{tw}_Y(a, b)$.

The twist numbers are additive under concatenation of annuli: Suppose that Y is decomposed as a union of annuli Y_1 and Y_2 along their common boundary and α_1 and α_2 are two arcs connecting the boundaries of Y so that $\alpha_{ij} \equiv \alpha_i \cap Y_j$ is an arc joining the boundaries of Y_j for each $i = 1, 2$ and $j = 1, 2$. The definition of twist numbers easily yields

$$(4.3) \quad \text{tw}_Y(\alpha_1, \alpha_2) = \text{tw}_{Y_1}(\alpha_{11}, \alpha_{21}) + \text{tw}_{Y_2}(\alpha_{12}, \alpha_{22}).$$

One can furthermore check that

$$|(\xi + \eta)' - \xi' - \eta'| \leq 1/2$$

and this yields the approximate additivity property:

$$(4.4) \quad |\text{tw}_Y(a, c) - \text{tw}_Y(a, b) - \text{tw}_Y(b, c)| \leq 1$$

for any three $a, b, c \in \mathcal{A}(Y)$.

The inequality (4.1) relating tw_Y to $d_{\mathcal{A}_1(Y)}$ can be easily verified from the definitions.

From (4.4) and (4.1) we conclude that, fixing any $a \in \mathcal{A}_0(Y)$, the map $b \mapsto \text{tw}_Y(a, b)$ induces a quasi-isometry from $\mathcal{A}(Y)$ to \mathbb{Z} . We can also use twisting to define a notion of *signed length* for a geodesic: If h is a directed geodesic in $\mathcal{A}_1(Y)$ beginning at a and terminating in b , we write

$$(4.5) \quad [h] = \begin{cases} |h| & \text{tw}_Y(a, b) \geq 0, \\ -|h| & \text{tw}_Y(a, b) < 0. \end{cases}$$

Note that

$$[\vec{ab}] = -[\vec{ba}]$$

except when $\text{tw}_Y(a, b) = 0$ and $d_Y(a, b) = 1$, in which case $[\vec{ab}] = [\vec{ba}] = 1$. This is an effect of coarseness.

Now consider an essential annulus $Y \subset S$, and let us define the subsurface projection π_Y in this case.

There is a unique cover of S corresponding to the inclusion $\pi_1(Y) \subset \pi_1(S)$, to which we can append a boundary using the circle at infinity of the universal cover of S to yield a closed annulus \hat{Y} (take the quotient of the compactified hyperbolic plane minus the limit set of $\pi_1(Y)$).

We define $\pi_Y : \mathcal{GL}(S) \rightarrow \mathcal{A}(\hat{Y}) \cup \{\emptyset\}$ as follows: Any lamination λ that crosses Y essentially lifts in \hat{Y} to a collection of disjoint arcs, some of which are essential. Hence we obtain a collection of vertices of $\mathcal{A}(\hat{Y})$ any finite subset of which form a simplex. Let $\pi_Y(\lambda)$ denote an arbitrary choice of vertex in this collection.

To simplify notation we often refer to $\mathcal{A}(\hat{Y})$ as $\mathcal{A}(\alpha)$, where α is the core curve of Y . We also let $d_\alpha(\beta, \gamma)$ and $d_Y(\beta, \gamma)$ denote $d_{\mathcal{A}_1(\hat{Y})}(\pi_Y(\alpha), \pi_Y(\beta))$. Similarly we denote $\text{tw}_{\hat{Y}}(\pi_Y(\beta), \pi_Y(\gamma))$ by $\text{tw}_\alpha(\beta, \gamma)$, and we also write $\pi_\alpha = \pi_Y$.

Projection bounds. It is evident that the projections π_W have the following 1-Lipschitz property: If u and v are vertices in $\mathcal{A}_0(S)$, both $\pi_W(u) \neq \emptyset$ and $\pi_W(v) \neq \emptyset$, and $d_S(u, v) = 1$, then $d_W(u, v) \leq 1$ as well.

As a map from $\mathcal{C}_0(S)$ to $\mathcal{A}_0(W)$, π_W has the same 1-Lipschitz property when $\xi(S) > 4$. If $\xi(S) = 4$ and $\xi(W) = 2$, then the property holds with a Lipschitz constant of 3 (in Lemma 2.3 [42] this is shown with a slight error that leads to a constant of 2, but 3 is correct because two curves in S that intersect minimally can give rise to two arcs in the annulus that intersect twice, and hence have distance 3. We are grateful to Hideki Miyachi for pointing out this error).

A stronger contraction property applies to projection images of geodesics, and plays an important role in the construction of hierarchies in [42]:

LEMMA 4.1 (Masur-Minsky [42]). *If g is a (finite or infinite) geodesic in $\mathcal{C}_1(S)$ such that $\pi_W(v) \neq \emptyset$ for each vertex v in g , then*

$$\text{diam}_W(g) \leq A$$

where A depends only on S .

The complexes $\mathcal{C}(W)$ and $\mathcal{A}(W)$ are in fact quasi-isometric when $\xi(W) \geq 4$. The inclusion $\iota : \mathcal{C}_0(W) \rightarrow \mathcal{A}_0(W)$ has a quasi-inverse $\psi : \mathcal{A}_0(W) \rightarrow \mathcal{C}_0(W)$ defined as follows: On $\mathcal{C}_0(W)$ let ψ be the identity. If a is a properly embedded arc in W , then the boundary of a regular neighborhood of $a \cup \partial W$ contains either one or two essential curves, and we let $\psi([a])$ be one of these (chosen arbitrarily).

In [42, Lemma 2.2] we show that ψ is a 2-Lipschitz map, with respect to $d_{\mathcal{A}_1(W)}$ and $d_{\mathcal{C}_1(W)}$.

Hyperbolicity and Klarreich's theorem. In Masur-Minsky [43], we proved

THEOREM 4.2. $\mathcal{C}(R)$ is δ -hyperbolic.

A geodesic metric space X is δ -hyperbolic if all triangles are “ δ -thin”. That is, for any geodesic triangle $[xy] \cup [yz] \cup [xy]$, each side is contained in a δ -neighborhood of the union of the other two. The notion of δ -hyperbolicity, due to Gromov [28] and Cannon [24], encapsulates some of the coarse properties of classical hyperbolic space as well as metric trees and a variety of Cayley graphs of groups. See also Alonso et al. [5], Bowditch [15] and Ghys-de la Harpe [27] for more about δ -hyperbolicity.

In particular a δ -hyperbolic space X has a well-defined *boundary at infinity* ∂X , which is roughly the set of asymptotic classes of quasi-geodesic rays. There is a natural topology on $\bar{X} \equiv X \cup \partial X$. When X is proper (bounded sets are compact) \bar{X} is compact, but in our setting $\mathcal{C}(R)$ is not locally compact, and hence $\partial \mathcal{C}(R)$ and $\bar{\mathcal{C}}(R)$ are not compact. Since $\mathcal{C}(R)$ and $\mathcal{A}(R)$ are naturally quasi-isometric, $\partial \mathcal{C}(R)$ can be identified with $\partial \mathcal{A}(R)$.

Klarreich showed in [34] that this boundary can be identified with $\mathcal{EL}(R)$:

THEOREM 4.3 (Klarreich [34]). *There is a homeomorphism*

$$k : \partial \mathcal{C}(R) \rightarrow \mathcal{EL}(R),$$

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

(Note that $\mathcal{C}_0(R)$ can be considered as a subset of $\mathcal{UML}(R)$, and hence the convergence $\beta_i \rightarrow k(\beta)$ makes sense.)

5. Hierarchies

In this section we will introduce the notion of hierarchies of tight geodesics, and discuss their basic properties. Most of the material here comes directly from [42], although the setting here is slightly more general in allowing infinite geodesics in the hierarchy. Thus, although we will mostly state definitions and facts, we will also need to indicate the changes to the arguments in [42] needed to treat the infinite cases. In particular Lemma 5.14 is an existence result for infinite geodesics, and Lemma 5.8 establishes a crucial property of “resolutions by slices” which was essentially immediate in the finite case.

This chapter is technical in spite of the author's efforts to simplify things. On a first pass, the reader is encouraged to ignore transversals and annulus geodesics. This makes markings into pants decompositions and avoids a number of special

cases, for example in the definition of subordinacy. The reader is also referred to [42] for examples and a motivating discussion.

5.1. Definitions and existence.

Generalized markings. The simplest kind of marking of a surface S is a system of curves that makes a simplex in $\mathcal{C}(S)$. In general we may want to include twist information for each curve, and we also want to include geodesic laminations instead of curves.

Thus we define a *generalized marking* of S as a lamination $\lambda \in \mathcal{UML}(S)$ together with a (possibly empty) set of *transversals* T , where each element t of T is a vertex of $\mathcal{A}(\alpha)$ for a closed component α of λ . For each closed component α of λ there is at most one transversal. The lamination λ is called $\text{base}(\mu)$, and the sublamination of λ consisting of closed curves is called $\text{simp}(\mu)$. Note that it is a simplex of $\mathcal{C}(S)$.

We further require that every nonclosed component ν of $\text{base}(\mu)$ is *filling* in the component R of $S \setminus \text{collar}(\text{simp}(\mu))$ that contains it; that is, $\nu \in \mathcal{EL}(R)$.

This is a generalization of the notion of marking used in [42], for which $\text{base}(\mu) = \text{simp}(\mu)$. We will call markings with the latter property *finite*.

Types of markings. When μ is a generalized marking let us call $\text{base}(\mu)$ *maximal* if it is not a proper subset of any element of $\mathcal{UML}(S)$. Equivalently, each complementary component Y of $\text{simp}(\mu)$ is either a 3-holed sphere or supports a component of $\text{base}(\mu)$ which is in $\mathcal{EL}(Y)$. In particular the base of a finite marking is maximal if and only if it is a pants decomposition.

We call μ itself *maximal* if it is not properly contained in any other marking – equivalently, if its base is maximal, and if every component of $\text{simp}(\mu)$ has a transversal.

It will also be important to consider *clean markings*: a marking μ is clean if $\text{base}(\mu) = \text{simp}(\mu)$, and if each transversal t for a component α of $\text{base}(\mu)$ has the form $\pi_\alpha(\bar{t})$, where $\bar{t} \in \mathcal{C}_0(S)$ is disjoint from $\text{simp}(\mu) \setminus \{\alpha\}$, and where $\alpha \cup \bar{t}$ fill a surface W with $\xi(W) = 4$ in which α and \bar{t} are adjacent as vertices of $\mathcal{C}_1(W)$. We will sometimes refer to \bar{t} as the transversal to α in this case.

If $Y \subset S$ is an essential subsurface and μ is a marking in S with $\text{base}(\mu) \in \mathcal{UML}(Y)$, then we call μ a *marking in Y* . (Note that the transversals are not required to stay in Y .) If Y is an annulus, then a marking in Y is just a simplex of $\mathcal{A}(\hat{Y})$.

We can extend the definition of the subsurface projections π_W to markings as follows: We let $\pi_W(\mu) = \pi_W(\text{base}(\mu))$ if the latter is nonempty. If W is an annulus and its core is a component α of $\text{base}(\mu)$, then we let $\pi_W(\mu)$ be the transversal for α , if one exists in μ . In all other cases $\pi_W(\mu) = \emptyset$.

With these definitions, if $\text{base}(\mu)$ is maximal, $\pi_W(\mu)$ is nonempty for any essential subsurface W with $\xi(W) \geq 4$, and for any three-holed sphere that is not a complementary component of $\text{simp}(\mu)$. If μ is maximal, $\pi_W(\mu)$ is nonempty for all essential subsurfaces W except three-holed spheres that are components of $S \setminus \text{simp}(\mu)$.

Tight geodesics. A pair of simplices α, β in a $\mathcal{C}(Y)$ are said to *fill* Y if all nontrivial nonperipheral curves in Y intersect at least one of γ_α or γ_β . If Y is an essential subsurface of S , then it also holds that any curve γ in S which essentially intersects a boundary component of Y must intersect one of γ_α or γ_β .

Given arbitrary simplices α, β in $\mathcal{C}(S)$, there is a unique essential subsurface $F(\alpha, \beta)$ (up to isotopy) which they fill: Namely, form a regular neighborhood of $\gamma_\alpha \cup \gamma_\beta$, and fill in all disks and one-holed disks. Note that F is connected if and only $\gamma_\alpha \cup \gamma_\beta$ is connected.

For a subsurface $X \subseteq Z$ let $\partial_Z(X)$ denote the *relative boundary* of X in Z , i.e. those boundary components of X that are nonperipheral in Z .

Definition 5.1. Let Y be an essential subsurface in S . If $\xi(Y) > 4$, a sequence of simplices $\{v_i\}_{i \in \mathcal{I}} \subset \mathcal{C}(Y)$ (where \mathcal{I} is a finite or infinite interval in \mathbb{Z}) is called *tight* if

- (1) For any vertices w_i of v_i and w_j of v_j where $i \neq j$, $d_{\mathcal{C}_1(Y)}(w_i, w_j) = |i - j|$,
- (2) Whenever $\{i - 1, i, i + 1\} \subset \mathcal{I}$, v_i represents the relative boundary

$$\partial_Y F(v_{i-1}, v_{i+1}).$$

If $\xi(Y) = 4$, then a tight sequence is just the vertex sequence of any geodesic in $\mathcal{C}_1(Y)$.

If $\xi(Y) = 2$, then a tight sequence is the vertex sequence of any geodesic in $\mathcal{A}(\hat{Y})$, with the added condition that the set of endpoints on $\partial \hat{Y}$ of arcs representing the vertices equals the set of endpoints of the first and last arc.

Note that condition (1) of the definition specifies that given any choice of components w_i of v_i the sequence $\{w_i\}$ is the vertex sequence of a geodesic in $\mathcal{C}_1(Y)$. It also implies that $\gamma_{v_{i-1}}$ and $\gamma_{v_{i+1}}$ always have connected union.

In the annulus case, the restriction on endpoints of arcs is of little importance, serving mainly to guarantee that between any two vertices there are only finitely many tight sequences.

With this in mind, a *tight geodesic* will be a tight sequence together with some additional data:

Definition 5.2. A *tight geodesic* g in $\mathcal{C}(Y)$ consists of a tight sequence $\{v_i\}_{\mathcal{I}}$, and two generalized markings in Y , $\mathbf{I} = \mathbf{I}(g)$ and $\mathbf{T} = \mathbf{T}(g)$, called its *initial* and *terminal* markings, such that:

If $i_0 = \inf \mathcal{J} > -\infty$, then v_{i_0} is a vertex of $\text{base}(\mathbf{I})$. If $i_\omega = \sup \mathcal{J} < \infty$, then v_{i_ω} is a vertex of $\text{base}(\mathbf{T})$.

If $\inf \mathcal{J} = -\infty$, then $\text{base}(\mathbf{I})$ is an element of $\mathcal{EL}(Y)$, and $\lim_{i \rightarrow -\infty} v_i = \text{base}(\mathbf{I})$ in $\mathcal{C}(Y) \cup \partial \mathcal{C}(Y)$, via the identification of Theorem 4.3. The corresponding limit holds for \mathbf{T} if $\sup \mathcal{J} = \infty$.

The length $|\mathcal{J}| \in [0, \infty]$ is called the length of g , usually written $|g|$. We refer to each of the v_i as *simplices* of g (in [42] we abused notation and called them “vertices”, and in case $\xi(D(g)) = 4$ we may still do so). Y is called the *domain* or *support* of g and we write $Y = D(g)$. We also say that g is *supported in* $D(g)$.

If Y is an annulus in S , then the markings $\mathbf{I}(g)$ and $\mathbf{T}(g)$ are just simplices in $\mathcal{A}(\hat{Y})$. We can also define the signed length $[g]$, as in (4.5).

We denote the obvious linear order in g as $v_i < v_j$ whenever $i < j$.

If v_i is a simplex of g define its *successor*

$$\text{succ}(v_i) = \begin{cases} v_{i+1} & v_i \text{ is not the last simplex} \\ \mathbf{T}(g) & v_i \text{ is the last simplex} \end{cases}$$

and similarly define $\text{pred}(v_i)$ to be v_{i-1} or $\mathbf{I}(g)$.

Subordinacy.

Restrictions of markings. If W is an essential subsurface in S and μ is a marking in S , then the *restriction* of μ to W , which we write $\mu|_W$, is constructed from μ in the following way: Suppose first that $\xi(W) \geq 3$. We let $\text{base}(\mu|_W)$ be the union of components of $\text{base}(\mu)$ that meet W essentially, and let the transversals of $\mu|_W$ be those transversals of μ that are associated to components of $\text{base}(\mu|_W)$.

If W is an annulus ($\xi(W) = 2$), then $\mu|_W$ is just $\pi_W(\mu)$.

Note in particular that, if all the base components of μ which meet W essentially are actually contained in W , then $\mu|_W$ is in fact a marking in W . If W is an annulus, then $\mu|_W$ is a marking in W whenever it is nonempty.

Component domains. Given a surface W with $\xi(W) \geq 4$ and a simplex v in $\mathcal{C}(W)$ we say that Y is a *component domain* of (W, v) if either Y is a component of $W \setminus \text{collar}(v)$, or Y is a component of $\text{collar}(v)$. Note that in the latter case Y is nonperipheral in W .

If g is a tight geodesic with domain $D(g)$, then we call $Y \subset S$ a *component domain* of g if for some simplex v_j of g , Y is a component domain of $(D(g), v_j)$. We note that g and Y determine v_j uniquely. In such a case, let

$$\mathbf{I}(Y, g) = \text{pred}(v_j)|_Y,$$

$$\mathbf{T}(Y, g) = \text{succ}(v_j)|_Y.$$

Note in particular that these are indeed markings in Y if nonempty, except when $\xi(Y) = 3$ (in which case they are just markings in S whose bases intersect Y).

If Y is a component domain of g and $\mathbf{T}(Y, g) \neq \emptyset$, then we say that Y is *directly forward subordinate* to g , or $Y \searrow^d g$. Similarly if $\mathbf{I}(Y, g) \neq \emptyset$ we say that Y is *directly backward subordinate* to g , or $g \swarrow^d Y$. Note that both $Y \searrow^d g$ and $g \swarrow^d Y$ can happen simultaneously, and often do (we may write $g \swarrow^d Y \searrow^d g$). To clarify this idea let us enumerate the possible cases when $\mathbf{T}(Y, g) \neq \emptyset$:

- (1) $\xi(Y) \geq 4$: In this case $\text{succ}(v_j)$ must have curves that are contained in Y . If $\text{succ}(v_j) = \mathbf{T}(g)$, then $\mathbf{T}(Y, g)$ will contain the base components of $\mathbf{T}(g)$ that are contained in Y , together with their transversals if any.
- (2) $\xi(Y) = 3$: Here we must have $\xi(D(g)) = 4$, and v_j cannot be the last vertex of g . $\mathbf{T}(Y, g)$ is the single vertex v_{j+1} .
- (3) $\xi(Y) = 2$: If v_j is not the last simplex of g , then again we must have $\xi(D(g)) = 4$, and $\mathbf{T}(Y, g) = \pi_Y(v_{j+1})$. If v_j is the last simplex, then $\mathbf{T}(g)$ must contain a transversal for the core curve of Y . This transversal becomes $\mathbf{T}(Y, g)$.

We can now define subordinacy for geodesics:

Definition 5.3. If k and g are tight geodesics, we say that k is *directly forward subordinate* to g , or $k \searrow^d g$, provided $D(k) \searrow^d g$ and $\mathbf{T}(k) = \mathbf{T}(D(k), g)$. Similarly we define $g \swarrow^d k$ to mean $g \swarrow^d D(k)$ and $\mathbf{I}(k) = \mathbf{I}(D(k), g)$.

We denote by *forward-subordinate*, or \searrow , the transitive closure of \searrow^d , and similarly for \swarrow . We let $h \supseteq k$ denote the condition that $h = k$ or $h \searrow k$, and similarly for $k \subseteq h$. We include the notation $Y \searrow f$ where Y is a subsurface to mean $Y \searrow^d f'$ for some f' such that $f' \supseteq f$, and similarly define $b \swarrow Y$.

For motivating examples of this structure in genus 1 and 2, see Section 1.5 of [42].

Definition of hierarchies.

Definition 5.4. A *hierarchy of geodesics* is a collection H of tight geodesics in S with the following properties:

- (1) There is a distinguished *main geodesic* g_H with domain $D(g_H) = S$. The initial and terminal markings of g_H are denoted also $\mathbf{I}(H), \mathbf{T}(H)$.
- (2) Suppose $b, f \in H$, and $Y \subset S$ is a subsurface with $\xi(Y) \neq 3$, such that $b \swarrow^d Y$ and $Y \searrow^d f$. Then H contains a unique tight geodesic k such that $D(k) = Y$, $b \swarrow^d k$ and $k \searrow^d f$.
- (3) For every geodesic k in H other than g_H , there are $b, f \in H$ (not necessarily distinct) such that $b \swarrow^d k \searrow^d f$.

Remark. The notation here differs from that in [42] only in the case of $\xi(Y) = 3$. Here we allow $\mathbf{T}(Y, g)$ and $\mathbf{I}(Y, g)$ to be nonempty, and hence $Y \searrow^d f$ and $b \not\searrow Y$ can occur, but we still explicitly disallow Y to be a domain of a geodesic in a hierarchy.

We will now investigate the structure of hierarchies, leaving the question of their existence until Section 5.5.

5.2. The descent theorem.

Footprints. If h is a tight geodesic in $D(h)$ and Y is an essential subsurface of S , we let the *footprint* of Y on h , denoted $\phi_h(Y)$, be the set of simplices of h that have no essential intersection with Y . If $Y \subset D(h)$, the triangle inequality implies that $\text{diam}_{\epsilon_1(D(h))}(\phi_h(Y)) \leq 2$. The condition of tightness implies that, if u, v, w are successive simplices of h and $u, w \in \phi_h(Y)$, then $v \in \phi_h(Y)$ as well. We remark that this is the only place where the tightness assumption is used.

Thus $\phi_h(Y)$, if nonempty, is a subinterval of 1, 2 or 3 successive simplices in h . Let $\min \phi_h(Y)$ and $\max \phi_h(Y)$ denote the first and last of these. Suppose Y is a component domain of $(D(h), v)$ for a simplex v of h . By definition, if $Y \searrow^d h$, then $v = \max \phi_h(Y)$, and similarly if $h \not\searrow Y$, then $v = \min \phi_h(Y)$. An inductive argument then yields

LEMMA 5.5. *If $Y \searrow f \searrow h$, then*

$$\max \phi_h(Y) = \max \phi_h(D(f)).$$

Similarly, if $h \not\searrow f \not\searrow Y$, then

$$\min \phi_h(Y) = \min \phi_h(D(f)).$$

This is used in [42] to control the structure of the “forward and backward sequences” of a geodesic in a hierarchy. That is, if $k \in H$, then by definition there exists $f \in H$ such that $k \searrow^d f$. It can be shown that this f is unique, so that there is a unique sequence $k \searrow^d f_0 \searrow^d f_1 \searrow^d \cdots \searrow^d g_H$, and similarly in the backward direction. The structure of these sequences is crucial to understanding and using hierarchies.

In particular it follows from Lemma 5.5 that if $k \searrow f$, then $D(k)$ intersects $\mathbf{T}(f)$ nontrivially. This condition can in fact capture all geodesics to which k is forward-subordinate. Let us define, for any essential subsurface Y of S ,

$$(5.1) \quad \Sigma_H^+(Y) = \{f \in H : Y \subseteq D(f) \text{ and } \mathbf{T}(f)|_Y \neq \emptyset\}$$

and similarly

$$(5.2) \quad \Sigma_H^-(Y) = \{b \in H : Y \subseteq D(f) \text{ and } \mathbf{I}(b)|_Y \neq \emptyset\}.$$

(We write Σ^\pm when H is understood.) Theorem 4.7 of [42] is a central result in that paper, and describes the structure of $\Sigma^\pm(Y)$. We present here a slight extension of that theorem which expands a little on the case of three-holed spheres.

THEOREM 5.6. *Let H be a hierarchy in S , and Y any essential subsurface of S .*

- (1) *If $\Sigma_H^+(Y)$ is nonempty, then it has the form $\{f_0, \dots, f_n\}$ where $n \geq 0$ and*

$$f_0 \searrow^d \cdots \searrow^d f_n = g_H.$$

Similarly, if $\Sigma_H^-(Y)$ is nonempty then it has the form $\{b_0, \dots, b_m\}$ with $m \geq 0$, where

$$g_H = b_m \swarrow^d \cdots \swarrow^d b_0.$$

- (2) *If $\Sigma^\pm(Y)$ are both nonempty and $\xi(Y) \neq 3$, then $b_0 = f_0$, and Y intersects every simplex of f_0 nontrivially.*
- (3) *If Y is a component domain in any geodesic $k \in H$, then*

$$f \in \Sigma^+(Y) \iff Y \searrow f,$$

and similarly,

$$b \in \Sigma^-(Y) \iff b \swarrow Y.$$

If, furthermore, $\Sigma^\pm(Y)$ are both nonempty and $\xi(Y) \neq 3$, then in fact Y is the support of $b_0 = f_0$.

- (4) *Geodesics in H are determined by their supports. That is, if $D(h) = D(h')$ for $h, h' \in H$, then $h = h'$.*

Proof. This theorem is taken verbatim from [42], except that part (3) is stated there only for $\xi(Y) \neq 3$. When $\xi(Y) = 3$, the definition in this paper of the relations $Y \searrow^d h$ and $h \swarrow^d Y$ is slightly different from the definition in [42] (see §5.1), and one can easily check that with this definition the proof in fact goes through verbatim in all cases (the relevant argument is to be found in Lemma 4.21 of [42]). \square

Completeness. A hierarchy H is k -complete if every component domain W in H with $\xi(W) \neq 3$ and $\xi(W) \geq k$ is the domain of some geodesic in H . H is complete if it is 2-complete.

If $\mathbf{I}(H)$ and $\mathbf{T}(H)$ are maximal markings, then $\mathbf{I}(H)|_W$ and $\mathbf{T}(H)|_W$ are nonempty for every W such that $\xi(W) \neq 3$. Theorem 5.6 part (3) then implies that H is complete.

If $\mathbf{I}(H)$ and $\mathbf{T}(H)$ only have maximal bases, then $\mathbf{I}(H)|_W$ is nonempty whenever $\xi(W) \neq 3$, except when $W = \mathbf{collar}(v)$ for a vertex $v \in \text{base}(\mathbf{I}(H))$ which has no transversal; and similarly for $\mathbf{T}(H)|_W$. It follows that H is 4-complete.

Slices and resolutions. A *slice* of a hierarchy H is a set τ of pairs (h, v) , where $h \in H$ and v is a simplex of h , satisfying the following properties:

- S1: A geodesic h appears in at most one pair in τ .
- S2: There is a distinguished pair (h_τ, v_τ) in τ , called the bottom pair of τ . We call h_τ the bottom geodesic.
- S3: For every $(k, w) \in \tau$ other than the bottom pair, $D(k)$ is a component domain of $(D(h), v)$ for some $(h, v) \in \tau$.

To a slice τ is associated a *marking*, μ_τ , whose base is the union $\{v : (h, v) \in \tau, \xi(D(h)) \geq 4\}$. The transversal curves are the vertices v for $(h, v) \in \tau$ with $\xi(D(h)) = 2$. We say that τ is *saturated* if it satisfies

- S4: Given $(h, v) \in \tau$, if $k \in H$ has $D(k)$ equal to a component domain of $(D(h), v)$, then there is a pair $(k, w) \in \tau$.

We say that τ is *full* if a stronger condition holds:

- S4': Given $(h, v) \in \tau$, if Y is a component domain of $(D(h), v)$ and $\xi(Y) \neq 3$, then there is a pair $(k, w) \in \tau$ with $D(k) = Y$.

(This terminology differs slightly from [42], but has the same content. Also, the $\xi \neq 3$ condition was mistakenly left out of the condition corresponding to (S4') in [42].)

A full slice whose bottom geodesic is equal to the main geodesic g_H will be called *maximal*. Note that if τ is maximal, then μ_τ is a maximal finite marking, and in particular $\text{base}(\mu_\tau)$ is a pants decomposition.

If H is complete, then every saturated slice is full.

Elementary moves on slices. We now define a *forward elementary move* on a saturated slice τ . Say that a pair (h, v) in τ is *forward movable* if:

- M1: v is not the last simplex of h . Let $v' = \text{succ}(v)$.
- M2: For every $(k, w) \in \tau$ with $D(k) \subset D(h)$ and $v'|_{D(k)} \neq \emptyset$, w is the last simplex of k .

When this occurs we can obtain a slice τ' from τ by replacing (h, v) with (h, v') , erasing all the pairs (k, w) that appear in condition (M2), and inductively replacing them (starting with component domains of $(D(h), v')$) so that the final τ' is saturated and satisfies

- M2': For every $(k', w') \in \tau'$ with $D(k') \subset D(h)$ and $v|_{D(k')} \neq \emptyset$, w' is the first simplex of k .

It is easy to see that τ' exists and is uniquely determined by this rule. We write $\tau \rightarrow \tau'$, and say that the move *advances* (h, v) to (h, v') .

The definition can also be reversed: A pair (h, v') in τ' is *backward movable* if it satisfies

M1': v' is not the first simplex of h

and the condition (M2'), with $v \equiv \text{pred}(v')$. We can then construct τ by erasing pairs appearing in (M2') and replacing them with pairs so that (M2) is satisfied.

These definitions come from [42], where we consider elementary moves only for maximal slices, but the definitions make sense in general.

Let $V(H)$ denote the set of saturated slices of H whose bottom geodesic is g_H . We remark that $V(H)$ is nonempty, since starting with any simplex of g_H we can successively add pairs in component domains to obtain a saturated slice. Note that if H is complete, then $V(H)$ is just the set of maximal slices. If H is 4-complete, then every slice $\tau \in V(H)$ has $\text{base}(\mu_\tau)$ a pants decomposition (but μ_τ may be missing transversals).

For a complete finite hierarchy, we show in [42] the existence of a *resolution*, which is a sequence $\{\tau_i\}_{i=0}^N$ in $V(H)$, such that $\tau_i \rightarrow \tau_{i+1}$, τ_0 is the initial slice of H , and τ_N is the terminal slice. (The initial slice is the unique slice $\tau \in V(H)$ such that for every pair $(h, v) \in \tau$, v is the first simplex of h . The terminal slice is defined similarly.) To do this, we show that for each $\tau \in V(H)$ there is at least one elementary move $\tau \rightarrow \tau'$, unless τ is the terminal slice. Beginning with the initial slice we successively apply elementary moves until we obtain the terminal slice. There is a certain partial order \prec_s on $V(H)$ such that $\tau \prec_s \tau'$ whenever $\tau \rightarrow \tau'$, and hence no τ can appear twice, and the process terminates by finiteness.

