

THE CLASSIFICATION OF KLEINIAN SURFACE GROUPS, II: THE ENDING LAMINATION CONJECTURE

JEFFREY F. BROCK, RICHARD D. CANARY, AND YAIR N. MINSKY

ABSTRACT. Thurston's Ending Lamination Conjecture states that a hyperbolic 3-manifold with finitely generated fundamental group is uniquely determined by its topological type and its end invariants. In this paper we prove this conjecture for Kleinian surface groups. The main ingredient is the establishment of a uniformly bilipschitz model for a Kleinian surface group. The first half of the proof appeared in [47], and a subsequent paper [15] will establish the Ending Lamination Conjecture in general.

CONTENTS

1. The ending lamination conjecture	1
2. Background and statements	8
3. Knotting and partial order of subsurfaces	25
4. Cut systems and partial orders	49
5. Regions and addresses	66
6. Uniform embeddings of Lipschitz surfaces	76
7. Insulating regions	91
8. Proof of the bilipschitz model theorem	99
9. Proofs of the main theorems	118
10. Corollaries	119
References	123

1. The ending lamination conjecture

In the late 1970's Thurston formulated a conjectural classification scheme for all hyperbolic 3-manifolds with finitely generated fundamental group. The picture proposed by Thurston generalized what had hitherto been understood, through the work of Ahlfors [3], Bers [10], Kra [34], Marden [37], Maskit [38], Mostow [51], Prasad [55], Thurston [67] and others, about geometrically finite hyperbolic 3-manifolds.

Thurston's scheme proposes *end invariants* which encode the asymptotic geometry of the ends of the manifold, and which generalize the Riemann

Date: December 7, 2004.

Partially supported by NSF grants DMS-0354288, DMS-0203698 and DMS-0203976.

surfaces at infinity that arise in the geometrically finite case. Thurston made this conjecture in [67]:

Ending Lamination Conjecture. *A hyperbolic 3-manifold with finitely generated fundamental group is uniquely determined by its topological type and its end invariants.*

This paper is the second in a series of three which will establish the Ending Lamination Conjecture for all topologically tame hyperbolic 3-manifolds. (For expository material on this conjecture, and on the proofs in this paper and in [47], we direct the reader to [43],[48] and [49]).

Together with the recent proofs of Marden’s Tameness Conjecture by Agol [2] and Calegari-Gabai [17], this gives a complete classification of all hyperbolic 3-manifolds with finitely-generated fundamental group.

In this paper we will discuss the *surface group* case. A Kleinian surface group is a discrete, faithful representation $\rho : \pi_1(S) \rightarrow \mathrm{PSL}_2(\mathbb{C})$ where S is a compact orientable surface, such that the restriction of ρ to any boundary loop has parabolic image. These groups arise naturally as restrictions of more general Kleinian groups to surface subgroups. Bonahon [11] and Thurston [66] showed that the associated hyperbolic 3-manifold $N_\rho = \mathbb{H}^3/\rho(\pi_1(S))$ is homeomorphic to $\mathrm{int}(S) \times \mathbb{R}$ and that ρ has a well-defined pair of end invariants (ν_+, ν_-) . Typically, each end invariant is either a point in the Teichmüller space of S or a geodesic lamination on S . In the general situation, each end invariant is a geodesic lamination on some (possibly empty) subsurface of S and a conformal structure on the complementary surface. We will prove:

Ending Lamination Theorem for Surface Groups *A Kleinian surface group ρ is uniquely determined, up to conjugacy in $\mathrm{PSL}_2(\mathbb{C})$, by its end invariants.*

The first part of the proof of this theorem appeared in [47], and we will refer to that paper for some of the background and notation, although we will strive to make this paper readable fairly independently.

See Section 1.3 for a discussion of the proof of the general Ending Lamination Conjecture, which will appear in [15].

Bilipschitz Model Theorem. The main technical result which leads to the Ending Lamination Theorem is the Bilipschitz Model Theorem, which gives a bilipschitz homeomorphism from a “model manifold” M_ν to the hyperbolic manifold N_ρ (See §2.7 for a precise statement). The model M_ν was constructed in Minsky [47], and its crucial property is that it depends only on the end invariants $\nu = (\nu_+, \nu_-)$, and not on ρ itself. (Actually M_ν is mapped to the “augmented convex core” of N_ρ , but as this is the same as N_ρ in the main case of interest, we will ignore the distinction for the rest of the introduction. See §2.7 for details.)

The proof of the Bilipschitz Model Theorem will be completed in Section 8, and the Ending Lamination Conjecture will be obtained as a consequence of this and Sullivan's rigidity theorem in Section 9. In that section we will also prove the Length Bound Theorem, which gives estimates on the lengths of short geodesics in N_ρ (see §2.7 for the statement).

1.1. Corollaries

A positive answer to the Ending Lamination Conjecture allows one to settle a number of fundamental questions about the structure of Kleinian groups and their deformation spaces. In the sequel, we will see that it gives (together with convergence theorems of Thurston, Ohshika, Kleineidam-Souto and Lecuire) a complete proof of the Bers-Sullivan-Thurston density conjecture, which predicts that every finitely generated Kleinian group is an algebraic limit of geometrically finite groups. In the surface group case, the density conjecture follows immediately from our main theorem and results of Thurston [65] and Ohshika [52]. We recall that $AH(S)$ is the space of conjugacy classes of Kleinian surface groups and that a surface group is quasifuchsian if both its ends are geometrically finite and it has no additional parabolic elements.

Density Theorem. *The set of quasifuchsian surface groups is dense in $AH(S)$.*

Marden [37] and Sullivan [63] showed that the interior of $AH(S)$ consists exactly of the quasifuchsian groups. Bromberg [16] and Brock-Bromberg [14] previously showed that freely indecomposable Kleinian groups without parabolic elements are algebraic limits of geometrically finite groups, using cone-manifold techniques and the bounded-geometry version of the Ending Lamination Conjecture in Minsky [46].

We also obtain a quasiconformal rigidity theorem that gives the best possible common generalization of Mostow [51] and Sullivan's [62] rigidity theorems.

Rigidity Theorem. *If two Kleinian surface groups are conjugate by an orientation-preserving homeomorphism of $\widehat{\mathbb{C}}$, then they are quasiconformally conjugate.*

We also establish McMullen's conjecture that the volume of the thick part of the convex core of a hyperbolic 3-manifold grows polynomially.

If x lies in the thick part of the convex core C_N , then let $B_R^{thick}(x)$ be the set of points in the ϵ_1 -thick part of C_N which can be joined to x by a path of length at most R lying entirely in the ϵ_1 -thick part.

Volume Growth Theorem. *If $\rho : \pi_1(S) \rightarrow PSL_2(\mathbb{C})$ is a Kleinian surface group and $N = \mathbb{H}^3/\rho(\pi_1(S))$, then for any x in the ϵ_1 -thick part of C_N ,*

$$\text{volume} \left(B_R^{thick}(x) \right) \leq cR^{d(S)},$$

where c depends only on the topological type of S , and

$$d(S) = \begin{cases} -\chi(S) & \text{genus}(S) > 0 \\ -\chi(S) - 1 & \text{genus}(S) = 0. \end{cases}$$

A different proof of the Volume Growth Theorem is given by Bowditch in [13]. We are grateful to Bowditch for pointing out an error in our original definition of $d(S)$.

Proofs of these corollaries are given in Section 10. Each of them admit generalizations to the setting of all finitely generated Kleinian groups and these generalizations will be discussed in [15]. In that paper, we will also use the solution of the Ending Lamination Conjecture and work of Anderson-Canary-McCullough [7] to obtain a complete enumeration of the components of the deformation space of a freely indecomposable, finitely generated Kleinian group.

Another corollary, which we postpone to a separate article, is a theorem which describes the topology and geometry of geometric limits of sequences of Kleinian surface groups – in particular such a limit is always homeomorphic to a subset of $S \times \mathbb{R}$. This was also shown by Soma [58].

1.2. Outline of the proof

The Lipschitz Model Theorem, from [47], provides a degree 1 homotopy equivalence from the model manifold M_ν to the hyperbolic manifold N_ρ (or in general to the augmented core of N_ρ , but we ignore the distinction in this outline), which respects the thick-thin decompositions of M_ν and N_ρ and is Lipschitz on the thick part of M_ν . (See §2.7.)

Our main task in this paper is to promote this map to a *bilipschitz* homeomorphism between M_ν and N_ρ , and this is the content of our main result, the Bilipschitz Model Theorem. The proof of the Bilipschitz Model Theorem converts the Lipschitz model map to a bilipschitz map incrementally on various subsets of the model. The main ideas of the proof can be summarized as follows:

Topological order of subsurfaces. In section 3 we discuss a “topological order relation” among embedded surfaces in a product 3-manifold $S \times \mathbb{R}$. This is the intuitive notion that one surface may lie “below” another in this product, but this relation does not in fact induce a partial order and hence a number of sticky technical issues arise.

We introduce an object called a “scaffold”, which is a subset of $S \times \mathbb{R}$ consisting of a union of unknotted solid tori and surfaces in $S \times \mathbb{R}$, each isotopic to a level subsurface, satisfying certain conditions. The main theorem in this section is the Scaffold Extension Theorem (3.8), which states that, under appropriate conditions (in particular an “order-preservation” condition), embeddings of a scaffold into $S \times \mathbb{R}$ can be extended to global homeomorphisms of $S \times \mathbb{R}$.

Much of the rest of the proof is concerned with analyzing this order in the model manifold, breaking the model up into pieces separated by scaffolds, and insuring that the model map satisfies the appropriate order-preserving condition.

Structure of the model: tubes, surfaces and regions. The structure of the model M_ν can be analyzed using the structure of “hierarchies in the complex of curves”, as developed in [41], applied in [47] and summarized here in §2.2. In particular, M_ν is homeomorphic to $S \times \mathbb{R}$, and contains a collection of unknotted solid tori which correspond to the Margulis tubes for short geodesics and cusps in N_ρ , and which the Lipschitz Model Theorem states are in fact mapped properly to those tubes. The model also contains a large family of “split-level surfaces”, which are surfaces isotopic to level surfaces in $S \times \mathbb{R}$ and bilipschitz-homeomorphic to bounded-geometry hyperbolic surfaces.

In Section 4 we discuss “cut systems” in the hierarchy that generates the model. A “cut system” gives rise to a family of split-level surfaces and we show, in Lemma 4.15, that one can impose spacing conditions on these cut systems so that the topological order relation restricted to the split-level surfaces coming from the cut system generates a partial order.

In Section 5 we will show how the surfaces of a cut system (together with the model tubes) cut the model into regions whose geometry is controlled. The collection of split-level surfaces and Margulis tubes which bound such a region form a scaffold.

Uniform embeddings of model surfaces. The restriction of the model map to a split-level surface is essentially a Lipschitz map of a bounded-geometry hyperbolic surface whose boundaries map to Margulis tubes (we call this an *anchored surface*). These surfaces are not themselves embedded, but we will show that they may be deformed in a controlled way to bilipschitz embeddings.

In general, a Lipschitz anchored surface may be “wrapped” around a deep Margulis tube in N_ρ and any homotopy to an embedding must pass through the core of this tube. In Theorem 6.1, we show that this wrapping phenomenon is the only obstruction to a controlled homotopy. The proof relies on a geometric limiting argument and techniques of Anderson-Canary [5]. In section 8.1, we check that we may choose the spacing constants for our cut system, so that the associated Lipschitz anchored surfaces are not wrapped, and hence can be uniformly embedded.

Preservation of topological order. In section 8.2 we show that any cut system may be “thinned out” in a controlled way to yield a new cut system, with uniform spacing constants, so that if two split level surfaces lie on the boundary of the same complementary region in M_ν , then their associated anchored embeddings in N_ρ (from §8.2) are disjoint. We adjust the model

map so that it is a bilipschitz embedding on collar neighborhoods of these slices.

In 8.3, we check that if two anchored surfaces are disjoint and ordered in the hyperbolic manifold, then their relative ordering agrees with the ordering of the associated split-level surfaces in the model. The idea is to locate “insulating regions” in the model which separate the two surfaces, and on which there is sufficient control to show that the topological order between the insulating region and each of the two surfaces is preserved. A certain transitivity argument can then be used to show that the order between the surfaces is preserved as well.

The insulating regions are of two types. Sometimes a model tube is available and it is fairly immediate from properties of the model map that its image Margulis tube has the correct separation properties. When such a tube is not available we show, in Theorem 7.1, that there exist certain “subsurface product regions”, which look roughly like pieces of bounded-geometry surface group manifolds based on lower-complexity surfaces. The control over these regions is obtained by a geometric limit argument.

Bilipschitz extension to the regions. The union of the split-level surfaces and the solid tori divide the model manifold up into regions bounded by scaffolds. The Scaffold Extension Theorem can be used to show that the embeddings on the split-level surfaces can be extended to embeddings of these complementary regions. An additional geometric limit argument, given in section 8.4, is needed to obtain bilipschitz bounds on each of these embeddings. Piecing together the embeddings, we obtain a bilipschitz embedding of the “thick part” of the model to the thick part of N_ρ . A final brief argument, given in section 8.5 shows that the map can be extended also to the model tubes in a uniform way. This completes the proof.

This outline ignores the case when the convex hull of N_ρ has nonempty boundary, and in fact most of the proof on a first reading is improved by ignoring this case. Dealing with boundary is mostly an issue of notation and some attention to special cases; but nothing essentially new happens. In Section 8 most details of the case with boundary are postponed to §8.6.

1.3. The general case of the Ending Lamination Conjecture

In this section we briefly discuss the proof of the Ending Lamination Conjecture in the general situation. Details will appear in [15].

Let N be a hyperbolic 3-manifold with finitely generated fundamental group, and let N^0 be the complement of the cusps of N . N^0 has a *relative compact core* K (see Kulkarni-Shalen [35] or McCullough [42]) which is a compact submanifold whose inclusion into N^0 is a homotopy equivalence, and which meets the boundary of each cusp in either a π_1 -injective annulus or a torus. Let $P \subset \partial K$ denote the union of these annuli and tori. The ends of N^0 are in one-to-one correspondence with the components of $\partial K \setminus P$ (see Proposition 1.3 in Bonahon [11].)

Incompressible boundary case. In the case where each component of $\partial K \setminus P$ is incompressible, the derivation of the Ending Lamination Conjecture from the surface group case is fairly straightforward. In this case the restriction of $\pi_1(N)$ to the fundamental group of any component of R of $\partial K \setminus P$ is a Kleinian surface group. Let N_R be the cover of N associated to $\pi_1(R)$. The Bilipschitz Model Theorem applies to give a model for N_R , and since a neighborhood of one of the ends of N_R projects homeomorphically to N we obtain bilipschitz models for each of the ends of N .

Two homeomorphic hyperbolic manifolds with the same end invariants must therefore have a bilipschitz correspondence between neighborhoods of their ends (the end invariant data specify the cusps so that after removing the cusps the manifolds remain homeomorphic.) Since what remains is compact, one may easily extend the bilipschitz homeomorphism on the neighborhoods of the ends to a global bilipschitz homeomorphism. One again applies Sullivan's rigidity theorem [62] to complete the proof.

Compressible boundary case. When some component R of $\partial K \setminus P$ is compressible, the restriction of ρ to $\pi_1(R)$ is no longer a Kleinian surface group. Quite recently, Agol [2] and Calegari-Gabai [17] proved that N is homeomorphic to the interior of K . Canary [19] showed that the ending invariants are well-defined in this setting.

The first step of the proof in this case is to apply Canary's branched-cover trick from [19]. That is, we find a suitable closed geodesic γ in N and a double branched cover $\pi : \hat{N} \rightarrow N$ over γ , such that $\pi_1(\hat{N})$ is freely indecomposable. The singularities on the branching locus can be smoothed locally to give a pinched negative curvature (PNC) metric on \hat{N} . Since N is topologically tame, one may choose a relative compact core K for N^0 containing γ , so that $\hat{K} = \pi^{-1}(K)$ is a relative compact core for \hat{N}^0 . Let $P = \partial N^0 \cap K$ and $\hat{P} = \partial \hat{N}^0 \cap \hat{K}$. If R is any component of $\partial K - P$, then $\pi^{-1}(R)$ consists of two homeomorphic copies \hat{R}_1 and \hat{R}_2 of R , each of which is incompressible.

Given a component R of $\partial K - P$, we consider the cover \hat{N}_R of \hat{N} associated to $\pi_1(\hat{R}_1)$. We then apply the techniques of [47] and this paper to obtain a bilipschitz model for some neighborhood of the end E_R of \hat{N}_R^0 "facing" \hat{R}_1 . In particular, we need to check that the estimates of [47] apply in suitable neighborhoods of E_R . The key tool we will need is a generalization of Thurston's Uniform Injectivity Theorem [68] for pleated surfaces to this setting. (See Miyachi-Ohshika [50] for a discussion of this line of argument in the "bounded geometry" case.) Once we obtain a bilipschitz model for some neighborhood of E_R it projects down to give a bilipschitz model for a neighborhood of the end of N^0 "facing" R . As before, we obtain a bilipschitz model for the complement of a compact submanifold of N^0 and the proof proceeds as in the incompressible boundary setting.

It is worth noting that this construction does not yield a *uniform* model for N , in the sense that the bilipschitz constants depend on the geometry of

N and not only on its topological type (for example on the details of what happens in the branched covering step). The model we develop here for the surface group case is uniform, and we expect that in the incompressible-boundary case uniformity of the model should not be too hard to obtain. Uniformity in general is quite an interesting problem, and would be useful for further applications of the model manifold.

2. Background and statements

In this section we will introduce and discuss notation and background results, and then in §2.7 we will state the main technical results of this paper, the Bilipschitz Model Theorem and the Length Bound Theorem.

2.1. Surfaces, notation and conventions

We denote by $S_{g,n}$ a compact oriented surface of genus g and n boundary components, and define a complexity $\xi(S_{g,n}) = 3g + n$. A subsurface $Y \subset X$ is essential if its boundary components do not bound disks in X and Y is not homotopic into ∂X . All subsurfaces which occur in this paper are essential. Note that $\xi(Y) < \xi(X)$ unless Y is isotopic to X .

As in [47], it will be convenient to fix standard representatives of each isotopy class of subsurfaces in a fixed surface S . This can be done by fixing a complete hyperbolic metric σ_0 of finite area on $\text{int}(S)$. Then if v is a homotopy class of simple, homotopically nontrivial curves, let γ_v denote the σ_0 -geodesic representative of v , provided v does not represent a loop around a cusp. In [47, Lemma 3.3] we fix a version of the standard collar construction to obtain an open annulus $\mathbf{collar}(v)$ (or $\mathbf{collar}(\gamma_v)$) which is tubular neighborhood of γ_v or a horospherical neighborhood in the cusp case. This collar has the additional property that the closures of two such collars are disjoint whenever the core curves have disjoint representatives. If Γ is a collection of simple homotopically distinct and nontrivial disjoint curves we let $\mathbf{collar}(\Gamma)$ be the union of collars of components.

Now let \widehat{S} denote a separate copy of $\text{int}(S)$ with the metric σ_0 . Embed S in \widehat{S} as the complement of $\mathbf{collar}(\partial S)$. Similarly for any essential subsurface $X \subset S$, our standard representative will be the component of $\widehat{S} \setminus \mathbf{collar}(\partial X)$ isotopic to X if X is not an annulus, and $\overline{\mathbf{collar}(\gamma)}$ if X is an annulus with core curve γ . We will from now on assume that any subsurface of S is of this form. Note that two such subsurfaces intersect if and only if their intersection is homotopically essential. We will use the term “overlap” to indicate homotopically essential intersection (see also §3 for the use of this term in three dimensions).

We will denote by $\mathcal{D}(S)$ the set of discrete, faithful representations $\rho : \pi_1(S) \rightarrow \text{PSL}_2(\mathbb{C})$ such that any loop representing a boundary component is taken to a parabolic element – that is, the set of Kleinian surface groups for the surface S . If $\rho \in \mathcal{D}(S)$ we denote by N_ρ its quotient manifold $\mathbb{H}^3 / \rho(\pi_1(S))$.

2.2. Hierarchies and partial orders

We refer to Minsky [47] for the basic definitions of hierarchies of geodesics in the complex of curves of a surface. These notions were first developed in Masur-Minsky [41]. We will recall the needed terminology and results here.

Complexes, subsurfaces and projections. We denote by $\mathcal{C}(X)$ the complex of curves of a surface X (originally due to Harvey [25, 26]) whose k -simplices are $(k + 1)$ -tuples of nontrivial nonperipheral homotopy classes of simple closed curves with disjoint representatives. For $\xi(X) = 4$, we alter the definition slightly, so that $[vw]$ is an edge whenever v and w have representatives that intersect once (if $X = S_{1,1}$) or twice (if $X = S_{0,4}$).

When X has boundary we define the “curve and arc complex” $\mathcal{A}(X)$ similarly, where vertices are proper nontrivial homotopy classes of properly embedded simple arcs *or* closed curves. When X is an annulus the homotopies are assumed to fix the endpoints.

If $X \subset S$ we have a natural map $\pi_X : \mathcal{A}(S) \rightarrow \mathcal{A}(X)$ defined using the essential intersections with X of curves in S . When X is an annulus π_X is defined using the lift to the annular cover of S associated to S . If v is a vertex of $\mathcal{C}(S)$, we let $\mathcal{A}(v)$ denote the complex $\mathcal{A}(\text{collar}(v))$.

We recall from Masur-Minsky [40] that $\mathcal{C}(X)$ is δ -hyperbolic (see also Bowditch [12] for a new proof), and from Klarreich [33] (see also Hamenstadt [24] for an alternate proof) that its Gromov boundary $\partial\mathcal{C}(X)$ can be identified with the set $\mathcal{EL}(X)$ of minimal filling geodesic laminations on X , with the topology inherited from Thurston’s space of measured laminations under the measure-forgetting map.

Markings. A (generalized) marking μ in S is a geodesic lamination $\text{base}(\mu)$ in $\mathcal{GL}(S)$, together with a (possibly) empty list of “transversals”. A transversal is a vertex of $\mathcal{A}(v)$ where v is a vertex of $\text{base}(\mu)$ (i.e. a simple closed curve component of the lamination). We also have a notion of *clean marking*, where the base is a simplex in $\mathcal{C}(S)$ and the transversals are more concrete: for each vertex v of the base, the transversal is a simple closed curve intersecting v either once or twice, and disjoint from the other base curves. (After projection to $\mathcal{A}(v)$ we obtain a transversal in the previous sense. See [47] for details).

A marking is called *maximal* if its base is maximal as a lamination, and if every closed curve component of the base has a nonempty transversal.

Tight geodesics and subordinacy. A *tight sequence* in (a non-annular) surface X is a (finite or infinite) sequence (w_i) of simplices in the complex of curves $\mathcal{C}(X)$, with the property that for any vertex $v_i \in w_i$ and $v_j \in w_j$ with $i \neq j$ we have $d_{\mathcal{C}(X)}(v_i, v_j) = |i - j|$, and the additional property that w_i is the boundary of the subsurface filled by $w_{i-1} \cup w_{i+1}$ (this is “tightness,” see Definition 5.1 of [47]).

If $X \subset S$ is a nonannular subsurface then a *tight geodesic* g in X is a tight sequence $\{v_i\}$ in $\mathcal{C}(X)$, together with two generalized markings $\mathbf{I}(g)$ and

$\mathbf{T}(g)$ such that the following holds: If the sequence $\{v_i\}$ has a first element, v_0 , we require that v_0 is a vertex of $\text{base}(\mathbf{I}(g))$; otherwise by Klarreich's theorem [33] v_i converge as $i \rightarrow -\infty$ to a unique lamination in $\mathcal{EL}(S)$. We choose $\text{base}(\mathbf{I}(g))$ to be this lamination (and $\mathbf{I}(g)$ has no transversals. A similar condition holds for $\mathbf{T}(g)$ and the forward direction of $\{v_i\}$. We call X the *domain of g* and write $X = D(g)$.

When X is an annulus a tight sequence is any finite geodesic sequence such that the endpoints on ∂X of all the vertices are contained in the set of endpoints of the first and last vertex. For a tight geodesic we define \mathbf{I} and \mathbf{T} to be simply the first and last vertices.

For convenience we define $\xi(g)$ to be $\xi(D(g))$ for a tight geodesic g .

Given a tight geodesic g and a subsurface $Y \subset D(g)$, we write $Y \searrow^d g$ if there is a simplex v of g such that Y is a complementary component of $D(g) \setminus \text{collar}(v)$ or a component of $\text{collar}(v)$ (for short we say that Y is a *component domain of $(D(g), v)$*), and the successor of v in g intersects Y nontrivially. (See [41] for a careful definition of “successor” when v is the last simplex of g). The simplex v is uniquely determined by Y and g ; to emphasize this relationship we say that $Y \searrow^d g$ at v . We define $g \not\searrow^d Y$ similarly with successor replaced by predecessor.

This relation yields a subordinacy relation among tight geodesics, namely that $g \searrow^d h$ when $D(g) \searrow^d h$ at v , and $\mathbf{T}(g)$ is the restriction to $D(g)$ of the successor of v ; and similarly $h \not\searrow^d g$ (again the definition here requires extra attention if $D(g)$ is an annulus). We let \searrow and $\not\searrow$ denote the transitive closures of \searrow^d and $\not\searrow^d$. Note that $Y \searrow g$ makes sense for a domain Y and geodesic g .

Hierarchies. A *hierarchy of tight geodesics* (henceforth just “hierarchy”) is a collection of tight geodesics in subsurfaces of S meant to “connect” two markings. There is a *main geodesic* g_H whose domain is $D(g_H) = S$, and all other geodesics are obtained by the rule that, if Y is a subsurface such that $b \not\searrow^d Y \searrow^d f$ for some $b, f \in H$, then there should be a (unique) geodesic $h \in H$ such that $D(h) = Y$ and $b \not\searrow^d h \searrow^d f$ (this determines $\mathbf{I}(h)$ and $\mathbf{T}(h)$ uniquely). The initial and terminal markings $\mathbf{I}(g_H)$ and $\mathbf{T}(g_H)$ are denoted $\mathbf{I}(H)$ and $\mathbf{T}(H)$ respectively, and we show in [41] that these two markings, when they are finite, determine H up to finitely many choices. In [47] (Lemma 5.13) we extend the construction to the case of generalized markings, and show that a hierarchy exists for any pair \mathbf{I}, \mathbf{T} of generalized markings such that no two infinite-leaved components of $\text{base}(\mathbf{I})$ and $\text{base}(\mathbf{T})$ are the same.

Hierarchy Structure Theorem. Theorem 4.7 of [41] (and its slight extension Theorem 5.6 of [47]) gives the basic structural properties of the subordinacy relations, and how they organize the hierarchy. In particular, it states that for $g, h \in H$ we have $g \searrow h$ if and only if $D(g) \subset D(h)$, and $\mathbf{T}(h)$ intersects $D(g)$ nontrivially (and similarly replacing \searrow with $\not\searrow$ and \mathbf{T} with \mathbf{I}). We quote the theorem here (as it appears in [47]) for the reader's

convenience, as we will use it often. For a subsurface $Y \subseteq S$ and hierarchy H let $\Sigma_H^+(Y)$ denote the set of geodesics $f \in H$ such that $Y \subseteq D(f)$ and $\mathbf{T}(f)$ intersects Y essentially (when Y is an annulus this means either that $\text{base}(\mathbf{T}(f))$ has homotopically nontrivial intersection with Y , or $\text{base}(\mathbf{T}(f))$ has a component equal to the core of Y , with a nonempty transversal). Similarly define $\Sigma_H^-(Y)$ with $\mathbf{I}(f)$ replacing $\mathbf{T}(f)$.

Theorem 2.1. (Descent Sequences) *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 $\Sigma_H^\pm(Y)$ 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 $k \in H$*

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

and similarly,

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

If, furthermore, $\Sigma_H^\pm(Y)$ 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 $h = h'$.*

When there is no chance of confusion, we will often denote Σ_H^\pm as Σ^\pm .

Slices and resolutions. A slice of a hierarchy is a combinatorial analogue of a cross-sectional surface in $S \times \mathbb{R}$. Formally, a slice is a collection τ of pairs (h, w) where h is a geodesic in H and w is a simplex in h , with the following properties. τ contains a distinguished “bottom pair” $p_\tau = (g_\tau, v_\tau)$. For each (k, u) in τ other than the bottom pair, there is an $(h, v) \in \tau$ such that $D(k)$ is a component domain of $(D(h), v)$; moreover any geodesic appears at most once among the pairs of τ . (See [47, §5.2] for more details).

We say a slice τ is *saturated* if, for every pair $(h, v) \in \tau$ and every geodesic $k \in H$ with $D(k)$ a component domain of $(D(h), v)$, there is some (hence exactly one) pair $(k, u) \in \tau$. It is easy to see by induction that for any pair (h, u) there is a saturated slice with bottom pair (h, u) . A slightly stronger condition is that, for every $(h, u) \in \tau$ and every component domain Y of $(D(h), u)$ with $\xi(Y) \neq 3$ there is, in fact, a pair $(k, u) \in \tau$ with $D(k) = Y$; we then say that τ is *full*. It is a consequence of Theorem 2.1 that, if $\mathbf{I}(H)$ and $\mathbf{T}(H)$ are maximal markings, then every saturated slice is full.

A slice is *maximal* if it is full and its bottom geodesic is the main geodesic g_H .

A *non-annular slice* is a slice in which none of the pairs $(k, u) \in \tau$ have annulus domains. A non-annular slice is saturated, full, or maximal if the conditions above hold with the exception of annulus domains. In particular, Theorem 2.1 implies that if $\text{base}(\mathbf{I}(H))$ and $\text{base}(\mathbf{T}(H))$ are maximal laminations, then every saturated non-annular slice is a full non-annular slice.

To a slice τ is associated a clean marking μ_τ , whose base is simply the union of simplices w over all pairs $(h, w) \in \tau$ with nonannular domains. The transversals of the marking are determined by the pairs in τ with annular domains, if any (see [47]). We also denote $\text{base}(\mu_\tau)$ by $\text{base}(\tau)$. We note that if τ is maximal then μ_τ is a maximal marking, and if τ is maximal non-annular then μ_τ is a pants decomposition.

We also refer to [41, §5] for the notion of “(forward) elementary move” on a slice, denoted $\tau \rightarrow \tau'$. The main effect of this move is to replace one pair (h, v) in τ with a pair (h, v') in τ' where v' is the successor of v in h . In addition certain pairs in τ whose domains lie in $D(h)$ are replaced with other pairs in τ' . The underlying curve system $\text{base}(\tau)$ stays the same except in the case that $\xi(h) = 4$. When $\xi(h) = 4$, v and v' intersect in a minimal way, and all other curves of $\text{base}(\tau)$ and $\text{base}(\tau')$ agree; if τ is maximal this amounts to a standard elementary move on pants decompositions.

A *resolution* of a hierarchy H is a sequence of elementary moves $(\dots \rightarrow \tau_n \rightarrow \tau_{n+1} \rightarrow \dots)$ (possibly infinite or biinfinite), where each τ_n is a saturated slice with bottom geodesic g_H , with the additional property that every pair (h, u) (with $h \in H$ and u a simplex of h) appears in some τ_n . Lemmas 5.7 and 5.8 of [47] guarantee that every hierarchy has a resolution with this property. A resolution is closely related to a “sweep” of $S \times \mathbb{R}$ by cross-sectional surfaces, and this will be exploited more fully in §4.

It is actually useful not to involve the annulus geodesics in a resolution. Thus given a hierarchy H one can delete all annulus geodesics to obtain a *hierarchy without annuli* H' (see [41, §8]) and a resolution of H' will be called a non-annular resolution.

Partial orders. In [41] we introduce several partial orders on the objects of a hierarchy H . In this section we extend the notion of “time order” \prec_t on geodesics to a time order on component domains, and we recall the properties of the partial order \prec_p on pairs.

First, for a subsurface Y of $D(g)$, define the *footprint* $\phi_g(Y)$ to be the set of simplices of g which represent curves disjoint from Y . Tightness implies that $\phi_g(Y)$ is always an interval of g , and the triangle inequality in $\mathcal{C}(D(g))$ implies it has diameter at most 2.

If X and Y are component domains arising in H , we say that $X \prec_t Y$ whenever there is a “comparison geodesic” $m \in H$ such that $D(m)$ contains

X and Y with nonempty footprints, and

$$\max \phi_m(X) < \min \phi_m(Y). \quad (2.1)$$

(Max and min are with respect to the natural order $v_i < v_{i+1}$ of the simplices of m). Note that (2.1) also implies that $Y \searrow m$ and $m \swarrow X$, by Theorem 2.1.

For geodesics g and h in H we can define $g \prec_t h$ if $D(g) \prec_t D(h)$, and this is equivalent to definition 4.16 in [41]. (We can similarly define $g \prec_t Y$ and $Y \prec_t h$.) In Lemma 4.18 of [41] it is shown, among other things, that \prec_t is a strict partial order on the geodesics in H , and it follows immediately that it is a strict partial order, with our definition, on all domains of geodesics in H . It is not hard to generalize this and the rest of that lemma to the set of *all* component domains in H , which for completeness we do here. (The main point of this generalization is to deal appropriately with 3-holed spheres, which can be component domains but never support geodesics.)

Lemma 2.2. *Suppose that H is a hierarchy with base($\mathbf{I}(H)$) or base($\mathbf{T}(H)$) maximal. The relation \prec_t is a strict partial order on the set of component domains occurring in H . Moreover, if Y and Z are component domains, then*

- (1) *If $Y \subseteq Z$ then Y and Z are not \prec_t -ordered.*
- (2) *Suppose that $Y \cap Z \neq \emptyset$ and neither domain is contained in the other. Then Y and Z are \prec_t -ordered.*
- (3) *If $b \swarrow Y \searrow f$ then either $b = f$, $b \searrow f$, $b \swarrow f$, or $b \prec_t f$.*
- (4) *If $Y \searrow m$ and $m \prec_t Z$ then $Y \prec_t Z$. Similarly if $Y \prec_t m$ and $m \swarrow Z$ then $Y \prec_t Z$.*

Proof. We follow the proof of Lemma 4.18 in [41], making adjustments for the fact that the domains may not support geodesics. Let us first prove the following slight generalization of Corollary 4.14 of [41], in which Y was assumed to be the domain of a geodesic in H .

Lemma 2.3. *If h is a geodesic in a hierarchy H , Y is a component domain in H , and $Y \not\subseteq D(h)$, then $\phi_h(Y)$ is nonempty.*

Proof. If Y fails to intersect $\mathbf{T}(h)$ then in particular it is disjoint from the last simplex of h , and hence $\phi_h(Y)$ is nonempty. If Y does intersect $\mathbf{T}(h)$ then, by Theorem 2.1, we have $Y \searrow h$. This means that there is some m with $Y \searrow m \searrow^d h$ and $\phi_h(D(m))$ is therefore nonempty, and is contained in $\phi_h(Y)$. \square

Now let us assume that base($\mathbf{I}(H)$) is a maximal lamination; the proof works similarly for $\mathbf{T}(H)$ (this assumption is used just once in the proof of part (2), and we suspect that the lemma should be true without it).

To prove part (1), suppose $Y \subseteq Z$. Then for any geodesic m with $Z \subseteq D(m)$, we have $\phi_m(Z) \subseteq \phi_m(Y)$. In particular the footprints of Y and Z can never be disjoint, and hence they are not \prec_t -ordered.

To prove part (2), let us first establish the following statement:

- (*) If $m \in H$ is a geodesic such that $Y \cup Z \subset D(m)$, where Y and Z are component domains in H which intersect but neither is contained in the other, and in addition we have either $m \not\prec Y$ or $m \not\prec Z$, then Y and Z are \prec_t -ordered.

The proof will be by induction on $\xi(m)$. If $\xi(m)$ equals $\xi(Y)$ or $\xi(Z)$ then $D(m)$ equals one of Y or Z and then the other is contained in it, but we have assumed this is not the case, so the statement holds vacuously.

Now assume that $\xi(m) > \max(\xi(Y), \xi(Z))$. Consider the footprints $\phi_m(Y)$ and $\phi_m(Z)$ (both nonempty by Lemma 2.3). If the footprints are disjoint then Y and Z are \prec_t -ordered with m the comparison geodesic, and we are done. If the footprints intersect then since they are intervals the minimum of one must be contained in the other. Let $v = \min \phi_m(Y)$ and $w = \min \phi_m(Z)$.

If $v < w$ then in particular $\phi_m(Z)$ does not include the first simplex of m , and so by Theorem 2.1 we have $m \not\prec Z$. This means that there is some m' with $m \not\prec m' \leq Z$ (where $m' \leq Z$ means either $m' \not\prec Z$ or $D(m') = Z$). $D(m')$ is a component domain of $(D(m), w)$, so since $w \in \phi_m(Y)$ and $Y \cap Z \neq \emptyset$ we find that $Y \subset D(m')$. Now by induction we may conclude that Y and Z are \prec_t -ordered.

If $w < v$ we of course apply the same argument with the roles reversed. If $w = v$ then we use the hypothesis that $m \not\prec Z$ or $m \not\prec Y$, and again repeat the previous argument. This concludes the proof of assertion (*).

To show that (2) follows from (*), it suffices to show the hypothesis holds for $m = g_H$. Suppose that $g_H \not\prec Y$ fails to hold. This means by Theorem 2.1) that Y does not intersect $\text{base}(\mathbf{I}(H))$, and since this is maximal Y must be either a 3-holed sphere with boundary in $\text{base}(\mathbf{I}(H))$, or an annulus with core in $\text{base}(\mathbf{I}(H))$. Now since Y and Z intersect, it follows that Z does intersect $\text{base}(\mathbf{I}(H))$ nontrivially. Hence again by Theorem 2.1 we have $g_H \not\prec Z$. Thus we may apply (*) to obtain (2).

To prove (3), suppose $b \neq f$. Suppose first that $D(b) \subset D(f)$. Since $Y \searrow f$, Theorem 2.1 implies that $f \in \Sigma^+(Y)$, and since $Y \subset D(b)$, we must have $f \in \Sigma^+(D(b))$ as well. Hence $b \searrow f$ by Theorem 2.1. Similarly if $D(f) \subset D(b)$ we have $b \not\prec f$. If neither domain is contained in the other, since they both contain Y we may apply (2) to conclude that they are \prec_t -ordered. Suppose by contradiction $f \prec_t b$, and let m be the comparison geodesic. Thus $f \searrow m \not\prec b$, and $\max \phi_m(f) < \min \phi_m(b)$. Since $Y \searrow f \searrow m$ we have by Corollary 4.11 of [41] that $\max \phi_m(Y) = \max \phi_m(f)$. Similarly $\min \phi_m(Y) = \min \phi_m(b)$. This contradicts $\max \phi_m(f) < \min \phi_m(b)$, so we conclude $b \prec_t f$.

To prove (4), consider the case $Y \searrow m \prec_t Z$ (the other case is similar). Let l be the comparison geodesic for m and Z . Then $\max \phi_l(D(m)) < \min \phi_l(Z)$. By Corollary 4.11 of [41], $Y \searrow m$ implies $\max \phi_l(Y) = \max \phi_l(D(m))$, and hence $Y \prec_t Z$.

Now we may finish the proof that \prec_t is a strict partial order. Suppose that $X \prec_t Y$ and $Y \prec_t Z$. Thus we have geodesics l and m such that $X \searrow l \swarrow Y \searrow m \swarrow Z$, and furthermore $\max \phi_l(X) < \min \phi_l(Y)$, and $\max \phi_m(Y) < \min \phi_m(Z)$.

Applying (3) to $l \swarrow Y \searrow m$ we find that $l = m$, $l \swarrow m$, $l \searrow m$, or $l \prec_t m$. Suppose first that $l = m$. Then $X \searrow m \swarrow Z$, and $\max \phi_m(X) < \min \phi_m(Y) \leq \max \phi_m(Y) < \min \phi_m(Z)$. Thus $X \prec_t Z$.

If $l \searrow m$ then $X \searrow m$ and by Corollary 4.11 of [41] we have $\max \phi_m(X) = \max \phi_m(D(l))$. Now since $D(l)$ contains Y we have $\phi_m(D(l)) \subset \phi_m(Y)$ and it follows that $\max \phi_m(D(l)) \leq \max \phi_m(Y) < \min \phi_m(Z)$. Thus again we have $X \prec_t Z$. The case $l \swarrow m$ is similar.

If $l \prec_t m$ then $X \searrow l \prec_t m \swarrow Z$ and applying (4) twice we conclude $X \prec_t Z$.

Hence \prec_t is transitive, and since by definition $X \prec_t X$ can never hold, it is a strict partial order. \square

There is a similar order on *pairs* (h, u) where h is a geodesic in H , and u is either a simplex of h , or $u \in \{\mathbf{I}(h), \mathbf{T}(h)\}$. We define a *generalized footprint* $\widehat{\phi}_m(g, u)$ to be $\phi_m(D(g))$ if $D(g) \subset D(m)$, and simply $\{u\}$ if $g = m$. We then say that

$$(g, u) \prec_p (h, v)$$

if there is an $m \in H$ with $D(h) \subseteq D(m)$, $D(g) \subseteq D(m)$, and

$$\max \widehat{\phi}_m(g, u) < \min \widehat{\phi}_m(h, v)$$

In particular if $g = h$ then \prec_p reduces to the natural order on $\{\mathbf{I}(g), u_0, \dots, u_N, \mathbf{T}(g)\}$ where $\{u_i\}$ are the simplices of g . This relation is also shown to be a partial order in Lemma 4.18 of [41].

In the proof of Lemma 5.1 in [41], the following fact is established which will be used in §4 and §5. It is somewhat analogous to part (1) of Lemma 2.2.

Lemma 2.4. *Any two elements (h, u) and (k, v) of a slice τ of H are not \prec_p -ordered.*

2.3. Margulis tubes

Tube constants. Let N_J for $J \subset \mathbb{R}$ denote the region $\{x \in N \mid 2 \operatorname{inj}_N(x) \in J\}$. Thus $N_{(0, \epsilon]}$ denotes the ϵ -thin part of a hyperbolic manifold N , and $N_{[\epsilon, \infty)}$ denotes the ϵ -thick part. Let ϵ_0 be a Margulis constant for \mathbb{H}^3 , so that for $\epsilon \leq \epsilon_0$, $N_{(0, \epsilon]}$ for a hyperbolic 3-manifold N is a disjoint union of standard closed tubular neighborhoods of closed geodesics, or horocyclic cusp neighborhoods. (See e.g. Benedetti-Petronio [9] or Thurston [69].)

We will call the interior of such a component an (open) ϵ -Margulis tube, and denote it by $\mathbb{T}_\epsilon(\gamma)$, where γ is the homotopy class of the core (or in the rank-2 cusp case, any nontrivial homotopy class in the tube). If Γ is a collection of simple closed curves or homotopy classes we will denote $\mathbb{T}(\Gamma)$ the union of the corresponding Margulis tubes.

Let $\epsilon_1 < \epsilon_0$ be chosen as in Minsky [47], so that the ϵ_0 -thick part of an essential pleated surface maps into the ϵ_1 -thick part of the target 3-manifold. This is the constant used in the Lipschitz Model Theorem (§2.7). It will be our “default” Margulis constant and we will usually denote \mathbb{T}_{ϵ_1} as just \mathbb{T} . (The only place we use ϵ_0 will be in the definition of the augmented convex core).

Let $\rho \in \mathcal{D}(S)$ be a Kleinian surface group, and $N = N_\rho$. Then N is homeomorphic to $\text{int}(S) \times \mathbb{R}$ (Bonahon [11]). More precisely, Bonahon showed that $N^1 \cong S \times \mathbb{R}$, where N^1 is $N_\rho \setminus \mathbb{T}(\partial S)$ – the complement of the open cusp neighborhoods associated to ∂S .

Thurston showed that sufficiently short primitive geodesic in N is homotopic to a simple loop in S . Otal proved the following stronger theorem in [53, 54]:

Theorem 2.5. *There exists $\epsilon_u > 0$ depending only on the compact surface S such that, if $\rho \in \mathcal{D}(S)$ and Γ is the set of primitive closed geodesics in $N = N_\rho$ of length at most ϵ_u , then Γ is unknotted and unlinked. That is, N^1 can be identified with $S \times \mathbb{R}$ in such a way that Γ is identified with a collection of disjoint simple closed curves of the form $c \times \{t\}$.*

We remark that Otal’s proof only explicitly treats the case that S is a closed surface, but the case with boundary is quite similar. One can also obtain this result, for any finite subcollection of Γ , by applying a special case of Souto [60].

Bounded homotopies into tubes. The next lemma shows that a bounded-length curve homotopic into a Margulis tube admits a controlled homotopy into the tube. It will be used at the end of Section 6.

Lemma 2.6. *Let \mathcal{U} be a collection of ϵ_1 -Margulis tubes in a hyperbolic 3-manifold N and let γ be an essential curve which is homotopic within $N \setminus \mathcal{U}$ to $\partial\mathcal{U}$.*

Then such a homotopy can be found whose diameter is bounded by a constant r depending only on ϵ_1 and the length $l_N(\gamma)$.

Proof. The choice of ϵ_1 strictly less than the Margulis constant for \mathbb{H}^3 implies that \mathcal{U} has an embedded collar neighborhood of definite radius, so possibly enlarging \mathcal{U} we may assume that the radius of each component is at least $R > 0$ (depending on ϵ_1). Let U be a component of \mathcal{U} with core geodesic c . Agol has shown in [1] (generalizing a construction of Kerckhoff) that there exists a metric g on $U \setminus c$ such that

- (1) g agrees with the metric of N on a neighborhood of ∂U .
- (2) g has all sectional curvatures between $[-\kappa_1, -\kappa_0]$, where $\kappa_1 > \kappa_0 > 0$ depend on R .
- (3) On some neighborhood of c , g is isometric to a rank-2 parabolic cusp.

Let \widehat{N} be the complete, negatively-curved manifold obtained by deleting the cores of components of \mathcal{U} and replacing the original metric by the metric

g in each one. The homotopy from γ to $\partial\mathcal{U}$ can be deformed to a *ruled annulus* $A : \gamma \times [0, 1] \rightarrow \widehat{N}$, i.e. a map such that $A(\cdot, 0) = id$, $A(\cdot, 1)$ has image in \mathcal{U} , and $A|_{x \times [0, 1]}$ is a geodesic. This is possible simply by straightening the trajectories of the original homotopy, since \widehat{N} is complete and negatively curved. Because a ruled surface has non-positive extrinsic curvature, the pullback metric on $\gamma \times [0, 1]$ must have curvatures bounded above by $-\kappa_0$. Furthermore, by pushing $A(\gamma \times \{1\})$ sufficiently far into the cusps of \widehat{N} , we can ensure that the total boundary length of the annulus is at most $l(\gamma) + 1$.

The area of the annulus is bounded by $Cl(\partial A) \leq C(l(\gamma) + 1)$, where C depends on the curvature bounds. Let $\epsilon = \epsilon_1/2$.

Let A' be the component of $A^{-1}(\widehat{N} \setminus \mathcal{U})$ containing the outer boundary $\gamma \times \{0\}$. This is a punctured annulus, and the punctures can be filled in by disks in $\gamma \times [0, 1]$. Let A'' denote the union of A' with these disks. The injectivity radius in A' is at least ϵ , and it follows that the same holds for A'' , since any essential loop passing through one of the added disks must also pass through A' . Let A''_ϵ be the complement in A'' of an ϵ -neighborhood of the outer boundary (in the induced metric). Any point in A''_ϵ is the center of an embedded disk of area at least $\pi\epsilon^2$, so the area bound on A implies that $\text{diam } A''_\epsilon \leq C'(l(\gamma) + 1)$. This gives a bound $\text{diam } A'' \leq C'(l(\gamma) + 1) + l(\gamma) + \epsilon$.

This bounds how far each of the disks of $A' \setminus A''$ reaches into the tubes \mathcal{U} , and hence bounds the distortion caused by pushing these disks back to $\partial\mathcal{U}$. Applying this deformation to A'' yields a homotopy of γ into $\partial\mathcal{U}$ with bounded diameter, as desired. \square

2.4. Geometric Limits

Let us recall the notion of geometric limits for hyperbolic manifolds with basepoints (N, x) . We say that $(N_i, x_i) \rightarrow (N, x)$ if for any $R > 0$ and $\epsilon > 0$ there is, for large enough i , a diffeomorphic embedding $f_i : B_R(x) \rightarrow N_i$ taking x to x_i so that f_i are ϵ -close, in C^2 , to local isometries. (here B_R denotes an R -neighborhood in N). We call f_i *comparison maps*.

Equivalently one can state the definition with the comparison maps going the other direction, $f_i : B_R(x_i) \rightarrow N$. It will be convenient for us to use both definitions.

It will help to have the following slight improvement of the behavior of these maps with respect to thin parts.

Lemma 2.7. *For a geometrically convergent sequence $(N_i, x_i) \rightarrow (N, x)$ we may choose the comparison maps f_i so that*

$$f_i((N_i)_{(0, \epsilon_1]} \cap B_R(x_i)) \subset N_{(0, \epsilon_1]}$$

and

$$f_i((N_i)_{[\epsilon_1, \infty)} \cap B_R(x_i)) \subset N_{[\epsilon_1, \infty)}.$$

The corresponding statement holds with the comparison maps going in the opposite direction.

Proof. Fix R . For sufficiently large i we may assume that the domain of f_i contains $B_{R+\epsilon_1}(x_i)$, and that its image contains $B_{R+\epsilon_1}(x)$. Choose also $\delta \ll 1$ and assume that f_i is $(1 + \delta)$ -bilipschitz. Now a loop of length ϵ_1 based at a point $p \in B_R(x_i)$ has image of length at most $(1 + \delta)\epsilon_1$, and this image must be homotopically nontrivial since if it bounds a disk it bounds a disk of diameter at most $(1 + \delta)\epsilon_1/2$ which we can assume is still contained in the image of f_i and hence pulls back to N_i . Thus we have

$$f_i((N_i)_{(0, \epsilon_1]} \cap B_R(x_i)) \subset N_{(0, (1+\delta)\epsilon_1]}$$

and a similar argument gives

$$f_i((N_i)_{[\epsilon_1, \infty)} \cap B_R(x_i)) \subset N_{[\epsilon_1/(1+\delta), \infty)}.$$

Now since the region $N_{[\epsilon_1/(1+\delta), \epsilon_1(1+\delta)]}$ is a bicollar neighborhood of radius $O(\delta)$ of ∂N_{ϵ_1} , and the map f_i is eventually δ -close to a local isometry in C^2 , the image of $\partial(N_i)_{(0, \epsilon_1]} \cap B_R(x_i)$ can be represented in the product structure of the collar as the graph of a nearly constant function over $\partial N_{(0, \epsilon_1]}$. We can then use the collar structure to adjust the map in this neighborhood so that it exactly respects the thick-thin decomposition, and is still close to isometric. □

One can also consider geometric limits from the point of view of groups: A sequence of Kleinian groups Γ_i converges geometrically to Γ if they converge in the Gromov-Hausdorff sense as subsets of $\mathrm{PSL}_2(\mathbb{C})$. This is equivalent to geometric convergence of $(\mathbb{H}^3/\Gamma_i, x_i)$ to $(\mathbb{H}^3/\Gamma, x)$, where x and x_i are the images of a fixed basepoint in \mathbb{H}^3 under the quotient by Γ and Γ_i , respectively. See Benedetti-Petronio [9] for more details.

