CUSP GEOMETRY OF FIBERED 3-MANIFOLDS

By DAVID FUTER and SAUL SCHLEIMER

Abstract. Let F be a surface and suppose that $\varphi\colon F \to F$ is a pseudo-Anosov homeomorphism, fixing a puncture p of F. The mapping torus $M = M_\varphi$ is hyperbolic and contains a maximal cusp C about the puncture p. We show that the area (and height) of the cusp torus ∂C is equal to the stable translation distance of φ acting on the arc complex $\mathcal{A}(F,p)$, up to an explicitly bounded multiplicative error. Our proof relies on elementary facts about the hyperbolic geometry of pleated surfaces. In particular, the proof of this theorem does not use any deep results from Teichmüller theory, Kleinian group theory, or the coarse geometry of $\mathcal{A}(F,p)$. A similar result holds for quasi-Fuchsian manifolds $N \cong F \times \mathbb{R}$. In that setting, we find a combinatorial estimate for the area (and height) of the cusp annulus in the convex core of N, up to explicitly bounded multiplicative and additive error. As an application, we show that covers of punctured surfaces induce quasi-isometric embeddings of arc complexes.

1. Introduction.

Following the work of Thurston, Mostow, and Prasad, it has been known for over three decades that almost every 3-manifold with torus boundary admits a hyperbolic structure [38], which is necessarily unique up to isometry [30, 33]. Thus, in principle, it is possible to translate combinatorial data about a 3-manifold into a detailed description of its geometry—and conversely, to use geometry to identify topological features. Indeed, given a triangulated manifold (up to over 100 tetrahedra) the computer program SnapPy can typically approximate the manifold's hyperbolic metric to a high degree of precision [17]. However, building an effective dictionary between combinatorial and geometric features, for all but the most special families of manifolds, has proven elusive. The prevalence of hyperbolic geometry makes this one of the central open problems in low-dimensional topology.

1.1. Fibered 3-manifolds.

In this paper, we attack this problem for the class of hyperbolic 3-manifolds that fiber over the circle. Let F be a connected, orientable surface with $\chi(F) < 0$, and (for this paper) with at least one puncture. Given an orientation-preserving homeomorphism $\varphi\colon F \to F$ we construct the *mapping torus*

$$M_\varphi := F \times [0,1] \Big/ (x,1) \sim (\varphi(x),0) \cdot$$

Manuscript received September 7, 2011.

Research of the first author supported in part by NSF grant DMS-1007221; research of the second author supported in part by EPSRC grant EP/I028870/1.

American Journal of Mathematics 136 (2014), 309–356.

Thus M_φ fibers over S^1, with fiber F and monodromy φ. Thurston showed that M_φ is hyperbolic if and only if φ is *pseudo-Anosov*: equivalently, if and only if $\varphi^n(\gamma)$ is not homotopic to γ for any $n \neq 0$ and any essential simple closed curve $\gamma \subset F$ [38, 39]. See also Otal [32]. In addition to the connection with dynamics, fibered 3-manifolds are of central importance in low-dimensional topology because every finite-volume, non-positively curved 3-manifold has a finite–sheeted cover that fibers [4, 34].

For those fibered 3-manifolds that are hyperbolic, the work of Minsky, Brock, and Canary on Kleinian surface groups provides a combinatorial, bi-Lipschitz model of the hyperbolic metric [29]. The bi-Lipschitz constants depend only on the fiber F [11]. However, the existence of these constants is proved using compactness arguments; as a result the constants are unknown.

Using related ideas, Brock established the following notable entry in the dictionary between combinatorics and geometry.

THEOREM 1.1. (Brock [10, 9]) *Let F be an orientable surface with $\chi(F) < 0$. Then there exist positive constants K_1 and K_2, depending only on F, such that the following holds. For every orientation-preserving, pseudo-Anosov homeomorphism $\varphi\colon F \to F$, the mapping torus M_φ is a hyperbolic 3-manifold satisfying*

$$K_1\, \overline{d}_{\mathcal{P}}(\varphi) \leq \mathrm{vol}(M_\varphi) \leq K_2\, \overline{d}_{\mathcal{P}}(\varphi).$$

Here $\mathcal{P}(F)$ is the adjacency graph of pants decompositions of F. Also $\overline{d}_{\mathcal{P}}(\varphi)$ is the stable translation distance of φ in $\mathcal{P}(F)$, defined in Equation (1.2) below.

The constant K_2 in the upper bound can be made explicit. Agol showed that the sharpest possible value for K_2 is $2v_8$, where $v_8 = 3.6638\ldots$ is the volume of a regular ideal octahedron [3]. On the other hand, the constant K_1 is only known in the special case when F is a punctured torus or 4-puncture sphere; see Guéritaud and Futer [23, Appendix B]. For all other surfaces, it remains an open problem to give an explicit estimate for K_1.

Brock's theorem is a template for obtaining combinatorial information; the pants graph $\mathcal{P}(F)$ is just one of many complexes naturally associated to a surface F. Others include the curve complex $\mathcal{C}(F)$ and the arc complex $\mathcal{A}(F)$; the latter is the main focus of this paper. Using $\mathcal{A}(F)$ we give *effective* two-sided estimates for the geometry of maximal cusps in M_φ.

Definition 1.2. In this paper, we use the symbols F and S to denote surfaces of negative Euler characteristic, connected and orientable, without boundary and with one or more punctures. Let F be such a surface. The *arc complex* $\mathcal{A}(F)$ is the simplicial complex whose vertices are proper isotopy classes of essential arcs from puncture to puncture. Simplices are collections of vertices admitting pairwise disjoint representatives. We engage in the standard abuse of notation by using the same symbol for an arc and its isotopy class.

The 1-skeleton $\mathcal{A}^{(1)}(F)$ has a combinatorial metric. For a pair of vertices $v, w \in \mathcal{A}^{(0)}(F)$, the distance $d(v,w)$ is the minimal number of edges required to connect v to w. This is well-defined, because $\mathcal{A}(F)$ is connected [24].

When F has a preferred puncture p, we define the subcomplex $\mathcal{A}(F,p) \subset \mathcal{A}(F)$ whose vertices are arcs with at least one endpoint at p. The 1-skeleton $\mathcal{A}^{(1)}(F,p)$ is again connected. The distance $d_{\mathcal{A}}(v,w)$ is the minimal number of edges required to connect v to w inside of $\mathcal{A}^{(1)}(F,p)$.

The mapping class group MCG(F) acts on $\mathcal{A}(F)$ by isometries. In fact, Irmak and McCarthy showed [26] that, apart from a few low-complexity exceptions, MCG$(F) \cong \operatorname{Isom} \mathcal{A}(F)$. Similarly, the subgroup of MCG(F) that fixes the puncture p acts on $\mathcal{A}(F,p)$ by isometries. We are interested in the geometric implications of this action.

Definition 1.3. Let $\varphi\colon F \to F$ be a homeomorphism fixing p. Define the *translation distance* of φ in $\mathcal{A}(F,p)$ to be

$$d_{\mathcal{A}}(\varphi) = \min\left\{ d_{\mathcal{A}}(v, \varphi(v)) \mid v \in \mathcal{A}^{(0)}(F,p) \right\}. \tag{1.1}$$

The same definition applies in any simplicial complex where MCG(F) acts by isometries.

We also define the *stable translation distance* of φ to be

$$\overline{d}_{\mathcal{A}}(\varphi) = \lim_{n\to\infty} \frac{d_{\mathcal{A}}(v, \varphi^n(v))}{n}, \quad \text{for an arbitrary vertex } v \in \mathcal{A}^{(0)}(F,p), \tag{1.2}$$

and similarly for other MCG-complexes. It is a general property of isometries of metric spaces that the limit in (1.2) exists and does not depend on the base vertex v [8, p. 230]. In addition, the triangle inequality implies that $\overline{d}_{\mathcal{A}}(\varphi) \le d_{\mathcal{A}}(\varphi)$.

Note that applying equation (1.2) to the pants graph $\mathcal{P}(F)$ gives the stable translation distance $\overline{d}_{\mathcal{P}}(\varphi)$ that estimates volume in Theorem 1.1. In the same spirit, one may ask the following.

Question 1.4. Let $\mathcal{S}(F)$ be a simplicial complex associated to a surface F, on which the mapping class group MCG(F) acts by isometries. How are the dynamics of the action of $\varphi \in$ MCG(F) on $\mathcal{S}(F)$ reflected in the geometry of the mapping torus M_φ?

We answer this question for the arc complex of a once-punctured surface F or, more generally, for the sub-complex $\mathcal{A}(F,p)$ of a surface with many punctures. Here, the stable distance $\overline{d}_{\mathcal{A}}(\varphi)$ predicts the cusp geometry of M_φ.

1.2. Cusp area from the arc complex. Let M be a 3-manifold whose boundary is a non-empty union of tori, such that the interior of M supports a complete hyperbolic metric. In this metric, every non-compact end of M is a *cusp*, homeomorphic to $T^2 \times [0,\infty)$. Geometrically, each cusp is a quotient of a horoball in $\mathbb{H}^3$ by a $\mathbb{Z} \times \mathbb{Z}$ group of deck transformations. We call this geometrically standard end a *horospherical cusp neighborhood* or *horocusp*.

Associated to each torus $T \subset \partial M$ is a *maximal cusp* $C = C_T$. That is, C is the closure of $C^\circ \subset M$, where C° is the largest embedded open horocusp about T. The same construction works in dimension 2: every punctured hyperbolic surface has a maximal cusp about each puncture.

In dimension 3, Mostow-Prasad rigidity implies that the geometry of a maximal cusp $C \subset M$ is completely determined by the topology of M. One may compute that $\text{area}(\partial C) = \frac{1}{2}\text{vol}(C)$. The Euclidean geometry of ∂C is an important invariant that carries a wealth of information about Dehn fillings of M. For example, if a *slope* s (an isotopy class of simple closed curve on ∂C) is sufficiently long, then Dehn filling M along s produces a hyperbolic manifold [2, 27], whose volume can be estimated in terms of the length $\ell(s)$ [20].

In our setting, M_φ is a fibered hyperbolic 3-manifold, with fiber a punctured surface F. The maximal cusp torus ∂C contains a canonical slope, called the *longitude* of C, which encircles a puncture of F. The Euclidean length of the longitude is denoted λ. Any other, non-longitude slope on ∂C must have length at least

$$\text{height}(\partial C) := \text{area}(\partial C)/\lambda. \tag{1.3}$$

As discussed in the previous paragraph, lower bounds on $\text{height}(\partial C)$ imply geometric control over Dehn fillings of M_φ.

Our main result in this paper uses the action of φ on $\mathcal{A}(F,p)$ to give explicit estimates on the area and height of the cusp torus ∂C.

THEOREM 1.5. *Let F be a surface with a preferred puncture p, and let $\varphi\colon F \to F$ be any orientation-preserving, pseudo-Anosov homeomorphism. In the mapping torus M_φ, let C be the maximal cusp that corresponds to p. Let $\psi = \varphi^n$ be the smallest positive power of φ with the property that $\psi(p) = p$. Then*

$$\frac{\overline{d}_{\mathcal{A}}(\psi)}{450\,\chi(F)^4} < \text{area}(\partial C) \le 9\,\chi(F)^2\,\overline{d}_{\mathcal{A}}(\psi).$$

Similarly, the height of the cusp relative to a longitude satisfies

$$\frac{\overline{d}_{\mathcal{A}}(\psi)}{536\,\chi(F)^4} < \text{height}(\partial C) < -3\,\chi(F)\,\overline{d}_{\mathcal{A}}(\psi).$$

If the surface F has only one puncture p, the statement of Theorem 1.5 becomes simpler in several ways. In this special case, we have $n = 1$, hence $\psi = \varphi$. There is only one cusp in M_φ, and $\mathcal{A}(F,p) = \mathcal{A}(F)$. In this special case, the area and height of the maximal cusp are estimated by the stable translation distance $\overline{d}_{\mathcal{A}}(\varphi)$, acting on $\mathcal{A}(F)$.

In the special case where F is a once-punctured torus or 4-punctured sphere, Futer, Kalfagianni, and Purcell proved a similar estimate, with sharper constants. See [21, Theorems 4.1 and 4.7]. Theorem 1.5 generalizes those results to *all* punctured hyperbolic surfaces.

We note that a non-effective version of Theorem 1.5 can be derived from Minsky's *a priori bounds* theorem for the length of curves appearing in a hierarchy [29, Lemma 7.9]. In fact, this line of argument was our original approach to estimating cusp area. In the process of studying this problem, we came to realize that arguments using the geometry and hierarchical structure of the curve complex $\mathcal{C}(F)$ can be replaced by elementary arguments focusing on the geometry of pleated surfaces. See Section 1.5 below for an outline of this effective argument.

1.3. Quasi-Fuchsian 3-manifolds. The methods used to prove Theorem 1.5 also apply to quasi-Fuchsian manifolds. We recall the core definitions; see Marden [28, Chapter 3] or Thurston [37, Chapter 8] for more details. A hyperbolic manifold $N = \mathbb{H}^3/\Gamma$ is called *quasi-Fuchsian* if the limit set $\Lambda(\Gamma)$ of Γ is a Jordan curve on $\partial\mathbb{H}^3$, and each component of $\partial\mathbb{H}^3 \smallsetminus \Lambda(\Gamma)$ is invariant under Γ. In this case, N is homeomorphic to $F\times\mathbb{R}$ for a surface F. The *convex core* of N, denoted $\mathrm{core}(N)$, is defined to be the quotient, by Γ, of the convex hull of the limit set $\Lambda(\Gamma)$. When N is quasi-Fuchsian but not Fuchsian, $\mathrm{core}(N) \cong F\times[0,1]$, and its boundary is the disjoint union of two surfaces $\partial_+\mathrm{core}(N)$ and $\partial_-\mathrm{core}(N)$, each intrinsically hyperbolic, and each pleated along a lamination. See Definition 2.2.

Although the quasi-Fuchsian manifold N has infinite volume, the volume of $\mathrm{core}(N)$ is finite. Each puncture of F corresponds to a rank one maximal cusp C (the quotient of a horoball by $\mathbb{Z}$), such that $C\cap\mathrm{core}(N)$ has finite volume and $\partial C\cap\mathrm{core}(N)\cong S^1\times[0,1]$ has finite area. Thus we may attempt to estimate the area and height of $\partial C\cap\mathrm{core}(N)$ combinatorially.

Definition 1.6. Let $N\cong F\times\mathbb{R}$ be a quasi-Fuchsian 3-manifold, and let p be a puncture of F. Define $\Delta_+(N)$ to be the collection of all shortest arcs from p to p in $\partial_+\mathrm{core}(N)$. By Lemma 3.4, the arcs in $\Delta_+(N)$ are pairwise disjoint, so $\Delta_+(N)$ is a simplex in $\mathcal{A}(F,p)$. Similarly, let $\Delta_-(N)$ be the simplex of shortest arcs from p to p in $\partial_-\mathrm{core}(N)$.

We define the *arc distance* of N relative to the puncture p to be

$$d_{\mathcal{A}}(N,p) = \min\{d_{\mathcal{A}}(v,w) \mid v\in\Delta_-(N),\ w\in\Delta_+(N)\}.$$

In words, $d_{\mathcal{A}}(N,p)$ is the length of the shortest path in $\mathcal{A}(F,p)$ from a shortest arc in the lower convex core boundary to a shortest arc in the upper boundary.

THEOREM 1.7. *Let F be a surface with a preferred puncture p, and let $N\cong F\times\mathbb{R}$ be a quasi-Fuchsian 3-manifold. Let C be the maximal cusp corresponding* *to p. Then*

$$\frac{d_{\mathcal{A}}(N,p)}{450\,\chi(F)^4} - \frac{1}{23\,\chi(F)^2} < \text{area}(\partial C \cap \text{core}(N))$$
$$< 9\,\chi(F)^2\, d_{\mathcal{A}}(N,p) + \left|12\chi(F)\ln|\chi(F)| + 26\chi(F)\right|$$

Similarly, the height of the cusp relative to a longitude satisfies

$$\frac{d_{\mathcal{A}}(N,p)}{536\,\chi(F)^4} - \frac{1}{27\,\chi(F)^2} < \text{height}(\partial C \cap \text{core}(N))$$
$$< -3\,\chi(F)\, d_{\mathcal{A}}(N,p) + 2\ln|\chi(F)| + 5.$$

We note that the multiplicative constants in Theorem 1.7 are exactly the same as in Theorem 1.5. However, in addition to multiplicative error, the estimates in Theorem 1.7 contain explicit additive error. This additive error is necessary: for example, if the limit set of N is sufficiently close to a round circle, one may have $d_{\mathcal{A}}(N,p) = 0$. (See Lemma 9.2 for a constructive argument.) On the other hand, $\text{area}(\partial C \cap \text{core}(N)) > 0$ whenever N is not Fuchsian.

Theorem 1.7 has an interesting relation to the work of Akiyoshi, Miyachi, and Sakuma [5]. For a quasi-Fuchian manifold N, they study the scale-invariant quantity

$$\text{width}(\partial C) := \text{height}(\partial C \cap \text{core}(N))/\lambda = \text{area}(\partial C \cap \text{core}(N))/\lambda^2$$

where λ is the Euclidean length of the longitude of the cusp annulus ∂C. Generalizing McShane's identity, they give an exact expression for $\text{width}(\partial C)$ as the sum of an infinite series involving the complex lengths of closed curves created by joining the endpoints of an arc. It seems reasonable that most of the contribution in this infinite sum should come from the finitely many arcs in F that are shortest in N. Theorem 1.7 matches this intuition, and indeed its lower bound is proved by summing the contributions of finitely many short arcs.

1.4. Covers and the arc complex. Theorem 1.7 has an interesting application to the geometry of arc complexes, whose statement does not involve 3-manifolds in any way.

Definition 1.8. Suppose $f\colon \Sigma \to S$ is an n-sheeted covering map of surfaces. We define a relation $\pi\colon \mathcal{A}(S) \to \mathcal{A}(\Sigma)$ as follows: $\alpha \in \pi(a)$ if and only if α is a component of $f^{-1}(a)$. In other words, $\pi(a) \subset \mathcal{A}(\Sigma)$ is the set of all n lifts of a, which span an $(n-1)$-simplex.

Definition 1.8 also applies to curve complexes, with the (inessential) difference that the number of lifts of a curve is not determined by the degree of the cover. Using tools from Teichmüller theory, Rafi and Schleimer proved that $\pi\colon \mathcal{C}(S) \to \mathcal{C}(\Sigma)$ is a quasi-isometric embedding [35]. That is, there exist constants $K \geq 1$ and

$C \geq 0$, such that for all $a, b \in \mathcal{C}^{(0)}(S)$ and for all $\alpha \in \pi(a)$, $\beta \in \pi(b)$, we have

$$d(a,b) \leq K\, d(\alpha,\beta) + C \quad \text{and} \quad d(\alpha,\beta) \leq K\, d(a,b) + C.$$

The constants K and C depend only on S and the degree of the cover, but are not explicit. Using tools from Kleinian groups, Tang gave an alternate proof, again without explicit constants [36].

In a similar spirit to Tang, we use Theorem 1.7 to prove a version of the Rafi-Schleimer theorem for arc complexes, with explicit constants.