We wish to extend this idea to hierarchies that contain infinite geodesics, and we also want to consider the case that $\mathbf{I}(H)$ and $\mathbf{T}(H)$ may not be maximal.

In this setting we can consider sequences $\{\tau_i\}_{i \in \mathcal{I}}$ in $V(H)$ where $\mathcal{I} \subset \mathbb{Z}$ is a possibly infinite interval. For every slice $\tau \in V(H)$, if it is not terminal we will be able to find at least one move $\tau \rightarrow \tau'$ and if it is not initial, at least one $\tau'' \rightarrow \tau$. Thus we can begin with any slice in $V(H)$ and apply moves in both the forward and backward direction. There is only one pitfall: when there is an infinite geodesic whose support is a proper subsurface, we must avoid making infinitely many moves within that subsurface while the complement gets “stuck”. Therefore we would like to consider sequences with the following property:

- R: If (h, v) is forward movable in τ_i , then there exists some $j \geq i$ such that the move $\tau_j \rightarrow \tau_{j+1}$ advances the pair (h, v) . Similarly if (h, v) is backward movable, then there exists $j \leq i$ such that the move $\tau_{j-1} \rightarrow \tau_j$ advances from $(h, \text{pred}(v))$ to (h, v) .

We call an elementary move sequence $\{\tau_i\}$ satisfying this condition a resolution.

LEMMA 5.7. *For any hierarchy H and any slice $\tau_0 \in V(H)$, there is a resolution of H containing τ_0 .*

Proof. Fix τ_0 in $V(H)$. We define the rest of the resolution sequence inductively.

Suppose τ_j has been defined for $0 \leq j \leq i$. If $(h, v) \in \tau_i$ is forward movable we have seen that there exists a move $\tau_i \rightarrow \tau'$ that advances (h, v) . If $(k, w) \in \tau$ is also forward movable, we claim that (k, w) is still in τ' and still forward movable. $D(k)$ and $D(h)$ are either disjoint or contained one in the other, since they are both in the same slice. If they are disjoint, then the claim is evident. Suppose $D(k) \subset D(h)$. Since (k, w) is forward movable, w is not the last simplex of k , so since (h, v) is forward movable and hence satisfies condition (M2), we must have $\text{succ}(v)|_{D(k)} = \emptyset$. Hence (k, w) remains in τ' , and so does any pair with domain in $D(k)$. It follows that (k, w) is forward movable in τ' . Finally if $D(h) \subset D(k)$, the same argument tells us that $\text{succ}(w)|_{D(h)} = \emptyset$. Since only domains in $D(h)$ are changed by the move, it follows that $(k, w) \in \tau'$ and is still forward movable. Thus the claim is established.

We can therefore assemble a finite list of all forward movable pairs in τ_i , and advance them one after the other (in any order). This yields a sequence $\tau_i \rightarrow \tau_{i+1} \rightarrow \cdots \rightarrow \tau_i'$, such that no pair that was movable in τ_i remains in τ_i' . We continue inductively, stopping only if we reach a slice with no movable pairs. The resulting sequence $\tau_0 \rightarrow \tau_1 \rightarrow \cdots$ satisfies the forward half of condition (R). The same argument works in the backward direction, yielding the backward half of condition (R) for $\cdots \rightarrow \tau_{-2} \rightarrow \tau_{-1} \rightarrow \tau_0$. We also need to check that the forward half of (R) works for the negative-index τ_i , and vice versa, but for this it suffices to note that a forward movable pair remains forward movable until it is advanced, and hence will eventually be taken care of on the positive side. \square

The basic important property of the resolution is that it sweeps through all parts of the hierarchy, in a monotonic way:

LEMMA 5.8. *Let H be a hierarchy and $\{\tau_i\}_{i \in \mathcal{J}}$ a resolution. For any pair (h, v) with $h \in H$ and v a simplex of h , there is a slice τ_i containing (h, v) .*

Furthermore, if v is not the last simplex of h , then there is exactly one elementary move $\tau_i \rightarrow \tau_{i+1}$ which advances (h, v) .

Proof. Let us first prove the following property of a resolution:

G: If $(k, w) \in \tau_i$, then there exists a finite interval $\tau_j \dots \tau_k$ in which all pairs $\{(k, w') : w' \text{ a simplex of } k\}$ are obtained.

If (k, w) is already forward movable, then condition (R) yields $j \geq i$ such that τ_j contains $(k, \text{succ}(w))$. Similarly if (k, w) is backward movable, then condition (R) yields $j \leq i$ such that τ_j contains $(k, \text{pred}(w))$.

Let us prove property (G) by induction on $\xi(D(k))$. If $\xi(D(k))$ is minimal among all $\{\xi(D(g)) : g \in H\}$, then (k, w) is automatically forward movable whenever w is not the last simplex, and backward movable whenever w is not the first simplex (because conditions (M2) and (M2') hold vacuously). Thus we can use condition (R) as in the previous paragraph to obtain all simplices of k in finitely many steps.

Now in general, suppose w is not the last simplex of k , but that (k, w) is not movable. Let $w' = \text{succ}(w)$. Then (M2) is violated so there is some $(p, u) \in \tau_i$ with $D(p) \subset D(k)$, $w'|_{D(p)} \neq \emptyset$, and u not the last simplex of p . For each such p , since $\phi_k(D(p))$ contains w (by definition of a slice), but not w' (since $w'|_{D(p)} \neq \emptyset$), we must have $\mathbf{T}(k)|_{D(p)} \neq \emptyset$. Thus $k \in \Sigma^+(D(p))$, and applying Theorem 5.6 we have $p \searrow k$. It follows that p is not infinite in the forward direction, for if it were, $\mathbf{T}(p)$ would be a lamination in $\mathcal{EL}(D(p))$ and then, inductively, for each f such that $p \searrow f$ we would have $\max \phi_f(D(p))$ equal to the last vertex of f – and this contradicts the fact that $w' \notin \phi_k(D(p))$.

Thus, by the inductive hypothesis we can advance such a p in finitely many moves to its last simplex. Apply this first to a p with $D(p)$ a component domain of $(D(k), w)$, reaching a slice $\tau_{i'}$ ($i' \geq i$) containing (p, u') with u' the last vertex of p . Now for all $j \geq i'$, as long as (k, w) has not been advanced to (k, w') , the pair (p, u') must still be in τ_j . This is because no other move can affect p : since $D(p)$ is a component domain of $(D(k), w)$, there is no intervening k' in the slice with $D(p) \subset D(k') \subset d(k)$. Thus we can apply the same argument to other component domains of $(D(k), w)$ (if any) without removing (p, u') . Once these component domains have had their pairs advanced to satisfy (M1), we repeat the argument successively in *their* component domains. In finitely many moves we will therefore reach a point where (k, w) is forward movable, and we can apply (R). The same argument applies in the backward direction. Repeating this, we can obtain all the simplices of k , which establishes (G).

Now let (h, v) be any pair, and let τ be any maximal slice (say, $\tau = \tau_0$). We claim that there is a unique pair $(k, w) \in \tau$ such that $D(h) \subseteq D(k)$ and $\phi_k(D(h))$ does not contain w . We find this pair by induction: let $(g, u) \in \tau$ be a pair with $D(h) \subseteq D(g)$ (the bottom pair has this property). If $u \notin \phi_g(D(h))$, then (g, u) satisfies our conditions, and no other pairs in τ with domain in $D(g)$ can contain $D(h)$. (Note, this case includes $h = g$, in which case $\phi_g(D(h)) = \emptyset$.) If $u \in \phi_g(D(h))$, then let W be the component domain of $(D(h), u)$ containing $D(h)$. Since $\mathbf{I}(H)|_{D(h)}$ and $\mathbf{T}(H)|_{D(h)}$ are nonempty (by Theorem 5.6), so are $\mathbf{I}(H)|_W$ and $\mathbf{T}(H)|_W$, and again by Theorem 5.6 W must be the support of some geodesic $g' \in H$. (A special case is when $D(h)$ is an annulus with core a component of u , and then $g' = h$.) Since τ is saturated, there is a pair $(g', u') \in \tau$. We can repeat the

argument with (g', u') . Thus starting with the bottom pair (g_H, u) of τ we arrive at the unique (k, w) as claimed.

If $h = k$, then we can apply (G) to advance τ by forward or backward elementary moves to a slice containing (h, v) , and we are done.

Suppose now that $D(h)$ is properly contained in $D(k)$. Now since $w \notin \phi_k(D(h))$, either $\phi_k(D(h)) = \emptyset$, or (without loss of generality) assume $w < \min \phi_k(D(h))$. In either case, $D(h)$ intersects the first vertex of k so $k \in \Sigma^-(D(h))$ and hence $k \not\prec h$ by Theorem 5.6. This in turn rules out the possibility that $\phi_k(D(h)) = \emptyset$. Now again using property (G), we can advance τ by elementary moves until we obtain a slice τ' containing (k, x) , where $x = \min \phi_k(D(h))$. The claim above implies that there is a new pair (k', w') in τ' , with $D(h) \subseteq D(k') \subset D(k)$, and $w' \notin \phi_{k'}(D(h))$. We can therefore repeat the argument with (k', w') replacing (k, w) . Since $\xi(D(k')) < \xi(D(k))$, the process must terminate in a finite number of steps.

This proves the first statement of the lemma, that every pair (h, v) is obtained. Now consider a transition (h, v) to (h, v') where $v' = \text{succ}(v)$. In [42] we introduce a strict partial order \prec_s on $V(H)$ (this is done there for a complete hierarchy, but the proof carries through in our setting as well), which has the following properties: First, if $\tau \prec_s \tau'$ and $(h, u) \in \tau$, $(h, v) \in \tau'$, then u precedes v in h . Second, if $\tau \rightarrow \tau'$ is an elementary move, then $\tau \prec_s \tau'$.

Now if the transition (h, v) to (h, v') occurs twice, then in particular there are $i < j < k$ such that $(h, v) \in \tau_i$, $(h, v') \in \tau_j$, and then again $(h, v) \in \tau_k$. But this is a contradiction to the two properties of \prec_s . \square

5.3. Projections and lengths. Lemma 6.2 of Masur-Minsky [42] states the following, for any hierarchy H .

LEMMA 5.9. *If Y is any essential subsurface in S and*

$$d_Y(\mathbf{I}(H), \mathbf{T}(H)) > M_2,$$

then Y is the support of a geodesic h in H .

Conversely if $h \in H$ is any geodesic with $Y = D(h)$,

$$||h| - d_Y(\mathbf{I}(H), \mathbf{T}(H))| \leq 2M_1.$$

The constants M_1 and M_2 depend only on S .

The proof of this lemma goes through in the infinite setting as well. The implications of the lemma are that it is possible to detect, just from the initial and terminal marking, which subsurfaces in S participate “strongly” in the hierarchy, and how long their supported geodesics are. In fact a slightly more refined fact is shown (in the course of the proof of this lemma):




Figure 5. A twist move in $S_{0,5}$.




Figure 6. A flip move in $S_{0,5}$.

LEMMA 5.10. *If $h \in H$ is any geodesic with $Y = D(h)$, then*

$$d_Y(\mathbf{I}(h), \mathbf{I}(H)) \leq M_1,$$

$$d_Y(\mathbf{T}(h), \mathbf{T}(H)) \leq M_1.$$

Thus the endpoints of the geodesics are determined up to bounded error, also. This implies, for annulus geodesics, a statement for *signed* lengths (see (4.5)):

LEMMA 5.11. *If Y is an annulus supporting a geodesic h in H , then*

$$|\text{tw}_Y(\mathbf{I}(H), \mathbf{T}(H)) - [h]| \leq M_3$$

where the constant M_3 depends only on S .

5.4. Elementary moves on clean markings. Let μ be a maximal clean marking, which we recall is determined by a pants decomposition (u_i) and a curve \bar{t}_i for each u_i which intersects u_i in a standard way and meets none of the other u_j . A *twist move* on μ changes one \bar{t}_i by a Dehn twist, or half-twist, around u_i , preserving all the other curves.

A *flip move* interchanges a pair u_i and \bar{t}_i , and adjusts the remaining \bar{t}_j by a surgery so that they are disjoint from \bar{t}_i . Figures 5 and 6 illustrate these moves in $S_{0,5}$. For more details see [42].

Any two markings are related by a sequence of elementary moves, and we let $d_{\text{el}}(\mu, \mu')$ be the length of the shortest such sequence.

In [42] we give ways of estimating $d_{\text{el}}(\mu, \mu')$ using hierarchies and subsurface projections. First let us introduce some notation. The relation

$$x \lesssim_{a,b} y$$

will denote $x \leq ay + b$, and $x \lesssim y$ denotes the same where a, b are understood as independent of the situation. Similarly let

$$x \approx_{a,b} y$$

denote $x \lesssim_{a,b} y$ and $y \lesssim_{a,b} x$. Let

$$\{x\}_K = \begin{cases} x & \text{if } x \geq K \\ 0 & \text{if } x < K \end{cases}$$

be the ‘‘threshold function’’.

Now, every maximal slice τ of a hierarchy has an associated maximal marking μ_τ . μ_τ may not be clean, but there is always a clean marking μ' with the same base and satisfying the property that for any u in $\text{base}(\mu)$ with transversal t , u has a transversal t' in μ' such that $d_{u_i}(t_i, t'_i) \leq 2$. We say that μ' is *compatible with* τ . An elementary move $\tau_1 \rightarrow \tau_2$ on maximal slices yields a bounded number of elementary moves on compatible clean markings (roughly, a move advancing a pair (h, v) gives rise to twist moves if $\xi(D(h)) = 2$, a flip move if $\xi(D(h)) = 4$, and nothing at all if $\xi(D(h)) > 4$).

Given two maximal clean markings μ, ν there exists (by Theorem 4.6 of [42]) a complete hierarchy H with $\mathbf{I}(H) = \mu$ and $\mathbf{T}(H) = \nu$. The length of any resolution of H is equal to the sum of lengths

$$|H| \equiv \sum_{h \in H} |h|,$$

and hence, by considering the sequence of compatible clean markings, we obtain an upper bound for $d_{\text{el}}(\mu, \nu)$. In fact in Theorem 6.10 of [42] we show that

$$(5.3) \quad d_{\text{el}}(\mu, \nu) \approx_{a,b} |H|$$

for a, b depending only on S .

Using the ideas of Lemma 5.9 we further deduce in [42] that $|H|$ can be estimated in terms of the set of subsurface projection distances $\{d_W(\mu, \nu)\}$, and in particular:

LEMMA 5.12 (Thm. 6.12 of [42]). *Given S there exists K_0 such that, for any $K \geq K_0$ there are a, b such that, for any pair of maximal clean markings μ, ν and*

hierarchy H connecting them,

$$(5.4) \quad d_{\text{el}}(\mu, \nu) \approx_{a,b} \sum_{W \subseteq S} \{d_W(\mu, \nu)\}_K.$$

This also follows from the counting results in Section 9.4.

5.5. Existence of infinite hierarchies. In Theorem 4.6 of [42] we showed that a hierarchy exists connecting any two finite markings. In fact with a closer look we can now obtain

LEMMA 5.13. *For any two generalized markings \mathbf{I}, \mathbf{T} of S , which do not share any infinite-leaf components, there exists a hierarchy H with $\mathbf{I}(H) = \mathbf{I}$ and $\mathbf{T}(H) = \mathbf{T}$.*

Proof. The main gap between the finite and infinite existence theorems is filled by this lemma:

LEMMA 5.14. *Let X be a surface with $\xi(X) \geq 4$. Let μ and ν be two distinct points in $\mathcal{C}_0(X) \cup \mathcal{EL}(X)$. There exists a tight geodesic g connecting μ and ν .*

In the case $\xi(X) = 2$ we need not consider vertices at infinity, and the existence of a tight geodesic in that case is trivial.

Proof of Lemma 5.14. The case where μ and ν are both vertices of $\mathcal{C}(X)$ is covered by [42, Lemma 4.5]. (Of course in this case we need not require them to be distinct.)

In a locally compact δ -hyperbolic space the other cases, which correspond to endpoints at infinity, would then follow from the finite case by a limiting argument. This limiting step requires a special argument in our setting.

Consider first the case that μ is a finite point, i.e. $\mu \in \mathcal{C}_0(X)$, and $\nu \in \mathcal{EL}(X)$ is a point at infinity. By definition there exists a sequence of points $\nu_i \in \mathcal{C}_0(X)$ such that $\nu_i \rightarrow \nu$, and we may form tight geodesics $[\mu, \nu_i]$ (not necessarily unique). Our goal is to extract a limiting ray $[\mu, \nu)$. By δ -hyperbolicity and the definition of the boundary at infinity, we have the following:

- (*) For each $R > 0$ there exists n so that for $i, j \geq n$ the initial segment of length R of in $[\mu, \nu_i]$ is within δ of $[\mu, \nu_j]$, and vice versa.

Thus, let us extend μ to a maximal clean marking \mathbf{I} , and let \mathbf{T}_i be the marking consisting just of the endpoint ν_i . By [42, Th. 4.6], there exists a hierarchy H_i with $\mathbf{I}(H_i) = \mathbf{I}$ and $\mathbf{T}(H_i) = \mathbf{T}_i$. In fact we may assume that the main geodesic of H_i is $[\mu, \nu_i]$.

We now reprise an argument used in Lemma 6.13 of [42]. Pick $R > 5\delta$, and let n be as in (*). For $j \geq n$, let ν' be a simplex of $[\mu, \nu_j]$ within distance $R/2$ of μ , and let ν be a simplex of $[\mu, \nu_n]$ such that (via (*)) $d_{\mathcal{C}(X)}(\nu, \nu') \leq \delta$. Let τ and τ' be

saturated slices of H_n and H_j whose bottom simplices are v and v' , respectively. In fact τ and τ' must be maximal for large n : Since \mathbf{I} is maximal, it intersects any subsurface, and since \mathbf{T}_j or \mathbf{T}_n are far away in $\mathcal{C}(X)$, they must intersect any subsurface $W \subset X$ for which $[\partial W]$ is at $\mathcal{C}_1(X)$ -distance 1 from v or v' . Hence (using Theorem 5.6) every component domain that arises in the construction of τ and τ' must support a geodesic, and it follows that τ and τ' are maximal. Let m and m' be maximal clean markings compatible with τ and τ' , respectively; we wish to bound the elementary-move distance $d_{\text{el}}(m, m')$.

Let J be a hierarchy with $\mathbf{I}(J) = m$ and $\mathbf{T}(J) = m'$ (J exists again by [42, Th. 4.6]). If W is a subsurface that occurs in J , then $[\partial W]$ is within δ of v . We claim that $d_W(v_j, v_n)$ is uniformly bounded, independently of $j \geq n$. Let w and w' be points on $[\mu, v_n]$ and $[\mu, v_j]$ that are at least $2\delta + 2$ further from μ than v and v' , respectively, and such that $d(w, w') \leq \delta$ (this is possible since R is large enough). We can connect w to w' with a geodesic all of whose vertices are distance at least 2 from $[\partial W]$, and hence by the Lipschitz property of π_W , we have $d_W(w, w') \leq \delta$ (if W is an annulus and $\xi(X) = 4$, then the bound is 3δ ; see §4 for a discussion of the Lipschitz property). Now $d_W(w, v_n)$ and $d_W(w', v_j)$ are each bounded by Lemma 4.1, and this gives us a bound of the form

$$(5.5) \quad d_W(v_n, v_j) = O(1).$$

It is a consequence of Lemmas 6.1 and 6.9 of [42] (see the proof of Lemma 6.7 of that paper) that given any hierarchy H and any subsurface $W \subseteq X$, the projection $\pi_W(v)$ for any vertex occurring in H , if nonempty, is in an M -neighborhood of a geodesic in $\mathcal{A}(W)$ connecting $\pi_W(\mathbf{I}(H))$ and $\pi_W(\mathbf{T}(H))$, where $M = M(X)$. Applying this to the hierarchy H_n we see that $\pi_W(m)$ is in an M -neighborhood of a geodesic in $\mathcal{A}(W)$ connecting $\pi_W(\mu)$ and $\pi_W(v_n)$ (and similarly using H_j , for $\pi_W(m')$, $\pi_W(\mu)$ and $\pi_W(v_j)$), and we conclude using (5.5) that

$$d_W(m, m') \leq d_W(\mu, v_n) + O(1).$$

This gives a uniform bound on the length of the hierarchy J for any $j \geq n$, and hence on $d_{\text{el}}(m, m')$ by (5.3).

We conclude that, fixing τ and m , there are only finitely many possibilities for m' , and in particular for v' . Thus, there are only finitely many possibilities for the initial segment of length $R/2$ of $[\mu, v_j]$, for all $j \geq n$. We can therefore take a subsequence for which these initial segments are all the same, and then increase R and make the usual diagonalization argument.

The limiting sequence must be tight (since tightness is a local property) and it must accumulate on v , since its initial segments are equal to the initial segments of $[\mu, v_j]$ for all large j (in the subsequence). Hence this is the desired geodesic $[\mu, v]$.

Now consider the case that both μ and ν are distinct points in $\mathcal{EL}(X)$. Let $\mu_i \rightarrow \mu$, and $\nu_i \rightarrow \nu$ be sequences in $\mathcal{C}_0(X)$, and form segments $[\mu_i, \nu_i]$. In this case a similar condition to (*) holds: There is a constant δ' and a sequence of points $x_i \in [\mu_i, \nu_i]$ which remain in a bounded subset of $\mathcal{C}(X)$, such that

(**) For any $R > 0$ there exists n such that for all $i \geq n$ the subsegments of $[\mu_i, \nu_i]$ of radius R centered on x_i are within δ' of each other.

The proof is left as an exercise in applying the notion of δ -thinness of triangles.

Now as before, we take hierarchies H_i with base geodesic $[\mu_i, \nu_i]$ and apply an argument similar to the previous case to argue that a subsequence converges to a biinfinite geodesic (μ, ν) . In fact this is the exact case treated in Lemma 6.13 of [42].

This concludes the proof of Lemma 5.14. \square

Continuing with the proof of Lemma 5.13, given generalized markings \mathbf{I} and \mathbf{T} with no common infinite-leaf components, we construct a hierarchy by following the argument of [42, Th. 4.6]: We begin by constructing a main geodesic g in $\mathcal{C}(S)$: If $\text{simp}(\mathbf{I})$ is nonempty we choose one of its vertices to be the initial vertex of the geodesic. If not, then \mathbf{I} is a lamination in $\mathcal{EL}(S)$ and gives a point at infinity. The same holds for \mathbf{T} , and we connect the resulting points with a tight geodesic via Lemma 5.14. We then inductively build up a sequence of “partial hierarchies” H_n . At each stage we find an “unutilized configuration” in H_n , which is a triple (W, b, f) where W is a component domain in some geodesic in H_n , $b, f \in H_n$ are tight geodesics such that $b \not\prec W \prec^d f$, but W is not the support of a geodesic in H_n . For such a domain we use Lemma 5.14 to get a tight geodesic h with $\mathbf{I}(h) = \mathbf{I}(W, b)$ and $\mathbf{T}(h) = \mathbf{T}(W, f)$, and let $H_{n+1} = H_n \cup \{h\}$.

In the finite case, this was sufficient: the process was guaranteed to terminate after finitely many steps and the last H_n is a hierarchy, because there are no more domains to fill in. In the general case, because of the possibility of infinite geodesics, the process of filling in unutilized configurations may be infinite. Thus we must order the process in such a way that the *union* of the H_n is a hierarchy, i.e. so that every unutilized configuration is filled in a finite number of steps.

To do this, we maintain an *order* on each H_n , and in addition fix a choice of basepoint v_h for each geodesic h . At the first step H_1 is just the main geodesic with an arbitrary choice of basepoint. At each step of the process, consider the first geodesic h in the order. If there are any unutilized configurations (W, b, f) with W a component domain of h , choose one minimizing the distance of $[\partial W]$ to v_h in $\mathcal{C}(D(h))$ (there are finitely many such). Construct a new tight geodesic k with $D(k) = W$ and $b \not\prec k \prec^d f$, pick an arbitrary basepoint v_k , and add k to H_n obtaining H_{n+1} , with the ordering unchanged except that h and k should now

be the last two elements. If there are no such unutilized configurations, adjust the order so that h becomes the last element.

Repeating this, we see that every time a geodesic is created or examined, it will be examined again in a finite time. The component domains are examined by order of distance from the basepoints, and hence every unutilized configuration will be filled after a finite number of steps. Thus the union of the H_n will have no unutilized configurations, and must therefore be a hierarchy. \square

5.6. Vertices, edges and 3-holed spheres. We can now consider in a little more detail how vertices appear in a hierarchy. An edge $e = [vw]$ in a geodesic $h \in H$ with $\xi(D(h)) = 4$ is called a 4-edge. We write $v = e^-$ and $w = e^+$, where $v < w$ in the natural order of h (note in this case both v and w are actually vertices, not general simplices).

We say that a vertex v “appears in H ” if it is part of a simplex in some geodesic in H . If v is not a vertex of $\text{simp}(\mathbf{I}(H))$ or $\text{simp}(\mathbf{T}(H))$, then it is called *internal*. A vertex of $\text{simp}(\mathbf{I}(H))$ which does not have a transversal in $\mathbf{I}(H)$ is called *parabolic* in $\mathbf{I}(H)$, and similarly for $\mathbf{T}(H)$.

LEMMA 5.15. *Let v be a vertex appearing in H . Then γ_v intersects $\text{base}(\mathbf{T}(H))$ if and only if there exists a 4-edge e_1 with $v = e_1^-$. Similarly γ_v intersects $\text{base}(\mathbf{I}(H))$ if and only if there exists a 4-edge e_2 with $v = e_2^+$.*

The edges e_1 and e_2 are unique.

Proof. Suppose γ_v intersects $\text{base}(\mathbf{T}(H))$. We must find a geodesic h with $\xi(D(h)) = 4$, such that v is a vertex in h but not the last, and we must show that h is unique.

The condition that v appears in H is equivalent to the condition that the annulus $A = \text{collar}(\gamma_v)$ is a component domain in H , and the condition that v intersects $\text{base}(\mathbf{T}(H))$ means that $\Sigma^+(A)$ is nonempty, since it contains the main geodesic g_H . By part (3) of Theorem 5.6, there exists a geodesic f such that $A \searrow f$, so by definition there must be some $h \in H$ with $A \searrow^d h$.

The condition $A \searrow^d h$ means that v appears as a vertex in a simplex w of h , and that $\text{succ}(w)|_A \neq \emptyset$. If w is not the last simplex, then, if $\xi(D(h)) > 4$, $\text{succ}(w)$ would be a simplex disjoint from w and hence from v , a contradiction. Thus $\xi(D(h)) = 4$, $v = w$, and $\text{succ}(w)$ is a vertex w' . This gives us our 4-edge, $e_2 = [vw']$.

We claim that w cannot be the last simplex of h . If, by way of contradiction, it is, then $\text{succ}(w) = \mathbf{T}(h)$ and $v \in \text{base}(\mathbf{T}(h))$, so $\mathbf{T}(h)$ must contain a transversal to v . If $h = g_H$, then γ_v is a base curve of $\mathbf{T}(H)$, contradicting the hypothesis of the lemma. Hence $h \searrow^d h'$ for some h' . Now since $\mathbf{T}(h)$ contains a transversal, $D(h)$ must be a component domain for the last simplex of h' , and $\mathbf{T}(h) = \mathbf{T}(h')|_{D(h)}$. Hence v must be a component of $\text{base}(\mathbf{T}(h'))$ with a transversal, and we repeat

this by induction until we terminate with the main geodesic g_H , and obtain our contradiction.

Uniqueness of e_2 is seen as follows. If there were two such edges, there would be two geodesics h, h' with $\xi = 4$, and such that $A \searrow^d h$, $A \searrow^d h'$. Hence both are in $\Sigma^+(A)$, but this contradicts part (1) of Theorem 5.6, which says that $\Sigma^+(A)$ is a sequence $f_0 \searrow^d f_1 \searrow^d \cdots$.

In the converse direction, if a 4-edge e exists with $v = e^-$, then v appears in a geodesic h_1 with $\xi = 4$, where v is not the last vertex. Thus γ_v intersects the last vertex of h_1 , and in particular it intersects $\text{base}(\mathbf{T}(h_1))$. Note that $A \searrow^d h_1$, and Theorem 5.6 gives us a sequence $h_1 \searrow^d h_2 \searrow^d \cdots \searrow^d g_H$. We claim by induction that γ_v intersects $\text{base}(\mathbf{T}(h_i))$ for each i . Let $w_i = \max \phi_{h_i}(D(h_{i-1}))$, so that $D(h_{i-1})$ is a component domain of $(D(h_i), w_i)$. By Lemma 5.5, w_i must equal $\max \phi_{h_i}(A)$. If w_i is not the last simplex, then the last simplex of h_i intersects γ_v , and hence does $\text{base}(\mathbf{T}(h_i))$. If w_i is the last simplex, then $\mathbf{T}(h_{i-1}) = \mathbf{T}(h_i)|_{D(h_{i-1})}$, so that the statement for h_i follows from the statement for h_{i-1} . Hence γ_v intersects $\text{base}(\mathbf{T}(H))$, which is what we wanted to show.

The argument for e_1 is identical, with directions reversed. \square

Now given v appearing in a hierarchy H , and fixing a resolution $\{\tau_i\}_{i \in \mathcal{J}}$ of H , with \mathcal{J} a subinterval of \mathbb{Z} , let

$$J(v) = \{i \in \mathcal{J} : v \in \text{base}(\mu_{\tau_i})\}.$$

The following fact will be fundamental for us:

LEMMA 5.16. *Let H be a 4-complete hierarchy and $\{\tau_i\}_{i \in \mathcal{J}}$ a resolution. If v is a vertex in H , then $J(v)$ is an interval in \mathbb{Z} .*

Recall that an interval can be finite, one-sided infinite, or bi-infinite.