If G is any finitely-generated group, the set $\mathcal{D}(G)$ of discrete, faithful representations $\rho : G \rightarrow \mathrm{PSL}_2(\mathbb{C})$ is given the natural topology of convergence on each element of G , also called the topology of algebraic convergence. Algebraic and geometric convergence are related but quite different. Part of this relationship is described by the following proposition, which is a relative version of Lemma 3.6 and Proposition 3.8 in [32].

Proposition 2.8. *Let G be a torsion-free group and let H be a non-abelian subgroup of G . Let $\{\rho_i\}$ be a sequence in $\mathcal{D}(G)$ such that $\{\rho_i|_H\}$ converges to $\rho \in \mathcal{D}(H)$. Then*

- (1) *If $\{g_i\}$ is a sequence in G and $\{\rho_i(g_i)\}$ converges to the identity, then g_i is equal to the identity for all large enough i .*
- (2) *There exists a subsequence $\{\rho_{i_n}\}$ of $\{\rho_i\}$ such that $\{\rho_{i_n}(G)\}$ converges geometrically to a torsion-free Kleinian group Γ .*

Proof. The key point here is that the sequence $\{\rho_i(G)\}$ is uniformly discrete. Let h_1 and h_2 be two non-commuting elements of H . Since $\{\rho_i|_H\}$ is algebraically convergent, there exists $K > 0$ such that $d(\rho_i(h_k)(0), 0) \leq K$ for all i and $k = 1, 2$. The Margulis lemma implies that given K there exists $\epsilon > 0$ such that if $\alpha, \beta \in \mathrm{PSL}_2(\mathbf{C})$ and $d(\alpha(0), 0) < \epsilon$ and $d(\beta(0), 0) \leq K$, then either α and β commute or the group they generate is indiscrete or has torsion. Thus, if $\gamma \in \rho_i(G)$, then $d(0, \gamma(0)) \geq \epsilon$, since $\rho_i(G)$ is discrete and torsion-free (for all i) and γ cannot commute with both $\rho_i(h_1)$ and $\rho_i(h_2)$. The first fact follows immediately, while the second fact follows from Theorem 3.1.4 in Canary-Epstein-Green [21]. \square

2.5. Ends and ending laminations

We recap here the definitions of ends and ending laminations. See [47] for more details.

Fix $\rho \in \mathcal{D}(S)$. Let $N = N_\rho$ and let C_N denote the convex core of N . Let Q denote the union of (open) ϵ_1 -Margulis tube neighborhoods of the cusps of N , and let $N^0 = N \setminus Q$. Let $Q_1 \subset Q$ be the union of tubes associated to ∂S (thus $Q_1 = \mathbb{T}(\partial S)$, and $N \setminus Q_1$ is N^1 , as defined in §2.3). Let $\partial_\infty N$ denote the conformal boundary of N at infinity, obtained as the quotient of the domain of discontinuity of $\rho(\pi_1(S))$.

As a consequence of Bonahon's tameness theorem [11], we can fix an orientation-preserving identification of N with $\widehat{S} \times \mathbb{R}$ in such a way that $Q_1 = \mathbf{collar}(\partial S) \times \mathbb{R}$, and furthermore so that $\mathcal{K} \equiv S \times [-1, 1]$ meets the closure of Q in a union of disjoint essential annuli $P = P_1 \cup P_+ \cup P_-$, where $P_1 = \partial S \times [-1, 1]$ and $P_\pm \subset \mathit{int}(S) \times \{\pm 1\}$. The pair (\mathcal{K}, P) is the *relative compact core* of N_0 , and the components of $N_0 \setminus \mathcal{K}$ are neighborhoods of the ends of N_0 .

For each component R of $\partial \mathcal{K} \setminus P$, there is an invariant ν_R defined as follows: If the end associated to R is geometrically finite then ν_R is the point in the Teichmüller space $\mathcal{T}(R)$ associated to the component of the conformal boundary $\partial_\infty N$ that faces R . If the end is geometrically infinite then (again by Bonahon [11] and by Thurston's original definition in [66]) ν_R is a geodesic lamination in $\mathcal{EL}(R)$, which is the unique limit (in the measure-forgetting topology) of sequences of simple closed curves in R whose geodesic representatives exit the end associated to R .

ν_+ then has a lamination part ν_+^L and a Riemann surface part ν_+^T : ν_+^L is the union of the core curves p_+ of the annuli P_+ (these are the "top parabolics") and the laminations ν_R for those components R of $S \times \{+1\} \setminus P_+$ which correspond to simply degenerate ends. ν_+^T is the union of $\nu_R \in \mathcal{T}(R)$ for those components R of $S \times \{+1\} \setminus P_+$ which correspond to geometrically finite ends. We define ν_- similarly. We let ν , or $\nu(\rho)$, denote the pair (ν_+, ν_-) .

Note in particular the special case that there are no parabolics except for P_1 , and both the $+$ and $-$ ends are degenerate. In this case ν_+ and ν_- are

both filling laminations in $\mathcal{EL}(S)$. This is called the *doubly degenerate case*, and it is helpful for most of this paper to focus just on this case.

2.6. Definition of the model manifold

We recall here the definition of the model manifold M_ν from [47].

Given a pair $\nu = (\nu_+, \nu_-)$ of end invariants, we construct as in [47, §7.1] a pair of markings μ_\pm which encode the geometric information in ν up to bilipschitz equivalence. In particular, when ν_+ is a filling lamination (the $+$ end is simply degenerate), $\text{base}(\mu_+) = \nu_+$. When ν_+ is a point in Teichmüller space, μ_+ is a minimal-length marking in the corresponding metric on S . In general $\text{base}(\mu_\pm)$ is a maximal lamination (maximal among supports of measured laminations) whose infinite-leaf components are ending laminations for ends of the manifold N_0 obtained from N by cutting along cusps. The closed-leaf components of $\text{base}(\mu_\pm)$ which are not equipped with transversals are exactly the (non-peripheral) parabolics of N .

Note also that a component can be common to $\text{base}(\mu_+)$ and $\text{base}(\mu_-)$ only if it is a closed curve, and has a transversal on at least one of the two. This is because a non-peripheral parabolic in N corresponds to a cusp in either side of the compact core or the other.

We let $H = H_\nu$ be a hierarchy such that $\mathbf{I}(H) = \mu_-$ and $\mathbf{T}(H) = \mu_+$.

The model manifold M_ν is identified with a subset of $\widehat{M} \equiv \widehat{S} \times \mathbb{R}$, and is partitioned into pieces called *blocks* and *tubes*. It is also endowed with a path metric (in fact piecewise smooth).

Doubly degenerate case. We give first the description of the model when both ν_\pm are filling laminations. In this case N has two simply degenerate ends and no non-peripheral parabolics, and the main geodesic g_H is doubly infinite.

The blocks are associated to *4-edges*, which are edges e of geodesics $h \in H$ with $\xi(h) = 4$. For each such e the block $B(e)$ is isotopic to $D(h) \times [-1, 1]$ in $S \times \mathbb{R}$. More precisely, we can identify each $B(e)$ abstractly with

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

That is, $B(e)$ is a product with solid-torus “trenches” dug out of its top and bottom corresponding to the vertices e^\pm . This abstract block is embedded in M *flatly*, which means that each connected leaf of the horizontal foliation $Y \times \{t\}$ is mapped to a level set $Y \times \{s\}$ in the image, with the map on the first factor being the identity. (In [47] we first build the abstract union of blocks and then prove it can be embedded).

The 3-holed spheres coming from $(D(e) \setminus \text{collar}(e^\pm)) \times \{\pm 1\}$ are called the *gluing boundaries* of the block. We show in [47] that every 3-holed sphere Y that arises as a component domain in H appears as a gluing boundary of exactly two blocks, and these blocks are in fact attached along these

boundaries via the identity map on Y . The resulting level surface $Y \times \{s\}$ in $\widehat{M} = \widehat{S} \times \mathbb{R}$ will always be denoted F_Y .

The complement of the blocks in \widehat{M} is a union of solid tori of the form $U(v) = \mathbf{collar}(v) \times I_v$, where v is a vertex in H or a boundary component of S , and I_v is an interval.

If v is a boundary component of S then $I_v = \mathbb{R}$. Otherwise, since we are describing the doubly degenerate case, I_v is always a compact interval.

Geometry and tube coefficients. The model is endowed with a metric in which the (non-boundary) blocks fall into a finite number of isometry classes (in fact two, depending on the topological type), and in which all the annuli in the boundaries are Euclidean, with circumference ϵ_1 . Thus every torus $\partial U(v)$ is equipped with a Euclidean metric.

This allows us to associate to $U(v)$ a coefficient $\omega(v) \in \mathbb{H}^2$ (in [47] denoted $\omega_M(v)$), defined as follows: $\partial U(v)$ comes with a preferred marking (α, μ) where α is the core curve of any of the annuli making up $\partial U(v)$ and μ is a meridian curve of the solid torus $U(v)$. This together with the Euclidean structure on $\partial U(v)$ determines a point in the Teichmüller space of the torus which is just \mathbb{H}^2 .

This information uniquely determines a metric on $U(v)$ (modulo isotopy fixing the boundary) which makes it isometric to a hyperbolic Margulis tube. The radius of this tube is given by

$$r = \sinh^{-1} \frac{\epsilon_1 |\omega|}{2\pi} \quad (2.2)$$

and the complex translation length of the element generating this tube is given by

$$\lambda = h_r \left(\frac{2\pi i}{\omega} \right) \quad (2.3)$$

where $h_r(z) = \operatorname{Re} z \tanh r + i \operatorname{Im} z$ (see §3.2 of [47]). Note in particular that r grows logarithmically with $|\omega|$, and that for large $|\omega|$, $2\pi i/\omega$ becomes a good approximation for λ .

When v corresponds to a boundary component of S (or, in the general case, to a parabolic component of base μ_{\pm}), we write $\omega(v) = i\infty$ and we make $U(v)$ isometric to a cusp associated to a rank 1 parabolic group.

We let $M_{\nu}[0]$ denote the union of the blocks, i.e. M_{ν} minus the interiors of the tubes. For any $k \in [0, \infty]$ we let $M_{\nu}[k]$ denote the union of $M_{\nu}[0]$ with the tubes $U(v)$ for which $|\omega(v)| < k$ (in particular note that $M_{\nu}[\infty]$ excludes exactly the parabolic tubes).

We let \mathcal{U} denote the union of all the tubes in the model, and let $\mathcal{U}[k]$ denote those tubes with $|\omega| \geq k$. Thus $M_{\nu}[k] = M_{\nu} \setminus \mathcal{U}[k]$.

The case with boundary. When N has geometrically finite ends, ν_{\pm} are not filling laminations, the main geodesic g_H is not bi-infinite, and the model manifold has some boundary. The construction then involves a finite number of “boundary blocks.”

A boundary block is associated to a geometrically finite end of N_0 . Let R be a subsurface of S homotopic to a component of $S \times \{1\} \setminus P_+$ which faces a geometrically finite end, 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

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

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*. 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

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

Similarly if R is a component of $S \times \{-1\} \setminus P_-$ associated to a geometrically finite end, we let $\mathbf{I}_R = \mathbf{I}(H_\nu) \cap R$ and define

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

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

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 .

To put a metric on a boundary block, we let σ^m denote the conformal rescaling of the Poincaré metric on $\partial_\infty N$ which makes the collars of curves of length less than ϵ_1 into Euclidean cylinders (and is the identity outside the collars). Identifying the outer boundary of the block with the appropriate component of $\partial_\infty N$ we pull back σ^m , and then extend using the product structure of the block. See §8.3 of [47] for details.

2.7. The bilipschitz model theorem

Lipschitz model theorem. We begin by describing the main theorem of [47]. Again, $\rho \in \mathcal{D}(S)$ is a Kleinian surface group with quotient manifold $N = N_\rho$ and end invariants ν .

If U is a tube of the model manifold, let $\mathbb{T}(U)$ denote the ϵ_1 -Margulis tube (if any) whose homotopy class is the image via ρ of the homotopy class of U . For $k > 0$ let $\mathbb{T}[k]$ denote the union of $\mathbb{T}(U)$ over tubes U in $\mathcal{U}[k]$.

The *augmented convex core* of N is $\widehat{C}_N = C_N^1 \cup N_{(0, \epsilon_0]}$ where C_N^1 is the closed 1-neighborhood of the convex core C_N of N . We show in [47] that this is homeomorphic to C_N , and hence to M_ν .

Definition 2.9. A (K, k) model map for ρ is a map $f : M_\nu \rightarrow \widehat{C}_N$ satisfying the following properties:

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

Lipschitz Model Theorem. [47] *There exist $K, k > 0$ depending only on S , so that for any $\rho \in \mathcal{D}(S)$ with end invariants ν there exists a (K, k) model map*

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

The exterior of the augmented core

In fact what we really want is a model for all of N , not just its augmented convex core. Thus we need a description of the exterior of \widehat{C}_N . This was done in Minsky [47], by a slight generalization of the work of Sullivan and Epstein-Marden in [22]. Let E_N denote the closure of $N \setminus \widehat{C}_N$ in N , $\partial_\infty N$ the conformal boundary at infinity of N , and $\bar{E}_N = E_N \cup \partial_\infty N$. The metric σ^m on $\partial_\infty N$ is defined as in §2.6, as the Poincaré metric adjusted conformally so that every thin tube and cusp becomes a Euclidean annulus.

Let E_ν denote a copy of $\partial_\infty N \times [0, \infty)$, endowed with the metric

$$e^{2r} \sigma^m + dr^2$$

where r is a coordinate for the second factor.

The boundary of M_ν is naturally identified with $\partial_\infty N$, and this enables us to form ME_ν as the union of M_ν with E_ν identifying $\partial_\infty N \times \{0\}$ with ∂M_ν . We attach a boundary at infinity $\partial_\infty N \times \{\infty\}$ to E_ν , obtaining a manifold with boundary \overline{ME}_ν . We also denote this boundary at infinity as $\partial_\infty ME_\nu$.

In Lemma 3.5 of [47] we give a uniformly bilipschitz homeomorphism of E_ν to E_N , which extends to conformal homeomorphism on the boundaries at infinity, and together with the Lipschitz Model Theorem gives the following (called the Extended Model Theorem in [47]):

Theorem 2.10. *There exists a proper degree 1 map*

$$f : ME_\nu \rightarrow N,$$

which is a (K, k) model map from M_ν to \widehat{C}_N , restricts to a K -bilipschitz homeomorphism $\varphi : E_\nu \rightarrow E_N$, and extends to a conformal isomorphism from $\partial_\infty ME_\nu$ to $\partial_\infty N$. The constants K and k depend only on the topology of N .

The main result of this paper will be the upgrading of this model map to a bilipschitz map:

Bilipschitz Model Theorem. *There exist $K, k > 0$ depending only on S , so that for any Kleinian surface group $\rho \in \mathcal{D}(S)$ with end invariants $\nu = (\nu_+, \nu_-)$ there is an orientation-preserving K -bilipschitz homeomorphism of pairs*

$$F : (M_\nu, \mathcal{U}[k]) \rightarrow (\widehat{C}_{N_\rho}, \mathbb{T}[k]).$$

Furthermore this map extends to a homeomorphism

$$\bar{F} : \overline{ME}_\nu \rightarrow \bar{N}$$

which restricts to a K -bilipschitz homeomorphism from ME_ν to N , and a conformal isomorphism from $\partial_\infty ME_\nu$ to $\partial_\infty N$.

Length estimates. The length $\ell_\rho(v)$ of a simple closed curve v was bounded above using end-invariant data in Minsky [45]. Lower bounds for ℓ_ρ and for the complex translation length λ_ρ were obtained in [47] using the Lipschitz Model Theorem. The following is a slight restatement of the second main theorem of [47], which incorporates this information.

Short Curve Theorem. *There exist $\bar{\epsilon} > 0$ and $c > 0$, and a function $\Omega : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, depending only on S , such that the following holds: Given a surface group $\rho \in \mathcal{D}(S)$, and any vertex $v \in \mathcal{C}(S)$,*

- (1) *If $\ell_\rho(v) < \bar{\epsilon}$ then v appears in the hierarchy H_{ν_ρ} .*
- (2) *(Lower length bounds) If v appears in H_{ν_ρ} then*

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

and

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

- (3) *(Upper length bounds) If v appears in H_{ν_ρ} and $\epsilon > 0$ then*

$$|\omega_M(v)| \geq \Omega(\epsilon) \implies \ell_\rho(v) \leq \epsilon.$$

The quantity $|\omega_M(v)|$ can be estimated from the lengths of the geodesics in the hierarchy whose domains border v . In particular, Theorem 9.1 and Proposition 9.7 of [47] give the following:

Lemma 2.11. *There exist positive constants b_1 and b_2 depending on S such that, for any hierarchy H and associated model, if v is any vertex of $\mathcal{C}(S)$,*

$$|\omega_M(v)| \geq -b_1 + b_2 \sum_{\substack{h \in H \\ v \in [\partial D(h)]}} |h| \tag{2.4}$$

Putting the Short Curve Theorem together with Lemma 2.11, we obtain:

Lemma 2.12. *There is a function $\mathcal{L} : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, depending only on S , such that, given $\rho \in \mathcal{D}(S)$, for any geodesic h in H_{ν_ρ} , and any $\epsilon > 0$,*

$$|h| \geq \mathcal{L}(\epsilon) \implies \ell_\rho(\partial D(h)) \leq \epsilon. \quad (2.5)$$

The Bilipschitz Model Theorem proved in this paper will allow us to give the following improvement of the Short Curve Theorem:

Length Bound Theorem. *There exist $\bar{\epsilon} > 0$ and $c > 0$ depending only on S , such that the following holds:*

Let $\rho : \pi_1(S) \rightarrow \mathrm{PSL}_2(\mathbb{C})$ be a Kleinian surface group and v a vertex of $\mathcal{C}(S)$, and let H_{ν_ρ} be an associated hierarchy.

- (1) *If $\ell_\rho(v) < \bar{\epsilon}$ then v appears in H_{ν_ρ} .*
- (2) *If v appears in H_{ν_ρ} then*

$$d_{\mathbb{H}^2} \left(\omega_M(v), \frac{2\pi i}{\lambda_\rho(v)} \right) \leq c$$

The distance estimate in part (2) is natural because ω_M is a Teichmüller parameter for the boundary torus of a Margulis tube, as is $2\pi/\lambda$, and $d_{\mathbb{H}^2}$ is the Teichmüller distance. A bound on $d_{\mathbb{H}^2}$ corresponds directly to a model for the Margulis tube that is correct up to bilipschitz distortion. See the discussion in Minsky [44, 47]. The proof of the Length Bound theorem will be given in §9.2.

3. Knotting and partial order of subsurfaces

In this section we will study the topological ordering of surfaces in a product manifold $M = S \times \mathbb{R}$, where S is a compact surface. We will define “scaffolds” in M , which are collections of embedded surfaces and solid tori satisfying certain conditions. Scaffolds arise naturally in the model manifold as unions of cut surfaces and tubes. Our main goal, encapsulated in Theorem 3.8 (Scaffold Extension), is to show that two scaffolds embedded in M and satisfying consistent topological order relations have homeomorphic complements. This will allow us, in the final part of the proof (§8.4), to adjust our model map to be a homeomorphism on selected regions.

In sections 3.7 and 3.8 we use the technology developed in the proof of Theorem 3.8 (Scaffold Extension) to develop technical lemmas which will be useful later in the paper.

3.1. Topological order relation

Recall that $M = S \times \mathbb{R}$ and $\widehat{M} = \widehat{S} \times \mathbb{R}$. Let $s_t : \widehat{S} \rightarrow \widehat{M}$ be the map $s_t(x) = (x, t)$, and let $\pi : \widehat{M} \rightarrow \widehat{S}$ be the map $\pi(x, t) = x$.

For $R \subseteq S$ an essential non-annular surface, let $\mathbf{map}(R, M)$ denote the homotopy class $[s_0|_R]$.

If R is a closed annulus we want $\mathbf{map}(R, M)$ to denote a certain collection of maps of solid tori into \widehat{M} . Thus, we consider proper maps of the form $f : V \rightarrow \widehat{M}$ where $V = R \times J$, J is a closed connected subset of \mathbb{R} , and for

any $t \in J$ $f \circ s_t : R \rightarrow M$ is in $[s_0|_R]$. If R is a nonperipheral annulus then J is a finite or half-infinite interval. If R is peripheral then we require $J = \mathbb{R}$. We say that these maps are of “annulus type”.

Let $\mathbf{map}(M)$ denote the disjoint union of $\mathbf{map}(R, M)$ over all essential subsurfaces R .

We say that $f \in \mathbf{map}(R_1, M)$ and $g \in \mathbf{map}(R_2, M)$ *overlap* if R_1 and R_2 have essential intersection.

We define a *topological order relation* \prec_{top} on $\mathbf{map}(R, M)$ as follows. First, we say that $f : R \rightarrow M$ is *homotopic to $-\infty$ in the complement of $X \subset M$* if for some r there is a proper map

$$F : R \times (-\infty, 0] \rightarrow M \setminus (X \cup S \times [r, \infty))$$

such that $F(\cdot, 0) = f$. We define *homotopic to $+\infty$ in the complement of X* similarly. (The definition when R is an annulus is similar, where we then consider the map of the whole solid torus $V = R \times J$.)

Now, given $f \in \mathbf{map}(R, M)$ and $g \in \mathbf{map}(Q, M)$ we write $f \prec_{\text{top}} g$ if and only if

- (1) f and g have disjoint images.
- (2) f is homotopic to $-\infty$ in the complement of $g(Q)$, but f is not homotopic to $+\infty$ in the complement of $g(Q)$.
- (3) g is homotopic to $+\infty$ in the complement of $f(R)$, but g is not homotopic to $-\infty$ in the complement of $f(R)$.

The next lemma states some elementary observations about this notion of ordering.

Lemma 3.1. *Let R and Q be essential subsurfaces of S which intersect essentially.*

- (1) *If $f \in \mathbf{map}(R, M)$ and $g \in \mathbf{map}(Q, M)$ have disjoint images and f is homotopic to $-\infty$ in the complement of $g(Q)$, then f cannot be homotopic to $+\infty$ in the complement of $g(Q)$.*
- (2) *Similarly if g is homotopic to $+\infty$ in the complement of $f(R)$, then g is not homotopic to $-\infty$ in the complement of $f(R)$.*
- (3) *For the level mappings $s_t(x) = (x, t)$, we have*

$$s_t|_R \prec_{\text{top}} s_r|_Q$$

if and only if $t < r$.

Proof. Since R and Q overlap, there exist curves α in R and β in Q that intersect essentially. If f is homotopic to both $+\infty$ and $-\infty$ in the complement of g then we may construct a map of $\alpha \times \mathbb{R}$ to M which is properly homotopic to the inclusion map and misses $g(\beta)$. Since $g(\beta)$ is homotopic to $\beta \times \{0\}$, this contradicts the essential intersection of α and β . This gives (1), and a similar argument gives (2).

For (3), it is clear that when $t < r$ $s_r|_Q$ is homotopic to $+\infty$ in the complement of $s_t(R)$, and that $s_t|_R$ is homotopic to $-\infty$ in the complement of $s_r(Q)$. The rest follows from (1) and (2). \square

Example 3.2. We note also that \prec_{top} does not extend to a partial order on $\mathbf{map}(M)$. In fact, we may construct a finite family of disjoint embeddings all of which are ordered with respect to one another, such that \prec_{top} is not transitive on the family. Let γ_1, γ_2 and γ_3 be three disjoint curves on a surface S of genus 4 such that the components A_1, A_2 and A_3 of $S \setminus \cup_{i=1}^3 \mathbf{collar}(\gamma_i)$ are all twice-punctured tori. Moreover, we may assume that $P = A_1 \cup \mathbf{collar}(\gamma_1) \cup A_2$, $Q = A_2 \cup \mathbf{collar}(\gamma_2) \cup A_3$ and $R = A_3 \cup \mathbf{collar}(\gamma_3) \cup A_1$, are all connected. Let $f : P \rightarrow M$ map A_1 to $A_1 \times \{0\}$, A_2 to $A_2 \times \{1\}$ and $\mathbf{collar}(\gamma_1)$ to an annulus in $\mathbf{collar}(\gamma_1) \times [0, 1]$. Similarly, let $g : Q \rightarrow M$ map A_2 to $A_2 \times \{0\}$, A_3 to $A_3 \times \{1\}$ and $\mathbf{collar}(\gamma_2)$ to an annulus in $\mathbf{collar}(\gamma_2) \times [0, 1]$, and let $h : R \rightarrow M$ map A_3 to $A_3 \times \{0\}$, A_1 to $A_1 \times \{1\}$ and $\mathbf{collar}(\gamma_3)$ to an annulus in $\mathbf{collar}(\gamma_3) \times [0, 1]$. It is clear that $f \prec_{\text{top}} h \prec_{\text{top}} g \prec_{\text{top}} f$.

Ordering disconnected surfaces. We can extend \prec_{top} to maps $f : R \rightarrow M$ where R is a disconnected subsurface of S , with a bit of care. We call R essential if all of its components are essential, and no two are isotopic. If R and T are essential subsurfaces and $f : R \rightarrow M, g : T \rightarrow M$ are in the homotopy class of s_0 , we say that $f \prec_{\text{top}} g$ provided that, for any pair R', T' of intersecting components of R and T , we have $f|_{R'} \prec_{\text{top}} g|_{T'}$. (A similar definition can be made for unions of annuli and the corresponding maps of solid tori).

It is easy to see, using Lemma 3.1, that

Lemma 3.3. *if $R' \subset R$ and $T' \subset T$, and R' and T' have some overlapping components, then $R \prec_{\text{top}} T$ implies $R' \prec_{\text{top}} T'$.*

3.1.1. Embeddings and Scaffolds. Let $\mathbf{emb}(R, M)$ be the set of images of those maps in $\mathbf{map}(R, M)$ that are embeddings. When R is a (closed) annulus we also require the map of the solid torus $R \times J$ to be proper and orientation-preserving. Define $\mathbf{emb}(M) = \cup_R \mathbf{emb}(R, M)$. Note that an embedding in $\mathbf{map}(M)$ is determined by its image, up to reparametrization of the domain of the map by a map homotopic to the identity, and it follows that we can extend \prec_{top} to a well-defined relation on $\mathbf{emb}(M)$.

A solid torus in $\mathbf{emb}(R, M)$ is *straight* if it is of the form $\mathbf{collar}(v) \times J$ for some v , where J is a closed connected subset of \mathbb{R} . If v is nonperipheral in S we allow J to be of the form $[a, b], [a, \infty)$ or $(-\infty, b]$. If v is a component of ∂S then we require $J = \mathbb{R}$.

We are now ready to define scaffolds, which are the primary object of study in this section.

Definition 3.4. *A scaffold $\Sigma \subset \widehat{M}$ is a union of two sets \mathcal{F}_Σ and \mathcal{V}_Σ , where*

- (1) \mathcal{F}_Σ is a finite disjoint union of elements of $\mathbf{emb}(M)$, of non-annulus type.
- (2) \mathcal{V}_Σ is a finite disjoint union of elements in $\mathbf{emb}(M)$ of annulus type (that is, solid tori).

- (3) \mathcal{V}_Σ is unknotted and unlinked: it is isotopic in M to a union of straight solid tori.
- (4) \mathcal{F}_Σ only meets \mathcal{V}_Σ along boundary curves of surfaces in \mathcal{F}_Σ , and conversely for every component F of \mathcal{F}_Σ , ∂F is embedded in $\partial\mathcal{V}_\Sigma$.
- (5) No two elements of \mathcal{V}_Σ are homotopic.

The components of \mathcal{F}_Σ and \mathcal{V}_Σ are called the pieces of Σ .

In a straight solid torus V let the *level homotopy class* denote the homotopy class in ∂V of the curves of the form $\gamma \times \{t\}$ for some $t \in J$. If V is isotopic to a straight solid torus we define the level homotopy class as the one containing the isotopes of the level curves. The following lemma guarantees that all the curves in $\mathcal{F} \cap \mathcal{V}$ are in the level homotopy class.

Lemma 3.5. *Let Σ be a scaffold in \widehat{M} . The intersection curves $\mathcal{F}_\Sigma \cap \mathcal{V}_\Sigma$ are in the level homotopy classes of the components of $\partial\mathcal{V}_\Sigma$.*

Proof. By property (3) of the definition, and the isotopy extension theorem [56], we may assume that \mathcal{V}_Σ is a union of straight solid tori. Let γ be an intersection curve of a component F of \mathcal{F}_Σ with a component $V = \overline{\text{collar}(v)} \times J$ of \mathcal{V}_Σ .

If v is peripheral in M then $J = \mathbb{R}$, ∂V is an annulus, and there is a unique nontrivial homotopy class of simple curves in ∂V , the level homotopy class. Hence we are done.

Thus we may assume v is nonperipheral, $J \neq \mathbb{R}$, and without loss of generality that $b = \sup J < \infty$. Let $\gamma_b \subset \partial V$ be the level curve $\gamma_v \times \{b\}$ where γ_v is the standard representative of v in S . Since γ is homotopic to γ_b in V , to show that they are homotopic in ∂V it suffices to show that the algebraic intersection number $\gamma \cdot \gamma_b$ vanishes.

Since F is homotopic to a subsurface of S , ∂F meets ∂V in either one or two components. Consider first the case that ∂F meets ∂V only in the single component γ .

Compactify M to get $\overline{M} = S \times [-\infty, \infty]$. Let \mathcal{V}_F denote the set of components of \mathcal{V}_Σ meeting boundary components of F , and let $X = \overline{M} \setminus \text{int}(\mathcal{V}_F)$. Let $B \subset \overline{M}$ be the annulus $\gamma_v \times [b, \infty]$. Since the components of ∂F do not overlap, the solid tori of \mathcal{V}_F have disjoint projections to S and it follows that B is contained in X . Let $A = \partial V \cup S \times \{\infty\}$. B determines a class $[B] \in H_2(X, A)$, and intersection number with B gives a cohomology class (its Lefschetz dual) $i_B \in H^1(X, \partial X - A)$.

If α is a closed curve in ∂V , then $i_B(\alpha)$ is the algebraic intersection number of α and ∂B on ∂V . In particular, since $\partial B \cap \partial V = \gamma_b$, we have

$$\gamma \cdot \gamma_b = i_B(\gamma).$$

Now since $\partial F \cap A = \gamma$, we find that $[\gamma]$ vanishes in $H_1(X, \partial X - A)$, and it follows that $i_B(\gamma) = 0$. This concludes the proof in this case.

If $\partial F \cap \partial V$ has two components, there is a double cover of M to which F has two lifts, each of which has no pair of homotopic boundary components.

Picking one of these lifts and a lift of V , we repeat the above argument in this cover. \square

A scaffold Σ is *straight* if every component of \mathcal{F}_Σ is in fact a level surface $R \times \{t\}$, and every solid torus in \mathcal{V}_Σ is straight.

Let $\prec_{\text{top}}|_\Sigma$ denote the restriction of the \prec_{top} relation to the pieces of Σ . We can capture the essential properties of being a straight scaffold with this definition (see Theorem 3.10).

Definition 3.6. *A scaffold Σ is combinatorially straight provided $\prec_{\text{top}}|_\Sigma$ satisfies these conditions:*

- (1) (*Overlap condition*) *Whenever two pieces p and q of Σ overlap, either $p \prec_{\text{top}} q$ or $q \prec_{\text{top}} p$,*
- (2) (*Acyclic condition*) *The transitive closure of $\prec_{\text{top}}|_\Sigma$ is a partial order.*

Notice that one may use a construction similar to the one in Example 3.2 to construct scaffolds which satisfy the overlap condition but not the acyclic condition.

3.2. Scaffold extensions and isotopies

Our technology will allow us to study “good” maps of scaffolds into M , those which, among other things, have a scaffold as image and preserve the topological ordering of pieces.

Definition 3.7. *A map $f : \Sigma \rightarrow \widehat{M}$ is a good scaffold map if the following holds:*

- (1) *f is homotopic to the identity.*
- (2) *$f(\Sigma)$ is a scaffold Σ' with $\mathcal{V}_{\Sigma'} = f(\mathcal{V}_\Sigma)$, and $\mathcal{F}_{\Sigma'} = f(\mathcal{F}_\Sigma)$*
- (3) *For each component V of \mathcal{V}_Σ , $f(V)$ is a component of $\mathcal{V}_{\Sigma'}$, and $f|_V : V \rightarrow f(V)$ is proper.*
- (4) *f is an embedding on \mathcal{F}_Σ .*
- (5) *f is order-preserving. That is, for any two pieces p and q of Σ which overlap,*

$$f(p) \prec_{\text{top}} f(q) \iff p \prec_{\text{top}} q.$$

The main theorem of this section gives that a well-behaved map F of \widehat{M} to itself which restricts to a good scaffold map of a combinatorially straight scaffold Σ is homotopic to a homeomorphism which agrees with F on Σ . In particular this implies that the complements of Σ and $F(\Sigma)$ are homeomorphic.

Theorem 3.8. (Scaffold Extension) *Let $\Sigma \subset \widehat{M}$ be a combinatorially straight scaffold, and let $F : \widehat{M} \rightarrow \widehat{M}$ be a proper degree 1 map homotopic to the identity, such that $F|_\Sigma$ is a good scaffold map, and $F(M \setminus \text{int}(\mathcal{V}_\Sigma)) \subset M \setminus \text{int}(F(\mathcal{V}_\Sigma))$.*

Then there exists a homeomorphism $F' : \widehat{M} \rightarrow \widehat{M}$, homotopic to F , such that

- (1) $F'|_{\mathcal{F}_\Sigma} = F|_{\mathcal{F}_\Sigma}$
- (2) On each component V of \mathcal{V}_Σ , $F'|_V$ is homotopic to $F|_V$ rel \mathcal{F}_Σ , through proper maps $V \rightarrow F(V)$.

We will derive the Scaffold Extension theorem from the Scaffold Isotopy theorem, which essentially states that the image of a good scaffold map can be ambiently isotoped back to the original scaffold. The proof of the Scaffold Isotopy theorem will be deferred to section 3.6.

Theorem 3.9. (Scaffold Isotopy) *Let Σ be a straight scaffold, and let $f : \Sigma \rightarrow \widehat{M}$ be a good scaffold map. There exists an isotopy $H : \widehat{M} \times [0, 1] \rightarrow \widehat{M}$ such that $H_0 = id$, $H_1 \circ f(\Sigma) = \Sigma$, and $H_1 \circ f$ is the identity on \mathcal{F}_Σ .*

3.3. Straightening

Now assuming Theorem 3.9, let us prove the following corollary, which allows us to treat combinatorially straight scaffolds as if they were straight.

Lemma 3.10. *A scaffold is combinatorially straight if and only if it is ambient isotopic to a straight scaffold.*

Proof. If Σ is straight, then by Lemma 3.1 two disjoint pieces are ordered whenever they overlap. Lemma 3.1 also implies that \prec_{top} is determined by the ordering of the \mathbb{R} coordinates, and must therefore be acyclic. Hence Σ is combinatorially straight. The property of being combinatorially straight is preserved by isotopy and hence holds for any scaffold isotopic to a straight scaffold.

Now suppose that Σ is combinatorially straight. Let \mathcal{P} denote the set of pieces of Σ . We will first construct a straight scaffold Σ_0 together with a bijective correspondence $c : \mathcal{P} \rightarrow \mathcal{P}_0$ taking pieces of Σ to pieces of Σ_0 , such that whenever p and q are overlapping pieces of Σ , we have $p \prec_{\text{top}} q \iff c(p) \prec_{\text{top}} c(q)$.

Let \prec'_{top} denote the transitive closure of $\prec_{\text{top}}|_{\Sigma}$, which by hypothesis is a partial order. It is then an easy exercise to show that there is a map $l : \mathcal{P} \rightarrow \mathbb{Z}$ which is order preserving, i.e. $p \prec'_{\text{top}} q \implies l(p) < l(q)$.

Now for each component F of \mathcal{F}_Σ let $c(F)$ be the level embedding $s_{l(F)}(F)$, and let \mathcal{F}_{Σ_0} be the union of these level embeddings. Two components of \mathcal{F}_Σ have the same l value only if they are unordered, and since \prec_{top} satisfies the overlap condition, this implies they have disjoint domains. It follows that these level embeddings are all disjoint.

We next construct the solid tori in \mathcal{V}_{Σ_0} . Let V be a component of \mathcal{V}_Σ , and let v be its homotopy class in S . Recall that V is isotopic to $\overline{\text{collar}}(v) \times J$ where $J \subset \mathbb{R}$ is closed and connected. The solid torus $c(V)$ will be $\overline{\text{collar}}(v) \times J_0$, where $J_0 = [a, b] \cap \mathbb{R}$, with $a = a(V)$ and $b = b(V)$ defined as follows. Let $\beta(V)$ be the set of l values for surfaces bordering V . If $\inf J = -\infty$ let $a = -\infty$, and if $\sup J = \infty$ let $b = \infty$. In all other cases, let

$$a(V) = \min \beta(V) \cup \{l(V)\} - 1/3$$

and

$$b(V) = \max \beta(V) \cup \{l(V)\} + 1/3.$$

The union of $c(V)$ over $V \in \mathcal{V}_\Sigma$ gives \mathcal{V}_{Σ_0} .

Note that this definition implies that, whenever any component F of \mathcal{F}_Σ and V of \mathcal{V}_Σ intersect along a boundary component γ of F , the corresponding $c(F)$ and $c(V)$ intersect along the boundary component γ' of $c(F)$ corresponding to γ .

Once we check that these are the only intersections between the pieces of Σ_0 , it will follow that Σ_0 is a (straight) scaffold. The order-preserving property of the correspondence will follow from the same argument.

First let us establish the following claim about the ordering. Suppose that p is a piece of Σ which overlaps $V \in \mathcal{V}_\Sigma$ and F is a piece of \mathcal{F}_Σ with a boundary component γ in ∂V (hence $l(F) \in \beta(V)$). We claim that

$$V \prec_{\text{top}} p \implies F \prec_{\text{top}} p \tag{3.1}$$

and

$$p \prec_{\text{top}} V \implies p \prec_{\text{top}} F. \tag{3.2}$$

Since V and p overlap, F and p must also overlap and hence are \prec_{top} -ordered because Σ is combinatorially straight. Thus we only have to rule out $V \prec_{\text{top}} p \prec_{\text{top}} F$ and $F \prec_{\text{top}} p \prec_{\text{top}} V$. Assume the former without loss of generality. $V \prec_{\text{top}} p$ implies that V is homotopic to $-\infty$ in the complement of p . Since V is homotopic into F by a homotopy taking V into itself, $p \prec_{\text{top}} F$ implies that V is homotopic to $+\infty$ in the complement of p . This is a contradiction by Lemma 3.1. Together with the corresponding argument for the case $F \prec_{\text{top}} p \prec_{\text{top}} V$, this establishes (3.1) and (3.2).

Now if p is a piece of \mathcal{F}_Σ overlapping $V \in \mathcal{V}_\Sigma$, suppose without loss of generality that $V \prec_{\text{top}} p$. Then $l(V) < l(p)$, and for each F with boundary component on V we have $F \prec_{\text{top}} p$ by (3.1), and hence $l(F) < l(p)$. Furthermore note that it is not possible to have $b(V) = \infty$ in this case because then $V \prec_{\text{top}} p$ could not hold. It follows that $b(V) < l(p)$, and therefore $c(V)$ and $c(p)$ are disjoint and $c(V) \prec_{\text{top}} c(p)$.

If V' is a component of \mathcal{V}_Σ overlapping V , and (without loss of generality) $V \prec_{\text{top}} V'$, a similar argument implies that $b(V) < a(V')$, and again $c(V)$ and $c(V')$ are disjoint, and $c(V) \prec_{\text{top}} c(V')$.

This establishes disjointness for overlapping pieces of Σ_0 . Disjointness for non-overlapping pieces is immediate from the definition of Σ_0 . We have also established one direction of order-preservation, namely $p \prec_{\text{top}} q \implies c(p) \prec_{\text{top}} c(q)$ for overlapping pieces p and q of Σ . For the opposite direction we need just observe that $c(p)$ and $c(q)$ must be \prec_{top} -ordered, so that the opposite direction follows from the forward direction with the roles of p and q reversed.

Next, we construct a good scaffold map $h : \Sigma_0 \rightarrow \Sigma$: On each component $F_0 = c(F)$ of \mathcal{F}_{Σ_0} , by construction, there is a homeomorphism to F , in the homotopy class of the identity. After defining h on \mathcal{F}_{Σ_0} , for each $V_0 = c(V)$ in \mathcal{V}_{Σ_0} h is already defined on the circles $\partial V_0 \cap \mathcal{F}_{\Sigma_0}$, and it is easy to extend this

to a proper map of V_0 to the corresponding solid torus V (although the map is not guaranteed to be an embedding). If V_0 does not meet any component of \mathcal{F}_{Σ_0} then we simply define the map $V_0 \rightarrow V$ to be a homeomorphism that takes level curves of ∂V_0 to the homotopy class of level curves on ∂V . This gives h , which satisfies all the conditions of a good scaffold map: property (5), order preservation, follows from the order-preserving property of the correspondence c , while the other properties are implicit in the construction.

Now apply Theorem 3.9 to the map h , producing an isotopy H with $H_0 = id$ and $H_1 \circ h(\Sigma_0) = \Sigma_0$. Thus $H_1(\Sigma) = \Sigma_0$, and we have exhibited the desired ambient isotopy from Σ to a straight scaffold. \square

3.4. Proof of the scaffold extension theorem

We now give the proof of Theorem 3.8, again assuming Theorem 3.9.

First we note that it suffices to prove the theorem when Σ is straight. For by Lemma 3.10, if Σ is combinatorially straight there is a homeomorphism $\Phi : M \rightarrow M$ isotopic to the identity such that $\Phi(\Sigma)$ is straight. We can apply the result for $\Phi(\Sigma)$ and conjugate the answer by Φ .

Denote $\mathcal{F} = \mathcal{F}_\Sigma$ and $\mathcal{V} = \mathcal{V}_\Sigma$. Apply the Scaffold Isotopy theorem 3.9 to the good scaffold map $F|_\Sigma$, obtaining an isotopy H of M with $H_0 = id$, $H_1 \circ F(\Sigma) = \Sigma$, and $H_1 \circ F$ equal to the identity on \mathcal{F} . Let $F_1 = H_1 \circ F$.

Our desired map will just be $F' = H_1^{-1}$. We have immediately that $F'(\Sigma) = F(\Sigma)$, and $F' = F$ on \mathcal{F} . It remains to show that $F_1|_V$ for any component V of \mathcal{V} is homotopic rel \mathcal{F} through proper maps of V to the identity. Composing this homotopy by H_1^{-1} , we will have that $F'|_V$ is homotopic rel \mathcal{F} through proper maps of V to F , as desired.

Fixing a component V of \mathcal{V} , let us first find a homotopy of $F_1|_{\partial V}$, through maps of $\partial V \rightarrow \partial V$ fixing $\mathcal{F} \cap \partial V$ pointwise, to the identity on ∂V .

We claim that $F_1|_{\partial V}$ preserves the level homotopy class. If ∂V meets \mathcal{F} , this is immediate since F_1 fixes \mathcal{F} and Σ is straight. If V is disjoint from \mathcal{F} , consider a level curve γ in ∂V . Since γ is homotopic to $+\infty$ in the complement of $int(V)$, and since F_1 takes $\widehat{M} \setminus int(V)$ to itself, it follows that $F_1(\gamma)$ is homotopic to $+\infty$ in the complement of $int(V)$. Adding a boundary $\widehat{S} \times \{\infty\}$ to \widehat{M} , we conclude that γ and $F_1(\gamma)$ are homotopic within $\widehat{M} \setminus int(V)$ to curves in this boundary, which are of course homotopic and hence have vanishing algebraic intersection number. Thus the intersection number $F_1(\gamma) \cdot \gamma$ on ∂V vanishes as well, so as in Lemma 3.5 it follows that $F_1(\gamma)$ and γ are homotopic in ∂V .

$F_1|_{\partial V}$ also takes meridians to powers of meridians, since it is a restriction of a self-map of V . Thus it is homotopic to the identity provided $F_1|_{\partial V} : \partial V \rightarrow \partial V$ has degree 1, or equivalently that $F_1|_V : V \rightarrow V$ has degree 1. Since F_1 takes $M \setminus int(\mathcal{V})$ to $M \setminus int(\mathcal{V})$ by the hypotheses of the theorem, and no two components of \mathcal{V} are homotopic, F_1 must take $M \setminus int(V)$ to $M \setminus int(V)$, and it follows that the degree of F_1 on V equals the degree of F_1 , which is 1 by hypothesis.

If ∂V meets no components of \mathcal{F} , this suffices to give the desired homotopy of $F_1|_{\partial V}$ to the identity. In the general case, \mathcal{F} may meet ∂V in one or more level curves, which break up ∂V into annuli. For each such annulus A we must show that $F_1|_A : A \rightarrow \partial V$ is homotopic to the identity rel ∂A . Let $\gamma \subset S$ be the curve in the homotopy class of V , so that $V = \mathbf{collar}(\gamma) \times J$.

There are two obstructions to this, which we call $d(A)$ and $t(A)$. Inducing an orientation on A from ∂V , the 2-chain $A - F_1(A)$ is closed, and determines a homology class in $H_2(\partial V)$. If V is compact then $H_2(\partial V) = \mathbb{Z}$ and we may define $d(A)$ by the equation $[A - F_1(A)] = d(A)[\partial V]$. If V is noncompact then ∂V is an annulus, $H_2(\partial V) = 0$, and we define $d(A) = 0$.

To define $t(A)$, choose an arc α connecting the boundary components of A , and note that $\alpha * F_1(\alpha)$ is a closed curve, which is homotopic in the solid torus to some multiple of γ . Let $t(A)$ be this multiple.

If $d(A) = 0$, then $F_1|_A$ is homotopic in ∂V , rel ∂A , to a homeomorphism from A to itself which is the identity on the boundary. The number $t(A)$ then measures the Dehn-twisting of this homeomorphism, and if $t(A) = 0$ then the homeomorphism is homotopic, rel boundary, to the identity. Thus we must establish that $d(A) = 0$ and $t(A) = 0$.

To prove $t(A) = 0$, recall that F_1 and F are homotopic. Since F in turn is homotopic to the identity, there is a homotopy $G : M \times [0, 1] \rightarrow M$ with $G_0 = id$ and $G_1 = F_1$. Since $F_1|_{\mathcal{F}}$ is the identity, the trajectories $G(x \times [0, 1])$ for $x \in \mathcal{F}$ are closed loops. We claim these loops are homotopically trivial: Let F be the component of \mathcal{F} containing x . Since F is not an annulus, x is contained in two loops ξ and η in F that are not commensurable (say that two elements in $\pi_1(M, x)$ are commensurable if they are powers of a common element.) $G(\xi \times [0, 1])$ is a torus and hence homotopically non essential in M , so $G(x \times [0, 1])$ is commensurable with ξ . Similarly it is commensurable with η . But since ξ and η are not commensurable, $G(x \times [0, 1])$ must be homotopically trivial.

Now place any complete, non-positively curved metric on M for which all the solid tori are convex (e.g. put a fixed hyperbolic metric on S and take the product metric on $S \times \mathbb{R}$) and deform the trajectories of G to their unique geodesic representatives. The result is a new homotopy G' from the identity to F_1 , which is constant on \mathcal{F} . It follows that the arc α , whose endpoints are on ∂A , is homotopic *rel endpoints*, inside V , to $F_1(\alpha)$. Hence the loop $\alpha * F_1(\alpha)$ bounds a disk in V , so that $t(A) = 0$.

Next we argue that $d(A) = 0$, which breaks down into several cases. We may assume that ∂V is a torus.

Suppose that A is of the form $\beta \times [s, t]$ with $[s, t] \subset J$ (where $V = \mathbf{collar}(\gamma) \times J$) and β is a boundary component of $\mathbf{collar}(\gamma)$.

If γ separates S , let R be the component of $S \setminus \mathbf{collar}(\gamma)$ which is adjacent to β , and let $Q = R \times [s, t]$. Let B be the vertical annulus $\gamma \times [\max J, \infty]$ in $\overline{M} = S \times [-\infty, \infty]$. With the natural orientation, the intersection $B \cap F_1(\partial Q)$ (which we may assume transverse) defines a class in $H_1(X)$. This is just the intersection pairing $H_2(X) \times H_2(X, \partial X) \rightarrow H_1(X)$, where $X = \overline{M} \setminus \text{int}(V)$

as in Lemma 3.5. In our case since all the intersection curves are trivial or homotopic to γ , this class is a multiple of $[\gamma]$. Let $i(F_1(\partial Q), B)$ denote this multiple. As F_1 maps $M \setminus \mathcal{V}$ to $M \setminus \mathcal{V}$, it follows that $F_1(Q) \subset X$, so $[F_1(\partial Q)] = 0$ in $H_2(X)$. Therefore, $i(F_1(\partial Q), B) = 0$.

We will show that $i(F_1(\partial Q), B) = \pm d(A)$, which implies that $d(A) = 0$. The components of ∂R other than β are in ∂S , and hence map to $\partial S \times \mathbb{R}$, and miss B . Hence $F_1^{-1}(B) \cap \partial Q$ is contained in the surfaces $\partial_+ Q = R \times \{t\}$, $\partial_- Q = R \times \{s\}$, and A . Now $\partial_+ Q$ and $\partial_- Q$ contain components Y and Z of \mathcal{F} which meet V at the boundary of A . Since F_1 is the identity on the (straight) pieces Y and Z , $F_1(Y)$ and $F_1(Z)$ are disjoint from B . Thus, any curve of $F_1^{-1}(B) \cap \partial_{\pm} Q$ lies in a component of $\partial_{\pm} Q$ which does not contain any curves homotopic to γ , so must be homotopically trivial. We conclude that $i(F_1(\partial Q), B) = i(F_1(A), B)$. But $F_1(A)$ only meets B in its boundary curve $\gamma \times \{q\}$, and the number of essential intersections, counted with signs, is exactly the degree with which $F_1(A)$ covers the complementary annulus $\partial V \setminus A$. Hence this is (up to sign) the degree with which $A - F_1(A)$ covers ∂V , which is $d(A)$.

Next consider the case that γ does not separate S . There is a double cover of S to which $\mathbf{collar}(\gamma)$ lifts to two disjoint copies C_1 and C_2 , which separate S into R_1 and R_2 . In the corresponding double cover of M we have two lifts V_1 and V_2 of V . Letting $Q = R_1 \times [s, t]$, we can repeat the above argument to obtain $d(A_1) = 0$, where A_1 is the lift of A to V_1 . Projecting back to M , we have $d(A) = 0$.

Another type of annulus A is one that contains the bottom annulus $\mathbf{collar}(\gamma) \times \{\min J\}$, and whose boundaries are of the form $\beta \times \{s\}$ and $\beta' \times \{s'\}$, where β' is the other boundary component of $\mathbf{collar}(\gamma)$, and $s, s' \in \text{int}J$ are heights where components of \mathcal{F} meet V on its two sides. Suppose again that $\mathbf{collar}(\gamma)$ separates S into two components R and R' adjacent to β and β' respectively. In this case we let Q be the region of M below $V \cup R \times \{s\} \cup R' \times \{s'\}$. Again we can show that $d(A) = \pm i(F_1(\partial Q), B) = 0$. If γ is non-separating we can again use a double cover. The case where A contains the top annulus $\mathbf{collar}(\gamma) \times \{\max J\}$ is treated similarly.