THEOREM 1.9. *Let Σ and S be surfaces with one puncture, and $f\colon \Sigma \to S$ a covering map of degree n. Let $\pi\colon \mathcal{A}(S) \to \mathcal{A}(\Sigma)$ be the lifting relation. Then, for all $a, b \in \mathcal{A}^{(0)}(S)$, we have*

$$\frac{d(a,b)}{4050\, n\, \chi(S)^6} - 2 < d(\alpha,\beta) \leq d(a,b)$$

where $\alpha \in \pi(a)$ and $\beta \in \pi(b)$.

1.5. An outline of the arguments. The proofs of Theorems 1.5 and 1.7 have a decidedly elementary flavor. The primary tool that we use repeatedly is the geometry of pleated surfaces, as developed by Thurston [37]. (See Bonahon [6] or Canary, Epstein, and Green [14] for a detailed description.) In our context, a pleated surface is typically a copy of the fiber F with a prescribed hyperbolic metric, immersed into M in a piecewise geodesic fashion, and bent along an ideal triangulation of F. In Sections 2 and 3 below, we give a detailed definition of pleated surfaces and discuss the geometry of cusp neighborhoods in such a surface. We also study a mild generalization of pleated surfaces, called simplicial hyperbolic surfaces, that are hyperbolic everywhere except for a single cone point with angle at least 2π.

The upper bounds of Theorems 1.5 and 1.7 are proved in Section 4 and 5, respectively. To sketch the argument in the fibered case, let τ be an ideal triangulation of the fiber F. Then F can be homotoped to a pleated surface, F_τ, in which every ideal triangle is totally geodesic. Using lemmas in Sections 2 and 3, we show that the intersection $F_\tau \cap \partial C$ gives a closed polygonal curve about the puncture p, whose length is bounded by $-6\chi(F)$. As a result, the pleated surface F_τ makes a bounded contribution to the area and height of ∂C. Summing up the contributions from a sequence of triangulations that "realize" the monodromy φ gives the desired upper bound of Theorem 1.5. The upper bound of Theorem 1.7 uses very similar ideas; the one added ingredient is a bound on how far a short arc in $\partial_\pm \operatorname{core}(N)$ drifts when it is pulled tight, making it geodesic in N.

The lower bounds on cusp area and height rely on the idea of a geometrically controlled *sweepout*. This is a degree-one map $\Psi\colon F \times [0,1]/\varphi \to M_\varphi$, in which every fiber $F \times \{t\}$ in the domain is mapped to a piecewise geodesic surface $F_t \subset$

M, which is either pleated or simplicial hyperbolic. The elementary construction of such a sweepout, which is due to Thurston [37] and Canary [13], is recalled in Section 6.

The lower bound of Theorem 1.5 is proved in Section 7. We show that every piecewise geodesic surface F_t in the sweepout of Section 6 must contain an an arc from cusp to cusp whose length is explicitly bounded above. As the parameter t moves around the sweepout, we obtain a sequence of arcs, representing a walk through the 1-skeleton $\mathcal{A}^{(1)}(F,p)$, such that each arc encountered has bounded length in M. This sequence of somewhat-short arcs in the fiber leads to a packing of the cusp torus ∂C by shadows of somewhat-large horoballs, implying a lower bound on $\text{area}(\partial C)$ and $\text{height}(\partial C)$.

It is worth emphasizing that the entire proof of Theorem 1.5 is elementary in nature. In particular, this proof does not rely on any deep results from Teichmüller theory, Kleinian groups, or the coarse geometry of the curve or arc complexes.

The lower bound of Theorem 1.7 is proved in Section 8, using very similar ideas to those of Theorem 1.5. Once again, we have a sweepout $\Psi \colon F \times [0,r] \to \text{core}(N)$ by simplicial hyperbolic surfaces. Once again, each surface F_t in the sweepout contains a somewhat-short arc from cusp to cusp, corresponding to a horoball whose shadow contributes area to ∂C. However, we also need to know that the pleated surfaces at the start and end of the sweepout can be chosen arbitrarily close to $\partial_\pm \text{core}(N)$. This fact, written down as Theorem A.1 in the appendix, is one of the few places in the paper where we need to reach into the non-elementary toolbox of Kleinian groups.

Finally, in Section 9, we prove Theorem 1.9. Given a cover $\Sigma \to S$, vertices a, b of $\mathcal{A}(S)$, and vertices $\alpha \in \pi(a)$, $\beta \in \pi(b)$, the upper bound on the distance $d(\alpha,\beta)$ is immediate because disjoint arcs lift to disjoint multi-arcs. To prove a lower bound, we construct a quasi-Fuchsian manifold $M \cong S \times \mathbb{R}$, so that a and b are the unique shortest arcs on its convex core boundaries. The hyperbolic metric on $M \cong S \times \mathbb{R}$ lifts to a quasi-Fuchsian structure on $N \cong \Sigma \times \mathbb{R}$. By applying Theorem 1.7 to both M and N, we will bound $d(\alpha,\beta)$ from below.

Acknowledgments. This project began at the University of Warwick symposium on *Low Dimensional Geometry and Topology*, in honor of David Epstein, and continued at the MSRI program in *Teichmüller Theory and Kleinian Groups*. We thank the organizers of both events for creating such a fertile ground for collaboration. We thank Ian Agol for numerous helpful conversations, and in particular for contributing the key idea of Lemma 3.8. We thank Dick Canary and Yair Minsky for clarifying a number of points about pleated surfaces, and for helping us sort out the proof of Theorem A.1. We thank Marc Lackenby for continually encouraging us to make our estimates effective.

2. Pleated surfaces and cusps. To prove the upper and lower bounds in our main theorems, we need a detailed understanding of the geometry of pleated surfaces in a hyperbolic 3-manifold. In this section, we survey several known results about pleated surfaces. We also describe the somewhat subtle geometry of the intersection between a pleated surface and a cusp neighborhood in a 3-manifold N. The study of pleated surfaces is continued in Section 3, where we obtain several geometric estimates.

References for this material include Bonahon [6] and Canary, Epstein, and Green [14].

Definition 2.1. Let S be a surface. (Recall the convention of Definition 1.2.) A *lamination* $L \subset S$ is a 1-dimensional foliation of a closed subset of S.

A special case of a lamination is the union of the edges of an ideal triangulation; this special case appears frequently in our setting.

Definition 2.2. Let N be a hyperbolic 3-manifold, and let S be a surface. Fix a proper map $f_0 \colon S \to N$, sending punctures to cusps. For a lamination $L \subset S$ a *pleating* of f_0 along L, or a *pleating map* for short, is a map $f \colon S \to N$, properly homotopic to f_0, such that

(1) f maps every leaf of L to a hyperbolic geodesic, and

(2) f maps every component of $S \smallsetminus L$ to a totally geodesic surface in N.

We say that f *realizes* the lamination L, and call its image $f(S)$ a *pleated surface*.

Note the existence of a pleating map places restrictions on L; for instance, every closed leaf of L must be essential and non-peripheral in S. The hyperbolic metric on N, viewed as a path-metric, pulls back via f to induce a complete hyperbolic metric on S. In this induced hyperbolic metric, every leaf of L becomes a geodesic, and the map $f \colon S \to N$ becomes a piecewise isometry, which is bent along the geodesic leaves of $f(L)$.

LEMMA 2.3. *Every pleated surface $f(S)$ is contained in the convex core of N.*

Proof. Adding leaves as needed to subdivide the totally geodesic regions, we can arrange for the complement $S \smallsetminus L$ to consist of ideal triangles. Fix g, a side of some ideal triangle of $S \smallsetminus L$. Thus g is a bi-infinite geodesic; each end of the image geodesic $f(g)$ either runs out a cusp of N or meets a small metric ball infinitely many times. In either case, the lift of $f(g)$ to the universal cover $\widetilde{N} = \mathbb{H}^3$ has both endpoints at limit points of N. Since an ideal triangle in $\mathbb{H}^3$ is the convex hull of its vertices, the surface $f(S)$ is contained in $\mathrm{core}(N)$. □

In a quasi-Fuchsian manifold N, the components of $\partial_\pm \, \mathrm{core}(N)$ are themselves pleated surfaces. In this paper, the convex core boundaries are the only examples of pleated surfaces where the pleating laminations are not ideal triangulations.

A foundational result is that every essential surface $S \subset N$ can be pleated along an arbitrary ideal triangulation. This was first observed by Thurston [37, Chapter

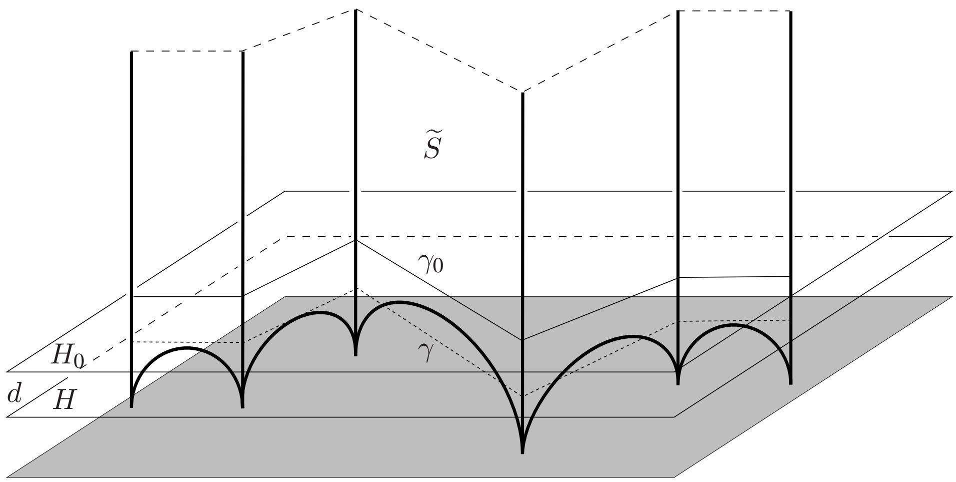

Figure 1. The intersection between a pleated surface and a horocusp. The intersection with horoball H_0 is standard, whereas the intersection with H may contain portions of the surface bent along geodesics whose endpoints are not in H. Graphic based on a design of Agol [2, Figure 1].

8]. For a more detailed account of the proof, see Canary, Epstein, and Green [14, Theorem 5.3.6] or Lackenby [27, Lemma 2.2].

PROPOSITION 2.4. *Let N be a cusped orientable hyperbolic 3-manifold. Let $f_0\colon S \to N$ be a proper, essential map, sending punctures to cusps. Then, for any ideal triangulation τ of S, the map f_0 is homotopic to a pleating map along τ.*

In other words, every ideal triangulation τ is realized by a pleating map $f_\tau\colon S \to N$.

Suppose that $C \subset N$ is a horospherical cusp neighborhood in N, and $f\colon S \to N$ is a pleating map. Our goal is to describe the geometry of $f(S) \cap C$. We offer Figure 1 as a preview of the geometric picture. The figure depicts a lift $\widetilde{S}$ of a pleated surface to $\mathbb{H}^3$. For a sufficiently small horocusp $C_0 \subset C$, which lifts to horoball H_0 in the figure, the intersection $f(S) \cap C_0$ is *standard*, meaning that $f^{-1}(C_0)$ is a union of horospherical cusp neighborhoods in S. The intersection $f(S) \cap C$ is more complicated, because the surface is bent along certain geodesics whose interior intersects $C \smallsetminus C_0$. Nevertheless, we can use the geometry of $f(S) \cap C_0$ to find certain cusp neighborhoods contained in $f^{-1}(C)$ (in Lemma 2.5), and certain geometrically meaningful closed curves in ∂C (in Lemma 2.6).

LEMMA 2.5. *Let N be a cusped orientable hyperbolic 3-manifold, with a horocusp C. Let $f\colon S \to N$ be a pleating map, such that n punctures of S are mapped to C. Suppose that a loop about a puncture of S is represented by a geodesic of length λ on ∂C. Then, in the induced hyperbolic metric on S, the preimage $f^{-1}(C) \subset S$ contains horospherical cusp neighborhoods $R_1, \dots, R_n$ with disjoint* *interiors, such that*

$$\ell(\partial R_i) = \text{area}(R_i) \geq \lambda \quad \textit{for each } i.$$

Our proof is inspired by an argument of Agol [2, Theorem 5.1].

Proof of Lemma 2.5. Without loss of generality, assume that the pleating lamination L cuts S into ideal triangles. (Otherwise, add more leaves to L.) Let $C_0 \subset C$ be a horocusp chosen sufficiently small so that $C_0 \cap f(L)$ is a union of non-compact rays into the cusp. Then $f^{-1}(C_0)$ is a union of tips of ideal triangles in S and consists of disjoint horospherical neighborhoods $R_1^0, \dots, R_n^0$, each mapped into C.

Lift N to its universal cover $\mathbb{H}^3$, so that C_0 lifts to a horoball H_0 about ∞ in the upper half-space model. Then $\widetilde{f(S)}$ intersects H_0 in vertical bands, as shown in Figure 1.

Let d be the distance in N between ∂C_0 and ∂C. Since the interior of C is embedded, this means that the shortest geodesic in N from C_0 to C_0 has length at least $2d$. Since the pleating map $f\colon S \to N$ is distance-decreasing, the shortest geodesic in S from $f^{-1}(C_0)$ to itself also has length at least $2d$. In other words, we may take a closed d-neighborhood of each R_i^0 and obtain a cusp neighborhood R_i, such that $R_1, \dots, R_n$ have disjoint interiors.

Consider the areas of these neighborhoods, along with their boundary lengths. A standard calculation in the upper half-plane model of $\mathbb{H}^2$ implies that the length of a horocycle in S equals the area of the associated cusp neighborhood. Furthermore, both quantities grow exponentially with d.

On ∂C_0, a Euclidean geodesic about a puncture of S has length $e^{-d}\lambda$. Since $f(S) \cap C_0$ may not be totally geodesic (in general, it is bent, as in Figure 1), each curve of ∂R_i^0 has length bounded below by $e^{-d}\lambda$. These lengths grow by e^d as we take a d-neighborhood of $\cup_i R_i^0$. We conclude that each component R_i satisfies

$$\ell(\partial R_i) = \text{area}(R_i) \geq e^d \cdot e^{-d} \cdot \lambda = \lambda. \tag{2.1}$$

It remains to show that $f(R_i) \subset C$ for each i. Suppose, without loss of generality, that R_1^0 is the component of $f^{-1}(C_0)$ whose lift is mapped to the horoball H_0. Then, by construction, the lift of R_1 is mapped into the d-neighborhood of H_0, which is a horoball H covering C. Thus $f(R_1) \subset C$. Note that the containment might be strict, because $f(S)$ might be bent along some geodesics in the region $C \setminus C_0$, as in the middle of Figure 1. □

The argument of Lemma 2.5 also permits the following construction, which is important for Section 4.

LEMMA 2.6. *Let N be a cusped orientable hyperbolic 3-manifold, with a horocusp C. Let $f\colon S \to N$ be a pleating map that realizes an ideal triangulation τ. Then, for each puncture p of S that is mapped to C, there is an immersed closed* *curve $\gamma = \gamma(f,p,C)$, piecewise geodesic in the Euclidean metric on ∂C, with the following properties:*

(1) *The loop γ is homotopic in C to a loop in $f(S)$ about p.*

(2) *The vertices of γ lie in $\partial C \cap f(\tau)$, and correspond to the endpoints of edges of τ at puncture p.*

(3) *$\ell(\gamma) = \ell(\partial R_i)$, where $R_i \subset S$ is one of the cusp neighborhoods of Lemma 2.5.*

A lift of γ to a horoball H covering C is shown, dotted, in Figure 1.

Proof of Lemma 2.6. We may construct γ as follows. Recall, from the proof of Lemma 2.5, that there is a horocusp $C_0 \subset C$ such that the intersection $f(S) \cap C_0$ is standard, consisting of tips of ideal triangles. Then $f^{-1}(C_0)$ is a disjoint union of horospherical cusp neighborhoods. Let R_i^0 be the component of $f^{-1}(C_0)$ that contains puncture p, and let $\gamma_0 = f(\partial R_i^0) \subset \partial C_0$.

Note that the curve γ_0 is piecewise geodesic in the Euclidean metric on ∂C_0, and that it is bent precisely at the intersection points $\partial C_0 \cap f(\tau)$, where the triangulation τ enters the cusp. See Figure 1.

We define γ to be the projection of γ_0 to the horospherical torus ∂C. Note that if C_0 and C are lifted to horoballs about ∞ in $\mathbb{H}^3$, as in Figure 1, this projection is just vertical projection in the upper half-space model.

Observe that while $\gamma_0 \subset f(S)$, its projection γ might not be contained in the pleated surface. Nevertheless, γ is completely defined by γ_0. The vertices where γ is bent are contained in $f(\tau)$.

Let d be the distance between ∂C_0 and ∂C. Then, as in Lemma 2.5, lengths grow by a factor of e^d as we pass from ∂C_0 to ∂C. Thus, by the same calculation as in (2.1),

$$\ell(\gamma) = e^d \cdot \ell(\gamma_0) = e^d \cdot \text{area}(R_i^0) = \text{area}(R_i) = \ell(\partial R_i),$$

where $R_i \supset R_i^0$ is the cusp neighborhood in S that is mapped into C, as in Lemma 2.5. □

Our final goal in this section is to provide a universal lower bound on the size of the cusp neighborhoods R_i. We do this using the following result of Adams [1].

LEMMA 2.7. *Let N be a non-elementary, orientable hyperbolic 3-manifold, and let C be a maximal horocusp in N. (This neighborhood may correspond to either a rank one or rank two cusp.) Let s be a simple closed curve on ∂C, which forms part of the boundary of an essential surface in N. Then $\ell(s) > 2^{1/4}$.*

Proof. This is a consequence of a theorem of Adams [1, Theorem 3.3]. He proved that every parabolic translation of the maximal cusp of any non-elementary hyperbolic 3-manifold has length greater than $2^{1/4}$, with exactly three exceptions: one parabolic each in the three SnapPea census manifolds `m004`, `m009`, and `m015`.

Each of the manifolds `m004`, `m009`, and `m015` is either a punctured torus bundle or a two-bridge knot complement. Hence the boundary slopes of incompressible surfaces in these manifolds are classified [19, 25]. In particular, none of the three slopes shorter than $2^{1/4}$ bounds an essential surface. □

As a result, we obtain

LEMMA 2.8. *Let N be a cusped orientable hyperbolic 3-manifold, with a maximal cusp C. Let $f(S) \subset N$ be a pleated surface, homotopic to a properly embedded essential surface, such that n punctures of S are mapped to C. Then $f^{-1}(C) \subset S$ contains n disjoint horospherical cusp neighborhoods $R_1, \dots, R_n$, such that*

$$\ell(\partial R_i) = \text{area}(R_i) > 2^{1/4} \quad \textit{for each } i.$$

Proof. This is immediate from Lemmas 2.5 and 2.7. □

3. Hyperbolic surfaces with one cone point. Recall from Definition 2.2 that every pleated surface carries an intrinsic hyperbolic metric. In this section, we prove several lemmas about the geometry of cusp neighborhoods and geodesic arcs in these surfaces. These estimates are used throughout the proofs of Theorems 1.5 and 1.7.