Proof. The only type of elementary move that can change the vertex set of a slice is a move along a 4-edge e , which replaces one vertex e^- by its successor e^+ . By Lemma 5.8, there is exactly one step in the resolution which advances along any given edge. By Lemma 5.15, there is at most one 4-edge, e_1 , for which $e_1^+ = v$, and at most one 4-edge, e_2 , for which $e_2^- = v$. Thus there is at most one move $\tau_{i-1} \rightarrow \tau_i$ at which v appears, and at most one move $\tau_j \rightarrow \tau_{j+1}$ at which v disappears. The possibilities for $J(v)$, depending on the existence of e_1 and e_2 , are therefore $[i, \sup \mathcal{J}]$, $[\inf \mathcal{J}, j]$, \mathcal{J} , $[i, j]$ (if $i \leq j$), or $[\inf \mathcal{J}, j] \cup [i, \sup \mathcal{J}]$ (if $j < i$). Only the last of these fails to be an interval and must be ruled out (note that $\inf \mathcal{J}$ may be $-\infty$ and $\sup \mathcal{J}$ may be ∞ , yielding infinite intervals).

In the last case where $j < i$, v must be in all slices after τ_i . We claim that v is therefore in $\text{simp}(\mathbf{T}(H))$. Let (g_H, w_k) be the bottom pair of τ_k . For $k > i$ we have $d_{\epsilon_1(S)}(w_k, v) \leq 1$, and on the other hand Lemma 5.8 implies that all vertices

of g_H are obtained in the resolution – hence g_H is finite in the forward direction, and for sufficiently large k we must have w_k equal to the last vertex of g_H , which is in $\text{simp}(\mathbf{T}(g_H)) = \text{simp}(\mathbf{T}(H))$. If $v \neq w_k$, then it is contained in a component W of $S \setminus \text{collar}(w_k)$, which must therefore support a geodesic g' (otherwise the slice τ_k cannot have any vertices in W), and we have $\mathbf{T}(g') = \mathbf{T}(g_H)|_W$. The same argument implies that g' is finite in the forward direction, and we proceed inductively. The process must terminate, and at that point v must be a vertex of $\text{simp}(\mathbf{T}(H))$.

However Lemma 5.15 then implies that v never appears as e^- for a 4-edge e , and this is a contradiction. This rules out the case $j < i$. \square

Three-holed spheres. The following result for three-holed spheres is analogous to Lemma 5.15.

LEMMA 5.17. *Let Y a component domain of a geodesic in H , and suppose $\xi(Y) = 3$. If $\mathbf{T}(H)|_Y \neq \emptyset$, then there exists a unique geodesic $f \in H$ with $\xi(D(f)) = 4$ and $Y \searrow^d f$. Similarly, if $\mathbf{I}(H)|_Y \neq \emptyset$, then there exists a unique geodesic $b \in H$ with $\xi(D(b)) = 4$ and $b \swarrow^d Y$.*

Note that $\mathbf{T}(H)|_Y \neq \emptyset$ just means that Y has an essential intersection with a base curve of $\mathbf{T}(H)$. In particular, if Y intersects both $\text{base}(\mathbf{T}(H))$ and $\text{base}(\mathbf{I}(H))$, then the lemma gives b, f with $\xi = 4$ and $b \swarrow^d Y \searrow^d f$. We call this a *gluing configuration*, and it will be used in the model manifold construction in Section 8.

Proof. If $\mathbf{T}(H)|_Y \neq \emptyset$, part (3) of Theorem 5.6 gives some f with $Y \searrow^d f$. Thus it only remains to verify that $\xi(D(f)) = 4$.

The condition $Y \searrow^d f$ means that there is a simplex w of f , with Y a component domain of $(D(f), w)$, and such that $\text{succ}(w)|_Y \neq \emptyset$. Suppose that w is not the last vertex of f . If $\xi(D(f)) > 4$, then $\text{succ}(w)$ is disjoint from w , and since a three-holed sphere can have no nontrivial nonperipheral simple curves, $\text{succ}(w)|_Y = \emptyset$. This contradiction implies $\xi(D(f)) = 4$.

If w is the last vertex, then $w \subset \text{base}(\mathbf{T}(f))$ and $\text{succ}(w) = \mathbf{T}(f)$. Since Y is a three holed sphere bounded by w in $D(f)$, it cannot contain other components of $\text{base}(\mathbf{T}(f))$, so we obtain $\text{succ}(w)|_Y = \emptyset$, again a contradiction.

Uniqueness of f follows from the sequential structure of $\Sigma^+(Y)$ (part (1) of Theorem 5.6). The backwards case is proved in the same way. \square

6. The coarse projection property

In this section we will begin the geometric argument that connects the geometry of a Kleinian surface group $\rho : \pi_1(S) \rightarrow \text{PSL}_2(\mathbb{C})$ with the structure of the complex of curves $\mathcal{C}(S)$. The main tool for this is the “short-curve projection” Π_ρ , which takes any element $a \in \mathcal{A}(S)$ to the set of short curves on pleated surfaces in

N mapping the curves associated to a geodesically. This projection was studied in [48] and [49], where it was shown to have properties analogous to a nearest-point projection to a convex set. Here we will refine these to obtain *relative* statements about projections π_Y to subsurfaces. In Section 7, we will use Π_ρ to generate our main geometric estimates.

We proceed to define the projection map. Fix $L \geq L_0$, where L_0 is Bers' constant (see §3.2). As in [49], we define for a hyperbolic structure σ on $\text{int}(S)$

$$\mathbf{short}_L(\sigma) = \{\alpha \in \mathcal{C}_0(S) : \ell_\sigma(\alpha) \leq L\}.$$

For any lamination $\lambda \in \mathcal{GL}(S)$, we define

$$\Pi_{\rho,L}(\lambda) = \bigcup_{f \in \text{pleat}_\rho(\lambda)} \mathbf{short}_L(\sigma_f).$$

(If L is understood we write just Π_ρ .) In particular any $x \in \mathcal{A}(S)$ determines a lamination λ_x consisting of leaves representing the vertices of the smallest simplex containing x , and we define $\Pi_{\rho,L}(x) \equiv \Pi_{\rho,L}(\lambda_x)$. This gives us a map

$$\Pi_{\rho,L} : \mathcal{A}(S) \rightarrow \mathcal{P}(\mathcal{C}(\rho, L))$$

where $\mathcal{P}(X)$ denotes the set of subsets of X .

We will establish the following ‘‘coarse projection’’ property, generalizing Lemma 3.2 of [49]:

THEOREM 6.1. *For a surface S and a constant $L \geq L_0$ there exist $D_0, D_1 > 0$ such that, if $\rho \in \mathfrak{D}(S)$ and $Y \subseteq S$ is an essential subsurface with $\xi(Y) \neq 3$, then:*

(P1) (*Coarse definition*). *If v is any vertex in $\mathcal{A}(S)$ with $\pi_Y(v) \neq \emptyset$, then*

$$\text{diam}_Y(\Pi_{\rho,L}(v)) \leq D_0.$$

(P2) (*Relative Coarse Lipschitz*). *If v, w are adjacent vertices in $\mathcal{A}(S)$ with $\pi_Y(v) \neq \emptyset$ and $\pi_Y(w) \neq \emptyset$, then*

$$\text{diam}_Y(\Pi_{\rho,L}(v) \cup \Pi_{\rho,L}(w)) \leq D_1.$$

(P3) (*Relative Coarse Idempotence*). *If $v \in \mathcal{C}(\rho, L)$ and $\pi_Y(v) \neq \emptyset$, then*

$$d_Y(v, \Pi_{\rho,L}(v)) = 0.$$

Note that property (P2) translates to a weak sort of Lipschitz condition for the map $\pi_Y \circ \Pi_{\rho,L}$: If v_0, v_1, \dots, v_n is a sequence of vertices in $\mathcal{A}(S)$ with v_i, v_{i+1} adjacent such that $\pi_Y(v_i) \neq \emptyset$, and in addition $\pi_Y(\Pi_{\rho,L}(v_i)) \neq \emptyset$, then repeated application of (P2) yields

$$\text{diam}_Y(\Pi_{\rho,L}(v_0) \cup \Pi_{\rho,L}(v_n)) \leq D_1 n.$$

The issue of when $\Pi_{\rho,L}$ is nonempty will be addressed in Lemma 7.4.

6.1. Train tracks. Before proceeding with the proof of Theorem 6.1, let us digress a bit to discuss the construction and properties of train tracks. Thurston [60] first introduced the notion of a train track in a surface as a finite approximation to a geodesic lamination. Although it has by now become a standard tool in the field, we will attempt to be careful with details here, especially because of the presence of thin parts.

Definitions (See Penner-Harer [53]). A *train track* in a surface S is an embedded 1-complex τ with a special structure at the vertices. It is convenient to describe this in terms of a foliation of a small neighborhood of τ by leaves (called “ties”) that are transverse to τ at every point. A tie passing through a vertex locally divides the ends of adjacent edges according to which side of the tie they are on. One thinks of these as “incoming” and “outgoing” sides, and both sets must be nonempty. Vertices of a train track are called “switches” and edges are called “branches.”

Two branch ends coming in to the same side of a switch bound between them a corner of a complementary region which we call a “cusp”. If a switch meets n branch ends, then a small neighborhood of the switch is cut by τ into n corners, $n - 2$ of which are cusps. The other two are called smooth. A further condition usually imposed on the train track is that each complementary region is *hyperbolic*, in the sense that it is not a disk with 2 or fewer boundary cusps, a once punctured disk with no boundary cusps, or an annulus with no boundary cusps.

When $\partial S \neq \emptyset$, we will also allow train tracks to have branches which terminate in the boundary. A regular neighborhood of a boundary component is cut by τ into regions which we think of as cusps for the purpose of the hyperbolicity condition.

A *train route* in τ is an immersion of a 1-manifold into τ which always traverses switches from one side to the other, and whose endpoints (if any) map to ∂S . An element of $\mathcal{A}(S)$ is *carried* by τ if it can be represented by a train route. If α is carried in τ it imposes a *measure* on its branches counting how many times each is traversed by α . For each branch b we denote this measure by $\alpha(b)$. We let

$$\ell_\tau(\alpha) = \sum_b \alpha(b)$$

denote the *combinatorial length* of α .

A train track in $\text{int}(S)$ is just the restriction of a train track in S to the interior. In particular branches that terminate in ∂S become branches that exit cusps of $\text{int}(S)$.

Collapsing curves to train tracks. Fix a Margulis constant $\bar{\epsilon} > 0$, and let σ be a complete (finite-area) hyperbolic metric on $\text{int}(S)$. If γ is a geodesic representative of a vertex in $\mathcal{A}(S)$, or in general any geodesic lamination, and $0 < \epsilon < \bar{\epsilon}$, we define

an ϵ -collapse of γ to a train track to be a map

$$q : \text{int}(S) \rightarrow S'$$

that is homotopic to a homeomorphism, and such that:

- (1) $q(\gamma)$ is a train track τ in S' , and q restricted to any leaf of γ is a train route.
- (2) There is a metric on τ such that $q|_\gamma$ is a local isometry.
- (3) For each $x \in \tau$ the preimage $q^{-1}(x)$ is a point or arc. If $q^{-1}(x)$ touches $S_{[\bar{\epsilon}, \infty)}$, then its length is at most ϵ .

LEMMA 6.2. *Given a surface S and $\bar{\epsilon} > 0$, for each $\epsilon < \bar{\epsilon}$, for any hyperbolic metric σ on S and any geodesic lamination γ there is an ϵ -collapse of γ to a train track. The total length of*

$$\tau_{\bar{\epsilon}} \equiv q(\gamma \cap S_{[\epsilon, \infty)})$$

is bounded by a constant K depending only on ϵ .

Proof. A version of this construction is discussed in Thurston [60, §8.9]. See Brock [17] for a complete discussion in the case without thin parts.

First, for simplicity let us add enough leaves to γ to obtain a lamination λ whose complement is a union of (interiors of) ideal hyperbolic triangles. Fixing a small ϵ , let \mathcal{F}_ϵ be the foliation of the ends of each of these triangles by horocyclic arcs of length less than ϵ . Thus \mathcal{F}_ϵ is supported in an open subset of $S \setminus \lambda$ we will call U_ϵ . Because the tangent directions of the horocycles form a Lipschitz line field, \mathcal{F}_ϵ may be extended across the leaves of λ to be a foliation of the interior of the closure \overline{U}_ϵ . Let us continue to call this foliation \mathcal{F}_ϵ . In fact \mathcal{F}_ϵ extends to give a decomposition of all of \overline{U}_ϵ into leaves, where the boundary consists of endpoints of leaves or boundary horocycles of length ϵ .

This foliation inherits a transverse measure from the length measure along leaves of λ . This is evident from the fact that, in each end of a triangle, the horocyclic flow preserves length for its family of orthogonal geodesics.

Let us establish the following claims:

- (1) Each cusp of S has a neighborhood that is foliated by closed leaves of \mathcal{F}_ϵ .
- (2) All leaves of \mathcal{F}_ϵ are compact.
- (3) There are constants C, C' such that, if $\epsilon < C\bar{\epsilon}$, then any leaf meeting the $\bar{\epsilon}$ -thick part of S is an arc, and has length bounded by $C'\epsilon$.

Proof. Claim (1) is fairly clear: There is a uniform δ such that the δ -Margulis tube associated to a cusp avoids all simple geodesics except those which go vertically all the way up the cusp. Since the complement of λ is a union of ideal triangles it penetrates every such neighborhood. Hence if we choose $\delta < \epsilon/2$, say,

the leaves of \mathcal{F}_ϵ in this neighborhood are complete horocycles going around the cusp.

Let a *foliated rectangle* of \mathcal{F}_ϵ denote a region on which \mathcal{F}_ϵ is equivalent to the foliation of $[0, 1]^2$ by vertical arcs, and for which the horizontal arcs are geodesics orthogonal to \mathcal{F}_ϵ (e.g. leaves of λ). The leaves of \mathcal{F}_ϵ identified with $\{0\} \times [0, 1]$ and $\{1\} \times [0, 1]$ are its *boundary leaves*, and the distance along λ between them is called the *width* $w(R)$. The lengths of the leaves vary by a factor of at most e^w . Note that the area of R is at least the smaller leaf length times the width.

We can make the same definitions for the lift $\tilde{\mathcal{F}}_\epsilon$ of \mathcal{F}_ϵ to the universal cover \mathbb{H}^2 . Let R be a foliated rectangle in the universal cover, with boundary leaves l_0, l_1 . Suppose that l_0 and l_1 can be extended indefinitely in one direction (say the upward vertical direction in the identification with $[0, 1]^2$). There is an a depending only on \mathcal{F}_ϵ such that, provided $w(R) < a$, the whole rectangle can be extended indefinitely in the upward direction – that is the region between the extensions of l_0 and l_1 can be foliated to make a foliated rectangle in \mathbb{H}^2 for arbitrarily long extensions. To see this, choose a less than the minimal length of an arc of λ in the boundary of a complementary region of $\tilde{\mathcal{F}}_\epsilon$. Suppose R has been extended to a rectangle R' with boundary leaves l'_0 and l'_1 . Since $w(R') = w(R) < a$, and by assumption l'_0 and l'_1 can be extended further, the top boundary of R' does not lie on the boundary of a complementary region, and hence the rectangle can be extended further.

This gives at least an immersed foliated rectangle. However this must in fact be an embedding (in \mathbb{H}^2), since a self-intersection would give rise to a disk whose boundary consists of one arc of $\tilde{\mathcal{F}}_\epsilon$ and one arc transverse to $\tilde{\mathcal{F}}_\epsilon$ – violating the index formula for line fields since \mathcal{F}_ϵ has an extension to a foliation of S with negative-index singularities.

Now suppose that l is a noncompact leaf of \mathcal{F}_ϵ . In view of (1), l must accumulate somewhere in S . Hence there must be two segments l_0 and l_1 close enough together to bound a foliated rectangle R , and such that l_0 and l_1 extend infinitely in one direction. In the universal cover, the lift of R must extend indefinitely by the above argument. The extension $R_\infty \cong [0, 1] \times [0, \infty)$ has infinite area, so its immersed image downstairs cannot be embedded. It follows that it is an annulus in which the half-leaf $\{0\} \times [0, \infty)$ (which is part of l) either spirals or maps to a closed leaf. In the former case the annulus has infinite area, and in the latter case l is a closed leaf after all. We conclude that there are no noncompact leaves. This proves claim (2).

For claim (3), if x is in the $\bar{\epsilon}$ thick part it is the center of an embedded $\bar{\epsilon}/2$ -disk. Consider the intersection m of a leaf of \mathcal{F}_ϵ passing through x with a ball of radius $\bar{\epsilon}/4$ around x . Every interval of $m \setminus \lambda$ is the boundary leaf of a foliated rectangle on one side or the other, of width at least $\bar{\epsilon}/4$, these rectangles are all disjoint, and hence m is in the boundary of a region of \mathcal{F}_ϵ with area at least $C|m|$, for a uniform




Figure 7. The local picture of \mathcal{F}'_ϵ , collapsing to a train track.

C (see Brock [17] for details). On the other hand the total area of \mathcal{F}_ϵ is at most $N\epsilon$, where N is the number of vertices of ideal triangles in the lamination, by an easy computation (and using the fact that $\text{area}(\lambda) = 0$; see [25]). It follows that $|m|$ is bounded by $C'\epsilon$ where C' depends on $\bar{\epsilon}$. In particular, if ϵ is sufficiently small, then the leaf containing m must terminate before reaching the boundary of the $\bar{\epsilon}/4$ disk, and claim (3) follows. \square

In order to get rid of the closed leaves of \mathcal{F}_ϵ , we make the following adjustment. If A is a cusp neighborhood foliated by closed leaves, then A is cut by leaves of λ into one or more cusp regions, each equal to an end of one of the complementary triangles of λ . Pick one of these regions and erase all the arcs of \mathcal{F}_ϵ inside the corresponding end. This will cut all the loops in A into arcs.

If A is the collar of a geodesic with length less than $\bar{\epsilon}$, and contains closed leaves of \mathcal{F}_ϵ , then there are leaves of λ passing all the way through A . Choose one complementary region R of $A \setminus \lambda$ – it is contained in the cusp end of a triangle in $S \setminus \lambda$. For that end choose ϵ' sufficiently small that, if we remove all the \mathcal{F}_ϵ arcs of length at least ϵ' in the cusp end containing R , then this will remove all the arcs in R , again cutting all the closed leaves.

By claim (3) there are no other closed leaves in \mathcal{F}_ϵ .

Call the new foliation \mathcal{F}'_ϵ . Now consider the quotient $q : S \rightarrow S' = S/\mathcal{F}'_\epsilon$ obtained by identifying each leaf of \mathcal{F}'_ϵ to a point. We claim this is our desired map, if ϵ is taken sufficiently small. Since all leaves of \mathcal{F}'_ϵ are arcs, one can see by Moore's theorem [50], [51] that S' is a surface homeomorphic to S and that q is homotopic to a homeomorphism. In fact if x is a point on a leaf of \mathcal{F}'_ϵ , we can explicitly describe a neighborhood of $q(x)$ by considering a neighborhood of the leaves passing near x (Figure 7).

The image $\tau = q(\lambda)$ in S' is a finite graph (with some edges going out the punctures). To see the train track structure on τ , extend the foliation \mathcal{F}'_ϵ slightly to obtain, in the image, a foliation in a neighborhood of τ transverse to τ . An interior leaf a of \mathcal{F}'_ϵ maps to an interior point of τ . This is because a terminates on two leaf segments of λ , and leaves sufficiently close to a terminate on the same segments;

hence producing a foliated rectangle that maps to a segment in τ . A boundary leaf b of \mathcal{F}_ϵ maps to a switch of τ : The leaves on either side of b map to arcs of τ on either side of $q(b)$, hence producing the switch structure.

The complementary regions of τ are all three-cusped disks, obtained directly from the triangular components of $S \setminus \lambda$ by the quotient map.

A branch b of τ is the q -image of a foliated rectangle R_b in S . The lengths of λ -leaves in R_b induce a metric on b with total length $w(R_b)$. Two leaves of λ in R_b lie on the boundary, and are also boundary arcs of triangles in $S \setminus \lambda$. The intersection of these boundary arcs with $S_{[\epsilon, \infty)}$ has total length bounded by some $K(\epsilon)$ (proportional to $\log 1/\epsilon$). This gives a bound on the total length of $\tau_{\bar{\epsilon}}$.

Now any point x of τ has preimage $q^{-1}(x)$ which is an arc of \mathcal{F}'_ϵ , with both endpoints on λ . If $q^{-1}(x)$ meets the $\bar{\epsilon}$ -thick part of S , then its length is bounded by $C'\epsilon$, as in the proof of claim (3).

Thus to obtain an ϵ -collapse, we take the quotient $S/\mathcal{F}'_{\epsilon/C'}$, and then restrict to the sub-track $\tau_\gamma = q(\gamma)$, noting that all the properties of a collapse are inherited by this track. \square

Nice representatives. If τ is obtained from γ by an ϵ -collapse q , and β is a curve carried on τ (more accurately, $q(\beta)$ is carried on τ), then β can be realized after homotopy as a chain of segments

$$\beta^\tau = \beta_1 * t_1 * \cdots * \beta_n * t_n,$$

where $n = \ell_\tau(\beta)$, β_i are arcs of γ and each t_i is a subarc of $q^{-1}(s)$ for a switch s . This is done by first choosing, for each branch b of τ , an arc of γ in $q^{-1}(b)$ which maps isometrically to b , and then for branches meeting at a switch, connecting their preimage arcs with the preimage of the switch. The total length of β^τ in the $\bar{\epsilon}$ -thick part of S is bounded by

$$(6.1) \quad \ell(\beta^\tau \cap S_{[\bar{\epsilon}, \infty)}) \leq (K(\epsilon) + \epsilon)\ell_\tau(\beta)$$

since each t_i has length at most ϵ and each β_i has length at most $K(\epsilon)$ in the thick part.

Short curves with prescribed intersection properties. This lemma will be used in the final stages of the proof in Section 6.2:

LEMMA 6.3. *For any compact surface S there is a number A such that, for any train track τ in S , if $\beta \in \mathcal{C}_0(S)$ has essential intersection with some $\gamma \in \mathcal{A}_0(S)$ carried on τ , then there exists $\alpha \in \mathcal{A}_0(S)$ carried on τ with*

$$\ell_\tau(\alpha) \leq A,$$

which also intersects β essentially.

Proof. Since, up to homeomorphisms of S , there are only finitely many train tracks in S (see e.g. Penner-Harer [53]), and since for any finite set of curves or arcs carried in τ there is an upper bound on ℓ_τ , it will suffice to prove the following statement: For any train track τ there is a finite set $\{\alpha_1, \dots, \alpha_k\} \subset \mathcal{A}_0(S)$ carried in τ so that, for any $\beta \in \mathcal{C}_0(S)$, if $i(\beta, \alpha_j) = 0$ for each j , then $i(\beta, \gamma) = 0$ for every $\gamma \in \mathcal{A}_0(S)$ carried on τ .

For any set X in $\mathcal{A}_0(S)$ there is an essential subsurface S_X filled by X , which is unique up to isotopy. To form S_X place representatives of X in minimal position (for example using geodesic representatives in a hyperbolic metric on S), form a regular neighborhood, and adjoin all complementary components which are disks or annuli with one boundary component in ∂S . If $\beta \in \mathcal{C}_0(S)$ intersects S_X essentially, then β must intersect one of the elements of X essentially.

Let $X(\tau)$ be the set of elements of $\mathcal{A}_0(S)$ carried on τ , and let $X_1 \subset X_2 \subset \dots$ be an exhaustion of $X(\tau)$ by finite sets. An Euler characteristic argument implies that the sequence $S_{X_1} \subseteq S_{X_2} \subseteq \dots$ is eventually constant (up to isotopy), and hence there exists some k for which $S_{X_k} = S_{X(\tau)}$.

X_k is therefore our desired finite set $\{\alpha_i\}$. □

(Note there ought to be a more concrete proof of this lemma, which gives estimates on A , but I have not found it.)

6.2. Proof of coarse projection. We can now proceed with the proof of Theorem 6.1.

Property (P3). This property is immediate: v is realized with its minimal ρ -length in any $f \in \mathbf{pleat}_\rho(v)$, and since this is at most L we have $v \in \mathbf{short}_L(\sigma_f) \subset \Pi_\rho(v)$. Since $\pi_Y(v)$ is not empty, it is also a component of $\pi_Y(\Pi_\rho(v))$ which establishes the property.

Property (P1) \implies Property (P2). Suppose first that $\xi(Y) \geq 4$. Then Y has essential intersection with any pants decomposition. Since $L \geq L_0$, $\Pi_\rho(x)$ contains a pants decomposition, and hence $\pi_Y(\Pi_\rho(x))$ is nonempty for any lamination x . Now if v and w are adjacent in $\mathcal{A}_0(S)$, the disjoint union of their representatives is a lamination which we denote $v \cup w$. We have

$$\pi_Y(\Pi_{\rho,L}(v \cup w)) \subset \pi_Y(\Pi_{\rho,L}(v)) \cap \pi_Y(\Pi_{\rho,L}(w)),$$

and in particular this intersection is nonempty. Thus Property (P1) for v and w immediately yields (P2), with $D_1 = 2D_0$.

Now suppose $\xi(Y) = 2$, so Y is an annulus. Let $u \in \mathcal{C}_0(S)$ represent the core curve of Y . If either $\pi_Y(\Pi_{\rho,L}(v))$ or $\pi_Y(\Pi_{\rho,L}(w))$ are empty there is nothing to prove; so assume they are nonempty. In particular for some $f \in \mathbf{pleat}_\rho(v)$, γ_u is intersected by a curve of σ_f -length at most L , and hence $\ell_{\sigma_f}(u) \geq \varepsilon(L)$. It follows

that $f(S)$ must stay out of the tube $\mathbb{T}_{\epsilon'}(u)$, where $\epsilon' = \epsilon_T(\epsilon(L))$ (see §3.2.2 and §3.2.3 for the definitions of ϵ_T and ϵ). Now for any $h \in \mathbf{pleat}_\rho(v \cup w)$, since it also maps γ_v to its geodesic representative, it follows that $\ell_{\sigma_h}(u) \geq \epsilon'$. Therefore there is a curve intersecting γ_u of σ_h -length at most $L' \equiv \max(L, \mathcal{L}(\epsilon'))$ (see §3.2.2), and thus $\pi_Y(\Pi_{\rho, L'}(v \cup w))$ is nonempty. As in the $\xi(Y) \geq 4$ case we therefore obtain a bound on

$$\text{diam}_Y(\Pi_{\rho, L'}(v) \cup \Pi_{\rho, L'}(w))$$

and since this contains $\Pi_{\rho, L}(v) \cup \Pi_{\rho, L}(w)$, we obtain Property (P2).

Proof of Property (P1). Let $\gamma = \gamma_v$. Since $\Pi_\rho(v)$ is the union of $\mathbf{short}(\sigma_f)$ for $f \in \mathbf{pleat}_\rho(\gamma)$, it suffices to bound

$$\text{diam}_Y(\mathbf{short}(\sigma_f) \cup \mathbf{short}(\sigma_g))$$

for any two $f, g \in \mathbf{pleat}_\rho(\gamma)$.

Note first that there is a $K(L)$ such that

$$(6.2) \quad \text{diam}_Y(\mathbf{short}_L(\sigma)) \leq K$$

for any hyperbolic metric σ on S . This is because two curves of bounded length in the same metric have a bound on their intersection number, and this readily gives a bound on the distance between the corresponding curve systems in $\mathcal{A}(Y)$. From now on we will assume that $\pi_Y(\mathbf{short}_L(\sigma_f))$ and $\pi_Y(\mathbf{short}_L(\sigma_g))$ are nonempty, since otherwise there is nothing to prove (this is of consequence only if Y is an annulus – in other cases it holds automatically).

Thus our goal will be to find a curve α , intersecting Y essentially, whose length in *both* σ_f and σ_g is bounded by some *a priori* $L'' \geq L$. This would give us a nonempty intersection $\pi_Y(\mathbf{short}_{L''}(\sigma_f)) \cap \pi_Y(\mathbf{short}_{L''}(\sigma_g))$, and hence by (6.2) bound the diameter of the union.

Bounding bridge arcs. Let us invoke some machinery developed in [48]. First, by Lemma 3.3 of [48], up to precomposing g with a homeomorphism isotopic to the identity, we may (and will) assume that f and g are *homotopic rel γ* : this means that $f|_\gamma = g|_\gamma$ and f and g are homotopic keeping the points of γ fixed. A *bridge arc* for γ is an arc with endpoints on γ , which is not homotopic rel endpoints into γ .

We have the following lemma, which is part (1) of Lemma 3.4 of [48]:

LEMMA 6.4. *Fixing $\hat{\epsilon}$, for any δ_1 there exists $\delta_0 \in (0, \delta_1)$ such that, if $f, g \in \mathbf{pleat}_\rho(\gamma)$ and are homotopic rel γ , then for any bridge arc t for γ in the $\hat{\epsilon}$ -thick part of σ_f we have*

$$\ell_{\sigma_f}(t) \leq \delta_0 \implies \ell_{\sigma_g}(t) \leq \delta_1.$$

This lemma is a fairly direct consequence of Thurston's Uniform Injectivity Theorem [62], which uses a compactness argument on the space of all pleated maps

to conclude that different leaves of a pleating lamination cannot line up too closely in the image, in a uniform sense.

Now let $\epsilon_2 = \min(\epsilon_1/2, \varepsilon(L))$. Let Γ denote the system of simple closed σ_f -geodesics in S whose σ_f -lengths are less than ϵ_2 , and which intersect γ essentially (on a first reading one should consider the case that σ_f is ϵ_2 -thick, and in particular $\Gamma = \emptyset$). For each component α of Γ , its f -image must be in the ϵ_2 thin part of N . Since α crosses γ , and f and g are homotopic rel γ , $g(\alpha)$ meets the ϵ_2 thin part of N as well. By the choice of ϵ_1 (see §3.2.2), the σ_g -length of the σ_g -geodesic representative of each component of Γ is bounded by ϵ_0 . Hence if Y crosses Γ essentially, we may let α be a component of Γ intersecting Y , and we are done.

Assume therefore that Y has no essential intersections with Γ . If $\xi(Y) \geq 4$, Y must be an essential subsurface in a component R of $S \setminus (\mathbf{collar}(\Gamma, \sigma_f) \cup \mathbf{collar}(\partial S, \sigma_f))$. If $\xi(Y) = 2$, then there is the possibility that Y is the collar of a component of Γ . However, since $\pi_Y(\mathbf{short}_L(\sigma_f)) \neq \emptyset$, there is a curve of length at most L crossing Y , and because $\epsilon_2 \leq \varepsilon(L)$ this is not true for any component of Γ . Thus Y is a nonperipheral annulus in R .