Finally it is possible that ∂V only meets \mathcal{F} on one side, say on $\beta \times J$, and hence there may be one annulus A which is the closure of the complement of $\bar{A} = \beta \times [s, t]$ for some $[s, t] \subset \text{int}J$ (possibly $s = t$, when ∂V meets \mathcal{F} in a unique circle). In this case, $[\bar{A} - F_1(\bar{A})] = 0$ since \bar{A} is a concatenation of annuli for which we have already proved $d = 0$ (or a single circle viewed as a singular 2-chain which is fixed by F_1). Since $[A - F_1(A)] + [\bar{A} - F_1(\bar{A})] = [\partial V - F_1(\partial V)]$, and we have already shown that F_1 takes ∂V to itself with degree 1, this is 0 and it follows that $d(A) = 0$ as well.

We conclude, then, that $F_1|_A$ can be deformed to the identity rel ∂A , for each annulus A . The resulting homotopy of $F_1|_{\partial V}$ to the identity can be extended to the interior of the solid torus V by a coning argument, yielding a homotopy to the identity through proper maps $V \rightarrow V$, fixing $\mathcal{F} \cap \partial V$.

Thus, we can now let H_1^{-1} be the desired map F' , and this concludes the proof.

3.5. Intersection patterns

Before we prove Theorem 3.9, we will need to consider carefully the ways in which embedded surfaces intersect in M .

Pockets. A pair (Y_1, Y_2) of connected, compact incompressible surfaces in M is a *parallel pair* if $\partial Y_1 = \partial Y_2$, $\text{int}(Y_1) \cap \text{int}(Y_2) = \emptyset$, and there is a homotopy $H : Y_1 \times [0, 1] \rightarrow M$ such that H_0 is the inclusion of Y_1 into M and $H_1 : Y_1 \rightarrow Y_2$ is a homeomorphism that takes each boundary component to itself. (Note that, if no two boundary components of Y_1 are homotopic in M , the last condition on H_1 is automatic).

Lemma 3.11. *A parallel pair (Y_1, Y_2) in M is the boundary of a unique compact region, which is homeomorphic to*

$$Y_1 \times [0, 1] / \sim$$

where $(x, t) \sim (x, t')$ for any $x \in \partial Y_1$ and $t, t' \in [0, 1]$, by a homeomorphism taking $Y_1 \times \{0\}$ to Y_1 and $Y_1 \times \{1\}$ to Y_2 .

This region is called a *pocket*, and the surfaces Y_1 and Y_2 are its *boundary surfaces*. We often denote a pocket by its boundary surfaces; e.g. an *annulus pocket* is a solid torus with annulus boundary surfaces.

Proof. Let the map $H : Y_1 \times [0, 1] \rightarrow M$ be as in the definition of parallel pair. Proposition 5.4 of Waldhausen [72] implies that if H is constant on ∂Y_1 , then the parallel pair bounds a compact region of the desired homeomorphism type. We will adjust H to obtain a homotopy H' which is constant on ∂Y_1 .

The map H_1 takes each component γ of ∂Y_1 to itself. We may assume that $H_1(x) = H_0(x) = x$ for some point $x \in \gamma$, and let t be $[H(x \times [0, 1])]$ in $\pi_1(M, x)$. If H_1 were to reverse orientation on γ , we would obtain a relation of the form $tat^{-1} = a^{-1}$ with $a = [\gamma]$, but this is impossible in the fundamental group of an orientable surface. Thus we must have $tat^{-1} = a$, and since a is primitive in $\pi_1(M)$ and S is not a torus, $t = a^m$ for some m . Hence, after possibly adjusting H_1 by a further twist in the collar of ∂Y_1 , we may assume that H_1 is the identity on γ , and furthermore that $m = 0$. Thus $H|_{\gamma \times [0, 1]}$ is homotopic rel boundary to the map $(x, s) \mapsto x$. A modification on a collar of $\partial Y_1 \times [0, 1]$ yields a homotopy H' which is constant on ∂Y_1 . This establishes existence of the pocket. Uniqueness follows from the non-compactness of M . \square

Given two homotopic embedded surfaces with common boundary, one might hope that the surfaces can be divided into subsurfaces bounding disjoint pockets. One could then use the pockets to construct a controlled homotopy which pushed the surfaces off of one another (except at their common boundary). The following lemma shows that, if the surfaces have

no homotopic boundary components, this is always the case unless there is one of three specific configurations of disk or annulus pockets.

Lemma 3.12. *Let R_1 and R_2 be two homotopic surfaces in $\mathbf{emb}(M)$ intersecting transversely such that $\partial R_1 = \partial R_2$. Suppose also that no two components of ∂R_1 are homotopic in M . Let $C = R_1 \cap R_2$. Then there exists a nonempty collection \mathcal{X} of pockets, such that each $X \in \mathcal{X}$ has boundary surfaces $Y_1 \subset R_1$ and $Y_2 \subset R_2$, so that Y_1 is the closure of a component of $R_1 \setminus C$. Furthermore at least one of the following holds:*

- (1) \mathcal{X} contains a disk pocket.
- (2) \mathcal{X} contains a pair of annulus pockets X and X' in the same, nontrivial, homotopy class, and their interiors are disjoint from each other, and from R_1 and R_2 . Furthermore, X and X' are on opposite sides of R_1 , as determined by its transverse orientation in M .
- (3) \mathcal{X} contains an annulus pocket in the homotopy class of a component of ∂R_1 .
- (4) \mathcal{X} is a decomposition into pockets: every component of $R_1 \setminus C$ is parallel to some component of $R_2 \setminus C$, and the interiors of the resulting pockets are disjoint.

Proof. First, if the intersection locus C has a component that is homotopically trivial, take such a component γ which is innermost on R_1 . Thus γ bounds a disk component Y_1 of $R_1 \setminus C$. On R_2 , γ must also bound a disk Y_2 , although $\text{int}(Y_2)$ may contain components of C . These two disks must bound a disk pocket since M is irreducible, so we have case (1).

From now on we will assume all components of C are homotopically non-trivial.

Let $H : R_1 \times [0, 1] \rightarrow M$ be a homotopy from the inclusion map H_0 of R_1 to a homeomorphism $H_1 : R_1 \rightarrow R_2$. We note that, since no two boundary components of R_1 are homotopic in M , and R_1 is homotopic to an essential subsurface of S , the homotopy class in M of any curve in R_1 determines its homotopy class in R_1 (and similarly for R_2). Thus, for any component β of C , β and $H_1(\beta)$ are homotopic in R_2 .

Parallel internal annuli. Suppose that a homotopy class $[\beta]$ in R_1 that is not peripheral contains at least 3 elements of C . Then we claim that annulus pockets as in conclusion (2) exist.

The union of all annuli in R_1 and R_2 bounded by curves in $[\beta]$ forms a 2-complex in M . There is a regular neighborhood of this complex which is a submanifold K of M all of whose boundaries are tori and which R_i intersects in a properly embedded annulus A_i for $i = 1, 2$. Each component T of ∂K , being a compressible but not homotopically trivial torus, bounds a unique solid torus V_T in M . If T is the boundary component containing ∂A_i then V_T must contain K , for otherwise it would contain $R_i \setminus A_i$, which is impossible (see Figure 1).

FIGURE 1. The solid torus V_T is obtained by crossing this picture with the circle.

A_1 cuts V_T into two solid tori; Let V_T^+ be one of them. A_2 meets V_T^+ in a union of annuli, and since there are at least 3 intersection curves of A_1 and A_2 , these annuli have at least 3 boundary components on A_1 . A_2 meets ∂V_T in only 2 circles, so there is at least one annulus of $A_2 \cap V_T^+$ whose boundary components are both in A_1 . An innermost such annulus in V_T^+ yields the desired pocket in V_T^+ . Repeating for the other component of $V_T \setminus A_1$, we have established conclusion (2).

Peripheral Annuli. Now suppose that a peripheral homotopy class $[\beta]$ in R_1 contains at least 2 elements of C . There is then an annulus Y_1 in $R_1 \setminus C$ whose boundary contains a boundary component of R_1 . Again by our assumption that no two boundary components of R_1 are homotopic in M , the two boundary components of Y_1 must also be homotopic in R_2 . Thus they bound an annulus Y_2 in R_2 . The two annuli bound an annulus pocket by Lemma 3.11, and this gives conclusion (3). (Note that Y_2 is allowed to have interior intersections with R_1).

Pocket decomposition. From now on, we will assume that each non-peripheral homotopy class in R_1 contains at most two elements of C , and each peripheral homotopy class in R_1 contains exactly one element.

The curves of C define partitions of R_1 and R_2 whose components are in one-to-one correspondence by homotopy class. In particular, if Y_1 is the closure of a component of $R_1 \setminus C$, its boundary ∂Y_1 must bound a unique surface Y_2 in R_2 which is homotopic to Y_1 . However we note that $\text{int}(Y_2)$ need *not* a priori be a component of $R_2 \setminus C$, since two elements of C may be homotopic. This is the main technical issue we must deal with now.

Now suppose that no non-annular component of $R_1 \setminus C$ has homotopic boundary components. If Y_1 is an annular component of $R_1 \setminus C$, then there is clearly a homotopy from Y_1 to Y_2 which takes each boundary to

itself. If Y_1 is a non-annular component, the existence of such a homotopy follows from the fact that Y_1 has no homotopic boundary components. In either case, Y_1 and Y_2 form a parallel pair, and by Lemma 3.11 they bound a pocket X_{Y_1} (note that $\text{int}(Y_1)$ and $\text{int}(Y_2)$ are disjoint by choice of Y_1).

Let γ be a component of $\partial Y_1 = \partial Y_2$ which is nonperipheral in R_1 . There are two possible local configurations for X_{Y_1} in a small regular neighborhood of γ , shown in Figure 2. X_{Y_1} meets the neighborhood in a solid torus whose boundary contains annuli of Y_1 and Y_2 adjacent to γ . $R_1 \setminus Y_1$ and $R_2 \setminus Y_2$ meet the neighborhood in two annuli which are either both outside of X_{Y_1} (case (a)) or inside of X_{Y_1} (case (b)). We will rule out case (b).

FIGURE 2. Local configurations for the corner of a pocket X_{Y_1} near a boundary curve γ . X_{Y_1} is shaded.

In case (b), there is a component W of $R_2 \setminus C$ contained in $\text{int}(X_{Y_1})$, meeting its boundary at γ . Thus W is homotopic into Y_1 , but this can only be if W is an annulus adjacent to Y_2 . The second boundary component of W can not be contained in $\text{int}(Y_2)$ or $\text{int}(Y_1)$ by definition, nor can it be in ∂Y_1 since then Y_1 and Y_2 would also be annuli (and there is one annulus in R_2 in each homotopy class). Thus we have a contradiction and case(b) cannot hold.

Now we can show that, in fact, $\text{int}(Y_2)$ cannot meet C . For if a component β of C were contained in $\text{int}(Y_2)$, there would be a component W' of $R_1 \setminus C$ contained in $\text{int}(X_{Y_1})$ with one boundary component in $\text{int}(Y_2)$. As in the previous paragraph W' has to be an annulus, in the homotopy class of some γ in ∂Y_1 . Since there are at most 2 components of C in a homotopy class, the second boundary component of W' must be γ itself, but this is impossible since we have ruled out case (b) of Figure 2.

From the above, we see that if Y_1 and Y'_1 are any two components of $R_1 \setminus C$, then the associated pockets X_{Y_1} and $X_{Y'_1}$ are disjoint except possibly along common boundary curves of Y_1 and Y'_1 ; any other intersection would lead to the case (b) configuration or to components of C in $\text{int}(Y_2)$. Thus we obtain the desired pocket decomposition of conclusion (4).

In general, some non-annular components of $R_1 \setminus C$ may have homotopic boundaries. There is a double cover of S such that each non-annular component of $R_1 \setminus C$ lifts to two homeomorphic components that do not have homotopic boundaries, and the surface in the isotopy class of R_1 has a connected lift. We can repeat the previous arguments in the corresponding double cover \widetilde{M} of M , and obtain a pocket decomposition $\widetilde{\mathcal{X}}$ there. Every pocket of $\widetilde{\mathcal{X}}$ must embed under the double cover, since there can be no new intersection curves. For the same reason, pockets downstairs do not intersect except in the expected way along curves of C . Therefore \mathcal{X} is a pocket decomposition. \square

3.6. Proof of Theorem 3.9

We are now prepared to give the proof of Theorem 3.9, which we restate for the reader's convenience.

Theorem 3.9 (Scaffold Isotopy) *Let Σ be a straight scaffold, and let $f : \Sigma \rightarrow \widehat{M}$ be a good scaffold map. There exists an isotopy $H : \widehat{M} \times [0, 1] \rightarrow \widehat{M}$ such that $H_0 = id$, $H_1 \circ f(\Sigma) = \Sigma$, and $H_1 \circ f$ is the identity on \mathcal{F}_Σ .*

Proof. Let $\mathcal{F} = \mathcal{F}_\Sigma$ and $\mathcal{V} = \mathcal{V}_\Sigma$. We first reduce to the case that $f(\mathcal{V}) = \mathcal{V}$.

By assumption, $\Sigma' = f(\Sigma)$ is a scaffold so there is a map $\Phi : \widehat{M} \rightarrow \widehat{M}$ isotopic to the identity, such that $\Phi(\mathcal{V}_{\Sigma'})$ is a union of tori of the form $\overline{\text{collar}}(v) \times J_v$ with J_v an interval. Since Σ is straight, \mathcal{V} is also a union of tori of the form $\overline{\text{collar}}(v) \times I_v$, where the set of homotopy classes v is the same, so that the only difference is that I_v may be different from J_v .

Now let $h_v : J_v \rightarrow I_v$ be an affine orientation preserving homeomorphism, and let $g_{v,t}(x) = (1-t)x + th_v(x)$. Thus $g_{v,t}$ ($t \in [0, 1]$) is an ‘‘affine slide’’ of J_v to I_v . This allows us to define a family of maps G_t on the tubes of $\Phi(\mathcal{V}_{\Sigma'})$, so that for each homotopy class v the restriction of G_t to $\overline{\text{collar}}(v) \times J_v$ is $G_t(p, x) = (p, g_{v,t}(x))$. This slides $\overline{\text{collar}}(v) \times J_v$ to $\overline{\text{collar}}(v) \times I_v$.

If v and w are disjoint, so are their collars and the G_t -images of the corresponding tubes do not collide.

Whenever v and w overlap, $J_v \cap J_w = \emptyset$, and $I_v \cap I_w = \emptyset$. Supposing $\max J_v < \min J_w$, the order-preserving property of f implies that $\max I_v < \min I_w$. Thus it follows that $g_{v,t}(J_v)$ and $g_{w,t}(J_w)$ are disjoint for any t , and again the images of the corresponding tubes do not collide.

By the isotopy extension theorem [56], this isotopy of $\Phi(\mathcal{V}_{\Sigma'})$ can be extended to an isotopy Ψ_t of \widehat{M} .

Thus, after replacing f with $\Psi_1 \circ \Phi \circ f$, we may from now on assume that $\mathcal{V}_{\Sigma'} = \mathcal{V}$. We will build an isotopy of maps of pairs

$$H : (\widehat{M} \times [0, 1], \mathcal{V} \times [0, 1]) \rightarrow (\widehat{M}, \mathcal{V})$$

such that $H_0 = id$ and $H_1 \circ f$ is the identity on \mathcal{F} . This will be done by induction on the pieces of \mathcal{F} .

In the inductive step, we may assume that on some union of components $\mathcal{E} \subset \mathcal{F}$, f is already equal to the identity. We let R be a component of $\mathcal{F} \setminus \mathcal{E}$, and build an isotopy of pairs $(\widehat{M}, \mathcal{V})$ which fixes pointwise a neighborhood of \mathcal{E} , and moves $f|_R$ to the identity.

Our first step is to apply an isotopy so that ∂R and $f(\partial R)$ are disjoint. Let γ be a boundary component of R , lying in the boundary of a solid torus V in \mathcal{V} , and let γ' be a component of $f(\partial R)$ lying in ∂V .

By Lemma 3.5, since both Σ and $f(\Sigma)$ are scaffolds, both γ and γ' are in the homotopy class of level curves and hence homotopic to each other in ∂V . We claim that they are isotopic, within $\partial V \setminus \mathcal{E}$, to disjoint curves. If they are in different components of $\partial V \setminus \mathcal{E}$ then they are already disjoint. If they are in the same component A , then A is either an annulus or a torus in which γ' and γ are homotopic, so γ' is isotopic within A to a curve disjoint from γ .

This isotopy may be extended to a small neighborhood of ∂V to have support in the complement of \mathcal{E} , so after applying this isotopy to f we may assume that ∂R and $f(\partial R)$ are disjoint.

Now, in order to apply Lemma 3.12, let us enlarge R and $f(R)$ to surfaces R_1 and R_2 , as follows. If V is a component of \mathcal{V} meeting just one boundary component γ of R , let A_1 and A_2 be embedded annuli in V joining γ and $f(\gamma)$, respectively, to a fixed core curve of V , and let A_1 and A_2 be disjoint except at the core. If V meets two components of ∂R , join them with an embedded annulus A_1 , and similarly join the corresponding pair of components of $f(\partial R)$ with an embedded A_2 , so that A_1 and A_2 intersect transversely and minimally – either not at all, or transversely in one core curve. Figure 3 illustrates these possibilities. Repeating for each V , let R_1 be the union of R with all the annuli A_1 , and let R_2 be the union of $f(R)$ with all the annuli A_2 .

Since R is a level surface, we may choose the annuli A_1 to be level, so that R_1 is still a level surface.

FIGURE 3. The three ways in which annuli are added to R and $f(R)$. To obtain the true picture, cross each diagram with S^1 .

Because we have joined homotopic pairs of boundary components, R_1 and R_2 satisfy the conditions of lemma 3.12. Let \mathcal{X} be the collection of pockets described in the lemma. For each of the possible cases we will describe how

to simplify the picture by an isotopy of (M, \mathcal{V}) . The general move, given a pocket X bounded by $Y_1 \subset R_1$ and $Y_2 \subset R_2$, is to apply an isotopy in a neighborhood of X which pushes Y_2 off R_1 using the product structure of the pocket. However we have to be careful to deal correctly with the possible intersections of X with \mathcal{E} and \mathcal{V} . In particular we will maintain inductively the property that the intersection of $R_1 \cup R_2$ with each solid torus is always one of the configurations in Figure 3.

1: Suppose \mathcal{X} contains a disk pocket X . Since X is a ball, no component of \mathcal{V} or \mathcal{E} can be contained in it. \mathcal{E} is disjoint from R_1 and R_2 and hence cannot intersect X at all. Since according to Figure 3 any intersection curve of V with $R_1 \cup R_2$ is homotopically nontrivial, \mathcal{V} cannot intersect ∂X , and hence X , either. Hence R_2 can be isotoped through X , reducing the number of intersections, and the isotopy is the identity on \mathcal{V} and \mathcal{E} . After a finite number of such moves we may assume there are no disk pockets.

2: Suppose \mathcal{X} contains two homotopic non-peripheral nnulus pockets X and X' , with interiors disjoint from each other and from R_1 and R_2 , and on opposite sides of R_1 .

If one of the annulus pockets misses \mathcal{V} then it cannot meet any component of \mathcal{E} , and hence we may isotope across it to reduce the number of intersection curves.

If both of the pockets meet \mathcal{V} , it must be in a single solid torus V in the same homotopy class. By the inductive hypothesis, the intersections with V must be as in Figure 3, so that X and X' each meet V in a solid torus. There are two intersection patterns in V , depicted in figure 4. In case (a), the product structure in X can be made so that $\partial V \cap X$ is vertical, and again we can isotope through X , preserving V , to reduce the number of intersection curves.

In case (b), we have to consider the picture more globally. Since X and X' meet R_1 on opposite sides, the possible configurations are as in figure 5. In each case, X meets V in a region bounded by a subset of R_2 , and X' meets V in a region bounded by a subset of R_1 . Outside V there is a solid torus Z bounded by three annuli, one in ∂V , one in ∂X and in one in $\partial X'$, and $\text{int}(Z)$ is disjoint from R_1 . In the top case, $\text{int}(Z)$ is disjoint from R_2 and we may push the annulus $\partial Z \cap \partial X$ across the rest of Z , thus reducing the number of intersection curves outside \mathcal{V} (although we introduce an intersection in V , as shown).

In the bottom case, we can find an innermost annulus of $R_2 \cap Z$, and push across the resulting pocket. Note that in these cases the isotopy can be done in the complement of \mathcal{E} since no component of \mathcal{E} can be contained in a solid torus.

3: If \mathcal{X} contains a peripheral annulus pocket X , let V denote the component of \mathcal{V} in the same homotopy class. The intersection of ∂V with R_1 and R_2 must be as in the first picture in Figure 3. We can adjust the product structure of the pocket so that ∂V is vertical, and hence we can isotope Y_2

FIGURE 4. The two ways that a solid torus V can intersect two homotopic annulus pockets.

FIGURE 5. The two possible moves in case (b)

off of Y_1 while preserving V (see figure 6). Again, this can be done in the complement of \mathcal{E} .

4: If \mathcal{X} is a pocket decomposition, we first claim that no pocket contains all of a component of \mathcal{V} or \mathcal{E} . Suppose that $X \in \mathcal{X}$ does contain such a component Z . We will obtain a contradiction to the order-preserving properties of the map f .

Z is homotopic into R_1 , and we claim it must also overlap R – the alternative is that Z is homotopic to one of the annuli in $R_1 \setminus R$ – but then it would have to be one of the solid tori that intersect R_1 , contradicting the choice of Z .

FIGURE 6. A peripheral annulus pocket can only intersect V as shown (multiplied by S^1). The isotopy move is shown as well.

Recall that R_1 is a level surface, let Y_1 be the subsurface of R_1 in the boundary of X , and assume without loss of generality that X is adjacent to R_1 from below. Because R_1 is a level surface Z can be pushed to $-\infty$ in the complement of R_1 . Since Z overlaps R and Σ is straight, they are ordered and so $Z \prec_{\text{top}} R$. Since f is a good scaffold map we have $f(Z) \prec_{\text{top}} f(R)$, but $f(Z) = Z$ so $Z \prec_{\text{top}} f(R)$.

There is an isotopy of M supported on a small neighborhood of the pockets *different from* X , which pushes all of them outside of the region above Y_1 , which we can call $Y_1 \times (t, \infty)$. Let Ψ be the end result of this isotopy. We can push Z through the product structure of X to just above Y_1 , and then to $+\infty$ through the region $\Psi^{-1}(Y_1 \times (t, \infty))$, avoiding R_2 and in particular $f(R)$ (figure 7). This contradicts $Z \prec_{\text{top}} f(R)$.

FIGURE 7. If Z is contained in a pocket, it contradicts order preservation.

The only remaining issue is that a pocket X might intersect, but not contain, one of the solid tori V . A priori there are six possible intersection patterns of V with the pockets, as in figure 8.

In case (1), V corresponds to a peripheral homotopy class in R_1 and X meets V in a solid torus with the core of V at its boundary. The product structure of X can be adjusted so that the annulus $\partial V \cap X$ is vertical. (This is similar to the peripheral annulus pocket case.)

In case (2), V meets two pockets X and X' and there is an intersection curve in the core of V . Again $\partial V \cap X$ can be made vertical.

FIGURE 8. The six intersection patterns of a pocket X with a tube V

In case (3), the local pattern is the same as in case (2) but we consider the possibility that both intersections are part of X . This case cannot occur: The orientation of X induced from M induces an orientation on each of the two boundary surfaces Y_1 and Y_2 . However in the local picture in V , each Y_i inherits inconsistent orientation from the two sides, since Y_1 and Y_2 intersect transversely.

In case (4), X meets V in two disjoint solid tori. This case can also be ruled out: Consider an essential annulus A in V with boundaries in the two components of $X \cap V$. Because X is a pocket, and the map of S to M is a homotopy-equivalence, the boundaries of A can also be joined by an annulus A' in X . This produces a torus that intersects the level surface R_1 in exactly one, essential, curve γ . The curve is in the homotopy class of V , which in this case is non-peripheral in R_1 . But this is impossible: Let S_1 be the full level surface containing R_1 . Any intersection of the torus with $S_1 \setminus R_1$ cannot be essential because then it would be homotopic to γ , which is nonperipheral in R_1 . The nonessential intersections can be removed by surgeries, yielding a torus that cuts through S_1 in just one curve – but S_1 separates M , a contradiction.

In case (5), V meets X and X' in disjoint solid tori. Hence V must be homotopic to a boundary component of X and of X' . This gives us a configuration which agrees, in a neighborhood of V , with the top picture in figure 5 – the only difference is that the pockets are not annulus pockets now. (Notice that the bottom picture in figure 5 cannot occur, since we have assumed that C never contains more than 2 homotopic curves.) The same

isotopy move as in figure 5 simplifies the situation by reducing the number of intersection curves outside of \mathcal{V} (and changing the local configuration at V to case (2)).

In case (6), V intersects ∂X in two disjoint annuli. The intersection $V \cap X$ is a solid torus Z , and the product structure of X can be adjusted so that the two annuli of $\partial Z \cap \partial V$ are vertical. (This is essentially because in a product $R \times [0, 1]$ there is only one isotopy class of embedded annulus for each isotopy class in R .)

In conclusion: cases (3) and (4) do not occur. Case (5) can be removed locally, yielding a simpler situation. Hence we can assume that all pockets only intersect \mathcal{V} in the patterns of cases (1), (2) and (6). The product structure of each pocket X can then be adjusted so that all the annuli of the form $\partial V \cap X$ are simultaneously vertical, and then a single push of R_2 through each pocket yields an isotopy of (M, \mathcal{V}) that takes R_2 to R_1 , and $f(R)$ to R , and fixes a neighborhood of \mathcal{E} . A final isotopy within R takes f to the identity. This completes the proof of Theorem 3.9. \square

3.7. Wrapping coefficients

In this section, we develop a criterion to guarantee that a surface is embedded which will be used in the proof of Theorem 6.1. Roughly, we show that if a surface is embedded off of an annulus immersed in the boundary of a tube, and is homotopic to $+\infty$ or $-\infty$ in the complement of that tube, then the surface is embedded.

Lemma 3.13. *Let $R \subseteq S$ be an essential subsurface and V a solid torus in $\mathbf{emb}(M)$ homotopic to a nonperipheral curve γ in R . Let $R' = R \setminus \mathbf{collar}(\gamma)$. Suppose that $h \in \mathbf{map}(R, M)$ is a map such that $h|_{R'}$ is an embedding into $M \setminus \mathit{int}(V)$ with $h^{-1}(\partial V) = \overline{\mathbf{collar}}(\gamma)$, $h(\partial R')$ is unknotted and unlinked, and $h|_{\mathbf{collar}(\gamma)}$ is an immersion of $\mathbf{collar}(\gamma)$ in ∂V .*

If h is homotopic to either $-\infty$ or to $+\infty$ in the complement of V then $h|_{\mathbf{collar}(\gamma)}$ (and hence h itself) is an embedding.

We first define an algebraic measure of wrapping. Suppose that V is a straight solid torus in M , $R \subseteq S$ is an essential subsurface, and the core of V is homotopic to a nonperipheral curve in R . Consider $f \in \mathbf{map}(R, M)$ whose image is disjoint from $\mathit{int}(V)$, and suppose that $f(\partial R)$ is *not linked with V* , which means that $f|_{\partial R}$ is homotopic to both $-\infty$ and $+\infty$ in the complement of V . (Note that it is sufficient to assume it is homotopic to $-\infty$ in the complement of V : since V is nonperipheral in R and straight, $\partial R \times \mathbb{R}$ is disjoint from V , so once $f(\partial R)$ is pushed far enough below it can be pushed above along $\partial R \times \mathbb{R}$).

We then define a “wrapping coefficient” $d_-(f, V) \in \mathbb{Z}$ as follows: Let $r \in \mathbb{R}$ be smaller than the minimal level of V , and let $G : R \times [0, 1] \rightarrow M$ be a homotopy with $G_0 = s_r : R \rightarrow R \times \{r\}$, and $G_1 = f$. By the unlinking assumption on $f(\partial R)$ we may choose G so that $G(\partial R \times [0, 1])$ is disjoint from

$\text{int}(V)$. Then the degree of G over $\text{int}(V)$ is well-defined, and we denote it $\deg(G, V)$.

If G' is another such map we claim $\deg(G, V) = \deg(G', V)$. Viewing G and G' as 3-chains, the difference $G - G'$ has boundary equal to $G(\partial R \times [0, 1]) - G'(\partial R \times [0, 1])$, a union of singular tori. Since $M \setminus V$ is atoroidal and these tori are not homotopic to ∂V , they must bound a union of (singular) solid tori W in the complement of V , and we can write $\partial(G - G' - W) = 0$. Since $H_3(M) = \{0\}$ we know $\deg(G - G' - W, V) = 0$ and since the contribution from W is 0 we must have $\deg(G, V) = \deg(G', V)$. Thus we are justified in defining $d_-(f, V) \equiv \deg(G, V)$. We also drop V from the notation, writing $d_-(f)$, $\deg(G)$ etc, where convenient.

It is clear that $d_-(s_r) = 0$ and, more generally, that if f is deformable to $-\infty$ in the complement of V then $d_-(f) = 0$. Furthermore if H is a homotopy such that $H(\partial R \times [0, 1])$ is disjoint from V , $H_0 = f$, $H_1 = g$, and $f(R)$ and $g(R)$ are disjoint from V , then

$$d_-(g) - d_-(f) = \deg(H). \quad (3.3)$$

We can define $d_+(f)$ similarly as $\deg(G)$, where G is a homotopy such that $G_0 = f$ and $G_1 = s_t$, with t larger than the top of V . Then $d_+(f)$ must be 0 for f to be deformable to $+\infty$ in the complement of V .

Proof of Lemma 3.13. Since $h(\partial R')$ is unknotted and unlinked we may attach solid tori to the boundary components of $h(\partial R)$ so that, together with V and $h(R')$, we have a scaffold Σ . Since Σ has no overlapping pieces, it is combinatorially straight. By Lemma 3.10 we may assume, after re-identifying M with $S \times \mathbb{R}$, that V is a straight solid torus and $h(R')$ is a level surface. Assume without loss of generality that h is homotopic to $-\infty$ in the complement of V .

Let A_- be the annulus in ∂V consisting of the part of the boundary below $h(R')$, and let $h_- : R \rightarrow M$ be an extension of $h|_{R'}$ that maps $\mathbf{collar}(\gamma)$ to A_- and is an embedding.

Clearly h_- is deformable to $-\infty$ in the complement of V , and hence $d_-(h_-) = 0$.

The difference between $h(\mathbf{collar}(\gamma))$ and $h_-(\mathbf{collar}(\gamma))$, considered as 2-chains, gives a cycle $k[\partial V]$ with $k \in \mathbb{Z}$. We immediately have $d_-(h) = d_-(h_-) + k = k$. Since h is assumed to be deformable to $-\infty$ in the complement of V , we conclude that $k = 0$. This implies that, since h and h_- are both immersions on $\mathbf{collar}(\gamma)$ which agree on $\partial \mathbf{collar}(\gamma)$, they must have the same image. It follows that $h|_{\mathbf{collar}(\gamma)}$ is an embedding. \square

3.8. Some ordering lemmas

We now apply the scaffold machinery to obtain some basic properties of the \prec_{top} relation. All these properties will be used in the proof of Lemma 8.4, which allows us to choose a cut system whose images under the, suitably altered, model map are correctly ordered.

The following lemma gives us a transitivity property for \prec_{top} in some special situations.

Lemma 3.14. *Let R_1 and R_2 be disjoint homotopic surfaces in $\mathbf{emb}(M)$. Let \mathcal{V} be an unlinked unknotted collection of solid tori in M with one component for each homotopy class of component of ∂R_1 , so that ∂R_1 and ∂R_2 are embedded in $\partial \mathcal{V}$.*

- (1) R_1 and R_2 are \prec_{top} -ordered.
- (2) Let $Q \in \mathbf{emb}(M)$ be disjoint from $R_1 \cup R_2 \cup \mathcal{V}$, so that the domain of Q is contained in the domain R of R_1 and R_2 .
Suppose that $R_1 \prec_{\text{top}} Q$ and $Q \prec_{\text{top}} R_2$. Then $R_1 \prec_{\text{top}} R_2$.

Proof. Let Σ denote the scaffold with $\mathcal{F}_\Sigma = R_1$ and $\mathcal{V}_\Sigma = \mathcal{V}$. Σ has no overlapping pieces, so it is combinatorially straight. Using the straightening lemma 3.10, after an isotopy we may assume that Σ is straight.

Now the Pocket Lemma 3.11 implies that there is a product region X homeomorphic to $R \times [0, 1]$ whose boundary consists of R_1 , R_2 , and annuli in the tubes of \mathcal{V}_Σ associated to their boundaries. Thus R_2 is isotopic to R_1 by an isotopy keeping the boundaries in \mathcal{V} , so that we may assume that X is actually equal to $R \times [0, 1]$ in $M = S \times \mathbb{R}$, with $R_1 \cup R_2 = R \times \{0, 1\}$. It follows from this that if $R_1 = R \times \{0\}$ then $R_1 \prec_{\text{top}} R_2$, and if $R_1 = R \times \{1\}$ then $R_2 \prec_{\text{top}} R_1$. Hence we have part (1).

It remains to show that, given the hypothesis on Q , $R_1 = R \times \{0\}$, from which $R_1 \prec_{\text{top}} R_2$ and hence (2) follows.

If $Q \subset X$ then we must have $R \times \{0\} \prec_{\text{top}} Q$ and $Q \prec_{\text{top}} R \times \{1\}$, so we are done.

Now let us suppose that Q is in the complement of X , and obtain a contradiction. The condition $R_1 \prec_{\text{top}} Q$ implies that there exists a homotopy $H : Q \times [0, \infty) \rightarrow M$ such that $H(Q \times \{t\})$ goes to $+\infty$ as $t \rightarrow +\infty$, and which avoids R_1 . We claim that H can be chosen to avoid R_2 as well. Let W be a neighborhood of X disjoint from Q . There is a homeomorphism ϕ taking $M \setminus R_1$ to $M \setminus X$, which is the identity outside W . Thus $\phi \circ H$ is the desired homotopy. This contradicts $Q \prec_{\text{top}} R_2$, so again we are done. \square

This lemma tells us that for disjoint overlapping non-homotopic non-annular surfaces with boundary in an unknotted and unlinked collection of solid tori, the \prec_{top} -ordering is determined by their boundaries.

Lemma 3.15. *Let $P, R \in \mathbf{emb}(M)$ be disjoint overlapping non-homotopic non-annular surfaces such that $\partial R \cup \partial P$ is embedded in a collection \mathcal{V} of unknotted, unlinked, homotopically distinct solid tori, so that each component of \mathcal{V} intersects ∂P or ∂R , and $\mathcal{V} \cap (\text{int}(R) \cup \text{int}(P)) = \emptyset$. Suppose that for each component α of ∂R that overlaps P we have $\alpha \prec_{\text{top}} P$, and for each component β of ∂P that overlaps R we have $R \prec_{\text{top}} \beta$. Then $R \prec_{\text{top}} P$.*

Proof. Assume, possibly renaming P and R and reversing directions, that the domain of R is not contained in the domain of P .

Let Σ be the scaffold with $\mathcal{F}_\Sigma = R$ and $\mathcal{V}_\Sigma = \mathcal{V}$. By hypothesis \mathcal{V} is isotopic to a union of straight solid tori, so that $\prec_{\text{top}}|_{\mathcal{V}}$ is acyclic and satisfies the overlap condition. Since the hypotheses also give us that $R \prec_{\text{top}} V$ for each component V of \mathcal{V} that overlaps R we conclude that $\prec_{\text{top}}|_\Sigma$ is still acyclic and satisfies the overlap condition. Hence Σ is combinatorially straight and by Lemma 3.10, after isotopy we may assume that Σ is straight.

In particular we may assume that $R = R' \times \{0\}$. Let $X = R' \times (-\infty, 0]$.

After pushing ∂P in $\partial \mathcal{V}$ (by an isotopy supported in a small neighborhood of \mathcal{V} and leaving Σ invariant) we may assume that ∂P is outside of X .

Now consider $P \cap \partial X = P \cap \partial R' \times (-\infty, 0]$ (we may assume this intersection is transverse). All inessential curves bound disks in both P and ∂X , and so an isotopy of P will remove them. The remaining curves are in the homotopy classes of the components of $\partial R'$. Let A be a component of $P \cap X$. If A is an annulus then it and an annulus in ∂X bound a solid torus (annulus pocket) in X , and taking an innermost such solid torus we may remove it by an isotopy of P without producing new intersection curves with ∂X . After finitely many moves we may assume there are no annular intersections.

Any non-annular A must be a subsurface of P , and hence is the image of a subsurface A' of S . On the other hand, $A \subset X$ implies that A' is homotopic into R' , and $\partial A \subset \partial X$ implies that $\partial A'$ is homotopic into $\partial R'$. Since A' is not an annulus, the only way this can happen is if A' is isotopic to R' (see, for example, Theorem 13.1 in [27].) But this contradicts the assumption that the domain of R is not contained in the domain of P . Thus there is no non-annular component A .

We conclude that, after an isotopy that does not move R , P may be assumed to lie outside of X . Thus P is homotopic to $+\infty$ avoiding X and hence R , and R is homotopic to $-\infty$ (through X) in the complement of P . R and P overlap, so we conclude by Lemma 3.1 that $R \prec_{\text{top}} P$. \square

The next lemma will allow us to check more easily that tubes are ordered with respect to non-annular embedded surfaces with boundary in a collection of straight solid tori.

Lemma 3.16. *Let $R \in \mathbf{emb}(M)$ be non-annular, with boundaries embedded in a collection \mathcal{V} of straight solid tori. Let U be a straight solid torus disjoint from $R \cup \mathcal{V}$ and overlapping R . If R is homotopic to $-\infty$ in the complement of U then $R \prec_{\text{top}} U$, and if R is homotopic to $+\infty$ in the complement of U then $U \prec_{\text{top}} R$.*

Proof. Assume without loss of generality that R is homotopic to $-\infty$ in the complement of U . Let $B = \gamma \times [t, \infty)$ where $\gamma \times \{t\}$ is embedded in ∂U and homotopic to the core of U . Since U and \mathcal{V} are straight and \mathcal{V} is homotopic to $-\infty$ in the complement of U (since R is), B must be disjoint from \mathcal{V} . We may assume that B intersects R transversely, so that components of $B \cap R$ are either homotopically trivial or homotopic to the core of U .

Homotopically trivial components may be removed by isotopy of B . The nontrivial components are signed via the natural orientation of B and R ,

and the fact that R is homotopic to $-\infty$ in the complement of U means that the signs sum up to 0. Two components of opposite signs that are adjacent on B can be removed by an isotopy of B , and we conclude that B can be isotoped away from R . Thus U is homotopic to $+\infty$ in the complement of R , and we conclude (invoking Lemma 3.1) that $R \prec_{\text{top}} U$. \square

4. Cut systems and partial orders

In this section, we link combinatorial information from the hierarchy H to topological ordering information of *split-level surfaces* in M_ν . A *split-level surface* is an embedded surface in the model manifold associated to a slice of H that is made up of level subsurfaces arising as the upper and lower boundaries of blocks in the model. As these split-level surfaces and their images in M will play an important role in what follows, we now develop some control over them and their interactions within the model, aiming in particular for a consistency result (Proposition 4.14) comparing topological ordering in M_ν and a more combinatorial ordering we define on corresponding slices in H .

This consistency result will not apply generally to all slices in H and their associated domains, but after a thinning procedure we arrive at a collection of slices called a *cut system* whose split-level surfaces are well behaved with respect to the topological partial ordering and divide the model into regions of controlled size as we will see in §5.

4.1. Split-level surfaces associated to slices

If a is a slice of H , we recall from Section 2.2 that g_a denotes its bottom geodesic and v_a is the bottom simplex of a . Then $p_a = (g_a, v_a)$ is the bottom pair of a . Let

$$D(a) = D(p_a) = D(g_a)$$

be the domain of g_a . If $D(a)$ is not an annulus, let

$$\check{D}(a) = D(a) \setminus \text{collar}(\text{base}(a))$$

be the complement in $D(a)$ of the standard annular neighborhoods of the curves in $\text{base}(a)$. When a is a saturated slice, the subsurface $\check{D}(a) \subset D(a)$ is a collection of pairwise disjoint 3-holed spheres. If $D(a)$ is an annulus, we let $\check{D}(a) = D(a)$.

Each slice in H gives rise to a properly embedded surface in $M_\nu[0]$ called *split-level surface*. Given a non-annular slice a of H , each 3-holed sphere $Y \subset \check{D}(a)$ admits a natural level embedding $F_Y \subset M_\nu[0]$. This embedded copy F_Y of Y lies in the top boundary and the bottom boundary of the two blocks that are glued along F_Y . The *split-level surface* F_a associated to a is obtained by letting

$$F_a = \bigcup_{Y \subset \check{D}(a)} F_Y.$$

Given a slice a and $v \in \text{base}(a)$, we say γ_v is a *hyperbolic base curve* for a if there is a solid torus $U(v)$ in $M_\nu[\infty]$ whose closure is compact; otherwise we say γ_v is a *parabolic base curve*. For each $v \in a$ with γ_v a hyperbolic base curve, we extend the above embedding of $\check{D}(a)$ across the annulus $\mathbf{collar}(v)$ to a map of $\mathbf{collar}(v)$ into $U(v)$: the core γ_v is sent to the core of the tube $U(v)$ with its model hyperbolic metric, and the pair of annuli $\mathbf{collar}(v) \setminus \gamma_v$ are mapped in such a way that radial lines in $\mathbf{collar}(v)$ map to radial geodesics in the tube $U(v)$. Given $v \in \text{base}(a)$ for which γ_v is a parabolic base curve, we extend across the corresponding annulus $\mathbf{collar}(\gamma_v)$ to any embedding of $\mathbf{collar}(\gamma_v)$ into $U(v)$.

We remark that given a slice a the only base curves that fail to be hyperbolic correspond to vertices $a \in \text{base}(a)$ for which v is a vertex in $\text{base}(I(H))$ without a transversal or v is a vertex in $\text{base}(T(H))$ without a transversal.

Extending over each annulus $\mathbf{collar}(v)$ for $v \in \text{base}(a)$ in this way, we obtain an embedding of $D(a)$ into M_ν whose image we denote by \widehat{F}_a . For each integer $k \in [0, \infty]$ we denote by $\widehat{F}_a[k]$ the intersection

$$\widehat{F}_a[k] = \widehat{F}_a \cap M_\nu[k].$$

We call the surfaces $\widehat{F}_a[k]$ *extended split-level surfaces*.

When a is an annular slice, there is a vertex v so that $D(a) = \overline{\mathbf{collar}(v)}$. We refer to this vertex v as the *core vertex* of a , and denote it by $v = \text{core}(a)$. Then we have the associated solid torus $U(v) \subset M_\nu$. In the interest of comparing all slices in C and their associated topological objects in M_ν , we adopt the convention that for each integer $k \in [0, \infty]$ we have

$$F_a[k] = \widehat{F}_a[k] = \widehat{F}_a = U(v).$$

4.2. Resolution sweeps of the model

A resolution $\{\tau_n\}$ of the hierarchy H_ν yields a “sweep” of the model manifold by split-level surfaces, which is monotonic with respect to the \prec_{top} relation. More specifically, in §8.2 of [47], the embedding of $M_\nu[0]$ in $S \times \mathbb{R}$ is constructed inductively using an exhaustion of $M_\nu[0]$ by submanifolds M_i^j that are unions of blocks. Each F_{τ_j} appears as the “top” boundary of M_i^j , so that $\text{int}(M_i^j)$ lies below the cross-sectional surface \widehat{F}_{τ_j} in the product structure of $S \times \mathbb{R}$. When $\tau_j \rightarrow \tau_{j+1}$ is a move associated to a 4-edge e_j , there is a block $B_j = B(e_j)$ which is appended above \widehat{F}_{τ_j} , and $F_{\tau_{j+1}}$ is obtained from F_{τ_j} by replacing the bottom boundary of B_j by its top boundary. We say that B_j “is appended at time j ” in the resolution. (For other types of elementary moves the surfaces \widehat{F}_{τ_j} and $\widehat{F}_{\tau_{j+1}}$ are the same.)

The following statement about \prec_{top} is an immediate consequence of this construction, and Lemma 3.1.

Lemma 4.1. *Fix a resolution $\{\tau_n\}$ of H . If $i < j$ and $W \subset \widehat{F}_{\tau_i}$ and $W' \subset \widehat{F}_{\tau_j}$ are essential subsurfaces which overlap and are disjoint, then $W \prec_{\text{top}} W'$.*

Similarly, if B is appended at time $j > i$ and W_B is the middle surface of B then $\widehat{F}_{\tau_i} \prec_{\text{top}} W_B$. If B is appended at time $j < i$ then $W_B \prec_{\text{top}} \widehat{F}_{\tau_i}$

4.3. Cut systems

Given a collection C of slices of H , we let

$$C|_h \equiv \{\tau \in C : g_\tau = h\}$$

denote the slices in C with bottom geodesic h . Let $d_1 < d_2$ be *a priori* positive integers, allowing the possibility that $d_2 = \infty$. Then the collection C is a *cut system* if it satisfies the following properties:

- (1) *Distribution of bottom pairs:* For each $h \in H$ with $\xi(D(h)) \geq 4$,
 - $C|_h$ is *well-distributed*: The set $\{v_\tau : \tau \in C|_h\}$ of bottom vertices on h cuts h into intervals of length at most d_2 , and, if non-empty, cuts h into at least three intervals of size at least d_1 . Furthermore, no two slices have the same bottom pair, and no v_τ is the first or last simplex of h .
 - $C|_h$ satisfies *deep placement*: d_1 is at least 5.
- (2) *Initial pairs:* For every pair $(h, w) \in \tau \in C$ that is not a bottom pair, w is the first simplex of h .
- (3) *Saturation:* Each slice $\tau \in C$ with non-annular bottom geodesic is a saturated non-annular slice.
- (4) *Annular cut slices:* For any annular geodesic g there is at most one slice $\tau \in C$ with $g_\tau = g$.

For a given pair of constants d_1 and d_2 , we say the cut system C satisfying (1) satisfies a (d_1, d_2) -*spacing condition*. (Note that the spacing condition puts no restriction on annular slices in C).

Lemma 4.2. *Given positive integers $d_1 \geq 5$ and $d_2 \geq 3d_1$ there is a cut system C satisfying a (d_1, d_2) -spacing condition.*

Proof. Given a geodesic $g \in H$ with non-annular domain, so that $|g| > d_2$ we may choose a non-empty collection of pairs along g satisfying condition (1). As in section 2.2, for each such pair (g, u) there is a choice of a saturated non-annular slice τ of H with (g, u) as its bottom pair. The choice of τ proceeds inductively by choosing pairs from geodesics in H whose domains arise as component domains of $(D(k), v)$ where (k, v) is a pair in τ . By choosing the initial pair at each stage we satisfy conditions (2) and (3).

In general any choice of a collection of slices on annular geodesics will satisfy condition (4) provided there is at most one slice for each annular geodesic, so we may make any such choice to conclude the proof of the lemma. \square

Next we prove the following lemma, which expands on the consequences of the deep placement condition in part 1 of the definition of a cut system.

Lemma 4.3. *Let H be a hierarchy, and let a be a slice of H with bottom pair $p_a = (g_a, v_a)$. If v_a has distance at least 3 from $\mathbf{I}(g_a)$ and $\mathbf{T}(g_a)$ along*

g_a then for any pair $p = (h, v) \in a$ we have $g_a \not\prec h$, $h \searrow g_a$, and

$$(g_a, \mathbf{I}(g_a)) \prec_p p \prec_p (g_a, \mathbf{T}(g_a)).$$

If a is an annular slice, then for any pair $p = (g_a, v)$ we have

$$(g_a, \mathbf{I}(g_a)) \preceq_p p \preceq_p (g_a, \mathbf{T}(g_a)).$$

Proof. Given a pair $p = (h, v)$ with $p \in a$, the footprint $\phi_{g_a}(D(h))$ has diameter at most 2 and contains v_a . If v_a lies distance at least 3 from $\mathbf{I}(g_a)$ and $\mathbf{T}(g_a)$, we have

$$\max \phi_{g_a}(D(h)) < \mathbf{T}(g_a) \quad \text{and} \quad \mathbf{I}(g_a) < \min \phi_{g_a}(D(h)).$$

It follows that $\mathbf{I}(g_a)|_{D(h)} \neq \emptyset$ and applying Theorem 2.1 we conclude that $g_a \not\prec h$. Likewise, since $\mathbf{T}(g_a)|_{D(h)} \neq \emptyset$ we have $h \searrow g_a$. We conclude that

$$(g_a, \mathbf{I}(g_a)) \prec_p p \prec_p (g_a, \mathbf{T}(g_a))$$

by the definition of \prec_p in §2.2.

If a is annular then \prec_p -order is just linear order on the pairs with bottom geodesic g_a so the second statement follows immediately. \square

4.4. Hierarchy partial order and split level-surfaces

Given a cut system C we have the topological ordering relation \prec_{top} on its associated extended split-level surfaces \widehat{F}_a , $a \in C$. Because the surfaces \widehat{F}_a are not themselves level surfaces, however, it does not immediately follow from the preceding section that the transitive closure of \prec_{top} on these split-level surfaces is a partial order: it could conceivably have cycles. We devote the remainder of this section to establishing that this is not the case.

To show there are no cycles, we employ the ordering properties inherent in the hierarchy H to construct an order relation on the slices in a cut system whose transitive closure is a partial order. We prove that this *cut ordering* is consistent with topological ordering of overlapping associated split level surfaces (Lemma 4.14) from which it follows directly that the transitive closure of \prec_{top} on the split-level surfaces associated to C is a partial order.

Let C be a cut system. We define some relations on the slices of C as follows. Given slices a and b in C ,

- (1) $a \vee b$ means that for all $p \in a$ and all $p' \in b$ we have

$$p \prec_p p'.$$

- (2) $a \dashv b$ means that $D(a) \subset D(b)$ and for some $p' \in b$ and all $p \in a$ we have

$$p \prec_p p'.$$

- (3) $a \vdash b$ means that $D(a) \supset D(b)$ and for some $p \in a$ and all $p' \in b$ we have

$$p \prec_p p'.$$

- (4) $a \mid b$ means that there exists a third slice x (called a comparison slice) such that

$$a \dashv x \vdash b.$$

Remark: These possibilities need not be mutually exclusive.

We define the ordering relation

$$a \prec_c b$$

to be the transitive closure of the relation generated by \vee, \dashv , and \vdash . Our main lemma is the following.

Lemma 4.4. *The order relation \prec_c defines a (strict) partial order. Furthermore, we have $a \prec_c b$ if and only if at least one of the following holds*

- $a \vee b$,
- $a \dashv b$,
- $a \vdash b$, or
- $a \mid b$.