In fact, we work in a slightly more general setting: namely, hyperbolic surfaces with a cone point, whose cone angle is at least 2π. These singular surfaces arise in sweepouts of a hyperbolic 3-manifold: see Section 6. Therefore, we derive length and area estimates for these singular surfaces, as well as non-singular ones.

Definition 3.1. A *hyperbolic cone surface* is a complete metric space S, homeomorphic to a surface of finite type. We require that S admits a triangulation into finitely many simplices, such that each simplex is isometric to a totally geodesic triangle in $\mathbb{H}^2$. The triangles are allowed to have any combination of ideal vertices (which correspond to punctures of S) and material vertices (which correspond to points in S). The triangles are glued by isometries along their edges.

Every point of S that is not a material vertex of the triangulation thus has a neighborhood isometric to a disk in $\mathbb{H}^2$. Every material vertex $v \in S$ has a neighborhood where where the metric (in polar coordinates) takes the form

$$ds^2 = dr^2 + \sinh^2(r)\, d\theta^2, \tag{3.1}$$

where $0 \le r < r_v$ and $0 \le \theta \le \theta_v$. Here θ_v is called the *cone angle* at v, and can be computed as the sum of the interior angles at v over all the triangles that meet v. Note that if $\theta_v = 2\pi$, equation (3.1) becomes the standard polar equation for the hyperbolic metric in a disk. The vertices of S whose cone angles are *not* equal to 2π are called the *cone points* or *singular points* of S; all remaining points are called *non-singular*.

If all singular points of S have cone angles $\theta_v > 2\pi$, another common name for S is a *simplicial hyperbolic surface*. Simplicial hyperbolic surfaces have played an important role in the study of geometrically infinite Kleinian groups [13, 22].

Just as with non-singular hyperbolic surfaces, cone surfaces have a natural geometric notion of a cusp neighborhood.

Definition 3.2. Let S be a hyperbolic cone surface, with one or more punctures, and let $R \subset S$ be a closed set. Then R is called an *equidistant cusp neighborhood* of a puncture of S if the following conditions are satisfied:

(1) The interior of R is homeomorphic to $S^1 \times (0,\infty)$.

(2) There is a closed subset $Q \subset R$, whose universal cover $\widetilde{Q}$ is isometric to a horoball in $\mathbb{H}^2$. This implies that the interior of Q is disjoint from all cone points.

(3) There is a distance $d > 0$, such that R is the closed d-neighborhood of Q.

R is called a *maximal cusp* if it is not a proper subset of any larger equidistant cusp neighborhood. Equivalently, R is maximal if and only if it is not homeomorphic to $S^1 \times [0,\infty)$.

LEMMA 3.3. *Let S be a hyperbolic surface with one cone point v, of angle $\theta_v \geq 2\pi$. Let $R \subset S$ be a non-maximal equidistant neighborhood of a puncture of S. Then*

(1) *There is a geodesic α that is shortest among all essential paths from R to R.*

(2) *The arc α is either embedded, or is the union of a segment and a loop based at v. In the latter case, there is an arbitrarily small homotopy in S making α embedded.*

(3) *If $2\pi \leq \theta_v < 4\pi$, and β is another shortest arc from R to R, there is an arbitrarily small homotopy making α and β disjoint.*

One way to interpret Lemma 3.3 is as follows. Let p be a puncture of S. Then any shortest arc relative to a cusp neighborhood about p gives a vertex of the arc complex $\mathcal{A}(S,p)$. If there are two distinct shortest arcs, they span an edge of $\mathcal{A}(S,p)$; more generally, if there are n distinct shortest arcs, they span an $(n-1)$-simplex. This is used in Sections 7 and 8 to construct a path in $\mathcal{A}(S,p)$.

Proof of Lemma 3.3. Let d be the infimal distance in the universal cover $\widetilde{S}$ between two different lifts of ∂R; since R is not a maximal cusp, $d > 0$. Furthermore, since distance to the nearest translate is an equivariant function on $\partial\widetilde{R}$, it achieves a minimum. Thus there are points x, y on distinct lifts of ∂R, whose distance is exactly d. By the Hopf-Rinow theorem for cone manifolds [16, Lemma 3.7], the distance between x and y is realized by a geodesic path α, and every distance-realizing path is a geodesic. This proves (1).

Suppose that $\alpha\colon [0,d] \to S$ is a unit-speed parametrization. The image of α is a graph $\Gamma = \mathrm{Im}(\alpha)$. If α is not embedded, then Γ has at least one vertex of valence

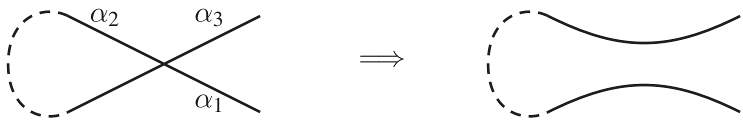

Figure 2. Exchanging and rounding off the arcs α_1 and α_3 produces a shorter path.

larger than 2. In this case, we will show that Γ is the union of a segment and a loop based at v.

Let w be a vertex of Γ that has valence larger than 2. Consider preimages $x, y \in \alpha^{-1}(w)$, where $x < y$. Then $[0, d]$ splits into sub-intervals

$$I_1 = [0, x], \quad I_2 = [x, y], \quad I_3 = [y, d].$$

Let α_i be the restriction of α to the sub-interval I_i.

We claim that α_1 must be homotopic to $\overline{\alpha_3}$ (the reverse of α_3): otherwise, cutting out the middle segment α_2 would produce a shorter essential path. We also claim that $\ell(\alpha_1) = \ell(\alpha_3)$: for, if $\ell(\alpha_1) < \ell(\alpha_3)$, we could homotope α_3 to $\overline{\alpha_1}$ while shortening the length of α. In particular, the last claim implies that x and y are the only preimages of w, and w has valence 3 or 4.

If $w \neq v$, then it is a non-singular point of S, hence a 4-valent vertex. This means that α_1 meets α_3 at a nonzero angle. Then, exchanging α_1 and $\overline{\alpha_3}$ by homotopy and rounding off the corner at w, as in Figure 2, produces an essential arc shorter than α. This is a contradiction.

We may now assume that the only vertex in the interior of $\Gamma = \mathrm{Im}(\alpha)$ occurs at the cone point v. This implies that v is not in the cusp neighborhood R, hence R contains no singular points and its universal cover $\widetilde{R}$ is isometric to a horoball. Thus, since horoballs are convex, there is a unique shortest path from v to R, in each homotopy class.

In particular, the homotopic arcs α_1 and $\overline{\alpha_3}$ must coincide. Thus $w = v$ is 3-valent, and is the only vertex of Γ. Hence α_2 is an embedded loop based at v, and α must be an "eyeglass" that follows α_1 from R to v, runs around the loop α_2, and returns to R by retracing α_1. In this case, even though α is not embedded, one component of the frontier of an ε-neighborhood of Γ is an embedded arc homotopic to α. This proves (2).

For future reference, we note an important feature of eyeglass geodesics. Suppose that α consists of an arc α_1 from v to R and a loop α_2 based at v. Then α_1 must be the *unique* shortest path from v to R. For if another geodesic α_3 from v to R has length $\ell(\alpha_3) \leq \ell(\alpha_1)$, then α_1 and α_3 must be in different homotopy classes. This means $\alpha_1 \cup \alpha_3$ is an essential arc from R to R, whose length is

$$\ell(\alpha_1) + \ell(\alpha_3) \leq 2\ell(\alpha_1) < 2\ell(\alpha_1) + \ell(\alpha_2) = \ell(\alpha),$$

contradicting the fact that α is shortest.

Figure 3. An eyeglass path and an embedded arc can be made disjointly embedded after a short homotopy. The dashed sections of α and β are schematics meant to indicate that the arcs are traveling through a distant part of the surface, while staying disjoint.

For part (3), suppose that α and β are two distinct shortest arcs from R to R. By statement (2), each of α and β is either an embedded arc or an eyeglass with a loop based at v. Suppose that α and β intersect, and let $\Gamma = \mathrm{Im}(\alpha) \cup \mathrm{Im}(\beta)$.

If Γ has a non-singular vertex w, then w cuts α into sub-arcs α_1, α_2 that run from w to R. Similarly, w cuts β into sub-arcs β_3, β_4 from w to R. Without loss of generality, say that α_1 is shortest among these four arcs. Then at least one of $\alpha_1 \cup \beta_1$ or $\alpha_1 \cup \beta_2$ is an essential arc from R to R, and both of these arcs are no longer than β. By rounding off the corner at w, we can make $\alpha_1 \cup \beta_1$ or $\alpha_1 \cup \beta_2$ into an essential arc shorter than β. This is a contradiction.

For the rest of the proof, we assume that the only vertex of Γ is at v. One consequence of this assumption is that R contains no singular points. Hence, as in the proof of (2), there is a unique shortest path from v to R in each homotopy class. There are three cases: (i) neither α nor β is an eyeglass, (ii), α is an eyeglass but β is not, and (iii) both α and β are eyeglasses.

If α and β are embedded arcs that intersect at v, consider the valence of v, which must be 3 or 4. If v is 3-valent, α and β must share the same path $\alpha_1 = \beta_1$ from R to v, then diverge. In this case, the ε-neighborhood of Γ contains disjointly embedded arcs homotopic to α and β.

If v is 4-valent, let $\gamma_1, \ldots, \gamma_4$ be the four geodesic sub-arcs of Γ from v to R. Since each γ_i is the unique shortest path in its homotopy class, any combination $\gamma_i \cup \gamma_j$ is an essential arc. Since $\alpha = \gamma_1 \cup \gamma_2$ and $\beta = \gamma_3 \cup \gamma_4$ are both shortest arcs in S, every γ_i must have the same length. But since v has cone angle $\theta_v < 4\pi$, there must be two sub-arcs γ_i, γ_j that meet at an angle less than π. Thus $\gamma_i \cup \gamma_j$ can be shortened by smoothing the corner at v, contradicting the assumption that α and β are shortest.

If α is an eyeglass, but β is not, let α_1 be the sub-arc of α from v to R. By the observation at the end of part (2), α_1 is the unique shortest geodesic from v to R. Let β_1, β_2 be the sub-arcs of β from v to R, where $\ell(\beta_1) \leq \ell(\beta_2)$. If α_1 is distinct from β_1, then $\alpha_1 \cup \beta_1$ would be an essential path that is shorter than

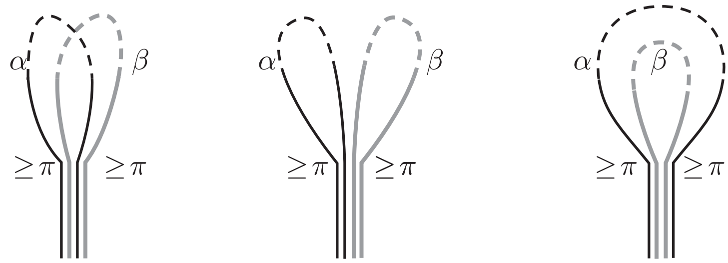

Figure 4. Two eyeglass paths that share the same stem can be made disjointly embedded after a short homotopy. Shown are the embedded versions of α and β, in the three possible interleaving configurations at vertex v.

β—contradiction. Thus $\alpha_1 = \beta_1$. In this case, homotoping α to an embedded arc makes it disjoint from β. See Figure 3.

Finally, if each of α and β is an eyeglass, the observation at the end of part (2) implies that each of α and β must contain the unique shortest geodesic from v to R. Thus each of α and β consists of the same geodesic arc γ_1 from v to R, as well as a loop based at v. Figure 4 shows that the ε-neighborhood of $\Gamma = \mathrm{Im}(\alpha) \cup \mathrm{Im}(\beta)$ contains disjointly embedded paths representing α and β. □

In the case where S is a non-singular surface, we have a stronger version of Lemma 3.3: not only are shortest arcs disjoint, but nearly-shortest arcs must be disjoint as well.

LEMMA 3.4. *Let S be a punctured hyperbolic surface, and let R be a horospherical neighborhood about one puncture. Let α and β be distinct geodesic arcs from R to R. If α and β intersect, then there is a third geodesic arc γ, satisfying*

$$\ell(\gamma) \leq \max\{\ell(\alpha), \ell(\beta)\} - \ln(2).$$

Here, all lengths are measured relative to the cusp neighborhood R.

In practice, we use the contrapositive statement: if both α and β are at most $\ln(2)$ longer than the shortest geodesic from R to R, then they must be disjoint.

Proof of Lemma 3.4. If we change the size of the cusp neighborhood R, then all geodesic arcs from R to R have their lengths changed by the same additive constant. Thus, without loss of generality, we may assume that R is small enough so that all intersections between α and β happen outside R. As S has no cone points, α and β meet transversely.

Orient both α and β. If we cut the geodesic α along its intersection points with β, we obtain a collection of segments. Let α_1 and α_2 be the first and last such segments, respectively, along an orientation of α. That is, α_1 (respectively α_2) is the sub-arc of α from ∂R to the first (last) point of intersection with β. Similarly, let β_1 and β_2 be the first and last segments of β, along an orientation of β.

Assume, without loss of generality, that α_1 is shortest among the segments $\alpha_1, \alpha_2, \beta_1, \beta_2$. Set $v = \alpha_1 \cap \beta$. Then the vertex v cuts β into sub-arcs β_3 and β_4, such that $\beta_1 \subset \beta_3$ and $\beta_2 \subset \beta_4$. Without loss of generality, we may also assume that β_3 (rather than β_4) is the sub-arc of β that meets α_1 at an angle of at most $\pi/2$.

Note that $\alpha_1 \cup \beta_3$ is an embedded arc, because (by construction) α_1 only intersects β at the vertex v. Furthermore, $\alpha_1 \cup \beta_3$ is topologically essential (otherwise, one could homotope β to reduce its length). The hypothesis that α_1 is shortest among $\alpha_1, \alpha_2, \beta_1, \beta_2$ implies that

$$\ell(\alpha_1 \cup \beta_3) \le \ell(\beta_2 \cup \beta_3) \le \ell(\beta_4 \cup \beta_3) = \ell(\beta).$$

Let γ denote the geodesic from R to R in the homotopy class of $\alpha_1 \cup \beta_3$. Then the geodesic extensions of α_1, β_3, and γ form a $2/3$ ideal triangle, with angle $\theta \le \pi/2$ at the material vertex v. Then, [15, Lemma A.3] gives

$$\ell(\gamma) = \ell(\alpha_1 \cup \beta_3) + \ln\left(\frac{1-\cos\theta}{2}\right) \le \ell(\beta) + \ln\left(\frac{1}{2}\right),$$

as desired. □

Next, we consider the area of equidistant cusp neighborhoods in S.

LEMMA 3.5. *Let S be a simplicial hyperbolic surface with at most one singular point. Let R_1 and R_2 be embedded equidistant neighborhoods of the same puncture of S, such that $R_1 \subset R_2 \subset S$, and d is the distance between ∂R_1 and ∂R_2. Then*

$$\operatorname{area}(R_2) \ge e^d \operatorname{area}(R_1).$$

Proof. Let $Q \subset R_1$ be a cusp neighborhood isometric to the quotient of a horoball. Then there is a number $m > 0$, such that for all $x \in [0,m]$, the closed x-neighborhood of Q is an equidistant cusp neighborhood $R(x)$. In particular, $R_1 = R(x_1)$ and $R_2 = R(x_2)$, where $x_2 = x_1 + d$. We shall explore the dependence of area$(R(x))$ on the parameter x.

Let v be the singular point of S. (If S is non-singular, let v be an arbitrary point of $S \setminus Q$, and consider it a cone point of cone angle 2π.) Let x_v be the distance from v to ∂Q. Then, for $x < x_v$, $R(x)$ is a non-singular neighborhood of a cusp, itself the quotient of a horoball. As mentioned in the proof of Lemma 2.5, the area of a horospherical cusp grows exponentially with distance. In symbols,

$$\operatorname{area}(R(x)) = e^x \operatorname{area}(Q) \quad \text{if} \quad x \le x_v.$$

For $x > x_v$, the cusp neighborhood $R(x)$ can be constructed from a horoball and a cone. More precisely: take a horospherical cusp, and cut it along a vertical slit of length $r = x - x_v$. Then, take a cone of radius r and angle $\theta = \theta_v - 2\pi$, and cut it open along a radius. Gluing these pieces together along the slits produces $R(x)$. The area of a hyperbolic cone with radius r and angle θ can be computed as $2\theta \sinh^2(r/2)$. Adding this to the area of a horospherical cusp, we obtain

$$\operatorname{area}(R(x)) = \begin{cases} e^x \operatorname{area}(Q) & \text{if } x < x_v, \\ e^x \operatorname{area}(Q) + 2\theta \sinh^2((x - x_v)/2) & \text{if } x \geq x_v. \end{cases}$$

To complete the proof, it suffices to check that the function $f(x) = \sinh^2(x/2)$ grows super-exponentially for $x \geq 0$:

$$\begin{aligned} \left(\sinh \frac{x+d}{2}\right)^2 &= \left(\sinh \frac{x}{2} \cosh \frac{d}{2} + \cosh \frac{x}{2} \sinh \frac{d}{2}\right)^2 \\ &> \left(\sinh \frac{x}{2} \cosh \frac{d}{2} + \sinh \frac{x}{2} \sinh \frac{d}{2}\right)^2 \\ &= \left(e^{d/2} \sinh \frac{x}{2}\right)^2 \\ &= e^d \sinh^2\left(\frac{x}{2}\right). \end{aligned}$$

Thus, since the area of $R(x)$ is the sum of two functions, each of which grows at least exponentially with x, it follows that $\operatorname{area}(R(x+d)) \geq e^d \operatorname{area}(R(x))$. □

LEMMA 3.6. *Let S be a simplicial hyperbolic surface with at most one singular point. Let $R_{\max} \subset S$ be a maximal cusp neighborhood of a puncture of S. Then*

$$\operatorname{area}(R_{\max}) \leq -2\pi\chi(S).$$

Furthermore, if S is non-singular, then

$$\operatorname{area}(R_{\max}) \leq -6\chi(S).$$

Proof. Let $\theta_v \geq 2\pi$ be the cone angle at the singular point. (As above, we take $\theta_v = 2\pi$ if the surface S is non-singular.) Then, by the Gauss-Bonnet theorem [16, Theorem 3.15],

$$\operatorname{area}(R_{\max}) \leq \operatorname{area}(S) = -2\pi\chi(S) + [2\pi - \theta_v] \leq -2\pi\chi(S). \tag{3.2}$$

If S is a non-singular surface, then horosphere packing estimates of Böröczky [7] imply that at most $3/\pi$ of the area of S can be contained in the cusp neighborhood $R_{\max}$. Thus the above estimate improves to $\operatorname{area}(R_{\max}) \leq -6\chi(S)$. □

Remark 3.7. By the Gauss-Bonnet theorem expressed in equation (3.2), the area of S decreases as the cone angle θ_v increases from 2π. Thus it seems reasonable that the area of a maximal cusp $R_{\max}$ would also decrease as θ_v increases from 2π. If this conjecture is true, then the estimate $\mathrm{area}(R_{\max}) \le -6\chi(S)$ would hold for singular hyperbolic surfaces as well as non-singular ones.