Train track. We now apply Lemma 6.2 from Section 6.1 to find a good train-track approximation of γ .

Setting $\hat{\epsilon} = \epsilon_2/2$ and $\delta_1 = \epsilon_2$, let δ_0 be the constant provided by Lemma 6.4. Now let $\epsilon = \min(\epsilon_2/4, \delta_0)$. Setting the Margulis constant $\bar{\epsilon}$ in Lemma 6.2 to ϵ_2 , we obtain an ϵ -collapse $q : \text{int}(S) \rightarrow S'$ (with respect to the metric σ_f) where S' is homeomorphic to $\text{int}(S)$, taking γ to a train track τ .

Now since τ carries γ and γ has an essential intersection with some curve β in Y (possibly a component of ∂Y), it follows from Lemma 6.3 that there exists α' carried in τ with $\ell_\tau(\alpha') \leq A$ (with A a bound depending only on S), such that α' also intersects β essentially. We will use α' to construct our curve of bounded length in both σ_f and σ_g .

Bridge arcs and nice representatives. Using the discussion in Section 6.1, α' is homotopic to

$$\alpha^\tau = \alpha_1 * t_1 * \cdots * \alpha_n * t_n,$$

where $n = \ell_\tau(\alpha') \leq A$, each α_i runs along γ , and each t_i is in the q -preimage of a switch. Lemma 6.2 gives us a uniform bound $K(\epsilon)$ on the length of the intersection of each α_i with the ϵ_2 -thick part of (S, σ_f) , and a bound of ϵ on each t_i that meets the ϵ_2 -thick part. We also note that, by definition of Γ , no α_i crosses a thin collar of a curve of σ_f -length $\leq \epsilon_2$, unless that curve is in Γ . It follows that $\alpha_R = \alpha^\tau \cap R$ remains in the ϵ_2 -thick part, and hence has uniformly bounded length. The endpoints of α_R lie in ∂R , which is part of $\partial \mathbf{collar}(\Gamma, \sigma_f) \cup \partial \mathbf{collar}(\partial S, \sigma_f)$. Since α' intersects Y essentially and Y is an essential subsurface of R , there must be some arc a of α_R that intersects Y essentially.




Figure 8. The surgery yielding the final curve α from the arc a . The two possible configurations are shown.

We already have a uniform bound on $\ell_{\sigma_f}(a)$. To view this arc in σ_g , note the following: Each arc α_i has exactly the same length in σ_g as in σ_f . Each t_i is a bridge arc for γ , and since the ones occurring in a touch the ϵ_2 -thick part of σ_f , their length is bounded by ϵ . Since $\epsilon < \epsilon_2/4$, the entire t_i is contained in the $\epsilon_2/2$ thick part. Now since $\epsilon \leq \delta_0$, then we may apply Lemma 6.4 to conclude that each t_i may be deformed rel endpoints to an arc of σ_g -length at most ϵ_2 . Thus, a may be deformed (rel endpoints if it is an arc) to have uniformly bounded σ_g -length.

If a is a closed curve, this is our desired curve α , and we are done.

If a is an arc, its endpoints lie in standard collars of Γ , whose boundaries have σ_f -length ϵ_1 . We may assume (possibly adjusting the curve α^τ slightly) that these endpoints lie on γ . Thus their g -images agree with their f -images, and lie in the ϵ_1 -thin parts of N ; hence the endpoints lie in the ϵ_0 -thin parts of σ_g . Now we can do the same mild surgery in both surfaces: form a small regular neighborhood of the union of a with the standard collar (or collars) meeting its endpoints. Let α be the boundary component of this neighborhood which passes close to a once or twice, and makes one or two additional trips along the collar boundaries. These arcs again are bounded in both metrics, and α still intersects β essentially. (See Figure 8.) This concludes the proof of Theorem 6.1. \square

6.2.1. *A variation on Theorem 6.1.* It will be useful to have the following minor variant of the Coarse Lipschitz property (P2) of Theorem 6.1, for a special case in the proof of Theorem 7.1. If $W \subset S$ and a is a vertex in $\mathcal{A}(W)$ represented by an arc, let λ_a be the lamination obtained by spinning this arc leftward around ∂W (as in Figure 2) and including ∂W . Using this, we can define $\mathbf{pleat}_\rho(a)$ to be $\mathbf{pleat}_\rho(\lambda_a)$, and this allows us to define $\Pi_{\rho,L}(a)$.

LEMMA 6.5. *Let $Y \subseteq W \subseteq S$, with $\partial W \neq \emptyset$. Let $\rho \in \mathcal{D}(S)$ be a Kleinian surface group such that $\ell_\rho(\partial W) < \epsilon_0$. If $v, w \in \mathcal{A}_0(W)$ and $d_W(v, w) \leq 1$, then*

$$\text{diam}_Y(\Pi_{\rho,L}(v) \cup \Pi_{\rho,L}(w)) \leq D_2$$

where D_2 depends only on S and L .

Proof. The proof proceeds as in Theorem 6.1, where the main point is to bound

$$\text{diam}_Y(\mathbf{short}_L(\sigma_f) \cup \mathbf{short}_L(\sigma_g))$$

for $f, g \in \mathbf{pleat}_\rho(v)$, where v is now a vertex of $\mathcal{A}(W)$. We construct a train track τ from λ_v in the metric σ_f , and note that τ has arcs that enter the collar of each component of ∂W that contains an endpoint of v . We can consider $\tau \cap W$ as a train track in W , and find as before an element of $\mathcal{A}(W)$ carried on τ which has bounded length, outside thin parts, in both σ_f and σ_g . To obtain a closed curve in $\mathcal{C}_0(W)$, we do the same surgery construction as in Figure 8, but using the collars of ∂W rather than those of ∂S . Because $\ell_\rho(\partial W) < \epsilon_0$, the surgered curve still has bounded length. \square

7. The projection bound theorem and consequences

In this section we will associate to a pair ν of end invariants a hierarchy H_ν , and prove Theorem 7.1, which in particular states that

$$d_Y(v, \Pi_{\rho, L}(v)),$$

where defined, is uniformly bounded above for all vertices v in $H_{\nu(\rho)}$ and essential subsurfaces $Y \subseteq S$. Theorem 7.1 generalizes Theorem 3.1 of [49], which only applies to the case $Y = S$.

This bound should be taken as an indication that the vertices in $H_{\nu(\rho)}$ and the bounded-length curves in N_ρ are somehow close to each other in a combinatorial sense. Indeed we will deduce the following two corollaries of this theorem:

The Tube Penetration Lemma 7.7 controls which pleated surfaces can penetrate deeply into a Margulis tube in N . In particular it states if we pleat along the curves of a pants decomposition coming from a slice of H_ν , then the resulting surface cannot enter any ϵ -tubes (for a certain ϵ) except those corresponding to the pants curves.

The Upper Bound Lemma 7.9 then shows that there is a uniform upper bound on the length $\ell_\rho(v)$ for every v appearing in H . This is the main step to obtaining Lipschitz bounds on the model map in Section 10.

7.1. From end invariants to hierarchy. Given a pair $\nu = \nu_\pm$ of end invariants we now produce a hierarchy H_ν . This is done by associating to ν_+ and ν_- a pair of generalized markings μ_+ and μ_- , and then applying Lemma 5.13.

The ending laminations in ν_\pm will be part of the base of the markings, so that what is left to do is encode the Teichmüller data. Note that ν_\pm are not uniquely recoverable from μ_\pm . In what follows let ν_s denote ν_+ or ν_- .

If R is a component of R_s^T (see §2.2 for notation), then the hyperbolic structure $\nu_R \in \mathcal{T}(R)$ admits a pants decomposition of total length at most L_0 (see 3.2), which includes all curves of length bounded by ϵ_0 . Let $\mu_s(R)$ be a maximal clean marking in R whose base is such a pants decomposition, and whose transversals \bar{t}_i are taken to be of minimal possible length (There is a bounded number of choices.) Note

that, by Bers' inequality relating lengths on the conformal boundary with lengths in the interior of a hyperbolic manifold (see [9]), we have $\ell_\rho(\text{base}(\mu_s(R))) \leq 2L_0$.

Recall that v_s^L is the union of parabolic curves p_s and ending laminations associated to the s side. Define μ_s to be the clean marking whose base is the union of v_s^L and $\text{base}(\mu_s(R))$ for all components R of R_s^T , and whose transversals are the transversals of the markings $\mu_s(R)$.

By the discussion in Section 2, μ_+ and μ_- do not share any infinite-leaf components. Thus we can apply Lemma 5.13 to conclude that there exists a hierarchy H_ν with $\mathbf{I}(H_\nu) = \mu_-$ and $\mathbf{T}(H_\nu) = \mu_+$.

Note that $\text{base}(\mu_\pm)$ is maximal, and therefore H_ν is 4-complete and in particular the base of every slice of H_ν is a pants decomposition (see discussion in §5.2).

Note also that H_ν is not uniquely defined by ν , as there were choices in the construction of μ_\pm , and there are choices in the construction of a hierarchy. However, our results will hold for any choice of H_ν , and we emphasize that no properties of the representation ρ other than its end invariants are used in the construction.

The main theorem of this section can now be stated:

THEOREM 7.1. *Fix a surface S . There exists $L_1 \geq L_0$ such that for every $L \geq L_1$ there exist $B, D_2 > 0$ such that, given $\rho \in \mathcal{D}(S)$, a hierarchy $H = H_{\nu(\rho)}$, and an essential subsurface Y in S with $\xi(Y) \neq 3$, the set*

$$\pi_Y(\mathcal{C}(\rho, L))$$

is B -quasiconvex in $\mathcal{A}(Y)$. Furthermore,

$$(7.1) \quad d_Y(v, \Pi_{\rho, L}(v)) \leq D_2$$

for every vertex v appearing in H such that the left-hand side is defined.

Note that the left-hand side of (7.1) is defined provided both $\pi_Y(v)$ and $\pi_Y(\Pi_{\rho, L}(v))$ are nonempty. The former is satisfied whenever γ_v intersects Y essentially. For the latter, since $\Pi_{\rho, L}(v)$ contains a pants decomposition, its π_Y -image is always nonempty provided $\xi(Y) \geq 4$.

7.2. Quasiconvexity. We first consider the bound (7.1) in a case that is only a slight perturbation of the result proved in [49]. After proving this version, we will prove Theorem 7.1 in the nonannulus case in Section 7.3. As a consequence of this we will obtain the Tube Penetration Lemma in Section 7.4, and then in Section 7.5 we will complete the proof of Theorem 7.1 for the case of annuli.

LEMMA 7.2. *Fix $L \geq L_0$, and suppose that $h \in H$ with $\xi(D(h)) \geq 4$ satisfies:*

$$(7.2) \quad d_{D(h)}(u, \Pi_{\rho, L}(u)) \leq d$$

whenever u is a vertex of $\text{simp}(\mathbf{I}(h))$ or $\text{simp}(\mathbf{T}(h))$. Then

$$(7.3) \quad d_{D(h)}(v, \Pi_{\rho, L}(v)) \leq d'$$

for all simplices v in h , where d' depends on d and L .

Proof. We will need the following lemma, which we proved in [49]:

LEMMA 7.3 (Lemma 3.3 of [49]). *Let \mathcal{X} be a δ -hyperbolic geodesic metric space and $\mathcal{Y} \subset \mathcal{X}$ a subset admitting a map $\Pi : \mathcal{X} \rightarrow \mathcal{Y}$ which is coarse-Lipschitz and coarse-idempotent. That is, there exists $C > 0$ such that*

(Q1) *If $d(x, x') \leq 1$, then $d(\Pi(x), \Pi(x')) \leq C$, and*

(Q2) *If $y \in \mathcal{Y}$, then $d(y, \Pi(y)) \leq C$.*

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

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

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

The proof uses a variation of the “stability of quasigeodesics” argument originating in Mostow’s rigidity theorem.

Proceeding with the proof of Lemma 7.2, let us consider first the case that h is a finite geodesic, with endpoints u, u' for which condition (7.2) holds. Let $Z = D(h)$. In this case we will not use the fact that h is in the hierarchy – the Lemma will hold for any finite geodesic h .

Short boundary case. Consider now the case that $\ell_\rho(\partial Z) < \epsilon_0$. If $x \in \mathcal{C}_0(Z)$, then, letting $f \in \mathbf{pleat}_\rho(x \cup \partial Z)$, the induced metric σ_f satisfies $\ell_{\sigma_f}(\partial Z) < \epsilon_0$, and therefore by our choice of L_0 (§3.2.2), there is a pants decomposition of S of length at most $L_0 \leq L$ which contains ∂Z as components, and hence also contains elements of $\mathcal{C}_0(Z)$. Hence

$$(7.4) \quad \Pi_{\rho, L}(x) \cap \mathcal{C}_0(Z) \neq \emptyset.$$

In order to apply Lemma 7.3, set

$$\begin{aligned} \mathcal{X} &= \mathcal{C}_1(Z), \\ \mathcal{Y} &= \mathcal{C}(\rho, L) \cap \mathcal{C}_0(Z), \end{aligned}$$

and define a map $\Pi : \mathcal{X} \rightarrow \mathcal{Y}$ by letting $\Pi(x)$ be an arbitrary choice of vertex in $\Pi_{\rho, L}(x) \cap \mathcal{C}_0(Z)$. Note that \mathcal{X} is δ -hyperbolic by Theorem 4.2.

Hypothesis (Q2) of Lemma 7.3 follows from property (P3) of Theorem 6.1, noting that any vertex y of \mathcal{Y} is also a vertex of $\mathcal{C}(\rho, L)$ that satisfies $\pi_Z(y) \neq \emptyset$.

If $\xi(Z) > 4$, then hypothesis (Q1) follows from property (P2) of Theorem 6.1, since two adjacent vertices of $\mathcal{C}(Z)$ are also adjacent vertices of $\mathcal{C}(S)$.

If $\xi(Z) = 4$, then two adjacent vertices x, x' of $\mathcal{C}(Z)$ represent curves that intersect, and hence are *not* adjacent in $\mathcal{A}(S)$, so we cannot apply Theorem 6.1 directly. However, there are vertices $a, a' \in \mathcal{A}_0(Z)$ (represented by arcs) such that $[x, a], [a, a']$ and $[a', x']$ are edges in $\mathcal{A}(Z)$. We can make sense of $\Pi_{\rho,L}(a)$ and $\Pi_{\rho,L}(a')$ as in Section 6.2.1, and apply Lemma 6.5 (with $Y = W = Z$) to bound $\text{diam}_Z(\Pi_{\rho,L}(v) \cup \Pi_{\rho,L}(w))$ for $(v, w) = (x, a), (a, a')$ and (a', x') . We then conclude (Q1) for x and x' via the triangle inequality.

For each endpoint w of h , of Theorem 6.1(P1) bounds $\text{diam}_Z(\Pi_{\rho,L}(w))$, and the hypotheses of Lemma 7.2 tell us that $d_Z(w, \Pi_{\rho,L}(w)) \leq d$. Together this bounds $d_{\mathcal{X}}(w, \Pi(w))$, and hence we can apply Lemma 7.3 to get the desired bound (7.3) on $d_Z(v, \Pi_{\rho,L}(v))$ for each simplex v in h .

Long boundary case. Now consider the case that $\ell_\rho(\partial Z) \geq \epsilon_0$. The main theorem of [48] states that there is a constant K depending only on L, ϵ_0 and S such that

$$\text{diam}_Z \mathcal{C}(\rho, L) \leq K.$$

This means, since the endpoints of h are within d of their $\pi_Z \circ \Pi_{\rho,L}$ images which are in $\pi_Z(\mathcal{C}(\rho, L))$, that the length of h is at most $2d + K$. The desired bound (7.3) now follows from the relative coarse Lipschitz property (P2) of Theorem 6.1 (note that $\pi_Z \circ \Pi_{\rho,L}$ is never empty since $\xi(Z) \geq 4$). This concludes the proof of the lemma when h is a finite geodesic.

Suppose that $\text{base}(\mathbf{T}(h))$ has no finite vertices, meaning it is an element of $\mathcal{EL}(Z)$. Then $\text{base}(\mathbf{T}(h))$ is a component of $\text{base}(\mathbf{T}(H))$, and in particular an ending lamination component of $\nu_+(\rho)$. The structure of ending laminations (§2.2) implies that ∂Z must be parabolic in ρ , and hence we are in the short boundary case above. Thurston's Theorem 2.2 gives us a sequence $\{\alpha_i\}_{i=1}^\infty$ in $\mathcal{C}_0(Z) \cap \mathcal{C}(\rho, L_0)$, such that $\alpha_i \rightarrow \mathbf{T}(h)$ as $i \rightarrow \infty$. A similar statement is true if $\mathbf{I}(h)$ has no finite vertices.

Note that, since the parabolics in ν_+ and ν_- must be distinct, if $\mathbf{T}(h)$ has no finite vertices, then $\mathbf{I}(h)$ has a finite vertex, and vice versa, unless $h = g_H$ and $Z = S$. Thus let us assume now that $\mathbf{I}(h)$ has a finite vertex α_0 .

The geodesics $h_i = [\alpha_0, \alpha_i]$ then satisfy the conditions of Lemma 7.3, and hence the bound (7.3) holds for all vertices of h_i . Because $\mathcal{C}(Z)$ is δ -hyperbolic and $\text{base}(\mathbf{T}(h))$ is a point in $\partial^{\text{c}}\mathcal{C}(Z)$ by Klarreich's Theorem 4.3, the h_i are fellow-travelers of h on larger and larger subsets. That is, for any simplex u in h , for large enough i , u is at most δ from a vertex u' in h_i .

Connecting u to u' by a geodesic in $\mathcal{C}_1(Z)$, we can apply the Coarse Lipschitz property (P2) of $\Pi_{\rho,L}$ to bound $d_Z(\Pi(u), \Pi(u'))$. We already have the bound (7.3) on $d_Z(u', \Pi(u'))$, so the triangle inequality then bounds $d_Z(u, \Pi(u))$.

The case that $\mathbf{I}(h)$ and $\mathbf{T}(h)$ are both infinite (and $Z = S$) is similar: there is a biinfinite sequence $\{\alpha_i\}_{i=-\infty}^\infty$ so that $h_i = [\alpha_{-i}, \alpha_i]$ are fellow travelers of h on

larger and larger segments, and $\alpha_i \in \mathcal{C}(\rho, L_0)$ for all i . The bound is then obtained in the same way. \square

7.3. Proof of the projection bound theorem: nonannulus case. To proceed with the proof of Theorem 7.1, we will need the following lemma:

LEMMA 7.4. *There exists L such that, for any essential subsurface Y with $\xi(Y) \neq 3$, and any vertex v in the hierarchy $H_{v(\rho)}$ with $\pi_Y(v) \neq \emptyset$, we have*

$$\pi_Y(\Pi_{\rho,L}(v)) \neq \emptyset.$$

This will allow us to effectively use the Coarse Lipschitz statement of Theorem 6.1.

If $\xi(Y) \geq 4$, the lemma evidently holds with $L \geq L_0$, since as we have already observed, $\Pi_{\rho,L}(v)$ always contains a pants decomposition, which must intersect Y essentially. The proof of the lemma for $\xi(Y) = 2$ will be postponed to Section 7.5.

Thus, although the rest of this section is written to be valid for any Y , we may only apply it for $\xi(Y) = 2$ after Lemma 7.4 has been established in that case. Let us henceforth assume that L has been given at least as large as the constant in Lemma 7.4, and denote $\Pi_\rho = \Pi_{\rho,L}$.

Bounds between levels. We next consider conditions on a geodesic h that allow us to establish (7.1) for $Y \subseteq D(h)$ and for simplices appearing in h .

LEMMA 7.5. *Suppose $h \in H$ with $\xi(D(h)) \geq 4$, $Y \subseteq D(h)$, and*

- (1) $d_{D(h)}(v, \Pi_\rho(v)) \leq d$ for all simplices v in h ,
- (2) $d_Y(u, \Pi_\rho(u)) \leq d$ if u is a vertex of $\text{simp}(\mathbf{I}(h))$ or $\text{simp}(\mathbf{T}(h))$ that intersects Y .

Then

$$(7.5) \quad d_Y(v, \Pi_\rho(v)) \leq d'$$

for all simplices v of h which intersect Y , where d' depends on d .

Proof. Let v be a simplex of h that intersects Y . Then v is not in $\phi_h(Y)$, and let us assume without loss of generality that $\max \phi_h(Y) < v$.

Suppose that the distance from v to the last simplex ω of h is no more than $2d + 1$. Let α be the segment of h from v to ω . Since every simplex in α crosses Y , the 1-Lipschitz property for π_Y implies

$$\text{diam}_Y(\alpha) \leq 2d + 1.$$

The Relative Coarse Lipschitz property (P2) of Π_ρ in Theorem 6.1 says that for every two successive simplices x, x' of α we have

$$\text{diam}_Y(\Pi_\rho(x) \cup \Pi_\rho(x')) \leq D_1.$$

By Lemma 7.4, $\pi_Y(\Pi_\rho(x))$ is nonempty for each x in α , so we can sum this over α to obtain

$$\text{diam}_Y(\Pi_\rho(\alpha)) \leq D_1(2d + 1).$$

Finally, we have by hypothesis (2) the bound

$$d_Y(\omega, \Pi_\rho(\omega)) \leq d.$$

Noting that $\pi_Y(\omega) \in \pi_Y(\alpha)$ and $\pi_Y(\Pi_\rho(\omega)) \subset \pi_Y(\Pi_\rho(\alpha))$, we can then put these three bounds together to obtain

$$\text{diam}_Y(\alpha \cup \Pi_\rho(\alpha)) \leq d + (D_1 + 1)(2d + 1).$$

In particular this bounds $d_Y(v, \Pi_\rho(v))$, and we have the desired statement.

Alternatively, suppose that the distance from v to the last simplex of h is at least $2d + 2$ (including the possibility that it is infinite). Let α be the segment of length $2d + 2$ beginning with v , and let w be the other endpoint of α . By hypothesis we have

$$d_{D(h)}(w, \Pi_\rho(w)) \leq d$$

and hence we can join w with $\pi_{D(h)}(\Pi_\rho(w))$ with a path in $\mathcal{A}_1(D(h))$ of length at most d , and whose first and last vertex (by Lemma 7.4) intersect Y . Using the 2-Lipschitz map $\psi : \mathcal{A}_0(D(h)) \rightarrow \mathcal{C}_0(D(h))$ described in Section 4, we can replace this with a path β with the same endpoints and whose interior vertices are in $\mathcal{C}(D(h))$, of length at most $2d$. Now recalling that $\max \phi_h(Y) < v$, we find that $d_{\mathcal{C}_1(D(h))}(w, \max \phi_h(Y)) \geq 2d + 3$. By the triangle inequality in $\mathcal{C}_1(D(h))$, every point in β is at least distance 3 from $\max \phi_h(Y)$ and hence at least distance 2 from $[\partial Y]$. In particular every interior vertex in β has nontrivial intersection with Y . Now the 1-Lipschitz property of π_Y again applies, to give us $\text{diam}_Y(\beta) \leq 2d$. In particular

$$d_Y(w, \Pi_\rho(w)) \leq 2d.$$

Now we can apply exactly the same argument as in the previous case. □

Inductive argument. We will now establish, by induction, the following claim:

LEMMA 7.6. *Let $h \in H$ be a geodesic with $\xi(D(h)) \geq 4$, and $Y \subset D(h)$. If v is a simplex of h or $\text{simp}(\mathbf{I}(h))$ or $\text{simp}(\mathbf{T}(h))$, and $\pi_Y(v) \neq \emptyset$, then*

$$d_Y(v, \Pi_\rho(v)) \leq d$$

where d depends on $\xi(D(h))$ and $\xi(S)$.

Proof. Consider first the case where $h = g_H$ and $Y = S$, and let us apply Lemma 7.2. If g_H is bi-infinite there are no conditions to check. If not, suppose that $\mathbf{I}(g_H) = \mathbf{I}(H)$ contains a finite vertex u . Then by definition of $\mathbf{I}(H)$ (§7.1) we have $\ell_\rho(u) \leq 2L_0$. This uniformly bounds $d_S(u, \Pi_\rho(u))$, since u and any

element of $\Pi_\rho(u)$ have lengths bounded by $2L_0$ and L_0 on the same pleated surface $f \in \mathbf{pleat}_\rho(u)$. Hence the condition of Lemma 7.2 holds for u . The same is true for $\mathbf{T}(g_H)$, and we therefore have

$$(7.6) \quad d_S(v, \Pi_\rho(v)) \leq d$$

for every simplex v of g_H , where d is a constant depending only on S .

We may now apply Lemma 7.5 to $h = g_H$, with every subsurface Y : Condition (1) of the lemma follows from (7.6) which we have just proved. For condition (2), note as above that for any finite vertex u in $\mathbf{I}(H)$ or $\mathbf{T}(H)$, $u \in \Pi_\rho(u)$ and hence if u intersects Y we have $d_Y(u, \Pi_\rho(u)) = 0$. Thus the conclusion of Lemma 7.6 holds for the geodesic g_H , and we have established the base case.

Now let $h \in H$ be any geodesic other than g_H , and suppose that Lemma 7.6 holds for all h' with $\xi(D(h')) > \xi(D(h))$, for some constant d . In order to apply Lemma 7.2 to h , we must consider the condition on its endpoints. Let b, f be such that $b \not\prec h \prec^d f$. If $\mathbf{T}(h)$ is in $\mathcal{EL}(D(h))$ (h is infinite in the forward direction), then there is nothing to check. Otherwise $\mathbf{T}(h)$ has a finite simplex w , which by definition of $h \prec^d f$ appears either in a simplex of f or in $\mathbf{T}(f)$. By the induction hypothesis applied to f we have $d_{D(h)}(w, \Pi_\rho(w)) \leq d$. The same reasoning applies to $\mathbf{I}(h)$, using b . Thus, we may apply Lemma 7.2 to obtain the bound (7.3) for all simplices of h .

This then establishes condition (1) of Lemma 7.5 for h . To obtain condition (2) for any $Y \subseteq D(h)$, we again use the fact that the endpoints of h are contained in b and f , and hence the inductive hypothesis for b and f yields this bound also (for suitable constants). Thus, we apply Lemma 7.5 to give us the statement of Lemma 7.6 for h . \square

Reduction to nested case. So far we have proved the bound on $d_Y(v, \Pi_\rho(v))$ in the case where v is in a geodesic h with $Y \subseteq D(h)$. It remains to show that the general case reduces to this one.

For any geodesic $h \in H$ with $\xi(D(h)) \geq 4$ and simplex v in h , let $v \cup \partial D(h)$ denote the simplex of $\mathcal{C}(S)$ corresponding to the disjoint union of curves γ_v and (the nonperipheral components of) $\partial D(h)$. Let us fix Y and prove, by induction on $\xi(D(h))$, that

$$(7.7) \quad d_Y(v \cup \partial D(h), \Pi_\rho(v \cup \partial D(h))) \leq C,$$

provided the left-hand side is defined (where C is a uniform constant). By Lemma 7.4, all that is necessary for the left-hand side to be defined is that $v \cup \partial D(h)$ have nontrivial intersection with Y .

For any h such that $Y \subseteq D(h)$, (7.7) reduces to what we have already proven. This applies in particular to $h = g_H$. Now given any other $h \in H$, suppose the claim is true for all h' with $\xi(D(h')) > \xi(D(h))$. Let $h \prec^d f$. Then $D(h)$ is a

component domain of $(D(f), w)$ for some simplex w in f , and hence $\partial D(h)$ is contained in $w \cup \partial D(f)$.

If $\partial D(h)$ intersects Y essentially, then so does $w \cup \partial D(f)$, and hence the inductive hypothesis bounds $d_Y(w \cup \partial D(f), \Pi_\rho(w \cup \partial D(f)))$. Now note that $v \cup \partial D(h)$ and $w \cup \partial D(f)$ have intersection $\partial D(h)$, and since this has nonempty projection $\pi_Y(\partial D(h))$ we have a bound on $\text{diam}_Y(v \cup \partial D(h) \cup w \cup \partial D(f))$. Furthermore $\Pi_\rho(\partial D(h))$ contains $\Pi_\rho(v \cup \partial D(h))$ and $\Pi_\rho(w \cup \partial D(f))$. Thus, invoking Theorem 6.1 we may deduce a bound on

$$d_Y(v \cup \partial D(h), \Pi_\rho(v \cup \partial D(h))).$$

If $\partial D(h)$ does not intersect Y essentially, but v does, then $Y \subseteq D(h)$, and we have already proven the bound.

This concludes the proof of Theorem 7.1, in the nonannulus case.

7.4. Tube penetration lemma. Applying the nonannulus case of Theorem 7.1, we are now going to control which pleated surfaces arising from a hierarchy can meet Margulis tubes in N_ρ .

LEMMA 7.7. *There exists $\epsilon_3 > 0$ with the following property: Let u be a vertex appearing in $H_{v(\rho)}$ and let $\alpha \in \pi_1(S)$ satisfy $\ell_\rho(\alpha) < \epsilon_3$. A map*

$$f \in \text{pleat}_\rho(u)$$

meets the Margulis tube $\mathbb{T}_{\epsilon_3}(\alpha)$ only if α represents a simple curve which has no essential intersection with γ_u .

COROLLARY 7.8. *Let μ be a pants decomposition whose components are vertices of H , and $\alpha \in \pi_1(S)$ a primitive element with $\ell_\rho(\alpha) < \epsilon_3$. A map*

$$f \in \text{pleat}_\rho(\mu)$$

meets the Margulis tube $\mathbb{T}_{\epsilon_3}(\alpha)$ if and only if α represents an element of μ .

The “if” direction of the corollary is obvious. For the “only if” direction, if f meets $\mathbb{T}_{\epsilon_3}(\alpha)$, then Lemma 7.7 implies that every component of μ has no essential intersection with the simple curve α . Since μ is a pants decomposition, α must be one of the components.

Proof of the Tube Penetration Lemma. We will choose $\epsilon_3 \leq \epsilon_1$, and assume that $f(S)$ meets $\mathbb{T}_{\epsilon_3}(\alpha)$. By the choice of ϵ_1 in Section 3.2.2 we know that α has length at most ϵ_0 in σ_f . Hence it is a multiple of the core of a thin collar and since we assumed it was primitive it must in fact be the core, and in particular represents a vertex of $\mathcal{C}(S)$.