Proof. We begin by proving a consistency lemma, which ensures that slices in C are usually comparable via these relations.

Lemma 4.5. *For any two distinct slices a and b in a cut system C , the following possibilities hold:*

- (1) *The geodesics g_a and g_b are the same. Then we have $a \vee b$ or $b \vee a$, and the \vee -ordering is consistent with the order of bottom simplices v_a and v_b along $g_a = g_b$.*
- (2) *If $D(a) \neq D(b)$, $D(a)$ and $D(b)$ overlap, and neither is strictly contained in the other, then $a \vee b$ or $b \vee a$.
More generally if $g_a \prec_t g_b$, even without intersection of the domains, then $a \vee b$.*
- (3) *If $D(a) \subset D(b)$ and $D(a) \not\subseteq \overline{\text{collar}}(\text{base}(b))$ then we have $a \dashv b$ or $b \vdash a$. Furthermore if $a \dashv b$ or $a \vdash b$ then for no pairs $p \in a$ and $p' \in b$ do we have $p' \prec_p p$.*
- (4) *If $\check{D}(a)$ and $\check{D}(b)$ do not overlap, and g_a and g_b are not \prec_t -ordered, then for all $p \in a$ and $p' \in b$, p and p' are not \prec_p -ordered.*

Proof. Proof of Part (1). If $g_a = g_b$ then $D(a) = D(b)$. Moreover, $D(a) = D(b)$ is not an annulus, since C contains at most a single slice on any annular geodesic and a and b are assumed distinct. Since non-annular slices of a cut system satisfy the (d_1, d_2) spacing condition with $d_1 \geq 5$, the bottom simplices v_a and v_b have distance at least 5 on the geodesic $g_a = g_b$. Assume that $v_a < v_b$.

Given any pairs $p \in a$ and $p' \in b$, we wish to show that $p \prec_p p'$. Since $\text{diam}(\hat{\phi}_{g_a}(p)) \leq 2$ and $\text{diam}(\hat{\phi}_{g_a}(p')) \leq 2$, and since $v_a \in \hat{\phi}_{g_a}(p)$ and $v_b \in \hat{\phi}_{g_a}(p')$, we have $\max \hat{\phi}_{g_a}(p) < \min \hat{\phi}_{g_a}(p')$. It follows that $p \prec_p p'$, and we conclude that $a \vee b$.

Proof of Part (2). Assume that $D(a) \neq D(b)$. If $g_a \prec_t g_b$, then by definition there is a geodesic $m \in H$ so that $g_a \searrow m \swarrow g_b$ and $\max \phi_m(D(g_a)) < \min \phi_m(D(g_b))$. In particular, we have $(g_a, \mathbf{T}(g_a)) \prec_p (g_b, \mathbf{I}(g_b))$. Let $p \in a$ and $p' \in b$ be pairs in the slices a and b . Applying Lemma 4.3 we have

$$p \preceq_p (g_a, \mathbf{T}(g_a)) \prec_p (g_b, \mathbf{I}(g_b)) \preceq_p p'.$$

By transitivity of \prec_p we conclude that $p \prec_p p'$.

If $D(a) \cap D(b) \neq \emptyset$ and neither domain is strictly contained in the other, then g_a and g_b are \prec_t -ordered by Lemma 2.2. It follows that either $a \vee b$ or $b \vee a$, and if $g_a \prec_t g_b$ then $a \vee b$.

Proof of Part (3). Suppose $D(a) \subset D(b)$; note that in particular this guarantees that b is non-annular. Let (h, v) be a pair in b with $D(a) \subset D(h)$ (the bottom pair p_b has this property). Then either v intersects $D(a)$, or $D(a)$ is contained in one of the component domains of $(D(h), v)$. Since b is a saturated non-annular slice, either this component domain supports a pair $(h', v') \in b$ or the component domain is an annulus in $\mathbf{collar}(\text{base}(b))$ and we have

$$D(a) \subseteq \overline{\mathbf{collar}}(\text{base}(b)).$$

Thus, provided $D(a) \not\subseteq \overline{\mathbf{collar}}(\text{base}(b))$, we may begin with p_b and proceed inductively to arrive at a unique $(h, v) \in b$ such that $D(a) \subseteq D(h)$ and $v \notin \phi_h(D(a))$. Since $\phi_h(D(a))$ is non-empty by Lemma 2.3 we may assume without loss of generality that $\max \phi_h(D(a)) < v$ which guarantees that $(g_a, \mathbf{T}(g_a)) \prec_p (h, v)$.

Applying Lemma 4.3, for any pair $p \in a$ we have $p \preceq_p (g_a, \mathbf{T}(g_a))$ so we may conclude that

$$p \prec_p (h, v)$$

for all $p \in a$. Now if there were some $p' \in b$ and $p \in a$ such that $p' \prec_p p$, then $p' \prec_p (h, v)$, contradicting Lemma 2.4, which guarantees the non-orderability of pairs in a single slice. This proves the second paragraph of Part (3).

Proof of Part (4). Assume g_a and g_b are not \prec_t -ordered, and that $\check{D}(a)$ and $\check{D}(b)$ do not overlap. This implies that either

- (1) $D(a)$ and $D(b)$ are disjoint domains,
- (2) $D(a)$ is an annulus with $D(a) \subseteq \overline{\mathbf{collar}}(\text{base}(b))$, or
- (3) $D(b)$ is an annulus with $D(b) \subseteq \overline{\mathbf{collar}}(\text{base}(a))$.

Suppose that $p \prec_p p'$ for some $p = (h, v) \in a$ and $p' = (h', v') \in b$. Then there is a comparison geodesic m , with $h \geq m \leq h'$ and $\max \hat{\phi}_m(p) < \min \hat{\phi}_m(p')$. Assume first that $D(a)$ and $D(b)$ do not overlap. We note that $D(m)$ contains both $D(h)$ and $D(h')$, while $D(h) \subseteq D(a)$ and $D(h') \subseteq D(b)$. Thus, disjointness of $D(a)$ and $D(b)$ implies that $D(m)$ is not contained in either $D(a)$ or $D(b)$. The proof of Lemma 4.3 guarantees $h \geq g_a$ and $g_b \leq h'$. It follows that g_a and m lie in the forward sequence $\Sigma^+(D(h))$ while g_b and m lie in the backward sequence $\Sigma^-(D(h'))$ (see Theorem 2.1). Since the domains in $\Sigma^+(D(h))$ are nested and likewise for $\Sigma^-(D(h'))$, we conclude

that $D(m)$ is either equal to, contained in, or contains each of $D(a)$ and $D(b)$.

It follows that $D(a) \subset D(m)$ and $D(b) \subset D(m)$. Since $\phi_m(D(a)) \subseteq \phi_m(D(h))$ and $\phi_m(D(b)) \subseteq \phi_m(D(h'))$, then, it follows that

$$\max \phi_m(D(a)) < \min \phi_m(D(b))$$

and hence $g_a \prec_t g_b$, contradicting our assumption.

Without loss of generality, the final case is that $D(a)$ is an annulus with $D(a) \subseteq \overline{\text{collar}}(\text{base}(b))$. Then $D(a)$ is an annulus component domain of a pair $(h, v) \in b$. Since b is a saturated non-annular slice, $D(a)$ supports no pair in b . Thus a can be added to b to form a slice τ in the hierarchy H . The non-orderability of pairs in a slice (Lemma 2.4) implies that for each p in a (here $p = p_a$) and each pair $p' \in b$, we have p and p' are not \prec_p ordered contrary to our assumption. This completes the proof of Part (4). \square

With this lemma in hand, we can reduce all expressions of length 3 in these relations to length 2:

- (1) $a \vee b \vee c \implies a \vee c$: By transitivity of \prec_p we have $p \prec_p p''$ for all $p \in a$ and $p'' \in c$. Thus, we have $a \vee c$.
- (2) $a \dashv b \dashv c \implies a \dashv c$: This also follows from transitivity of \prec_p and of the inclusion relation.
- (3) $a \dashv b \vee c \implies a \vee c$: By definition there exists $p' \in b$ such that for all $p \in a$ we have $p \prec_p p'$. For all $p'' \in c$ we also have $p' \prec_p p''$. Thus, by transitivity of \prec_p , we have $a \vee c$.
- (4) $a \vdash b \vee c \implies a \vdash c, a \dashv c$ or $a \vee c$: Let $p \in a$ be such that $p \prec_p p'$ for all $p' \in b$. Then by transitivity $p \prec_p p''$ for all $p'' \in c$. If $D(a) \supset D(c)$ we may conclude $a \vdash c$. In all other cases, we use Lemma 4.5: if $\check{D}(a)$ and $\check{D}(b)$ do not overlap then option (4) of that lemma guarantees that $g_a \prec_t g_b$ since $p \prec_p p''$ for at least one $p \in a$ and $p'' \in c$. Thus, applying the second paragraph of option (2) of Lemma 4.5, we may conclude $a \vee c$ in this case.

Options (1), (2), and (3) remain: either $D(a)$ and $D(c)$ overlap and neither is strictly contained in the other, or we have $D(a) \subset D(c)$ and $D(a) \not\subseteq \overline{\text{collar}}(\text{base}(c))$. These cases have the consistency property: all orderings between pairs in a and pairs in c are in the same direction. Thus, the existence of $p \in a$ and $p'' \in c$ with $p \prec_p p''$, precludes the cases $c \vdash a$ and $c \vee a$, leaving $a \dashv c$ and $a \vee c$ as the remaining possibilities.

- (5) $a \dashv b \vdash c \implies a \mid c$: by definition.
- (6) $a \vdash b \dashv c \implies a \vdash c, a \dashv c$, or $a \vee c$: Again there is some $p \in a$ and $p'' \in c$ with $p \prec_p p''$. Thus, the reasoning in case (4) using Lemma 4.5 guarantees that either $a \vdash c, a \dashv c$, or $a \vee c$.

All other reductions using the relations \dashv, \vdash and \vee are just like these up to change of direction (e.g. $a \vee b \dashv c$ reduces like $c \vdash b \vee a$ does).

Reductions of expressions with $|$ then follow. For example (letting $*$ denote \neg, \vdash or \vee), $a | b * c$ is $a \neg x \vdash b * c$, which reduces to $a \neg x \neg c$ or $a \neg x \vdash c$ or $a \neg x \vee c$, and each of these reduce to one relation by the above list. A similar reduction works for $a | b | c$.

We conclude the latter part of the lemma: the transitive closure of \neg, \vdash and \vee is obtained just by adjoining the relation $|$.

The fact that \prec_c is a partial order now reduces to checking that $a \prec_c a$ is never true. For $a \neg a$, $a \vdash a$ and $a \vee a$ this follows from Lemma 2.4. If $a | a$, then for some x we have $a \neg x \vdash a$, so that there exists $p \in x$ so that for each $p' \in a$, $p' \prec_p p \prec_p p'$. This contradicts the fact that \prec_p is a strict partial order. \square

We deduce the following as a corollary.

Corollary 4.6. *If $a \prec_c b$ then, for every $p \in a$ and $p' \in b$ that are \prec_p -ordered we have $p \prec_p p'$.*

Proof. If $a \vee b$, $a \neg b$ or $a \vdash b$, this already follows from Lemma 4.5. The remaining possibility is that $a | b$, so that for some $x \in C$ we have $a \neg x \vdash b$. There is some $q \in x$ such that $p \prec_p q$ for all $p \in a$, and there is some $q' \in x$ such that $q' \prec_p p'$ for all $p' \in b$. Thus, if $p' \prec_p p$ we have $q' \prec_p q$. This contradicts Lemma 2.4, concluding the proof. \square

4.5. Topological partial order

Given a cut system C , each slice $a \in C$ determines either a split-level surface F_a as the disjoint union of the 3-holed spheres $\hat{D}(a)$ in M_ν , or, if a is annular, a determines a solid torus $U(v)$ where $v = \text{core}(a)$. We will now relate the \prec_c -order on the slices of a cut system to the \prec_{top} -order on their associated split-level surfaces and solid tori in M_ν . (We remind the reader that the ordering \prec_{top} is defined on disconnected subsurfaces of S in Section 3.1, and therefore \prec_{top} applies to the split-level surfaces F_a).

To begin with, we relate \prec_t -ordering properties in the hierarchy H_ν of the 3-holed spheres Y arising as component domains in H_ν and the annulus geodesics k_v arising for each vertex $v \in H_\nu$ to the topological ordering \prec_{top} applied to the level surfaces F_Y and solid tori $U(v)$ in M_ν . We will then use these ordering relations to relate the \prec_c -ordering to \prec_{top} -order on the surfaces \hat{F}_a , for slices a in a cut system C .

In order to discuss this relationship, fix a resolution $\{\tau_i\}$ for the hierarchy H . The following definitions allow us to keep track of the parts of the resolution sequence where certain objects appear. Let v denote a simplex whose vertices appear in H , Y a 3-holed sphere in S , k a geodesic in H , and

$p = (h, w)$ a geodesic-simplex pair in H . Define:

$$\begin{aligned} J(v) &= \{i : v \subset \text{base}(\tau_i)\}. \\ J(p) &= \{i : p \in \tau_i\}. \\ J(Y) &= \{i : Y \subset S \setminus \text{collar}(\text{base}(\tau_i))\}. \\ J(k) &= \{i : \exists v \in k, (k, v) \in \tau_i\}. \end{aligned}$$

To relate the appearance of these intervals in \mathbb{Z} to partial orderings in the hierarchy H , we record the following consequence of the *slice order* \prec_s in [41, §5].

Lemma 4.7. *Let τ_i and τ_j be slices in a resolution $\{\tau_n\}$ of H with $i < j$. Then if $p \in \tau_i$ and $p' \in \tau_j$ are \prec_p -ordered, we have*

$$p \prec_p p'.$$

Proof. The slices in the resolution $\{\tau_n\}$ are ordered with respect to the order \prec_s on complete slices with bottom geodesic the main geodesic of H . By [41, Lemma 5.3], we have $\tau_i \prec_s \tau_j$ and therefore that each pair $q \in \tau_i$ either also lies in τ_j or there is a $q' \in \tau_j$ with $q \prec_p q'$ (by the definition of \prec_s).

Assume that $p' \prec_p p$. Then p and p' cannot both lie in τ_i by Lemma 2.4, so it follows that there is a p'' in τ_j with $p \prec_p p''$. By transitivity of \prec_p we have $p' \prec_p p''$, with p' and p'' both in τ_j , which contradicts Lemma 2.4 applied to τ_j . \square

Clearly $J(Y) = J([\partial Y])$, and if $p = (h, w)$ and $v \subset w$ then $J(p) \subset J(v)$. We also have:

Lemma 4.8. *Let Y , $p = (h, w)$ and $v \subset w$ be as above. Then $J(Y)$ and $J(v)$ are intervals. If h has positive length then $J(h)$ and $J(p)$ are intervals as well.*

Proof. The conclusion for $J(v)$ was proven in [47, Lemma 5.6]. For any simplex w , $J(w) = \bigcap_{v \in w} J(v)$ so $J(w)$ is an interval too. For a 3-holed sphere Y , then, we use the fact that $J(Y) = J([\partial Y])$.

In any elementary move $\tau_i \rightarrow \tau_{i+1}$, some pair $(h, v) \in \tau_i$ is “advanced” to $(h, \text{succ}(v))$, some pairs (k, u) where u is the last simplex are erased, and some pairs (k', u') where u' is the first simplex, are created. Thus a geodesic can only appear at its beginning, advance monotonically, and disappear at its end. If a geodesic h of positive length makes two such appearances, then choosing two simplices $v < v'$ on h , they appear in slices with both orders along the resolution. This contradiction of Lemma 4.7 above implies that a positive-length geodesic occurs only once.

For a pair $p = (h, w)$, we know by hypothesis that $|h| > 0$ so we may repeat the idea of the first paragraph. If w is not the last in h then after p disappears we must have $p' = (h, \text{succ}(w))$, so that if p reappears the ordering is reversed, giving a contradiction. If w is the last simplex in h and it disappears and reappears there must also be a recurrence of (h, w_0) where

w_0 is the first vertex of h ; hence again we have a contradictory reversal of order. \square

Now we consider ordering relations for vertices and their associated solid tori in M_ν . Given a vertex v , note that the interval $J(v)$ is precisely the interval $J(k_v)$ corresponding to the annulus geodesic k_v associated to v .

Lemma 4.9. *Let v, v' be vertices appearing in H , whose corresponding curves γ_v and $\gamma_{v'}$ intersect non-trivially in S . The following are equivalent:*

- (1) $U(v) \prec_{\text{top}} U(v')$.
- (2) $\max J(v) < \min J(v')$
- (3) $k_v \prec_t k_{v'}$
- (4) *For any pairs $p = (h, w)$ and $p' = (h', w')$, if $v \in w$ or $h = k_v$, and if $v' \in w'$ or $h' = k_{v'}$ we have $p \prec_p p'$.*

Furthermore, either these relations or their opposites hold.

Proof. We note first that since γ_v and $\gamma_{v'}$ intersect non-trivially, the tori $U(v)$ and $U(v')$ must be \prec_{top} -ordered. Further, since v and v' cannot appear simultaneously in any slice in the resolution, the intervals $J(v)$ and $J(v')$ must be disjoint. Likewise, since the domains of the geodesics k_v and $k_{v'}$ overlap, Lemma 2.2 guarantees that k_v and $k_{v'}$ must be \prec_t -ordered. By Lemma 4.1 the resolution $\{\tau_i\}$ yields a sweep through the model M_ν by split-level surfaces which is monotonic in the sense that if two overlapping level surfaces W and W' appear in the sweep with W occurring first, then $W \prec_{\text{top}} W'$. This applies both to level surfaces and to solid tori with respect to their associated annular domains. Hence (1) and (2) are equivalent.

As observed above, $J(v) = J(k_v)$, so for $i \in J(v)$ we must have a pair of the form (k_v, u) in τ_i . For $j \in J(v')$ we must have some $(k_{v'}, u')$ in τ_j . The \prec_t -ordering of k_v and $k_{v'}$ is consistent with their ordering in the resolution. Thus if (2) holds then $i < j$, and hence $k_v \prec_t k_{v'}$. Reversing the roles of v and v' we obtain the other implication, so (2) and (3) are equivalent.

Now for $p = (h, w)$ with $v \in w$ or $h = k_v$ and $p' = (h', w')$ with $v' \in w'$ or $h' = k_{v'}$, note that $J(p) \subseteq J(v)$ and $J(p') \subseteq J(v')$. Since γ_v and $\gamma_{v'}$ have essential intersections, p and p' cannot appear simultaneously in slices of the resolution, and hence $J(p)$ and $J(p')$ must be disjoint. Since $J(v)$ and $J(v')$ are intervals by Lemma 4.8, it follows immediately that $\max J(p) < \min J(p')$ if and only if $\max J(v) < \min J(v')$. Now if $\max J(p) < \min J(p')$, we cannot have $p' \prec_p p$ by Lemma 4.7.

To conclude that $p \prec_p p'$, then, we need only prove that p and p' are \prec_p -ordered. Since γ_v and $\gamma_{v'}$ intersect nontrivially, the domains $D(h)$ and $D(h')$ intersect. If one is not inside the other then h and h' are \prec_t -ordered by [41, Lemma 4.18] and hence p and p' are \prec_p -ordered. If $D(h) = D(h')$ are equal then, since $w \neq w'$, either $w < w'$ or $w > w'$ and again they are \prec_p -ordered. If, say, $D(h) \subset D(h')$ then, since γ_v intersects $\gamma_{v'}$, we have $w' \notin \phi_{h'}(D(h))$. If $w' > \max \phi_{h'}(D(h))$, then in particular $D(h)$ intersects $\mathbf{T}(h')$ and hence

$h \searrow h'$ by Theorem 2.1. It follows by definition that $p \prec_p p'$. The case $w' < \min \phi_{h'}(D(h))$ is similar.

We hence conclude that (2) \implies (4). Again for the opposite direction reverse the roles of v and v' . \square

For 3-holed spheres, we have the following:

Lemma 4.10. *Let Y, Y' be 3-holed spheres appearing as component domains in H , and suppose that Y and Y' intersect essentially. Then the following are equivalent:*

- (1) $F_Y \prec_{\text{top}} F'_{Y'}$
- (2) $\max J(Y) < \min J(Y')$
- (3) $Y \prec_t Y'$.

Proof. The equivalence of (1) and (2) follows from the same argument as in Lemma 4.9.

We now show that (3) implies (2). Suppose $Y \prec_t Y'$. Recall from section 2.2 that this implies there exists a geodesic m in H such that

$$Y \searrow m \swarrow Y'$$

and that

$$\max \phi_m(Y) < \min \phi_m(Y').$$

Furthermore, recall that a 3-holed sphere can only be directly subordinate to a 4-geodesic. It follows that there exist 4-geodesics f, b' (possibly the same) such that

$$Y \searrow^d f \supseteq m \subseteq b' \swarrow^d Y'.$$

Let v be the vertex of f such that Y is a component domain of $(D(f), v)$, and let v' be the vertex of b' such that Y' is a component domain of $(D(b'), v')$.

There is exactly one elementary move in the resolution that replaces (f, v) with $(f, \text{succ}(v))$ (note that $\text{succ}(v)$ exists by definition of $Y \searrow^d f$). Before this move Y is a complementary domain of the slice marking, and afterwards it is not, since $\text{succ}(v)$ intersects Y . Hence,

$$\max J(Y) = \max J(v).$$

The same logic gives us

$$\min J(Y') = \min J(v').$$

We claim that $k_v \prec_t k_{v'}$. The annulus $D(k_v)$ is a component domain of $(D(f), v)$ and likewise $D(k_{v'})$ is a component domain of $(D(b'), v')$. It follows that $k_v \searrow^d f$ and $b' \swarrow^d k_{v'}$, and thus

$$k_v \searrow^d f \supseteq m \subseteq b' \swarrow^d k_{v'}.$$

The claim follows provided the footprints of $D(k_v)$ and $D(k_{v'})$ on m are disjoint and correctly ordered.

In the case that $m = f$ we note that since f is a 4-geodesic the vertex $(f, \text{succ}(v))$ intersects v and Y . Thus we have

$$v = \max \phi_m(Y) = \max \phi_m(D(k_v)),$$

and likewise when $m = b'$ then we have $v' = \min \phi_m(Y') = \min \phi_m(D(k_{v'}))$.

When $f \searrow m$, we may apply Lemma 4.11 of [41] or Lemma 5.5 of [47] to conclude that

$$\max \phi_m(D(k_v)) = \max \phi_m(Y) = \max \phi_m(D(f)).$$

Likewise if $m \swarrow b'$ then we similarly conclude that

$$\min \phi_m(D(k_{v'})) = \min \phi_m(Y') = \min \phi_m(D(b'))$$

We also know, in particular, that all these footprints are non-empty.

Since we have $\max \phi_m(Y) < \min \phi_m(Y')$, it follows that $\max \phi_m(D(k_v)) < \min \phi_m(D(k_{v'}))$ and thus that $k_v \prec_t k_{v'}$.

Applying Lemma 4.9, then, we have

$$\max J(v) < \min J(v'),$$

so we may conclude that $\max J(Y) < \min J(Y')$ and this establishes (2). To show that (2) implies (3), note that since Y and Y' intersect they must be \prec_t -ordered. Hence, if (3) is false we have $Y' \prec_t Y$ and we apply the above argument to reach a contradiction. \square

Lemma 4.11. *Let Y be a 3-holed sphere and v' a vertex appearing in H such that $\gamma_{v'}$ and Y overlap. Then the following are equivalent:*

- (1) $F_Y \prec_{\text{top}} U(v')$,
- (2) $\max J(Y) < \min J(v')$, and
- (3) $Y \prec_t k_{v'}$.

Symmetric conditions hold for $U(v) \prec_{\text{top}} F_{Y'}$, and either these relations or their opposites hold.

Proof. As in the previous lemmas, we first show that (1) is equivalent to (2). If $F_Y \prec_{\text{top}} U(v')$ then there is a $v \in [\partial Y]$ so that γ_v and $\gamma_{v'}$ intersect nontrivially. Thus, as in lemma 4.9, the solid tori $U(v)$ and $U(v')$ must be \prec_{top} -ordered and it follows that the height intervals they determine in the vertical coordinate of M_ν are disjoint. Likewise, since $F_Y \prec_{\text{top}} U(v')$, its height in M_ν must be less than that of the minimum height of $U(v')$ in M_ν . Since we have $F_Y \cap \partial U(v) \neq \emptyset$, we conclude that $U(v) \prec_{\text{top}} U(v')$.

By lemma 4.9 we have $\max J(v) < \min J(v')$, but this guarantees that $\max J(Y) < \min J(v')$ since $J(Y) \subset J(v)$. Assume that $\max J(Y) < \min J(v')$. Since $J(v)$ and $J(v')$ are disjoint, $J(v)$ is an interval and $J(Y) \subset J(v)$, we have $\max J(v) < \min J(v')$. Applying Lemma 4.9 we have $U(v) \prec_{\text{top}} U(v')$ and therefore that $F_Y \prec_{\text{top}} U(v')$ (since F_Y is a level surface).

Finally we show the equivalence of (3) and (2). Arguing as above, if $Y \prec_t k_{v'}$, then there is a geodesic m so that

$$Y \searrow^d f \supseteq m \subseteq b' \swarrow^d k_{v'}.$$

and $\max \phi_m(Y) < \min \phi_m(D(k_{v'}))$. Again, let $v \in [\partial Y]$ be such that γ_v is the component of ∂Y that intersects $\gamma_{v'}$. Then, once again, since v and v' intersect, the geodesics k_v and $k_{v'}$ are \prec_t -ordered, and since $\phi_m(Y) \subset \phi_m(D(k_v))$ their \prec_t -ordering must be consistent: we have $k_v \prec_t k_{v'}$.

Thus, by Lemma 4.9, we have $\max J(v) < \min J(v')$ and so since $J(Y) \subset J(v)$, we have

$$\max J(Y) < \min J(v').$$

As before, to show the converse, we simply observe that since $\gamma_{v'}$ and Y overlap, we have that Y and $k_{v'}$ are \prec_t -ordered. If $k_{v'} \prec_t Y$, we apply the above argument to conclude $\max J(v') < \min J(Y)$, which is a contradiction. This shows the equivalence of (2) and (3), concluding the proof. \square

Now let us go back to considering slices in a cut system C . Let a and b be two distinct slices in the cut system C whose domain surfaces $D(a)$ and $D(b)$ overlap. We would like to ensure that the surfaces \widehat{F}_a and \widehat{F}_b in the model are \prec_{top} -ordered if and only if a and b are consistently \prec_c -ordered.

Before commencing the proof, we argue that distinct non-annular slices a and b in a cut system have no underlying curves in common.

Lemma 4.12. *If a and b are two distinct non-annular slices in a cut system C with spacing lower bound $d_1 \geq 5$, then $\text{base}(a)$ and $\text{base}(b)$ have no vertices in common.*

Proof. If $D(a)$ and $D(b)$ are disjoint then the bases must be disjoint, and we are done. Thus from now on we may assume that $D(a)$ and $D(b)$ intersect non-trivially. Suppose by way of contradiction that there is a vertex v common to $\text{base}(a)$ and $\text{base}(b)$.

If $D(a) = D(b)$ then $g_a = g_b$ and v_a and v_b are simplices on g_a spaced at least 5 apart. However since v is distance 1 from both in $\mathcal{C}(D(a))$, this is a contradiction. From now on we assume that $D(a) \neq D(b)$. Thus by Lemma 4.5 a and b are \prec_c -ordered, and without loss of generality we may assume $a \prec_c b$.

Let $p_1 = (h_1, u_1)$ and $p_2 = (h_2, u_2)$ be the pairs of a and b , respectively, such that u_1 and u_2 contain the vertex v . We claim that there is a pair $q = (k, w)$ in the hierarchy such that

$$p_1 \prec_p q \prec_p p_2$$

and γ_w intersects γ_v non-trivially.

To see this, note first that we may assume either $a \dashv b$ or $a \vee b$ (the case $a \vdash b$ follows by a symmetrical argument). In either case there is a pair $p' = (h', x') \in b$ such that $p_a \prec_p p'$, and in fact (see Lemma 4.5) for any simplex u of g_a we have $(g_a, u) \prec_p p'$. Let $k = g_a$, and let w be a simplex in k such that $v_a < w$ and $d_{D(a)}(v_a, w) > 3$. This is possible because of the lower spacing constant $d_1 \geq 5$, and $q = (k, w)$ will be our desired pair: Since $d_{D(a)}(v_a, v) \leq 1$, we have that γ_v must intersect γ_w . Since the footprint $\phi_k(D(h_1))$ contains v_a and has diameter at most 2, we also have $\max \phi_k(D(h_1)) < w$, and so $p_1 \prec_p q$. Because γ_w and γ_v have non-trivial intersection, q and p_2 must be \prec_p -ordered, by Lemma 4.9. Thus, to show that $q \prec_p p_2$ it suffices to rule out $p_2 \prec_p q$. But since $q \prec_p p'$, if $p_2 \prec_p q$ then $p_2 \prec_p p'$, and this contradicts the non-orderability of different pairs in a slice (Lemma 2.4). We conclude that q has the desired properties.

Now fixing a resolution for the hierarchy, let $J(q)$ and $J(p_i)$ be defined as before. Since γ_w intersects γ_v we must have that $J(q)$ is disjoint from both $J(p_1)$ and $J(p_2)$. Since $p_1 \prec_p q \prec_p p_2$, we can apply Lemma 4.9 to obtain

$$\max J(p_1) < \min J(q) \leq \max J(q) < \min J(p_2).$$

On the other hand, both $J(p_1)$ and $J(p_2)$ are contained in $J(v)$, which is an interval disjoint from $J(q)$. This is a contradiction, and Lemma 4.12 is established. \square

We remark that as a consequence of Lemma 4.12, if a and b are each non-annular slices in a cut system C and $Y \subset \check{D}(a)$ and $Y' \subset \check{D}(b)$ are 3-holed spheres, then Y and Y' are distinct; otherwise there would be some common vertex in $\text{base}(a)$ and $\text{base}(b)$ in their common boundary.

4.6. Comparing topological and cut ordering

To relate \prec_c -order of a and b to topological ordering in the model, we relate the order properties we have obtained for the constituent annuli and 3-holed spheres making up $D(a)$ and $D(b)$ to the corresponding level 3-holed spheres and annuli in \widehat{F}_a and \widehat{F}_b or the corresponding solid tori when either a or b is annular. We call level 3-holed spheres and annuli in \widehat{F}_a *pieces* of \widehat{F}_a , when a is non-annular, and likewise for \widehat{F}_b . Then our first task will be to demonstrate that whenever pieces of \widehat{F}_a and \widehat{F}_b overlap (or the solid tori, when a or b is annular) they are topologically ordered consistently with the cut ordering on the slices a and b .

Lemma 4.13. *Let a and b be non-annular slices in a cut system C such that $a \prec_c b$ and $D(a)$ and $D(b)$ overlap. Let $v \in \text{base}(a)$ and $v' \in \text{base}(b)$ be distinct vertices of H , and let $Y \subset \check{D}(a)$ and $Y' \subset \check{D}(b)$ be 3-holed spheres. Then the following holds:*

- (1) if $\gamma_v \cap \gamma_{v'} \neq \emptyset$ then $U(v) \prec_{\text{top}} U(v')$,
- (2) if $\gamma_v \cap Y' \neq \emptyset$ then $U(v) \prec_{\text{top}} F_{Y'}$,
- (3) if $Y \cap \gamma_{v'} \neq \emptyset$ then $F_Y \prec_{\text{top}} \gamma_{v'}$, and
- (4) if $Y \cap Y' \neq \emptyset$ then $F_Y \prec_{\text{top}} F_{Y'}$.

If either a or b is annular, statements (1)-(3) hold taking $v = \text{core}(a)$ or $v' = \text{core}(b)$, and statement (1) holds with $v = \text{core}(a)$ and $v' = \text{core}(b)$ if both a and b are annular.

Proof. First assume neither a nor b is annular. We note that by Lemma 4.5 we may assume up to reversal of order and containment that we are in one of two cases: either

- (i) $D(a) \subset D(b)$ and $a \dashv b$, or
- (ii) $D(a)$ and $D(b)$ are not nested and $a \vee b$.

The structure of the arguments for each case will be different so we trace them through the arguments for each part of the lemma.

Consider first part (1) of the lemma and assume we are in case (i). Let $q = (h, w) \in a$ be a pair so that v lies in w , and let $q' = (h', w') \in b$ be a

pair so that v' lies in w' . Lemma 4.9 guarantees that $U(v)$ and $U(v')$ are \prec_{top} ordered. We claim that $U(v) \prec_{\text{top}} U(v')$.

Otherwise we have $U(v') \prec_{\text{top}} U(v)$, and we note that by the same lemma this implies that $\max J(v') < \min J(v)$ and that $q' \prec_p q$.

Since we are in case (i), we have $a \dashv b$, however, and this guarantees that there is some $p' \in b$ so that $q \prec_p p'$. But this implies that

$$q' \prec_p q \prec_p p'$$

where q' and p' are both pairs in b , contradicting Lemma 2.4; we conclude that $U(v) \prec_{\text{top}} U(v')$ and that $\max J(v) < \min J(v')$.

In our treatment of part (1) in case (ii), we actually prove a slightly stronger statement for future reference in the discussion of parts (2)-(4). Rather than simply considering vertices $v \in \text{base}(a)$ and $v' \in \text{base}(b)$ we verify part (1), case (ii), for any pair of vertices

$$v \in [\partial D(a)] \cup \text{base}(a) \quad \text{and} \quad v' \in [\partial D(b)] \cup \text{base}(b)$$

for which $\gamma_v \cap \gamma_{v'} \neq \emptyset$.

We claim there is a pair $q \in a$ so that $J(q) \subseteq J(v)$ and likewise there is a $q' \in b$ so that $J(q') \subseteq J(v')$; we argue for a .

If v lies in $\text{base}(a)$ then we let $q = (h, w)$ be the pair in a so that $v \in w$. Then in each case we have $J(q) \subseteq J(v)$.

If v lies in $[\partial D(a)]$, we claim that for each $p \in a$ we have

$$J(p) \subseteq J(v).$$

To see this, we note that since v represents a curve in the boundary of $D(a)$, the vertex v is present in the base of any complete slice containing a pair with g_a as its geodesic; in particular we have that

$$J(g_a) \subseteq J(v).$$

By Lemma 4.3, we have

$$(g_a, \mathbf{I}(g_a)) \preceq_p p \preceq_p (g_a, \mathbf{T}(g_a)).$$

Applying Lemma 4.7 we have

$$\min J(g_a, \mathbf{I}(g_a)) \leq \min J(p) \quad \text{and} \quad \max J(p) \leq \max J(g_a, \mathbf{T}(g_a))$$

and so we conclude that

$$\min J(v) \leq \min J(g_a) \leq \min J(p) \leq \max J(p) \leq \max J(g_a) \leq \max J(v)$$

and thus $J(p) \subseteq J(v)$, since $J(v)$ is an interval by Lemma 4.8.

Let $q \in a$ and $q' \in b$ be chosen as above so that that $J(q) \subseteq J(v)$ and $J(q') \subseteq J(v')$. Since γ_v and $\gamma_{v'}$ intersect, Lemma 4.9 guarantees that $J(v) \cap J(v') = \emptyset$, and so we have $J(q) \cap J(q') = \emptyset$ as well.

Since we have the relation $a \vee b$, we may conclude that $q \prec_p q'$, from which it follows that $\max J(q) < \min J(q')$ by Lemma 4.7. Thus, we have

$$\max J(v) < \min J(v').$$

It follows that we have

$$U(v) \prec_{\text{top}} U(v')$$

as before.

This establishes part (1), and we have that if $v \in \text{base}(a)$, and $v' \in \text{base}(b)$ or $v = v_b$ with $\gamma_v \cap \gamma_{v'} \neq \emptyset$ then $\max J(v) < \min J(v')$ by Lemma 4.9. Moreover, in case (ii) the proof of part (1) shows that if v lies in $[\partial D(a)] \cup \text{base}(a)$ and v' lies in $[\partial D(b)] \cup \text{base}(b)$, then $\gamma_v \cap \gamma_{v'} \neq \emptyset$ guarantees that $\max J(v) < \min J(v')$.

Now consider part (2). If $v \in \text{base}(a)$ and there is a $Y' \subset \check{D}(b)$ with $\gamma_v \cap Y' \neq \emptyset$, then by Lemma 4.12 either

- (i) $D(a) \subset D(b)$ and there is a $v' \in \text{base}(b) \cap [\partial Y']$ with $\gamma_v \cap \gamma_{v'} \neq \emptyset$, or
- (ii) $D(a)$ and $D(b)$ are non-nested and there is a $v' \in [\partial Y']$ for which $\gamma_v \cap \gamma_{v'} \neq \emptyset$.

(Note that in case (ii) we must allow that the vertex v' lies in $[\partial D(b)]$). In each case, applying the proof of part 1, we have

$$\max J(v) < \min J(v') \leq \min J(Y'),$$

and therefore that

$$U(v) \prec_{\text{top}} F_{Y'}$$

by Lemma 4.11. The argument for part (3) is symmetrical.

Finally, for part (4), if there are 3-holed spheres $Y \subset \check{D}(a)$ and $Y' \subset \check{D}(b)$ that overlap then once again by Lemma 4.12 either

- (i) $D(a) \subset D(b)$ and there is a $v \in \text{base}(a) \cap [\partial Y]$ and a $v' \in \text{base}(b) \cap [\partial Y']$ with $\gamma_v \cap \gamma_{v'} \neq \emptyset$, or
- (ii) $D(a)$ and $D(b)$ are non-nested and there are vertices $v \in [\partial Y]$ and $v' \in [\partial Y']$ with $\gamma_v \cap \gamma_{v'} \neq \emptyset$.

Again, applying the proof of part (1) in each case, we have

$$\max J(Y) \leq \max J(v) < \min J(v') \leq \min J(Y'),$$

so we may apply Lemma 4.10 to conclude that

$$F_Y \prec_{\text{top}} F_{Y'}$$

as desired.

Now consider the case that only a is annular and assume that $v = \text{core}(a)$. First consider part (1). In case (i), we simply take $q = (k_v, w)$ to be the unique pair in the slice a and the arguments go through verbatim. In case (ii), the same choice for q satisfies $J(q) \subseteq J(v)$, which is the only condition required to complete the proof in this case. For parts (2) and (3), the arguments go through without alteration with $v = \text{core}(a)$, and the symmetric arguments go through for parts (1)-(3) when only b is annular and $v' = \text{core}(b)$.

When both a and b are annular, taking $v = \text{core}(a)$ and $v' = \text{core}(b)$ we are necessarily in part (1), case (ii), and again the argument goes through without alteration. (One may argue more directly in this case by noting

that $a \vee b$ and applying Lemma 4.9 to the unique pairs $q \in a$ and $q' \in b$. \square

We are now ready to prove the following.

Proposition 4.14. *If a and b are slices in a cut system C with overlapping domains, then*

$$\widehat{F}_a \prec_{\text{top}} \widehat{F}_b \iff a \prec_c b.$$

Proof. Assume that $a \prec_c b$. Since $D(a)$ and $D(b)$ overlap, we know that $\check{D}(a)$ and $\check{D}(b)$ also overlap, since otherwise either

$$D(a) \subset \overline{\text{collar}}(\text{base}(b)) \quad \text{or} \quad D(b) \subset \overline{\text{collar}}(\text{base}(a))$$

and in this case a and b are not \prec_c -ordered by Lemma 4.5. It follows that all vertices of $\text{base}(a)$ and $\text{base}(b)$, or the vertices corresponding to the cores of $D(a)$ or $D(b)$ if either is annular, satisfy the hypotheses of Lemma 4.13. Thus, whenever pieces or the solid tori making up \widehat{F}_a and \widehat{F}_b overlap they are consistently topologically ordered in M_ν .

Given $t \in \mathbb{R}$, let $T_t: S \times \mathbb{R} \rightarrow S \times \mathbb{R}$ be the translation $T_t(x, s) = (x, s + t)$ in the vertical (\mathbb{R}) direction, and consider the embeddings of \widehat{F}_a and \widehat{F}_b into $S \times \mathbb{R}$ as subsets of M_ν (see section 2.6). Then the consistent topological ordering guarantees that for each positive s and t we have

$$T_{-s}(\widehat{F}_a) \cap T_t(\widehat{F}_b) = \emptyset$$

(recall that when a is not annular, each annular piece of \widehat{F}_a is contained in a solid torus $U(v)$ for $v \in \text{base}(a)$).

Thus, these translations provide homotopies of \widehat{F}_a to $-\infty$ in the complement of \widehat{F}_b and of \widehat{F}_b to $+\infty$ in the complement of \widehat{F}_a . Applying Lemma 3.1, it follows that

$$\widehat{F}_a \prec_{\text{top}} \widehat{F}_b.$$

Conversely, we assume that $\widehat{F}_a \prec_{\text{top}} \widehat{F}_b$. Since $D(a)$ and $D(b)$ are overlapping domains, Lemma 4.5 guarantees that either

- a and b are \prec_c -ordered, or
- we have $D(a) \subset \overline{\text{collar}}(\text{base}(b))$ or $D(b) \subset \overline{\text{collar}}(\text{base}(a))$.

In the latter case, if $D(a) \subset \overline{\text{collar}}(\text{base}(b))$ then a is an annular slice, and $D(a) = \overline{\text{collar}}(v)$ for some $v \in \text{base}(b)$. Then \widehat{F}_b intersects $U(v) = \widehat{F}_a$ in an annulus, and so \widehat{F}_b and \widehat{F}_a are not \prec_{top} -ordered, which is a contradiction. The symmetric argument rules out $D(b) \subset \overline{\text{collar}}(\text{base}(a))$.

Thus a and b are \prec_c ordered. If $b \prec_c a$, then the previous argument guarantees that

$$\widehat{F}_b \prec_{\text{top}} \widehat{F}_a$$

contradicting the hypothesis. This completes the proof. \square

We conclude with the following consequence, guaranteeing that topological order on the extended split-level surfaces arising from a cut system is a partial order.

Proposition 4.15. (Topological Partial Order) *The relation \prec_{top} on the components of $\{\widehat{F}_a : a \in C\}$ has no cycles, and hence its transitive closure is a partial order.*

Proof. We have shown that \prec_{top} is equivalent to the relation \prec_c restricted to surfaces $\{\widehat{F}_a : a \in C\}$ whose domains overlap.

Thus the transitive closure of \prec_{top} over all the cut surfaces is a subrelation of \prec_c (which was already transitive). Since \prec_c is a partial order, \prec_{top} has no cycles. \square

5. Regions and addresses

In this section we will explore the way in which a cut system divides the model manifold into complementary regions, whose size and geometry are bounded in terms of the spacing constants of the cut system.

For the remainder of the section we fix a cut system C . The split-level surfaces $\{F_\tau : \tau \in C\}$ divide $M_\nu[0]$ into components which we call *complementary regions* of C (or just *regions*).

In §5.2, we will define the *address* of a block in $M_\nu[0]$ in terms of the way the block is nested among the split-level surfaces of C . In §5.3 we will then describe the structure of each subset $\mathcal{X}(\alpha) \subset M_\nu[0]$ consisting of blocks with address α . In particular Lemma 5.6 will show that, roughly speaking, $\mathcal{X}(\alpha)$ can be described as a product region bounded between two split-level surfaces, minus a union of smaller product regions (and tubes). We will also prove Lemma 5.7, which shows that each complementary region of C lies in a unique $\mathcal{X}(\alpha)$.

In §5.4 we will bound the size (i.e. number of blocks) of each $\mathcal{X}(\alpha)$, and hence of each complementary region of C .

In §5.5 we will extend the discussion to the filled model $M_\nu[k]$ with $k \in [0, \infty]$. The filled cut surfaces $\widehat{F}_\tau[k]$ cut $M_\nu[k]$ into connected components, and we shall show in Proposition 5.9 that, under appropriate assumptions on the spacing constants of C , these components correspond in a simple way to the components in $M_\nu[0]$.

In the rest of the section, for an internal block B let W_B denote the “halfway surface” $D(B) \times \{0\}$ in the parametrization of B as a subset of $D(B) \times [-1, 1]$. If B is a boundary block let W_B denote its outer boundary.

5.1. More ordering lemmas

Before we get started let us prove three lemmas involving slice surfaces and \prec_{top} .

The first is another “transitivity” lemma.

Lemma 5.1. *Let c and d be two slices in a cut system C , and let B be a block with $D(B) \subset D(c) \cap D(d)$. If the halfway surface W_B satisfies*

$$\widehat{F}_c \prec_{\text{top}} W_B \quad \text{and} \quad W_B \prec_{\text{top}} \widehat{F}_d$$

then

$$\widehat{F}_c \prec_{\text{top}} \widehat{F}_d.$$

Although this statement does not seem surprising we note that, since \prec_{top} is not in general transitive, and W_B is not itself a cut surface, we must take care in the proof.

Proof. Assume first that B is an internal block. A cut surface \widehat{F}_τ is a union of level surfaces (3-holed sphere gluing surfaces) and annuli embedded in straight solid tori. Call the level surfaces and the solid tori the “pieces” associated to \widehat{F}_τ . W_B is also a level surface, and moreover avoids (the interiors of) all solid tori and gluing surfaces in M_ν . It is therefore \prec_{top} -ordered with any piece which it overlaps. For overlapping pieces x, y, z it is easy to see that $x \prec_{\text{top}} y$ and $y \prec_{\text{top}} z$ implies $x \prec_{\text{top}} z$. Now let x and y be pieces associated with c and d , respectively, which overlap each other and W_B . These exist since $D(B) \subset D(c) \cap D(d)$, and the projections of the pieces of c and d to $D(c)$ and $D(d)$, respectively, decompose them into essential subsurfaces. These subsurfaces cover all of $D(B)$, and hence must intersect each other there.

From the hypotheses of the lemma we conclude that $x \prec_{\text{top}} W_B$, and $W_B \prec_{\text{top}} y$, and therefore $x \prec_{\text{top}} y$.

Now since c and d have overlapping domains they are \prec_c -ordered by Lemma 4.5, and by Lemma 4.14 we may conclude that either $\widehat{F}_c \prec_{\text{top}} \widehat{F}_d$ or $\widehat{F}_d \prec_{\text{top}} \widehat{F}_c$. The latter implies $y \prec_{\text{top}} x$, which contradicts $x \prec_{\text{top}} y$. We conclude that $\widehat{F}_c \prec_{\text{top}} \widehat{F}_d$.

If B is a boundary block then the theorem is vacuous, since W_B is part of the boundary of M_ν , and is embedded in $\widehat{S} \times \mathbb{R}$ in such a way that nothing in M_ν lies above it (if it is in the top boundary) or below it (if it is in the bottom). \square

The following lemma tells us that we can compare blocks and cut surfaces, whenever they overlap.

Lemma 5.2. *Let B be any block and let τ be any saturated nonannular slice. If W_B and \widehat{F}_τ overlap, then they are \prec_{top} -ordered.*

Proof. Again the lemma is immediate for boundary blocks so we may assume that B is internal. The first step is to extend τ to a maximal nonannular slice. Note that since $\text{base}(\mathbf{I}(H))$ and $\text{base}(\mathbf{T}(H))$ are maximal laminations, any saturated slice τ is full (see §2.2). Hence if the bottom geodesic g_τ is g_H , we are done. If not, there is some h such that $g_\tau \stackrel{d}{\prec} h$, and a simplex w in h such that $D(g_\tau)$ is a component domain of $(D(h), w)$. Add (h, w) to τ , and successively fill in the components of $D(h) \setminus D(g_\tau)$ to obtain a

saturated nonannular slice τ' with $g_{\tau'} = h$. Repeat this inductively until we get a saturated nonannular slice τ_0 with bottom geodesic g_H , hence a maximal nonannular slice.

Now by Lemma 5.7 of [47], there exists a (nonannular) resolution with τ_0 as one of its slices. If we now consider the sweep through M_ν determined by this resolution (see §4.2), we see that there is some moment when the block B is appended. Applying Lemma 4.1, for any slice τ_i that occurs in the resolution before this moment we have $\widehat{F}_{\tau_i} \prec_{\text{top}} W_B$, and for any τ_i that occurs after, we have $W_B \prec_{\text{top}} \widehat{F}_{\tau_i}$. Since \widehat{F}_τ is an essential subsurface of \widehat{F}_{τ_0} , and \widehat{F}_τ and W_B overlap, this implies (using Lemma 3.1) that they are \prec_{top} -ordered. \square

The next lemma allows us to compare tubes and cut surfaces, and will be used to prove the “unwrapping property” at the end of the proof in Section 8. It shows, in particular, that a slice surface \widehat{F}_c and a disjoint tube U can be moved to $-\infty$ and $+\infty$, respectively (or vice versa) without intersecting each other. In Section 8 we will apply this to their images in a hyperbolic 3-manifold N to conclude that certain surfaces cannot be wrapped around Margulis tubes, and this will allow us to construct controlled embedded surfaces in N .

Lemma 5.3. *Let τ be any saturated nonannular slice in H_ν and let w be a vertex of H_ν , such that $\mathbf{collar}(w)$ and $\check{D}(\tau)$ have non-trivial intersection.*

Then either $\widehat{F}_\tau \prec_{\text{top}} U(w)$ or $U(w) \prec_{\text{top}} \widehat{F}_\tau$.

Proof. As in Lemma 5.2, we extend τ to a maximal slice τ_0 , and fix a resolution of H that includes τ_0 . The assumption that $\mathbf{collar}(w)$ and $\check{D}(\tau)$ intersect implies that \widehat{F}_{τ_0} does not intersect $U(w)$. Thus in the sweep through M_ν defined by the resolution, \widehat{F}_{τ_0} is reached either before or after $U(w)$, and it follows as in Lemma 5.2 that they are \prec_{top} -ordered. \square

5.2. Definition of addresses

An *address pair* for a block B in M_ν is a pair of cuts (c, c') with $D(B) \subset D(c) = D(c')$, such that

$$\widehat{F}_c \prec_{\text{top}} W_B \prec_{\text{top}} \widehat{F}_{c'}.$$

We say that an address pair (c, c') is *nested within* a different address pair (d, d') if $d \preceq_c c$ and $c' \preceq_{c'} d'$. We say that an address pair is *innermost* if it is minimal with respect to the relation of nesting among address pairs for B .

Lemma 5.4. *If B has at least one address pair then it has a unique innermost address pair (c, c') and furthermore (c, c') is nested within every other address pair for B .*

Proof. Let (c, c') and (d, d') be address pairs for B . We first claim that one of $D(c)$ and $D(d)$ must be contained in the other, and that if $D(c) \subsetneq D(d)$ then (c, c') is nested within (d, d') .

Since $D(B) \subseteq D(c)$ and $D(B) \subseteq D(d)$, the domains $D(c)$ and $D(d)$ intersect. First assume that neither $D(c)$ nor $D(d)$ is contained in the other. In this case the bottom geodesics g_c and g_d are \prec_t -ordered (by Lemma 2.2), and without loss of generality we may assume $g_c \prec_t g_d$. Note also that $g_{c'} = g_c$.