LEMMA 3.8. *Let S be a simplicial hyperbolic surface with at most one singular point. Let $R \subset S$ be an embedded equidistant neighborhood of a puncture of S. Then there exists a geodesic arc α from R to R, satisfying*

$$\ell(\alpha) \le 2\ln|2\pi\,\chi(S)/\,\mathrm{area}(R)|.$$

Furthermore, if S is non-singular, then

$$\ell(\alpha) \le 2\ln|6\,\chi(S)/\,\mathrm{area}(R)|.$$

Proof. By Lemma 3.3, there is a geodesic arc α that is shortest among all essential arcs from R to R. Let $R_{\max}$ be the maximal cusp neighborhood containing R. Then, by construction, $R_{\max}$ meets itself at the midpoint of α. Thus the distance from ∂R to $\partial R_{\max}$ is $d = \ell(\alpha)/2$. By Lemma 3.5, this implies

$$e^d\,\mathrm{area}(R) \le \mathrm{area}(R_{\max}),$$

which simplifies to

$$\ell(\alpha) = 2d \le 2\ln\left(\mathrm{area}(R_{\max})/\,\mathrm{area}(R)\right).$$

Substituting the bound on $\mathrm{area}(R_{\max})$ from Lemma 3.6 completes the proof. □

4. Upper bound: fibered manifolds. In this section, we prove the upper bounds of Theorem 1.5. We begin with a slightly simpler statement:

THEOREM 4.1. *Let F be an orientable hyperbolic surface with a preferred puncture p, and let $\psi\colon F \to F$ be an orientation-preserving, pseudo-Anosov homeomorphism such that $\psi(p) = p$. In the mapping torus M_ψ, let C be the maximal cusp that corresponds to p. Then*

$$\mathrm{area}(\partial C) \le 9\,\chi(F)^2\,d_{\mathcal{A}}(\psi) \quad \textit{and} \quad \mathrm{height}(\partial C) < -3\,\chi(F)\,d_{\mathcal{A}}(\psi).$$

Theorem 4.1 differs from the upper bound of Theorem 1.5 in two relatively small ways. First, Theorem 4.1 restricts attention to monodromies that fix the puncture p. (Given an arbitrary pseudo-Anosov φ, one can let ψ be the smallest power of φ such that $\psi(p) = p$.) Second, Theorem 4.1 estimates cusp area and height in terms of the translation distance $d_{\mathcal{A}}(\psi)$, rather than the stable translation distance $\overline{d}_{\mathcal{A}}(\psi)$. We shall see at the end of the section that this simpler statement quickly implies the upper bound of Theorem 1.5.

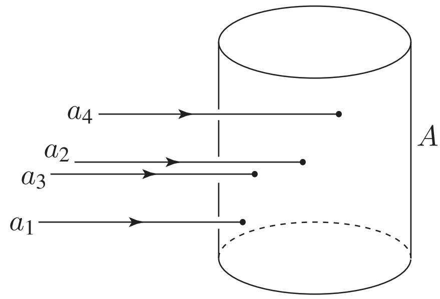

Figure 5. The edges a_i enter the cusp annulus $A \subset N_\psi$ at well-defined heights.

Proof of Theorem 4.1. The proof involves a direct construction. Suppose that $d_{\mathcal{A}}(\psi) = k$. Then, by Definition 1.3, there is a vertex $a_0 \in \mathcal{A}^{(0)}(F,p)$, that is an isotopy class of arc in F meeting the puncture p, so that $d_{\mathcal{A}}(a_0, \psi(a_0)) = k$. Fix a geodesic segment in $\mathcal{A}^{(1)}(F,p)$ with vertices

$$a_0, a_1, \ldots, a_k = \psi(a_0).$$

By Definition 1.2, the arcs representing a_{i-1} and a_i are disjoint. Thus, for every $i = 1, \ldots, k$, we can choose an ideal triangulation τ_i of F that contains a_{i-1} and a_i. In the arc complex $\mathcal{A}(F,p)$, the geodesic segment from a_0 to a_k extends to a bi-infinite, ψ-invariant, piecewise geodesic. Similarly, the sequence of ideal triangulations $\tau_1, \ldots, \tau_k$ extends to a bi-infinite sequence of triangulations in which $\tau_{i+k} = \psi(\tau_i)$.

To prove the upper bounds on cusp area and height, it is convenient to work with the infinite cyclic cover of M_ψ. This is a hyperbolic manifold $N_\psi \cong F \times \mathbb{R}$, in which the torus cusps of M lift to annular, rank one cusps. Let $A \subset N_\psi$ be the lift of ∂C that corresponds to the puncture p of F.

We choose geodesic coordinates for the Euclidean metric on $A \cong S^1 \times \mathbb{R}$, in which the non-trivial circle (a longitude about p) is horizontal, and the $\mathbb{R}$ direction is vertical. We also choose an orientation for every arc a_i, such that the oriented edge a_i points into the preferred puncture p. In the 3-manifold N_ψ, the edge a_i is homotopic to a unique oriented geodesic. Given our choices, every arc a_i has an associated *height* $h(a_i)$, namely the vertical coordinate of the point of A where the oriented geodesic representing a_i enters the cusp. For simplicity, we may assume that $h(a_0) = 0$ and $h(a_k) = h(\psi(a_0)) > 0$. See Figure 5.

To estimate distances on the annulus A, we place many copies of the fiber F into pleated form. That is, for each ideal triangulation τ_i, $i \in \mathbb{Z}$, Proposition 2.4 implies that the fiber $F \times \{0\} \subset N_\psi = F \times \mathbb{R}$ can be homotoped in N_ψ to a pleated surface F_i realizing the triangulation τ_i. Recall that every ideal triangle of τ_i is totally geodesic in F_i.

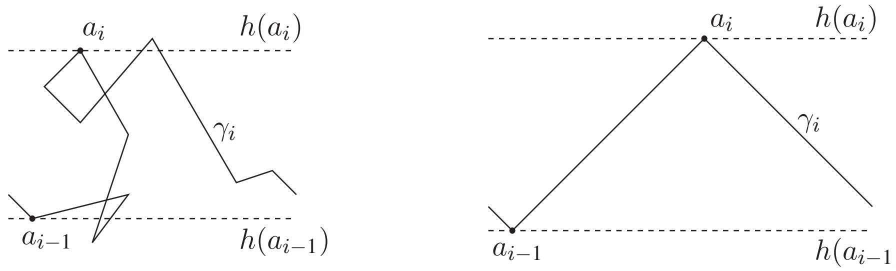

Figure 6. Left: the polygonal closed curve $\gamma_i \subset A$. Right: the shape of γ_i that maximizes the area of the band B_i between heights $h(a_{i-1})$ and $h(a_i)$.

For each pleated surface F_i, Lemma 2.6 gives a possibly self-intersecting, piecewise geodesic closed curve $\gamma_i \subset A$, which is homotopic to a loop about the puncture p. The curve γ_i is not necessarily contained in $F_i \cap A$, but we do know that the vertices where it bends are the endpoints of edges of τ_i meeting the annulus A. See Figure 1 for a review.

LEMMA 4.2. *Each piecewise linear closed curve $\gamma_i \subset A$, determined by the pleated surface F_i, has length $\ell(\gamma_i) \leq -6\chi(F)$.*

Proof. Lemma 2.6 states that $\ell(\gamma_i) = \ell(\partial R_i) = \text{area}(R_i)$, where $R_i \subset F_i$ is an embedded horospherical neighborhood of the puncture p. By Lemma 3.6, $\text{area}(R_i) \leq -6\chi(F)$. □

See Agol [2, Theorem 5.1] or Lackenby [27, Lemma 3.3] for a very similar statement, on which Lemma 4.2 is based.

Applying Lemma 4.2 to the pleated surface F_i gives a height estimate.

LEMMA 4.3. *The heights of consecutive arcs satisfy*

$$|h(a_i) - h(a_{i-1})| < -3\chi(F).$$

Proof. By construction, the arcs a_{i-1} and a_i have endpoints at the puncture p. Additionally, the geodesic representatives of both arcs are contained in the triangulation τ_i along which F_i is bent. Thus the piecewise geodesic closed curve γ_i, containing the vertices at the forward endpoints of a_{i-1} and a_i, must visit heights $h(a_{i-1})$ and $h(a_i)$. See Figure 6, left. Since $\ell(\gamma_i) \leq -6\chi(F)$, and this closed curve covers the distance between heights $h(a_{i-1})$ and $h(a_i)$ at least twice, we conclude that

$$|h(a_i) - h(a_{i-1})| < -3\chi(F).$$

The inequality is strict because γ_i must also travel around a horizontal loop in A. □

Lemma 4.2 also leads to an area estimate.

LEMMA 4.4. *Let $B_i \subset A$ be the band whose boundary consists of horizontal circles at heights $h(a_{i-1})$ and $h(a_i)$. Then* $\mathrm{area}(B_i) \leq 9\,\chi(F)^2$.

Proof. As in the proof of Lemma 4.3, we study the piecewise geodesic closed curve $\gamma_i \subset A$. Since γ_i meets both a_{i-1} and a_i, it must meet both boundary components of B_i. The goal is to determine the shape of γ_i that allows the largest possible area for B_i.

Without loss of generality, we may assume that γ_i contains exactly two geodesic segments connecting the two boundary circles of B_i: otherwise, one can straighten γ_i while stretching B_i. Such a piecewise-linear loop consisting of two segments splits B_i into two isometric triangles: one triangle below γ_i, and the other triangle above γ_i. See Figure 6, right.

At this point, we have reduced to the classical calculus problem of building a triangular corral adjacent to a river. As is well-known, the optimal shape for γ_i is one where the two segments have the same length and meet at right angles. By Lemma 4.2, the total length of these two equal segments is at most $-6\,\chi(F)$. Therefore, the maximum possible area for B_i is $9\,\chi(F)^2$. □

We can now complete the proof of Theorem 4.1. A fundamental domain for the torus ∂C is the portion of A between height $h(a_0) = 0$ and height $h(a_k) = h(\psi(a_0))$. This fundamental domain is contained in $B_1 \cup \ldots \cup B_k$. (The containment might be strict, since there is no guarantee that the sequence $h(a_i)$ is monotonically increasing; see Figure 5.) Thus, by Lemma 4.4,

$$\mathrm{area}(\partial C) \leq \sum_{i=1}^{k} \mathrm{area}(B_i) \leq 9k\,\chi(F)^2.$$

Similarly, by Lemma 4.3,

$$\mathrm{height}(\partial C) = h(a_k) - h(a_0) \leq \sum_{i=1}^{k} |h(a_i) - h(a_{i-1})| < -3k\,\chi(F).$$

Recalling that $d_{\mathcal{A}}(\psi) = k$ completes the proof. □

COROLLARY 4.5. *Let F be an orientable hyperbolic surface with a preferred puncture p, and let $\psi\colon F \to F$ be an orientation-preserving, pseudo-Anosov homeomorphism such that $\psi(p) = p$. In the mapping torus M_ψ, let C be the maximal cusp that corresponds to p. Then*

$$\mathrm{area}(\partial C) \leq 9\,\chi(F)^2\,\overline{d}_{\mathcal{A}}(\psi) \quad \textit{and} \quad \mathrm{height}(\partial C) < -3\,\chi(F)\,\overline{d}_{\mathcal{A}}(\psi).$$

Corollary 4.5 differs from Theorem 4.1 in that $d_{\mathcal{A}}$ has been replaced by $\overline{d}_{\mathcal{A}}$. Since $\overline{d}_{\mathcal{A}}(\psi) \leq d_{\mathcal{A}}(\psi)$ by triangle inequalities, the statement of Corollary 4.5 is slightly sharper.

Proof of Corollary 4.5. Let $n \geq 1$. The maximal cusp C of M_ψ lifts to an embedded horocusp in M_{ψ^n}, whose area is $n \cdot \text{area}(\partial C)$. Applying Theorem 4.1 to M_{ψ^n}, we obtain

$$n \cdot \text{area}(\partial C) \leq 9\,\chi(F)^2\, d_{\mathcal{A}}(\psi^n).$$

Thus

$$\begin{aligned} \text{area}(\partial C) &\leq 9\,\chi(F)^2 \inf_{n \geq 1} \frac{d_{\mathcal{A}}(\psi^n)}{n} \\ &\leq 9\,\chi(F)^2 \liminf_{n \to \infty} \frac{d_{\mathcal{A}}(\psi^n)}{n} \\ &\leq 9\,\chi(F)^2\, \overline{d}_{\mathcal{A}}(\psi). \end{aligned}$$

The identical calculation goes through for $\text{height}(\partial C)$. □

Proof of Theorem 1.5, upper bound. Let $\varphi\colon F \to F$ be a pseudo-Anosov homeomorphism, and let $\psi = \varphi^n$ be the smallest power of φ that fixes the puncture p. Let C be the maximal cusp of M_φ corresponding to p. Then C lifts to a (not necessarily maximal) horocusp $C' \subset M_\psi$, which is a one-sheeted cover of C. Corollary 4.5 gives upper bounds on the area and height of the maximal cusp of M_ψ, which implies upper bounds on the (possibly smaller) area and height of C. □

5. Upper bound: quasi-Fuchsian manifolds. In this section, we prove the upper bounds of Theorem 1.7. The proof strategy is nearly the same as the proof of Theorem 4.1, with the quasi-Fuchsian manifold $N \cong F \times \mathbb{R}$ playing the same role as N_ψ in the previous section. The main geometric difference is that $\text{core}(N_\psi)$ is the whole manifold, whereas $\text{core}(N)$ has finite volume and finite cusp area.

Let $C \subset N$ be the maximal cusp corresponding to the puncture p of F. As in Section 4, we choose geodesic coordinates on $A = \partial C \cong S^1 \times \mathbb{R}$, in which the $\mathbb{R}$ direction is vertical. Since $\text{core}(N)$ is convex, the intersection $\text{core}(N) \cap A$ must be a compact annulus whose boundary is a pair of horizontal circles. We choose the orientation on $\mathbb{R}$ so that $\partial_+ \text{core}(N)$ is higher than $\partial_- \text{core}(N)$. Then every oriented essential arc $a_i \subset N$ whose forward endpoint is at C has a well-defined height $h(a_i)$, namely the vertical coordinate of the point on A where the geodesic homotopic to a_i enters the cusp. See Figure 5.

Following Definition 1.6, let $\Delta_\pm(N)$ be the simplex in $\mathcal{A}(F,p)$ consisting of all shortest arcs from p to p in $\partial_\pm \text{core}(N)$. Let $a_0 \in \Delta_-(N)$ and $a_k \in \Delta_+(N)$ realize the distance between these simplices, so that

$$k = d_{\mathcal{A}}(a_0, a_k) = d_{\mathcal{A}}(N, p).$$

Since a_0 has both of its endpoints at p, we may choose the orientation on a_0 so that the point where a_0 enters the cusp is the lower of the two endpoints. Similarly, we choose the orientation on a_k so that the point where a_k enters the cusp is the higher of the two endpoints.

LEMMA 5.1. *Let $B \subset A$ be the compact annular band whose boundary consists of horizontal circles at heights $h(a_0)$ and $h(a_k)$. Then*

$$\operatorname{area}(B) \leq 9\,\chi(F)^2\, d_{\mathcal{A}}(N,p) \quad \textit{and} \quad |h(a_k)-h(a_0)| < -3\,\chi(F)\, d_{\mathcal{A}}(N,p).$$

Proof. This follows from the results of Section 4. Let $a_0, a_1, \dots, a_k$ be the vertices of a geodesic in $\mathcal{A}^{(1)}(F,p)$ between a_0 and a_k. Then, for every i, let $B_i \subset A$ be the annular band whose boundary consists of horizontal circles at heights $h(a_{i-1})$ and $h(a_i)$. By Lemmas 4.3 and 4.4,

$$\operatorname{area}(B_i) \leq 9\,\chi(F)^2 \quad \text{and} \quad \big|h(a_i) - h(a_{i-1})\big| < -3\,\chi(F).$$

Adding up these estimates as i ranges from 1 to k gives the result. □

To prove the upper bounds of Theorem 1.7, it remains to estimate the area and height of the part of $\operatorname{core}(N) \cap A$ that is *not* contained in the band B. To make this region more precise, define $h_\pm(N)$ to be the vertical coordinate of the circle $\partial_\pm \operatorname{core}(N) \cap A$. Then we may define $B_- = B_-(N)$ to be the band whose boundary consists of horizontal circles at heights $h_-(N)$ and $h(a_0)$, and similarly $B_+ = B_+(N)$ to be the band between heights $h(a_k)$ and $h_+(N)$. Note that the orientations of a_0 and a_k were chosen precisely so as to minimize the size of $B_-(N)$ and $B_+(N)$, respectively.

Recall that $\partial_- \operatorname{core}(N)$ is an intrinsically hyperbolic surface, pleated along a lamination. The arc a_0 has a geodesic representative in $\partial_- \operatorname{core}(N)$; in fact, by definition this geodesic is shortest in $\partial_- \operatorname{core}(N)$ among all arcs from p to p. Then $h_-(N)$ is the height at which the geodesic representative of a_0 in $\partial_- \operatorname{core}(N)$ enters the cusp C, and $h(a_0)$ is the height at which the geodesic representative of a_0 in N enters the cusp C. The difference $|h(a_0) - h_-(N)|$ is the height of $B_-(N)$. We control $|h(a_0) - h_-(N)|$ via the following proposition.

PROPOSITION 5.2. *Let $\gamma \subset N$ be an oriented, essential arc from cusp C back to C, which is disjoint from the interior of C. Let $g \subset N$ be the geodesic in the homotopy class of γ. Let $h(\gamma)$ be the height at which γ enters C, and $h(g)$ be the height at which g enters C.*

Assume that the orientation of γ has been chosen to minimize $|h(g) - h(\gamma)|$. Then either $|h(g) - h(\gamma)| \leq \sqrt{2}$, or

$$|h(g) - h(\gamma)| \leq \frac{\ell(\gamma) - \ln\left(3 + 2\sqrt{2}\right) + 2\sqrt{2}}{2} = \frac{\ell(\gamma)}{2} + 0.5328\dots, \tag{5.1}$$

where $\ell(\gamma)$ is the arclength of γ.

Proof. Lift γ to an arc $\widetilde{\gamma} \subset \mathbb{H}^3$. The oriented arc $\widetilde{\gamma}$ runs from horoball H' to horoball H. Let $\widetilde{g}$ be the corresponding lift of g, namely the oriented geodesic from H' to H. Let $d_+ = d_+(g,\gamma)$ be the distance along ∂H between the endpoints of $\widetilde{g}$ and $\widetilde{\gamma}$ on H, and similarly let $d_- = d_-(g,\gamma)$ be the distance along $\partial H'$ between the endpoints of $\widetilde{g}$ and $\widetilde{\gamma}$ on H'. Since the orientation of γ has been chosen to minimize $|h(g) - h(\gamma)|$, we have

$$|h(g) - h(\gamma)| \leq \min\{d_-(g,\gamma),\, d_+(g,\gamma)\} \leq \frac{d_- + d_+}{2}. \tag{5.2}$$

Thus, to bound $|h(g) - h(\gamma)|$, it will suffice to bound the average of d_- and d_+.