The vertex u must appear in some simplex v of a geodesic $h \in H$, with $\xi(D(h)) \geq 4$. Assuming γ_v intersects α essentially, we will prove that $f(S)$ can penetrate no further than a distance d into $\mathbb{T}_{\epsilon_1}(\alpha)$, where d depends on $\xi(D(h))$.

The proof will be by downward induction on $\xi(D(h))$ with $D(h) = S$ being the base case.

Since $v \notin \phi_h(\alpha)$, we may assume without loss of generality that $v > \max \phi_h(\alpha)$ (otherwise we reverse the directions in the rest of the argument). Let J denote the largest segment of h beginning with v so that for every vertex x of J its geodesic representative x^* in N_ρ meets $\mathbb{T}_{\epsilon_1}(\alpha)$. For any $x \in J$, every $f \in \mathbf{pleat}_\rho(x)$ has nontrivial intersection with $\mathbb{T}_{\epsilon_1}(\alpha)$. As above this means $\ell_{\sigma_f}(\alpha) \leq \epsilon_0$, and in particular

$$\alpha \in \Pi_\rho(x).$$

Since $\xi(D(h)) \geq 4$ we may apply the nonannulus case of Theorem 7.1 to obtain

$$(7.8) \quad d_{D(h)}(x, \alpha) \leq b$$

where b is a constant derived from that theorem using the bound D_0 supplied by Theorem 6.1 for $\text{diam}_{D(h)}(\Pi_\rho(x))$. We conclude, since J is geodesic in $\mathcal{C}_1(D(h))$, a bound of the form

$$(7.9) \quad |J| \leq b'.$$

If J is the entire portion of h following v , then let w denote the last simplex of J and hence the last of h . If $h = g_H$ (the base case of the induction), then $w \subset \text{base}(\mathbf{T}(H))$ and hence represents curves in N_ρ of length at most $2L_0$. This bounds by L_0 the distance which w can penetrate into $\mathbb{T}_{\epsilon_1}(\alpha)$, since w^* cannot be completely contained in this Margulis tube (if it were it would be α itself, but we have $w > \max \phi_h(\alpha)$).

If $h \neq g_H$, then there exists $f \in H$ such that $h \searrow^d f$, and hence w is contained in a simplex of f . By the induction, w^* cannot penetrate further than $d_{\xi(D(f))}$ into $\mathbb{T}_{\epsilon_1}(\alpha)$.

If the last simplex of J is not the last simplex of h , then let w denote its successor, so by definition w^* does not meet $\mathbb{T}_{\epsilon_1}(\alpha)$. Define $J' = J \cup \{w\}$ (so $|J'| \leq |J| + 1$).

We will now construct a path in N_ρ joining any point $q_v \in v^* \cap \mathbb{T}_{\epsilon_1}(\alpha)$ to the boundary of $\mathbb{T}_{\epsilon_1}(\alpha)$, whose length will be bounded. There are two possible cases:

Case 1: $\xi(D(h)) > 4$. Consider any two successive simplices x, x' in J' , and let q be a point of $x^* \cap \mathbb{T}_{\epsilon_1}(\alpha)$. Since x and x' determine two disjoint curves in $D(h)$, there is a pleated surface $f \in \mathbf{pleat}_\rho([xx'])$.

Again by the choice of ϵ_1 , q is the f -image of a point $p \in \gamma_x$ in the ϵ_0 -thin part of σ_f associated to α . By our construction both x and x' cross α , so there is a point p' on $\gamma_{x'}$ in this thin part, a σ_f -distance of at most ϵ_0 from p . Thus $q' = f(p')$ is connected to q by a path of length at most ϵ_0 .

Apply this successively to all the vertices of J' , beginning with q_v . Each new point is either outside $\mathbb{T}_{\epsilon_1}(\alpha)$, in which case we stop, or inside it, in which

case we continue, and reach w^* in at most $|J'|$ steps. Since w^* is either disjoint from $\mathbb{T}_{\epsilon_1}(\alpha)$ or penetrates no more than $d_{\xi(D(h))+1}$ into it, we conclude that the distance from q to the boundary of $\mathbb{T}_{\epsilon_1}(\alpha)$ is at most $|J'|\epsilon_0 + d_{\xi(D(h))+1} \leq b'\epsilon_0 + d_{\xi(D(h))+1} \equiv d_{\xi(D(h))}$.

Case 2: $\xi(D(h)) = 4$. Now successive vertices x, x' in J' are *not* disjoint and hence we cannot form $\mathbf{pleat}_\rho(x \cup x')$. Instead, extend $x \cup \partial D(h)$ and $x' \cup \partial D(h)$ to pants decompositions p, p' which differ by an elementary move, and consider the corresponding surfaces $g \in \mathbf{pleat}_\rho(p)$, $g' \in \mathbf{pleat}_\rho(p')$, and the halfway surface $f = f_{p,p'}$ (see §3.1). The halfway surface f is pleated along a lamination containing a leaf (or two leaves) l that is part of the pleating locus of g in $D(h)$, and a leaf or two l' that is part of the pleating locus of g' in $D(h)$. The discussion in Section 3.1 tells us that α intersects l essentially since it intersects γ_x , and intersects l' since it intersects $\gamma_{x'}$.

Thus in σ_g both γ_x and l pass through the collar of α , in σ_f both l and l' pass through it, and in $\sigma_{g'}$ both l' and $\gamma_{x'}$ do. It follows that we can apply the argument of case 1 to produce a path as before, but with three times the number of steps.

Thus we have shown that v^* cannot penetrate more than a certain $d_{\xi(D(h))}$ into $\mathbb{T}_{\epsilon_1}(\alpha)$. Now using (3.7) this implies that for a certain uniform ϵ_3 , v^* cannot meet the tube $\mathbb{T}_{\epsilon_3}(\alpha)$. \square

7.5. Projection bounds in the annulus case. We are now ready to complete the proof of Theorem 7.1 in the case that Y is an annulus. It suffices to give the proof of Lemma 7.4 in this case:

Proof. Let α denote the core of the annulus Y . Let ϵ_3 be the constant in Lemma 7.7, and take $L \geq \mathcal{L}(\epsilon_3)$. Then $\pi_Y(v) \neq \emptyset$ implies that, for any $f \in \mathbf{pleat}_\rho(v)$, $f(S)$ does not meet $\mathbb{T}_{\epsilon_3}(\alpha)$. Thus, $\ell_{\sigma_f}(\alpha) \geq \epsilon_3$ and it follows by definition of the function $\mathcal{L}()$ (see §3.2) that there is a pants decomposition of total length at most L which crosses α essentially. Thus, $\pi_Y(\Pi_{\rho,L}(v)) \neq \emptyset$, which gives the statement of Lemma 7.4. \square

The rest of the proof of Theorem 7.1 proceeds exactly as in Section 7.3.

7.6. Upper length bounds. Our final application of the projection mechanism will be to obtain an *a priori* upper bound on the length of every curve that appears in the hierarchy.

LEMMA 7.9. *For every vertex v in the hierarchy $H_{v(\rho)}$,*

$$\ell_\rho(v) \leq D_3$$

where D_3 is a constant depending only on the topological type of S .

Proof. Any vertex v in H is contained in some maximal slice τ , by Lemma 5.8. Let $\mu = \mu_\tau$ be the associated marking, and fix

$$f \in \mathbf{pleat}_\rho(\text{base}(\mu)) \subset \mathbf{pleat}_\rho(v).$$

Let ϵ_3 be the constant in Lemma 7.7 and Corollary 7.8. Applying Corollary 7.8, we find that all curves α in S with $\ell_{\sigma_f}(\alpha) \leq \epsilon_3$ must be components of $\text{base}(\mu)$. If $\ell_{\sigma_f}(v) \leq \epsilon_3$, then we already have our desired bound, and we are done. Thus assume $\ell_{\sigma_f}(v) > \epsilon_3$. Since γ_v is disjoint from the other curves of $\text{base}(\mu)$, it must be contained, nonperipherally, in a component R of the ϵ_3 -thick part of σ_f .

Since $\epsilon_3 < \epsilon_0$, there is a pants decomposition η of R of σ_f -length at most L_0 , and in particular

$$\eta \subset \Pi_{\rho, L_0}(\text{base}(\mu)).$$

Since R is ϵ_3 -thick, there is a constant L depending on ϵ_3 so that each component of η is crossed by a transversal curve of σ_f -length at most L , which misses the rest of η . Let μ_2 be the clean marking with $\text{base}(\mu_2) = \eta$ and these bounded-length transversal curves.

By Theorem 7.1, if Y is any subsurface of R that has essential intersection with both η and $\text{base}(\mu)$ (including $Y = R$), then

$$d_Y(\text{base}(\mu), \eta) \leq D_2.$$

The only subsurfaces of R excluded by this are annuli associated with components of $\text{base}(\mu)$ or η . Let Y be an annulus whose core is a component of η , but not a component of $\text{base}(\mu)$. The transversal t in μ_2 crossing Y , since it has σ_f -length at most L , is contained in $\Pi_{\rho, L}(\text{base}(\mu))$. Thus applying Theorem 7.1 again, there is a D'_2 such that

$$d_Y(\text{base}(\mu), t) \leq D'_2.$$

To deal with the remaining annuli, extend $\text{base}(\mu)$ to a new clean marking μ' by adding for each component α a transversal t satisfying $d_\alpha(t, \mu_2) \leq 2$ (this is always possible by picking some transversal and applying Dehn twists). Thus for all annuli Y with cores in $\text{base}(\mu) = \text{base}(\mu')$ we also have

$$d_Y(\mu', \mu_2) \leq 2.$$

In other words we have shown that the two clean markings satisfy

$$d_Y(\mu', \mu_2) \leq D$$

for a uniform D , and *all* domains $Y \subseteq R$. By Lemma 5.12, this gives an upper bound E on the elementary move distance $d_{\text{el}}(\mu_2, \mu')$ (to obtain the bound, set $K = \max(K_0, D + 1)$ in the lemma).

Since the base curves and transversals of μ_2 have σ_f lengths bounded by L , and since an elementary move can change the lengths of components of a maximal

clean marking by at most a bounded factor, we obtain a bound on the σ_f -lengths of the curves of μ' , in terms of L and E . In particular we have an upper bound on $\ell_\rho(v)$, and we are done. \square

8. The model manifold

In this section we will construct an oriented metric 3-manifold M_ν associated to the end invariants $\nu = \nu_\pm(\rho)$. M_ν is intended to be a model for the geometry of the augmented convex core \widehat{C}_N of N_ρ . The following is a summary of the structure of M_ν .

- (1) M_ν is properly embedded in $\widehat{S} \times \mathbb{R}$, and is homeomorphic to \widehat{C}_N .
- (2) An open subset $\mathcal{U} \subset M_\nu$ is called the set of *tubes* of M_ν . Each component $U \subset \mathcal{U}$ is of the form

$$\mathbf{collar}(v) \times I$$

where v is either a vertex of H_ν or a boundary component of S , and $I \subset \mathbb{R}$ is an open interval. I is bounded except for the finitely many v corresponding to parabolics. The correspondence $U \leftrightarrow v$ is bijective.

- (3) Let $M_\nu[0]$ be $M_\nu \setminus \mathcal{U}$. $M_\nu[0]$ is a union of standard “blocks” of a finite number of topological types.
- (4) Except for finitely many blocks adjacent to ∂M_ν , all blocks fall into a predetermined finite number of isometry types.
- (5) Each tube U is isometric to a hyperbolic or a parabolic tube, and with respect to a natural marking has boundary parameters $(\omega_M(U), \epsilon_1)$. (See §3.2.) We call $\omega_M(U)$ the *meridian coefficient* of U .

Let $H = H_\nu$ be associated to ν as in Section 7.1. We will begin by constructing $M_\nu[0]$ abstractly as a union of blocks. We will then show how to embed it in $\widehat{S} \times \mathbb{R}$, and in its complement we will find the tubes \mathcal{U} which we adjoin to obtain M_ν .

8.1. Blocks and gluing. The typical blocks from which we build $M_\nu[0]$ are called *internal blocks*. There are some special cases of blocks associated to the boundary of the convex core, but these can be ignored on a first reading (and do not appear, for example, in the doubly degenerate case).

Given a 4-edge e in H , let g be the 4-geodesic containing it, and let $D(e)$ be the domain, $D(g)$. Recall (§5.6) that e^- and e^+ denote the initial and terminal vertices of e .

To each e we will associate a block $B(e)$, defined as follows:

$$B(e) = (D(e) \times [-1, 1]) \setminus (\mathbf{collar}(e^-) \times [-1, -1/2) \cup \mathbf{collar}(e^+) \times (1/2, 1]).$$




Figure 9. Constructing an internal block $B(e)$. If $D(e)$ is a one-holed torus, then $B(e)$ is obtained by gluing face A to face A' , and B to B' . The curved vertical faces become $\partial D(e) \times [-1, 1]$. If $D(e)$ is a 4-holed sphere, then $B(e)$ is obtained by doubling this object along A, A', B and B' .

That is, $B(e)$ is the product $D(e) \times [-1, 1]$, with solid-torus trenches dug out of its top and bottom boundaries, corresponding to the two vertices of e . See Figure 9.

We remark that, in this construction, we think of blocks of distinct edges as disjoint (for example think of the intervals $[-1, 1]$ as disjoint copies of a standard interval). Afterwards we will glue them together using specific rules, and embed the resulting manifold in $\hat{S} \times \mathbb{R}$.

We break up the boundary of $B(e)$ into several parts (see Figure 10 for a schematic). The *gluing boundary* of $B(e)$ is

$$\partial_{\pm} B(e) \equiv (D(e) \setminus \mathbf{collar}(e^{\pm})) \times \{\pm 1\}.$$

Note that the gluing boundary is always a union of three-holed spheres.

The rest of the boundary is a union of annuli, with

$$\partial_{\parallel} B(e) \equiv \partial D(e) \times [-1, 1]$$

being the *outer annuli*.

The *inner annuli* $\partial_i^{\pm} B(e)$ are the boundaries of the removed solid tori. That is,

$$\partial_i^{\pm} B(e) = \partial B(e) \cap \partial(\mathbf{collar}(e^{\pm}) \times \pm(1/2, 1])$$

(where $+(a, b]$ denotes $(a, b]$ and $-(a, b] = [-b, -a)$). These annuli break up into a *horizontal part*

$$\partial_{ih}^{\pm} B(e) = \overline{\mathbf{collar}(e^{\pm})} \times \{\pm 1/2\}$$

and a *vertical part*

$$\partial_{iv}^{\pm} B(e) = \partial \mathbf{collar}(e^{\pm}) \times \pm[1/2, 1].$$

Boundary blocks. Recall from Section 2.2 that R_+^T denotes the union of sub-surfaces in the top of the relative compact core that face geometrically finite ends.




Figure 10. Schematic diagram of an internal block.




Figure 11. Schematic for a boundary block

Let R be a subsurface of S homotopic to a component of R_+^T , and let ν_R be the associated component of ν_+^T in $\mathcal{T}(R)$. We construct a block $B_{\text{top}}(\nu_R)$ as follows: Let \mathbf{T}_R be the set of curves of $\text{base}(\mathbf{T}(H_\nu)) = \text{base}(\mu_+)$ that are contained in R . Define

$$B'_{\text{top}}(\nu_R) = R \times [-1, 0] \setminus (\text{collar}(\mathbf{T}_R) \times [-1, -1/2))$$

and let

$$B_{\text{top}}(\nu_R) = B'_{\text{top}}(\nu_R) \cup \partial R \times [0, \infty).$$

This is called a *top boundary block* (see Figure 11). Its *outer boundary* $\partial_o B_{\text{top}}(\nu_R)$ is $R \times \{0\} \cup \partial R \times [0, \infty)$, which we note is homeomorphic to $\text{int}(R)$. This will correspond to a boundary component of \widehat{C}_N . The gluing boundary of this block lies on its bottom: it is

$$\partial_- B_{\text{top}}(\nu_R) = (R \setminus \text{collar}(\mathbf{T}_R)) \times \{-1\}.$$

Similarly if R is a component of R_-^T we let $\mathbf{I}_R = \mathbf{I}(H_\nu) \cap R$ and define

$$B'_{\text{bot}}(\nu_R) = R \times [0, 1] \setminus \text{collar}(\mathbf{I}_R) \times (1/2, 1].$$

and the corresponding *bottom boundary block*

$$B_{\text{bot}}(\nu_R) = B'_{\text{bot}}(\nu_R) \cup \partial R \times (-\infty, 0].$$

The gluing boundary here is $\partial_+ B_{\text{bot}}(\nu_R) = (R \setminus \text{collar}(\mathbf{I}_R)) \times \{1\}$.

The vertical annulus boundaries are now $\partial_{\parallel} B_{\text{top}}(\nu_R) = \partial R \times [-1, \infty)$ and the internal annuli ∂_i^{\pm} are a union of possibly several component annuli, one for each component of \mathbf{T}_R or \mathbf{I}_R .

Since the interior of each block is a subset of $Z \times [-1, 1]$ for a subsurface Z of S , it inherits a natural orientation from the fixed orientation of S and \mathbb{R} .

Gluing instructions. We obtain $M_v[0]$ by taking the disjoint union of all blocks and identifying them along the three-holed spheres in their gluing boundaries. The rule is that whenever two blocks B and B' have the same three-holed sphere Y appearing in both $\partial^+ B$ and $\partial^- B'$, we identify these boundaries using the identity on Y . The hierarchy will serve to organize these gluings and insure that they are consistent.

By definition, all these three-holed spheres are component domains in the hierarchy. Conversely, we will now check that every component domain Y in H with $\xi(Y) = 3$ must occur in the gluing boundary of exactly two blocks.

Lemma 5.17 tells us that $\mathbf{T}(H)|_Y \neq \emptyset$ if and only if there exists a 4-geodesic $f \in H$ with $Y \not\prec^d f$, and f is unique if it exists. When f exists, there is an edge e_f in f with Y a component domain of $(D(f), e_f^-)$, and hence $Y \times \{-1\}$ is a component of $\partial_- B(e_f)$.

If f does not exist and $\mathbf{T}(H)|_Y = \emptyset$, Y must be a component domain of $\text{base}(\mathbf{T}(H))$, and hence $Y \times \{-1\}$ occurs on the gluing boundary $\partial_- B_{\text{top}}(\nu_R)$ for some top boundary block.

Similarly, if $\mathbf{I}(H)|_Y \neq \emptyset$, then $b \not\prec^d Y$ for a unique 4-geodesic b , and $Y \times \{1\}$ is in $\partial_+ B(e_b)$, and if $\mathbf{I}(H)|_Y = \emptyset$, then $Y \times \{1\}$ appears in the gluing boundary of a bottom boundary block.

We conclude that each Y serves to glue exactly two blocks, and in particular $M_v[0]$ is a manifold. The orientations of the blocks extend consistently to an orientation of $M_v[0]$.

8.2. Embedding in $\widehat{S} \times \mathbb{R}$. The interior of each block $B(e)$ inherits a 2-dimensional ‘‘horizontal foliation’’ from the foliation of the product $D(e) \times [-1, 1]$ by surfaces $D(e) \times \{t\}$. Let us call the connected leaves of this foliation the *level surfaces* of $B(e)$. (For boundary blocks we do similarly, and for the added vertical annuli of the form $\partial R \times [0, \infty)$ or $\partial R \times (-\infty, 0]$, we also include the level circles as leaves of this foliation.) An embedding $f : M_v[0] \rightarrow S \times \mathbb{R}$ will be called *flat* if each connected leaf of the horizontal foliation $Y \times \{t\}$ in $M_v[0]$ is mapped to a level set $Y \times \{s\}$ in the image, with the map on the first factor being the identity.

THEOREM 8.1. *$M_v[0]$ admits a proper flat orientation-preserving embedding $\Psi : M_v[0] \rightarrow S \times \mathbb{R}$.*

Remark. Once we have fixed the embedding Ψ we will adopt notation where $M_v[0]$ is identified with its image in $S \times \mathbb{R}$, and Ψ is the identity.

Proof. In the course of the proof we will fix a certain exhaustion of $M_v[0]$ (minus its boundary blocks) by subsets M_i^j , where M_i^j and M_i^{j+1} differ by the addition of one block on the “top”, and M_{i-1}^j and M_i^j differ by the addition of a block on the bottom. The map will be built inductively on M_i^j .

Let $\{\tau_i\}_{i \in \mathcal{I}}$ be a resolution of H , as in Lemma 5.8.

The pants decompositions $\text{base}(\mu_{\tau_i})$ form a sequence that may have adjacent repetitions – that is, $\text{base}(\mu_{\tau_i})$ and $\text{base}(\mu_{\tau_{i+1}})$ may be the same because the move $\tau_i \rightarrow \tau_{i+1}$ involves a $\xi = 2$ geodesic (twists in an annulus complex) or a $\xi > 4$ geodesic (reorganization moves). If we remove such repetitions we obtain a sequence of pants decompositions $\{\eta_i\}_{i \in \mathcal{I}'}$, where \mathcal{I}' is a new index set, so that each step $\eta_i \rightarrow \eta_{i+1}$ corresponds to an edge in a $\xi = 4$ geodesic. For a vertex v appearing in H , define

$$J'(v) = \{i \in \mathcal{I}' : v \in \eta_i\}.$$

Lemma 5.16 implies that $J(v)$ is an interval in \mathbb{Z} , and it follows that $J'(v)$ is an interval as well.

To each η_i we associate a subsurface F_i of $M_v[0]$ whose components are level surfaces of blocks, as follows. A complementary component Y of η_i is necessarily a three-holed sphere which appears as a component domain in H . Thus there are blocks B_1 and B_2 such that Y is isotopic to a component of $\partial_+ B_1$ and to a component of $\partial_- B_2$, which are identified in $M_v[0]$. Let this identified subsurface be a component of F_i . Repeating this for all complementary components of η_i we obtain all of F_i , which we call a *split-level surface*.

If η_i and η_{i+1} differ by a move in a $\xi = 4$ geodesic k , then there is a block B_i with domain $D(k)$, so that F_{i+1} is obtained from F_i by removing $\partial_- B_i$ and replacing it with $\partial_+ B_i$.

Now define $M_0^0 = F_0$, and inductively define

$$M_i^{j+1} = M_i^j \cup B_j$$

and

$$M_{i-1}^j = M_i^j \cup B_{i-1}.$$

Thus we are building up $M_v[0]$ by successively adding blocks above and below. Since the resolution $\{\tau_i\}$ contains an elementary move for every 4-edge of H by Lemma 5.8, every internal block is included in $\{M_i^j\}$.

Now the map Ψ can easily be defined inductively on $M_{-\infty}^\infty = \cup_{i,j \in \mathbb{Z}} M_i^j$. Begin by mapping $F_0 = M_0^0$ to $S \times \{0\}$ by the map that restricts to the identity on the surface factors. Now suppose we wish to add a block B_j to M_i^j . By induction the boundary components of $\partial_- B_j$, which are part of F_j , are already mapped flatly by Ψ . We extend Ψ to B_j so that it is an orientation-preserving flat embedding. Since there may be two components of $\partial_- B_j$ which are mapped to different heights, we

may have to stretch the two “legs” of B_j by different factors, but the map can be piecewise-affine on the vertical directions, and the identity on the surface factors. (See Figure 12 for a schematic example.) We always choose Ψ on B_j so that the image of every annulus in $\partial_{\parallel} B_j$ has height at least 1. This will guarantee that Ψ is proper.

To define the map on the boundary blocks, note that the gluing boundary of a boundary block must be part of the boundary of $M_{-\infty}^{\infty}$, and is mapped by Ψ so that each component is mapped flatly. Hence we can extend Ψ as an orientation-preserving flat embedding on each boundary block, making sure that the map on the added annuli $\partial R \times [0, \infty)$ (or $\partial R \times (-\infty, 0]$) is proper.

We observe that $\Psi(M_i^j)$ always lies “below” $\Psi(F_j)$, in the following sense: First, for each component Y of F_j , if $\Psi(Y) = Y \times \{t\}$, then $Y \times (t, \infty)$ does not meet $\Psi(M_i^j)$. Second, let A be a component annulus of $\mathbf{collar}(\eta_j)$. Let s be the minimum image height of the (one or two) subsurfaces of F_j which project to complementary components of $\mathbf{collar}(\eta_j)$ adjacent to A . Then $A \times (s, \infty)$ does not meet $\Psi(M_i^j)$. This property holds trivially for M_0^0 , and it is easy to see that it is preserved with the addition of each block. The corresponding property holds in the opposite direction, i.e. $\Psi(M_i^j)$ lies “above” $\Psi(F_i)$. From this we may conclude that Ψ is an embedding. \square

We next wish to keep track of how the solid tori $\{U(v)\}$ arise in the complement of $\Psi(M_v[0])$.

LEMMA 8.2. *Let v be vertex of H and let $\mathcal{F}(v)$ be the union of annuli in the boundaries of blocks of $M_v[0]$ that are in the homotopy class of v . Then $\mathcal{F}(v)$ is a torus or an annulus, and $\Psi|_{\mathcal{F}(v)}$ is an embedding with image*

$$\partial(\mathbf{collar}(v) \times [s_1, s_2])$$

if v is not parabolic in either $\mathbf{I}(H)$ or $\mathbf{T}(H)$,

$$\partial(\mathbf{collar}(v) \times [s_1, \infty))$$

if v is parabolic in $\mathbf{T}(H)$, and

$$\partial(\mathbf{collar}(v) \times (-\infty, s_2])$$

if v is parabolic in $\mathbf{I}(H)$.

Proof. If v intersects $\text{base}(\mathbf{I}(H))$, then Lemma 5.15 gives us a unique 4-edge e_1 with $v = e_1^+$. Thus the block $B(e_1)$ has inner annulus $\partial_{i^+}^+ B(e_1)$ in the homotopy class of v . Let us call the horizontal subannulus $\partial_{ih}^+ B(e_1)$ the *bottom annulus* of v . By the definition, Ψ maps this annulus to $\mathbf{collar}(v) \times \{s_1\}$ for some $s_1(v) \in \mathbb{R}$, and the vertical part $\partial_{iv}^+ B(e_1)$ is mapped to annuli of the form $\alpha \times [s_1, t]$ where α is a boundary component of $\mathbf{collar}(v)$ and $t > s_1$.

The elementary move associated to e_1 introduces v into the resolution of the hierarchy, so that all other annuli in this homotopy class occur after e_1 .

If v is a component of $\text{base}(\mathbf{I}(H))$, and if it has a transversal, then a similar description holds, where $B(e_1)$ is replaced by a bottom boundary block. If $v \in \text{base}(\mathbf{I}(H))$ but has no transversal, it is parabolic in $\mathbf{I}(H)$, and there is no block B with an inner annulus homotopic to v .

The same discussion holds with regard to $\mathbf{T}(H)$, yielding a unique *top annulus* for v in the bottom boundary of an appropriate block, unless v is parabolic in $\mathbf{T}(H)$. We note that v cannot be parabolic in both $\mathbf{I}(H)$ and $\mathbf{T}(H)$. All other annuli in the homotopy class of v must occur as outer annuli, i.e. in the sides of blocks whose domains have boundary components homotopic to v . These blocks occur in the resolution in the interval $J(v)$, that is between the move that introduces v in the resolution and the one that takes it out. Since Ψ is locally an embedding (by the orientation-preserving condition) these annuli must fit together into a torus or annulus, as described in the statement of the lemma. \square

We note also that these annuli and tori, together with the outer boundaries of boundary blocks, form the entire boundary of $M_\nu[0]$. From the description of $\Psi(\mathcal{F}(v))$ it clearly bounds an open solid torus in $S \times \mathbb{R}$, namely $\mathbf{collar}(v) \times (s_1, s_2)$, $\mathbf{collar}(v) \times (s_1, \infty)$, or $\mathbf{collar}(v) \times (-\infty, s_2)$ in the various cases, and we denote this solid torus $U(v)$.

Each solid torus $U(v)$ is in fact disjoint from $\Psi(M_\nu[0])$: For $j = \inf J'(v)$, the block B_{j-1} contains the bottom annulus of v , and $\Psi(M_i^j)$ is below the image annulus $\mathbf{collar}(v) \times \{s_1(v)\}$ as in the proof of Theorem 8.1, and in particular disjoint from $U(v)$ which lies above this annulus. For $\inf J'(v) < j \leq \sup J'(v)$, B_{j-1} is disjoint from $\mathbf{collar}(v) \times \mathbb{R}$ and in particular from $U(v)$. For $j = \sup J'(v) + 1$, B_j contains the top annulus of v and lies above it in the image, and all subsequent blocks are placed above this, and hence disjoint from $U(v)$.

Peripheral tubes. It remains to describe the tubes $U(v)$ for v a component of ∂S . So far, both $M_\nu[0]$ and the tubes $U(v)$ for vertices v have embedded, disjointly, into $S \times \mathbb{R}$. Now for a boundary component v , the annulus $\mathbf{collar}(v)$ is a component of $\hat{S} \setminus S$. We let $U(v)$ be simply $\mathbf{collar}(v) \times \mathbb{R}$.

The intersection of $\partial U(v)$ with $\Psi(M_\nu)$ is the union of all vertical annuli of blocks which are in the homotopy class of v . As we observed above for vertices, these annuli cover all of $\partial U(v)$.

Let \mathcal{U} be the union of the open tubes $U(v)$ we have described, so that \mathcal{U} is disjoint from $\Psi(M_\nu[0])$ and $\partial \mathcal{U}$ is contained in $\partial \Psi(M_\nu[0])$. From now on we identify $M_\nu[0]$ with its Ψ -image, and define $M_\nu = M_\nu[0] \cup \mathcal{U}$.

8.3. The model metric. We will describe a metric for each of the finitely many block types of which $M_\nu[0]$ is constructed, and use this to piece together




Figure 12. A schematic of a flat embedding of M_i^j into $S \times \mathbb{R}$. Note that some portions of blocks are stretched vertically. This picture is a fairly accurate rendition of the case $S = S_{0,5}$.

a metric on all of $M_v[0]$. Then we will discuss the structure of the tubes \mathcal{U} , and extend the metric to them as well.