By Lemma 4.5, we have $c' \prec_c d$ which implies that $\widehat{F}_{c'} \prec_{\text{top}} \widehat{F}_d$ by Lemma 4.14. On the other hand, by definition of address pairs, we have

$$\widehat{F}_d \prec_{\text{top}} W_B \quad \text{and} \quad W_B \prec_{\text{top}} \widehat{F}_{c'},$$

which by Lemma 5.1 then implies $\widehat{F}_d \prec_{\text{top}} \widehat{F}_{c'}$. But this contradicts $\widehat{F}_{c'} \prec_{\text{top}} \widehat{F}_d$, by definition of \prec_{top} . We conclude that one of the domains is contained in the other.

Suppose that $D(c) \subsetneq D(d)$. We claim that in this case we must have

$$d \prec_c c \prec_c c' \prec_c d' \tag{5.1}$$

so that (c, c') is nested within (d, d') .

To see this, note that by Lemma 4.5 we have that either

$$c \dashv d \quad \text{or} \quad d \vdash c$$

in the partial order on cuts. Suppose first that $c \dashv d$. Then there is some $p \in d$ such that for the bottom pair p_c of c , $p_c \prec_p p$, and in fact the proof of lemma 4.5 shows that $(g_c, \mathbf{T}(g_c)) \prec_p p$. Lemma 4.3 then shows that for any pair $q \in c'$, $q \prec_p p$. This implies that

$$c' \dashv d.$$

By Lemma 4.14 we have

$$\widehat{F}_{c'} \prec_{\text{top}} \widehat{F}_d,$$

and since (c, c') and (d, d') are address pairs we have both

$$\widehat{F}_d \prec_{\text{top}} W_B \quad \text{and} \quad W_B \prec_{\text{top}} \widehat{F}_{c'}.$$

By Lemma 5.1, this implies $\widehat{F}_d \prec_{\text{top}} \widehat{F}_{c'}$, again a contradiction.

Thus we have ruled out $c \dashv d$, and it follows that $d \vdash c$. By the same argument with directions reversed, we may also conclude that $c' \dashv d'$. This establishes the nesting claim (5.1).

Now suppose that $D(c) = D(d)$. The relation \prec_{top} on surfaces $\{\widehat{F}_\tau : \tau \in C, g_\tau = g_c\}$ is the same as the linear order on their bottom simplices $\{v_\tau\}$. Thus by Lemma 5.1 the sets $\{\tau \in C : g_\tau = g_c, W_B \prec_{\text{top}} \widehat{F}_\tau\}$ and $\{\tau \in C : g_\tau = g_c, \widehat{F}_\tau \prec_{\text{top}} W_B\}$ form disjoint intervals in this order, and there is a unique innermost pair.

Since we have shown that the domains of address pairs are linearly ordered by inclusion, there is a unique domain of minimal complexity, and among

the pairs with that domain there is a unique innermost one. This is the desired address pair. \square

We are now justified in making the following definition.

Definition 5.5. *If (c, c') is the innermost address pair for B then we say B has address $\langle c, c' \rangle$. If B has no address pairs, we say that B has the empty address denoted $\langle \emptyset \rangle$.*

We let $D(\langle c, c' \rangle)$ denote $D(c) = D(c')$ and let $D(\langle \emptyset \rangle) = S$. Note that if $\langle c, c' \rangle$ is an address then c and c' are successive in the \prec_c order on $C|_{g_c}$.

5.3. Structure of address regions

Having shown that each block has a well defined address, let $\mathcal{X}(\alpha)$ denote the union of blocks with address $\langle \alpha \rangle$. We will now describe the structure of $\mathcal{X}(\alpha)$ as, roughly speaking, a product region minus a union of smaller product regions.

If (c, c') is any address pair, note (e.g. by Lemma 4.14) that \widehat{F}_c and $\widehat{F}_{c'}$ are disjoint unknotted properly embedded surfaces in $D(c) \times \mathbb{R} \subset S \times \mathbb{R}$, which are isotopic to level surfaces, and transverse to the \mathbb{R} direction. Hence they cut off from $D(c) \times \mathbb{R}$ a region $\mathcal{B} = \mathcal{B}(c, c')$ homeomorphic to $D(c) \times [0, 1]$ containing (the closure of) all points above \widehat{F}_c and below $\widehat{F}_{c'}$ (in the \mathbb{R} coordinate). Define also $\mathcal{B}(c, \cdot)$ to be the set of all points in $D(c) \times \mathbb{R}$ that are above \widehat{F}_c , and define $\mathcal{B}(\cdot, c')$ similarly (these are useful for considering infinite geodesics).

The boundary of \mathcal{B} in $S \times \mathbb{R}$ is therefore the union of $\widehat{F}_c \cup \widehat{F}_{c'}$ with annuli in $\partial D(c) \times \mathbb{R}$, one for each component of $\partial D(c)$. Indeed these annuli lie in $\partial \mathcal{U}(\partial D(c))$, and their boundaries are the circles $\partial \widehat{F}_c$ and $\partial \widehat{F}_{c'}$.

It is clear from this description that a block B is contained in $\mathcal{B}(c, c')$ if and only if (c, c') is an address pair for B . If we define $\mathcal{B}(\emptyset)$ to be all of M_ν , we can generally say that $\mathcal{X}(\alpha) \subset \mathcal{B}(\alpha)$.

Furthermore if (d, d') is any address pair that is nested within (c, c') then $\mathcal{B}(d, d')$ has interior disjoint from $\mathcal{X}(\alpha)$. Similarly for any (d, d') , $\mathcal{B}(d, d')$ has interior disjoint from $\mathcal{X}(\emptyset)$.

In fact it follows from the definitions that $\mathcal{X}(\alpha)$ is obtained by deleting from the product region $\mathcal{B}(\alpha)$ all (interiors of) such product regions $\mathcal{B}(d, d')$, as well as the tubes \mathcal{U} .

If g is a geodesic with $C|_g$ nonempty, let a_g and z_g be the first and last slices of $C|_g$, if they exist (g may be infinite in either direction) and define $\mathcal{B}(g) = \mathcal{B}(a_g, z_g)$ if g is finite, $\mathcal{B}(g) = \mathcal{B}(a_g, \cdot)$ if a_g exists but not z_g , and $\mathcal{B}(g) = \mathcal{B}(\cdot, z_g)$ if z_g exists but not a_g .

We call a geodesic g an *inner boundary geodesic* for α if $D(g) \subsetneq D(\alpha)$, g supports slices $d, d' \in C|_g$ which are nested within α , and $D(g)$ is maximal by inclusion among such geodesics. For $\alpha = \langle \emptyset \rangle$, the same definition holds with the convention that every pair (d, d') is said to be nested in \emptyset . The interior of $\mathcal{B}(g)$ is then excluded from $\mathcal{X}(\alpha)$, but its boundary, and in particular F_{a_g}

and F_{z_g} (when a_g and z_g exist, respectively), must be contained in $\partial\mathcal{X}(\alpha)$. We can formulate this as a lemma:

Lemma 5.6. *For an address α of C , $\mathcal{X}(\alpha)$ is equal to*

$$\mathcal{B}(\alpha) \setminus (\mathcal{U} \cup \bigcup_h \text{int}(\mathcal{B}(h)))$$

where the union is over inner boundary geodesics h for α .

Given α , these $\mathcal{B}(h)$ are disjoint for different inner boundary geodesics h for α . Any $h \in H$ with nonempty $C|_h$ is an inner boundary geodesic for exactly one address α .

When $\alpha = \langle c, c' \rangle$ we call F_c and $F_{c'}$ the outer boundaries of $\mathcal{X}(\alpha)$. The surfaces F_{a_h} and F_{z_h} for any inner boundary geodesic h are called inner boundary subsurfaces. When $\alpha = \langle \emptyset \rangle$, the outer boundaries of $\mathcal{X}(\alpha)$ are the outer boundaries of M_ν . The boundary of $\mathcal{X}(\alpha)$, minus the annuli and tori of $\partial\mathcal{U}$, consists of these inner and outer boundary surfaces.

Proof. Most of this statement is evident from the definitions and the discussion above, but the second paragraph requires a bit of explanation.

If h and g are inner boundary geodesics for α , a nonempty intersection of $\mathcal{B}(h)$ and $\mathcal{B}(g)$ implies that, for some pairs $a, a' \in C|_h$ and $b, b' \in C|_g$, there is a block for which (a, a') and (b, b') are both address pairs. However as in the proof of lemma 5.4 this implies that one of $D(g)$ and $D(h)$ must be strictly contained in the other, and this then implies that one of (a, a') and (b, b') is nested in the other, which contradicts the definition of an inner boundary geodesic.

To show the last statement, suppose that c and c' are in $C|_h$ and there is a pair (d, d') for which $d \vdash c \dashv d'$. The argument in the proof of Lemma 5.4 implies that $d \vdash c' \dashv d'$ as well. It follows that either there is no (d, d') such that $d \vdash c \dashv d'$ for any $c \in C|_h$, or there is a nonempty collection \mathcal{D} of pairs (d, d') for which $d \vdash c \dashv d'$ for all $c \in C|_h$. In the first case, h must be an inner boundary geodesic for $\alpha = \langle \emptyset \rangle$. In the second case it follows, as in the proof of Lemma 5.4, that there is a unique innermost (d, d') in \mathcal{D} and this must be the unique address α for which h is an inner boundary geodesic. \square

With this picture in mind, we can relate addresses to connected components of the complement of the surfaces of C :

Lemma 5.7. *All blocks in a complementary region of C have the same address.*

Proof. Any two blocks B and B' in the same connected component are connected by a chain $B = B_0, \dots, B_n = B'$ where B_i and B_{i+1} are adjacent along a gluing surface which does not lie in any of the cut surfaces $\{F_c : c \in C\}$. It thus suffices to consider the case that B and B' are adjacent along gluing surfaces that are not in the cuts.

Let us show that any address pair (c, c') for B is also an address pair for B' (and, by symmetry, vice versa). This will imply that the innermost pairs, and hence the addresses, are the same.

The region $\mathcal{B}(c, c')$ contains B . Since $\partial\mathcal{B}(c, c')$ consists of F_c, F'_c and portions of the boundaries of tubes, the gluing surface connecting B to B' is not in this boundary. It follows that B' is also contained in $\mathcal{B}(c, c')$, and hence (c, c') is an address pair for B' . This completes the proof. \square

5.4. Sizes of regions

Our next lemma will bound the number of blocks in any $\mathcal{X}(\alpha)$. As an immediate consequence of Lemma 5.7, we also get a bound on the size of any complementary region of the cut system.

Lemma 5.8. *The number of blocks in $\mathcal{X}(\alpha)$ for any address α is bounded by a constant K depending only on S and d_2 .*

Proof. Fix an address α . If $\alpha = \langle c, c' \rangle$, let $g_\alpha = g_c = g_{c'}$ be the bottom geodesic for c and c' . If $\alpha = \langle \emptyset \rangle$, let $g_\alpha = g_H$.

Let $\mathcal{Z} = \mathcal{Z}_\alpha$ be the set of all three-holed spheres Y such that F_Y is a component of $\partial_- B$ for some internal block B in $\mathcal{X}(\alpha)$. Since every internal block B has nonempty $\partial_- B$ and every F_Y is in the ∂_- gluing boundary of at most one block, it follows that the number of internal blocks in $\mathcal{X}(\alpha)$ is at most $|\mathcal{Z}|$. Since there is a bound on the number of boundary blocks depending only on S , it will suffice to find a bound on $|\mathcal{Z}|$ that depends only on S and d_2 .

For a geodesic h , let $\mathcal{Y}(h)$ be the set of all three-holed spheres Y for which $Y \searrow h$. For each geodesic $h \supseteq g_\alpha$, we define

$$J_{\mathcal{Z}}(h) = \{v \in h : v = \max \phi_h(Y) \text{ for some } Y \in \mathcal{Z}, Y \subset D(h)\},$$

the set of landing points on h of forward sequences for $Y \in \mathcal{Z}$.

The bound on $|\mathcal{Z}|$ will follow from the following four claims:

- (1) $\mathcal{Z} \subset \mathcal{Y}(g_\alpha)$
- (2) If $h \supseteq g_\alpha$ and $\xi(h) > 4$ then

$$\mathcal{Z} \cap \mathcal{Y}(h) = \bigsqcup_{\substack{k \searrow h \\ \max \phi_h(k) \in J_{\mathcal{Z}}(h)}} (\mathcal{Z} \cap \mathcal{Y}(k))$$

- (3) If $h \supseteq g_\alpha$ and $\xi(h) = 4$ then

$$|\mathcal{Z} \cap \mathcal{Y}(h)| \leq 2|J_{\mathcal{Z}}(h)|$$

- (4) For any $h \supseteq g_\alpha$,

$$|J_{\mathcal{Z}}(h)| \leq m,$$

where m is a constant depending only on d_2 .

Assuming these claims hold, we can prove a bound $|\mathcal{Z} \cap \mathcal{Y}(h)| \leq K_{\xi(h)}$ by induction on $\xi(h)$. For $\xi(h) = 4$, we obtain the bound $K_4 = 2m$ by claims (3) and (4). In the induction step suppose we already have a bound $K_{\xi(h)-1}$.

For any interior simplex $v \in h$ there is a unique geodesic $k \searrow^d h$ with $\max \phi_h(D(k)) = v$ ($D(k)$ can only be the component domain of $(D(h), v)$ that intersects the successor of v). If v is the first or last simplex of h then it is a vertex so there are at most two non-annular component domains of $(D(h), v)$ and hence at most two non-annular $k \searrow^d h$ with $\max \phi_h(D(k)) = v$. Thus the union in claim (2) has at most $|J_{\mathcal{Z}}(h)| + 2$ terms. Together with claim (4) we obtain a bound of $K_{\xi(h)} = K_{\xi(h)-1}(m + 2)$. Claim (1) then gives us our desired bound $|\mathcal{Z}| \leq K_{\xi(g_\alpha)}$.

Before proving the claims, we note the following facts: If Y and Y' are three-holed spheres in the hierarchy and are both contained in $D(h)$ for some h , then

$$\max \phi_h(Y) < \min \phi_h(Y') \implies Y \prec_t Y'. \quad (5.2)$$

This follows directly from the definition of \prec_t , noting that if $\max \phi_h(Y) < \min \phi_h(Y')$ then in particular the last vertex of h is not in $\phi_h(Y)$ and the first is not in $\phi_h(Y')$, so that $Y \searrow h \swarrow Y'$.

From the contrapositive, with Y and Y' interchanged, we obtain

$$Y \prec_t Y' \implies \min \phi_h(Y) \leq \max \phi_h(Y'). \quad (5.3)$$

Now we prove claim (1). If $\alpha = \langle \emptyset \rangle$ then $g_\alpha = g_H$ and (1) is immediate since $\mathcal{Y}(g_H)$ contains all three-holed spheres which are component domains of the hierarchy except those that are component domains of $\mathbf{T}(H)$, and those are excluded from \mathcal{Z} (they correspond to ∂_- gluing surfaces of boundary blocks). Assume $\alpha = \langle c, c' \rangle$. Let B be a block in $\mathcal{X}(\alpha)$ and F_Y a component of $\partial_- B$. Since $W_B \prec_{\text{top}} \widehat{F}_{c'}$ and the interior of B is disjoint from $\widehat{F}_{c'}$, we must have $F_Y \prec_{\text{top}} \widehat{F}_{c'}$.

Let Y' be a component domain of $(D(g_\alpha), \text{base}(c'))$ which overlaps Y . It follows (applying Lemma 3.1) that $F_Y \prec_{\text{top}} F_{Y'}$. Lemma 4.10 tells us that

$$Y \prec_t Y'.$$

Thus by (5.3) we have $\min \phi_{g_\alpha}(Y) \leq \max \phi_{g_\alpha}(Y')$. (These footprints are nonempty by Lemma 2.3.) Letting (g_α, v') denote the bottom pair of c' , since Y' is a complementary component of $\text{base}(c')$ it follows that $\phi_{g_\alpha}(Y')$ contains v' . Since the lower spacing bound d_1 for the cut system is at least 5, and footprints have diameters at most 2, it follows that $\max \phi_{g_\alpha}(Y')$ is at least 3 away from the last simplex of g_α , and hence $\max \phi_{g_\alpha}(Y)$ is at least 1 away. Thus $Y \searrow g_\alpha$, or $Y \in \mathcal{Y}(g_\alpha)$. This establishes claim (1).

This discussion also proves claim (4) for $h = g_\alpha$ when $\alpha = \langle c, c' \rangle$: if (g_α, v) is the bottom pair of c then v and v' are at most d_2 apart. The above argument shows that $\max \phi_{g_\alpha}(Y)$ is at most 4 forward of v' , and the same argument run in the opposite order (with c replacing c') shows that $\max \phi_{g_\alpha}(Y)$ occurs no further back than 2 behind v . This restricts it to an interval of diameter $d_2 + 6$, which gives us claim (4) for g_α provided $m > d_2 + 6$.

Now consider claim (4) for $h \searrow g_\alpha$, or for $h = g_\alpha$ when $\alpha = \langle \emptyset \rangle$. We claim that if $Y \searrow h$ and $Y \in \mathcal{Z}$ then $\max \phi_h(Y)$ occurs within $d_2 + d_1/2 + 3$

of the endpoints of h . Suppose this is not the case and let us look for a contradiction. The length of h is then greater than $2d_2 + d_1$ (possibly it is infinite), which means that there must be at least two slices of C based on h . There exist slices $d, d' \in C|_h$ such whose bottom simplices v, v' satisfy $v < \max \phi_h(Y) < v'$ and are at least $d_1/2 + 3$ away from $\max \phi_h(Y)$: These can be the first and last slices of h , if these exist, since they are within d_2 of the endpoints of h ; or if h is infinite in the backward or forward direction a sufficiently far away slice will do for d or d' respectively. Note that $d_1/2 + 3 > 5$. For any component Y' of $D(h) \setminus \mathbf{collar}(\text{base}(d'))$, $\phi_h(Y')$ contains v' and it follows that $\max \phi_h(Y) < \min(\phi_h(Y'))$ so that $Y \prec_t Y'$ by (5.2) and hence $F_Y \prec_{\text{top}} F_{Y'}$ by Lemma 4.10. It follows, as in the proof of Proposition 4.14, that $F_Y \prec_{\text{top}} \widehat{F}_d$. A similar argument yields $\widehat{F}_d \prec_{\text{top}} F_Y$.

Now let B be a block in $\mathcal{X}(\alpha)$ with $F_Y \subset \partial_{\pm} B$. By Lemma 5.2, W_B and \widehat{F}_d must be \prec_{top} -ordered, and similarly for W_B and $\widehat{F}_{d'}$. Since the interior of B does not meet \widehat{F}_d or $\widehat{F}_{d'}$, the ordering we've established for F_Y implies that $\widehat{F}_d \prec_{\text{top}} W_B \prec_{\text{top}} F_{d'}$.

We also note that $D(B) \subset D(h)$, as follows. The block B is associated to an edge in a geodesic k with $D(k) = D(B)$. Assume without loss of generality that $Y \subset \partial_- B$. Thus if e is the edge of k defining B we have that Y is component domain of $(D(h), e^-)$ which intersects e^+ , and in particular $Y \searrow^d k$. We also have $Y \searrow h$, since $\phi_h(Y)$ is far from the ends of h . Hence h and k are in the forward sequence $\Sigma^+(Y)$, so one is contained in the other. Since $\xi(k) = 4$, we must have $D(k) \subset D(h)$.

We can therefore conclude that (d, d') is an address pair for B . If $\alpha = \langle \emptyset \rangle$ then this is already a contradiction. If not, then since the domain of d is strictly smaller than that of c we must have (d, d') nested within (c, c') by (the proof of) Lemma 5.4, a contradiction to the assumption that (c, c') is the innermost address pair for B .

This contradiction establishes our claim, so that $\max \phi_h(Y)$ is confined to a pair of intervals of total length $2d_2 + d_1 + 6 \leq 3d_2 + 6$. This gives the desired bound on $|J_{\mathcal{Z}}(h)|$ for Claim (4).

For Claim (2), let $h \supseteq g_{\alpha}$ with $\xi(h) > 4$. Suppose that $Y \in \mathcal{Z} \cap \mathcal{Y}(h)$. Then $Y \searrow h$ but we cannot have $Y \searrow^d h$ since $\xi(h) > 4$. Thus there is a $k \in \Sigma^+(Y)$ such that $k \searrow^d h$, and hence (by Corollary 4.11 of [41]) $\max \phi_h(D(k)) = \max \phi_h(Y)$. In particular $\max \phi_h(D(k)) \in J_{\mathcal{Z}}(h)$. Since $Y \searrow k$ we also have $Y \subset \mathcal{Z} \cap \mathcal{Y}(k)$. Note that Y cannot be in $\mathcal{Z} \cap \mathcal{Y}(k')$ for a different $k' \searrow^d h$ by the uniqueness of the forward sequence $\Sigma^+(Y)$. Thus we obtain the partition of $\mathcal{Z} \cap \mathcal{Y}(h)$ described in Claim (2).

For Claim (3), let $h \supseteq g_{\alpha}$ with $\xi(h) = 4$. Now $Y \searrow h$ exactly if $Y \searrow^d h$, and this occurs when Y is a complementary component of $D(h) \setminus \mathbf{collar}(v)$ for $v = \max \phi_h(Y)$. There is one such component for each v when $D(h)$ is a one-holed torus, and two when $D(h)$ is a four-holed sphere. The inequality of claim (3) follows.

Thus we have established the bound $|\mathcal{Z}| \leq K_{\xi(S)}$, where $K_{\xi(S)}$ depends only on S and d_2 . \square

5.5. Filled regions

We will also need to consider regions determined by a cut system C in the *filled model* $M_\nu[k]$ for some constant $k > 0$. If C is a cut system, then the surfaces $\{\widehat{F}_\tau[k] : \tau \in C\}$ again decompose the model $M_\nu[k]$ into regions. We wish to verify that these regions in a filled model differ from the regions determined by $\{F_\tau : \tau \in C\}$ only by filling in certain tubes whose boundaries lie entirely in a given region. More precisely, let

$$\mathcal{W}_i = M_\nu[i] \setminus \bigcup_{c \in C} \widehat{F}_c[i].$$

Thus the components of \mathcal{W}_0 are the complementary regions in $M_\nu[0]$ of the cut system, which we have been considering up to now.

Proposition 5.9. *Given $k > 0$ there is a constant d_1 such that, if C is a cut system with a spacing lower bound of at least d_1 , then the connected components of \mathcal{W}_0 are precisely the connected components of \mathcal{W}_k minus the tubes of size $|\omega| < k$.*

In particular, all blocks in a connected component of \mathcal{W}_k have the same address.

The main step in the proof of Proposition 5.9 is the following lemma, which restricts how many slices can meet a tube with large $|\omega|$.

Lemma 5.10. *Given k there exists $d_1 > 5$ so that for any cut system C with spacing lower bound of at least d_1 and each tube $U(v)$ in $M_\nu[k]$, there is at most one nonannular $a \in C$ such that $\partial U(v)$ meets F_a .*

Proof. Let v be a vertex in H so that $|\omega(v)| < k$, and hence $U(v) \subset M_\nu[k]$. Suppose $\partial U(v)$ meets a cut surface F_a for some $a \in C$. This implies that $v \in [\partial \check{D}(a)]$, so either

- (1) $v \in [\partial D(a)]$, or
- (2) $v \in \text{base}(a)$.

The lower spacing bound on C means that the bottom geodesic h_a has length at least $3d_1$, so if $v \in [\partial D(a)]$ this yields a lower bound on $|\omega_M(v)|$. In particular letting b_1 and b_2 be the constants in Lemma 2.11, if we have chosen $d_1 \geq (k + b_1)/3b_2$, Lemma 2.11 would imply $|\omega_M(v)| \geq k$ which is a contradiction, and hence case (1) cannot occur.

Now suppose that there are two slices $a, b \in C$ such that $\partial U(v)$ meets F_a and F_b , and hence that $v \in \text{base}(a)$ and $v \in \text{base}(b)$. This possibility is ruled out by Lemma 4.12, and this completes the proof of Lemma 5.10. \square

We can now complete the proof of the Proposition:

Proof of Proposition 5.9. Choose d_1 to be the same constant as given by Lemma 5.10. \mathcal{W}_k is obtained from \mathcal{W}_0 by attaching, for each tube U with

$|\omega(U)| < k$, the set

$$U \setminus \bigcup_{c \in C} \widehat{F}_c.$$

By Lemma 5.10, U meets at most one surface \widehat{F}_c with $c \in C$, and if it does so then the intersection is a single annulus. Thus each component of $U \setminus \bigcup_{c \in C} \widehat{F}_c$ is a solid torus which either meets ∂U in the entire boundary, or in a single annulus. In either case each component meets ∂U in a connected set. This means that the components of $U \setminus \bigcup_{c \in C} \widehat{F}_c$ cannot connect different components of \mathcal{W}_0 .

It follows that a connected component of \mathcal{W}_k is equal to a connected component of \mathcal{W}_0 union the adjacent pieces of tubes.

The final statement of the proposition is an immediate consequence of Lemma 5.7. □

6. Uniform embeddings of Lipschitz surfaces

The main theorem of this section is Theorem 6.1, which proves that a Lipschitz map of a surface with bounded geometry into the manifold N_ρ can be deformed to an embedding in a controlled way, provided it satisfy a number of conditions, the most important being an “unwrapping condition” that rules out the possibility that the homotopy will be forced to go through a deep Margulis tube.

We begin by introducing a series of definitions which allow us to describe the type of surfaces we allow and to express what it means to deform to an embedding in a controlled way.

A compact hyperbolic surface X (possibly disconnected) with geodesic boundary is said to be *L-bounded* (or has a *L-bounded metric*) if no homotopically non-trivial curve in X has length less than $1/L$ and no boundary component has length more than L . A map $f : X \rightarrow N$ is *L-bounded* if X is *L-bounded* and f is *L-Lipschitz*.

An *anchored surface* (or map) is a map of pairs

$$f : (X, \partial X) \rightarrow (N \setminus \mathbb{T}(\partial X), \partial \mathbb{T}(\partial X))$$

where $X \subseteq S$ is an essential subsurface and f is in the homotopy class determined by ρ . An anchored surface is *ϵ -anchored* if $\ell_\rho(\gamma) < \epsilon$ for each component γ of ∂X .

If X has a hyperbolic metric, an anchored surface $f : X \rightarrow N$ is *$(K, \hat{\epsilon})$ -uniformly embeddable* if there exists a homotopy, called a *$(K, \hat{\epsilon})$ -uniform homotopy*,

$$H : (X \times [0, 1], \partial X \times [0, 1]) \rightarrow (N \setminus \mathbb{T}(\partial X), \partial \mathbb{T}(\partial X)) \quad (6.1)$$

with $H(\cdot, 0) = f$ such that

- H is K -Lipschitz

- H restricted to $X \times [1/2, 1]$ is a K -bilipschitz C^2 embedding with the norm of the second derivatives bounded by K .
- For all $t \in [1/2, 1]$, $H(\partial X \times \{t\})$ is a collection of geodesic circles in $\partial\mathbb{T}(\partial X)$.
- $H(X \times [1/2, 1])$ avoids all ϵ_1 -Margulis tubes with core length less than $\hat{\epsilon}$.
- $H(X \times [0, 1])$ avoids all $\hat{\epsilon}$ -Margulis tubes.

The following theorem gives a condition for an anchored bounded map to be uniformly embeddable.

Theorem 6.1. *Let S be an oriented compact surface. Given $\delta_1 \leq \epsilon_u$ and $\delta_2, L > 0$, there exist ϵ and $\hat{\epsilon}$ in $(0, \min\{\delta_1, \delta_2\})$, and $K > 0$, so that the following holds.*

Let $\rho \in \mathcal{D}(S)$, $R \subset S$ an essential subsurface, and Γ a simplex in $\mathcal{C}(R)$, and let $X = R \setminus \mathbf{collar}(\Gamma)$. Suppose that there is a δ_1 -anchored, L -bounded surface

$$f : X \rightarrow C_{N_\rho} \cap N_\rho^1$$

in the homotopy class determined by ρ , and there exists an extension

$\bar{f} : R \rightarrow N_\rho$ of f such that

- (1) \bar{f} takes $\mathbf{collar}(\Gamma)$ into $\mathbb{T}(\Gamma)$,
- (2) \bar{f} satisfies this unwrapping condition: For any $\alpha \in \mathcal{C}_0(R) \setminus \Gamma$, if $\ell_\rho(\alpha) < \delta_2$, then \bar{f} is homotopic to either $+\infty$ or $-\infty$ in $N_\rho^0 \setminus \mathbb{T}(\alpha)$.
- (3) \bar{f} is ϵ -anchored.

Then f is $(K, \hat{\epsilon})$ -uniformly embeddable, the uniform homotopy H has image in $C_{N_\rho}^{1/2} \cup (N_\rho)_{(0, \epsilon_2]}$ where $\epsilon_2 = (\epsilon_0 + \epsilon_1)/2$, and each component of $H_1(\partial X)$ is a geodesic in the intrinsic metric on $\partial\mathbb{T}(\partial X)$.

In the first stage of the proof, we will prove a general result (Proposition 6.3) which produces embedded surfaces in geometric limits of certain sequences of representations. In the second stage (§6.2) we will apply this to an alleged sequence of counterexamples and then obtain a contradiction.

6.1. Embedding in geometric limits

In this section we show that given a sequence $\{\rho_n\}$ of representations which converge on a subsurface F of S so that $\{\ell_{\rho_n}(\partial F)\}$ converges to 0 and the limits have no non-peripheral parabolics, we can produce an anchored embedding of F into the geometric limit of $\{N_{\rho_n}\}$.

Let F be a (possibly disconnected) subsurface of S which has no annulus components. Note that given a component F_i of F there is a family of homomorphisms $\sigma : \pi_1(F_i) \rightarrow \pi_1(S)$ consistent with the inclusion map (depending on choice of basepoints and arcs connecting them), and any two of these differ by conjugation in $\pi_1(S)$.

Definition 6.2. *A sequence $\{\rho_n\}$ in $\mathcal{D}(S)$ is convergent on F if, for each i there is a sequence $\{\sigma_n^i : \pi_1(F_i) \rightarrow \pi_1(S)\}$ consistent with the inclusion map*

so that the sequence of representations

$$\rho_n^i = \rho_n \circ \sigma_n^i$$

converges to a representation $\rho^i : \pi_1(F_i) \rightarrow PSL_2(\mathbb{C})$.

We call the ρ^i limit representations on F of $\{\rho_n\}$ (but note that they depend on the choice of σ_n^i).

Proposition 6.3. *Let S be an orientable surface and let F be an essential subsurface with components $\{F_i\}$, none of which are annuli. Suppose that $\{\rho_n\}$ is a sequence in $\mathcal{D}(\pi_1(S))$ such that*

- (1) $\{\rho_n\}$ is convergent on F with limit representations $\rho^i \in \mathcal{D}(\pi_1(F_i))$
- (2) $\rho^i(g)$ is parabolic if and only if g is peripheral in $\pi_1(F_i)$, and
- (3) $\{\rho_n(\pi_1(S))\}$ converges geometrically to Γ .

Then letting $\widehat{N} = \mathbb{H}^3/\Gamma$, there exists an anchored embedding

$$h : F \rightarrow \widehat{N}$$

such that $h|_{F_i}$ is in the homotopy class determined by ρ^i for each F_i .

In the proof of Proposition 6.3 we will need to consider separately the components F_i for which $\rho^i(\pi_1(F_i))$ is geometrically finite, and those for which it is geometrically infinite. The geometrically finite subsurfaces will be handled using (relative versions of) machinery developed by Anderson-Canary-Culler-Shalen [6] and Anderson-Canary [5], while the geometrically infinite subsurfaces will be handled using Thurston's covering theorem.

6.1.1. The limit set machine. We first establish Proposition 6.3 when the algebraic limits are quasifuchsian:

Proposition 6.4. *Let S be an orientable surface and let F be an essential subsurface with components $\{F_i\}$, none of which are annuli. Suppose that $\{\rho_n\}$ is a sequence in $\mathcal{D}(\pi_1(S))$ such that*

- (1) $\{\rho_n\}$ is convergent on F with limit representations $\rho^i \in \mathcal{D}(\pi_1(F_i))$
- (2) ρ^i is a quasifuchsian representation of F_i for all i , and
- (3) $\{\rho_n(\pi_1(S))\}$ converges geometrically to Γ .

Then letting $\widehat{N} = \mathbb{H}^3/\Gamma$, there exists an anchored embedding

$$h : F \rightarrow \widehat{N}$$

such that $h|_{F_i}$ is in the homotopy class determined by ρ^i for each F_i .

Let us give an outline of the proof of Proposition 6.4. (The actual proof proceeds in the opposite order to the outline.) Since limit sets of quasifuchsian groups are Jordan curves, any essential intersection of F_i and F_j (or essential self-intersection of F_i) would result in limit sets of conjugates of $\rho^i(\pi_1(F_i))$ and $\rho^j(\pi_1(F_j))$ which cross (see Lemma 6.8). A result of Susskind [64], see Theorem 6.5, implies that the intersection of the limit sets of two geometrically finite subgroups Φ_1 and Φ_2 of a Kleinian group consists of the limit set of their intersection $\Phi_1 \cap \Phi_2$ along with certain parabolic fixed points $P(\Phi_1, \Phi_2)$. Therefore, it suffices to prove that the intersection of

$\rho^i(\pi_1(F_i))$ and a distinct conjugate of $\rho^j(\pi_1(F_j))$ has at most one point in its limit set (Lemma 6.6) and that there are no problematic parabolic fixed points (see Proposition 6.7.)

We first recall Susskind's result which describes the intersections of the limit sets of two geometrically finite subgroups of a Kleinian group. (Soma [59] and Anderson [8] have generalized Susskind's result to allow the subgroups to be topologically tame.) Given a pair Θ and Θ' of subgroups of a Kleinian group Γ , let $P(\Theta, \Theta')$ be the set of points $x \in \Lambda(\Gamma)$ such that the stabilizers of x in Θ and Θ' are rank one parabolic subgroups which generate a rank two parabolic subgroup of Γ .

Theorem 6.5. (Susskind [64]) *Let Γ be a Kleinian group and let Φ_1 and Φ_2 be nonelementary, geometrically finite subgroups of Γ . Then,*

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

We next show that the intersection of $\rho^i(\pi_1(F_i))$ and a distinct conjugate of $\rho^j(\pi_1(F_j))$ has at most one point in its limit set. This generalizes Lemma 2.4 from [6], in the setting of surface groups.

Lemma 6.6. *Let $\{\rho_n\}$ be a sequence in $\mathcal{D}(S)$ which is convergent on an essential subsurface F , with nonannular components F_i and limit representations ρ^i . Suppose that $\{\rho_n(\pi_1(S))\}$ converges geometrically to Γ , and that $\{\ell_{\rho_n}(\partial F)\}$ converges to 0.*

If $\gamma \in \Gamma$ and either $i \neq j$ or $i = j$ and $\gamma \notin \rho^i(\pi_1(F_i))$, then

$$\gamma \rho^i(\pi_1(F_i)) \gamma^{-1} \cap \rho^j(\pi_1(F_j))$$

is purely parabolic.

Proof. Let $\{\sigma_n^i\}$ be the sequences of maps, as in Definition 6.2, such that $\{\rho_n^i\} = \{\rho_n \circ \sigma_n^i\}$ converges to ρ^i for each i .

Let $\gamma \in \Gamma$ and suppose that $\{\rho_n(h_n)\}$ converges to γ . Suppose that $\alpha \in \gamma \rho^i(\pi_1(F_i)) \gamma^{-1} \cap \rho^j(\pi_1(F_j))$ is non-trivial. Then there exist nontrivial $a \in \pi_1(F_i)$ and $b \in \pi_1(F_j)$ such that

$$\rho^j(b) = \gamma \rho^i(a) \gamma^{-1} = \alpha. \quad (6.2)$$

Our goal is to prove that α must be parabolic.

Since $\{\rho_n(\sigma_n^j(b))\}$ and $\{\rho_n(h_n \sigma_n^i(a) h_n^{-1})\}$ both converge to α , Proposition 2.8 (part 1) implies that

$$\sigma_n^j(b) = h_n \sigma_n^i(a) h_n^{-1}$$

for all sufficiently large n .

In particular a and b represent the same free homotopy class in S . If $i \neq j$, a and b must represent boundary components of F_i and F_j that are freely homotopic to each other, and since we have assumed $\{\ell_{\rho_n}(\partial F)\}$ converge to 0, we conclude that α is parabolic.

If $i = j$, we may re-mark the sequence $\{\rho_n|_{\pi_1(F_i)}\}$ by precomposing with σ_n^i , so that from now on we may fix an inclusion of $\pi_1(F_i)$ in $\pi_1(S)$, and set

$\sigma_n^i = id$. After dropping finitely many terms from the sequence, we have

$$b = h_n a h_n^{-1} \quad (6.3)$$

for all n . If $h_n \notin \pi_1(F_i)$ for some n , then a and b must represent boundary components of F_i that are homotopic in the complement of F_i . Again, since $\{\ell_{\rho_n}(\partial F)\}$ converges to 0, we may conclude that α is parabolic.

Thus, suppose that $h_n \in \pi_1(F_i)$ for all n . We will show that, if α is not parabolic, then h_n is eventually constant. Equation (6.3) implies that $h_m h_n^{-1}$ centralizes b for all m, n . Letting $m = 1$, applying ρ_n , and taking a limit as $n \rightarrow \infty$, we find that $\rho^i(h_1)\gamma^{-1}$ centralizes $\rho^i(b)$. Since we are assuming that $\rho^i(b)$ is hyperbolic, its centralizer in Γ is infinite cyclic, so there exist non-zero integers k and l such that

$$(\rho^i(h_1)\gamma^{-1})^k = \rho^i(b)^l.$$

Again Proposition 2.8 (part 1) implies that $(h_1 h_n^{-1})^k = b^l$ for all large enough n . Since elements of torsion-free Kleinian groups have unique roots, we conclude that $\{h_1 h_n^{-1}\}$, and hence $\{h_n\}$, is eventually constant. Therefore $\gamma = \lim \rho_n(h_n)$ lies in $\rho^i(\pi_1(F_i))$, which contradicts the hypotheses of the lemma. We conclude that α must be parabolic. \square

In order to show that the limit set of $\rho^i(\pi_1(F_i))$ and a distinct conjugate of $\rho^j(\pi_1(F_j))$ do not cross it remains to check that there are no problematic parabolic fixed points. Our proof generalizes the argument of Proposition 2.7 of [6].

Proposition 6.7. *Let $\{\rho_n\}$ be a sequence in $\mathcal{D}(S)$ which is convergent on an essential subsurface F , with nonannular components F_i and quasifuchsian limit representations ρ^i . Suppose also that $\{\rho_n(\pi_1(S))\}$ converges geometrically to Γ .*

If $\gamma \in \Gamma$ and either $i \neq j$ or $i = j$ and $\gamma \notin \rho^i(\pi_1(F_i))$, then the intersection of limit sets

$$\Lambda(\gamma \rho^i(\pi_1(F_i))\gamma^{-1}) \cap \Lambda(\rho^j(\pi_1(F_j)))$$

contains at most one point.

If one makes use of Soma and Anderson's generalization of Susskind's result and Bonahon's tameness theorem, one may replace the assumption in Proposition 6.7 that the limit representations are quasifuchsian with the weaker assumption that $\{\ell_{\rho_n}(\partial F)\}$ converges to 0.

Proof. The hypothesis that the ρ^i are quasifuchsian representations of F_i implies in particular that the lengths $\{\ell_{\rho_n}(\partial F)\}$ converge to 0, and hence we may apply Lemma 6.6.

Fixing γ , i and j , let $\Phi_1 = \gamma \rho^i(\pi_1(F_i))\gamma^{-1}$, and $\Phi_2 = \rho^j(\pi_1(F_j))$. Lemma 6.6 implies that $\Phi_1 \cap \Phi_2$ is a purely parabolic subgroup, and so has at most 1 limit point (0 if it is trivial). Thus, the proposition follows from Theorem 6.5 once we establish that $P(\Phi_1, \Phi_2) = \emptyset$.

Let $\{\rho_n(h_n)\}$ be a sequence converging to γ , and suppose there is a point $x \in P(\Phi_1, \Phi_2)$. The stabilizer $\text{stab}_{\Phi_2}(x)$ is generated by some $\rho^j(b)$, and

$\text{stab}_{\Phi_1}(x)$ is generated by $\gamma\rho^i(a)\gamma^{-1}$ where a and b are primitive elements of $\pi_1(F_i)$ and $\pi_1(F_j)$, respectively. Since these two elements must commute, Proposition 2.8 implies that $h_n\sigma_n^i(a)h_n^{-1}$ commutes with $\sigma_n^j(b)$ for sufficiently large n . (Here σ_n^i are as in the proof of Lemma 6.6). Since a and b are primitive and all abelian subgroups of $\pi_1(S)$ are cyclic,

$$h_n\sigma_n^i(a)h_n^{-1} = (\sigma_n^j(b))^{\pm 1}$$

for sufficiently large n . Applying ρ_n , and taking a limit we conclude that

$$\gamma\rho^i(a)\gamma^{-1} = \rho^j(b)^{\pm 1}$$

but this contradicts the assumption that $\gamma\rho^i(a)\gamma^{-1}$ and $\rho^j(b)$ generate a rank 2 group. Thus $P(\Phi_1, \Phi_2)$ must be empty and the proposition follows. \square

In order to convert these conclusions about limit sets to conclusions about embedded surfaces, let us recall from [5] that a collection Φ_1, \dots, Φ_n of non-conjugate quasifuchsian subgroups of a Kleinian group Γ is called *precisely embedded* if $\text{stab}_{\Gamma}(\Lambda(\Phi_i)) = \Phi_i$ for each i , and if every translate of $\Lambda(\Phi_i)$ by an element of Γ is contained in the closure of a component of $\widehat{\mathbb{C}} \setminus \Lambda(\Phi_j)$, for each i and j .

A system of *spanning disks* $\{D_1, \dots, D_n\}$ for $\{\Phi_i\}$ are disks properly embedded in $\mathbb{H}^3 \cup \widehat{\mathbb{C}}$ such that $\partial D_i = \Lambda(\Phi_i)$, and any two translates by Γ of these disks either coincide or are disjoint. Thus, such disks would project in \mathbb{H}^3/Γ to embedded, disjoint surfaces D_i/Φ_i .

Anderson and Canary observe, in Lemma 6.3 of [5] and the remark that follows (p. 766), that

Lemma 6.8. *Any precisely embedded system $\{\Phi_1, \dots, \Phi_n\}$ of quasifuchsian subgroups of a Kleinian group Γ admits a system of spanning disks $\{D_1, \dots, D_n\}$. Furthermore, one may choose the spanning disks so that there exists $\epsilon > 0$ such that each component of the intersection of any embedded surfaces D_i/Φ_i with a non-compact component of \mathbb{T}_{ϵ} is a properly embedded, totally geodesic half-open annulus.*

One may either prove Lemma 6.8 using the traditional methods of cut-and-paste topology, see Theorem VII.B.16 in Maskit [39], or by using the least area surface techniques of Freedman-Hass-Scott [23] and Ruberman [57].

We are now ready to complete the proof of the embedding theorem in the quasifuchsian case.

Proof of Proposition 6.4. Let S , F , and $\{\rho_n\}$ be as in the statement of the proposition. Let Γ be the geometric limit of $\{\rho_n(\pi_1(S))\}$ and consider the quasifuchsian limit representations $\rho^i : \pi_1(F_i) \rightarrow \Gamma$.

Proposition 6.7 implies that the limit sets of any two distinct conjugates of $\rho^i(\pi_1(F_i))$ and $\rho^j(\pi_1(F_j))$ for components F_i and F_j are disjoint, or intersect in exactly one point. These limit sets are all Jordan curves since the groups are quasifuchsian, and thus any one of them is contained in the

closure of a complementary disk of any other. The conclusion of Proposition 6.7, applied to the conjugates of a single group $\rho^i(\pi_1(F_i))$, imply that $\text{stab}_\Gamma(\Lambda(\rho^i(\pi_1(F_i))))$ must be $\rho^i(\pi_1(F_i))$ itself. Thus $\{\rho^i(\pi_1(F_i))\}$ form a precisely embedded system of quasifuchsian groups in Γ , and we can apply Lemma 6.8 to obtain a system of spanning disks D_i for these groups. Note that the quotients $D_i/\rho^i(\pi_1(F_i))$ may be identified with $\text{int}(F_i)$, so that the resulting embeddings $h_i : \text{int}(F_i) \rightarrow \widehat{N}$ are disjoint, are in the homotopy classes determined by ρ^i , and so that there exists $\epsilon > 0$ so that intersection of $h_i(F_i)$ with each component of $\mathbb{T}_\epsilon(h_{i*}([\partial F_i]))$ is a properly embedded totally geodesic half-open annulus. Therefore, truncating the maps at $\mathbb{T}_\epsilon(h_{i*}([\partial F_i]))$ yields embeddings that are anchored with respect to the ϵ -Margulis tubes, whose domains can be identified with the compact surfaces F_i .

Finally, composing with a diffeomorphism of \widehat{N} that takes the ϵ -Margulis tubes $\mathbb{T}_\epsilon(h_{i*}([\partial F_i]))$ to the ϵ_1 -Margulis tubes $\mathbb{T}(h_{i*}([\partial F_i]))$ and is homotopic to the identity, we obtain the desired anchored embedding h . This concludes the proof of Proposition 6.4.

6.1.2. Using the covering theorem. We now consider the case where the algebraic limits are geometrically infinite. The main statement we need is the following, whose proof is an application of Thurston's Covering Theorem.

Proposition 6.9. *Let S be an orientable surface and let R be a connected essential non-annular subsurface. Let $\{\rho_n\}$ be a sequence in $\mathcal{D}(S)$ that is convergent on R with limit representation $\hat{\rho} : \pi_1(R) \rightarrow \text{PSL}_2(\mathbb{C})$, and suppose that $\{\rho_n(\pi_1(S))\}$ converges geometrically to Γ .*

Suppose that

- (1) $\hat{\rho}(g)$ is parabolic if and only if g represents a boundary component of R , and
- (2) $\hat{\rho}(\pi_1(R))$ is geometrically infinite.

If K is a compact subset of $\widehat{N} = \mathbb{H}^3/\Gamma$, then there exists an anchored embedding $h : R \rightarrow N$ in the homotopy class determined by $\hat{\rho}$, whose image does not intersect K .

Proof. The following generalization of Thurston's covering theorem [66] is established in [20].

Theorem 6.10. *Let N be a topologically tame hyperbolic 3-manifold which covers an infinite volume hyperbolic 3-manifold \widehat{N} by a local isometry $\pi : N \rightarrow \widehat{N}$. If E is a geometrically infinite end of N^0 , then E has a neighborhood U such that π is finite-to-one on U .*

(Here we recall that N^0 denotes N minus the ϵ_1 -Margulis tubes of its cusps).

Since R has only one component, after remarking $\{\rho_n\}$ by a sequence of inner automorphisms we may assume that $\{\rho_n|_{\pi_1(R)}\}$ converges to $\hat{\rho}$.

Let $N = \mathbb{H}^3/\hat{\rho}(\pi_1(R))$. By the assumptions and Bonahon's theorem, N^0 may be identified with $R \times \mathbb{R}$, and has a geometrically infinite end E . Let $\pi : N \rightarrow \hat{N}$ be the covering map associated to the inclusion $\hat{\rho}(\pi_1(R)) \subset \Gamma$. Note that \hat{N} has infinite volume since it is the geometric limit of infinite-volume hyperbolic 3-manifolds. Theorem 6.10 then implies that there exists a neighborhood of E , on which π is finite-to-one.

Suppose that there does not exist a neighborhood of E on which π is one-to-one. An argument in Proposition 5.2 of [5] then implies that there exists a primitive element $\alpha \in \hat{\rho}(\pi_1(R))$ which is a k -th power of some $\gamma \in \Gamma$, with $k > 1$. Let $\alpha = \hat{\rho}(a)$ and $\gamma = \lim \rho_n(g_n)$ for $\{g_n \in \pi_1(S)\}$. By Lemma 2.8, for large enough n we must have $a = g_n^k$. However, a primitive element of $\pi_1(R)$ must also be primitive in $\pi_1(S)$, and this is a contradiction.

Thus there is a neighborhood U of E on which π is an embedding, and hence there is a $t \in \mathbb{R}$ such that $R \times \{t\} \subset U$, and $\pi(R \times \{t\})$ avoids K . This gives our desired anchored embedding. \square

6.1.3. Proof of Proposition 6.3. The proof of the limit embedding theorem in the general case now follows from Propositions 6.4 and 6.9. Let F be the surface on which $\{\rho_n\}$ converges in the sense of Definition 6.2. The assumption that $\rho^i(g)$ is parabolic if and only if g is peripheral in $\pi_1(F_i)$ implies that $\rho^i(\pi_1(F_i))$ is either quasifuchsian or geometrically infinite with no non-peripheral parabolics. Let $F' \subset F$ be the union of the quasifuchsian components. Proposition 6.4 gives us an anchored embedding $h' : F' \rightarrow \hat{N}$. Enumerate the components of $F \setminus F'$ as F_1, \dots, F_k . Let $K_0 = h'(F')$. Applying Proposition 6.9 gives us an anchored embedding $h_1 : F_1 \rightarrow \hat{N}$ avoiding K_0 . Now inductively define $K_i = K_0 \cup \bigcup_{j \leq i} h_j(F_j)$, and apply Proposition 6.9 to obtain h_{i+1} avoiding K_i . The union of maps h', h_1, \dots, h_k is the desired anchored embedding of F .

6.2. Proof of Theorem 6.1

Proof. Fix S, L, δ_1 and δ_2 , and suppose by way of contradiction that the theorem fails. Then there is a sequence $\{(\rho_n, R_n, f_n, \Gamma_n, \epsilon_n, \hat{\epsilon}_n, K_n)\}$ with $\epsilon_n \rightarrow 0$, $\hat{\epsilon}_n \rightarrow 0$, and $K_n \rightarrow \infty$ for which the hypotheses of the theorem hold, but for which the conclusions fail.

Possibly precomposing ρ_n and f_n with a sequence of homeomorphisms of S and passing to a subsequence, we may assume that all R_n are equal to a fixed surface R , and Γ_n are equal to a fixed curve system Γ . The surface $X = R \setminus \mathbf{collar}(\Gamma)$ is equipped with a sequence of L -bounded metrics, which we may assume (again after remarking and passing to a subsequence) converge to an L -bounded metric ν . Then, possibly adjusting L slightly, we may assume that f_n is L -bounded with respect to ν for each n . For each n we have an extension $\tilde{f}_n : R \rightarrow N_n$ with the properties given in the statement of Theorem 6.1, notably the unwrapping condition. The failure

of the conclusions means that there is no homotopy

$$H_n : (X \times [0, 1], \partial X \times [0, 1]) \rightarrow (N_n - \mathbb{T}(\partial X), \partial \mathbb{T}(\partial X))$$

with $H_n(\cdot, 0) = f_n$, such that

- (1) H_n is K_n -Lipschitz,
- (2) H_n is a smooth K_n -bilipschitz embedding of $X \times [1/2, 1]$ with the norm of the second derivatives bounded by K_n ,
- (3) For all $t \in [1/2, 1]$, $H(\partial X \times \{t\})$ is a collection of geodesic circles in $\partial \mathbb{T}(\partial X)$.
- (4) $H_n(X \times [1/2, 1])$ does not intersect any ϵ_1 -Margulis tubes with core length less than $\hat{\epsilon}_n$, and
- (5) $H_n(X \times [0, 1])$ avoids the $\hat{\epsilon}_n$ -thin part of N_n ,
- (6) $H_n(X \times [0, 1])$ lies in $C_{N_n}^{1/2} \cup (N_n)_{(0, \epsilon_2]}$, and
- (7) each component of $(H_n)_1(\partial X)$ is a geodesic in the intrinsic metric on $\partial \mathbb{T}(\partial X)$.