Next, we reduce the problem from three to two dimensions, as follows. Consider cylindrical coordinates (r,θ,z) on $\mathbb{H}^3$, with the geodesic $\widetilde{g}$ at the core of the cylinder. Thus r measures distance from $\widetilde{g}$, while θ is the rotational parameter, and z measures distance along $\widetilde{g}$. With these coordinates, the hyperbolic metric becomes

$$ds^2 = dr^2 + \sinh^2(r)\,d\theta^2 + \cosh^2(r)\,dz^2. \tag{5.3}$$

We claim that no generality is lost by assuming $\widetilde{\gamma}$ lies in the half-plane corresponding to $\theta = 0$. This is because the expression for the metric in (5.3) is diagonalized, hence the map $(r,\theta,z) \mapsto (r,0,z)$ is distance–decreasing. Thus replacing $\widetilde{\gamma}$ by its image in this half-plane only makes it shorter. Furthermore, horoballs H and H' are rotationally symmetric about $\widetilde{g}$, hence the new planar curve is still disjoint from their interiors. Finally, observe that rotation about $\widetilde{g}$ keeps the endpoints of $\widetilde{\gamma}$ a constant distance from $\widetilde{g} \cap H'$ and $\widetilde{g} \cap H$, respectively. Thus the quantities $d_-(g,\gamma)$ and $d_+(g,\gamma)$ are unchanged when we replace $\widetilde{\gamma}$ by its rotated image. From now on, we will assume that $\widetilde{\gamma}$ lies in the copy of $\mathbb{H}^2$ for which $\theta \in \{0,\pi\}$.

Recall that $\widetilde{\gamma}$ is disjoint from the interiors of H and H'. When its endpoints are sufficiently far apart from $\widetilde{g}$, the geodesic between those points would pass through the interiors of the horoballs. Instead, the shortest path that stays outside H and H' follows the boundary of H', then tracks a hyperbolic geodesic tangent to H' and H, then follows the boundary of H. (See Figure 7.) The following lemma estimates the length of the geodesic segment in the middle of this path.

LEMMA 5.3. *Let H and H' be horoballs in $\mathbb{H}^2$ with disjoint interiors. Let $\alpha \subset \mathbb{H}^2$ be a hyperbolic geodesic such that H and H' are both tangent to α, on the same side of α. Let $\beta \subset \mathbb{H}^2$ be a hyperbolic geodesic perpendicular to both H and H'. If ℓ_1 denotes the length along α from $H \cap \alpha$ to $H' \cap \alpha$, and ℓ_2 denotes the length along ∂H from $H \cap \alpha$ to $H \cap \beta$, then*

$$\ell_1 \geq \ln\left(3 + 2\sqrt{2}\right) \quad \text{and} \quad \ell_2 \leq \sqrt{2}.$$

Each inequality becomes equal if and only if H is tangent to H'.

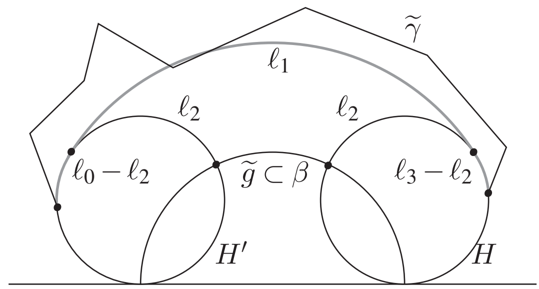

Figure 7. The setup of Proposition 5.2. When the endpoints of $\widetilde{\gamma}$ are sufficiently far apart, the shortest path that stays disjoint from the interiors of H and H' is the three-piece thick arc, of length $(\ell_0 - \ell_2) + \ell_1 + (\ell_3 - \ell_2)$.

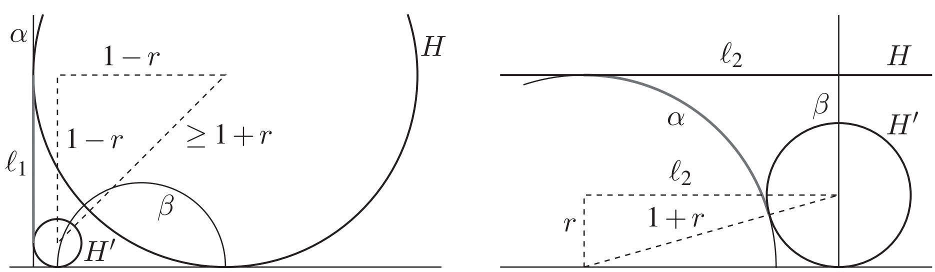

Figure 8. The setup of Lemma 5.3. Left: making geodesic α vertical helps estimate ℓ_1. Right: making geodesic β vertical helps estimate ℓ_2. In both panels, the horoball H' has Euclidean radius r.

Proof. Let g be the geodesic segment of α whose length is ℓ_1. For the first inequality, apply an isometry of $\mathbb{H}^2$ so that $g \subset \alpha$ is vertical in the upper half-plane model, so that H is the larger horoball, and so that the Euclidean radius of H is 1. Then the point of tangency $\alpha \cap H$ is at Euclidean height 1. (See Figure 8, left.) A calculation with the Pythagorean theorem then implies that the Euclidean radius of H' must be $r \leq 1/(3+2\sqrt{2})$, with equality if and only if H is tangent to H'. Since one endpoint of g is at height 1 and the other endpoint is at height r, we have $\ell_1 = \ln(1/r) \geq \ln(3+2\sqrt{2})$.

For the second inequality, apply an isometry of $\mathbb{H}^2$ so that ∂H is a horizontal line at Euclidean height 1. Then α is a Euclidean semicircle of radius 1, and the horoball H' must have Euclidean radius $r \leq 1/2$. (See Figure 8, right.) Again, a calculation with the Pythagorean theorem implies that $\ell_2 \leq \sqrt{2}$, with equality if and only if $r = 1/2$. □

Returning to the proof of Proposition 5.2, we import the notation of Lemma 5.3. That is: the geodesic α is tangent to H and H', while the geodesic β contains $\widetilde{g}$. Let ℓ_1 and ℓ_2 be as in the lemma. Let $\ell_0 = d_-(g,\gamma)$ be the distance along $\partial H'$ between the $\widetilde{g} \cap \partial H'$ and $\widetilde{\gamma} \cap \partial H'$, and let $\ell_3 = d_+(g,\gamma)$.

If one of $\ell_0 = d_-$ or $\ell_3 = d_+$ is no longer than $\sqrt{2}$, then equation (5.2) gives $|h(g) - h(\gamma)| \le \sqrt{2}$ as well, and the proof is complete. Otherwise, if ℓ_0 and ℓ_3 are both longer than $\sqrt{2}$, then Lemma 5.3 implies that they are longer than ℓ_2. As a consequence, the shortest path between the endpoints of $\widetilde{\gamma}$ that stays disjoint from H and H' will need to track horoball H' for distance $(\ell_0 - \ell_2)$, then follow geodesic α for distance ℓ_1, then track horoball H for distance $(\ell_3 - \ell_2)$. See Figure 7.

Thus we may compute:

$$\begin{aligned}
\ell(\gamma) &\ge (\ell_0 - \ell_2) + \ell_1 + (\ell_3 - \ell_2) && \text{by construction of } \widetilde{\gamma} \\
&= (d_- + d_+) - 2\ell_2 + \ell_1 && \text{by the definition of } \ell_0 \text{ and } \ell_3 \\
&\ge (d_- + d_+) - 2\sqrt{2} + \ln(3 + 2\sqrt{2}) && \text{by Lemma 5.3,} \\
&\ge 2|h(g) - h(\gamma)| - 2\sqrt{2} + \ln(3 + 2\sqrt{2}) && \text{by (5.2),}
\end{aligned}$$

implying (5.1). □

We can now prove the upper bounds of Theorem 1.7.

Proof of Theorem 1.7, upper bound. Let $C \subset N$ be a maximal cusp corresponding to the puncture p of F. As above, $\partial C \cap \mathrm{core}(N)$ decomposes into three compact annular bands: the band B_- between heights $h_-(N)$ and $h(a_0)$, the band B between heights $h(a_0)$ and $h(a_k)$, and the band B_+ between heights $h(a_k)$ and $h_+(N)$.

The area and height of B were bounded in Lemma 5.1. As for B_-, let $a_0 \in \Delta_-(N)$ be one of the arcs from C to C that is shortest on $\partial_- \mathrm{core}(N)$. Let γ be the geodesic in $\partial_- \mathrm{core}(N)$ in the homotopy class of a_0. By Lemmas 3.8 and 2.8,

$$\ell(\gamma) \le 2\ln\left|6\chi(F)/2^{1/4}\right| = 2\ln|\chi(F)| + \ln\left(18\sqrt{2}\right). \tag{5.4}$$

Note that $\ln\left(18\sqrt{2}\right) \approx 3.2369 > \sqrt{2}$, hence the larger upper bound in Proposition 5.2 is the one in equation (5.1). Thus, by plugging estimate (5.4) into (5.1), we obtain

$$\mathrm{height}(B_-) = |h(a_0) - h(\gamma)| \le \ln|\chi(F)| + \frac{\ln\left(18\sqrt{2}\right)}{2} + 0.54 < \ln|\chi(F)| + 2.16.$$

By Lemma 4.2, the circumference of B_- (which is a longitude of the cusp C) satisfies $\lambda \le -6\chi(F)$. Thus

$$\mathrm{area}(B_-) = \lambda \cdot \mathrm{height}(B_-) < \left|6\chi(F)\ln|\chi(F)| + 13\chi(F)\right|.$$

The top band $B_+ = B_+(N)$ satisfies the same estimates.

Combining these estimates with Lemma 5.1, we obtain

$$\begin{aligned}\text{area}(\partial C \cap \text{core}(N)) &= \text{area}(B_- \cup B \cup B_+) \\ &< 9\chi(F)^2\, d_{\mathcal{A}}(N,p) + \left|12\chi(F)\ln|\chi(F)| + 26\chi(F)\right|.\end{aligned}$$

Similarly,

$$\begin{aligned}\text{height}(\partial C \cap \text{core}(N)) &= \text{height}(B_- \cup B \cup B_+) \\ &< -3\,\chi(F)\, d_{\mathcal{A}}(N,p) + 2\ln|\chi(F)| + 5,\end{aligned}$$

completing the proof. □

6. Sweepouts. In this section, we describe an important geometric and topological construction needed for the lower bounds in Theorems 1.5 and 1.7.

Definition 6.1. Let N be a hyperbolic 3-manifold, F a surface, and $f_0\colon F \to N$ a map sending punctures to cusps. Fix a connected set $J \subset \mathbb{R}$. A *sweepout through* f_0 is a map $\Psi\colon F \times J \to N$, thought of as a one-parameter family of maps $\Psi_t\colon F \to N$, each homotopic to f_0.

A sweepout Ψ is called *geometric* if every Ψ_t is a *simplicial hyperbolic map*: that is, for every $t \in J$, the image $F_t = \Psi_t(F)$ is a hyperbolic cone surface with at most one cone point of angle $2\pi \le \theta_t < 4\pi$. Note that a pleating map along an ideal triangulation is a special case of a simplicial hyperbolic map.

Let g_t be the hyperbolic cone metric on F induced by Ψ_t. Then the continuity of Ψ implies that g_t varies continuously with t. In particular, the lengths of geodesic realizations of homotopy classes of arcs and curves (with respect to g_t) vary continuously with t.

Definition 6.2. Let $M = M_\psi$ be a fibered hyperbolic 3-manifold with fiber F and monodromy ψ. Let $N = N_\psi$ be the infinite cyclic cover of M, with primitive deck transformation $Z\colon N \to N$. Fix $r > 0$ and define $z\colon F \times \mathbb{R} \to F \times \mathbb{R}$ by $z(x,t) = (\psi(x), t+r)$. We say a sweepout $\Psi\colon F \times \mathbb{R} \to N$ is *equivariant* if each Ψ_t is properly homotopic to the fiber, and

$$Z \circ \Psi = \Psi \circ z.$$

Note that equivariance implies that Ψ descends to a degree-one sweepout of M.

PROPOSITION 6.3. *Let $M = M_\psi$ be a fibered hyperbolic 3-manifold with fiber F and monodromy ψ. Let $N = N_\psi$ be the infinite cyclic cover. Then there is a geometric, equivariant sweepout $\Psi\colon F \times \mathbb{R} \to N$.*

This result originates in the work of Thurston [37, Theorem 9.5.13]; see page 9.47 in particular. A careful account of the proof was also written down by Canary

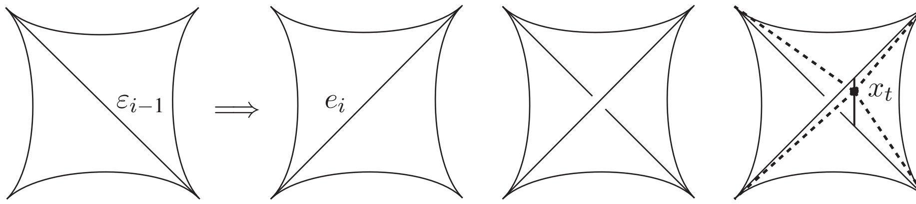

Figure 9. Left: a diagonal exchange in a quadrilateral. Center right: a diagonal exchange between two pleated surfaces creates a 3-dimensional tetrahedron Δ_i. Right: a singular quadrilateral Q_t interpolates between the top and bottom pleated sides of the tetrahedron.

[13, Sections 4–5]. What follows below is a review of their argument, adapted to ideal triangulations.

Proof of Proposition 6.3. Let τ_0 be an ideal triangulation of F. Hatcher proved [24] that the triangulations τ_0 and $\psi(\tau_0)$ can be connected by a sequence of diagonal exchanges, as in Figure 9. Thus we have a sequence of ideal triangulations, $\tau_0, \tau_1, \dots, \tau_r = \psi(\tau_0)$, where each τ_i differs from τ_{i-1} by a diagonal exchange. We extend the sequence of triangulations τ_i to a bi-infinite sequence $\{\tau_i \mid i \in \mathbb{Z}\}$, such that $\tau_{i+r} = \psi(\tau_i)$.

Fix an embedding $f_0 \colon F \to N$ isotopic to the fiber. By Proposition 2.4, for each $i \in \mathbb{Z}$ we may take $\Psi_i \colon F \to N$ to be a pleating of f_0 along τ_i. We choose these pleating maps so that $Z \circ \Psi_i = \Psi_{i+r} \circ \psi$. Define $F_i = \mathrm{Im}(\Psi_i)$ and notice that, since τ_i differs from τ_{i-1} by a diagonal exchange, F_i differs from F_{i-1} by an ideal tetrahedron Δ_i.

Fix $i \in \mathbb{Z}$. Let ε_{i-1} be the edge of τ_{i-1} that is exchanged for the edge e_i of τ_i. Because all six edges of the tetrahedron Δ_i lift to hyperbolic geodesics, the edges ε_{i-1} and e_i lift to hyperbolic geodesics with no shared endpoints at infinity. In $\mathbb{H}^3$, this pair of geodesics is joined by a unique geodesic segment γ that meets ε_{i-1}, e_i perpendicularly. (In the special case where ε_{i-1} and e_i intersect, γ has length 0.)

Now, for every $t \in [i-1, i]$, we take x_t to be the point on γ that is distance $\ell(\gamma)(t-i+1)$ from ε_{i-1} and distance $\ell(\gamma)(i-t)$ from e_i. In other words, x_t is obtained by linear interpolation between the points where γ meets ε_{i-1} and e_i. We construct a (singular) ideal quadrilateral Q_t by coning x_t to the four edges of $\Delta_i \smallsetminus (e_i \cup \varepsilon_{i-1})$. (See Figure 9, right.) Finally, let F_t be the surface that includes the quadrilateral Q_t inside the tetrahedron Δ_i, and agrees with F_i everywhere else. Recall that F_{i-1} agrees with F_i outside Δ_i.

For $t \in (i-1, i)$, we take $\Psi_t \colon F \to N_\psi$ to be a simplicial hyperbolic map with image F_t. We choose the maps Ψ_t so that $\Psi|_{F \times [i-1,i]}$ is continuous and so that $Z \circ \Psi_t = \Psi_{t+r} \circ \psi$.

Now consider the geometry of $F_t = \Psi_t(F)$. Wherever this surface agrees with F_i, it is built out of ideal triangles, and inherits an intrinsically hyperbolic metric from N_ψ. Meanwhile, the quadrilateral Q_t where F_t disagrees with F_i is constructed out of four $(2/3)$-ideal triangles that share a vertex at x_t. Thus the surface F_t has a smooth hyperbolic metric everywhere except at x_t. At this cone point, observe that the singular quadrilateral Q_t is not contained in any hyperbolic half-space through x_t. As a result, a lemma of Canary [13, Lemma 4.2] implies that the cone angle at x_t is $\theta_t \geq 2\pi$. Also, because each of the four triangles meeting at x_t has an interior angle less than π, we have $\theta_t < 4\pi$.

Since the triangulations τ_i satisfy $\tau_{i+r} = \psi(\tau_i)$, we have

$$Z \circ \Psi_t = \Psi_{t+r} \circ \psi, \quad \text{for all } t \in \mathbb{R}.$$

Therefore, Ψ is the desired equivariant, geometric sweepout of N. □

The above construction extends nicely to quasi-Fuchsian manifolds.

PROPOSITION 6.4. *Let $N \cong F \times \mathbb{R}$ be a cusped quasi-Fuchsian 3-manifold. Let τ, τ' be ideal triangulations of F. Then there exists a geometric sweepout $\Psi\colon F \times [0,r] \to N$, such that Ψ_0 is the pleating map along τ and Ψ_r is the pleating map along τ'.*

Proof. We repeat the proof of Proposition 6.3, without needing to worry about equivariance. Let $\tau = \tau_0, \tau_1, \ldots, \tau_r = \tau'$ be a sequence of ideal triangulations of F, with each τ_i differing from τ_{i-1} by a diagonal exchange. Then each τ_i can be realized by a pleated surface F_i, and we may interpolate from F_{i-1} to F_i by a 1-parameter family of simplicial hyperbolic surfaces, as in Figure 9. □

LEMMA 6.5. *Let $\Psi\colon F \times J \to N$ be a geometric sweepout in a hyperbolic 3-manifold N. Let C be an embedded horocusp in N, with longitude of length λ. Then, for every surface $F_t = \Psi_t(F)$ in the sweepout, $C \cap F_t$ contains an equidistant cusp neighborhood whose area is at least λ.*

Recall that an *equidistant cusp* of a hyperbolic cone surface is defined in Definition 3.2.