Internal blocks. Fix one copy W_1 of a four-holed sphere, and one copy W_2 of a one-holed torus. Mark W_1 with a pair v_1^-, v_1^+ of adjacent vertices in $\mathcal{C}(W_1)$, and similarly v_2^-, v_2^+ for $\mathcal{C}(W_2)$. This determines two blocks $\hat{B}_k = B([v_k^- v_k^+])$, $k = 1, 2$. That is, \hat{B}_k is constructed from $W_k \times [-1, 1]$ by removing solid torus neighborhoods of $\gamma_{v_k^-} \times \{-1\}$ and $\gamma_{v_k^+} \times \{1\}$, as in the construction at the beginning of the section.

Let us define some *standard metrics* on the surfaces $S_{0,2}$, $S_{0,3}$, $S_{0,4}$ and $S_{1,1}$:

Call an annulus *standard* if it is isometric to $\mathbf{collar}(\gamma)$ for a geodesic γ of length $\epsilon_1/2$. Fix a three-holed sphere Y' with a hyperbolic metric so that $\partial Y'$ is geodesic with components of length $\epsilon_1/2$. Call a three-holed sphere *standard* if it is isometric to $Y' \setminus \mathbf{collar}(\partial Y')$.

Now fix a surface W'_k homeomorphic to W_k , and endowed with a fixed hyperbolic metric σ for which $\partial W'_k$ is geodesic with components of length $\epsilon_1/2$, and fix an identification of W_k with $W'_k \setminus \mathbf{collar}(\partial W'_k)$. We can do this, for specificity, in such a way that the curves $\gamma_{v_k^\pm}$ are identified with a pair of orthogonal geodesics of equal length. Finally, choose the identification so that the collars $\mathbf{collar}(v_k^\pm)$ in W_k (fixed by our global convention in §3.3) are identified with $\mathbf{collar}(v_k^\pm, \sigma)$ in W'_k . We call this metric *standard* on W_k .

Now we may fix a metric on \hat{B}_k with the following properties:

- (1) The metric restricts to standard metrics on $W_k \times \{0\}$, and on each 3-holed sphere in $\partial_\pm \hat{B}_k$.

- (2) Each annulus component of $\partial_{\parallel} \widehat{B}_k$ (respectively $\partial_i^{\pm} \widehat{B}_k$) is isometric to $S^1 \times [0, \epsilon_1]$ (resp. $S^1 \times [0, \epsilon_1/2]$), with S^1 normalized to length ϵ_1 , and this product structure agrees with the product structure imposed by the inclusion in $W_k \times [0, 1]$.

(The details of the construction do not matter, just that these properties hold and that a fixed choice is made. The length of ϵ_1 for the cores of all the Euclidean annuli is made possible by the definition of collars in §3.2.3.)

Note that, by definition of a standard metric, each component of the gluing boundary $\partial_{\pm} \widehat{B}_k$ admits an orientation-preserving isometry group realizing all six permutations of the three boundary components. This will enable us to glue the blocks via isometries.

We will call these two specific blocks the “standard blocks.” Every block $B(e)$ in $M_{\nu}[0]$ associated to a 4-edge e can be identified with one of \widehat{B}_1 or \widehat{B}_2 , depending on the homeomorphism type of $D(e)$, by a map that takes e_{\pm} to v_k^{\pm} . This identification is unique up to isotopy preserving the various parts of the boundary, and any such identification yields what we call a standard metric on $B(e)$.

Boundary blocks. Consider a top boundary block $B = B_{\text{top}}(\nu_R)$. Recall that the outer boundary $\partial_o B$ was constructed as $R \times \{0\} \cup \partial R \times [0, \infty)$, which is homeomorphic to $\text{int}(R)$. Endow it with the Poincaré metric σ_{∞} representing ν_R , in such a way that $\text{collar}(\partial R, \sigma_{\infty})$ is identified with $\partial R \times (0, \infty)$, and $\text{collar}(\mathbf{T}_R, \sigma_{\infty})$ is identified with $\text{collar}(\mathbf{T}_R) \times \{0\}$. Let σ_m be the conformal rescaling described in Section 3.4, which makes the collars of curves of length less than ϵ_1 into Euclidean cylinders. Note that (by definition of \mathbf{T}_R) these collars are either components of $\text{collar}(\mathbf{T}_R)$ or of $\text{collar}(\partial R)$.

Let σ'_{∞} be a hyperbolic metric on $\text{int}(R)$ for which every component of \mathbf{T}_R has length less than $\epsilon_1/2$, and which differs from σ_{∞} by a uniformly bilipschitz distortion (the constant can be chosen to depend only on ϵ_1/L_0). Let σ'_m be the conformal rescaling of σ'_{∞} that makes component of $\text{collar}(\mathbf{T}_R \cup \partial R)$ a Euclidean cylinder, and equals σ'_{∞} elsewhere. Then σ'_m/σ_m is uniformly bounded above and below.

Now we can transport $\sigma'_m|_{R \times \{0\}}$ to $R \times \{-1/2\}$ via the identity on the first factor. Extend these to a metric on $R \times [-1/2, 0]$ which is in uniformly bounded ratio with the product metric of σ_m and dt (here $t \in [-1, 0]$). The rest of the block, $(R \setminus \text{collar}(\mathbf{T}_R)) \times [-1, -1/2]$, may be metrized as follows: The restriction of σ'_m to each component Y of $R \setminus \text{collar}(\mathbf{T}_R)$ is uniformly bilipschitz equivalent to a standard metric on a three-holed sphere, since all the curves of \mathbf{T}_R have σ_{∞} -length at most L_0 by the choice of \mathbf{T}_R (see §7.1). Thus we may place a standard metric on the corresponding gluing surface $Y \times \{-1\}$, and interpolate between them on $Y \times [-1, -1/2]$ so that the resulting metric is uniformly bilipschitz equivalent to the product metric of $\sigma_{\infty}|_Y$ and dt .

We may do this, as for the internal blocks, in such a way that the annulus boundaries are given a flat Euclidean structure, and the natural product structure agrees with that inherited from $S \times \mathbb{R}$.

The metric on bottom boundary blocks is defined analogously.

LEMMA 8.3. *There is a metric on $M_v[0]$ which restricts on each block to a standard metric.*

Proof. By definition, each block admits a homeomorphism to a standard block, and we can pull back the metric via this identification. A gluing surface Y will then be mapped to the gluing surfaces of the standard block (which are all standard 3-holed spheres) in two possibly different ways by the identifications of the blocks on its two sides. However, in each isotopy class of self-homeomorphisms of a standard 3-holed sphere there is an isometry, by construction. Therefore, after adjustment by an appropriate isotopy, we may assume that the metrics on the gluing surfaces match. \square

Meridian coefficients. Let v be a nonparabolic vertex in H . The torus $\partial U(v)$ inherits a Euclidean metric from M_v , and a boundary orientation from $U(v)$. It also has a natural marking (α, μ) , where α is the homotopy class of the cores of the annuli making up $\partial U(v)$, homotopic to γ_v in S , and μ is the meridian class of $\partial U(v)$. To describe the meridian explicitly, recall that we represent $U(v)$ as a product $\mathbf{collar}(v) \times (s, t)$. If a is any simple arc in $\overline{\mathbf{collar}}(v)$ connecting the boundaries, then

$$(8.1) \quad \partial(a \times [s, t])$$

is a meridian. Note that the choice of a does not affect the isotopy class of this curve in $\partial U(v)$.

As in Section 3.2, we can describe the geometry of this oriented marked Euclidean torus with parameters (ω, t) . In this case $t = \epsilon_1$, since the circumference of the annuli in $\partial U(v)$ is ϵ_1 by construction. The Teichmüller parameter $\omega \in \mathbb{H}^2$ will be called the meridian coefficient of v , and denoted $\omega_M(v)$.

Note that $\epsilon_1 |\omega_M(v)|$ is the length of the meridian, and also that the imaginary part $\epsilon_1 \operatorname{Im} \omega$ is simply the sum of the heights of the annuli that make up $\partial U(v)$. We have $\operatorname{Im} \omega \geq 1$ since $\partial U(v)$ contains at least the annuli from its bottom and top blocks.

If $U(v)$ is a parabolic tube, we define $\omega_M = i\infty$.

Metriizing the tubes. For each nonparabolic tube $U(v)$, Lemma 3.2 gives us a unique hyperbolic tube $\mathbb{T}(\lambda, r)$ whose boundary parameters with respect to the natural marking, are $(\omega_M(v), \epsilon_1)$, so we identify this tube with $U(v)$ via a marking preserving isometry on the boundary. There is clearly a unique way to do this up to isotopy of $U(v)$.

$\partial U(v)$ for a parabolic v or a boundary component of S is an infinite Euclidean cylinder of circumference ϵ_1 , and there is a unique (up to isometry) rank-1 parabolic tube with circumference ϵ_1 , so we impose that metric on $U(v)$.

This completes the definition of the model metric on all of M_ν .

We remark that our eventual goal (in [18]) is to show that these hyperbolic tubes are in fact bilipschitz equivalent to the tubes in the hyperbolic manifold N_ρ . In this paper we will only obtain a Lipschitz map, and only for those tubes with $|\omega_M|$ bounded above by a certain constant.

Fillings of $M_\nu[0]$. As in the introduction, we define

$$\mathfrak{U}[k] = \bigcup_{|\omega_M(v)| \geq k} U(v)$$

and

$$M_\nu[k] \equiv M_\nu[0] \cup \bigcup_{|\omega_M(v)| < k} U(v).$$

Note that $M_\nu[\infty]$ indicates the inclusion of all the nonparabolic tubes.

9. Comparing meridian coefficients

We introduced the meridian coefficient $\omega_M(v)$ in Section 8, to describe the geometry of the model torus associated to a vertex v of a hierarchy H . Now we will describe two more ways of estimating this coefficient: one, $\omega_H(v)$, will be defined using the data in the hierarchy H_ν , and the third, $\omega_\nu(v)$, using the subsurface projection maps π_W applied directly to the end invariants ν .

We will prove that these invariants are close in the following sense:

THEOREM 9.1. *For any pair of end invariants ν and associated hierarchy H and model manifold M , and for any nonparabolic vertex v in H , we have bounds*

$$(9.1) \quad d_{\mathbb{H}^2}(\omega_H(v), \omega_M(v)) \leq D$$

and

$$(9.2) \quad d_{\mathbb{H}^2}(\omega_H(v), \omega_\nu(v)) \leq D$$

where D depends only on the topological type of S . Here $d_{\mathbb{H}^2}$ refers to the Poincaré metric in the upper half-plane.

Note that we are interpreting the ω 's as Teichmüller parameters, so this estimate is natural since $d_{\mathbb{H}^2}$ can be identified with the Teichmüller distance in Teichmüller space of the torus.

If v is a vertex of $\mathcal{C}(S)$, let us denote the two boundary components of $\text{collar}(v)$ (arbitrarily) as v_l and v_r . Now for $\alpha = v_r$ or v_l , define

$$(9.3) \quad X_\alpha = \{h \in H : \alpha \subset \partial D(h)\}.$$

Note that the same h can be in both X_{v_l} and X_{v_r} , if its domain borders $\mathbf{collar}(v)$ from both sides. We similarly define

$$(9.4) \quad X_{\alpha,k} = \{h \in H : h \in X_\alpha, \xi(D(h)) = k\}$$

and

$$(9.5) \quad X_{\alpha,k+} = \{h \in H : h \in X_\alpha, \xi(D(h)) \geq k\}.$$

Coefficients for internal vertices. From now until Section 9.7, we will assume that v is an *internal* vertex of H – that is, it is not a vertex of $\mathbf{simp}(\mathbf{I}(H))$ or $\mathbf{simp}(\mathbf{T}(H))$. Let h_v be the annulus geodesic in H with domain $\mathbf{collar}(v)$. We can then define:

$$(9.6) \quad \omega_H(v) = [h_v] + i \left(1 + \sum_{\alpha=v_l, v_r} \sum_{h \in X_{\alpha,4}} |h| \right).$$

(Here $[h_v]$ is the *signed* length of h_v , defined as in (4.5).)

For the next definition, again with $\alpha = v_l$ or v_r , let

$$(9.7) \quad \mathfrak{Y}_\alpha = \{Y \subset S : \alpha \subset \partial Y\}$$

(where we recall our convention of only taking the standard representatives of isotopy classes of surfaces, whose boundaries are collar boundaries). Define also $\mathfrak{Y}_{\alpha,4}$ and $\mathfrak{Y}_{\alpha,4+}$ in analogy with (9.4) and (9.5).

We now define

$$(9.8) \quad \omega_v(v) = \mathrm{tw}_v(v_-, v_+) + i \left(1 + \sum_{\alpha=v_l, v_r} \sum_{Y \in \mathfrak{Y}_{\alpha,4+}} \{d_Y(v_-, v_+)\}_K \right)$$

where K will be determined later, and $\{x\}_K$ is the threshold function, defined in Section 5.4. Here for convenience we have defined $d_Y(v_-, v_+) \equiv d_Y(\mu_-, \mu_+)$ and similarly for tw_Y , where μ_\pm are the generalized markings derived from v_\pm in Section 7.1.

9.1. Shearing. Let us recall from [49] the notion of *shearing outside a collar* for two hyperbolic metrics on a surface, and some of its properties. Let $\gamma \in \mathcal{C}_0(S)$, and suppose that σ_1 and σ_2 are two hyperbolic metrics on S for which $\mathbf{collar}(\gamma, \sigma_1) = \mathbf{collar}(\gamma, \sigma_2)$ (therefore denote both as $\mathbf{collar}(\gamma)$). We emphasize that σ_i are *actual metrics* rather than isotopy classes.

We define a quantity

$$\mathbf{shear}_\gamma(\sigma_1, \sigma_2)$$

as follows. Let \hat{Y} denote the compactified annular cover of S associated to γ (as in §4), let \hat{B} be the annular lift of $\mathbf{collar}(\gamma)$ to this cover, and let $\hat{\sigma}_1$ and $\hat{\sigma}_2$ denote

the lifts of σ_1 and σ_2 to $\text{int}(\widehat{Y})$. Define $\mathbf{shear}_\gamma(\sigma_1, \sigma_2)$ to be

$$\sup_{E, g_1, g_2} d_{\mathcal{A}(E)}(g_1 \cap E, g_2 \cap E).$$

Here E varies over the two complementary annuli E_1, E_2 of \widehat{B} in \widehat{Y} , and g_i (for $i = 1, 2$) varies over all arcs in \widehat{Y} connecting the two boundaries which are $\widehat{\sigma}_i$ -geodesic in $\text{int}(\widehat{Y})$. Thus we are measuring the relative twisting, outside \widehat{B} , of any two geodesics in the two metrics.

Note that the shear is *not* an approximation to the twisting parameter tw_γ . Rather, it is used for bounding the error in twisting that can accumulate outside of a collar, and our purpose will be to bound it from above (via Lemma 9.2), enabling us to measure twisting by restricting to a collar.

If α_i are simple closed σ_i -geodesics in S crossing γ (for $i = 1, 2$) and a_i are components of $\alpha_i \cap \overline{\mathbf{collar}}(\gamma)$ that cross $\mathbf{collar}(\gamma)$, then we will show

$$(9.9) \quad |\text{tw}_{\overline{\mathbf{collar}}(\gamma)}(a_1, a_2) - \text{tw}_\gamma(\alpha_1, \alpha_2)| \leq 2 \mathbf{shear}_\gamma(\sigma_1, \sigma_2) + 2$$

where we recall from Section 4 that $\text{tw}_{\overline{\mathbf{collar}}(\gamma)}(a_1, a_2)$ is a quantity that depends on the exact arcs of intersection with the collar, whereas $\text{tw}_\gamma(\alpha_1, \alpha_2)$ depends only on the homotopy classes of α_1 and α_2 .

Proof of (9.9). Let a'_i for $i = 1, 2$ be the σ_i -geodesic arc in \widehat{Y} that represents $\pi_Y(\alpha_i)$ – that is, a choice of (geodesic) lift of α_i to the annulus that connects the boundaries. By definition,

$$\text{tw}_\gamma(\alpha_1, \alpha_2) = \text{tw}_{\widehat{Y}}(a'_1, a'_2).$$

The right-hand side decomposes as the sum

$$\text{tw}_{E_1}(a'_1 \cap E_1, a'_2 \cap E_1) + \text{tw}_{\overline{B}}(a'_1 \cap \overline{B}, a'_2 \cap \overline{B}) + \text{tw}_{E_2}(a'_1 \cap E_2, a'_2 \cap E_2),$$

by two applications of the additivity property (4.3) (where \overline{B} is the closure of \widehat{B} , i.e. the lift of $\overline{\mathbf{collar}}(\gamma)$ to \widehat{Y}). The absolute values of the first and third terms are bounded by $\mathbf{shear}_\gamma(\sigma_1, \sigma_2)$, by the definition and inequality (4.1).

The arcs a_1 and a_2 in $\mathbf{collar}(\gamma)$ lift to two arcs in \overline{B} , which we still call a_1 and a_2 , and which are disjoint from or equal to $a'_1 \cap \overline{B}$ and $a'_2 \cap \overline{B}$, respectively. Thus $\text{tw}_{\overline{B}}(a_i, a'_i \cap \overline{B}) = 0$, and using the additivity inequality (4.4) twice, we conclude

$$|\text{tw}_{\overline{\mathbf{collar}}(\gamma)}(a_1, a_2) - \text{tw}_{\overline{B}}(a'_1 \cap \overline{B}, a'_2 \cap \overline{B})| \leq 2.$$

The estimate (9.9) follows. □

We can bound the shear between two metrics in the following setting:

LEMMA 9.2 ([49]). *Suppose R is a subsurface of S each component of which is convex in two hyperbolic metrics σ and τ , and that σ and τ are locally K -bi-lipschitz in the complement of R . Suppose that one component of R is an annulus*

B which is equal to both $\mathbf{collar}(\gamma, \sigma)$ and $\mathbf{collar}(\gamma, \tau)$ for a certain curve γ . Then

$$\mathbf{shear}_\gamma(\sigma, \tau) \leq \delta_0 K,$$

where δ_0 depends only on the topological type of S .

We remark that this lemma in [49] is stated for a slightly different definition of $\mathbf{collar}(\gamma, \sigma)$. However, the same argument applies for our definition, resulting in a different constant δ_0 .

Another property of shearing that follows immediately from the definition and the triangle inequality is subadditivity: If γ has the same collar with respect to three metrics σ_1, σ_2 and σ_3 , then

$$(9.10) \quad \mathbf{shear}_\gamma(\sigma_1, \sigma_3) \leq \mathbf{shear}_\gamma(\sigma_1, \sigma_2) + \mathbf{shear}_\gamma(\sigma_2, \sigma_3).$$

9.2. Sweeping through the model. Let us elaborate on the sequence of split-level surfaces $\{F_i\}_{i \in \mathcal{I}'}$ constructed in the proof of Theorem 8.1. Each F_i corresponds via the projection $S \times \mathbb{R} \rightarrow S$ to a subsurface $Z_i \subset S$ which is the complement of a union of annuli $\mathbf{collar}(\eta_i)$, where η_i is a pants decomposition coming from the resolution of the hierarchy.

We can add “middle” surfaces $F_{i+1/2}$ as follows: Each transition $F_i \rightarrow F_{i+1}$ corresponds to an edge e_i in a 4-geodesic, and hence a block $B_i = B(e_i)$, so that F_i and F_{i+1} agree except on $F_i \cap B_i = \partial_- B_i$ and $F_{i+1} \cap B_i = \partial_+ B_i$. Define the surface $F_{i+1/2}$ to be $F_i \cap F_{i+1}$ union the middle surface $D(e_i) \times \{0\}$. The projection of this to S , which we call $Z_{i+1/2}$, is the complement of $\mathbf{collar}(\eta_{i+1/2})$, where $\eta_{i+1/2}$ is $\eta_i \cap \eta_{i+1}$.

Now for s either integer or half-integer, Let σ'_s be the metric on Z_s inherited from F_s . By construction, this metric is “standard” on each component (see §8.3) and hence can be extended to a hyperbolic metric σ_s on all of S , such that the components of $\mathbf{collar}(\eta_s)$ are standard collars with respect to σ_s . Note that σ_s is not unique but we will fix some choice for each s .

Now consider the interval $J'(v) \subset \mathcal{I}'$, consisting of those i for which v is a component of η_i . Since v is internal, $t = \max J'(v) < \sup \mathcal{I}'$ and $b = \min J'(v) > \inf \mathcal{I}'$. Thus there is a transition $\eta_{b-1} \rightarrow \eta_b$ that replaces some vertex u by v , and a transition $\eta_t \rightarrow \eta_{t+1}$ replacing v by some w . The block $B_{b-1} = B(e_{b-1})$ contains the bottom annulus of $\partial U(v)$, and the block $B_t = B(e_t)$ contains the top annulus of $\partial U(v)$.

In order to understand how $\omega_M(v)$ is related to $\omega_H(v)$, we will have to understand something about the relation between the metrics $\sigma_{b-1/2}$ and $\sigma_{t+1/2}$:

LEMMA 9.3. $\mathbf{shear}_v(\sigma_{b-1/2}, \sigma_{t+1/2}) \leq K \operatorname{Im} \omega_M(v)$ where K depends only on the topological type of S .

Proof. We will apply Lemma 9.2 to the transitions $\sigma_s \rightarrow \sigma_{s+1/2}$, where s (for the rest of the proof) is an integer or half-integer in $[b - 1/2, t]$.

For every integer $s \in [b, t]$, $Z_{s+1/2}$ has a component $W(e_s)$ of type $\xi = 4$ which is the union of one or two components of Z_s and a collar. The metrics σ_s and $\sigma_{s+1/2}$ agree pointwise on $Z_s \setminus W(e_s) = Z_{s+1/2} \setminus W(e_s)$. On each component Y of Z_s which is contained in $W(e_s)$, the metrics σ_s and $\sigma_{s+1/2}$ are related by a uniform bilipschitz constant L , which only depends on our model construction. This is because these two metrics on Y come from subsurfaces of one of the blocks, and the identification is inherited from one of our finite number of standard blocks, where there is *some* bilipschitz constant. The corresponding statement holds for s and $s - 1/2$.

For each $s \in [b, t]$, let V_s denote the union of components (one or two) of Z_s that are adjacent to **collar**(v). Let Q denote the union of intervals $[s, s + 1/2]$ in $[b, t]$ for which $V_s = V_{s+1/2}$.

Let $[x, y]$ be a component of Q . Then $V_x = V_y$. Let R be the subsurface $S \setminus V_x$. Then R contains **collar**(v) as a component, and is convex in both metrics σ_x and σ_y , since each boundary component γ of R is a boundary component of the standard **collar**(γ) which is contained in R . Since every transition $s \rightarrow s + 1/2$ in $[x, y]$ does not involve $V_s = V_x$, we conclude that σ_x and σ_y are pointwise identical in the complement of R . Thus we may apply Lemma 9.2 to conclude that

$$(9.11) \quad \mathbf{shear}_v(\sigma_x, \sigma_y) \leq \delta_0.$$

Let $s \in [b - 1/2, t]$ be such that $(s, s + 1/2)$ is not in Q , and suppose first that s is an integer. Let R be the complement of V_s . This is still convex in both metrics σ_s and $\sigma_{s+1/2}$ just as in the previous paragraph, and has **collar**(v) as a component. This time the transition does involve V_s , so at least one component of V_s is contained in the component $W(e_s)$ of $Z_{s+1/2}$, and the two metrics on this component are related by a uniform bilipschitz constant L . On all other components they are identical, so we have a bilipschitz bound of L on the complement of R . Again applying Lemma 9.2 we obtain

$$(9.12) \quad \mathbf{shear}_v(\sigma_s, \sigma_{s+1/2}) \leq \delta_0 L.$$

(Note that this goes through correctly if $s = t$, even though $W(e_t)$ contains **collar**(v) in that case.)

The same bound holds if s is a half-integer, since then $s + 1/2$ is an integer in $[b, t]$ and we can apply the same argument in the opposite direction.

Thus, applying the subadditivity property (9.10), we conclude that

$$(9.13) \quad \mathbf{shear}_v(\sigma_{b-1/2}, \sigma_{t+1/2}) \leq \delta_0 \#Q + \delta_0 LN$$

where $\#Q$ is the number of components of Q , and N is the number of intervals $(s, s + 1/2)$ outside of Q . We note now that every such $(s, s + 1/2)$ with $s \neq b - 1/2, t$ corresponds to half of a block B with the property that $\partial_{\parallel} B$ contains an annulus on $\partial U(v)$ (in other words the domain subsurface of B is adjacent to $\mathbf{collar}(v)$). The intervals $(b - 1/2, b)$ and $(t, t + 1/2)$ correspond to the blocks meeting $U(v)$ on the bottom and top, respectively, in annuli of height $\epsilon_1/2$. The sum of heights of all annuli of $U(v)$ is exactly $\epsilon_1 \operatorname{Im} \omega_M(v)$, and hence we obtain $N \leq 2\epsilon_1 \operatorname{Im} \omega_M(v)$.

It is evident that $\#Q \leq N - 1$, and we conclude

$$(9.14) \quad \mathbf{shear}_v(\sigma_{b-1/2}, \sigma_{t+1/2}) \leq 2\delta_0(1 + L)\epsilon_1 \operatorname{Im} \omega_M(v). \quad \square$$

9.3. Comparing ω_H and ω_M . Recall now that u is the predecessor of v on the 4-edge e_{b-1} , and let u^- denote the geodesic representative of u in the metric $\sigma_{b-1/2}$. Note that the length of u^- in this metric is one of two values and hence uniformly bounded above and below. Similarly w is the successor of v on e_t , and we let w^+ be its geodesic representative in the metric $\sigma_{t+1/2}$, also with the same length bounds.

Let a^- be a component (there may be two) of the intersection of u^- with $\mathbf{collar}(v)$, and let a^+ be a component of the intersection of w^+ with $\mathbf{collar}(v)$. Writing $U(v)$ as $\mathbf{collar}(v) \times [p, q]$, we can use the curve

$$\mu = \partial(a^+ \times [p, q])$$

as a meridian. The real part of $\omega_M(v)$ is the amount of twisting of μ around the v direction in the torus, and this gives rise to the following estimate:

LEMMA 9.4.

$$|\operatorname{tw}_{\overline{\mathbf{collar}(v)}}(a^-, a^+) - \operatorname{Re} \omega_M(v)| = O(1).$$

Proof. Recalling the definition of the marked torus parameters in Section 3.2, and the discussion in Section 8.3, we have an orientation-preserving identification of the torus $\partial U(v)$ as the quotient of $\mathbb{C}/(\mathbb{Z} + \omega_M(v)\mathbb{Z})$, which is an isometry with respect to $1/\epsilon_1$ times the model metric on $\partial U(v)$. After possibly translating we may assume that the annulus $\mathbf{collar}(v) \times \{p\}$ lifts to a horizontal strip $V_1 = \{z : \operatorname{Im} z \in [0, k_1]\}$ and $\mathbf{collar}(v) \times \{q\}$ lifts to a horizontal strip $V_2 = \{z : \operatorname{Im} z \in [k_2, k_3]\}$ with $0 < k_1 < k_2 < k_3 < \operatorname{Im} \omega_M(v)$ (See Figure 13.)

The meridian μ lifts to four arcs in the cover (and their translates): $a^+ \times \{p\}$ lifts to an arc a_1^+ in V_1 connecting 0 to a point in $\mathbb{R} + ik_1$, $a^+ \times \{q\}$ lifts to an arc a_2^+ in V_2 connecting $\mathbb{R} + ik_2$ to $\mathbb{R} + ik_3$, and the arcs $\partial a^+ \times [p, q]$ lift to vertical arcs connecting a_1^+ to a_2^+ , and a_2^+ to $\omega_M(v)$, as shown in Figure 13.




Figure 13. A representative of the meridian of $\partial U(v)$, lifted to \mathbb{C} .

Thus, it is immediate that $\text{Re } \omega_M(v) = b_1 + b_2$, where b_i is the real part of the vector from the bottom to the top of a_i^+ . Since the length of a_i^+ in the metric $\sigma_{t+1/2}$ is $O(1)$, the same is true for its length in the annulus $\text{collar}(v) \times \{q\}$, and hence $b_2 = O(1)$. On the other hand, b_1 may be large because in the lower annulus $\text{collar}(v) \times \{p\}$, it is a^- that has length $O(1)$, and a^+ twists around a^- possibly many times. The number of twists is determined by $|\text{tw}_{\overline{\text{collar}(v)}}(a^-, a^+)|$, and the sign convention is such (see §4) that $b_1 = \text{tw}_{\overline{\text{collar}(v)}}(a^-, a^+) + O(1)$. The lemma follows. \square

Now we note, applying (9.9) and Lemma 9.3, that

$$(9.15) \quad \begin{aligned} |\text{tw}_v(u, w) - \text{tw}_{\overline{\text{collar}(v)}}(a^-, a^+)| &= O(\text{shear}_v(\sigma_{b-1/2}, \sigma_{t+1/2})) \\ &= O(\text{Im } \omega_M(v)). \end{aligned}$$

We also have

$$(9.16) \quad |\text{tw}_v(u, w) - [h_v]| \leq 1$$

by the definition of annulus geodesics in hierarchies, and inequality (4.1). Recall that $[h_v] = \text{Re } \omega_H(v)$.

Combining these with Lemma 9.4, we have

$$(9.17) \quad |\text{Re } \omega_M(v) - \text{Re } \omega_H(v)| = O(\text{Im } \omega_M(v)).$$

(Here we are using $\text{Im } \omega_M(v) \geq 1$ to subsume $O(1)$ terms into the $O(\text{Im } \omega_M(v))$ term.)

We also have

$$(9.18) \quad \text{Im } \omega_H(v) = \text{Im } \omega_M(v).$$

This is because the sums $\sum_{h \in X_{\alpha,4}} |h|$ in equation (9.6) for ω_H count the number of 4-edges associated to domains bordering $\text{collar}(v)$, counting an edge once for each side that its domain borders. This therefore gives exactly the number of vertical annuli in $\partial U(v)$, each of which has height ϵ_1 . Counting also the annuli from the bottom and top blocks for $U(v)$, whose heights add to ϵ_1 , and recalling that $\epsilon_1 \text{Im } \omega_M(v)$ is the sum of these heights, we have the equality (9.18).

It follows immediately from (9.17) and (9.18) that the distance in \mathbb{H}^2 between $\omega_M(v)$ and $\omega_H(v)$ is bounded. This gives inequality (9.1) of Theorem 9.1.

9.4. Counting in a hierarchy. In order to compare ω_v and ω_H , we must consider more carefully the structure of a hierarchy, and prove some counting lemmas that allow us to estimate the “size” of a hierarchy or certain subsets of it by various approximations.