Possibly restricting to a subsequence again, we may assume that for each $\gamma \in \Gamma$ the lengths $\{\ell_{\rho_n}(\gamma)\}$ converge. Let Γ' be the set of $\gamma \in \Gamma$ whose lengths $\{\ell_{\rho_n}(\gamma)\}$ converge to 0, and let $\Delta = \Gamma \setminus \Gamma'$. It will be convenient to suppose that in the metric ν , all boundary components of X have the same length. This can be done by changing ν by a bilipschitz distortion, and then altering the constant L appropriately.

Let $D = X \cup \mathbf{collar}(\Delta) = R \setminus \mathbf{collar}(\Gamma')$. The metric on X can be extended across $\mathbf{collar}(\Delta)$ to a metric which makes each component isometric to $S^1 \times [0, 1]$ (with S^1 isometric to a component of ∂X). Extend each f_n to a map \hat{f}_n which, on each component of $\mathbf{collar}(\Delta)$, takes the intervals $\{p\} \times [0, 1]$ to geodesics whose maximal length is shortest among all such maps. This length is uniformly bounded above because the ρ_n -lengths of components of Δ converge to positive constants, and hence the distance from the cores of their Margulis tubes to their boundaries is bounded. Notice that $\hat{f}_n(\mathbf{collar}(\Delta)) \subset \mathbb{T}(\Delta, n)$, $\hat{f}_n(D) \subset C_{N_n} \cap N_n^1$, and \hat{f}_n is homotopic to $\bar{f}_n|_D$ by a homotopy supported on $\mathbf{collar}(\Delta)$ whose image lies in $\mathbb{T}(\Delta, n)$. (Here if B is a collection of simple closed curves on R such that if $\beta \in B$ then $\ell_{\rho_n}(\beta) < \epsilon_1$, then $\mathbb{T}(B, n)$ denotes the union of ϵ_1 -Margulis tubes in N_n associated to $\bar{f}_n(B)$.) In particular, \hat{f}_n satisfies the same unwrapping condition as \bar{f}_n .

After remarking the ρ_n by Dehn twists supported on $\mathbf{collar}(\Delta)$, we may assume the extended maps \hat{f}_n are in the homotopy classes determined by ρ_n . We now have a fixed metric on D and a sequence of L' -Lipschitz maps \hat{f}_n for some constant L' .

Outline. Now that we have completed the normalizations and book-keeping we can provide a brief outline of the rest of the proof. The basic strategy is to construct an anchored embedding of X into a geometric limit of $\{N_n\}$ and then pull back to N_n to obtain a contradiction for large values of n . Since the

components of $\widehat{f}_n(D)$ may be pulling apart, we may actually need to consider a collection of geometric limits. We first observe that, after conjugation and passing to a subsequence, $\{\rho_n\}$ is convergent on each component of D . Each component of D is contained in a maximal collection D^J of components of D such that, after conjugation and passage to a subsequence, $\{\rho_n\}$ is convergent on D^J . (These are essentially the components of D whose images under \widehat{f}_n stay a bounded distance from one another.) We then obtain a limiting map $\widehat{f}^J : D^J \rightarrow N^J$ into the appropriate geometric limit of $\{N_n\}$.

We'd like to apply the embedding result from the previous section to \widehat{f}^J , but we can't guarantee that there are no unexpected parabolic elements. Let P^J be a maximal collection of disjoint non-peripheral curves on D^J which are homotopic into cusps of N^J . If we let $F^J = D^J - \mathbf{collar}(P^J)$, then Proposition 6.3 guarantees the existence of an anchored embedding

$$\widehat{h}^J : (F^J, \partial F^J) \rightarrow (N^J \setminus \mathbb{T}(\partial F^J), \partial \mathbb{T}(\partial F^J))$$

which is homotopic to $\widehat{f}^J|_{F^J}$. We then use the unwrapping property and Lemma 2.6 to extend \widehat{h}^J to an embedding \bar{h}^J defined on all of D^J and homotopic to \widehat{f}^J . However, what we want is an anchored embedding of $X^J = D^J \cap X = D^J - \mathbf{collar}(\Delta)$, so we must "reanchor" on Δ . We apply a result of Freedman, Hass, and Scott to produce an embedding $g^J : D^J \rightarrow N^J$ whose image is contained in a regular neighborhood of $\widehat{f}^J(D^J)$ and misses $\mathbb{T}(\Delta)$. Finally, we apply the Annulus Theorem to produce an embedded annulus joining $g^J(D^J)$ to each component of $\mathbb{T}(\Delta)$. The usual surgery argument produces the desired anchored embedding. This anchored embedding has a bilipschitz collar neighborhood and is homotopic to \widehat{f}^J , so we can pull back the resulting homotopy to find, for all large n , a uniform homotopy (with uniform constants) of $f_n|_{X^J}$ to an anchored embedding of X^J into N_n . Since we can do this for each collection D^J , we can combine the resulting uniform homotopies to obtain a contradiction for large values of n .

Geometric limits. Fix a basepoint x^j for each component D^j of D , and let $y_n^j = f_n(x^j)$. Since the metric on D^j is fixed, each \widehat{f}_n is L' -Lipschitz, and $\pi_1(D^j)$ is non-abelian, a standard application of the Margulis lemma gives an uniform lower bound on the injectivity radius of N_n at y_n^j for all n and j . We may therefore pass to a subsequence such that, for any fixed j , the sequence of maps

$$\widehat{f}_n|_{D^j} : (D^j, x^j) \rightarrow (N_n, y_n^j)$$

converges geometrically to a map $\widehat{f}^j : (D^j, x^j) \rightarrow (N^j, y^j)$ and $\{(C(N_n), x^j)\}$ converges to (C_{N^j}, y^j) . After further restriction to a subsequence we may assume that, for each pair (j, j') , the distances $\{d(y_n^j, y_n^{j'})\}$ converge to some $d_{jj'} \in [0, \infty]$. The relation $d_{jj'} < \infty$ is an equivalence relation; fix an equivalence class J .

For $j, j' \in J$, we may identify N^j with $N^{j'}$, naming it N^J . Notice that $d_{N^J}(y^j, y^{j'}) = d_{jj'}$. Let $D^J = \cup_{j \in J} D^j$ and let $\widehat{f}^J : D^J \rightarrow N^J$ denote the union of the maps \widehat{f}^j for all $j \in J$. Since $\widehat{f}_n(D^J) \subset C_{N_n}$ for all n , $\widehat{f}^J(D^J) \subset C_{N^J}$.

Our goal now is to apply Proposition 6.3 to deform \widehat{f}^J to an embedding in N^J . We must first obtain an algebraically convergent sequence in the sense of Definition 6.2.

Fix a point θ^J in N^J and let Θ^J be an embedded tree in N^J formed by joining each y^j to θ^J with an arc. For all large enough n the pullback of Θ^J to N_n may be deformed slightly to give an embedded tree Θ_n^J all of whose edges join the pullback θ_n^J of θ^J to y_n^j for some n . (We must deform slightly since the endpoints of the pullback of Θ^J are only guaranteed to be near to the y_n^j .) Therefore the total length of Θ_n^J is bounded for all large n . Using paths in the tree we obtain homomorphisms

$$\rho_n^j : \pi_1(D^j, x^j) \rightarrow \pi_1(N_n, \theta_n^J)$$

consistent with the maps $\widehat{f}_n|_{D^j}$. (More explicitly, if e_n^j is the edge in Θ_n^J joining θ_n^j to y_n^j , then ρ_n^j takes $[\alpha]$ to $[e_n^j * f(\alpha) * \overline{e_n^j}]$.) Fixing an origin for \mathbb{H}^3 and possibly conjugating ρ_n in $\mathrm{PSL}_2(\mathbb{C})$, we may assume the origin maps to θ_n^J and hence consider ρ_n as an isomorphism from $\pi_1(S)$ to $\pi_1(N_n, \theta_n^J)$. Thus we can define

$$\sigma_n^j : \pi_1(D^j, x^j) \rightarrow \pi_1(S)$$

by $\sigma_n^j = \rho_n^{-1} \circ \rho_n^j$.

Now, $\{\rho_n^j\}$ converges, after restricting to a subsequence, because each $\widehat{f}_n|_{D^j}$ is L -lipschitz and Θ_n has bounded total length, so the images of any fixed element of $\pi_1(D^j)$ are represented by loops of uniformly bounded length, and hence move the origin in \mathbb{H}^3 a uniformly bounded amount.

Thus, repeating for all $j \in J$ and restricting to an appropriate subsequence, $\{\rho_n\}$ converges on the subsurface D^J , in the sense of Definition 6.2, using the maps σ_n^j as defined above. The limiting representation ρ^j , for a component D^j where $j \in J$, corresponds to the homotopy class of the limiting map $\widehat{f}^J|_{D^j}$.

Anchoring on parabolics. These limiting representations may have non-peripheral parabolics. Let P^J denote a maximal set of disjoint homotopically distinct simple closed nonperipheral curves in D^J whose images under the limiting representations are parabolic. Let $F^J = D^J \setminus \mathbf{collar}(P^J)$, and for each component F^i of F^J contained in a component D^j of D^J , fix an injection $\pi_1(F^i) \rightarrow \pi_1(D^j)$ consistent with the inclusion map. Then, with the same $\{\rho_n^j\}$ as before, restricted to $\pi_1(F^i)$, we have convergence of $\{\rho_n\}$ on F^J , and the limiting representations $\widehat{\rho}^i$ (which are the restrictions of ρ^j to $\pi_1(F^i)$) have no nonperipheral parabolics.

We can now apply Proposition 6.3 to F^J , obtaining an anchored embedding

$$\widehat{h}^J : (F^J, \partial F^J) \rightarrow (N^J, \mathbb{T}(\partial F^J))$$

such that, for each component F^i of F^J , $\widehat{h}^J|_{F^i}$ is in the homotopy class determined by $\widehat{\rho}^i$, which is the same as the homotopy class of $\widehat{f}^J|_{F^i}$. Since each component of $\mathbb{T}(\partial F^J)$ is a cusp, it is easy to see that $\widehat{h}^J|_{F^i}$ is properly homotopic to $\widehat{f}^J|_{F^i}$ within $N^J \setminus \mathbb{T}(\partial F^J)$.

Resewing along parabolics. We next want, for each component α of P^J , to add an embedded annulus on $\partial\mathbb{T}(\alpha)$ to the image of \widehat{h}^J , thus obtaining an anchored embedding of D^J which is homotopic to $\widehat{f}^J|_{D^J}$. The unwrapping property of each \bar{f}_n , and hence each \widehat{f}_n , will guarantee the existence of such an annulus.

Notice that since all the curves in P^J are homotopic into cusps of N^J , the unwrapping condition implies in particular that, for all large enough n , the image of \bar{f}_n , and hence of \widehat{f}_n , does not intersect $\mathbb{T}(P^J, n)$. Therefore, the image of \widehat{f}^J does not intersect $\mathbb{T}(P^J)$. Let

$$H : (F^J \times [0, 1], \partial F^J \times [0, 1]) \rightarrow (N^J \setminus \mathbb{T}(\partial F^J), \partial\mathbb{T}(\partial F^J)),$$

be a homotopy with $H_0 = \widehat{f}^J|_{F^J}$ and $H_1 = \widehat{h}^J$.

We now explain how to use H to extend \widehat{h}^J across the annuli $\mathbf{collar}(P^J)$ to obtain a map \bar{h}^J which takes these annuli to $\partial\mathbb{T}(P^J)$.

In the union of solid tori $\mathbf{collar}(P^J) \times [0, 1]$, let φ be a bilipschitz homeomorphism from the top annuli $\mathbf{collar}(P^J) \times \{1\}$ to the remainder of the boundary, $(\mathbf{collar}(P^J) \times \{0\}) \cup (\partial\mathbf{collar}(P^J) \times [0, 1])$, which is the identity on the intersection curves $\partial\mathbf{collar}(P^J) \times \{1\}$ and is homotopic rel boundary to the identity map. Extend H to all of $D^J \times \{0\}$ to be equal to \widehat{f}^J , and then consider the map $H \circ \varphi$ on the annuli $\mathbf{collar}(P^J) \times \{1\}$ which maps into the complement of $\mathbb{T}(P^J)$. On each annulus component A , this map is homotopic rel boundary, in the exterior of $\mathbb{T}(P^J)$, to a unique ‘‘straight’’ map to $\partial\mathbb{T}(P^J)$. Here ‘‘straight’’ means that geodesics orthogonal to the core of the annulus are taken to straight lines in the Euclidean metric of $\partial\mathbb{T}(P^J)$. In particular the map is an immersion. Define \bar{h}^J to be the extension of \widehat{h}^J to $\mathbf{collar}(P^J)$ by this straight map. Use the homotopy between \bar{h}^J and \widehat{f}^J on $\mathbf{collar}(P^J)$ to extend H across the solid tori to a proper (in $N^J \setminus \mathbb{T}(\partial D^J)$) homotopy \bar{H} between \bar{h}^J and \widehat{f}^J which is defined on $D^J \times [0, 1]$ and avoids $\mathbb{T}(P^J)$.

We recall from the definition of geometric limits that there exist C_n -bilipschitz maps $c_n : B_{N^J}(\theta^J, R_n) \rightarrow N_n$ such that $c_n(\theta^J) = \theta_n^J$, $R_n \rightarrow \infty$, $C_n \rightarrow 1$ and, by Lemma 2.7, whenever a component \mathbb{T} of $(N^J)_{thin(\epsilon_1)}$ intersects the domain of c_n , then

$$c_n(\partial\mathbb{T} \cap B_{N^J}(\theta^J, R_n)) \subset \partial(N_n)_{thin(\epsilon_1)}.$$

For large enough n , c_n is defined on a region that includes the image of \overline{H} .

We will now apply Lemma 3.13 to show that \overline{h}^J is an embedding on each component α of $\mathbf{collar}(P^J)$. Let $\overline{h}_n^J = c_n \circ \overline{h}^J$ be the pullback of \overline{h}^J and let $\overline{H}_n = c_n \circ \overline{H}$ be the pull back of \overline{H} . The unwrapping property implies that $\widehat{f}_n(D^J)$ can be pushed to $+\infty$ or $-\infty$ disjointly from $\mathbb{T}(\alpha, n)$. Since $\{\widehat{f}_n\}$ converges geometrically to \widehat{f}^J , for large enough n there is a very short homotopy from \widehat{f}_n to $\overline{H}_n|_{D^J \times \{0\}}$. Since the image of \overline{H}_n does not intersect $\mathbb{T}(P^J, n)$, we conclude that \overline{h}_n^J can also be pushed to either $+\infty$ or $-\infty$ in $N_n^0 \setminus \mathbb{T}(\alpha, n)$.

Since \overline{h}_n^J is δ_1 -anchored and $\delta_1 \leq \epsilon_u$, Otal's Theorem 2.5 implies that $\mathbb{T}(\partial F^J, n)$ is unknotted and unlinked in N_n^0 . Therefore, we can apply Lemma 3.13 to the restriction of \overline{h}_n^J to the union of $\mathbf{collar}(\alpha)$ with the components of F^J which are adjacent to it, concluding that \overline{h}_n^J restricted to $\mathbf{collar}(\alpha)$ is an embedding into $\partial\mathbb{T}(\alpha, n)$. Thus in the geometric limit $\overline{h}^J|_{\mathbf{collar}(\alpha)}$ is an embedding into $\partial\mathbb{T}(\alpha)$. Applying this to all components of P^J we conclude that the map \overline{h}^J is an embedding into N^J .

Reanchoring on Δ . The embedding \overline{h}^J is defined on D^J , whereas we need an anchored embedding whose domain is

$$X^J = X \cap D^J = D^J \setminus \mathbf{collar}(\Delta).$$

Restricting \overline{h}^J to X^J is not sufficient, since it would not be anchored on $\mathbb{T}(\Delta)$.

Thus, consider again the map \widehat{f}^J , which meets $\mathbb{T}(\Delta)$ only in the embedded annuli $\widehat{f}^J(\mathbf{collar}(\Delta))$. Deform these intersection annuli to the boundary of a small regular neighborhood of $\mathbb{T}(\Delta)$, obtaining a map $\underline{f}^J : D^J \rightarrow N^J$ which misses $\mathbb{T}(\Delta)$ and is properly homotopic to \overline{h}^J within $N^J \setminus \mathbb{T}(\partial X^J)$.

Let Z be a compact, irreducible submanifold of $N^J \setminus \mathbb{T}(\partial X^J)$ which contains a regular neighborhood Y of \underline{f}^J , within $N^J \setminus \mathbb{T}(\partial X^J \cup \Delta)$, such that \underline{f}^J is homotopic to \overline{h}^J within Z . In section 7 of Freedman-Hass-Scott [23], it is shown that a least area surface in Z in the proper homotopy class associated to \underline{f}^J is embedded in any metric on Z where the boundary has non-negative mean curvature. Corollary 3 of Jaco-Rubinstein [28] establishes the same result in the PL-setting. In a remark at the end of section 2 of Canary-Minsky [18], it is explained how one may choose a PL-metric on Z so that the least area surface in the proper homotopy class of \underline{f}^J lies within Y . Therefore, there exists an anchored embedding $g^J : D^J \rightarrow N^J$ which is properly homotopic to \underline{f}^J whose image lies in Y . In particular, $g^J(D^J)$ misses $\mathbb{T}(\Delta)$. (Notice that we can't simply obtain g^J by naively pushing \overline{h}^J off of $\mathbb{T}(\Delta)$, since we have no a priori control over the intersection of $\overline{h}^J(D^J)$ with $\mathbb{T}(\Delta)$.)

Let $g_n^J = c_n \circ g^J$ be the pullback of g^J by the comparison maps to N_n (defined for n sufficiently large).

We claim that for each core β of a component of Δ there is a homotopy from $g_n^J(\beta)$ to an embedded longitudinal curve on $\partial\mathbb{T}(\beta, n)$ that avoids $\mathbb{T}(\Delta \cup \partial D^J, n)$. Notice first, that since f_n is δ_1 -anchored and $\delta_1 \leq \epsilon_u$, Otal's Theorem 2.5 implies that $\mathbb{T}(\partial X) = \mathbb{T}(\Gamma \cup \partial R)$ are unknotted and unlinked in N_n^1 . Hence N_n^0 can be identified with $S \times \mathbb{R}$ in such a way that $\beta_n^* = \beta \times \{0\}$, $\mathbb{T}(\beta) = \text{collar}(\beta) \times [a, b]$ and $B = \beta \times \mathbb{R}$ is disjoint from $\mathbb{T}(\Delta - \{\beta\}, n) \cup \mathbb{T}(\partial D^J, n)$, and in particular from $g_n^J(\partial D^J)$ since g_n^J is anchored. Since β is non-peripheral in the essential surface D^J and g_n^J is in the homotopy class determined by ρ , $g_n^J(D^J)$ must intersect B . After proper isotopy of B we may assume the intersections of B with $g_n^J(D^J)$ are essential circles, and so the closest one to $\beta \times \{0\}$ yields the desired homotopy.

In order to use c_n to transport this homotopy to N^J we must first bound its diameter. As the bilipschitz constants of c_n converge to 1, $\ell_{N_n}(g_n^J(\beta)) \leq C$ for some uniform constant C . By Lemma 2.6, there is a homotopy from $g_n^J(\beta)$ to $\partial\mathbb{T}(\beta, n)$ which avoids $\mathbb{T}(\Delta \cup \partial D^J, n)$ and lies in an $r(C)$ -neighborhood of $g_n^J(D^J)$.

Now for n large enough, this $r(C)$ -neighborhood lies in the image of c_n , so $g^J(\beta)$ is homotopic to $\partial\mathbb{T}(\beta)$ in $N^J \setminus (\mathbb{T}(\Delta \cup \partial D^J))$. Denote this homotopy by Q_β . We will next apply a version of the Annulus Theorem to conclude that there is an *embedded* annulus \widehat{Q}_β in the complement of $\mathbb{T}(\Delta \cup \partial D^J)$ joining $g^J(\beta)$ to $\partial\mathbb{T}(\beta)$, whose interior is disjoint from $g^J(D^J)$.

Since $g^J(D^J)$ is embedded, $Q_\beta^{-1}(g^J(D^J))$ is a union of embedded curves in the domain annulus. The inessential ones may be removed by a homotopy, and the remainder are isotopic to the boundary. Hence by restricting to a complementary component of the remaining curves of intersection we obtain a new homotopy which meets $g^J(D^J)$ only on a boundary curve. The image of this curve may not be embedded in $g^J(D^J)$, but since it is homotopic to β we may deform it in $g^J(D^J)$ to a simple curve. Shifting this deformation slightly away from $g^J(D^J)$ we obtain a new homotopy Q'_β which meets $g^J(D^J)$ in a simple curve.

Let Z' be a compact, irreducible submanifold of $N^J \setminus \mathbb{T}(\Delta \cup \partial D^J)$ which contains the $2r(C)$ neighborhood of $g^J(D^J)$. Remove from Z' a regular neighborhood Y' of $g^J(D^J)$, to obtain a Haken manifold W . If Y' is chosen small enough then $Q'_\beta \cap W$ is a proper singular annulus with one boundary embedded in $\partial Y'$ and the other in $\partial\mathbb{T}(\beta)$. Now we may apply the Annulus Theorem (see [30] and [29, Thm VIII.13]) in $Z' - Y'$ to conclude that there is an embedded annulus \widehat{Q}_β joining $g^J(\beta)$ to $\partial\mathbb{T}(\beta)$ whose interior avoids $g^J(D^J)$ and $\mathbb{T}(\Delta \cup \partial D^J)$.

Repeat this for every component of Δ . The resulting embedded annuli may intersect but only in inessential curves, so after surgery we obtain a union \widehat{Q}_Δ of embedded annuli.

A surgery using a regular neighborhood of \widehat{Q}_Δ then yields a smooth embedding $\widehat{g}^J : X^J \rightarrow N^J$ which is anchored on $\mathbb{T}(\partial X^J)$, and is homotopic to $f^J = \widehat{f}^J|_{X^J}$ via a homotopy of pairs

$$H^J : (X^J \times [0, 1], \partial X^J \times [0, 1]) \rightarrow (N^J \setminus \mathbb{T}(\partial X^J), \partial \mathbb{T}(\partial X^J))$$

such that $H_0^J = f^J$. Since N^J is homeomorphic to $\text{int}(C_{N^J}^{1/2}) \cup N_{(0, \epsilon_2)}^J$, by a homeomorphism which is the identity on $C_{N^J}^{1/4} \cup N_{(0, \epsilon_1)}^J$, we may assume that both \widehat{g}^J and the homotopy H^J to f^J lie entirely in $\text{int}(C_{N^J}^{1/2}) \cup N_{(0, \epsilon_2)}^J$. We may further assume that the restriction of H^J to $X^J \times [1/2, 1]$ is an C^2 -embedding and that for all $t \in [1/2, 1]$, $H^J(\partial X^J \times \{t\})$ is a collection of geodesic circles in $\partial \mathbb{T}(\partial X^J)$.

Obtaining the contradiction. As its image is compact, the homotopy H^J between \widehat{g}^J and f^J avoids the $\hat{\epsilon}^J$ -thin part for some $\hat{\epsilon}^J > 0$. Let $\{\mathbb{T}_1, \dots, \mathbb{T}_n\}$ be the components of the ϵ_1 -thin part which $\widehat{g}^J(X^J)$ intersects and which are either cusps or have core curves of length less than $\hat{\epsilon}^J$. Notice that no \mathbb{T}_i is a component of $\mathbb{T}(\partial X^J)$. For each i , there is a regular neighborhood \mathcal{U}_i of \mathbb{T}_i which is contained in $N_{(0, \epsilon_2)}^J$ and a diffeomorphism $\Upsilon_i : \mathcal{U}_i \setminus (\mathbb{T}_i \cap N_{(0, \hat{\epsilon}^J)}^J) \rightarrow \mathcal{U}_i \setminus \mathbb{T}_i$ which is the identity on $\partial \mathcal{U}_i$. Extend the collection of Υ_i , via the identity outside $\cup \mathbb{T}_i$, to an embedding $\Upsilon : N_{[\hat{\epsilon}^J, \infty)} \rightarrow N$. We may replace H^J with $\Upsilon \circ H^J$ which has the additional property of avoiding ϵ_1 -Margulis tubes with core length less than $\hat{\epsilon}^J$.

Pulling H^J back to N_n and deforming slightly, we obtain, for all large enough n , a $(2K^J, \frac{\hat{\epsilon}^J}{2})$ -uniform homotopy H_n^J from $f_n|_{X^J}$ to an anchored embedding. Condition (3) in the definition, that the images of the boundary circles $\partial X \times \{t\}$ are geodesic in $\partial \mathbb{T}(\partial X)$ for $t \in [1/2, 1]$, may be obtained by noticing that, because the C^2 bounds on the comparison maps converge to 0, for large enough n the images are nearly geodesic circles and hence can be expressed as graphs of nearly constant functions over the geodesic circles to which they are homotopic. Hence the map can be adjusted to satisfy condition (3).

For all large enough n , we may also assume the resulting homotopy H_n^J lies within $(C_{N_n}^{1/2} \cup N_{(0, \epsilon_2)}^n) \setminus \mathbb{T}(\partial X^J, n)$.

For each equivalence class J' , we obtain a sequence of homotopies $H_n^{J'}$ in the same way. Since the distance $d_{jj'}$ between basepoints converges to ∞ if $j \in J, j' \in J'$, these maps eventually have disjoint images. Combining we obtain for all large n a $(K, \hat{\epsilon})$ -uniform homotopy H_n which shows that f_n is $(K, \hat{\epsilon})$ -uniformly embeddable, where $K = \max\{2K^J\}$ and $\hat{\epsilon} = \min\{\frac{\hat{\epsilon}^J}{2}\}$. If n is chosen large enough that $K_n > K$ and $\hat{\epsilon}_n < \hat{\epsilon}$, then we have obtained a contradiction. \square

7. Insulating regions

In this section we will establish the existence of long ‘bounded-geometry product regions’ in the hyperbolic manifold N when the hierarchy satisfies certain conditions. Roughly, if H contains a very long geodesic h supported in some non-annular domain R , and if there are no very long geodesics subordinate to h , then there is a big region in N isotopic to $R \times [0, 1]$, whose geometry is prescribed by h . Furthermore the model map can be adjusted to be an embedding on this region, without disturbing too much the structure outside of it. In order to quantify this more carefully, let us make the following definition:

A segment γ of a geodesic $h \in H$ is said to be (k_1, k_2) -thick, where $0 < k_1 < k_2$, provided:

- (1) $|\gamma| > k_2$
- (2) For any $m \in H$ with $D(m) \subset D(h)$ and $\phi_h(D(m)) \cap \gamma \neq \emptyset$, $|m| < k_1$.

Let τ_1 and τ_2 be two (full) slices with the same bottom geodesic h , and suppose that the bottom simplices v_{τ_1} and v_{τ_2} are spaced by at least 5, and $v_{\tau_1} < v_{\tau_2}$. As in §5.3, there is a region $\mathcal{B}(\tau_1, \tau_2) \subset M_\nu$, homeomorphic to $D(h) \times [0, 1]$ and bounded by \hat{F}_{τ_1} and \hat{F}_{τ_2} and the tori $U(\partial D(h))$. It is the geometry of the model map on such regions that we will control.

Theorem 7.1. *Fix a surface S . Given positive constants K, k, k_1 and Q , there exist k_2 and L such that, if $f : M_\nu \rightarrow N$ is a (K, k) model map, γ is a (k_1, k_2) -thick segment of $h \in H_\nu$ and $\xi(h) \geq 4$, then there exist slices $\tau_{-2}, \tau_{-1}, \tau_0, \tau_1, \tau_2$ with bottom geodesic h and bottom simplices v_{τ_i} in γ satisfying*

$$v_{\tau_{-2}} < v_{\tau_{-1}} < v_{\tau_0} < v_{\tau_1} < v_{\tau_2},$$

with spacing of at least 5 between successive simplices, and so that

- (1) f can be deformed, by a homotopy supported on the union of $\mathcal{B}_2 = \mathcal{B}(\tau_{-2}, \tau_2)$ and the tubes $U(\partial D(h))$, to an L -Lipschitz map f' which is an orientation-preserving embedding on $\mathcal{B}_1 = \mathcal{B}(\tau_{-1}, \tau_1)$, and
- (2) f' takes $M_\nu \setminus \mathcal{B}_1$ to $N \setminus f'(\mathcal{B}_1)$, and
- (3) the distance from $f'(F_{\tau_0})$ to $f'(F_{\tau_{-1}})$ and $f'(F_{\tau_1})$ is at least Q .

Proof. Suppose, by way of contradiction, that the theorem fails. Then there exist K, k, Q and k_1 and sequences $k_n \rightarrow \infty$ and $L_n \rightarrow \infty$, representations $\rho_n \in \mathcal{D}(S)$ with associated hierarchies H_n , (K, k) model maps $f_n : M_{\nu_n} \rightarrow N_n$, and (k_1, k_n) -thick segments $\gamma_n \subset h_n \in H_n$, but for which the model maps f_n cannot be deformed to an L_n -Lipschitz map satisfying the conclusions of the theorem.

We will extract and study a certain geometric limit in order to obtain a contradiction.

Convergence of hierarchies. A sequence of tight geodesic segments γ_n in $\mathcal{C}(R)$ converges to a sequence of simplices h_∞ if (after appropriately indexing the simplices $(\gamma_n)_i$ of γ_n and $(h_\infty)_i$ of h_∞) we have for each i that $(h_\infty)_i$ is

defined if and only if, for sufficiently high n , $(\gamma_n)_i$ is defined and equal to $(h_\infty)_i$. (This is equivalent to the notions used in [41, §6.5] and [47, §5.5].) Note that h_∞ is automatically a tight geodesic.

After passing to a subsequence and remarking, we can assume $D(h_n)$ is a constant surface R . We will show that, after passing to a further subsequence, we can make the sequence $\{\gamma_n\}$ converge in $\mathcal{C}(R)$ to a bi-infinite geodesic h_∞ . The geodesics of H_n subordinate to h_n will then converge, in an appropriate sense, to a hierarchy in $\mathcal{C}(R)$ with an associated model map.

Choose a basepoint $v_{n,0}$ for γ_n which is distance at least $k_n/3$ from each endpoint of γ_n . Let $\tau_{n,0}$ be a maximal slice of H_n containing the pair $(h_n, v_{n,0})$, and let $\mu_{n,0}$ be its associated clean marking. Since there are only finitely many clean markings in S up to homeomorphism, we may assume after remarking and extracting a subsequence that all the $\mu_{n,0}$ are equal to a fixed μ_0 , and $v_{n,0} \equiv v_0$.

Fix $E > 0$ and suppose that n is large enough that $E < k_n/6$. We claim that there is a finite set of possibilities, *independently of n* , for the simplices of h_n that are within distance E of v_0 . To see this, let w be such a simplex in h_n . By Lemmas 5.7 and 5.8 of [47], there is a resolution of H_n containing $\tau_{n,0}$ and passing through some slice τ containing the pair (h_n, w) . Now by the monotonicity property of resolutions (see Lemma 4.8) every slice in the resolution between $\tau_{n,0}$ and τ contains a pair (h_n, u) with $v_0 \leq u \leq w$. Therefore, any geodesic appearing in this part of the resolution and supported in R must have footprint in h_n that intersects the interval $[v_0, w]$. Because of the (k_1, k_n) -thick property, all these geodesics have length bounded by k_1 .

The sum of the lengths of all these geodesics can then be bounded by $O(Ek_1^\alpha)$ where $\alpha \leq \xi(S)$, using an inductive counting argument similar to the one in Section 4: First, the segment $[v_0, w]$ has length bounded by E . Thus there are at most $O(E)$ geodesics directly subordinate to h_n with footprint intersecting this interval. Each of these has length bounded by k_1 , so there are $O(Ek_1)$ geodesics directly subordinate to these geodesics. We continue inductively, and note that the complexity ξ decreases with each step. Since every geodesic with footprint in h_n is obtained in this way (by the definition of a hierarchy), this gives us the bound we wanted.

Each elementary move in the resolution that takes place in R corresponds to an edge in one of these geodesics, so we conclude that the markings $\mu_0|_R$ and $\mu_\tau|_R$ are separated by $O(Ek_1^\alpha)$ elementary moves. (Note that we do not obtain or need a bound on the number of moves that occur outside of R .) Since the number of possible elementary moves on any given complete clean marking of R is finite, this gives a finite set, independent of n , containing all possible simplices of h_n within distance E of v_0 .

We conclude that there are only finitely many possibilities for the segment of length E of h_n around v_0 , and after extracting a subsequence these can be assumed constant. Enlarging E and diagonalizing, we may assume that h_n converges to a bi-infinite geodesic h_∞ .

Now fix E again and consider all geodesics m with $D(m) \subset R$ and $d_R([\partial D(m)], v_0) \leq E$. For large enough n , all such m are forward and backward subordinate to h_n , since $D(m)$ intersects $\mathbf{I}(h_n)$ and $\mathbf{T}(h_n)$ by the triangle inequality in $\mathcal{C}(R)$. In particular, if $h_n \not\prec m \not\prec h_n$ then $\mathbf{I}(m)$ and $\mathbf{T}(m)$ are simplices of h_n , and they determine m up to a finite number of choices by Corollary 6.14 of [41]. Proceeding inductively we see that there are only finitely many possibilities for any geodesic m in this set. Thus by the same diagonalization argument as the previous paragraph we may assume that this set of geodesics eventually stabilizes. This gives us a limiting collection H_∞ of tight geodesics supported in subsurfaces of R , and furthermore the subordinacy relations of H_n among such geodesics survive to give H_∞ the structure of a hierarchy. Note that H_∞ has a biinfinite main geodesic h_∞ , and every other geodesic has length at most k_1 .

Convergence of models. The hierarchy H_∞ has an associated model manifold $M_\infty \cong \widehat{R} \times \mathbb{R}$ (where, as in §2.1, \widehat{R} is an open surface containing R such that $R = \widehat{R} \setminus \mathbf{collar}(\partial R)$). We claim that this is a geometric limit of the models M_n for H_n , with appropriate basepoints.

Let us first establish the following fact:

Lemma 7.2. *Let B be a block in M_n and F_Y a gluing surface for B , such that the corresponding 3-holed sphere Y is contained in R , and ∂Y intersects $\text{base}(\mathbf{I}(h_n))$ and $\text{base}(\mathbf{T}(h_n))$. Then $D(B) \subseteq R$.*

Proof. Let f be the 4-geodesic containing the edge e corresponding to B . Then $D(f) = D(B)$, and so $Y \not\prec f$ or $f \not\prec Y$ (often both). Suppose the former. The assumption on ∂Y implies that $h_n \in \Sigma^+(Y)$, so by Theorem 2.1 $Y \not\prec f \not\prec h_n$. In particular $D(B) \subseteq D(h_n) = R$. \square

Now for each n , let x_n be a basepoint on the part of $F_{\tau_{n,0}}$ which lies over R . Then x_n is contained in a three-holed sphere F_Y where Y is a complementary component of $\text{base}(\tau_{n,0})$. Let B_n be one of the two blocks for which F_Y is a gluing surface. By Lemma 7.2, $D(B_n) \subseteq R$.

Let M_{H_n, h_n} denote the union of blocks and tubes in M_n whose associated forward or backward sequences pass through h_n – in particular this is contained in $R \times \mathbb{R}$. The previous paragraph shows that $B_n \subset M_{H_n, h_n}$.

We claim that, for any fixed r , the r -neighborhood $\mathcal{N}_r(x_n)$ in M_n is contained in $M_{H_n, h_n} \cup \mathcal{U}_n(\partial R)$ for sufficiently large n .

First, after deforming paths in tubes to tube boundaries, we note that any point in $\mathcal{N}_r(x_n) \setminus \mathcal{U}_n$ is reachable from x_n through a sequence of $s = O(e^r)$ blocks.

Let B and B' be adjacent blocks such that $D(B) \subset R$ and $\phi_{h_n}(D(B))$ (nonempty by lemma 2.3) is distance at least 4 from the $\text{base}(\mathbf{I}(h_n))$ and $\text{base}(\mathbf{T}(h_n))$. Then B and B' share a gluing surface F_Y such that ∂Y must intersect the curves associated to $\text{base}(\mathbf{I}(h_n))$ and $\text{base}(\mathbf{T}(h_n))$ (or laminations if h_n is infinite). Lemma 7.2 implies that $D(B') \subseteq R$ as well.

Consider first the case that $\xi(R) > 4$. We have that $\phi_{h_n}(D(B))$ and $\phi_{h_n}(D(B'))$ are both in $\phi_{h_n}(Y)$, which has diameter at most 2. Thus, if n is large enough that v_0 is more than $2s + 8$ from the ends of h_n , then any block B that is reachable in s steps from B_n is still in M_{H_n, h_n} , and moreover $d_{\mathcal{C}(R)}(v_0, \partial D(B)) \leq 2s$.

We conclude that any block meeting $\mathcal{N}_r(x_n)$ is in M_{H_n, h_n} , and moreover the boundary of its associated domain is a bounded distance from v_0 in $\mathcal{C}(R)$, and hence for large enough n its associated 4-geodesic is equal to a geodesic in H_∞ .

Now let U be a tube in \mathcal{U}_n meeting $\mathcal{N}_r(x_n)$, such that U is not a component of $\mathcal{U}_n(\partial R)$. Since U meets $\mathcal{N}_r(x_n)$ it is adjacent to at least one block B with $D(B) \subset R$ and $d_{\mathcal{C}(R)}(v_0, \partial D(B)) \leq 2s$. Thus $\text{core}(U)$ is contained in R and $d_{\mathcal{C}(R)}(v_0, \text{core}(U)) \leq 2s + 1$. We claim that in fact *all* blocks B' adjacent to U have $D(B') \subset R$ and $d_{\mathcal{C}(R)}(v_0, \partial D(B')) \leq 2s + 2$. To see this, note first that any block B' adjacent to U must either have a boundary component in the homotopy class of U , or contain $\text{core}(U)$ in $D(B')$, and hence if $D(B') \subset R$ we have $d_{\mathcal{C}(R)}(v_0, \partial D(B')) \leq 2s + 2$. Now if B_1 is a block adjacent to U with $D(B_1) \subset R$, and B_2 is adjacent to U and to B_1 , then (for large enough n) we can again apply Lemma 7.2 to conclude that $D(B_2) \subset R$ as well. It follows by connectivity of ∂U that, in fact, all blocks adjacent to it have domain contained in R .

We can use this to show that, for high enough n , $\omega(U)$ is bounded. Recall from [47, §§8.3, 9.3] that $\text{Im } \omega(U)$, or the ‘‘height’’ of ∂U , is ϵ_1 times the number of annuli in ∂U , and this is estimated up to bounded ratio by the total number of blocks adjacent to U . The footprint $\phi_{h_n}(D(B))$ for any such block is contained in $\phi_{h_n}(\text{core}(U))$, and so if $D(B) \cong m \overset{d}{\ll} h_n$ there are a finite number of possibilities for m , independently of n . The length of m is bounded by k_1 by the (k_1, k_n) -thick condition. By the same inductive counting argument as used above in the discussion of limits of hierarchies, we therefore know the total number of such blocks is bounded by $O(k_1^\alpha)$. We conclude that $\text{Im } \omega(U)$ is uniformly bounded.

The magnitude $|\text{Re } \omega(U)|$ is estimated by the length $|l_U|$ of the annulus geodesic l_U associated to U ; more precisely $||\text{Re } \omega(U)| - |l_U|| \leq C |\text{Im } \omega(U)|$ (see (9.6) and (9.17) of [47]). Since the footprint of this annulus domain is also at most $2s$ from v_0 , for high enough n the (k_1, k_n) condition implies that $|l_U| \leq k_1$. We conclude that $|\omega(U)|$ is bounded, and hence so is the diameter of U .

Thus, fixing r and letting n grow, we find that $\mathcal{N}_r(x_n)$ is contained in $M_{H_n, h_n} \cup \mathcal{U}_n(\partial R)$, and that the geometry of the blocks and tubes (other than $\mathcal{U}_n(\partial R)$) eventually stabilize. It follows that (choosing the $x_n \in F_{\tau_n, 0}$ so that they converge) the geometric limit of (M_{H_n, h_n}, x_n) is the model M_∞ , minus the parabolic tubes associated to ∂R , with basepoint x_∞ .

Furthermore, for each γ in ∂R we have $\text{Im } \omega_n(\gamma) \rightarrow \infty$, because $|h_n| > k_n \rightarrow \infty$. Thus $U_n(\gamma)$ converge geometrically to a rank-1 parabolic tube. Thus in fact the geometric limit of (M_n, x_n) is M_∞ .

It remains to consider the case when $\xi(R) = 4$. Now the blocks of M_{H_n, h_n} are all associated with edges of h_n , and hence are organized in a linear sequence with each one glued to its successor. Each tube U in M_{H_n, h_n} is adjacent to exactly two blocks, so that $\text{Im } \omega(U)$ is uniformly bounded, and $\text{Re } \omega(U)$ is bounded by the lengths of the associated annulus geodesic, as before. Hence the geometric limit is an infinite sequence of blocks and tubes, and the rest of the conclusions follow easily in this case too.

Limit model map. Let τ_0 be the slice of H_∞ obtained as the limit of the restrictions of $\tau_{n,0}$ to R . Since the model maps are uniformly Lipschitz, after fixing an identification of R with \widehat{F}_{τ_0} in M_∞ and possibly conjugating ρ_n in $\text{PSL}_2(\mathbb{C})$, we may extract from $\rho_n|_{\pi_1(R)}$ a convergent subsequence. Denote the limit by ρ_∞ , and let $G_\infty = \rho_\infty(\pi_1(R))$.

Since the boundary curves of R have $|\omega_n| \rightarrow \infty$, their lengths in N_n go to zero (by the Short Curve Theorem of [47]) so they must be parabolic in the limit. So G_∞ is a Kleinian surface group.

Let $\lambda_\pm \in \mathcal{EL}(R)$ be the endpoints of h_∞ (by Klarreich's theorem, see §2.2). The vertices v_i of h_∞ converging to λ_\pm (as $i \rightarrow \pm\infty$) all have bounded length in G_∞ , so their geodesic representatives leave every compact set in the quotient. We conclude that λ_\pm are the ending laminations, and G_∞ is doubly degenerate.

After restricting to a further subsequence we may assume (Lemma 2.8) that $\{\rho_n(\pi_1(S))\}$ converges geometrically to a group Γ_∞ . We claim that $\Gamma_\infty = G_\infty$. The proof is similar to an argument made by Thurston [66] in a slightly different context (see also [20]). Since both ends of G_∞ are degenerate, Thurston's covering theorem (see Theorem 6.10) tells us that the covering map $\mathbb{H}^3/G_\infty \rightarrow \mathbb{H}^3/\Gamma_\infty$ is finite-to-one, and hence $[\Gamma_\infty, G_\infty] < \infty$. If $\gamma \in \Gamma_\infty \setminus G_\infty$ then for some finite k we have $\gamma^k \in G_\infty$. Let $\gamma = \lim \rho_n(g_n)$ with $g_n \in \pi_1(S)$, and let $\gamma^k = \rho_\infty(h)$ with $h \in \pi_1(R)$. By Lemma 2.8, since $\rho_n(h^{-1}g_n^k)$ converges to the identity, we must have $h = g_n^k$ for large enough n . Since k -th roots are unique in $\pi_1(S)$, we find that g_n is eventually constant, and since $\pi_1(R)$ contains all of its roots in $\pi_1(S)$, g_n is eventually contained in $\pi_1(R)$. Thus $\gamma \in G_\infty$ after all.

Let $N_\infty = \mathbb{H}^3/G_\infty = \mathbb{H}^3/\Gamma_\infty$. Again restricting to a subsequence, the model maps f_n converge on M_∞ minus the tubes to a K -Lipschitz map into N_∞ . Since f_n are model maps and the tubes of M_∞ have bounded $|\omega|$, it follows (property (5) of Definition 2.9) that f_n are uniformly Lipschitz on these tubes as well, so that the maps in fact converge everywhere to a map $f_\infty : M_\infty \rightarrow N_\infty$ which is a homotopy-equivalence (since it induces an isomorphism on π_1)

We can see that f_∞ is proper as follows: Any block B in M_∞ meets some slice surface \widehat{F}_τ , and each \widehat{F}_τ meets a representative $\gamma_u \subset M_\infty$ of some vertex

u of the geodesic h_∞ such that γ_u has bounded length. If a sequence B_i of blocks leaves every compact set, the corresponding vertices u_i go to ∞ in $\mathcal{C}(R)$, and so the images of γ_{u_i} in N_∞ , whose lengths remain bounded, must leave every compact set. Since the surfaces \hat{F}_τ have bounded diameter (because there is a uniform bound on $|\omega|$ for all tubes in M_∞), this means that B_i leaves every compact set as well. Each non-peripheral tube lies in a bounded neighborhood of some block, so the images of these tubes are properly mapped as well. The map on $\mathcal{U}(\partial R)$ is a straightened map on rank-1 cusps, so it is proper because its restriction to the cusp boundary is proper. Thus f_∞ is proper.

In the remainder of the proof we will use the following lemma several times. It is essentially a uniform properness property for the sequence of model maps. Let $y_n = f_n(x_n) \in N_n$, and for the limiting basepoint $x_\infty \in M_\infty$ let $y = f_\infty(x_\infty)$.

Lemma 7.3. *For each $r > 0$ there exist $n(r)$ and $d(r)$ such that, for $n \geq n(r)$, $f_n^{-1}(\mathcal{N}_r(y_n))$ is contained in the $d(r)$ -neighborhood of x_n in $M_{H_n, h_n} \cup \mathcal{U}_n(\partial R)$.*

Proof. Suppose by way of contradiction that the lemma is false. Then there exists $r > 0$ such that, after possibly restricting again to a subsequence, there is a sequence $z_n \in M_n$ such that $f_n(z_n) \in \mathcal{N}_r(y_n)$, but $d(z_n, x_n) \rightarrow \infty$.

The tricky point here is that, a priori, the geometric limiting process only controls the maps f_n on large neighborhoods of x_n in M_{H_n, h_n} , so we have to rule out the possibility that z_n is in an entirely different part of the model M_n .

Assume without loss of generality that y is in the ϵ_1 -thick part of N_∞ . If $f_n(z_n)$ is contained in a Margulis tube then there is a point on the boundary of that tube which is still in $\mathcal{N}_r(y)$. Since f_n takes deep model tubes of M_n properly to deep Margulis tubes of N_n , we may assume without loss of generality that whenever z_n lies in a deep model tube it in fact lies on the boundary of the tube. Thus we have a sequence of blocks $\{B_n\}$ that remain a bounded distance from z_n for all n .

Since the maps f_n are uniformly Lipschitz, the image $f_n(B_n)$ remains a bounded distance from y_n . For large enough n , its image must be in the compact set $Y_n \subset N_n$ where the comparison map $\varphi_n : Y_n \rightarrow N_\infty$ is defined. Identifying N_∞ with $R \times \mathbb{R}$, we find that $\varphi_n(f_n(B_n))$ is homotopic into $R \times \{0\}$. For large enough n this homotopy is contained in the comparison region and can be pulled back to N_n . Since f_n is a homotopy equivalence, we conclude that B_n is homotopic to R within M_n .

Thus the domain $D(B_n)$ is a subsurface of R . Let e_n be the 4-edge associated to B_n . Now we claim that the distances $d_{\mathcal{C}(R)}(e_n, v_0)$ are unbounded. Otherwise there is a bounded subsequence and as in the discussion on convergence of hierarchies, for n in the subsequence we can reach B_n from x_n in a bounded distance, using elementary moves from the initial marking μ_0 to

a marking containing the vertex e_n^- . This contradicts the assumption that $d(z_n, x_n) \rightarrow \infty$.

The sequence $\{e_n^-\}$, being unbounded in $\mathcal{C}(R)$, contains infinitely many distinct elements. However all of these are vertices of the model and hence the f_n -images of the corresponding curves in B_n have uniformly bounded length in N_n . The comparison maps take these curves to curves $\{\alpha_n\}$ of bounded length in a compact subset of N_∞ , and this means they fall into finitely many homotopy classes in N_∞ . However, if α_m and $\alpha_{m'}$ are homotopic in N_∞ , then for large enough n the homotopy pulls back and their preimages are homotopic in N_n , and hence in S . This contradicts the fact that there are infinitely many distinct $\{e_n^-\}$, and this contradiction establishes the lemma. \square

As a consequence of this lemma we can show that f_∞ has degree 1. Indeed, let us show that $\deg f_\infty = \deg f_n$ for sufficiently high n . Let X_n denote the region in M_n where the comparison map to the geometric limit $\psi_n : X_n \rightarrow M_\infty$ is defined, and recall that $\varphi_n : Y_n \rightarrow N_\infty$ are the comparison maps for the geometric limit of N_n to N_∞ , and that X_n and Y_n contain arbitrarily large neighborhoods of the basepoints for large enough n .

If $W \subset M_\infty$ is a compact submanifold containing $f_\infty^{-1}(y)$ (which is compact since f_∞ is proper), the degree of $f_\infty|_W$ over y is equal to $\deg f_\infty$.

Now let $d(0)$ and $n(0)$ be given by Lemma 7.3, so that $f_n^{-1}(y_n) \subset \mathcal{N}_{d(0)}(x_n)$ for all $n > n(0)$. We may choose W large enough that, for large enough n , $\psi_n^{-1}(W)$ contains $\mathcal{N}_{d(0)+1}(x_n)$. Thus the degree of $f_n|_{\psi_n^{-1}(W)}$ over y_n is equal to $\deg f_n$.

Now choosing W according to the previous two paragraphs, we know by definition of geometric limits that the maps $\varphi_n \circ f_n \circ \psi_n^{-1}$ are eventually defined on W and converge to $f_\infty|_W$, so that for large enough n the degree of $f_n|_{\psi_n^{-1}(W)}$ over y_n equals the degree of $f_\infty|_W$ over y . Hence $\deg f_\infty = \deg f_n$, and since $\deg f_n = 1$, we are done.

Product regions in the limit. In order to finish the proof we need a topological lemma about deforming proper homotopy equivalences of pairs. Let V be the 3-manifold $R \times \mathbb{R}$, with $\partial V = \partial R \times \mathbb{R}$. Let $C_s = R \times [-s, s]$, which we note is a relative compact core for $(V, \partial V)$.

Lemma 7.4. *Suppose that a map of pairs $f : (V, \partial V) \rightarrow (V, \partial V)$ is a proper, degree 1 map homotopic to the identity. Then there exists a homotopy of f to f' through maps of pairs, such that*

- (1) *The homotopy is compactly supported,*
- (2) *$f'|_{C_1}$ is the identity,*
- (3) *$f'(V \setminus C_1) \subset V \setminus C_1$.*

The proof of this lemma is fairly standard and we omit it.