Proof of Lemma 6.5. This is identical to the proof of Lemma 2.5, with pleated surfaces replaced by simplicial hyperbolic surfaces. First, take a horocusp $C_0 \subset C$ small enough so that $F_t \cap C_0$ is a non-singular, horospherical cusp R_t^0. If d denotes the distance between ∂C_0 and ∂C, then $\operatorname{area}(R_t^0) \geq e^{-d}\lambda$. Then $F_t \cap C$ contains an equidistant d-neighborhood of R_t^0. By Lemma 3.5, this equidistant neighborhood R_t satisfies

$$\operatorname{area}(R_t) \geq e^d \operatorname{area}(R_t^0) \geq \lambda. \qquad \square$$

7. Lower bound: fibered manifolds. In this section, we prove the lower bound of Theorem 1.5. We begin with a slightly more restricted statement:

THEOREM 7.1. *Let F be an orientable hyperbolic surface with a preferred puncture p, and let $\psi\colon F \to F$ be an orientation-preserving, pseudo-Anosov* *homeomorphism such that $\psi(p) = p$. In the mapping torus M_ψ, let C be a horocusp corresponding to p, whose longitude has length $\lambda = 2^{1/4}$. Then there exists an integer $n \geq 1$, such that*

$$\text{area}(\partial C) > \frac{d_{\mathcal{A}}(\psi^n)}{450\,\chi(F)^4} \quad \textit{and} \quad \text{height}(\partial C) > \frac{d_{\mathcal{A}}(\psi^n)}{536\,\chi(F)^4}.$$

Theorem 7.1 differs from the lower bound of Theorem 1.5 in several small ways. First, it restricts attention to horocusps that have longitude of length $2^{1/4}$. (This choice of longitude is justified by Lemma 2.7, which represents the best available lower bound on the longitude.) Second, Theorem 7.1 restricts attention to monodromies that fix the puncture p. (Given an arbitrary pseudo-Anosov φ, one can let ψ be the smallest power of φ such that $\psi(p) = p$.) Finally, Theorem 7.1 estimates cusp area and height in terms of the translation distance $d_{\mathcal{A}}(\psi^n)$ for an unspecified integer n, rather than the stable translation distance $\overline{d}_{\mathcal{A}}(\psi)$. We shall see at the end of the section that this restricted statement quickly implies the lower bound of Theorem 1.5.

Proof of Theorem 7.1. As in Sections 4 and 6, it is convenient to work with the infinite cyclic cover of M_ψ, namely $N_\psi \cong F \times \mathbb{R}$. The horocusps of M lift to annular, rank one cusps. Let $A \subset N_\psi$ be the lift of ∂C that corresponds to the puncture p of F. Then A is an annulus, with longitude of length $\lambda = 2^{1/4}$, which covers the torus ∂C.

By Proposition 6.3, there is a geometric, equivariant sweepout $\Psi\colon F \times \mathbb{R} \to N_\psi$. In particular, for every t, we have $F_t = \Psi_t(F)$ is a hyperbolic cone surface with at most one singular point of cone angle $2\pi \leq \theta_t < 4\pi$.

Definition 7.2. Let F_t be a simplicial hyperbolic surface in N (such as one occurring in the sweepout). We say that an arc $a \in \mathcal{A}^{(0)}(F,p)$ is *short* on F_t if a runs from p to p, and if its geodesic representative in the singular hyperbolic metric of F_t is shortest among all such arcs.

Note that a given surface F_t can have multiple short arcs. However, all short arcs on F_t have disjoint representatives, by Lemma 3.3.

For each surface F_t, Lemma 3.8 gives an explicit upper bound on the length of a short arc in F_t (hence, also on its length in N_ψ). On the cusp annulus A, each short arc gives a shadow of a horoball, with area bounded below. More precisely, we obtain the following quantitative estimate.

LEMMA 7.3. *Let $a \in \mathcal{A}^{(0)}(F,p)$ be an arc that is short on F_t for some t. Then, on the cusp annulus $A \subset N_\psi$, the arc a corresponds to a pair of disjoint disks, each of radius*

$$r = \frac{\sqrt{2}}{8\pi^2\,\chi(F)^2}.$$

Furthermore, if an arc b is short on $F_{t'}$ for some t', and $a \neq b \in \mathcal{A}^{(0)}(F,p)$, then the disks corresponding to a and b are disjoint on A.

Proof. By convention (and by Lemma 2.7), the longitude of A has length $\lambda = 2^{1/4}$. Thus, by Lemma 6.5, the intersection between F_t and the horocusp of N_ψ contains an equidistant cusp neighborhood R_t, of area at least $2^{1/4}$. Therefore, by Lemma 3.8, the length of a short arc a in $F_t \smallsetminus R_t$ is

$$\ell(a) < 2\ln\left|2\pi\,\chi(F)/2^{1/4}\right|.$$

Since F_t is immersed in N_ψ as a piecewise geodesic union of hyperbolic triangles, this immersion is distance-decreasing. Thus, in N_ψ, the geodesic g_a in the homotopy class of a must also be shorter than the above estimate.

Lift N_ψ to the universal cover $\mathbb{H}^3$ in the upper half-space model, so that the cusp annulus A lifts to a horizontal horosphere at Euclidean height 1. Let H_∞ be the horoball above this horosphere. This means that g_a lifts to a vertical geodesic that starts at height 1 and ends at the top of a horoball H_a, of diameter

$$e^{-\ell(a)} \geq \frac{\sqrt{2}}{4\pi^2\,\chi(F)^2}. \tag{7.1}$$

In $\mathbb{H}^3$, there is a covering transformation for N_ψ that maps H_a to H_∞, and maps H_∞ to another horoball H'_a. The diameter of H'_a must be the same as that of H_a because they lie at the same distance from H_∞ (namely, the length of g_a). This new horoball H'_a cannot belong to the same orbit as H_a under the parabolic subgroup $\pi_1(A) = \mathbb{Z}$ preserving H_∞: for, this parabolic subgroup preserves the orientation on all the lifts of g_a, but one lift of g_a is oriented downward toward H_a whereas the other is oriented upward from H'_a toward H_∞. Thus the shadows of H_a and H'_a are disjoint disks on the horosphere at height 1, which project to disjoint disks D_a and D'_a on A because H_a and H'_a are in different orbits. After shrinking D_a and D'_a if necessary, we obtain a pair of disjoint disks of radius $\sqrt{2}/8\pi^2\,\chi(F)^2$.

Now, suppose that b is another arc from p to p, not isotopic to a, such that b is short on $F_{t'}$ for some t'. Then, performing the same construction as for a, we obtain a pair of horoballs H_b and H'_b, whose heights also satisfy equation (7.1). The four horoballs H_a, H'_a, H_b, H'_b must lie in distinct orbits of the parabolic $\mathbb{Z}$ subgroup preserving ∞, because a and b are in distinct homotopy classes on F. Thus the four horoballs are disjoint in $\mathbb{H}^3$. Their shadows on H_∞ are not necessarily disjoint. However, if we shrink all four horoballs until their diameter is exactly $\sqrt{2}/4\pi^2\,\chi(F)^2$, then the shadows of disjoint horoballs of the same size are themselves disjoint.

Since H_a, H'_a, H_b, H'_b lie in distinct orbits under $\pi_1(A)$, we obtain four disjoint disks D_a, D'_a, D_b, D'_b in A, each of radius $\sqrt{2}/8\pi^2\,\chi(F)^2$. □

To obtain a lower bound on the area and height of ∂C, we need to find a sequence of arcs in F, each of which is short in a surface F_t for some t.

LEMMA 7.4. *There is a sequence of real numbers* $0 = t_0, t_1, \ldots, t_k = r$ *and an associated sequence* $a_0, a_1, \ldots, a_k$ *of arcs embedded in* F, *with the following properties:*

(1) *Each* a_i *is short on* F_{t_i}.

(2) *The first and last arcs satisfy* $a_k = \psi(a_0)$.

(3) *Each* a_i *is disjoint from* a_{i-1}. *In other words,* $[a_{i-1}, a_i]$ *is an edge of* $\mathcal{A}(F, p)$.

Proof. Let a be an embedded arc in F from p to p. Define

$$S(a) := \{t \in \mathbb{R} \mid a \text{ is shortest on } F_t \text{ among all arcs from } p \text{ to } p\}.$$

In other words, $S(a)$ consists of those values of t for which a is short (as in Definition 7.2). According to the definition of a geometric sweepout (Definition 6.1) the length of any arc varies continuously with t. Since being shortest is a closed condition it follows that the set $S(a)$ is closed. Also, since every surface F_t in the sweepout has a short arc, the line $\mathbb{R}$ is covered by the sets $S(a)$, as a varies over the vertices of $\mathcal{A}(F, p)$.

We claim that the arcs a for which $S(a) \neq \emptyset$ belong to finitely many ψ-orbits. This is because every ψ-orbit of arcs in F descends to a single arc in M_ψ, with distinct orbits descending to distinct arcs. The two endpoints of a geodesic $g_a \subset M_\psi$ representing the arc a must be distinct in ∂C (otherwise, a deck transformation of M_ψ would reverse the orientation of a lift of g_a, fixing a point in the middle). The two endpoints of g_a on ∂C are the centers of disjoint disks guaranteed by Lemma 7.3. Thus, by Lemma 7.3, every orbit of arcs that is shortest on some F_t makes a definite contribution to area(∂C). On the other hand, area(∂C) is bounded above (e.g., by Theorem 4.1), hence there can be only finitely many ψ-orbits of arcs for which $S(a) \neq \emptyset$.

Next, we claim that each $S(a)$ is compact. Fix an arc a, and let g_a be the geodesic in N_ψ that is homotopic to a. Note that the sweepout of N_ψ must eventually *exit*: that is, as $|t| \to \infty$, the surface F_t must leave any compact set in N_ψ. Thus, as $|t| \to \infty$, the distance from F_t to the geodesic g_a becomes unbounded (outside the horocusp C). However, any path in F_t that is homotopic to g_a but remains outside an s-neighborhood of g_a must be extremely long (with length growing exponentially in s). Therefore, when $|t| \gg 0$, the geodesic representative of a on F_t must be very long, and in particular cannot be the shortest arc on F_t. This means that $S(a)$ is bounded, hence compact.

We conclude there are only finitely many arcs a for which $S(a) \cap [0, r] \neq \emptyset$. By the first claim above, these arcs belong to finitely many ψ-orbits. Within each orbit, the compact sets $S(a)$ and $S(\psi(a))$ differ by a translation by r. Thus only finitely many sets in each ψ-orbit can intersect $[0, r]$.

Let a_0 be an arc that is short on F_0; that is, $0 \in S(a_0)$. Then $\psi(a_0)$ is short on F_r. Since the connected interval $[0,r]$ is covered by finitely many closed sets $S(a)$, a lemma in point-set topology (Lemma B.1 in the Appendix) implies that one may "walk" from $S(a_0)$ to $S(\psi(a_0))$ by intersecting sets: there is a real number $t_1 \in S(a_0) \cap S(a_1)$, a number $t_2 \in S(a_1) \cap S(a_2)$, and so on, for arcs $a_0, \dots, a_k = \psi(a_0)$.

By definition, $t_i \in S(a_{i-1}) \cap S(a_i)$ means that both a_{i-1} and a_i are short on F_{t_i}. Thus, by Lemma 3.3, $[a_{i-1}, a_i]$ is an edge of $\mathcal{A}(F,p)$. We have therefore constructed a walk through the 1-skeleton of $\mathcal{A}(F,p)$, from a_0 to $a_k = \psi(a_0)$, through arcs a_i that are each short on some simplicial hyperbolic surface. □

LEMMA 7.5. *The sequence of arcs $a_0, a_1, \dots, a_k$ in Lemma 7.4 contains a subsequence $b_0, b_1, \dots, b_m$ with the following properties:*

(1) *Each b_i is short on F_{t_i}.*
(2) *The arcs $b_1, \dots, b_m$ are all in distinct ψ-orbits.*
(3) *The first and last arcs satisfy $b_m = \psi^n(b_0)$, for some integer $n \neq 0$.*
(4) *Each b_i is disjoint from b_{i-1}. In other words, $[b_{i-1}, b_i]$ is an edge of $\mathcal{A}(F,p)$.*

Proof. The arcs $a_0, a_1, \dots, a_k$ in Lemma 7.4 constitute a walk through the 1-skeleton of $\mathcal{A}(F,p)$, from a_0 to $a_k = \psi(a_0)$. Given that this walk exists, one can excise some of the a_i if necessary to form a loop-erased walk a_0 to $a_k = \psi(a_0)$. That is, one may walk from a_0 to $a_k = \psi(a_0)$ through some subcollection of the a_i, without visiting the same isotopy class more than once.

Next, suppose that there are indices $i < j$, such that a_i and a_j belong to the same ψ-orbit. Without loss of generality, assume that i, j are an innermost pair with this property. This means that $a_j = \psi^n(a_i)$ for some $n \neq 0$, and $a_{i+1}, \dots, a_j$ are all in distinct ψ-orbits. Now, we simply restrict attention to the subsequence from i to j. That is, let $b_0 = a_i$, $b_1 = a_{i+1}$, and so on, until $b_m = a_j$ for $m = j - i$. This subsequence satisfies the lemma. □

We can now complete the proof of Theorem 7.1. Notice that in Lemma 7.5, $b_0, b_1, \dots, b_m$ are the vertices of a path through $\mathcal{A}^{(1)}(F,p)$ from b_0 to $\psi^n(b_0)$. Thus $m \geq d_{\mathcal{A}}(\psi^n)$. By Lemma 7.3, each arc b_i corresponds to two disjoint disks $D_i, D_i' \subset A$, each of radius $\sqrt{2}/8\pi^2\chi(F)^2$. Furthermore, for $j \neq i$, the disks D_i, D_i', D_j, D_j' are all disjoint. Thus we have at least $2d_{\mathcal{A}}(\psi^n)$ disjoint disks altogether. Since the arcs $b_1, \dots, b_m$ are all in different ψ-orbits, these disks project to $2d_{\mathcal{A}}(\psi^n)$ disjoint disks on the cusp torus $\partial C \subset M_\psi$.

To obtain a lower bound on area(∂C), we sum up the areas of these disjoint disks, and multiply by the circle packing constant of $2\sqrt{3}/\pi$ (see [7, Theorem 1]). Thus

$$\text{area}(\partial C) \geq \frac{2\sqrt{3}}{\pi} \cdot 2d_{\mathcal{A}}(\psi^n) \cdot \pi \left(\frac{\sqrt{2}}{8\pi^2 4\,\chi(F)^2} \right)^2 = \frac{\sqrt{3}\, d_{\mathcal{A}}(\psi^n)}{8\pi^4\,\chi(F)^4} > \frac{d_{\mathcal{A}}(\psi^n)}{450\,\chi(F)^4}.$$

Finally, since $\text{area}(\partial C) = \lambda \cdot \text{height}(\partial C)$, and we have normalized the horocusp so that $\lambda = 2^{1/4}$, we have

$$\text{height}(\partial C) \geq \frac{\sqrt{3}\, d_{\mathcal{A}}(\psi^n)}{2^{1/4}\, 8\pi^4\, \chi(F)^4} > \frac{d_{\mathcal{A}}(\psi^n)}{536\, \chi(F)^4}. \qquad \square$$

Proof of Theorem 1.5, lower bound. Let $\varphi\colon F \to F$ be a pseudo-Anosov homeomorphism, and let $\psi = \varphi^n$ be the smallest power of φ that fixes the puncture p. Let C be an embedded cusp of M_φ corresponding to p, whose longitude has length $\lambda = 2^{1/4}$. By Lemma 2.7, an embedded cusp of this size exists, and is smaller than the maximal cusp of M_φ. Thus all lower bounds on C also apply to the maximal cusp.

In the cover M_ψ, the horocusp C lifts to an embedded horocusp $C_1 \subset M_\psi$, which is a one-sheeted cover of C. Furthermore, for every integer $m \geq 1$, the mapping torus M_{ψ^m} contains an embedded horocusp C_m whose longitude has length $2^{1/4}$ and which forms an m-fold cover of C. Consider what Theorem 7.1 says about the geometry of C_m.

By Theorem 7.1, there exists an integer $n(m) \geq m$, such that

$$\text{area}(\partial C_m) > \frac{d_{\mathcal{A}}(\psi^{n(m)})}{450\, \chi(F)^4} \quad \text{and} \quad \text{height}(\partial C_m) > \frac{d_{\mathcal{A}}(\psi^{n(m)})}{536\, \chi(F)^4}.$$

Because C_m is an m-fold cover of C, both its area and its height are m times larger than those of C. Thus, for all $m \geq 1$,

$$\text{area}(\partial C) > \frac{d_{\mathcal{A}}(\psi^{n(m)})}{m \cdot 450\, \chi(F)^4} \geq \frac{d_{\mathcal{A}}(\psi^{n(m)})}{n(m) \cdot 450\, \chi(F)^4} \geq \inf_{r \geq n(m)} \frac{d_{\mathcal{A}}(\psi^r)}{r \cdot 450\, \chi(F)^4}.$$

Since every $m \geq 1$ gives rise to an integer $n(m) \geq m$ satisfying the above inequality, we have

$$\text{area}(\partial C) \geq \liminf_{n \to \infty} \frac{d_{\mathcal{A}}(\psi^n)}{n \cdot 450\, \chi(F)^4} = \frac{\overline{d}_{\mathcal{A}}(\psi)}{450\, \chi(F)^4}.$$

The last equality holds because the limit in equation (1.2), which defines the stable translation distance, always exists [8, p. 230]. An identical calculation goes through for $\text{height}(\partial C)$. $\square$

8. Lower bound: quasi-Fuchsian manifolds. In this section, we prove the lower bounds of Theorem 1.7. The argument uses almost exactly the same ingredients as the proof of Theorem 7.1. The one additional ingredient is the ability to approximate the convex core boundary by surfaces pleated along triangulations.

PROPOSITION 8.1. *Let $N \cong F \times \mathbb{R}$ be a cusped quasi-Fuchsian 3-manifold, and p a puncture of F. As in Definition 1.6, let $\Delta_-(N)$ be the simplex of $\mathcal{A}(F,p)$ whose vertices are the shortest arcs from p to p in the lower core boundary*

$\partial_-\operatorname{core}(N)$. *Then there is an ideal triangulation* τ *of* F*, such that in the pleated surface* F_τ *pleated along* τ*, each shortest arc from* p *to* p *is distance at most* 1 *from* $\Delta_-(N)$.

The same statement holds for the simplex $\Delta_+(N)$.

Proof. Let $R_0 \subset \partial_-\operatorname{core}(N)$ be a horospherical cusp neighborhood about puncture p, such that $\operatorname{area}(R_0) = 1$. Note that by Lemma 2.8, the neighborhood R_0 is embedded. For any arc a from puncture p to p in F, let $\ell_0(a)$ be the length of the geodesic representing a, in the hyperbolic metric of $\partial_-\operatorname{core}(N)$, outside the cusp neighborhood R_0. Let $\ell_N(a)$ be the length of the geodesic representing a in the 3-manifold N, outside a cusp neighborhood of longitude 1. Define the set

(8.1)
$$T(N) = \left\{ a \in \mathcal{A}^{(0)}(F,p) \mid \text{both ends of } a \text{ are at } p,\ \ell_N(a) < 2\ln|6\chi(F)| + \ln(2) \right\}.$$

Note that by Lemmas 2.5 and 3.8, every arc $a \in \Delta_-(N)$ has length

$$\ell_N(a) \le \ell_0(a) \le 2\ln|6\chi(F)|.$$

Thus the simplex $\Delta_-(N)$ of shortest arcs in $\partial_-\operatorname{core}(N)$ must be contained in $T(N)$. Also note that $T(N)$ must be a finite set: one way to see this is to recall (e.g., from Lemma 7.3) that every arc of bounded length makes a definite contribution to $\operatorname{area}(\partial C \cap \operatorname{core}(N))$.