Recall that if $f \searrow^d g$, then there is a unique simplex v of g for which $D(f)$ is a component domain of $(D(g), v)$. Let us say in this situation that $f \searrow^d g$ at v . Call v an *interior simplex* of g if it is neither the first or the last. If v is interior, then there exactly one $f \searrow^d g$ at v , namely the one whose domain contains the predecessor and successor of v . If v is last, then it is a single vertex (by construction; see Definition 5.2) and there are at most three component domains of $(D(g), v)$ supporting a geodesic $f \searrow^d g$ (including the annulus $\text{collar}(v)$). If v is the first vertex, then there are at most 2 such f 's, since the option of $f \searrow^d g$ supported on $\text{collar}(v)$ can only occur when $\xi(g) = 4$. This gives us:

$$(9.19) \quad |g| - 1 \leq \#\{h \searrow^d g\} \leq |g| + 4.$$

Now consider subsets $X \subseteq H$ with the following property, for some fixed number M :

- (*) If $g \notin X$, then there are at most M geodesics $h \searrow^d g$ for which there exists $h' \in X$, $h' \succeq h$.

The two main examples of X satisfying this property are $X = H$ (trivially), and, of interest to us:

LEMMA 9.5. *For any vertex v in $\mathcal{C}(S)$, $\alpha = v_r$ or v_l , and $k \geq 4$, the sets X_α and $X_{\alpha,k+}$ satisfy property (*) with M depending only on S .*

Proof. Let $g \notin X_\alpha$. If $h' \succeq h \searrow^d g$, then $D(h') \subset D(g)$, so if $h' \in X_\alpha$, then α is in $\partial D(h')$ and hence in $D(g)$, where it must be nonperipheral since $g \notin X_\alpha$. Now $D(h)$ is a component domain for some simplex u of g , and since $D(h') \subseteq D(h)$, we have $d_{D(g)}(u, v) \leq 1$. It follows that u is restricted to an interval of diameter 2

in g , so there are at most 3 possibilities for u . The discussion leading to (9.19) then yields at most $M = 6$ possibilities for h .

Now to prove the property for $X_{\alpha,k+}$: if $g \notin X_{\alpha,k+}$, then either $g \notin X_\alpha$, in which case the estimate follows from property (*) for X_α together with the fact that $X_{\alpha,k+} \subset X_\alpha$, or $\xi(D(g)) < k$ in which case there is no $h \succeq g$ in $X_{\alpha,k+}$, and the estimate is trivial. \square

Recalling the notation “ $x \approx_{a,b} y$ ” from Section 5.4, we can state:

LEMMA 9.6. *If $X \subseteq H$ satisfies (*) and $\varphi : X \rightarrow \mathbb{R}_+$ is a function satisfying*

$$\varphi(h) \approx_{a,b} |h|$$

for all $h \in X$, then

$$\sum_{h \in X} |h| \approx_{A,B} \sum_{h \in X} \varphi(h)$$

where A, B depend only on a, b, S and the constant M in (*).

The point of this lemma is that, although the additive errors b can accumulate when X has many members, the additive errors at one level are swallowed up by the multiplicative constant at a lower level.

Proof. Define

$$(9.20) \quad \beta(g) = \sum_{\substack{h \in X \\ h \succeq g}} |h|, \quad \beta'(g) = \sum_{\substack{h \in X \\ h \succeq g}} \varphi(h).$$

We will show inductively for each $m \leq \xi(S)$ that $\beta(g) \approx_{a',b'} \beta'(g)$ for all g with $\xi(D(g)) \leq m$, where a', b' depend on m . The lemma then follows from setting $m = \xi(S)$. The base case, $m = 2$, is immediate from the hypothesis $\varphi(g) \approx_{a,b} |g|$, with $(a', b') = (a, b)$.

Now note that, by Theorem 5.6, whenever $f \searrow g$ there exists a unique h with $f \succeq h \searrow^d g$. This allows us to inductively decompose β and β' :

$$(9.21) \quad \beta(g) = \begin{cases} |g| + \sum_{h \searrow^d g} \beta(h) & g \in X \\ \sum_{h \searrow^d g} \beta(h) & g \notin X \end{cases}$$

and similarly with β' replacing β and $\varphi(g)$ replacing $|g|$.

For both β and β' , in case $g \in X$ there are at most $|g| + 4$ terms in the summation, by (9.19). In case $g \notin X$, there are at most M nonzero terms in the summation, by property (*).

Now we can compare β and β' . Suppose first that $g \in X$. We have:

$$\beta(g) = |g| + \sum_{h \preceq g} \beta(h)$$

and by the inductive hypothesis:

$$\leq |g| + \sum_{h \preceq g} a' \beta'(h) + b';$$

then by (9.19):

$$\leq (1 + b')|g| + 4b' + a' \sum_{h \preceq g} \beta'(h)$$

and using $\varphi(g) \approx |g|$,

$$\begin{aligned} &\leq (a\varphi(g) + b)(1 + b') + 4b' + a' \sum_{h \preceq g} \beta'(h) \\ &= a(1 + b')\varphi(g) + a' \sum_{h \preceq g} \beta'(h) + b(1 + b') + 4b' \\ &\leq a'' \beta'(g) + b'' \end{aligned}$$

where the last line follows from the β' version of (9.21), using

$$a'' = \max(a(1 + b'), a') \text{ and } b'' = b(1 + b') + 4b'.$$

In case $g \notin X$, we have:

$$\begin{aligned} \beta(g) &= \sum_{h \preceq g} \beta(h) \\ &\leq \sum_{\substack{h \preceq g \\ \beta(h) > 0}} a' \beta'(h) + b' \end{aligned}$$

and by property (*):

$$\begin{aligned} &\leq b'M + a' \sum_{h \preceq g} \beta'(h) \\ &= a' \beta'(g) + b'M. \end{aligned}$$

The inequality $\beta'(g) \lesssim \beta(g)$ is obtained in the same way (with slightly different constants). \square

Counting with top-level domains. In the following proposition we show that the size of $X_{\alpha,4+}$ can be estimated by the size of $X_{\alpha,4}$. This is in keeping with the intuition that the moves in the level 4 domains are the places where “real” change happens, and all the rest is a bounded amount of bookkeeping.

PROPOSITION 9.7. *For any $x \in \mathcal{C}_0(S)$ and $\alpha = x_1$ or x_r , we have*

$$\sum_{h \in X_{\alpha,4}} |h| \approx \sum_{h \in X_{\alpha,4+}} |h|.$$

Proof. The proof has the same inductive structure as the proof of Lemma 9.6.

Let

$$\gamma(g) = \sum_{\substack{h \rhd g \\ h \in X_{\alpha,4}}} |h|, \quad \gamma'(g) = \sum_{\substack{h \rhd g \\ h \in X_{\alpha,4+}}} |h|.$$

We shall inductively prove $\gamma(g) \approx_{a,b} \gamma'(g)$ for constants a, b depending on $\xi(D(g))$. First, if $\xi(D(g)) \leq 4$, then $\gamma(g) = \gamma'(g)$ is obvious. It is also obvious that $\gamma \leq \gamma'$. Now assume $\gamma'(h) \lesssim_{a,b} \gamma(h)$ for $\xi(D(h)) < \xi(D(g))$ and let us prove it for g .

We first need a few lemmas.

LEMMA 9.8. *Suppose that $h \rhd^d g$ at an interior simplex v . If $\xi(D(h)) > 4$ $|h| \geq 3$, and if $\xi(D(h)) = 4$, then $|h| \geq 1$.*

Proof. By assumption, the predecessor and successor u and w of v must be contained in $D(h)$ and by tightness of g , they fill $D(h)$. If $\xi(D(h)) > 4$, this implies that their distance in $\mathcal{C}_1(D(h))$ is at least 3. If $\xi(D(h)) = 4$ it just means they are distinct, so the distance is at least 1. \square

LEMMA 9.9. *Suppose $g \in X_{\alpha}$ and $\xi(D(g)) > 4$. If $m \rhd^d g$ at v and $m' \rhd^d g$ at v' , where v' is the successor of v and both are interior in g , then at least one of m and m' is in $X_{\alpha,4+}$.*

Proof. Note that $D(m)$ and $D(m')$ cannot be annuli since v and v' are interior vertices and $\xi(D(g)) > 4$. Thus it suffices to show one of them is in X_{α} . Suppose $m' \notin X_{\alpha}$, i.e. $\partial D(m')$ does not contain α . Since v' is not last or first, v must be contained in $D(m')$. It follows that v' separates v from α . This means that the component domain of $(D(g), v)$ that meets α must contain v' . This domain is $D(m)$, so we conclude $m \in X_{\alpha}$. \square

Finally we have:

LEMMA 9.10. *If $h \in X_{\alpha,4+}$ and $|h| \geq 3$, then $\gamma(h) \geq 1$.*

Proof. We proceed by induction. If $\xi(D(h)) = 4$ the statement is obvious, in fact $\gamma(h) = |h|$. Suppose $\xi(D(h)) > 4$. Since $|h| \geq 3$ it has at least two interior simplices, and by Lemma 9.9 for at least one of them there is a $k \rhd^d h$ with $k \in X_{\alpha,4+}$. If $\xi(D(k)) > 4$, then, by Lemma 9.8, $|k| \geq 3$. Thus by induction $\gamma(k) \geq 1$. If $\xi(D(k)) = 4$, then $\gamma(k) = |k|$, and Lemma 9.8 gives $|k| \geq 1$. Since clearly $\gamma(h) \geq \gamma(k)$, we are done. \square

We return to the proof of Proposition 9.7. Recall that we are in the case $\xi(D(g)) > 4$, so in particular

$$(9.22) \quad \gamma(g) = \sum_{h \triangleleft^d g} \gamma(h).$$

Now suppose that $g \in X_\alpha$. The recursive formula for γ' gives

$$\gamma'(g) = |g| + \sum_{h \triangleleft^d g} \gamma'(h).$$

By the inductive hypothesis this is

$$\leq |g| + \sum_{h \triangleleft^d g} a\gamma(h) + b.$$

Using (9.19), we obtain

$$(9.23) \quad \begin{aligned} &\leq 4b + |g|(b+1) + \sum_{h \triangleleft^d g} a\gamma(h) \\ &= 4b + |g|(b+1) + a\gamma(g). \end{aligned}$$

By (9.22). If $|g| \leq 2$ this becomes

$$\leq a\gamma(g) + 2 + 6b$$

and we are done. Now suppose that $|g| \geq 3$. By Lemma 9.9, for least $\lfloor \frac{|g|-1}{2} \rfloor$ of the interior simplices of g , the corresponding $h \triangleleft^d g$ are in $X_{\alpha,4+}$. Applying Lemmas 9.8 and 9.10, we have $\gamma(h) \geq 1$ for those h . This tells us that

$$\begin{aligned} \gamma(g) &= \sum_{h \triangleleft^d g} \gamma(h) \geq \#\{h \triangleleft^d g : \gamma(h) \geq 1\} \\ &\geq \frac{|g|-3}{2}. \end{aligned}$$

Hence $|g| \leq 3 + 2\gamma(g)$, so that (9.23) becomes

$$\gamma'(g) \leq 7b + 3 + (2(b+1) + a)\gamma(g).$$

Finally if $g \notin X_\alpha$, we use property (*) as in the previous section to argue

$$\begin{aligned} \gamma'(g) &= \sum_{h \triangleleft^d g} \gamma'(h) \\ &\leq \sum_{\substack{h \triangleleft^d g \\ \gamma'(h) > 0}} a\gamma(h) + b \\ &\leq Mb + a\gamma(g). \end{aligned}$$

This completes the inductive step, and establishes the proposition. \square

9.5. Comparing ω_H and ω_v . A uniform bound on the difference between real parts, $|\operatorname{Re} \omega_H - \operatorname{Re} \omega_v|$, follows directly from Lemma 5.11, which compares $[h_v]$ to $\operatorname{tw}_v(\mathbf{I}(H), \mathbf{T}(H)) = \operatorname{tw}_v(v_-, v_+)$. Since the imaginary parts are at least 1, this gives us

$$(9.24) \quad |\operatorname{Re} \omega_v(v) - \operatorname{Re} \omega_H(v)| = O(\operatorname{Im} \omega_H(v)).$$

If we can establish a bound of the form

$$(9.25) \quad \frac{1}{c} \leq \frac{\operatorname{Im} \omega_H(v)}{\operatorname{Im} \omega_v(v)} \leq c$$

for some uniform c , then a bound on $d_{\mathbb{H}^2}(\omega_v(v), \omega_H(v))$ will follow.

It will suffice to establish

$$(9.26) \quad \sum_{Y \in \mathfrak{Y}_{\alpha, 4+}} \{d_Y(v_-, v_+)\}_K \approx \sum_{h \in X_{\alpha, 4}} |h|$$

where $\alpha = v_l$ or v_r . For convenience let $d_Y \equiv d_Y(v_-, v_+)$ throughout this proof. Choose $K > M_2$, the constant in Lemma 5.9, so that by that lemma if $d_Y \geq K$, then Y is the support of some geodesic in H . With this choice, we note that

$$\sum_{Y \in \mathfrak{Y}_{\alpha, 4+}} \{d_Y\}_K = \sum_{h \in X_{\alpha, 4+}} \{d_{D(h)}\}_K.$$

Now since $\{d_{D(h)}\}_K \approx d_{D(h)}$ and $d_{D(h)} \approx |h|$ by Lemma 5.9, and since $X_{\alpha, 4+}$ satisfies property $(*)$ (Lemma 9.5), Lemma 9.6 gives us

$$\approx \sum_{h \in X_{\alpha, 4+}} |h|$$

and then Proposition 9.7 gives us

$$\approx \sum_{h \in X_{\alpha, 4}} |h|.$$

This establishes (9.26), and since both sides of this estimate are positive and $\operatorname{Im} \omega_v$ and $\operatorname{Im} \omega_H$ are obtained from them by adding 1, (9.25) follows.

9.6. Global projection bounds. The same type of counting arguments that led to Theorem 9.1 can give us the following *a priori* bound on $|\omega_H|$ from bounds on projections $d_Y(v_+, v_-)$. Define $\mathfrak{Y}_v = \mathfrak{Y}_{v_l} \cup \mathfrak{Y}_{v_r}$.

THEOREM 9.11. *Given end invariants v and an associated hierarchy H , for any internal vertex v in H we have*

$$|\omega_H(v)| \leq C \left(1 + \sup_{Y \in \mathfrak{Y}_v} d_Y(v_-, v_+) \right)^a$$

where the constants $a, C > 0$ depend only on S .

Proof. Let

$$B = 1 + \sup_{Y \in \mathcal{Y}_v} d_Y(v_-, v_+).$$

It suffices to find a bound of the form

$$(9.27) \quad \beta(h) \leq c_j B^j$$

where $j = \xi(D(h)) - 1$, and β is defined as in (9.20) with $X = X_\alpha$ ($\alpha = v_l$ or v_r). Since $|\operatorname{Re} \omega_H(v)|$ is the $\xi = 2$ term of $\beta(g_H)$ and $\operatorname{Im} \omega_H(v)$ is the sum of terms from $\beta(g_H)$ with $\xi > 2$, the bound (9.27) applied to $\beta(g_H)$ gives us a bound of the form $|\omega_H(v)| = O(B^{\xi(S)-1})$, which proves Theorem 9.11.

By choice of B , and Lemma 5.9, we have

$$|g| \leq B + 2M_1$$

for each $g \in X_\alpha$. Choose c_1 such that $B + 2M_1 + 4 < c_1 B$ (for all $B \geq 1$). Then for $\xi(D(h)) = 2$ we have

$$\beta(h) = |h| \leq B + 2M_1 < c_1 B$$

if $D(h) = \mathbf{collar}(v)$ and $\beta(h) = 0$ otherwise. This establishes the base case of (9.27).

Now assume (9.27) for $\xi(D(h)) < \xi(D(g))$ and let us prove it for $\xi(D(g))$. Let $j = \xi(D(g)) - 1$. Suppose first $g \in X_\alpha$. Using the recursive formula for $\beta(g)$ and then (9.19),

$$\begin{aligned} \beta(g) &= |g| + \sum_{h \preceq g} \beta(h) \\ &\leq B + 2M_1 + (|g| + 4)c_{j-1}B^{j-1} \\ &\leq B + 2M_1 + (B + 2M_1 + 4)c_{j-1}B^{j-1} \\ &\leq c_1 B + c_1 c_{j-1} B^j \\ &\leq 2c_1 c_{j-1} B^j. \end{aligned}$$

For $g \notin X_\alpha$, we have, using property (*),

$$\begin{aligned} \beta(g) &= \sum_{h \preceq g} \beta(h) \\ &\leq M c_{j-1} B^{j-1}. \end{aligned}$$

Thus by induction we have established (9.27), and the theorem. \square

9.7. Noninternal vertices. If v is not an internal vertex of H , the definitions of the coefficients must be adjusted somewhat before proving Theorem 9.1. Let us first dispense with the cases *not* included in that theorem, and which correspond to cusps in the hyperbolic 3-manifold:

- If v is *parabolic* in $\mathbf{I}(H)$ or $\mathbf{T}(H)$, then $U(v)$ is an unbounded solid torus, and we define $\omega_M(v) = \omega_H(v) = \omega_v(v) = i\infty$.
- A component of ∂S cannot be a vertex of H , but it does have an associated solid torus U , and it is convenient again to define all the coefficients to be $i\infty$.

We are left with the case that v is a vertex of $\text{base}(\mathbf{I}(H))$ which *does* have a transversal, or similarly for $\text{base}(\mathbf{T}(H))$ (or both).

In these cases v appears as a curve of length at most L_0 in the top or bottom (or both) of $\partial_\infty N$. Recall from Section 2.2 the conformal structure v_+^T on the top conformal boundary, and let σ_m be the rescaled metric, as in Section 3.4, which makes the thin collars Euclidean. In particular, if v is sufficiently short in the top boundary, then $\text{collar}(v, v_+^T)$ is Euclidean in σ_m , and we let $r_+(v)$ denote $1/\epsilon_1$ times its height. If not, then let $r_+(v) = 0$. Define $r_-(v)$ similarly. Now we can redefine $\omega_v(v)$ and $\omega_H(v)$, by adding the term

$$i(r_+(v) + r_-(v))$$

to the expressions in (9.6) and (9.8). The definition of ω_M is unchanged from Section 8.3.

The bound (9.2) on $d_{\mathbb{H}^2}(\omega_H(v), \omega_v(v))$ is now immediate, since we have added the same thing to both imaginary parts.

In order to bound $d_{\mathbb{H}^2}(\omega_H(v), \omega_M(v))$, we have to reconsider the sweeping discussion of Section 9.2:

Suppose that $v \in \text{base}(\mathbf{T}(H))$, and let R denote the component of R_+^T (see §2.2) that contains v . Thus there is a top block $B = B_{\text{top}}(v_R)$ associated to R . The interval $J'(v)$, as defined in Section 9.2, may be infinite because there may be infinitely many $i \in \mathcal{I}'$ such that v is a component of η_i . However, as in Section 8.2, there is a first point $i \in \mathcal{I}'$ such that F_i contains all of $\partial_- B$. Redefine $J'(v)$ so that $t = \max J'(v)$ is this value of i . The block B is then attached to F_t independently of the rest of the sweep, and so we define $F_{t+1/2}$ to be $F_t \setminus \partial_- B_{\text{top}}(v_R)$, union $\partial_o B'$ (where we recall from §8.1 that B' is B minus the cusp annuli $\partial R \times [0, \infty)$), and $\partial_o B' = R \times \{0\}$). Thus $\partial_o B$, in the isotopy class of R , can play the role of the middle surface $D(e_t) \times \{0\}$ in Section 9.2. We project $F_{t+1/2}$ to $Z_{t+1/2} \subset S$, and define $\sigma'_{t+1/2}$ on the R component of $Z_{t+1/2}$ to be the projection of $\sigma_\infty|_{\partial_o B}$, the Poincaré metric associated to v_R , as in Section 8.3. We define $\sigma_{t+1/2}$ as before by extending across collars, though we may have to make a small (uniform) adjustment to σ' first near ∂R , since σ_∞ has cusps rather than compact collars associated with ∂R .

If $v \in \text{base}(\mathbf{I}(H))$ we may use the bottom blocks to analogously define F_b , $F_{b-1/2}$, $\sigma_{b-1/2}$, etc. (and otherwise we use the same definition as in §9.2).

The shear computation that yields Lemma 9.3 now goes through as before – the argument, and particularly the appeal to Lemma 9.2, is insensitive to the fact that $\mathbf{collar}(v)$ may now have large radius in $\sigma_{b-1/2}$ and/or $\sigma_{t+1/2}$.

In order to complete the comparison of $\omega_M(v)$ and $\omega_H(v)$ we need to consider changes to Section 9.3, and particularly to Lemma 9.4. If $v \in \text{base}(\mathbf{T}(H))$ then the curve w should be chosen to be the transversal of v in the marking $\mathbf{T}(H)$ – recall that this means that w is a minimal-length curve crossing v with respect to the σ_∞ metric. Letting w^+ be its geodesic representative in $\sigma_{t+1/2}$ and letting a_+ be a component of $w^+ \cap \mathbf{collar}(v)$, we note that we no longer have an upper bound on the length of a_+ , but that its twist in $\mathbf{collar}(v)$ with respect to a geodesic arc orthogonal to the boundaries is at most 2, by the minimal-length choice of w . This means that, lifting $\partial U(v)$ to \mathbb{C} and using the notation of Section 9.3, that the real part b_2 of the vector associated to a_2^+ is still $O(1)$ (although the imaginary part may be large). A similar argument applies when $v \in \text{base}(\mathbf{I}(H))$ and we choose u to be the transversal to v from μ_- . The proof of Lemma 9.4 then proceeds as before. Thus we obtain (9.17) as before, namely

$$|\text{Re } \omega_M(v) - \text{Re } \omega_H(v)| = O(\text{Im } \omega_M(v)).$$

The equality (9.18), slightly adjusted to

$$\text{Im } \omega_H(v) = \text{Im } \omega_M(v) \pm O(1),$$

follows as before, using the additional information that the top (resp. bottom) annulus of $\partial U(v)$ has width $r_+(v)$ (resp. $r_-(v)$), and that these are the quantities added by definition to $\text{Im } \omega_H(v)$. Thus the estimate on $d_{\mathbb{H}^2}(\omega_H(v), \omega_M(v))$ follows as before.

Theorem 9.11, which is stated for the internal case, is easily generalized to yield

$$(9.28) \quad |\omega_H(v)| \leq C \left(\left(1 + \sup_{\substack{Y \subset \widehat{S} \\ v \in [\partial Y]}} d_Y(v_-, v_+) \right)^a + r_+(v) + r_-(v) \right).$$

The proof is exactly the same counting argument as for the internal case, with the added imaginary part in the definition of $\omega_H(v)$ yielding the extra terms.

10. The Lipschitz model map

We are now ready to prove the Lipschitz Model Theorem, whose statement appears in the introduction. We will build the map f in several stages.

Step 0. Let $f_0 : \widehat{S} \times \mathbb{R} \rightarrow N$ be any map in the homotopy class determined by ρ . Via the embedding $M_\nu \subset \widehat{S} \times \mathbb{R}$ we restrict f_0 to M_ν .

Step 1. On each three-holed sphere Y appearing as the gluing boundary of a block, we find a homotopy from f_0 to f_1 so that f_1 is uniformly Lipschitz. This

is done as follows: There is a map $p_Y : Y \rightarrow N$, homotopic to $f_0|_Y$, and which is pleated on ∂Y (with the usual proviso if a component of ∂Y corresponds to a cusp). Let σ_Y be the metric induced by p_Y on $\text{int}(Y)$, and let $Y_0 = Y \setminus \mathbf{collar}(\partial Y, \sigma_Y)$. Then, since Lemma 7.9 gives us a uniform upper bound for the σ_Y -lengths of ∂Y , there is a uniform K_0 such that there is a homeomorphism $\varphi_Y : Y \rightarrow Y_0$ homotopic to the identity, which is K_0 -bilipschitz with respect to the model metric on Y and the induced metric σ_Y on Y_0 . Now define

$$f_1|_Y = p_Y \circ \varphi_Y.$$

This is clearly a K_0 -Lipschitz map, and homotopic to $f_0|_Y$.

Step 2. Extend f_1 to a map f_2 , defined on the middle surfaces of internal blocks, and on the outer boundaries of boundary blocks.

For any internal block $B = B(e)$, consider the middle subsurface $W = D(e) \times \{0\}$ in B (in its original identification as a subset of $D(e) \times [-1, 1]$). Extend ∂W to a curve system v_0 which cuts S into components that are either 3-holed spheres, or W itself. Then $v^- = v_0 \cup e^-$ and $v^+ = v_0 \cup e^+$ are pants decompositions differing by an elementary move, and we can define a *halfway surface* $g_{v^-, v^+} \in \mathbf{pleat}_\rho(v_0)$, as in Section 3.1. Let σ_e be the metric induced by this map on $\text{int}(W)$. Now, Lemma 7.9 gives us a uniform bound

$$\ell_\rho(x) \leq D$$

where x is any component of ∂W , or e^\pm , since these are all vertices in the hierarchy. Lemma 3.1 then gives us uniform upper bounds in the induced metric σ_e :

$$\ell_{\sigma_e}(e^\pm) \leq D + C.$$

This immediately gives lower bounds on $\ell_{\sigma_e}(e^\pm)$ as well, since the curves cross each other. Furthermore it means that W under σ_e is not far from being our “standard” hyperbolic structure on W . More precisely, letting $W_0 = W \setminus \mathbf{collar}(\partial W, \sigma_e)$ there is an identification $\varphi_W : W \rightarrow W_0$ which is K_1 -bilipschitz with respect to the model metric on W and σ_e on W_0 . Hence as in the previous paragraph,

$$f_2|_W = g_{v^-, v^+} \circ \varphi_W$$

is the desired definition, and is clearly homotopic to $f_1|_W$.

Let B be a boundary block associated to a subsurface R , and let us define f_2 on its outer boundary $\partial_o B$. Recall that $\partial_o B$ is identified with a component R_∞ of $\partial_\infty N$, and its metric is equal to the rescaled Poincaré metric, σ_m . In Lemma 3.4 we obtain a map $\varphi : E_v \rightarrow E_N$ which, restricted to the bottom boundary of E_v (identified with $\partial_\infty N$) gives a uniformly bilipschitz homeomorphism from R_∞ to the corresponding component of $\partial \widehat{C}_N$.




Figure 14. A homotopy of a bounded figure-8 cannot be too long

By composing this map with the identification $\partial_o B \rightarrow R_\infty$, we obtain a homeomorphism from $\partial_o B$ to the corresponding component of $\partial \widehat{C}_N$, satisfying uniform bilipschitz bounds. This is our desired map f_2 .

Step 3. We can extend f_2 to a continuous map f_3 defined on all of M_ν , and still homotopic to f_0 . This is possible by an elementary argument because each of the surfaces where f_2 is defined so far is collared in M_ν . By appropriately defining the extension on these collars using the homotopy from f_2 to the restriction of f_0 , we can then extend to the whole manifold.

Step 4. We “straighten” f_3 on $M_\nu[0]$ to obtain a map f_4 , as follows:

Fix an internal block $B = B(e)$. Let f_4 agree with f_3 on the gluing boundaries $\partial_\pm B$, and on the middle surface $W = D(e) \times \{0\}$.

On $D(e) \times [-1/2, 1/2]$, define $f_4 = f_3 \circ q$ where $q(y, t) = (y, 0)$. This is certainly Lipschitz with a uniform constant. Now consider a component of $\partial_+ B$, which has the form $Y \times \{1\}$ where Y is a component of $D(e) \setminus \text{collar}(e^+)$. f_4 is already defined on $Y \times \{1/2, 1\}$, so for each $x \in Y$ extend f_4 to $\{x\} \times [1/2, 1]$ as the unique constant speed parametrization of the geodesic connecting the endpoints in the homotopy class determined by f_3 (negative curvature is used here for the definition to be unique, and for f_4 to be continuous and homotopic to f_3). This extends f_4 to $Y \times [1/2, 1]$.

To prove that $f_4|_{Y \times [1/2, 1]}$ is uniformly Lipschitz it suffices to find an upper bound on the lengths of the geodesics $f_4(\{x\} \times [1/2, 1])$. This is exactly the “figure-8” argument of [47, Lemma 9.3]. That is, since f_4 is already K_1 -Lipschitz on $Y \times \{1/2, 1\}$, there is a figure-8 $X \subset Y$, or bouquet of two circles, which is a deformation retract of Y with bounded tracks (see Figure 14), and such that $f_4(X \times \{1/2\})$ and $f_4(X \times \{1\})$ have uniformly bounded lengths (depending on K_1). If $f_4|_{X \times [1/2, 1]}$ has a very long trajectory then at some middle point $t \in [1/2, 1]$ the two loops of $f_4(X \times \{t\})$ are either both very short or nearly parallel to each other. In either case discreteness, or Jørgensen’s inequality, is violated.

For a component $Y \times \{-1\}$ of $\partial_- B$ the definition of f_4 on $Y \times [-1, -1/2]$ is analogous.

The very same discussion works for boundary blocks, except that the outer boundary (minus its vertical annuli) takes the role of the middle surface.

The map f_4 is now defined on all of $M_\nu[0]$, uniformly Lipschitz (with respect to the path metric on $M_\nu[0]$), and homotopic to f_0 .

Step 5. We adjust f_4 slightly to obtain a map f_5 satisfying:

- (1) $f_5(M_\nu[0]) \subset \widehat{C}_N$
- (2) Whenever $\ell_\rho(v) \leq \epsilon_1$ for a vertex v of H_ν , $f_5(\partial U(v)) \subset \mathbb{T}_{\epsilon_1}(v)$.

f_4 restricted to the middle surfaces, gluing surfaces and outer boundaries of boundary blocks already has image contained in \widehat{C}_N . For internal blocks, the image is in fact already in C_N itself, and since the extension was done with geodesic arcs f_4 takes each internal block into C_N by convexity. The f_4 -image of a boundary block could leave \widehat{C}_N , but only by a bounded amount because of the Lipschitz bound. In Section 3.4 we show in particular that the geometry of $\partial\widehat{C}_N$ is relatively tame: it consists of parts of ∂C_N^1 and Margulis tube boundaries, meeting together with outward dihedral angles of more than $\pi/2$. It follows that there is a uniformly sized collar neighborhood of the boundary and a uniform Lipschitz retraction of this neighborhood back into ∂C_N^1 . Composing f_4 with this retraction yields a map with image in \widehat{C}_N .