Now to apply this to our situation let $M'_\infty = M_\infty \setminus \mathcal{U}(\partial R)$ and $N'_\infty = N_\infty \setminus \mathbb{T}(\partial R)$ be the complements of the peripheral model tubes and Margulis

tubes, respectively, and note that there are orientation-preserving identifications

$$\Phi_M : M'_\infty \rightarrow V$$

and

$$\Phi_N : N'_\infty \rightarrow V$$

so that the map $F = \Phi_N \circ f_\infty \circ \Phi_M^{-1}$ satisfies the conditions of Lemma 7.4. We will need to choose these identifications a bit more carefully.

First note that in M_∞ every surface \widehat{F}_τ for a slice τ is isotopic to a level surface. Constructing a cut system as in §4 and using Lemma 4.14, we may choose an ordered sequence of slices $\{c_i\}_{i \in \mathbb{Z}}$ whose base simplices v_i are separated by at least 5 in h_∞ , and adjust Φ_M so that $\Phi_M(\widehat{F}_{c_i}) = R \times \{i\}$.

We may choose the identification Φ_N so that $\Phi_N^{-1}(C_1)$ contains a $(Q+1)$ -neighborhood of $\Phi_N^{-1}(R \times \{0\})$, where Q is the constant in part (3) of the theorem.

Now let $F' : (V, \partial V) \rightarrow (V, \partial V)$ be the map homotopic to F given by Lemma 7.4 and let $f'_\infty = \Phi_N^{-1} \circ F' \circ \Phi_M$. A remaining minor step is to extend f'_∞ to a map (still called f'_∞) which is defined on all of M_∞ , and homotopic to f_∞ by a homotopy of pairs $(M_\infty, \mathcal{U}(\partial R)) \rightarrow (N_\infty, \mathbb{T}(\partial R))$ which is supported on a compact set. Choose a positive integer s so that the homotopy from F to F' is supported in $C_s = R \times [-s, s]$, pull the annuli $\partial C_s \cap \partial V$ back to annuli in $\partial \mathcal{U}(\partial R)$ via Φ_M^{-1} , and pick collar neighborhoods in $\mathcal{U}(\partial R)$ of these annuli. The extension of the homotopy to one supported in the union of C_s and these collar neighborhoods is elementary.

Let G denote the final homotopy from f_∞ to f'_∞ and Z its support in M_∞ . Choose slices τ_i so that $\tau_0 = c_0$, $\tau_{\pm 1} = c_{\pm 1}$, and $\tau_{\pm 2} = c_{\pm s}$. Hence $\mathcal{B}_2 = \mathcal{B}(\tau_{-2}, \tau_2) = \Phi_M^{-1}(C_s)$ together with $\mathcal{U}(\partial R)$ contain the support of the homotopy, and f'_∞ has the properties described in the conclusions of the theorem. It remains to pull this picture back to the approximating manifolds.

Since $G(Z)$ is compact it is contained, for large enough n , in the comparison region for the geometric limit $N_n \rightarrow N_\infty$, and similarly Z is contained in the comparison region for the geometric limit $M_n \rightarrow M_\infty$. Thus we can pull f'_∞ and G back in Z to the approximants, obtaining maps f'_n with homotopies supported in the pullbacks of Z .

The slices $\tau_{\pm i}$ ($i = 0, 1, 2$) give, for large enough n , slices in H_n that define regions $\mathcal{B}_i(n)$ which converge, under the comparison maps, to \mathcal{B}_i . Thus after a small adjustment we can assume that the f'_n are orientation preserving embeddings on $\mathcal{B}_1(n)$, and the homotopies are supported in $\mathcal{B}_2(n) \cup \mathcal{U}_n(\partial R)$. All that remains is to show that, for n large enough, $f'_n(M_n \setminus \mathcal{B}_1(n))$ is disjoint from $f'_n(\mathcal{B}_1(n))$.

To see this, let r be such that $\Phi_N^{-1}(C_1)$ is contained in $\mathcal{N}_r(y)$, and let $n(r)$ and $d(r)$ be the constants given by Lemma 7.3.

If n is chosen greater than $n(r)$ and also sufficiently large that the comparison region X_n in M_n contains $\mathcal{N}_{d(r)}(x_n)$ and $\mathcal{B}_2(n)$, then Lemma 7.3

guarantees that $f_n(M_n \setminus X_n)$, which equals $f'_n(M_n \setminus X_n)$, is disjoint from $f'_n(\mathcal{B}_1(n))$, and the properties of f'_∞ guarantee that $f'_n(X_n \setminus \mathcal{B}_1(n))$ is disjoint from $f'_n(\mathcal{B}_1(n))$. Thus f'_n has the desired property.

Finally, we note that the entire construction can be performed so that all maps are Lipschitz with some constant L , simply by using piecewise-smooth maps in the geometric limit. Thus the sequence of maps f_n does admit f'_n which satisfy the conclusions of the theorem with Lipschitz constant L , and as soon as $L_n > L$ this contradicts our original choice of sequence. This contradiction establishes the theorem. \square

8. Proof of the bilipschitz model theorem

We are now ready to put together the ingredients of the previous sections and complete the proof of our main technical theorem, which we restate here:

Bilipschitz Model Theorem. *There exist $K', k' > 0$ depending only on S , so that for any Kleinian surface group $\rho \in \mathcal{D}(S)$ with end invariants $\nu = (\nu_+, \nu_-)$ there is an orientation-preserving K' -bilipschitz homeomorphism of pairs*

$$F : (M_\nu, \mathcal{U}[k']) \rightarrow (\widehat{C}_{N_\rho}, \mathbb{T}[k']).$$

We will first prove the theorem in the special case where ν_\pm are both laminations in $\mathcal{EL}(S)$ (the doubly degenerate case). The remaining cases will be treated in Section 8.6.

8.1. Embedding an individual cut

Let $f : M_\nu \rightarrow \widehat{C}_N$ be the (K, k) model map provided by the Lipschitz Model Theorem (see §2.7).

Recall the Otal constant ϵ_u from Theorem 2.5, and the functions \mathcal{L} from Lemma 2.12 and Ω from the Short Curve Theorem in §2.7. Let $k_u = \max(k, \Omega(\epsilon_u))$.

Each surface $\widehat{F}_\tau[k_u]$ is composed of standard 3-holed spheres attached to bounded-geometry annuli. Hence it admits an r -bounded hyperbolic metric σ_τ with geodesic boundary which is r -bilipschitz equivalent to its original metric, for some r depending on k_u . We will henceforth consider these surfaces with these adjusted metrics. This together with the K -Lipschitz bounds on the model map f tells us that $f|_{\widehat{F}_\tau[k_u]}$ is an L_0 -bounded map (as in §6), where L_0 depends on k_u and K .

Lemma 8.1. *There exist d_0, K_1 and $\hat{\epsilon}$ (depending only on S) such that, if τ is a full slice in H_ν such that the length $|g_\tau|$ of its base geodesic is at least d_0 , then $f|_{\widehat{F}_\tau[k_u]}$ is ϵ_u -anchored, and $(K_1, \hat{\epsilon})$ -uniformly embeddable (with respect to the metric σ_τ on $\widehat{F}_\tau[k_u]$).*

Proof. We will check, applying the Short Curve Theorem, that the conditions of Theorem 6.1 hold, and thereby obtain the uniform embeddability.

Let $\delta_1 = \epsilon_u$, and let δ_2 be sufficiently small that, if $v \in \mathcal{C}(S)$ has $\ell_\rho(v) < \delta_2$, then v is in the hierarchy H_ν and $|\omega(v)| > k$. Such a δ_2 is guaranteed by the Short Curve Theorem, parts (1) and (2) (see §2.7). Let ϵ be the constant provided by Theorem 6.1 for these values of δ_1 and δ_2 , and $L = L_0$. Now let $d_0 = \mathcal{L}(\epsilon)$. Let $R = D(\tau)$ and let Γ be the set of vertices v in τ with $|\omega_M(v)| \geq k_u$. Thus $X = R \setminus \mathbf{collar}(\Gamma)$ can be identified with $\widehat{F}_\tau[k_u]$, and this gives it an L_0 -bounded metric.

By part (3) of the Short Curve Theorem and the choice of d_0 and Γ , we have $\ell_\rho(v) \leq \epsilon_u$ for $v \in \Gamma$. Thus the map $f|_{\widehat{F}_\tau[k_u]}$, or $f|_X$, is ϵ_u -anchored (hence δ_1 -anchored). Moreover, by Lemma 2.12 and $|g_\tau| \geq d_0 = \mathcal{L}(\epsilon)$, we also have $\ell_\rho(\partial D(\tau)) < \epsilon$. Hence, letting $\tilde{f} = f|_{\widehat{F}_\tau}$, we find that \tilde{f} is ϵ -anchored, satisfying condition (3) of Theorem 6.1.

Condition (1) of Theorem 6.1 follows from the properties of the (K, k) model map, and the choice of k_u and Γ .

Next we establish the unwrapping condition (2) of Theorem 6.1. If w is any vertex such that $\ell_\rho(w) < \delta_2$, we have $|\omega(w)| > k$ and hence f takes $U(w)$ to $\mathbb{T}(w)$ and $M_\nu \setminus U(w)$ to $N \setminus \mathbb{T}(w)$. Applying Lemma 5.3 to $U(w)$ and \widehat{F}_τ , we have in particular that \widehat{F}_τ is homotopic to either $+\infty$ or $-\infty$ in the complement of $U(w)$, so $f|_{\widehat{F}_\tau}$ is homotopic to either $+\infty$ or $-\infty$ in the complement of $\mathbb{T}(w)$. This is exactly the unwrapping condition (2).

(When M_ν has nonempty boundary, we interpret “homotopic to $\pm\infty$ in the model” by considering $M' = M_\nu \setminus (\partial M_\nu \cup \mathcal{U}(\partial S))$, which is homeomorphic to $S \times \mathbb{R}$. In N , we consider $C' = \widehat{C}_N \setminus (\partial \widehat{C}_N \cup \mathbb{T}(\partial S))$. Since the model map takes M' properly to C' , we can make the same arguments.)

Having verified that the conditions of Theorem 6.1 hold, we obtain the desired uniform embeddability of $f|_{\widehat{F}[k_u]}$. \square

For a nonannular cut c in a cut system C with spacing lower bound at least d_0 , let $G_c : \widehat{F}_c[k_u] \times [0, 1] \rightarrow N$ be the homotopy provided by Lemma 8.1, which is a bilipschitz embedding on $\widehat{F}_c[k_u] \times [1/2, 1]$. Let $f_c : \widehat{F}_c[k_u] \rightarrow N$ be the embedding defined by

$$f_c(x) = G_c(x, 3/4). \quad (8.1)$$

We can extend this map to the annuli of $\widehat{F}_c \setminus \widehat{F}_c[k_u]$, if we are willing to drop the Lipschitz bounds:

Corollary 8.2. *The homotopy G_c provided by Lemma 8.1 can be extended to a map $\tilde{G}_c : \widehat{F}_c \times [0, 1] \rightarrow N$ so that, on each annulus A in $\widehat{F}_c \setminus \widehat{F}_c[k_u]$ we have $\tilde{G}_c(A \times [0, 1])$ contained in the corresponding Margulis tube $\mathbb{T}(A)$, and $\tilde{f}_c(x) = G_c(x, 3/4)$ is still an embedding.*

Proof. Note first that, by our choice of k_u , each A indeed has core curve whose length in N is sufficiently short that $\mathbb{T}(A)$ is non-empty. The model map f already takes A into $\mathbb{T}(A)$. Since G_c is a homotopy through anchored

embeddings, it takes $\partial A \times [0, 1]$ to $\partial\mathbb{T}(A)$, and so the existence of \bar{G}_c is a simple fact about mappings of annuli into solid tori. \square

Annular cuts. If c is an annular cut let $\omega(c)$ denote the meridian coefficient of the corresponding tube $U(c)$, and if $|\omega(c)| > k_u$ let $\mathbb{T}(c)$ denote the Margulis tube associated to the homotopy class of the annulus. For notational consistency, we let

$$f_c = \mathbb{T}(c) \tag{8.2}$$

when $|\omega(c)| > k_u$. As in §3 we are blurring the distinction between a map and its image here.

8.2. Thinning the cut system

Lemma 8.1 allows us, after bounded homotopy of the model map, to embed individual slices of a cut system, but the images of these embeddings may intersect in unpredictable ways. We will now show that, by thinning out a cut system in a controlled way we can obtain one for which the cuts that border any one complementary region have disjoint f_c -images.

Because the model manifold is built out of standard pieces, for any nonannular slice c there is a paired bicollar neighborhood E_c^0 for $(\widehat{F}_c[k_u], \partial\widehat{F}_c[k_u])$ in $(M_\nu[k_u], \partial\mathcal{U}[k_u])$ which is uniformly bilipschitz equivalent to a standard product. That is, there is a bilipschitz piecewise smooth orientation-preserving homeomorphism

$$\varphi_c : E_c^0 \rightarrow \widehat{F}_c[k_u] \times [-1, 1]$$

which restricts to $x \mapsto (x, 0)$ on $\widehat{F}_c[k_u]$, where the bilipschitz constant depends only on the surface S . We are here taking $\widehat{F}_c[k_u] \times [-1, 1]$ with the product metric $\sigma_c \times dt$, where σ_c is the hyperbolic metric defined in §8.1. The relative boundary $\varphi_c^{-1}(\partial\widehat{F}_c[k_u] \times [-1, 1])$ is the intersection of E_c^0 with $\partial\mathcal{U}[k_u]$. Moreover we may choose the collars so that, if $\widehat{F}_c[k_u]$ and $\widehat{F}_{c'}[k_u]$ are disjoint, so are E_c^0 and $E_{c'}^0$. Let E_c denote the subcollars $\varphi_c^{-1}(\widehat{F}_c[k_u] \times [-\frac{1}{2}, \frac{1}{2}])$; our final map will be an embedding on each E_c , for an appropriate set of slices $\{c\}$.

Lemma 8.3. *Given a cut system C with spacing lower bound $d_1 \geq d_0$ and upper bound $3d_1$, there exists a cut system $C' \subset C$ with spacing upper bound $d_2 = d_2(S, d_1)$, and $K_2 = K_2(S)$, such that there is a map of pairs*

$$f' : (M_\nu[k_u], \partial\mathcal{U}[k_u]) \rightarrow (\widehat{C}_{N_\rho} \setminus \mathbb{T}[k_u], \partial\mathbb{T}[k_u])$$

which is homotopic through maps of pairs to the restriction of the model map $f|_{M_\nu[k_u]}$, such that

- (1) f' is a (K_2, k_u) model map, and its homotopy to f is supported on the union of collars $\cup E_c^0$ over nonannular $c \in C'$.
- (2) Inside each subcollar E_c for nonannular $c \in C'$, f' is an orientation-preserving K_2 -bilipschitz embedding, and

$$f'|_{\widehat{F}_c[k_u]} = f_c.$$

Furthermore $f' \circ \varphi_c^{-1}$ restricted to $\varphi_c(E_c)$ has norms of second derivatives bounded by K_2 , with respect to the metric $\sigma_c \times dt$.

- (3) For each complementary region $W \subset M_\nu[k_u]$ of C' , the subcollars E_c of nonannular cut surfaces on ∂W have disjoint f' -images.

Proof. For each $c \in C$ we will use the homotopy to an embedding provided by Lemma 8.1 (via Theorem 6.1) to redefine f in E_c^0 . The resulting map will immediately satisfy the conclusions of Lemma 8.3 except possibly for the disjointness condition (3). In order to satisfy (3) we will have to “thin” the cut system.

To define the map in each collar, fix a nonannular $c \in C$, and let $G_c : \widehat{F}_c[k_u] \times [0, 1] \rightarrow N \setminus \mathbb{T}[k_u]$ denote the proper homotopy given by Lemma 8.1 and Theorem 6.1, where we recall that G_c restricted to $\widehat{F}_c[k_u] \times [1/2, 1]$ is a K_1 -bilipschitz embedding with norm of second derivatives bounded by K_1 , and $f_c(x) = G_c(x, 3/4)$.

Define $\sigma : [-1, 1] \rightarrow [-1, 1]$ so that it is affine in the complement of the ordered 6-tuple $(-1, -3/4, -1/2, 1/2, 3/4, 1)$, and takes the points of the 6-tuple, in order, to $(-1, 0, 1/2, 1, 0, 1)$ (see Figure 9). Note that σ is orientation-reversing in $(1/2, 3/4)$ and orientation-preserving otherwise.

FIGURE 9. The graph of the reparametrization function $\sigma(t)$

Define a map $g_c : \widehat{F}_c[k_u] \times [-1, 1] \rightarrow N$ via

$$g_c(x, t) = \begin{cases} G_c(x, \sigma(t)) & |t| \leq 3/4 \\ f(\varphi_c^{-1}(x, \sigma(t))) & |t| \geq 3/4. \end{cases}$$

Now given a subset $C' \subset C$, let f' restricted to E_c^0 for $c \in C'$ be $g_c \circ \varphi_c$, and let $f' = f$ on the complement of $\cup_{c \in C'} E_c^0$. Since the collars are pairwise disjoint for all $c \in C$, this definition makes sense.

One easily verifies that f' satisfies conclusions (1) and (2) of the lemma. Note in particular that, since σ takes $[-1/2, 1/2]$ to $[1/2, 1]$ by an orientation-preserving homeomorphism, f' on E_c is a reparametrization of the embedded part of G_c , and that since $\sigma(0) = 3/4$ we have $f'|_{\widehat{F}_c[k_u]} = f_c$.

We will now explain how to choose C' so that f' will satisfy the disjointness condition (3) as well.

The Margulis lemma gives a bound $n(L)$ on the number of loops of length at most L through any one point in a hyperbolic 3-manifold, if the loops represent distinct primitive homotopy classes. This, together with the Lipschitz bound on f' and the fact that all vertices in different slices are distinct, gives a bound $\beta(r)$ on the number of $f'(E_c)$ that can touch any given r -ball in N .

Let r_0 be an upper bound for the diameter of any embedded collar $f'(E_c)$ (following from the Lipschitz bound on f').

The pigeonhole principle then yields the following observation, which will be used repeatedly below:

- (*) Given a set $Z \subset C$ of at most k slices, and a set $Y \subset C$ of at least $k\beta(r_0) + 1$ slices, there exists $c \in Y$ such that $f'(E_c)$ is disjoint from $f'(E_{c'})$ for all $c' \in Z$.

We will describe a nested sequence $C = C_0 \supset C_1 \supset \cdots \supset C_{\xi(S)-3}$ of cut systems, so that C_j satisfies an upper spacing bound $d_2(j)$, and the following condition, where for an address α of C_j we let $\mathcal{X}_j(\alpha)$ denote the union of blocks with address α .

- (**) For any two slices $c, c' \in C_j$ whose cut surfaces are on the boundary of $\mathcal{X}_j(\alpha)$, and whose complexities are greater than $\xi(S) - j$, $f'(E_c)$ and $f'(E_{c'})$ are disjoint.

Thus $C_{\xi(S)-3}$ will be the desired cut system C' , with $d_2 = d_2(\xi(S) - 3)$. Note that by assumption $d_2(0) = 3d_1$, and that C_0 satisfies condition (**) vacuously.

We obtain C_1 from C_0 by letting $C_1|_h = C_0|_h$ for all $h \neq g_H$, and removing slices on the main geodesic g_H : If $C_0|_{g_H}$ has at most $\beta(r_0) + 1$ slices, set $C_1|_{g_H} = \emptyset$. In this case $|g_H|$ is at most $3d_1(\beta(r_0) + 2)$.

If $C_0|_{g_H}$ has at least $\beta(r_0) + 2$ slices, we partition it into a sequence of "intervals" $\{J_i\}_{i \in \mathcal{I}}$, indexed by an interval $\mathcal{I} \subset \mathbb{Z}$ containing 0; by this we mean that, using the cut order \prec_c on slices, each J_i contains all slices of $C_0|_{g_H}$ between $\min J_i$ and $\max J_i$, and that $\min J_{i+1}$ is the successor to $\max J_i$. Furthermore, we may do this so that J_0 is a singleton, and for $i \neq 0$ J_i has size $\beta(r_0) + 1$, except for the largest positive i and smallest negative i (if any), for which J_i has between $\beta(r_0) + 1$ and $2\beta(r_0) + 1$ elements. Note that \mathcal{I} is infinite if g_H is, for example in the doubly degenerate case. If g_H and hence $C_0|_{g_H}$ are finite, the condition on sizes of J_i is easily arranged by elementary arithmetic.

Let $J_0 = \{c_0\}$. Proceeding inductively, we select some c_i in J_i for each positive i so that $f'(E_{c_i})$ is disjoint from $f'(E_{c_{i-1}})$. Our constraints on the sizes of J_i , together with observation (*), guarantees that this choice is always possible. Similarly for $i < 0$ we choose c_i such that $f'(E_{c_i})$ is disjoint from $f'(E_{c_{i+1}})$.

In this case we define C_1 so that $C_1|_{g_H} = \{c_i\}_{i \in \mathcal{I}}$. The spacing upper bound in this case is at most $d_2(1) \equiv (3\beta(r_0) + 1)3d_1$. We note that C_1 has the property that for two successive c, c' in $C_1|_{g_H}$, $f'(E_c)$ and $f'(E_{c'})$ are disjoint. Since \widehat{F}_c and $\widehat{F}_{c'}$ are in the boundary of the same region if and only if they are consecutive, this establishes the inductive hypothesis for C_1 .

We proceed by induction. We will construct C_{j+1} from C_j by applying a similar thinning process to every geodesic h with complexity $\xi(h) = \xi(S) - j$, and setting $C_{j+1}|_m = C_j|_m$ for every m with $\xi(m) \neq \xi(S) - j$. For any address α , the total number of blocks in $\mathcal{X}_j(\alpha)$ is bounded by b_j depending on $d_2(j)$, by Lemma 5.8, and hence there is a bound s_j on the number of cut surfaces on the boundary of $\mathcal{X}_j(\alpha)$. (certainly s_j is no more than $4b_j$ but a better bound can probably be found).

Enumerate the geodesics with complexity $\xi(S) - j$ as h_1, h_2, \dots , and thin them successively: At the k th stage we have already, for each h_m with $m < k$, thinned $C_j|_{h_m}$ to obtain $C_{j+1}|_{h_m}$. If h_k contains fewer than $2s_j\beta(r_0) + 2$ slices, remove them all, obtaining $C_{j+1}|_{h_k} = \emptyset$. Otherwise, recall from Lemma 5.6 that there is a unique address α such that h_k is an inner boundary geodesic for $\mathcal{X}_j(\alpha)$. Let Q_1 be the set of all slices $c \in C_j$ with F_c in $\partial\mathcal{X}_j(\alpha)$ and $\xi(c) > \xi(h_k)$. Let Q_2 be the set of all slices that arise as first and last slices in $C_{j+1}|_{h_m}$ for each $m < k$ such that h_m is an inner boundary geodesic for $\mathcal{X}_j(\alpha)$. Thus each element of Q_2 is the ‘‘replacement’’ for a boundary surface of $\mathcal{X}_j(\alpha)$ of complexity $\xi(h_k)$ that may have been removed by the thinning process. The bound on the number of boundary surfaces of $\mathcal{X}_j(\alpha)$ implies that the union $Q_1 \cup Q_2$ has at most $s_j - 1$ elements.

Partition $C_j|_{h_k}$ into a (possibly infinite) sequence of consecutive, contiguous subsets $\{J_i\}_{i \in \mathcal{I}}$, such that the first and the last (if they exist) have length at least $s_j\beta(r_0) + 1$ and at most $(s_j + 1)\beta(r_0) + 1$, and the rest have length $2\beta(r_0) + 1$.

Now let $C_{j+1}|_{h_k}$ be the union of one cut from each J_i , selected as follows: Supposing that there is a first J_{i_p} , choose a cut $c_{i_p} \in J_{i_p}$ such that $f'(E_{c_{i_p}})$ is disjoint from $f'(E_b)$ for each $b \in Q_1 \cup Q_2$. If there is also a last J_{i_q} , choose c_{i_q} such that $f'(E_{c_{i_q}})$ is disjoint from $f'(E_b)$ for each $b \in Q_1 \cup Q_2 \cup \{c_{i_p}\}$. Now for each $i_p < i < i_q - 1$ we successively choose $c_i \in J_i$ so that $f'(E_{c_i})$ is disjoint from $f'(E_{c_{i-1}})$. If $i = i_q - 1$, we choose c_i so that $f'(E_{c_i})$ is disjoint from $f'(E_{c_{i-1}})$ and $f'(E_{c_{i_q}})$. Note that all these selections are possible by the choice of sizes of the J_i , the bound on the size of $Q_1 \cup Q_2$, and observation (*). If there is a last J_i but not a first, we proceed similarly but in the opposite direction. (There must be either a last or a first since g_H is the only geodesic in H that can be biinfinite).

We can then set $d_2(j+1) \equiv d_2(j)(2(s_j+1)\beta(r_0)+1)$ to be the upper spacing bound for C_{j+1} .

To verify that C_{j+1} satisfies the condition (**), consider $\mathcal{X}_{j+1}(\alpha)$ for any address pair α occurring in C_{j+1} . Let h be g_α (as in the proof of Lemma 5.8), and denote $\xi(\alpha) = \xi(h)$.

If $\xi(\alpha) < \xi(S) - j$ then there is nothing to check since the boundary surfaces of $\mathcal{X}_{j+1}(\alpha)$ all have complexity less than $\xi(S) - j$.

If $\xi(\alpha) = \xi(S) - j$ then only the two outer boundary surfaces of $\mathcal{X}_{j+1}(\alpha)$, namely F_c and $F_{c'}$, have complexity greater than $\xi(S) - (j+1)$. Since in this case h participated in the thinning step we just completed, and c and c' are successive slices in $C_{j+1}|_h$, we have $f'(E_c)$ disjoint from $f'(E_{c'})$.

If $\xi(\alpha) > \xi(S) - j$, then the address α is also an address of C_j , since h was not thinned in the construction of C_{j+1} . The outer boundary surfaces of $\mathcal{X}_j(\alpha)$ and $\mathcal{X}_{j+1}(\alpha)$ are the same, namely F_c and $F_{c'}$. Now consider any inner boundary geodesic m for $\mathcal{X}_{j+1}(\alpha)$. If $\xi(m) > \xi(S) - j$ then m was not thinned in this step, and hence $C_{j+1}|_m = C_j|_m$. If $\xi(m) = \xi(S) - j$ then $C_{j+1}|_m \subset C_j|_m$. In either case, m is an inner boundary geodesic for $\mathcal{X}_j(\alpha)$ as well, since any address pair of C_j in which slices on m are nested must have $\xi > \xi(m)$ and hence was not removed in this step.

Now for any boundary geodesics m and m' of $\mathcal{X}_{j+1}(\alpha)$ with complexities at least $\xi(S) - j$, and slices d on m and d' on m' corresponding to boundary surfaces of $\mathcal{X}_{j+1}(\alpha)$, we must show $f'(E_d)$ and $f'(E_{d'})$ are disjoint. If $\xi(m) > \xi(S) - j$ and $\xi(m') > \xi(S) - j$ then d and d' correspond to boundary surfaces of $\mathcal{X}_j(\alpha)$, and we have disjointness by induction. If $\xi(m) = \xi(S) - j$ and $\xi(m') > \xi(S) - j$ then, when the slices on m are thinned to yield $C_{j+1}|_m$, the slice d' is in Q_1 , and by the construction we have disjointness. If $\xi(m) = \xi(S) - j$ and $\xi(m') = \xi(S) - j$, then if $m = m'$ then d and d' are the first and last slices of $C_{j+1}|_m$, so again the construction makes $f'(E_d)$ and $f'(E_{d'})$ disjoint. Finally if $m \neq m'$ we may suppose that m' is thinned before m , and then at the point that m is thinned we have d' as one of the slices in Q_2 , so again the construction gives us disjointness.

This gives the disjointness property for all boundary surfaces of $\mathcal{X}_{j+1}(\alpha)$ of complexity at least $\xi(S) - j$, which establishes property (**) for C_{j+1} . \square

8.3. Preserving order of embeddings

Let C be a cut system with spacing lower bound at least d_0 , and such that $|\omega(c)| > k_u$ for each annular $c \in C$. Let f_c and G_c be as in §8.1. The following lemma states that, if the spacing of C is large enough then for slices with overlapping domains and with disjoint f_c -images, topological order in the image is equivalent to the cut order \prec_c .

Lemma 8.4. *There exists a $d_1 \geq d_0$ such that, if C is a cut system with spacing lower bound of d_1 and $|\omega(c)| > k_u$ for annular slices, c and c' are two slices in C such that $\tilde{D}(c) \cap \tilde{D}(c') \neq \emptyset$ and f_c and $f_{c'}$ are disjoint, then*

$$c \prec_c c' \implies f_c \prec_{\text{top}} f_{c'}.$$

Proof. Consider first the case where both c and c' are nonannular. Since G_c and $G_{c'}$ are $(K_1, \hat{\epsilon})$ uniform homotopies, f_c and $f_{c'}$ avoid $\mathbb{T}(\gamma)$ whenever γ has length less than $\hat{\epsilon}$, and the homotopies G_c and $G_{c'}$ stay out of the $\hat{\epsilon}$ -Margulis tubes.

Set $k_1 = (k + b_1)/b_2$ where b_1 and b_2 are the constants in Lemma 2.11, and let k_2 be the constant produced by Theorem 7.1 for this k_1 , and for $Q = K_1$. Let us set $d_1 = \max(d_0, k_2 + 14, \mathcal{L}(\hat{\epsilon}))$.

Case 1: $D(c) = D(c')$. Thus the slices have a common bottom geodesic h , and the base simplices satisfy $v_c < v_{c'}$.

The idea now is to find an intermediate subset in M_ν between $\widehat{F}_c[k_u]$ and $\widehat{F}_{c'}[k_u]$, whose image will separate their images in N , and force them to be in the correct order. This separator will either be a Margulis tube with large coefficient ω , in case 1a below, or a “product region” isotopic to $D(c) \times [0, 1]$ in case 1b.

Case 1a: Suppose that there is some geodesic m with $D(m) \subset D(h)$, $|m| > k_1$ and such that $\phi_h(D(m))$ is at least 5 forward of v_c and at least 5 behind $v_{c'}$. There is at least one boundary component w of $D(m)$ which is nonperipheral in $D(h)$. Let a be an annular slice with domain $\mathbf{collar}(w)$. Then $\{c, c', a\}$ satisfies the conditions of a cut system, and we claim that $c \prec_c a \prec_c c'$. This is because $\phi_h(w)$ contains $\phi_h(D(m))$, so by choice of m and the fact that footprints have diameter at most 2, we know that v_c is at least 3 behind $\min \phi_h(w)$ and $v_{c'}$ is at least 3 ahead of $\max \phi_h(w)$. This implies that $c \vdash a \dashv c'$, hence $c \prec_c a \prec_c c'$.

Now Lemma 4.14 implies that

$$\widehat{F}_c \prec_{\text{top}} U(w)$$

and

$$U(w) \prec_{\text{top}} \widehat{F}_{c'}.$$

Since $|m| > k_1$ we have $|\omega_M(w)| > k$ by Lemma 2.11. This guarantees that $f(U(w)) = \mathbb{T}(w)$ and that $f(M_\nu \setminus U(w)) = \widehat{C}_N \setminus \mathbb{T}(w)$. Therefore, $f|_{\widehat{F}_c}$ is homotopic to $-\infty$ in the complement of $\mathbb{T}(w)$, and $f|_{\widehat{F}_{c'}}$ is homotopic to $+\infty$ in the complement of $\mathbb{T}(w)$. (We make sense of this in the case when $\partial M_\nu \neq \emptyset$ just as in the proof of Lemma 8.1.) Now since G_c and $G_{c'}$ miss $\mathbb{T}_{\hat{\epsilon}}(w)$, we find that also f_c is homotopic to $-\infty$ in the complement of $\mathbb{T}(w)$, and $f_{c'}$ is homotopic to $+\infty$ in the complement of $\mathbb{T}(w)$. Let \bar{G}_c and \bar{f}_c be the extensions of G_c and f_c to \widehat{F}_c given by Corollary 8.2. Since they differ from G_c and f_c only in tubes associated to vertices of c , which are all disjoint from U , we may conclude that \bar{f}_c is homotopic to $-\infty$ in the complement of $\mathbb{T}(w)$. Define $\bar{f}_{c'}$ similarly and note that it is homotopic to $+\infty$ in the complement of $\mathbb{T}(w)$. Now since \bar{f}_c and $\bar{f}_{c'}$ are embedded surfaces anchored on Margulis tubes which are unknotted and unlinked by Otal’s theorem, we may apply Lemma 3.16 to conclude that

$$\bar{f}_c \prec_{\text{top}} \mathbb{T}(w)$$

and

$$\mathbb{T}(w) \prec_{\text{top}} \bar{f}_{c'}.$$

Now apply Lemma 3.14, with $R_1 = \bar{f}_c$, $R_2 = \bar{f}_{c'}$, $\mathcal{V} = \overline{\mathbb{T}(\partial D(h))}$, and $Q = \overline{\mathbb{T}(w)}$, to conclude

$$\bar{f}_c \prec_{\text{top}} \bar{f}_{c'}.$$

Thus by Lemma 3.3

$$f_c \prec_{\text{top}} f_{c'}$$

Case 1b: If Case 1a does not hold, then for every geodesic m with $D(m) \subset D(h)$ such that $\phi_h(D(m))$ is at least 5 forward of v_c and at least 5 behind $v_{c'}$, we must have $|m| \leq k_1$. Let γ be the subsegment of $[v_c, v_{c'}]$ which excludes 7-neighborhoods of the endpoints. Then since $d_1 \geq k_2 + 14$, γ satisfies the (k_1, k_2) -thick condition of Theorem 7.1. Thus, Theorem 7.1 provides slices $\tau_{-2}, \tau_{-1}, \tau_0, \tau_1, \tau_2$ with bottom geodesic h and bottom simplices in γ satisfying

$$v_{\tau_{-2}} < v_{\tau_{-1}} < v_{\tau_0} < v_{\tau_1} < v_{\tau_2},$$

with spacing of at least 5 between successive simplices, so that f can be deformed, by a homotopy supported on $\mathcal{B}_2 = \mathcal{B}(\tau_{-2}, \tau_2)$, to an L -Lipschitz map f' such that f' is an orientation-preserving embedding on $\mathcal{B}_1 = \mathcal{B}(\tau_{-1}, \tau_1)$, and f' takes $M_\nu \setminus \mathcal{B}_1$ to $N \setminus f'(\mathcal{B}_1)$.

Since $v_{\tau_{-2}}$ and v_{τ_2} are at least 5 away from v_c and $v_{c'}$ we may conclude that $\{c, \tau_{-2}, \dots, \tau_2, c'\}$ form a cut system, and that $c \prec_c \tau_{-2}$ and $\tau_2 \prec_c c'$. Lemma 4.14 now implies that $\widehat{F}_c \prec_{\text{top}} \widehat{F}_{\tau_{-2}}$ and $\widehat{F}_{\tau_2} \prec_{\text{top}} \widehat{F}_{c'}$.

This implies that \widehat{F}_c can be pushed to $-\infty$ in M_ν in the complement of \mathcal{B}_2 . Applying f' , we find that $f'|_{\widehat{F}_c} = f|_{\widehat{F}_c}$ may be pushed to $-\infty$ in N in the complement of $f'(\mathcal{B}_1)$.

Since we invoked Theorem 7.1 with $Q = K_1$, part (3) of that theorem tells us that $f'(\mathcal{B}_1)$ contains a K_1 neighborhood of $f'(\widehat{F}_{\tau_0})$ in $N \setminus \mathbb{T}(\partial D(h))$, and since the tracks of the homotopy G_c have length at most K_1 , we may conclude that G_c avoids $f'(\widehat{F}_{\tau_0})$.

Again let \bar{G}_c and \bar{f}_c be the extensions of G_c and f_c to \widehat{F}_c given by Corollary 8.2. Each annulus A of $\widehat{F}_c \setminus \widehat{F}_c[k_u]$ corresponds to an annular slice a such that, similarly to the argument in case 1a, $\{a, \tau_{-2}\}$ form a cut system where $a \prec_c \tau_{-2}$. Thus by Lemma 4.14 we have $U(a) \prec_{\text{top}} \widehat{F}_{\tau_{-2}}$ and it follows that $U(a)$ lies outside \mathcal{B}_2 . Hence, its image tube $\mathbb{T}(A)$ lies outside $f'(\mathcal{B}_1)$.

It follows that the extended homotopy \bar{G}_c avoids $f'(F_{\tau_0})$. Thus \bar{f}_c can be pushed to $-\infty$ in the complement of $f'(\widehat{F}_{\tau_0})$.

Since $f'|_{\widehat{F}_{\tau_0}}$ and \bar{f}_c are disjoint homotopic embeddings anchored on the tubes of $\mathbb{T}(\partial D(h))$, part (1) of Lemma 3.14 implies they are \prec_{top} -ordered, and so the homotopy of \bar{f}_c to $-\infty$ tells us that

$$\bar{f}_c \prec_{\text{top}} f'|_{F_{\tau_0}}.$$

Arguing similarly with c' , we obtain

$$f'|_{F_{\gamma_0}} \prec_{\text{top}} \bar{f}_{c'}.$$

Now we apply part (2) of Lemma 3.14 to conclude that

$$\bar{f}_c \prec_{\text{top}} \bar{f}_{c'}.$$

It follows by Lemma 3.3 that

$$f_c \prec_{\text{top}} f_{c'}.$$

Case 2: $D(c)$ and $D(c')$ intersect but are not equal. In this case we will obtain the correct order by looking at the tubes on the boundaries of $D(c)$ and $D(c')$.

Since $c \prec_c c'$, we have $\widehat{F}_c \prec_{\text{top}} \widehat{F}_{c'}$ by Lemma 4.14. Thus, for any component γ of $\partial D(c')$ which overlaps \widehat{F}_c , we can deform \widehat{F}_c to $-\infty$ in the complement of $U(\gamma)$. Since $d_1 > k_2 \geq k_1$, we have $|\omega_M(\gamma)| \geq k$ by Lemma 2.11. It follows from the properties of the model map that $f|_{\widehat{F}_c[k_u]}$ can be deformed to $-\infty$ in the complement of $\mathbb{T}(\gamma)$. Since $d_1 \geq \mathcal{L}(\hat{\epsilon})$, by Lemma 2.12 we have $\ell_\rho(\gamma) < \hat{\epsilon}$ and so the homotopy G_c avoids the core of Margulis tubes $\mathbb{T}(\gamma)$. We can conclude that the embedding f_c can also be deformed to $-\infty$ in the complement of $\mathbb{T}(\gamma)$.

Now the extended homotopy \bar{G}_c (from Corollary 8.2) takes each annulus A of $\widehat{F}_c \setminus \widehat{F}_c[k_u]$ to $\mathbb{T}(A)$, and hence is still disjoint from $\mathbb{T}(\gamma)$. Thus \bar{f}_c can also be deformed to $-\infty$ in the complement of $\mathbb{T}(\gamma)$, and Lemma 3.16 implies that

$$\bar{f}_c \prec_{\text{top}} \mathbb{T}(\gamma).$$

Similarly if γ is a component of $\partial D(c)$ which intersects $D(c')$, we find that

$$\mathbb{T}(\gamma) \prec_{\text{top}} \bar{f}_{c'}.$$

Now applying Lemma 3.15, we conclude that

$$\bar{f}_c \prec_{\text{top}} \bar{f}_{c'}.$$

Again by Lemma 3.3 we conclude

$$f_c \prec_{\text{top}} f_{c'}.$$

Annular cuts. It remains to consider the case that at least one of c and c' are annular. Suppose that both are. Since $\check{D}(c) = D(c)$ and $\check{D}(c') = D(c')$ intersect but are not the same, Lemma 4.14 implies that $U(c) \prec_{\text{top}} U(c')$. Thus $U(c)$ can be pushed to $-\infty$ in $M_\nu \setminus U(c')$ and $U(c')$ can be pushed to $+\infty$ in $M_\nu \setminus U(c)$. As before, we use the fact that f takes $M_\nu \setminus U(c)$ to $N \setminus \mathbb{T}(c)$, and similarly for c' , to conclude that $\mathbb{T}(c)$ and $\mathbb{T}(c')$ can be pushed to $-\infty$ and $+\infty$, respectively, in the complement of each other. It follows (Lemma 3.1) that $\mathbb{T}(c) \prec_{\text{top}} \mathbb{T}(c')$ or equivalently. $f_c \prec_{\text{top}} f_{c'}$.

Suppose that c is nonannular but c' is annular. Since $c \prec_c c'$ and the domains overlap, we may apply Lemma 4.14 to conclude that $\widehat{F}_c \prec_{\text{top}} U(c')$. Now again since f takes $M_\nu \setminus U(c)$ to $N \setminus \mathbb{T}(c)$, we may conclude, using

the same argument as in Case 1b above, that $f_c \prec_{\text{top}} \mathbb{T}(c')$, or equivalently $f_c \prec_{\text{top}} f_{c'}$. The case where c is annular is similar. \square

8.4. Controlling complementary regions

In this section and Section 8.5, we will assume that ρ is *doubly degenerate*, i.e. that there are no nonperipheral parabolics or geometrically finite ends, $N = \widehat{C}_N$, and the model has no boundary blocks. The difference between this and the general case essentially involves taking care with notation and boundary behavior, and in Section 8.6 we will explain how to address these issues.

From now on, assume d_1 has been chosen to be at least as large as the constant d_1 given by Lemma 8.4, *and* the constant d_1 given by Proposition 5.9 when $k = k_u$. Let C be a $(d_1, 3d_1)$ cut system, which exists as in §4, and which furthermore satisfies the condition that its annular cuts are exactly those annular slices a such that $|\omega(a)| > k_u$. Let C' be the (d_1, d_2) cut system obtained by applying Lemma 8.3 to C . Note that the annular slices of C' are the same as those of C .

Let W be (the closure of) a complementary region of the union of \widehat{F}_c and $U(c)$ for all slices c in C' . Note that, because of our choice of annular slices, W is also (the closure of) a complementary region of the union of nonannular surfaces $\{\widehat{F}_c[k_u] : c \in C'\}$ and solid tori $\mathcal{U}[k_u]$. That is, $\text{int}(W)$ is the closure of a connected component of \mathcal{W}_{k_u} as defined in §5.5. By Proposition 5.9, $\text{int}(W) \cap M_\nu[0]$ is a component of \mathcal{W}_0 , and in particular every block in it has the same address by Lemma 5.7. The number of blocks in W is uniformly bounded by Lemma 5.8.

Let Σ be the scaffold in M_ν whose surfaces \mathcal{F}_Σ are those cut surfaces $\widehat{F}_c[k_u]$ ($c \in C'$) that meet ∂W , and whose solid tori \mathcal{V}_Σ are the closures of \mathcal{U}_Σ , which are those components of $\mathcal{U}[k_u]$ whose closures meet ∂W .

Lemmas 4.5 and 4.14 imply that $\prec_{\text{top}}|_\Sigma$ satisfies the overlap condition, and by Lemma 4.15, the transitive closure of $\prec_{\text{top}}|_\Sigma$ is a partial order. Hence Σ is combinatorially straight.

We want to consider $f'|_\Sigma$ as a good scaffold map. The first step is to identify M_ν with N by an orientation-preserving homeomorphism in the homotopy class of f' , so that from now on we may consider f' to be homotopic to the identity. By Lemma 8.3, f' is an embedding on \mathcal{F}_Σ , and the images of components of \mathcal{F}_Σ are all disjoint. $f'(\mathcal{V}_\Sigma)$ is a subcollection of the closed Margulis tubes $\overline{\mathbb{T}}[k_u]$ which we denote $\overline{\mathbb{T}}_\Sigma$, and is unknotted and unlinked by Otal's theorem. Hence $f'(\Sigma) = f'(\mathcal{F}_\Sigma) \cup \overline{\mathbb{T}}_\Sigma$ is a scaffold.

Finally, Lemma 8.4 tells us that $f'|_\Sigma$ is order preserving. To see this, let p and q be two overlapping pieces of Σ and let us show that $p \prec_{\text{top}} q \iff f'(p) \prec_{\text{top}} f'(q)$. p and q are components of $\widehat{F}_c[k_u]$ and $\widehat{F}_{c'}[k_u]$ for two slices $c, c' \in C'$, respectively (where if p or q is a tube then the corresponding slice is annular). The overlap implies that $\check{D}(c)$ and $\check{D}(c')$ overlap, and hence c and c' are \prec_c -ordered by Lemma 4.5. If $c \prec_c c'$ then $\widehat{F}_c \prec_{\text{top}} \widehat{F}_{c'}$ by Lemma

4.14, and $f_c \prec_{\text{top}} f_{c'}$ by Lemma 8.4. For the components p and q this implies $p \prec_{\text{top}} q$ and $f'(p) \prec_{\text{top}} f'(q)$. If $c' \prec_c c$ then the opposite orders hold in both the model and the image. Therefore $f'|_{\Sigma}$ is order preserving.

This establishes all the properties of Definition 3.7, and hence $f'|_{\Sigma}$ is a good scaffold map.

By the properties of the model map, we also know that $f'(M_\nu \setminus \mathcal{U}_\Sigma)$ is contained in $N \setminus \mathbb{T}_\Sigma$, and that f' is proper and has degree 1. We can therefore apply Theorem 3.8 (Scaffold Extension) to find a homeomorphism of pairs

$$f'' : (M_\nu, \mathcal{V}_\Sigma) \rightarrow (N, \overline{\mathbb{T}}_\Sigma)$$

which agrees with f' on \mathcal{F}_Σ and is homotopic to it, rel \mathcal{F}_Σ , on each component of \mathcal{V}_Σ (through proper maps to the corresponding component of $\overline{\mathbb{T}}_\Sigma$).

We can now use the existence of f'' to obtain maps with geometric control. We will find maps Φ and Ψ from a neighborhood of W to N homotopic to $f''|_W$, such that:

- Φ is an embedding, agrees with f'' on \mathcal{F}_Σ , is isotopic to f'' on $\partial\mathcal{V}_\Sigma$ rel \mathcal{F}_Σ , satisfies a uniform bilipschitz bound on a uniform bicollar of ∂W , and respects the horizontal foliations on $\partial\mathcal{V}_\Sigma$ and $\partial\mathbb{T}_\Sigma$.
- Ψ agrees with Φ on ∂W , satisfies a uniform bilipschitz bound on a uniform bicollar of ∂W , and is uniformly Lipschitz in W .

Here a “uniform bound” is a bound independent of any of the data except the topological type of S . A uniform bicollar is the image of a piecewise-smooth embedding of $\partial W \times [-1, 1]$ into N with uniform bilipschitz bounds, so that $\partial W \times \{0\}$ maps to ∂W and $\partial W \times [0, 1]$ maps to W . Recall that the horizontal foliation on $\partial\mathcal{V}_\Sigma$ is the foliation by Euclidean geodesic circles homotopic to the cores of the constituent annuli, and the geodesic circles homotopic to their images form the horizontal foliation of $\partial\mathbb{T}_\Sigma$.

We remark that Φ is an embedding but not Lipschitz, whereas Ψ is Lipschitz but not an embedding. Converting these two maps into a uniformly bilipschitz embedding will be our goal after constructing them.

Construction of Φ . To construct Φ from f'' , we begin with $\partial\mathcal{V}_\Sigma$. Let V denote a component of \mathcal{V}_Σ and \mathbb{T}_V its image under f'' . We claim that $f''|_{\partial V}$ is homotopic, through maps $\partial V \rightarrow \partial\mathbb{T}_V$, to a uniformly bilipschitz homeomorphism, where the homotopy is constant on $\mathcal{F}_\Sigma \cap \partial V$,

Consider first a component annulus A of $\partial V \setminus \mathcal{F}_\Sigma$. $f''|_A$ is an embedding into $\partial\mathbb{T}_V$ which is homotopic to $f'|_A$ rel boundary. The height of A in M_ν is uniformly bounded since W consists of boundedly many blocks by Lemma 5.8. Since f' is uniformly Lipschitz, this bounds the height of $f''(A)$ from above. Since f' on ∂A is a (uniformly) bilipschitz bicollared embedding, the height of $f''(A)$ is also uniformly bounded below. We conclude that $f''|_A$ is isotopic rel ∂A to a bilipschitz embedding with uniform constant. Since f'' already takes $\partial\mathcal{F}_\Sigma$ to geodesics in $\partial\mathbb{T}_\Sigma$ by Theorem 6.1, this bilipschitz embedding can be chosen to respect the horizontal foliations. We let $\Phi|_A$ be this embedding. Piecing together over all the components of $\partial V \cap \partial W \setminus \mathcal{F}_\Sigma$,

we obtain a map which is an embedding into $\partial\mathbb{T}_V$, because $\Phi(A) = f''(A)$ for each component A , and f'' is a homeomorphism.

Now consider the possibility that ∂U does not meet \mathcal{F}_Σ . We claim that the Euclidean tori ∂U and $\partial\mathbb{T}_U$ admit uniformly bilipschitz affine identifications with the standard torus $\mathbb{R}^2/\mathbb{Z}^2$. For ∂U this follows because it is composed of a bounded number of standard annuli (again because of the bound on the size of W). Since $f' : \partial U \rightarrow \partial\mathbb{T}_U$ is a Lipschitz map that is homotopic to a homeomorphism (namely $f''|_{\partial U}$), the diameter of $\partial\mathbb{T}_U$ is uniformly bounded from above, and on the other hand its area is uniformly bounded from below since it is an ϵ_1 -Margulis tube boundary (see e.g. [44, Lemma 6.3]). It follows that $\partial\mathbb{T}_U$ is also uniformly bilipschitz equivalent to the standard torus. These identifications conjugate $f'|_{\partial U}$ to a uniformly Lipschitz self-map of the standard torus which is homotopic to a homeomorphism. It is now elementary to check that such a map can be deformed to a uniformly bilipschitz affine map. In fact the homotopy can be chosen so that the tracks of all points are uniformly bounded.

Thus we have defined Φ on ∂W , so that it is isotopic to $f''|_{\partial W}$ through maps taking ∂W to $f''(\partial W)$, and constant on \mathcal{F}_Σ .

In order to extend the map to W , we observe first that ∂W has a uniform bicollar in M_ν (by the explicit construction of the model manifold), and next that $\Phi(\partial W) = f''(\partial W)$ also has a uniform bicollar in N . To see the latter, note that the cut surface images $f''(\mathcal{F}_\Sigma)$ have uniform bicollars which are just the images by f' of the collars $\{E_c\}$ in M_ν given by Lemma 8.3. The boundary tori $\partial\mathbb{T}_\Sigma$ have uniform bicollars because of the choice of ϵ_1 . These collars can be fitted together to obtain a uniform bicollar of all of $\Phi(\partial W)$ because the pieces of $\Phi(\partial W)$ fit together at angles that are bounded away from 0 (due to the uniformity of f' on the bicollars E_c). By standard methods we may now use the isotopy between f'' and Φ on ∂W to extend Φ to a map on W which is an embedding isotopic to f'' , with uniform bilipschitz bounds on a uniform subcollar of the boundary.