Now, we apply Theorem A.1: there exists a sequence of triangulations τ_i of F, such that the hyperbolic metrics on the pleated surfaces F_{τ_i} converge in the Teichmüller space $\mathcal{T}(F)$ to $\partial_-\operatorname{core}(N)$. For each i, let $R_i \subset F_{\tau_i}$ be an embedded cusp neighborhood of area 1. For each arc $a \subset F$, let $\ell_i(a)$ be the length of the geodesic representing a, in the induced hyperbolic metric on F_{τ_i}, relative to the cusp neighborhood R_i. Note that $\ell_N(a) \le \ell_i(a)$ by Lemma 2.5. Then, because the metrics on F_{τ_i} converge to that on $\partial_-\operatorname{core}(N)$, the length of any arc also converges. In particular, because $T(N)$ is a finite set of arcs, there is some $k \gg 0$ such that

$$|\ell_k(a) - \ell_0(a)| < \ln(2)/3, \quad \forall a \in T(N). \tag{8.2}$$

Let $\tau = \tau_k$, and let b be any shortest arc on $F_\tau = F_{\tau_k}$. By Lemma 3.8, we have $\ell_k(b) \le 2\ln|6\chi(F)|$. Furthermore, $\ell_N(b) \le \ell_k(b)$, since the pleating map that produces F_{τ_k} is 1-Lipschitz. Then, by equation (8.1), it follows that $b \in T(N)$. Thus, for any arc a that is shortest on $\partial_-\operatorname{core}(N)$ (that is, $a \in \Delta_-(N)$), we obtain

$$\ell_0(b) - \ln(2)/3 < \ell_k(b) \le \ell_k(a) < \ell_0(a) + \ln(2)/3, \tag{8.3}$$

where the outer inequalities follow by (8.2) and the middle inequality holds because b is shortest on F_τ. A simpler way to state the conclusion of (8.3) is that on the hyperbolic surface $\partial_- \operatorname{core}(N)$,

$$\ell_0(b) < \ell_0(a) + 2\ln(2)/3.$$

Since a is shortest on $\partial_- \operatorname{core}(N)$, and b is nearly shortest, Lemma 3.4 implies that a and b are disjoint (or the same arc) for any $a \in \Delta_-(N)$. Thus b is distance at most 1 from $\Delta_-(N)$. □

We can now begin proving the lower bound of Theorem 1.7. Applying the same ideas as in Section 7 gives the following analogue of Lemma 7.4.

LEMMA 8.2. *Let $N \cong F \times \mathbb{R}$ be a cusped quasi-Fuchsian 3-manifold, and p a puncture of F. There is a sequence $a_0, a_1, \ldots, a_k$ of arcs embedded in F, and an associated sequence of simplicial hyperbolic surfaces $F_{t(a_i)} \subset N$ with at most one singular point of cone angle $2\pi \leq \theta_t < 4\pi$, such that the following hold:*

(1) *Each a_i is short on $F_{t(a_i)}$, in the sense of Definition 7.2.*
(2) *The arcs $a_0, \ldots, a_k$ are distinct up to isotopy.*
(3) *$F_{t(a_0)}$ is the lower core boundary $\partial_- \operatorname{core}(N)$, and $F_{t(a_k)} = \partial_+ \operatorname{core}(N)$.*
(4) *Each a_i is disjoint from a_{i-1}. In other words, $[a_{i-1}, a_i]$ is an edge of $\mathcal{A}(F, p)$.*

Proof. As in Definition 1.6, let $\Delta_-(N)$ be the simplex of $\mathcal{A}(F, p)$ whose vertices are the short arcs on the lower boundary $\partial_- \operatorname{core}(N)$. By Proposition 8.1, there is a triangulation τ of F, and a pleated surface F_τ pleated along τ, such that a short arc on this pleated surface either belongs to the simplex $\Delta_-(N)$, or is at distance 1 from some vertex of $\Delta_-(N)$. Similarly, there is a triangulation τ' of F, and a pleated surface $F_{\tau'}$ pleated along τ', whose short arc either belongs to the simplex $\Delta_+(N)$, or is at distance 1 from some vertex of $\Delta_+(N)$.

By Proposition 6.4, there is a geometric sweepout $\Psi \colon F \times [0, r] \to N$, where each $F_t = \Psi(F \times \{t\})$ is a hyperbolic cone surface with at most one singular point of cone angle $2\pi \leq \theta_t < 4\pi$. Furthermore, $F_0 = F_\tau$ and $F_r = F_{\tau'}$, for the given triangulations τ and τ'.

Next, we apply the argument of Lemma 7.4. In the quasi-Fuchsian setting, the proof simplifies in several ways. There is no need to worry about equivariance, and the set of arcs $\{a : S(a) \neq \emptyset\}$ is finite because it is contained in $T(N)$ from equation (8.1). We thus obtain a sequence of arcs $a_1, \ldots, a_{k-1}$ with the following properties:

- Each a_i is short on some surface $F_{t(a_i)}$ in the sweepout.
- Arc a_1 is short on $F_0 = F_\tau$, and a_{k-1} is short on $F_r = F_{\tau'}$.
- Each a_i is disjoint from a_{i-1}. In other words, $[a_{i-1}, a_i]$ is an edge of $\mathcal{A}(F, p)$.

Next, we extend this sequence of arcs to the convex core boundary. If $a_1 \notin \Delta_-(N)$, then there is an arc $a_0 \in \Delta_-(N)$, which is shortest on $\partial_- \operatorname{core}(N)$ by definition, and such that $[a_0, a_1]$ is an edge of $\mathcal{A}(F, p)$. Otherwise, if $a_1 \in \Delta_-(N)$, then we simply shift indices by 1, so that a_1 becomes a_0. Similarly, if $a_{k-1} \notin \Delta_+(N)$, then we add an arc $a_k \in \Delta_+(N)$, which is shortest on $\partial_+ \operatorname{core}(N)$. Otherwise, if $a_{k-1} \in \Delta_+(N)$, then we simply redefine $k := k-1$, and stop the sequence there.

We now have a sequence of arcs $a_0, \dots, a_k$, with associated simplicial hyperbolic surfaces $F_{t(a_i)}$, so that this sequence satisfies all the conclusions of the lemma except possibly (2). That is, some of the a_i might be in the same isotopy class. But if the arcs $a_0, \dots, a_k$ are the vertices of a path in $\mathcal{A}^{(1)}(F,p)$, then some subcollection of the a_i give an embedded path. This means that (2) is satisfied, and the proof is complete. □

LEMMA 8.3. *Let $N \cong F \times \mathbb{R}$ be a cusped quasi-Fuchsian manifold, and p a puncture of F. Let $C \subset N$ be a horocusp corresponding to the puncture p, whose longitude is $\lambda = 2^{1/4}$. Then, for some $k \geq d_{\mathcal{A}}(N,p)$, the annulus $A = \partial C$ contains $2k+2$ disjoint disks, each of radius*

$$r = \frac{\sqrt{2}}{8\pi^2 \chi(F)^2},$$

such that the center of each disk is in the convex core $\operatorname{core}(N)$.

Recall that by Lemma 2.7, there is indeed an embedded horocusp of longitude $\lambda = 2^{1/4}$. Recall as well, from Definition 1.6, that $d_{\mathcal{A}}(N,p)$ is defined to be the shortest distance in $\mathcal{A}(F,p)$ between a vertex of $\Delta_-(N)$ and a vertex of $\Delta_+(N)$.

Proof of Lemma 8.3. The sequence of arcs $a_0, \dots, a_k$, constructed in Lemma 8.2, is a walk through the 1-skeleton of $\mathcal{A}(F,p)$ from a vertex of $\Delta_-(N)$ to a vertex of $\Delta_+(N)$. Thus, by Definition 1.6, $k \geq d_{\mathcal{A}}(N,p)$.

For each $i \in \{0, \dots, k\}$, let g_{a_i} be the geodesic in N in the homotopy class of a_i. Then, Lemma 7.3 guarantees that there is a pair of disjoint disks D_{a_i} and D'_{a_i}, of radius $r = \sqrt{2}/8\pi^2 \chi(F)^2$, whose centers are the endpoints of g_{a_i} on the cusp annulus A. Since the geodesic g_{a_i} is contained in the convex core of N, the centers of D_{a_i} and D'_{a_i} lie in $\operatorname{core}(N)$ as well.

Finally, Lemma 7.3 also implies that if $j \neq i$, the disks of a_i are disjoint from those of a_j. Thus we have at least $2k+2$ disks in total. □

Proof of Theorem 1.7, lower bound. Let p be a puncture of F, and let $C \subset N$ be a horospherical cusp corresponding to the puncture p, whose longitude has length $\lambda = 2^{1/4}$. Note that by Lemma 2.7, C is contained in the maximal cusp about puncture p. Thus lower bounds on the area and height of $\partial C \cap \operatorname{core}(N)$ also apply to the maximal cusp.

As in Section 4 and 5, we may place Euclidean coordinates on the annulus $A = \partial C \cong S^1 \times \mathbb{R}$, in which the $\mathbb{R}$ direction is vertical. Let $B \subset A$ be a compact annular band, with boundary consisting of horizontal circles, which is the smallest such band that contains all of the $2k+2$ disks of Lemma 8.3. Since each disk has radius $r = \sqrt{2}/8\pi^2 \chi(F)^2$, Böröczky's estimate on the density of a circle packing

[7, Theorem 1] implies that

$$\begin{aligned}
\operatorname{area}(B) &\geq \frac{2\sqrt{3}}{\pi}\cdot(2d_{\mathcal{A}}(N)+2)\cdot\pi\left(\frac{\sqrt{2}}{8\pi^2\chi(F)^2}\right)^2\\
&= \frac{\sqrt{3}\,(d_{\mathcal{A}}(N)+1)}{8\pi^4\chi(F)^4}\\
&> \frac{d_{\mathcal{A}}(N)+1}{450\,\chi(F)^4}.
\end{aligned}$$

Similarly, since $\operatorname{area}(B)=\lambda\cdot\operatorname{height}(B)$, and we have normalized the horocusp so that $\lambda = 2^{1/4}$, we have

$$\operatorname{height}(B) \geq \frac{\sqrt{3}\,(d_{\mathcal{A}}(N)+1)}{2^{1/4}\,8\pi^4\chi(F)^4} > \frac{(d_{\mathcal{A}}(N)+1)}{536\,\chi(F)^4}.$$

To complete the proof, it remains is to bound the difference in area (or height) between B and $\partial C \cap \operatorname{core}(N)$. Note that by Lemma 8.3, each of the $(2k+2)$ disks has its center inside $\operatorname{core}(N)$. Therefore, the upper boundary of B is at most $r = \sqrt{2}/8\pi^2\chi(F)^2$ higher than $\partial_+\operatorname{core}(N)$, and the lower boundary of B is at most r lower than $\partial_-\operatorname{core}(N)$. This implies that

$$\begin{aligned}
\operatorname{area}(\partial C\cap\operatorname{core}(N)) &\geq \operatorname{area}(B) - 2\lambda r\\
&> \frac{d_{\mathcal{A}}(N)}{450\,\chi(F)^4} - \frac{2\cdot 2^{3/4}}{8\pi^2\chi(F)^2}\\
&> \frac{d_{\mathcal{A}}(N)}{450\,\chi(F)^4} - \frac{1}{23\,\chi(F)^2}.
\end{aligned}$$

Similarly,

$$\operatorname{height}(\partial C\cap\operatorname{core}(N)) \geq \operatorname{height}(B) - 2r > \frac{d_{\mathcal{A}}(N)}{536\,\chi(F)^4} - \frac{1}{27\,\chi(F)^2}. \qquad \square$$

9. Covers and the arc complex. In this section, we will apply Theorem 1.7 to prove Theorem 1.9, which relates the arc complex of a surface S to that of its cover Σ. The proof uses some classical results in Kleinian groups to construct a quasi-Fuchsian manifold with prescribed short arcs on its convex core boundary (see Lemma 9.2). We begin by recalling some terminology and results, while pointing the reader to Marden [28, Chapter 3] for a more detailed reference.

Let Γ be a Kleinian group with limit set Λ. The *domain of discontinuity* is $\Omega = \partial\mathbb{H}^3 \smallsetminus \Lambda$. When $N = \mathbb{H}^3/\Gamma$ is quasi-Fuchsian, Ω is the disjoint union of two open disks Ω_+ and Ω_-, each of which admits a conformal, properly discontinuous action by Γ. The quotients $S_\pm = \Omega_\pm/\Gamma$ are Riemann surfaces, called the (top and bottom) *conformal boundary* of $N = \mathbb{H}^3/\Gamma$.

For each quasi-Fuchsian manifold N, there is a natural "nearest point retraction" map $r\colon S_\pm \to \partial_\pm \operatorname{core}(N)$. Sullivan proved that, if $S_\pm$ is given the unique hyperbolic metric in its conformal class, the map $r\colon S_\pm \to \partial_\pm \operatorname{core}(N)$ is K-Lipschitz, for a universal constant $K > 1$. Much more recently, Epstein, Marden, and Markovic [18] showed that the optimal Lipschitz constant is 2. We will use their result for concreteness, while emphasizing that Sullivan's original K-Lipschitz statement is all that is truly needed.

We also recall some facts from the geometry of surfaces. Let ε_2 be the 2-dimensional Margulis constant. For any simple closed geodesic γ in a hyperbolic surface S, of length $\ell = \ell(\gamma) < \varepsilon_2$, the ε_2-thin region of S containing γ is an embedded collar of radius $r(\ell)$. The function $r(\ell)$ is monotonically decreasing, and $r(\ell) \to \infty$ as $\ell \to 0$. See Buser [12] for explicit estimates on ε_2 and $r(\ell)$.

Definition 9.1. Let S be a hyperbolic surface, and γ a simple closed geodesic on S. We say that γ is *sufficiently thin* if its length $\ell = \ell(\gamma)$ is short enough that the ε_2-thin collar about γ has radius

$$r(\ell) > \ln|6\chi(S)/\varepsilon_2|. \tag{9.1}$$

This requirement on collar radius is motivated by Lemma 3.8.

Given this background, we can prove the following constructive lemma.

LEMMA 9.2. *Let F be a surface with a puncture p, and let $a_-, a_+ \in \mathcal{A}^{(0)}(F,p)$ be arcs from p to p. Then there exists a quasi-Fuchsian manifold $N \cong F \times \mathbb{R}$, such that $a_\pm$ is the unique shortest arc from p to p on $\partial_\pm \operatorname{core}(N)$.*

Proof. Let $R_0 \subset F$ be a neighborhood of the puncture p, and let $R(a_+)$ be a regular neighborhood of $R_0 \cup a_+$. This is topologically a pair of pants, whose frontier in F consists of a pair of essential closed curves c_+, c'_+. (If F is a once-punctured torus, then c_+ is isotopic to c'_+; this will not affect our arguments.) Similarly, let c_-, c'_- be closed curves that form the frontier of a regular neighborhood $R(a_-)$.

Choose hyperbolic metrics $X_\pm$ on F, in which the geodesic representatives of $c_\pm$ and $c'_\pm$ have less than half the length required to be sufficiently thin. By Bers simultaneous uniformization [28, page 136], there is a quasi-Fuchsian manifold $N \cong F \times \mathbb{R}$ whose top conformal boundary is X_+ and whose bottom conformal boundary is X_-. Thus, by Epstein, Marden, and Markovic [18], the geodesic representatives of $c_\pm$ and $c'_\pm$ are sufficiently thin on $\partial_\pm \operatorname{core}(N)$.

We claim that a_+ is the unique shortest arc from p to p on $\partial_+ \operatorname{core}(N)$. For concreteness, we will measure lengths relative to the horospherical cusp neighborhood $Q(p) \subset \partial_+ \operatorname{core}(N)$ whose boundary has length ε_2. Since horoballs are convex, it follows that $Q(p)$ is ε_2-thin. Then, by Lemma 3.8, there must be a geodesic α_+ from p to p whose length relative to this cusp neighborhood satisfies

$$\ell(\alpha_+) \le 2\ln|6\chi(S)/\varepsilon_2|.$$

Now, let b_+ be any arc from p to p, other than a_+. Since b_+ is not isotopic into the neighborhood $R(a_+)$, it must cross $c_+ \cup c'_+$. By the Margulis lemma, the ε_2-thin collars about those curves are disjoint from the ε_2-thin cusp neighborhood $Q(p)$. Thus, since c_+ and c'_+ are sufficiently thin, and the width of a collar is twice the radius, equation (9.1) implies that $\ell(b_+) > \ell(\alpha_+)$. Therefore $a_+ = \alpha_+$, the unique shortest arc from p to p.

By the same argument, a_- is the unique short arc on $\partial_- \operatorname{core}(N)$. □

We can now complete the proof of Theorem 1.9, which we restate.

THEOREM 1.9. *Let Σ and S be surfaces with one puncture, and $f\colon \Sigma \to S$ a covering map of degree n. Let $\pi\colon \mathcal{A}(S) \to \mathcal{A}(\Sigma)$ be the lifting relation induced by f. Then, for all $a,b \in \mathcal{A}^{(0)}(S)$, we have*

$$\frac{d(a,b)}{4050\, n\, \chi(S)^6} - 2 < d(\alpha,\beta) \le d(a,b)$$

where $\alpha \in \pi(a)$ and $\beta \in \pi(b)$.

Proof. Recall that by Definition 1.8, $\pi(a)$ is the set of n vertices in $\mathcal{A}(\Sigma)$ representing arcs that project to a. These vertices form a simplex in $\mathcal{A}(\Sigma)$, since the arcs that comprise $f^{-1}(a)$ are disjoint. Similarly, if a and a' are distance 1 in $\mathcal{A}(S)$, then all the $2n$ lifts of a and a' are disjoint in Σ, hence every vertex $\alpha \in \pi(a)$ is distance 1 in $\mathcal{A}(\Sigma)$ from every vertex $\alpha' \in \pi(a')$. Thus, by induction on distance, we have

$$d(\alpha,\beta) \le d(a,b)$$

for any $\alpha \in \pi(a)$ and any $\beta \in \pi(b)$.

To prove the other inequality in Theorem 1.9, assume the cover is non-trivial: that is, $n > 1$. Fix $a,b \in \mathcal{A}^{(0)}(\Sigma)$. By Lemma 9.2, there is a quasi-Fuchsian manifold $M \cong S \times \mathbb{R}$, such that a is the unique shortest arc on $\partial_- \operatorname{core}(M)$ and b is the unique shortest arc on $\partial_+ \operatorname{core}(M)$.