Let v be a vertex with $\ell_\rho(v) < \epsilon_1$. The tube boundary $\partial U(v)$ maps by a uniformly Lipschitz map and is homotopic into $\mathbb{T}_{\epsilon_1}(v)$. It can therefore be pushed into $\mathbb{T}_{\epsilon_1}(v)$ by a uniformly Lipschitz map which comes from the orthogonal projection to a lift of $\mathbb{T}_{\epsilon_1}(v)$ to \mathbb{H}^3 . Using this homotopy in a collar neighborhood of $\partial U(v)$, we can adjust f_4 to get a map, still Lipschitz with a uniform constant, that satisfies (2). This map is f_5 .

Step 6. We extend f_5 to all the nonparabolic tubes in \mathcal{U} . The resulting map f_6 is still homotopic to f_0 . Note that some continuous extension exists because f_5 is homotopic to $f_3|_{M_\nu[0]}$. To get geometric control we extend using a ‘‘coning’’ argument: For each U (which we recall is isometric to a hyperbolic tube) we take a totally geodesic meridian disk \mathcal{D} . Foliating \mathcal{D} by geodesic segments emanating from one boundary point, we extend f_6 to the interior of \mathcal{D} to be totally geodesic on each of these segments. $U \setminus \mathcal{D}$ is now isometric to a convex region in \mathbb{H}^3 which we can again foliate by geodesics emanating from a single point, and repeat.

The extension has these properties:

- (1) Whenever $\ell_\rho(v) < \epsilon_1$ for a vertex v of H , $f_6(U(v))$ is inside $\mathbb{T}_{\epsilon_1}(v)$.
- (2) For any k there exists $L(k)$ such that f_6 is L -Lipschitz on $M_\nu[k]$.
- (3) Given $k > 0$ there exists $\epsilon(k) \in (0, \epsilon_0)$ such that $f_6(M_\nu[k])$ avoids the $\epsilon(k)$ -thin part of N .

For property (1), we apply property (2) of f_5 to see that $f_5(\partial U(v)) \subset \mathbb{T}_{\epsilon_1}(v)$ whenever $\ell_\rho(v) < \epsilon_1$. The extension of f_6 to U using geodesics is therefore also contained in $\mathbb{T}_{\epsilon_1}(v)$.

For property (2), we use the fact that if $|\omega_M(v)| < k$, then the hyperbolic tube $U(v)$ comes from a compact set of possible isometry types. The boundary length of the meridian disk, for example, is bounded by $k\epsilon_0$, and this together with the negative curvature of the target and our method of extension by coning gives us some Lipschitz bound L . Note that property (2) also implies part (5) of the statement of the Lipschitz Model Theorem – that is, that the map on each tube U is Lipschitz with a constant that depends only $|\omega_M(U)|$.

Property (3) follows from property (2) and the following argument: First note (following an argument of Thurston in [59]) that through every point $x \in M_\nu[0]$ there is a pair of loops α, β that generate a non-abelian subgroup of $\pi_1(S)$, and have uniformly bounded lengths. The images of these therefore cannot be contained in a Margulis tube. This, together with the Lipschitz bound on f_5 , gives a uniform $r_0 > 0$ so that $f_5(M_\nu[0])$ cannot penetrate more than r_0 into the ϵ_1 -Margulis tubes of N . Because of the $L(k)$ -Lipschitz property of f_6 on the tubes of $\mathcal{U} \cap M_\nu[k]$, and their bounded geometric type, there is a uniform $r_1(k)$ so that $f_6(M_\nu[k])$ cannot penetrate more than r_1 into the ϵ_0 -Margulis tubes. This gives the desired $\epsilon(k)$, by inequality (3.7).

Step 7. We next adjust f_6 to obtain a map f_7 such that, for a certain k_1 , the image $f_7(M_\nu[k_1])$ avoids the interiors of the “large” tubes $\mathbb{T}[k_1]$ in N (recall from the introduction that $\mathbb{T}[k]$ is the union of ϵ_1 -Margulis tubes in N , if any, corresponding to the homotopy classes of the model tubes $\mathcal{U}[k]$). We will need the following lemma:

LEMMA 10.1. *Given ϵ , there exists $k'(\epsilon)$ such that if $|\omega_M(v)| \geq k'$, then $\ell_\rho(v) \leq \epsilon$.*

Proof. Suppose first that v is an internal vertex of H . In [48] we prove that, if $\ell_\rho(v) > \epsilon$, then

$$d_Y(v_-, v_+) \leq B$$

for each $Y \subset S$ such that $v \in [\partial Y]$, where B depends only on ϵ and S . (This is a combination of Theorem B of [48] with the proof of Theorem A.) By Theorem 9.11, this implies

$$|\omega_H(v)| \leq n,$$

where n depends only on the topology of S and on B . An upper bound on $|\omega_M(v)|$ then follows from Theorem 9.1, and this gives the k' of the lemma.

Suppose v is noninternal. If it is parabolic, then $|\omega_M(v)| = \infty$ and $\ell_\rho(v) = 0$, so the lemma continues to hold. If v is not parabolic, then we have the definition of

$\omega_M(v)$ and $\omega_H(v)$ from Section 9.7, which has an added term $i(r_+(v) + r_-(v))$. If $\ell_\rho(v) > \epsilon$, then we obtain (via Bers' inequality [12]) a lower bound on its length in the conformal boundary of N_ρ , which yields upper bounds on $r_\pm(v)$. Combining this with the previous argument and the restatement (9.28) of Theorem 9.11 for the noninternal case, we obtain an upper bound on $|\omega_M(v)|$. \square

Now let $k_0 = k'(\epsilon_1/2)$, let $\epsilon_4 = \epsilon(k_0)/2$ (where ϵ is the function in property (3) of Step 6), and let $k_1 = k'(\epsilon_4)$. Consider a homotopy class of simple curves v with $\ell_\rho(v) \leq \epsilon_4$ – In particular v can be any vertex of H with $|\omega_M(v)| \geq k_1$, by Lemma 10.1, but v can also be a parabolic homotopy class, or a vertex of $\mathcal{C}(S)$ which *a priori* may not appear in H . In the latter case let $U(v) = \emptyset$.

We claim that $M_v \setminus U(v)$ must all be mapped outside of $\mathbb{T}_{2\epsilon_4}(v)$ by f_6 . For $M_v[k_0]$ this follows from property (3) in Step 6 and the choice of ϵ_4 .

If $w \neq v$ is parabolic or a vertex of H with $|\omega_M(w)| \geq k_0$, by Lemma 10.1 we have $\ell_\rho(w) \leq \epsilon_1/2$, and by Step 6, $U(w)$ is mapped into $\mathbb{T}_{\epsilon_1}(w)$, which is disjoint from $\mathbb{T}_{\epsilon_1}(v)$. This establishes the claim.

Since $\ell_\rho(v) \leq \epsilon_4$, using the “outward” orthogonal projection of the tube $\mathbb{T}_{\epsilon_1}(v)$ minus its geodesic core to its outer boundary $\partial\mathbb{T}_{\epsilon_1}(v)$, we can obtain a map $\varphi_v : \mathbb{T}_{\epsilon_1}(v) \rightarrow \mathbb{T}_{\epsilon_1}(v)$ which is Lipschitz with uniform constant, homotopic to the identity, and equal on the collar $\mathbb{T}_{\epsilon_1}(v) \setminus \mathbb{T}_{2\epsilon_4}(v)$ to the outward orthogonal projection. Let $\Phi : N_\rho \rightarrow N_\rho$ be equal to φ_v on $\mathbb{T}_{\epsilon_1}(v)$ for each v with $\ell_\rho(v) \leq \epsilon_4$ and to the identity elsewhere.

Let $f_7 = \Phi \circ f_6$. This map takes each (nonparabolic) component of $\mathcal{U}[k_1]$ to its corresponding component of $\mathbb{T}[k_1]$, and the complement of all these to the complement of $\mathbb{T}[k_1]$. It is Lipschitz with a uniform constant on $M_v[k_1]$. f_7 is not defined on the parabolic tubes, but their boundaries are mapped (because of the retraction Φ) to the boundaries of the corresponding ϵ_1 -cusps.

Step 8. To obtain f_8 we extend the definition of f_7 to the parabolic tubes of \mathcal{U} . Let U be such a tube and \mathbb{T} its corresponding ϵ_1 -cusp neighborhood. We already know that f_7 maps ∂U to $\partial\mathbb{T}$. Let us identify U with $\partial U \times [0, \infty)$ and \mathbb{T} with $\partial\mathbb{T} \times [0, \infty)$. We can then define $f_8|_U$ with the rule $f_8(x, t) = (f_7(x), t)$. This has the property that, if $f_7|_{\partial U} : \partial U \rightarrow \partial\mathbb{T}$ is a proper map, then so is $f_8|_U : U \rightarrow \mathbb{T}$.

f_8 is our final map f . It takes each component of $\mathcal{U}[k_1]$ to the corresponding Margulis tube $\mathbb{T}_{\epsilon_1}(v)$, and the complement $M_v[k_1]$ to the complement of these tubes. On $M_v[k_1]$ it is still Lipschitz with a uniform constant.

Proper and degree 1. We wish to show that f is proper in both senses: that it takes ∂M_v to $\partial\widehat{C}_N$, and that the inverse images of compact sets are compact. By our construction (Step 2), if ∂M_v is nonempty, then f maps it to $\partial\widehat{C}_N$ by orientation-preserving homeomorphism. Hence the first notion of properness holds.

Let us show that $f|_{M_v[0]}$ is proper. If $\{x_i \in M_v[0]\}$ leaves every compact set, let us show the same for $f(x_i)$. Each internal block is compact, and the noncompact pieces of the finitely many boundary blocks are the added vertical annuli, which by construction are mapped properly. Hence we may assume that each x_i is contained in a different internal block B_i . Writing $B_i = B(e_i)$ for a 4-edge e_i , let $v_i = e_i^+$. This is a sequence of distinct vertices of H . Each v_i has a representative γ_i in B_i of length bounded by a constant, and since f is Lipschitz we obtain a bound on the lengths of $f(\gamma_i)$. This means that $f(\gamma_i)$ must leave every compact set in N . Since the blocks have bounded diameters, the same is true for $f(B_i)$ and hence $f(x_i)$.

Now let U_i be any infinite sequence of distinct, nonparabolic tubes in \mathcal{U} . The generator curves of ∂U_i map to curves of bounded length in N , in distinct homotopy classes, and hence must leave every compact set. For tubes satisfying $|\omega_M(U_i)| < k_1$ there is an upper bound on image diameter, and so their images must leave every compact set as well. The images of tubes with $|\omega_M(U_i)| \geq k_1$ are the corresponding Margulis tubes, which are all distinct and hence again must leave every compact set.

Let U be a parabolic tube. The boundary ∂U is in $M_v[0]$ and hence maps properly to the boundary of the corresponding cusp tube in N . The extension that we constructed in Step 7 is therefore automatically proper.

We conclude that f is proper. The boundary ∂M_v , if nonempty, is mapped by an orientation-preserving homeomorphism to $\partial \hat{C}_N$, and in particular has degree 1. Hence in this case f has degree 1.

In the remaining cases M_v is all of $S \times \mathbb{R}$. Since f is a homotopy-equivalence, to show that it has degree 1 it suffices to show that it preserves the ordering of ends (determined by our orientation convention in both domain and range).

Let B_i be a sequence of blocks of $M_v[0]$ going to infinity in $S \times \mathbb{R}$, in the positive direction (without loss of generality). The associated 4-edges e_i must eventually be subordinate to one of the infinite geodesics of H , which terminates in the forward direction in a lamination component ν_R of ν_+ . Thus the vertices v_i converge to ν_R in $\mathcal{UML}(S)$. By Thurston's theorem on ending laminations (§2.2), the geodesic representatives $f(\gamma_{v_i})^*$ in N must exit the end of N_0 associated with ν_R , which is a “+” end. Since $f(B_i)$ is a bounded distance from either $f(\gamma_{v_i})^*$ or its Margulis tube, the block images exit this end as well. Hence f has degree 1.

This establishes the Lipschitz Model Theorem, where K is the uniform Lipschitz constant obtained by this argument, and $k = k_1$. \square

The extended model map. It is now a simple matter to prove the Extended Model Theorem, as stated in the introduction.

We note that the outer boundary of M_v is naturally identified with the “bottom” boundary of E_v , namely $\partial_\infty N \times \{0\}$, and furthermore that the metric with which

both these boundaries were endowed is the same σ_m . Thus we can glue ∂M_v to the bottom boundary of E_v , obtaining the extended model ME_v . The model map f restricted to ∂M_v is also exactly the same as the map φ . Thus we can combine the maps together to obtain a map $f' : ME_v \rightarrow N$. The desired properties of this map follow immediately from Lemma 3.4 and the Lipschitz Model Theorem.

11. Length bounds

We can now complete the proof of the Short Curve Theorem, stated in the introduction. To prove part (1), let $\bar{\epsilon} = \epsilon_4$. We saw in the proof of the Lipschitz Model Theorem that $f(M_v[k_1])$ avoids the ϵ_4 -thin part, and that each component of $\mathcal{U}[k_1]$ maps into the Margulis tube associated to its homotopy class. Since $f(M_v)$ covers all of \hat{C}_N and all Margulis tubes are contained in \hat{C}_N , it follows that all ϵ_4 -Margulis tubes are in fact covered by tubes of $\mathcal{U}[k_1]$.

Part (2) is exactly the statement of Lemma 10.1. Thus it remains to prove Part (3), which is the content of this lemma:

PROPOSITION 11.1. *There exists a constant c depending only on S , so that*

$$(11.1) \quad |\lambda_\rho(v)| \geq \frac{c}{|\omega_M(v)|}$$

for any vertex v of H . Furthermore the real part $\ell_\rho(v)$ satisfies

$$(11.2) \quad \ell_\rho(\gamma_v) \geq \frac{c}{|\omega_M(v)|^2}.$$

Proof. If v is a parabolic vertex, then $\lambda_\rho(v) = \ell_\rho(v) = 0$ and $|\omega_M(v)| = \infty$, so the inequalities hold in the natural extended sense.

Let v be a nonparabolic vertex of H . If $\ell_\rho(v) \geq \epsilon_4$, then, since $|\omega_M(v)|$ is always at least 1, our inequalities hold for appropriate choice of c .

Assume now that $\ell_\rho(v) < \epsilon_4$. Then as we saw in the proof of the Lipschitz Model Theorem, the model map f maps the complement of $U(v)$ to the complement of $\mathbb{T}_{\epsilon_1}(v)$. Since f has degree 1, it must take $U(v)$ to $\mathbb{T}_{\epsilon_1}(v)$ with degree 1.

It follows that the image of the meridian of $U(v)$ is a meridian of $\mathbb{T}_{\epsilon_1}(v)$. The meridian of $U(v)$ has length $\epsilon_1|\omega_M(v)|$ by definition of ω_M , and the model map is K -lipschitz. On the other hand we know that the meridian of a hyperbolic tube of radius r is at least $2\pi \sinh r$ (see §3.2). Hence

$$2\pi \sinh r(v) \leq K\epsilon_1|\omega_M(v)|.$$

The Brooks-Matelski inequality (3.5) then implies

$$2\pi \sinh \left(\log \frac{1}{|\lambda(\gamma_v)|} - c_1 \right) \leq K\epsilon_1|\omega_M(v)|$$

and the Meyerhoff inequality (3.6) gives

$$2\pi \sinh \left(\frac{1}{2} \log \frac{1}{\ell(\gamma_v)} - c_2 \right) \leq K\epsilon_1 |\omega_M(v)|.$$

A brief calculation yields (11.1) and (11.2). \square

This concludes the proof of the Short Curve Theorem, and gives us *a priori* bounds on our Lipschitz model for the thick part of a surface group. What remains to be done to prove the Ending Lamination Conjecture for surface groups is to promote this Lipschitz map to a bilipschitz map (enabling an application of Sullivan's rigidity theorem). The argument for this involves deeper analysis of the structure of level surfaces in the model manifold, and a local embedding argument that relies heavily on passage to geometric limits. This will be carried out in [18].

References

- [1] W. ABIKOFF, Kleinian groups—geometrically finite and geometrically perverse, in *Geometry of Group Representations* (Boulder, CO, 1987), *Contemp. Math.* **74**, Amer. Math. Soc., Providence, RI, 1988, pp. 1–50. MR 89j:30061 Zbl 0662.30043
- [2] I. AGOL, Tameness of hyperbolic 3-manifolds, preprint, arXiv:math.GT/0405568, 2004.
- [3] L. AHLFORS and L. BERS, Riemann's mapping theorem for variable metrics, *Ann. of Math.* **72** (1960), 385–404. MR 22 #5813 Zbl 0104.29902
- [4] L. V. AHLFORS, Finitely generated Kleinian groups, *Amer. J. Math.* **86** (1964), 413–429. MR 29 #4890 Zbl 0133.04201
- [5] J. M. ALONSO and ET AL., Notes on word hyperbolic groups, in *Group Theory from a Geometrical Viewpoint* (Trieste, 1990), World Sci. Publ., River Edge, NJ, 1991, (H. Short, ed.), pp. 3–63. MR 93g:57001 Zbl 0849.20023
- [6] J. W. ANDERSON and R. D. CANARY, Cores of hyperbolic 3-manifolds and limits of Kleinian groups, *Amer. J. Math.* **118** (1996), 745–779. MR 97k:57015 Zbl 0863.30048
- [7] R. BENEDETTI and C. PETRONIO, *Lectures on Hyperbolic Geometry*, *Universitext*, Springer-Verlag, New York, 1992. MR 94e:57015 Zbl 0768.51018
- [8] L. BERS, Simultaneous uniformization, *Bull. Amer. Math. Soc.* **66** (1960), 94–97. MR 22 #2694 Zbl 0090.05101
- [9] ———, On boundaries of Teichmüller spaces and on Kleinian groups. I, *Ann. of Math.* **91** (1970), 570–600. MR 45 #7044 Zbl 0197.06001
- [10] ———, Spaces of Kleinian groups, in *Several Complex Variables*, (*Proc. Conf., Univ. Maryland, College Park, Md.*, 1970), Springer-Verlag, New York, 1970, pp. *Lecture Notes in Math.* **155**, 9–34. MR 0271333
- [11] ———, Spaces of degenerating Riemann surfaces, in *Discontinuous Groups and Riemann Surfaces* (*Proc. Conf., Univ. Maryland, College Park, Md.*, 1973), *Ann. of Math. Studies* **79**, Princeton Univ. Press, Princeton, NJ, 1974, pp. 43–55. MR 50 #13497 Zbl 0294.32016
- [12] ———, An inequality for Riemann surfaces, in *Differential Geometry and Complex Analysis*, Springer-Verlag, New York, 1985, pp. 87–93. MR 0780038 Zbl 0575.30039
- [13] F. BONAHOON, Bouts des variétés hyperboliques de dimension 3, *Ann. of Math.* **124** (1986), 71–158. MR 88c:57013 Zbl 0671.57008

- [14] ———, Geodesic laminations on surfaces, in *Laminations and Foliations in Dynamics, Geometry and Topology* (Stony Brook, NY, 1998), *Contemp. Math.* **269**, Amer. Math. Soc., Providence, RI, 2001, pp. 1–37. MR 2001m:57023 Zbl 0996.53029
- [15] B. H. BOWDITCH, Notes on Gromov’s hyperbolicity criterion for path-metric spaces, in *Group Theory from a Geometrical Viewpoint (Trieste, 1990)*, World Sci. Publ., River Edge, NJ, 1991, pp. 64–167. MR 93h:57002 Zbl 0843.20031
- [16] ———, End invariants of hyperbolic 3-manifolds, preprint, 2005.
- [17] J. F. BROCK, Continuity of Thurston’s length function, *Geom. Funct. Anal.* **10** (2000), 741–797. MR 2001g:57028 Zbl 0968.57011
- [18] R. BROOKS, R. D. CANARY, and Y. N. MINSKY, The classification of Kleinian surface groups II: the ending lamination conjecture, preprint, arXiv:math.GT/0412006.
- [19] R. BROOKS and J. P. MATELSKI, Collars in Kleinian groups, *Duke Math. J.* **49** (1982), 163–182. MR 83f:30039 Zbl 0484.30029
- [20] P. BUSER, *Geometry and Spectra of Compact Riemann Surfaces*, *Progr. Math.* **106**, Birkhäuser, Boston, MA, 1992. MR 93g:58149 Zbl 0770.53001
- [21] D. CALEGARI and D. GABAI, Shrinkwrapping and the taming of hyperbolic 3-manifolds, *J. Amer. Math. Soc.* **19** (2006), 385–446. MR 2006g:57030 Zbl 1090.57010
- [22] R. D. CANARY, D. B. A. EPSTEIN, and P. GREEN, Notes on notes of Thurston, in *Analytical and Geometric Aspects of Hyperbolic Space (Coventry/Durham, 1984)*, *London Math. Soc. Lecture Note Ser.* **111**, Cambridge Univ. Press, Cambridge, 1987, pp. 3–92. MR 89e:57008 Zbl 0612.57009
- [23] R. D. CANARY, Ends of hyperbolic 3-manifolds, *J. Amer. Math. Soc.* **6** (1993), 1–35. MR 93e:57019 Zbl 0810.57006
- [24] J. W. CANNON, The theory of negatively curved spaces and groups, in *Ergodic Theory, Symbolic Dynamics, and Hyperbolic Spaces (Trieste, 1989)*, *Oxford Sci. Publ.*, Oxford Univ. Press, New York, 1991, pp. 315–369. MR 1130181 Zbl 0764.57002
- [25] A. J. CASSON and S. A. BLEILER, *Automorphisms of Surfaces after Nielsen and Thurston*, *London Math. Soc. Stud. Texts* **9**, Cambridge Univ. Press, Cambridge, 1988. MR 89k:57025 Zbl 0649.57008
- [26] D. B. A. EPSTEIN and A. MARDEN, Convex hulls in hyperbolic space, a theorem of Sullivan, and measured pleated surfaces, in *Analytical and Geometric Aspects of Hyperbolic Space (Coventry/Durham, 1984)*, *London Math. Soc. Lecture Note Ser.* **111**, Cambridge Univ. Press, Cambridge, 1987, pp. 113–253. MR 89c:52014 Zbl 0612.57010
- [27] E. GHYS and P. DE LA HARPE (eds.), *Sur les Groupes Hyperboliques d’après Mikhael Gromov*, *Progr. Math.* **83**, Birkhäuser, Boston, MA, 1990, (Papers from the Swiss Seminar on Hyperbolic Groups held in Bern, 1988). MR 92f:53050 Zbl 0731.20025
- [28] G. GROMOV, *Essays in Group Theory*, *Math. Sci. Res. Inst. Publ.* **8**, Springer-Verlag, New York, 1987, (S. M. Gersten, ed.). MR 88e:20004 Zbl 0547.00016
- [29] W. J. HARVEY, Boundary structure of the modular group, in *Riemann Surfaces and Related Topics: Proc. 1978 Stony Brook Conference (State Univ. New York, Stony Brook, NY, 1978)*, *Ann. of Math. Stud.* **97**, Princeton Univ. Press, Princeton, N.J., 1981, pp. 245–251. MR 83d:32022 Zbl 0461.30036
- [30] A. HATCHER and W. THURSTON, A presentation for the mapping class group of a closed orientable surface, *Topology* **19** (1980), 221–237. MR 81k:57008 Zbl 0447.57005

- [31] S. HERSONSKY, A generalization of the Shimizu-Leutbecher and Jørgensen inequalities to Möbius transformations in \mathbf{R}^N , *Proc. Amer. Math. Soc.* **121** (1994), 209–215. MR 94m:30085 Zbl 0812.30017
- [32] T. JØRGENSEN, On discrete groups of Möbius transformations, *Amer. J. Math.* **98** (1976), 739–749. MR 55 #658 Zbl 0336.30007
- [33] D. A. KAZHDAN and G. A. MARGULIS, A proof of Selberg’s conjecture, *Math. USSR Sb.* **4** (1968), 147–152. MR 0223487 Zbl 0241.22024
- [34] E. KLARREICH, The boundary at infinity of the curve complex and the relative Teichmüller space, preprint, <http://www.nasw.org/users/klarreich/research.htm>.
- [35] Y. KOMORI, V. MARKOVIC, and C. SERIES (eds.), *Kleinian Groups and Hyperbolic 3-Manifolds*, *London Math. Soc. Lecture Note Ser.* **299**, Cambridge Univ. Press, Cambridge, 2003. MR 2004j:57001 Zbl 1031.30002
- [36] I. KRA, On spaces of Kleinian groups, *Comment. Math. Helv.* **47** (1972), 53–69. MR 46 #5611 Zbl 0239.30020
- [37] R. S. KULKARNI and P. B. SHALEN, On Ahlfors’ finiteness theorem, *Adv. Math.* **76** (1989), 155–169. MR 91e:57067 Zbl 0684.57019
- [38] A. MARDEN and B. MASKIT, On the isomorphism theorem for Kleinian groups, *Invent. Math.* **51** (1979), 9–14. MR 80f:30037 Zbl 0399.30037
- [39] A. MARDEN, The geometry of finitely generated kleinian groups, *Ann. of Math.* **99** (1974), 383–462. MR 50 #2485 Zbl 0282.30014
- [40] B. MASKIT, Self-maps on Kleinian groups, *Amer. J. Math.* **93** (1971), 840–856. MR 45 #544 Zbl 0227.32007
- [41] ———, *Kleinian Groups*, *Grundle Math. Wissen.* **287**, Springer-Verlag, New York, 1988. MR 90a:30132 Zbl 0627.30039
- [42] H. A. MASUR and Y. N. MINSKY, Geometry of the complex of curves. II. Hierarchical structure, *Geom. Funct. Anal.* **10** (2000), 902–974. MR 2001k:57020 Zbl 0972.32011
- [43] ———, Geometry of the complex of curves. I. Hyperbolicity, *Invent. Math.* **138** (1999), 103–149. MR 2000i:57027 Zbl 0941.32012
- [44] D. MCCULLOUGH, Compact submanifolds of 3-manifolds with boundary, *Quart. J. Math. Oxford Ser.* **37** (1986), 299–307. MR 88d:57012 Zbl 0628.57008
- [45] R. MEYERHOFF, A lower bound for the volume of hyperbolic 3-manifolds, *Canad. J. Math.* **39** (1987), 1038–1056. MR 88k:57049 Zbl 0694.57005
- [46] Y. N. MINSKY, A geometric approach to the complex of curves on a surface, in *Topology and Teichmüller Spaces (Katinkulta, 1995)*, World Sci. Publ., River Edge, NJ, 1996, pp. 149–158. MR 2000g:32016 Zbl 0937.30027
- [47] ———, The classification of punctured-torus groups, *Ann. of Math.* **149** (1999), 559–626. MR 2000f:30028 Zbl 0939.30034
- [48] ———, Kleinian groups and the complex of curves, *Geom. Topol.* **4** (2000), 117–148. MR 2006k:57051 Zbl 0953.30027
- [49] ———, Bounded geometry for Kleinian groups, *Invent. Math.* **146** (2001), 143–192. MR 2004f:30032 Zbl 1061.37026
- [50] R. L. MOORE, Concerning upper semi-continuous collections of continua, *Trans. Amer. Math. Soc.* **27** (1925), 416–428. MR 1501320
- [51] ———, On the foundations of plane analysis situs, *Trans. Amer. Math. Soc.* **17** (1916), 131–164. MR 1501033

- [52] G. D. MOSTOW, Quasi-conformal mappings in n -space and the rigidity of hyperbolic space forms, *Inst. Hautes Études Sci. Publ. Math.* (1968), 53–104. MR 38 #4679 Zbl 0189.09402
- [53] R. C. PENNER and J. L. HARER, *Combinatorics of Train Tracks*, *Annals of Math. Studies* **125**, Princeton Univ. Press, Princeton, NJ, 1992. MR 94b:57018 Zbl 0765.57001
- [54] G. PRASAD, Strong rigidity of \mathbf{Q} -rank 1 lattices, *Invent. Math.* **21** (1973), 255–286. MR 52 #5875 Zbl 0264.22009
- [55] M. REES, The ending laminations theorem direct from Teichmüller geodesics, 2004. arXiv math.GT/0404007
- [56] G. P. SCOTT, Compact submanifolds of 3-manifolds, *J. London Math. Soc.* **7** (1973), 246–250. MR 48 #5080 Zbl 0266.57001
- [57] T. SOMA, Geometric approach to Ending Lamination Conjecture, arXiv:0801.4236, 2008.
- [58] D. SULLIVAN, On the ergodic theory at infinity of an arbitrary discrete group of hyperbolic motions, in *Riemann Surfaces and Related Topics: Proc. 1978 Stony Brook Conference (State Univ. New York, Stony Brook, NY, 1978)*, *Ann. of Math. Studies* **97**, Princeton Univ. Press, Princeton, NJ, 1981, pp. 465–496. MR 83f:58052 Zbl 0567.58015
- [59] W. P. THURSTON, Hyperbolic structures on 3-manifolds, II: surface groups and manifolds which fiber over the circle, preprint, arXiv:math.GT/9801045.
- [60] ———, The geometry and topology of 3-manifolds, online at <http://www.msri.org/publications/books/gt3m>, 1982.
- [61] ———, Three-dimensional manifolds, Kleinian groups and hyperbolic geometry, *Bull. Amer. Math. Soc.* **6** (1982), 357–381. MR 83h:57019 Zbl 0496.57005
- [62] ———, Hyperbolic structures on 3-manifolds. I. Deformation of acylindrical manifolds, *Ann. of Math.* **124** (1986), 203–246. MR 88g:57014 Zbl 0668.57015
- [63] ———, *Three-Dimensional Geometry and Topology. Vol. 1*, *Princeton Math. Series* **35**, Princeton Univ. Press, Princeton, NJ, 1997 (Silvio Levy, ed.). MR 97m:57016 Zbl 0873.57001
- [64] P. L. WATERMAN, Möbius transformations in several dimensions, *Adv. Math.* **101** (1993), 87–113. MR 95h:30056 Zbl 0793.15019

(Received April 10, 2003)

E-mail address: yair.minsky@yale.edu

DEPARTMENT OF MATHEMATICS, YALE UNIVERSITY, NEW HAVEN CT 06520-8283,
UNITED STATES

<http://www.math.yale.edu/users/yair>