Construction of Ψ . We observe that the homotopy from f' to Φ on ∂W can be made to have uniformly bounded tracks, simply by taking the straight-line homotopy in the Euclidean metric on $\partial\mathcal{U}_\Sigma$. First note that the homotopy is constant except on $\partial W \cap \partial\mathcal{U}_\Sigma$. Let U be a component of \mathcal{U}_Σ . If ∂U does not meet \mathcal{F}_Σ it has bounded geometry, and the boundedness of tracks was already noted above. In general, the homotopy is constant by construction on a subset X of ∂U which is a uniformly bounded distance from any point in ∂U , and are homotopic rel X . If $y \in \partial U$ let α denote a shortest arc from y to X . The union of $f'(\alpha)$ and $\Phi(\alpha)$ has uniformly bounded length since both maps are uniformly Lipschitz, and this serves to bound the shortest homotopy from f' to Φ .

Now let Ξ be a uniform collar of ∂W , such that there is a uniformly bilipschitz homeomorphism $h : W \setminus \Xi \rightarrow W$ isotopic to the inclusion (this is possible because the geometry of W is uniformly bounded). Define $\Psi|_{W \setminus \Xi} =$

$f' \circ h$, then extend Ψ to Ξ using the bounded-track homotopy between $f'|_{\partial W}$ and $\Phi|_{\partial W}$. This map agrees with Φ on ∂W , and satisfies a uniform Lipschitz bound. Using the uniform collar structure for ∂W and $\Phi(\partial W)$, as in the construction of Φ , we can arrange for Ψ to also satisfy uniform bilipschitz bounds in a uniform collar of the boundary.

Uniformity via geometric limits. We now have a uniform bilipschitz embedding of ∂W which extends, by Φ , to an embedding without geometric control, and by Ψ to a uniformly Lipschitz map which may not be an embedding. We claim next that $\Phi|_{\partial W} = \Psi|_{\partial W}$ can be extended to an embedding of W in N with *uniform bilipschitz constant*.

If this is false, then there is a sequence of examples $\{(M_{\nu_n}, W_n, N_n)\}$ where the best bilipschitz constant goes to infinity (we index our maps as Φ_n, Ψ_n , etc.). We shall reach a contradiction by extracting a geometric limit.

As before, W_n contain a bounded number of blocks. Since the tubes whose interiors meet W_n must have bounded coefficient $|\omega| < k_u$, we may assume, after restricting to a subsequence, that they have the same combinatorial structure and tube coefficients. After applying a sequence of homeomorphisms to the model manifolds we may assume that the W_n 's are all equal to a fixed W . Choose a basepoint $x \in W$ and let $y_n = \Psi_n(x)$. After taking subsequences we may assume that $(N_n, y_n) \rightarrow (N_\infty, y_\infty)$ geometrically, and $\Psi_n \rightarrow \Psi_\infty$ geometrically (the latter because of the uniform Lipschitz bounds on Ψ_n).

Because $\Psi_n|_{\partial W}$ are uniformly bicollared embeddings, their limit $\Psi_\infty|_{\partial W}$ is an embedding.

Since $\Psi_\infty(W)$ is a compact 3-chain with boundary $\Psi_\infty(\partial W)$, we know that $\Psi_\infty(\partial W)$ bounds some compact region $W'_\infty \subset N_\infty$. Similarly let W'_n be the compact region bounded by $\Psi_n(\partial W)$ (note that $W'_n = \Phi_n(W)$).

By definition of geometric convergence, given R and n large enough there is a map $h_n : \mathcal{N}_R(y_n) \rightarrow N_\infty$ which is an embedding with bilipschitz constant going to 1, and taking the basepoints y_n to y_∞ . Geometric convergence of the maps means, taking R larger than the diameter of W'_∞ , that $h_n \circ \Psi_n$ converge pointwise to Ψ_∞ .

In fact we can arrange things so that eventually $h_n \circ \Psi_n = \Psi_\infty$ on the boundary: note that $\partial W'_\infty$ is composed of finitely many pieces (images of cut surfaces and annuli in Margulis tubes) which are C^2 -embedded, and meet transversely along boundary circles. Thus it has a collar neighborhood which is smoothly foliated by intervals which $\partial W'_\infty$ intersects transversely. Since the convergence of $h_n \circ \Psi_n(\partial W)$ is C^2 on each cut surface and annulus piece, they are eventually transverse to this foliation too, and hence after adjusting h_n by small isotopies of this collar neighborhood we may assume that $h_n \circ \Psi_n = \Psi_\infty$ on ∂W . With this adjustment, we have $h_n(W'_n) = W'_\infty$, with h_n still satisfying a uniform bilipschitz bound.

Now given (large enough) m we note that the embedding $\Phi_m : W \rightarrow W'_m$ can be assumed to be bilipschitz with *some* constant depending on m . Fix

a value of m , and let $g_\infty = h_m \circ \Phi_m$. This is a K -bilipschitz embedding of W to W'_∞ , for some K_m , which restricts to Ψ_∞ on the boundary. Finally, let $g_n = h_n^{-1} \circ g_m$. Fixing m and letting n vary, we have a *uniformly* bilipschitz sequence of embeddings taking W to W'_n and restricting to Ψ_n on the boundary, and so contradicts the choice of sequence.

With this contradiction we therefore conclude that in fact there is a uniformly bilipschitz extension of $\Phi|_{\partial W}$ to W , as desired. Denote this map by $\Theta : W \rightarrow N$.

Degree of the map. We claim that Θ maps with degree 1 onto its image. Consider first the case that \mathcal{V}_Σ is non-empty, and let A be the intersection of $\partial\mathcal{V}_\Sigma$ with ∂W . The map f'' , since it is globally defined and of degree 1, must map A with degree 1 (and homeomorphically) to its image in $\partial\mathbb{T}_\Sigma$. Since Θ is isotopic to f'' on A , it also must map with degree 1. Any embedding of oriented manifolds $g : X \rightarrow Y$ which maps a nonempty subset of ∂X with degree 1 to its image in ∂Y must have degree 1 to its image in Y . Applying this to $\Theta : W \rightarrow N \setminus \mathbb{T}_\Sigma$, we conclude that Θ has degree 1 to its image.

Now if \mathcal{V}_Σ is empty, W only meets components of \mathcal{F}_Σ , and hence these components must have no nonperipheral boundary. Thus W is the region between two slices $\widehat{F}_c[k_u]$ and $\widehat{F}_{c'}[k_u]$ with domain equal to all of S . Assume that $\widehat{F}_c \prec_{\text{top}} \widehat{F}_{c'}$. If Θ does not have degree 1, it must switch the order of the boundaries, that is $\Theta(\widehat{F}_{c'}) \prec_{\text{top}} \Theta(\widehat{F}_c)$. Since Θ is equal to f' on these surfaces, this contradicts Lemma 8.4.

Putting together the maps. The embeddings Θ_W can be pieced together over all regions W to yield a global map $F : M_\nu[k_u] \rightarrow N$. This is because different regions meet only along the cut surfaces $\widehat{F}_c[k_u]$ ($c \in C'$), and on these each Θ_W is equal to the original f' .

For each tube U in $\mathcal{U}[k_u]$, $F|_{\partial U}$ is homotopic to $f'|_{\partial U}$ through maps to $\partial\mathbb{T}_U$: this is because for each region W the homotopy from Θ_W to f' is constant on the boundary circles of $\partial W \cap \partial U$, so the homotopies can be pieced together. Thus since f' was defined on U we can extend F to U (without any geometric control at this point). The resulting map $F : M_\nu \rightarrow \widehat{C}_N$ takes $\mathcal{U}[k_u]$ to $\mathbb{T}[k_u]$, and $M_\nu[k_u]$ to $N \setminus \mathbb{T}[k_u]$.

We next check that F is in the right homotopy class. This is not automatic, because in the geometric limiting step that produced the maps Θ_W , we did not keep track of homotopy class. However, we note that F agrees with f' on each of the cut surfaces, and is homotopic to f' on the union of the cut surfaces with the tube boundaries $\partial\mathcal{U}$.

Since we are in the doubly degenerate case g_H is infinite, and hence there exists a slice $c \in C'$ with $D(c) = S$. We then have a cut surface F_c which projects to S minus the collar of a pants decomposition. The missing annuli can be found on the boundaries of the tubes adjacent to ∂F_c . Adjoining these to F_c , we find a surface $S' \subset M_\nu$ which projects to all of S . Thus

M_ν is homotopy-equivalent to S' , and since $F|_{S'}$ is homotopic to $f'|_{S'}$ we conclude that F is homotopic to f' .

Note that F is a proper map, since the cut surfaces and tubes cannot accumulate in N , and the diameters of images of the regions W are uniformly bounded. Thus F has a well-defined degree. Since each Θ_W has degree 1 to its image, F has positive degree. Since it is a homotopy equivalence, the degree must be 1. The restriction to $M_\nu[k_u]$ is then a uniformly bilipschitz (with respect to path metrics) orientation-preserving homeomorphism to $N \setminus \mathbb{T}[k_u]$.

8.5. Control of Margulis tubes

It remains to adjust F on the tubes $\mathcal{U}[k_u]$ so that it is a global bilipschitz homeomorphism.

If \mathbb{T} is a hyperbolic tube with marking (α, μ) (where μ is a meridian and α represents the core curve; see §2.6 and [47]) we let the α -foliation be the foliation of $\partial\mathbb{T}$ whose leaves are Euclidean geodesics in the homotopy class of α .

Lemma 8.5. *Let \mathbb{T}_1 and \mathbb{T}_2 be hyperbolic ϵ_1 -Margulis tubes with markings (α_1, μ_1) and (α_2, μ_2) (where μ_i are meridians and α_i are minimal representatives of the core curves), and let $h : \partial\mathbb{T}_1 \rightarrow \partial\mathbb{T}_2$ be a marking-preserving K -bilipschitz homeomorphism which takes the α_1 -foliation of $\partial\mathbb{T}_1$ to the α_2 -foliation of $\partial\mathbb{T}_2$. Suppose that the radii of the tubes are at least $a > 0$.*

Then h can be extended to a K' -bilipschitz homeomorphism $\hat{h} : \mathbb{T}_1 \rightarrow \mathbb{T}_2$, where K' depends on K and a .

Proof. It will be convenient to recall Fermi coordinates (z, r, θ) around a geodesic, where z denotes length along the geodesic and (r, θ) are polar coordinates in orthogonal planes. The hyperbolic metric is given by

$$\cosh^2 r dz^2 + dr^2 + \sinh^2 r d\theta^2. \quad (8.3)$$

This metric descends to any hyperbolic tube quotient (where the geodesic $(z, 0, 0)$ descends to the core) in the usual way.

We begin by extending h to all but bounded neighborhoods of the cores of the tubes.

Let $r_i \geq a$ be the radius of \mathbb{T}_i , and m_i the length of its meridian. Because h is marking-preserving we have $m_1/m_2 \in [1/K, K]$, and hence

$$\sinh r_1 / \sinh r_2 \in [1/K, K] \quad (8.4)$$

since $\sinh r_i = m_i$ using (8.3).

By hypothesis, $r_1, r_2 > a$. Letting $\mathbb{T}_i(r)$ denote the r -neighborhood of the core in \mathbb{T}_i , we extend h to a map

$$h_1 : \mathbb{T}_1 \setminus \mathbb{T}_1(a/2) \rightarrow \mathbb{T}_2 \setminus \mathbb{T}_2(a/2)$$

using the foliations \mathcal{R}_i of \mathbb{T}_i minus its core by geodesics perpendicular to the core. More precisely, choose an increasing K' -bilipschitz homeomorphism $s : [a/2, r_1] \rightarrow [a/2, r_2]$ satisfying $\sinh s(r) / \sinh r \in [1/K', K']$, where K'

depends on K and a (one can easily do this with an affine map s , using a comparison of $\sinh(x)$ to $e^x/2$). Let h_1 be the unique extension of h which takes \mathcal{R}_1 to \mathcal{R}_2 and takes $\partial\mathbb{T}_1(r)$ to $\partial\mathbb{T}_2(s(r))$. The projection $\partial\mathbb{T}_i(r) \rightarrow \partial\mathbb{T}_i(a/2)$ along the foliation \mathcal{R}_i is (uniformly) quasiconformal and contracts by a factor between $\cosh(r)/\cosh(a/2)$ and $\sinh(r)/\sinh(a/2)$ (using (8.3)). Thus, the properties of s imply that the extension is bi-Lipschitz.

It remains to extend h_1 to $h_2 : \mathbb{T}_1(a/2) \rightarrow \mathbb{T}_2(a/2)$. The restriction of h_1 to $\partial\mathbb{T}_1(a/2)$ is bilipschitz with constant $K''(K, a)$, and we note that $\partial\mathbb{T}_1(a/2)$ is a torus with *bounded diameter*. This is true because both generators in the boundary markings are bounded at radius $a/2$: α was bounded at radius r_1 because it is a minimal representative of the generator and the tubes are Margulis tubes in a surface group (see Lemma 6.3 of [44]), and it only gets shorter inside. The meridian length at radius $a/2$ is bounded automatically by $2\pi \sinh a/2$, via (8.3).

We then use the following lemma:

Lemma 8.6. *Let T be a Euclidean torus of diameter at most 1. Let $f : T \rightarrow T$ be a K -bilipschitz homeomorphism homotopic to the identity, which preserves a linear foliation on T . Then there exists a map*

$$F : T \times [0, 1] \rightarrow T \times [0, 1]$$

such that $F(\cdot, 0) = id$ and $F(\cdot, 1) = f$, and F is K' -bilipschitz for K' depending only on K .

Remark: One would expect that the condition of preserving a linear foliation is not necessary in this lemma. However this seems to be a nontrivial matter. Luukkainen [36] has proven such a “bi-Lipschitz isotopy” lemma when f is a self-map of \mathbb{R}^n with a bound on $d(x, f(x))$ for $x \in \mathbb{R}^n$, building on work of Sullivan, Tukia, and Väisälä [70, 71, 61]. One could try to obtain the result for the torus by considering the universal cover, but getting equivariance for the isotopy with control of the bi-Lipschitz constant seems to be difficult.

At any rate with our added condition the proof is elementary:

Proof of Lemma 8.6. Consider first this one-dimensional version: Let $h : \mathbb{R} \rightarrow \mathbb{R}$ be a K -bilipschitz homeomorphism satisfying also $|h(s) - s| < C$ for all $s \in \mathbb{R}$. The map

$$H(s, t) = ((1 - t)h(s) + ts, t)$$

is then a homeomorphism from $\mathbb{R} \times [0, 1]$ to itself satisfying a K' -bilipschitz bound (where K' depends on K and C) and such that $H(\cdot, 0) = h$ and $H(\cdot, 1) = id$.

Now given our map f , let F be a lift of f to \mathbb{R}^2 . Since f is homotopic to the identity and $diam(T) \leq 1$, F can be chosen so that $|F(p) - p| \leq K + 2$ for all $p \in \mathbb{R}^2$. F preserves a foliation which we can assume is the horizontal foliation, so we can express it as

$$F(x, y) = (\xi(x, y), \eta(y))$$

with $\eta : \mathbb{R} \rightarrow \mathbb{R}$ K -bilipschitz, and $\xi(x, y)$ K -bilipschitz in x for each y , and K -Lipschitz in y for each x .

Now after applying the one-dimensional case to η we may assume $\eta(y) = y$, and applying it again to $\xi(x, y)$ for each fixed y , we have our desired bilipschitz isotopy.

Since this construction is evidently invariant under isometries of \mathbb{R}^2 , it can be projected back to the torus T . \square

Using this lemma we can extend h_1 to a K''' -bilipschitz homeomorphism from the collar $\mathbb{T}_1(a/2) \setminus \mathbb{T}_1(a/4)$ to $\mathbb{T}_2(a/2) \setminus \mathbb{T}_2(a/4)$, so that on the inner boundary it is an *affine map* in the Euclidean metric. We can then extend the map, again using the radial foliation, to the rest of the solid torus. The bilipschitz control in this last step follows from a simple calculation in the Fermi coordinates (8.3), and depends on the fact that the map on $\partial\mathbb{T}_1(a/4)$ is affine. It does not hold for a general bilipschitz boundary map; this was the reason we needed to apply Lemma 8.6. \square

Our model map, restricted to the boundary of each model tube, satisfies the conditions of Lemma 8.5. (Note that the condition of preserving a linear foliation was supplied in the construction, which respected the horizontal foliations on model tube boundaries and their images.) Thus we have the desired bilipschitz extension.

The resulting map is now a locally bilipschitz homeomorphism from M_ν to \widehat{C}_N (which in the doubly degenerate case is all of N). Thus it is globally bilipschitz, and the Bilipschitz Model Theorem is established, in the doubly degenerate case.

8.6. The mixed-end case

We will now consider the case of a Kleinian surface group that is not necessarily doubly degenerate.

The boundary blocks of M_ν , as described in §2.6, have *outer boundaries* which are the boundary components of M_ν . These outer boundaries behave essentially like cuts in a cut system. In particular in the proofs of Lemmas 5.1 and 5.2 we observe that their topological ordering properties in $M_\nu \subset \widehat{S} \times \mathbb{R}$ are as we would expect – i.e. an outer boundary associated to a top boundary block lies above all overlapping cut surfaces, and vice versa for a bottom boundary block. The set $\mathcal{X}(\emptyset)$ of blocks with address $\langle \emptyset \rangle$ is nonempty in the case with boundary, and in fact contains all of the boundary blocks (see §5.3).

Theorem 6.1 provides us with uniform collars for the cut surfaces that lie in \widehat{C}_N , at a distance of at least a from the boundary, where a is a uniform constant. The original model map $f : M_\nu \rightarrow \widehat{C}_N$ is already K -bilipschitz on the boundaries. Because each boundary component has a uniform collar in M_ν and in \widehat{C}_N , we may adjust the map to satisfy a uniform bilipschitz bound in these collars. We may assume that the uniformly embedded collar

obtained in \widehat{C}_N is within an α -neighborhood of the boundary. Thus the collars of the cut surfaces are disjoint from the boundary collars.

This tells us that the topological ordering of overlapping cut surfaces and boundary surfaces is preserved by the adjusted model map f' (generalizing Lemma 8.4).

The argument in §8.4 controlling the map on complementary regions requires a few remarks. The complementary regions contained in $\mathcal{X}(\emptyset)$ will have outer boundary components in their boundary, so these should be taken as components of \mathcal{F}_Σ for the scaffold Σ . The map f'' should take $(M_\nu, \mathcal{V}_\Sigma)$ to $(\widehat{C}_N, \overline{\mathbb{T}}_\Sigma)$, and again the appeal to Theorem 3.8 (Scaffold Extension) is by way of first identifying the interiors of both manifolds with $S \times \mathbb{R}$. The homotopy from f' to f'' can be assumed to be constant on the uniform collars of the outer boundaries. The same holds for the construction of Φ and Ψ , on the regions contained in $\mathcal{X}(\emptyset)$. In the geometric limit step, when W_n contain boundary blocks we cannot assume that they are all identical after passage to a subsequence. Boundary blocks do have finitely many combinatorial types, so we may assume that these are constant on a subsequence. The geometry of a block can degenerate: the curves of $\mathbf{I}(H)$ or $\mathbf{T}(H)$ supported on the block can have lengths going to zero. The geometric limit of a sequence of such blocks can be described as a union of blocks based on smaller subsurfaces, where the curves whose lengths vanish give rise to parabolic tubes in the limit. Thus we may assume that the W_n minus these tubes are eventually combinatorially equivalent to a fixed W and geometrically converge to it. This suffices to make the argument work.

Section 8.5 on the extension of the bilipschitz map to Margulis tubes goes through without change, noting that in the general case there may be parabolic tubes that are not associated to ∂S , but that extension to these is no harder. Thus, we obtain a bilipschitz homeomorphism of degree 1

$$F : M_\nu \rightarrow \widehat{C}_N.$$

Checking that F is homotopic to f is again done by exhibiting a surface S' in M_ν which projects to S and on which F is known to be homotopic to f . In the general case there may not be a single slice c in the cut system with $D(c) = S$; however we can piece S' together from slices and outer boundaries. Let P_+ denote the annuli corresponding to parabolics facing the top of the compact core, as in §2.5. A component Z of $S \setminus P_+$ is either associated to a top outer boundary of M_ν , or supports a filling lamination component of ν_+ , and hence a forward-infinite geodesic in H . In the latter case there is a cut c_Z with domain Z and an associated surface \widehat{F}_{c_Z} . the union of these boundary surfaces and cut surfaces, joined together with annuli along the parabolic model tubes associated to P_+ , gives our desired surface S' , and the argument goes through as in the doubly degenerate case. This gives the desired map from M_ν to \widehat{C}_N . Since the map has not changed on ∂M_ν , we can use the same extension to the exterior E_ν as given in Theorem 2.10, so

that we obtain the desired map from \overline{ME}_ν to \bar{N} . This completes the proof of the Bilipschitz Model Theorem.

9. Proofs of the main theorems

9.1. The Ending Lamination Theorem

At this point the Ending Lamination Theorem follows via an application of Sullivan's rigidity theorem [62].

Given $\rho_1, \rho_2 \in \mathcal{D}(S)$ with the same end invariants ν , we obtain from the Bilipschitz Model Theorem the homeomorphisms $\bar{F}_i : \overline{ME}_\nu \rightarrow \bar{N}_i$. The composition $f_2 \circ f_1^{-1}$ lifts to a K -bilipschitz homeomorphism of \mathbb{H}^3 that conjugates ρ_1 to ρ_2 , and furthermore (by the properties of \bar{F}_i on $\partial_\infty ME_\nu$) is conformal on the domain of discontinuity. Sullivan's theorem implies that this map is conformal on the sphere at infinity, and it follows that it is homotopic to an isometry on the interior.

9.2. Proof of the Length Bound Theorem

The Short Curve Theorem (§2.7) already contains part (1) of the Length Bound Theorem. It remains to prove part (2), the length estimate,

$$d_{\mathbb{H}^2} \left(\omega(v), \frac{2\pi i}{\lambda(\rho(v))} \right) \leq c$$

for a uniform c , whenever v is a vertex in H_ν . For simplicity we suppress v in the proof, writing ω , λ , etc.

Suppose first that $|\omega| \in (k, \infty)$ where k is the constant in the Bilipschitz Model Theorem. Then the tube $U = U(v)$ is in $\mathcal{U}[k]$ and F takes U to the corresponding Margulis tube $\mathbb{T}(v)$ by a K -bilipschitz map. This restricts to a map on the boundary tori which respects their natural markings. Letting $\omega_{\mathbb{T}}$ denote the Teichmüller parameter of $\mathbb{T}(v)$ with its natural marking, it follows that

$$d_{\mathbb{H}^2}(\omega, \omega_{\mathbb{T}}) \leq \log K \tag{9.1}$$

where \mathbb{H}^2 , the upper half plane, is identified with the Teichmüller space of the torus, and $d_{\mathbb{H}^2}$ is the Teichmüller metric and the Poincaré metric.

Now let $\omega_\infty = 2\pi i/\lambda$, and let $\lambda_{\mathbb{T}} = 2\pi i/\omega_{\mathbb{T}}$. Equation (2.3), which comes from §3.2 of [47], says

$$\lambda = h_r(\lambda_{\mathbb{T}}) \tag{9.2}$$

where r is the radius of $\mathbb{T}(v)$ and where $h_r(z) = \operatorname{Re} z \tanh r + i \operatorname{Im} z$. Note that h_r preserves the right half plane $\mathbb{H}' = \{z : \operatorname{Re} z > 0\}$ and, letting $d_{\mathbb{H}'}$ denote the Poincaré metric on \mathbb{H}' , $d_{\mathbb{H}'}(z, h_r(z))$ is uniformly bounded above by a constant $C(r)$ which shrinks as r grows. Now the bilipschitz correspondence between $U(v)$ and $\mathbb{T}(v)$ tells us that r grows with the radius of $U(v)$, and this is equal to $\sinh^{-1}(\epsilon_1 |\omega|/2\pi)$ by (2.2). Since $|\omega| > k$ we obtain a uniform bound on the translation distance of h_r , and hence by (9.2) on $d_{\mathbb{H}'}(\lambda, \lambda_{\mathbb{T}})$. Now since the map $z \rightarrow 2\pi i/z$ is an isometry in the

Poincaré metric from \mathbb{H}^1 to \mathbb{H}^2 , we conclude that there is a uniform bound on $d_{\mathbb{H}^2}(\omega_\infty, \omega_\mathbb{T})$. Together with (9.1), we have the desired distance bound.

If $|\omega(v)| \leq k$ then we have uniform lower and upper bounds on $\lambda(\rho(v))$ by the Short Curve Theorem (§2.7), and the estimate is immediate.

10. Corollaries

In this section we give proofs of the corollaries mentioned in the introduction. All of these corollaries have analogues in the general setting which will be discussed in the sequel [15].

Our first corollary is the resolution of the Bers-Sullivan-Thurston Density conjecture for $AH(S)$.

Corollary 10.1. (Density Theorem) *The set of quasifuchsian surface groups is dense in $AH(S)$.*

Proof. The Ending Lamination Theorem asserts that a Kleinian surface group is determined by its ending invariants. Ohshika [52] used Thurston's Double Limit Theorem [65] to prove that every allowable collection of end invariants arises as the end invariants of a surface group which is a limit of quasifuchsian groups. The Density Theorem follows. \square

The proof of our rigidity theorem is somewhat more involved as we must observe that a topological conjugacy can detect the ending invariants.

Corollary 10.2. (Rigidity Theorem) *If two Kleinian surface groups ρ_1 and ρ_2 are conjugate by an orientation-preserving homeomorphism ϕ of $\widehat{\mathbb{C}}$, then they are quasiconformally conjugate. Moreover, if ϕ is conformal on $\Omega(\rho_1)$, then ϕ is conformal.*

Proof. Let $\phi : \widehat{\mathbb{C}} \rightarrow \widehat{\mathbb{C}}$ be the orientation-preserving homeomorphism conjugating ρ_1 to ρ_2 (i.e. $\phi \circ \rho_1(g) \circ \phi^{-1} = \rho_2(g)$ for all $g \in \pi_1(S)$). Since $\phi(\Omega(\rho_1)) = \Omega(\rho_2)$, it induces a homeomorphism between $\partial_\infty N_1$ and $\partial_\infty N_2$, where $N_i = \mathbb{H}^3/\rho_i(\pi_1(S))$. Ahlfors' Finiteness Theorem [4] assures us that $\partial_\infty N_i$ is a Riemann surface of finite type, so we may deform ϕ so that it is quasiconformal on $\Omega(\rho_1)$. One may use the Measurable Riemann Mapping Theorem [3] to construct a quasiconformal map $\psi : \widehat{\mathbb{C}} \rightarrow \widehat{\mathbb{C}}$ such that $\rho'_2 = \psi \circ \rho_2 \circ \psi^{-1}$ is a Kleinian surface group and $\psi \circ \phi$ is conformal on $\Omega(\rho_1)$.

So, ρ_2 is quasiconformally conjugate to a Kleinian surface group ρ'_2 and there exists an orientation-preserving homeomorphism $\psi \circ \phi$ conjugating ρ_1 to ρ'_2 which is conformal on $\Omega(\rho_1)$. Therefore, it suffices to prove that, in this situation, $\psi \circ \phi$ is conformal.

For the remainder of the argument we will assume that ϕ is conformal on $\Omega(\rho_1)$. Let (K_1, Q_1) be a relative compact core for N_1^0 and let (K_2, Q_2) be a relative compact core for N_2^0 . Since ϕ identifies $\rho_1(\pi_1(S))$ with $\rho_2(\pi_1(S))$, it induces a homotopy equivalence $\bar{\phi}$ from K_1 to K_2 . Recall that $\rho_i(g)$ is parabolic if and only if it has exactly one fixed point in $\widehat{\mathbb{C}}$. Therefore, $\rho_1(g)$

is parabolic if and only if $\rho_2(g)$ is parabolic. Thus, $\bar{\phi}(Q_1)$ is homotopic to Q_2 .

The end invariants of ρ_i consist of Q_i together with for each component R of $\partial K_i - Q_i$, either a conformal structure, i.e. a point in $\mathcal{T}(R)$, or an ending lamination, i.e. a point in $\mathcal{EL}(R)$. Let G_i be the union of the components of $\partial K_i - Q_i$ which are associated to geometrically finite ends of N_i . One may identify G_i with $\partial_\infty N_i$ and assume that $\bar{\phi}$ is conformal on G_1 .

It is easy to check directly that one may deform $\bar{\phi}$ to a pared homeomorphism from (K_1, Q_1) to (K_2, Q_2) . This fact also follows from Johannson's version of Waldhausen's Theorem (Proposition 3.4 in [31]) or can be derived from the Scaffold Isotopy Theorem.

Suppose that λ is a minimal geodesic lamination on S which is not a closed curve. Let $R(\lambda)$ be the minimal essential surface containing λ . The lamination λ is the ending lamination of some end of N_i^0 if and only if there exists a sequence $\{\gamma_j\}$ of simple closed curves on $R(\lambda)$ such that $\{[\gamma_j]\}$ converges in $\mathcal{PL}(S)$ to a measured lamination with support λ and the geodesic representatives $\{(\gamma_j^i)^*\}$ of γ_j in N_i leaves every compact subset of N_i .

For this paragraph, we will assume that we are working in the ball model for \mathbb{H}^3 . Let b_i be the projection of the origin to N_i . Given a simple closed non-peripheral curve γ on S , let $e_i(\gamma) = \infty$ if γ is homotopic to a cusp in N_i . Otherwise, let $e_i(\gamma)$ be the distance from the geodesic representative $(\gamma^i)^*$ of γ in N_i to b_i . Then λ is the ending lamination of some end of N_i^0 if and only if there exists a sequence $\{\gamma_j\}$ of simple closed curves on $R(\lambda)$ such that $\{e_i(\gamma_j)\}$ converges to ∞ . Let $d_i(\gamma)$ denote the maximal distance between the fixed points of $\rho_i(g)$ where g is an element in the conjugacy class determined by γ . Notice that $\{e_i(\gamma_j)\}$ converges to ∞ if and only if $\{d_i(\gamma_j)\}$ converges to 0. However, since $\bar{\phi}$ takes the fixed point set of $\rho_1(g)$ to the fixed point set of $\rho_2(g)$, one easily checks that $\{d_1(\gamma_j)\}$ converges to 0 if and only if $\{d_2(\gamma_j)\}$ converges to 0. Therefore, λ is an ending lamination of an end of N_1^0 if and only if it is an ending lamination of an end of N_2^0 .

If λ is an ending lamination of an end of N_0^i , then $R(\lambda)$ is homotopic to exactly one component of $\partial K_i - Q_i$ unless (K_i, Q_i) is homeomorphic to $(S \times I, \partial S \times I)$. So, we may assume that $\bar{\phi}$ takes the component of $\partial K_1 - Q_1$ associated to λ to the component of $\partial K_2 - Q_2$ associated to λ . (If (K_1, Q_1) is homeomorphic to $(S \times I, \partial S \times I)$, we may have to alter our original $\bar{\phi}$ by reflecting in the I factor.) Therefore, $\bar{\phi}$ is a pared homeomorphism which preserves the ending invariants. The Ending Lamination Theorem then gives rise to a homeomorphism F from \bar{N}_1 to \bar{N}_2 which is conformal on ∂N_1 , an isometry on N_1 and in the homotopy class determined by $\bar{\phi}$.

Let $\phi' : \widehat{\mathbb{C}} \rightarrow \widehat{\mathbb{C}}$ be the map which is the extension of the lift $\tilde{F} : \mathbb{H}^3 \rightarrow \mathbb{H}^3$ of F . Then ϕ' is either conformal or anti-conformal and conjugates ρ_1 to ρ_2 . Notice that since fixed points of elements of $\rho_1(\pi_1(S))$ are dense in $\Lambda(\rho_1)$, ϕ and ϕ' agree on $\Lambda(\rho_1)$. If our initial map ϕ was conformal, then ϕ and ϕ' must

agree on $\Omega(\rho_1)$ and hence on \widehat{C} . Therefore, since ϕ is orientation-preserving and ϕ' is either conformal or anti-conformal, $\phi = \phi'$ is conformal. \square

We next turn our attention to:

Corollary 10.3. (Volume Growth Theorem) *If $\rho : \pi_1(S) \rightarrow PSL_2(\mathbb{C})$ is a Kleinian surface group and $N = \mathbb{H}^3/\rho(\pi_1(S))$, then for any x in the ϵ_1 -thick part of C_N ,*

$$\text{volume} \left(B_R^{\text{thick}}(x) \right) \leq cR^{d(S)},$$

where c depends only on the topological type of S .

Here recall

$$d(S) = \begin{cases} -\chi(S) & \text{genus}(S) > 0 \\ -\chi(S) - 1 & \text{genus}(S) = 0 \end{cases}$$

and $B_R^{\text{thick}}(x)$ is an R -neighborhood of x in the ϵ_1 -thick part of C_N , with respect to the path metric.

Proof. We can replace C_N by \widehat{C}_N , which contains it. The ϵ_1 -thick part of \widehat{C}_N is almost the same as $\widehat{C}_N \setminus \mathbb{T}[k]$; the latter may include some ϵ_1 -Margulis tubes with ω coefficients bounded (in terms of k and the bounds of the model theorem). Since all such tubes have uniformly bounded diameters and volumes, it suffices to prove the theorem for $\widehat{C}_N \setminus \mathbb{T}[k]$. Now since the Bilipschitz Model Theorem gives a uniformly bilipschitz homeomorphism of $M_\nu[k]$ to $\widehat{C}_N \setminus \mathbb{T}[k]$, it suffices to prove the theorem for $M_\nu[k]$. Finally, this is equivalent to proving the theorem for $M_\nu[0]$, again because the difference consists of tubes with bounded diameters and volumes. This is what we will do.

Fix a cut system C , and recall from §5.3 the definition of the product regions $\mathcal{B}(h) \subset M_\nu$ where $h \in H$ and $C|_h$ is nonempty. Each $\mathcal{B}(h)$ is isotopic to $D(h) \times I$ for an interval I , and is defined as the region between the first and last slices a_h and z_h in $C|_h$ (note that a_h or z_h could fail exist if h is infinite in the backward or forward direction, in which case $\mathcal{B}(h)$ is defined accordingly).

Define also $\mathcal{B}_0(h) = \mathcal{B}(h) \cap M_\nu[0]$. For $x \in M_\nu[0]$, let $\mathcal{N}_r(x)$ denote the r -neighborhood of x with respect to the path metric in $M_\nu[0]$.

We shall prove the following statement by induction on $d(D(h))$:

- (*) For any $h \in H$ with $C|_h \neq \emptyset$ and $d(D(h)) \geq 1$, given $x \in \mathcal{B}_0(h)$ the volume of $\mathcal{B}_0(h) \cap \mathcal{N}_r(x)$ is bounded by $cr^{d(D(h))}$, where c depends only on $d(D(h))$.

Note that the boundary of $\mathcal{B}(h)$ consists of the bottom and top slice surfaces \widehat{F}_{a_h} and \widehat{F}_{z_h} (these could be empty if h is infinite) together with tube-boundary annuli associated to $\partial D(h)$. Thus the frontier of $\mathcal{B}_0(h)$ in $M_\nu[0]$ is just the surfaces F_{a_h} and F_{a_z} , each of which have at most $-\chi(S)$ components, with uniformly bounded diameter. Since $\mathcal{B}_0(h) \cap \mathcal{N}_r(x)$ is contained in the union of r -neighborhoods, in the path metric of $\mathcal{B}_0(h)$, of x and

the frontier of $\mathcal{B}_0(h) \cap \mathcal{N}_r(x)$, this implies that proving (*) for $\mathcal{B}_0(h) \cap \mathcal{N}_r(x)$ is equivalent to proving it for the r -neighborhood of x in $\mathcal{B}_0(h)$ in the path metric of $\mathcal{B}_0(h)$, which we can denote $\mathcal{N}_{r,h}(x)$. We proceed to do this.

When $d(D(h)) = 1$, $D(h)$ can only be a one-holed torus or a 4-holed sphere. In this case, each block in $\mathcal{B}(h)$ is associated with exactly one edge of the geodesic h and is isotopic to a sub-product region $D(h) \times J$ – in particular it separates $\mathcal{B}(h)$. It follows immediately that

$$\text{vol}(\mathcal{N}_{r,h}(x)) \leq 2v_0r/r_0$$

where v_0 is an upper bound on the volume of a block and r_0 a lower bound on the distance between the top and bottom boundaries of a block. This establishes (*) for $d(D(h)) = 1$.

Now for $d(D(h)) = k > 1$, consider the set E of slices $e \in C$ such that $F_e \subset \mathcal{B}_0(h)$, and $d(D(e)) = k$. It follows from the definition of the function d that for each $e \in E$, $D(e)$ is equal to $D(h)$ minus a (possibly empty) disjoint union of annuli.

Therefore \widehat{F}_e is isotopic to $D(h) \times \{t\}$ in the product structure of $\mathcal{B}(h)$, and hence these surfaces are topologically ordered in $\mathcal{B}(h)$, so we can number them $\cdots \widehat{F}_{e_i} \prec_{\text{top}} \widehat{F}_{e_{i+1}} \cdots$ and let P_i denote the product region between \widehat{F}_{e_i} and $\widehat{F}_{e_{i+1}}$.

Each \widehat{F}_{e_j} separates P_i from P_k for $i < j \leq k$. Since $\widehat{F}_e \setminus F_e$ is a union of annuli in the tubes of M_ν , it follows that F_{e_j} also separates $P_i[0] = P_i \cap M_\nu[0]$ from $P_k[0] = P_k \cap M_\nu[0]$.

Note that r_0 is a lower bound on the distance in $P_i[0]$ between F_{e_i} and $F_{e_{i+1}}$, since the slices cannot have any pieces in common (Lemma 4.12) and hence must be separated by a layer at least one block thick. It follows that $\mathcal{N}_{r,h}(x)$ can meet at most $2r/r_0$ different regions $P_i[0]$.

It remains to estimate the volume of $P_i[0] \cap \mathcal{N}_{r,h}(x)$. Inside P_i there may be block regions $\mathcal{B}(m)$, where $m \in H$ such that, necessarily, $C|_m \neq \emptyset$ and $d(D(m)) < k$. Any slice surface \widehat{F}_b for $b \in C$ which meets $\text{int}(P_i)$ must in fact be contained in one of these $\mathcal{B}(m)$'s, and hence the complement in P_i of all such $\mathcal{B}(m)$ is contained in a complementary region of C , and hence has a uniformly bounded number of blocks by Lemma 5.8. This bounds its volume by a constant v_1 and also implies that the number of $\mathcal{B}(m)$ contained in P_i is bounded by some n_1 . For each $\mathcal{B}_0(m)$ we have, by induction,

$$\text{vol}(\mathcal{N}_{r,m}(y)) \leq cr^{d(D(m))} \leq cr^{k-1}$$

for any $y \in \mathcal{B}_0(m)$. Since the frontier of $\mathcal{B}_0(m)$ in P_i consists of a top and a bottom surface, each of bounded diameter, we may conclude that a bound of the form $c'r^{k-1}$ holds for $\mathcal{B}_0(m) \cap \mathcal{N}_{r,h}(x)$, whether $x \in \mathcal{B}_0(m)$ or not. Summing these over all $\mathcal{B}(m)$ in P_i and including the rest of P_i we have a bound

$$\text{vol}(P_i \cap \mathcal{N}_{r,h}(x)) \leq v_1 + n_1c'r^{k-1} \leq c''r^{k-1}$$

for some c'' depending only on $d(D(h))$. Now summing this over all the (at most $2r/r_0$) P_i that meet the r -neighborhood of x gives

$$\text{vol}(\mathcal{N}_{r,h}(x)) \leq c''' r^k$$

which establishes (*). (Note that the union of closures of P_i fills up $\mathcal{B}(h)$ except possibly if $h = g_H$ and the top or bottom boundary of M_ν is nonempty; at any rate the exterior of $\cup \bar{P}_i$ is contained in a single address region, and has bounded volume by Lemma 5.8.)

Now applying (*) with $h = g_H$ yields the desired growth estimate for $M_\nu[0]$. □

References

- [1] I. Agol, *Volume change under drilling*, *Geometry & Topology* **6** (2002), 905–916.
- [2] ———, *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.
- [4] Lars V. Ahlfors, *Finitely generated Kleinian groups*, *Amer. J. Math.* **86** (1964), 413–429.
- [5] J. Anderson and R. Canary, *Cores of hyperbolic 3-manifolds and limits of Kleinian groups*, *Amer. J. Math.* **118** (1996), 745–779.
- [6] J. Anderson, R. Canary, M. Culler, and P. Shalen, *Free Kleinian groups and volumes of hyperbolic 3-manifolds*, *J. Differential Geom.* **43** (1996), no. 4, 738–782.
- [7] J. Anderson, R. Canary, and D. McCullough, *The topology of deformation spaces of Kleinian groups*, *Ann. of Math. (2)* **152** (2000), no. 3, 693–741.
- [8] J. W. Anderson, *Intersections of topologically tame subgroups of Kleinian groups*, *Journal d'Analyse* **65** (1995), 77–94.
- [9] R. Benedetti and C. Petronio, *Lectures on hyperbolic geometry*, Springer-Verlag Universitext, 1992.
- [10] L. Bers, *Spaces of Kleinian groups*, Maryland conference in Several Complex Variables I, Springer-Verlag Lecture Notes in Math, No. 155, 1970, pp. 9–34.
- [11] F. Bonahon, *Bouts des variétés hyperboliques de dimension 3*, *Ann. of Math.* **124** (1986), 71–158.
- [12] B. Bowditch, *Intersection numbers and the hyperbolicity of the curve complex*, Preprint, Southampton.
- [13] ———, *Systems of bands in hyperbolic 3-manifolds*, Southampton preprint, 2003.
- [14] J. Brock and K. Bromberg, *On the density of geometrically finite Kleinian groups.*, Preprint 2002.
- [15] J. Brock, R. Canary, and Y. Minsky, *The classification of finitely-generated Kleinian groups*, in preparation.
- [16] K. Bromberg, *Projective structures with degenerate holonomy and the Bers density conjecture*, preprint.
- [17] D. Calegari and D. Gabai, *Shrinkwrapping and the taming of hyperbolic 3-manifolds*, Preprint.
- [18] R. Canary and Y. Minsky, *On limits of tame hyperbolic 3-manifolds*, *J. Differential Geom.* **43** (1996), 1–41.
- [19] R. D. Canary, *Ends of hyperbolic 3-manifolds*, *J. Amer. Math. Soc.* **6** (1993), 1–35.
- [20] ———, *A covering theorem for hyperbolic 3-manifolds and its applications*, *Topology* **35** (1996), 751–778.

- [21] R. D. Canary, D. B. A. Epstein, and P. Green, *Notes on notes of Thurston*, Analytical and Geometric Aspects of Hyperbolic Space, Cambridge University Press, 1987, London Math. Soc. Lecture Notes Series no. 111, pp. 3–92.
- [22] D. B. A. Epstein and A. Marden, *Convex hulls in hyperbolic space, a theorem of Sullivan, and measured pleated surfaces*, Analytical and Geometric Aspects of Hyperbolic Space, Cambridge University Press, 1987, London Math. Soc. Lecture Notes Series no. 111, pp. 113–254.
- [23] M. Freedman, J. Hass, and P. Scott, *Least area incompressible surfaces in 3-manifolds*, Invent. Math. **71** (1983), no. 3, 609–642.
- [24] U. Hamenstadt, *Train tracks and the Gromov boundary of the complex of curves*, preprint, 2004.
- [25] W. J. Harvey, *Boundary structure of the modular group*, Riemann Surfaces and Related Topics: Proceedings of the 1978 Stony Brook Conference (I. Kra and B. Maskit, eds.), Ann. of Math. Stud. 97, Princeton, 1981.
- [26] ———, *Modular groups and representation spaces*, Geometry of group representations (Boulder, CO, 1987), Amer. Math. Soc., 1988, pp. 205–214.
- [27] J. Hempel, *3-manifolds*, Annals of Math. Studies no. 86, Princeton University Press, 1976.
- [28] W. Jaco and H. Rubinstein, *Pl minimal surfaces in 3-manifolds*, J. Differential Geom. **27** (1988), 493–524.
- [29] W. H. Jaco, *Lectures on three-manifold topology*, CBMS Regional Conference Series no. 43, Amer. Math. Soc., 1980.
- [30] W. H. Jaco and P. B. Shalen, *Seifert fibered spaces in 3-manifolds*, vol. 21, Memoirs of the Amer. Math. Soc., no. 220, A.M.S., 1979.
- [31] K. Johannson, *Topology and combinatorics of 3-manifolds*, Lecture Notes in Mathematics, vol. 1599, Springer-Verlag, Berlin, 1995.
- [32] T. Jørgensen and A. Marden, *Algebraic and geometric convergence of Kleinian groups*, Math. Scand. **66** (1990), 47–72.
- [33] E. Klarreich, *The boundary at infinity of the curve complex and the relative Teichmüller space*, Preprint, <http://www.nasw.org/users/klarreich/research.htm>.
- [34] I. Kra, *On spaces of Kleinian groups*, Comment. Math. Helv. **47** (1972), 53–69.
- [35] R. S. Kulkarni and P. B. Shalen, *On Ahlfors’ finiteness theorem*, Adv. Math. **76** (1989), no. 2, 155–169.
- [36] J. Luukkainen, *Bi-Lipschitz concordance implies bi-Lipschitz isotopy*, Monatsh. Math. **111** (1991), no. 1, 35–46.
- [37] A. Marden, *The geometry of finitely generated Kleinian groups*, Ann. of Math. **99** (1974), 383–462.
- [38] B. Maskit, *Self-maps of Kleinian groups*, Amer. J. of Math. **93** (1971), 840–856.
- [39] ———, *Kleinian groups*, Springer-Verlag, 1988.
- [40] H. A. Masur and Y. Minsky, *Geometry of the complex of curves I: Hyperbolicity*, Invent. Math. **138** (1999), 103–149.
- [41] ———, *Geometry of the complex of curves II: Hierarchical structure*, Geom. Funct. Anal. **10** (2000), 902–974.
- [42] D. McCullough, *Compact submanifolds of 3-manifolds with boundary*, Quart. J. Math. Oxford **37** (1986), 299–306.
- [43] Y. Minsky, *On Thurston’s ending lamination conjecture*, Proceedings of Low-Dimensional Topology, May 18-23, 1992, International Press, 1994.
- [44] ———, *The classification of punctured-torus groups*, Annals of Math. **149** (1999), 559–626.
- [45] ———, *Kleinian groups and the complex of curves*, Geometry and Topology **4** (2000), 117–148.
- [46] ———, *Bounded geometry in Kleinian groups*, Invent. Math. **146** (2001), 143–192, arXiv:math.GT/0105078.

- [47] ———, *The classification of Kleinian surface groups I: models and bounds*, preprint, 2002.
- [48] ———, *Combinatorial and geometrical aspects of hyperbolic 3-manifolds*, Kleinian Groups and Hyperbolic 3-Manifolds (V. Markovic Y. Komori and C. Series, eds.), London Math. Soc. Lec. Notes, vol. 299, Cambridge Univ. Press, 2003, pp. 3–40.
- [49] ———, *End invariants and the classification of hyperbolic 3-manifolds*, Current developments in mathematics, 2002, Int. Press, Somerville, MA, 2003, pp. 181–217.
- [50] H. Miyachi and K. Ohshika, *On topologically tame Kleinian groups with bounded geometry*, Preprint, 2004.
- [51] G. D. Mostow, *Quasiconformal mappings in n -space and the rigidity of hyperbolic space forms*, Publ. I.H.E.S. **34** (1968), 53–104.
- [52] K. Ohshika, *Ending laminations and boundaries for deformation spaces of Kleinian groups*, J. London Math. Soc. **42** (1990), 111–121.
- [53] J.-P. Otal, *Sur le nouage des géodésiques dans les variétés hyperboliques.*, C. R. Acad. Sci. Paris Sér. I Math. **320** (1995), no. 7, 847–852.
- [54] ———, *Les géodésiques fermées d'une variété hyperbolique en tant que nœuds*, Kleinian groups and hyperbolic 3-manifolds (Warwick, 2001), London Math. Soc. Lecture Note Ser., vol. 299, Cambridge Univ. Press, Cambridge, 2003, pp. 95–104.
- [55] G. Prasad, *Strong rigidity of \mathbf{Q} -rank 1 lattices*, Invent. Math. **21** (1973), 255–286.
- [56] C. P. Rourke and B. J. Sanderson, *Introduction to piecewise-linear topology*, Springer-Verlag, New York, 1972, Ergebnisse der Mathematik und ihrer Grenzgebiete, Band 69.
- [57] D. Ruberman, *Mutation and volumes of knots in S^3* , Invent. Math. **90** (1987), no. 1, 189–215.
- [58] T. Soma, *Geometric limits of quasi-Fuchsian groups*, preprint.
- [59] ———, *Function groups in Kleinian groups*, Math. Annalen **292** (1992), 181–190.
- [60] J. Souto, *Short curves in hyperbolic manifolds are not knotted*, preprint, 2004.
- [61] D. Sullivan, *Hyperbolic geometry and homeomorphisms*, Geometric topology (Proc. Georgia Topology Conf., Athens, Ga., 1977), Academic Press, New York, 1979, pp. 543–555.
- [62] D. Sullivan, *On the ergodic theory at infinity of an arbitrary discrete group of hyperbolic motions*, Riemann Surfaces and Related Topics: Proceedings of the 1978 Stony Brook Conference, Ann. of Math. Stud. 97, Princeton, 1981.
- [63] ———, *Quasiconformal homeomorphisms and dynamics II: Structural stability implies hyperbolicity for Kleinian groups*, Acta Math. **155** (1985), 243–260.
- [64] P. Susskind, *Kleinian groups with intersecting limit sets*, J. d'Analyse Math. **52** (1989), 26–38.
- [65] W. Thurston, *Hyperbolic structures on 3-manifolds, II: surface groups and manifolds which fiber over the circle*, preprint, arXiv:math.GT/9801045.
- [66] ———, *The geometry and topology of 3-manifolds*, Princeton University Lecture Notes, online at <http://www.msri.org/publications/books/gt3m>, 1982.
- [67] ———, *Three dimensional manifolds, Kleinian groups and hyperbolic geometry*, Bull. Amer. Math. Soc. **6** (1982), 357–381.
- [68] ———, *Hyperbolic structures on 3-manifolds, I: deformation of acylindrical manifolds*, Ann. of Math. **124** (1986), 203–246.
- [69] ———, *Three-Dimensional Geometry and Topology*, Princeton University Press, 1997, (S. Levy, ed.).
- [70] P. Tukia and J. Väisälä, *Lipschitz and quasiconformal approximation and extension*, Ann. Acad. Sci. Fenn. Ser. A I Math. **6** (1981), no. 2, 303–342 (1982).
- [71] J. Väisälä, *Quasiconformal concordance*, Monatsh. Math. **107** (1989), no. 2, 155–168.
- [72] F. Waldhausen, *On irreducible 3-manifolds which are sufficiently large*, Ann. of Math. **87** (1968), 56–88.

BROWN UNIVERSITY

UNIVERSITY OF MICHIGAN

YALE UNIVERSITY