Since S has a unique puncture p, we have $\mathcal{A}(S) = \mathcal{A}(S,p)$. Let $C \subset M$ be the maximal horospherical cusp corresponding to this unique puncture. By Definition 1.6, $d_{\mathcal{A}}(M,p) = d(a,b)$. Thus, by the lower bound of Theorem 1.7,

$$\frac{d(a,b)}{450\,\chi(S)^4} - \frac{1}{23\,\chi(S)^2} < \operatorname{area}(\partial C \cap \operatorname{core}(M)). \tag{9.2}$$

Use the n-fold covering map $f\colon \Sigma \to S$ to lift the hyperbolic metric on $M \cong S \times \mathbb{R}$ to a quasi-Fuchsian structure on $N \cong \Sigma \times \mathbb{R}$. The convex core of N covers the convex core of M. The horocusp $C \subset M$ lifts to a horocusp $D \subset N$, which is an n-sheeted cover of C. Thus

$$n \cdot \operatorname{area}(\partial C \cap \operatorname{core}(M)) = \operatorname{area}(\partial D \cap \operatorname{core}(N)). \tag{9.3}$$

Observe that every arc $c \in \partial_+ \operatorname{core}(M)$ lifts to n disjoint arcs in $\partial_+ \operatorname{core}(N)$, each of which has the same length as c (outside the horocusps C and D, respectively). Thus the n arcs of $\pi(b)$ are shortest on $\partial_+ \operatorname{core}(N)$, and similarly the n arcs of $\pi(a)$ are shortest on $\partial_- \operatorname{core}(N)$. By Definition 1.6, this implies $d_{\mathcal{A}}(N,p) \le d(\alpha,\beta)$ for any $\alpha \in \pi(a)$ and any $\beta \in \pi(b)$. Therefore, the upper bound of Theorem 1.7 implies

$$\text{area}(\partial D \cap \operatorname{core}(N)) < 9\chi(\Sigma)^2\, d(\alpha,\beta) + \Big|12\chi(\Sigma)\ln|\chi(\Sigma)| + 26\chi(\Sigma)\Big|. \tag{9.4}$$

Combining equations (9.2), (9.3) and (9.4), we obtain

$$\frac{n\cdot d(a,b)}{450\,\chi(S)^4} - \frac{n}{23\,\chi(S)^2} < 9\chi(\Sigma)^2\, d(\alpha,\beta) + \Big|12\chi(\Sigma)\ln|\chi(\Sigma)| + 26\chi(\Sigma)\Big|,$$

which can be rearranged, using $\chi(\Sigma) = n\chi(S)$, to give

$$\frac{d(a,b)}{4050\,n\,\chi(S)^6} - \frac{1}{23\cdot 9\,n\,\chi(S)^4} - \frac{\Big|12\chi(\Sigma)\ln|\chi(\Sigma)| + 26\chi(\Sigma)\Big|}{9\chi(\Sigma)^2} < d(\alpha,\beta).$$

Since Σ is a once-punctured surface that non-trivially covers S, it has Euler characteristic $|\chi(\Sigma)| \ge 3$. It follows that the additive error on the left-hand side is bounded above by 2. This completes the proof. □

Appendix A. Approximating the convex core boundary. The goal of this appendix is to write down a proof of the following result, which is needed in the argument of Section 8.

THEOREM A.1. *Let $N \cong F \times \mathbb{R}$ be a cusped quasi-Fuchsian 3-manifold. Then there is a sequence τ_i of ideal triangulations of F, such that the induced hyperbolic metrics on the pleated surfaces F_{τ_i} converge in the Teichmüller space $\mathcal{T}(F)$ to the hyperbolic metric on the lower core boundary $\partial_- \operatorname{core}(N)$. Furthermore, the pleating maps for the F_{τ_i} converge in the compact-open topology to a pleating map for $\partial_- \operatorname{core}(N)$.*

The same statement holds for the upper core boundary $\partial_+ \operatorname{core}(N)$.

The statement of Theorem A.1 is entirely unsurprising, and morally it should fit into the toolbox of well-known results about laminations and pleated surfaces [14, Chapters 4 and 5]. Indeed, the standard toolbox of Kleinian group theory leads to a relatively quick proof of the theorem. However, this short proof is also somewhat technical, as it requires passing between several different topologies on spaces of laminations and pleated surfaces.

The main reference for the following argument is Canary, Epstein, and Green [14]. See also Thurston [37, 39], Bonahon [6], and Ohshika [31].

Definition A.2. Let S be a punctured hyperbolic surface of finite area. Let $\mathcal{GL}(S)$ denote the set of geodesic laminations on S: that is, laminations where each leaf is a geodesic. We equip $\mathcal{GL}(S)$ with the *Chabauty topology*. In this topology, a sequence $\{L_i\}$ converges to $L \in \mathcal{GL}(S)$ if and only if:

(1) If a subsequence $x_{n_i} \in L_{n_i}$ converges to $x \in S$, then $x \in L$.
(2) For all $x \in L$, there exists a sequence $x_i \in L_i$, such that $x_i \to x$.

The Chabauty topology is metrizable [14, Proposition 3.1.2]. In fact, when restricted to compact laminations, the Chabauty topology reduces to be the Hausdorff topology (induced by the Hausdorff distance between compact sets). See [14, Section 3.1] for more details.

Definition A.2 makes use of a hyperbolic metric on S, but in an inessential way. If we modify a metric d to a new hyperbolic metric d', each lamination $L \in \mathcal{GL}(S)$ that is geodesic in d can be straightened to a geodesic lamination of d'. This straightening does not affect convergence of laminations. Thus the space $\mathcal{GL}(S)$ only depends on the topology of S.

We use the term *curve* to denote a simple closed geodesic in S. The following lemma is a good example of convergence in the Chabauty topology.

LEMMA A.3. *For every curve $\alpha \subset S$, there is a sequence of ideal triangulations τ_i, converging in the Chabauty topology to a lamination $\alpha' \supset \alpha$.*

Proof. Fix any ideal triangulation τ. Let $D = D_\alpha$ be a Dehn twist about α. Then $\tau_i = D^i(\tau)$ converges to the desired α'. □

Definition A.4. Let S be a punctured hyperbolic surface of finite area. Then $\mathcal{ML}(S)$ denotes the space of compact, transversely measured laminations. Every point of $\mathcal{ML}(S)$ is a pair (L,μ) where L is a compact geodesic lamination and μ is a transverse measure of full support. That is, for each arc α intersecting L transversely, $\mu(\alpha)$ is a positive number that stays invariant under an isotopy preserving the leaves of L. The natural topology on $\mathcal{ML}(S)$ is called the *measure topology*.

Let $\mathcal{PML}(S)$ denote the projectivization of $\mathcal{ML}(S)$, in which nonzero measures that differ by scaling become identified. The measure topology on $\mathcal{ML}(S)$ descends to $\mathcal{PML}(S)$.

If L is a disjoint union of arcs and closed curves in S, an example of a transverse measure is the *counting measure*, where $\mu(\alpha) = |\alpha \cap L|$. For another example, suppose that $S = \partial_+ \operatorname{core}(N)$ is the upper boundary of the convex core in a quasi-Fuchsian 3-manifold. Then the pleating lamination L has a *bending measure*, where $\mu(\alpha)$ is the integral of the bending of α as it crosses leaves of L.

The measure-forgetting map $\mathcal{PML}(S) \to \mathcal{GL}(S)$ is not continuous, but it has the following convenient property.

Fact A.5. Suppose that $(L_i,\mu_i) \to (L,\mu) \in \mathcal{PML}(S)$, in the measure topology. Then, after passing to a subsequence, there is a lamination $L' \in \mathcal{GL}(S)$ so that $L \subset L'$ and $L_i \to L'$ in the Chabauty topology.

With these facts in hand, we can prove Theorem A.1. The proof contains two steps, the first of which deals with laminations only.

LEMMA A.6. *Let $L \in \mathcal{GL}(S)$ be a measurable lamination. Then there is a sequence of ideal triangulations τ_i converging in the Chabauty topology to a lamination $L'' \supset L$.*

Proof. Pick μ, a measure of full support on L, so that $(L, \mu) \in \mathcal{PML}(S)$.

Thurston proved that curves, equipped with the counting measure, are dense in $\mathcal{PML}(S)$ [6, Proposition 15]. Thus we may pick a sequence of curves $\{\alpha_i\}$ converging to (L, μ) in the measure topology. By Fact A.5, we may pass to a subsequence and reindex so that α_i converges, in the Chabauty topology, to a lamination L' containing L.

By Lemma A.3, we may choose $\{\tau_{i,j}\}$, a sequence of sequences of ideal triangulations, so that for all i

$$\tau_{i,j} \to \alpha_i' \supset \alpha_i \quad \text{as } j \to \infty.$$

Now, choose an increasing function $j \colon \mathbb{N} \to \mathbb{N}$, so that for all i, $\tau_{i,j(i)}$ and α_i' have distance at most $1/i$, in the metric that induces the Chabauty topology.

Claim. The sequence $\tau_{i,j(i)}$ contains a subsequence converging to $L'' \supset L'$.

Proof of claim. As $\mathcal{GL}(S)$ is compact [14, Proposition 4.1.6], the sequence $\tau_{i,j(i)}$ contains a convergent subsequence. In an abuse of notation, denote this convergent subsequence by $\tau_i = \tau_{i,j(i)}$. Let L'' be the limit of the τ_i. Note that L'' is not compact.

Recall that the Chabauty distance between τ_i and α_i' is at most $1/i$. Thus the sequence α_i' has the same limit as τ_i, namely L''. Since $\alpha_i \subset \alpha_i'$, it follows by [14, Lemma 4.1.8] that the limit of α_i must be contained in the limit of α_i'. That is, $L' \subset L''$, as desired. □

Since $L \subset L'$, Lemma A.6 is proven. □

The second step of the argument connects the above discussion of laminations to pleated surfaces. Following Definition 2.2, we call a lamination $L \in \mathcal{GL}(S)$ *realizable* if it is realized by a pleating map $f \colon S \to N$ homotopic to a prescribed map f_0.

LEMMA A.7. *Let $N \cong S \times \mathbb{R}$ be a cusped quasi-Fuchsian 3-manifold, and fix an embedding $f_0 \colon S \to S \times \{0\}$. Let $L_i \in \mathcal{GL}(S)$ be a sequence of laminations on S, which are realizable by pleating maps homotopic to f_0. Then, if $L_i \to L'$ in the Chabauty topology, and $L \subset L'$ is also realizable, the hyperbolic metrics d_i induced by pleating along L_i converge in $\mathcal{T}(S)$ to the hyperbolic metric d induced by pleating along L.*

In fact, the pleating maps $f_i\colon S \to N$ converge to the pleating map $f\colon S \to N$ that realizes L, in the space $\mathcal{MPS}(S,N)$ of marked pleated surfaces homotopic to f_0.

We refer the reader to [14, Definition 5.2.14] for the full definition of $\mathcal{MPS}(S,N)$. It suffices to note that convergence in $\mathcal{MPS}(S,N)$ involves *both* convergence of metrics in $\mathcal{T}(S)$ and convergence of pleating maps in the compact-open topology.

Proof of Lemma A.7. Let C be an embedded neighborhood of the cusps of N. Then $K := \operatorname{core}(N) \smallsetminus C$ is a compact set, and every pleated surface homotopic to f_0 must intersect K.

With this notation, [14, Theorem 5.2.18] implies that the space $\mathcal{MPS}(S,N) = \mathcal{MPS}(S,K)$ of marked pleated surfaces that intersect K is compact. In particular, the pleating maps $f_i\colon S \to N$, which pleat along lamination L_i, have a convergent subsequence, $f_{n_i} \to f_\infty$. Let L_∞ be the pleating lamination of f_∞.

Now, recall that $L_i \to L'$ in the Chabauty topology, and $L \subset L'$ is realizable by a pleating map $f\colon S \to N$ homotopic to f_0. Since every component of $f(S \smallsetminus L)$ is totally geodesic in N, the leaves of $L' \smallsetminus L$ are mapped into these totally geodesic regions, hence L' is realized by the same map f. In this setting, Ohshika [31, Lemma 1.3] notes that $f_\infty = f$ is the same pleating map that realizes the limiting lamination L'.

Therefore, every convergent subsequence of f_i must limit to the same map $f_\infty = f$ that realizes the lamination L. Since $\mathcal{MPS}(S,K)$ is compact, this means that $f_i \to f$ in the topology on $\mathcal{MPS}(S,K)$. This means that $f_i \to f$ in the compact-open topology, and also that the hyperbolic metrics d_i induced by f_i converge in $\mathcal{T}(S)$ to the hyperbolic metric d induced by f. □

Proof of Theorem A.1. Let $N \cong F \times \mathbb{R}$ be a cusped quasi-Fuchsian 3-manifold, as in the statement of the theorem. Let L be the pleating lamination on the lower core boundary $\partial_- \operatorname{core}(N)$. The bending measure μ on $\partial_- \operatorname{core}(N)$ is a transverse measure of full support, so L is a measurable lamination. By Lemma A.6, there is a sequence of ideal triangulations $\tau_i \to L'$ in the Chabauty topology, where $L \subset L'$. Now, by Lemma A.7, the pleating maps $f_i\colon F \to N$ that pleat along τ_i induce hyperbolic metrics on F that converge in $\mathcal{T}(F)$ to the induced metric on $\partial_- \operatorname{core}(N)$. By the definition of $\mathcal{MPS}(S,N)$, these pleating maps also converge in the compact–open topology to a pleating map for $\partial_- \operatorname{core}(N)$. □

Appendix B. A lemma in point-set topology. Let X be a topological space, and let $\mathcal{S}$ be a cover of X by closed sets. Define a *discrete walk of length k* through sets of $\mathcal{S}$ to be a sequence of points $x_0, x_1, \ldots, x_k$, such that $x_0, x_1 \in S_1$, $x_1, x_2 \in S_2$, and so on, for sets $S_1, \ldots, S_k$ that belong to $\mathcal{S}$. We say this sequence is a walk from x_0 to x_k.

The following observation is needed in the proof of Lemma 7.4.

LEMMA B.1. *Let X be a connected topological space, and let $S_1, \dots, S_n$ be closed sets whose union is X. Then, for any pair of points $x, y \in X$, there is a discrete walk from x to y through sets in the collection $\{S_1, \dots, S_n\}$.*

Proof. Define an equivalence relation $\equiv$ on X, where

$$x \equiv y \iff \text{there exists a discrete walk from } x \text{ to } y \text{ through } \{S_1, \dots, S_n\}.$$

Reversing and concatenating walks proves this is an equivalence relation. Furthermore, each closed set S_i must be entirely contained in an equivalence class, because its points are connected by a walk of length 1. Thus each equivalence class is closed. Since the connected space X cannot be expressed as a disjoint union of finitely many closed sets, all of X must be in the same equivalence class. □

DEPARTMENT OF MATHEMATICS, TEMPLE UNIVERSITY, PHILADELPHIA, PA 19122
E-mail: dfuter@temple.edu

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF WARWICK, COVENTRY CV4 7AL, UK
E-mail: s.schleimer@warwick.ac.uk

REFERENCES

[1] C. C. Adams, Waist size for cusps in hyperbolic 3-manifolds II, preprint.
[2] I. Agol, Bounds on exceptional Dehn filling, *Geom. Topol.* **4** (2000), 431–449.
[3] ———, Small 3-manifolds of large genus, *Geom. Dedicata* **102** (2003), 53–64.
[4] ———, The virtual Haken conjecture. With an appendix by Ian Agol, Daniel Groves, and Jason Manning, *Doc. Math.* **18** (2013), 1045–1087.
[5] H. Akiyoshi, H. Miyachi, and M. Sakuma, Variations of McShane's identity for punctured surface groups, *Spaces of Kleinian Groups*, *London Math. Soc. Lecture Note Ser.*, vol. 329, Cambridge Univ. Press, Cambridge, 2006, pp. 151–185.
[6] F. Bonahon, Geodesic laminations on surfaces, *Laminations and Foliations in Dynamics, Geometry and Topology (Stony Brook, NY, 1998)*, *Contemp. Math.*, vol. 269, Amer. Math. Soc., Providence, RI, 2001, pp. 1–37.
[7] K. Böröczky, Packing of spheres in spaces of constant curvature, *Acta Math. Acad. Sci. Hungar.* **32** (1978), no. 3-4, 243–261.
[8] M. R. Bridson and A. Haefliger, *Metric Spaces of Non-Positive Curvature*, *Grundlehren Math. Wiss.*, vol. 319, Springer-Verlag, Berlin, 1999.
[9] J. F. Brock, The Weil-Petersson metric and volumes of 3-dimensional hyperbolic convex cores, *J. Amer. Math. Soc.* **16** (2003), no. 3, 495–535.
[10] ———, Weil-Petersson translation distance and volumes of mapping tori, *Comm. Anal. Geom.* **11** (2003), no. 5, 987–999.
[11] J. F. Brock, R. D. Canary, and Y. N. Minsky, The classification of Kleinian surface groups, II: The ending lamination conjecture, *Ann. of Math. (2)* **176** (2012), no. 1, 1–149.

[12] P. Buser, *Geometry and Spectra of Compact Riemann Surfaces*, *Progr. Math.*, vol. 106, Birkhäuser, Boston, MA, 1992.

[13] R. D. Canary, A covering theorem for hyperbolic 3-manifolds and its applications, *Topology* **35** (1996), no. 3, 751–778.

[14] R. D. Canary, D. B. A. Epstein, and P. Green, Notes on notes of Thurston, *Analytical and Geometric Aspects of Hyperbolic Space (Coventry/Durham, 1984)*, *London Math. Soc. Lecture Note Ser.*, vol. 111, Cambridge Univ. Press, Cambridge, 1987, pp. 3–92.

[15] D. Cooper, D. Futer, and J. S. Purcell, Dehn filling and the geometry of unknotting tunnels, *Geom. Topol.* **17** (2013), no. 3, 1815–1876.

[16] D. Cooper, C. D. Hodgson, and S. P. Kerckhoff, *Three-Dimensional Orbifolds and Cone-Manifolds. With a Postface by Sadayoshi Kojima*, *MSJ Memoirs*, vol. 5, Mathematical Society of Japan, Tokyo, 2000.

[17] M. Culler, N. M. Dunfield, and J. R. Weeks, SnapPy, a computer program for studying the geometry and topology of 3-manifolds, http://www.math.uic.edu/t3m/SnapPy/.

[18] D. B. A. Epstein, A. Marden, and V. Markovic, Quasiconformal homeomorphisms and the convex hull boundary, *Ann. of Math. (2)* **159** (2004), no. 1, 305–336.

[19] W. Floyd and A. Hatcher, Incompressible surfaces in punctured-torus bundles, *Topology Appl.* **13** (1982), no. 3, 263–282.

[20] D. Futer, E. Kalfagianni, and J. S. Purcell, Dehn filling, volume, and the Jones polynomial, *J. Differential Geom.* **78** (2008), no. 3, 429–464.

[21] ———, Cusp areas of Farey manifolds and applications to knot theory, *Internat. Math. Res. Notices IMRN* **2010** (2010), no. 23, 4434–4497.

[22] D. Gabai, Hyperbolic geometry and 3-manifold topology, *Low Dimensional Topology*, *IAS/Park City Math. Ser.*, vol. 15, Amer. Math. Soc., Providence, RI, 2009, pp. 73–103.

[23] F. Guéritaud, On canonical triangulations of once-punctured torus bundles and two-bridge link complements. With an appendix by D. Futer, *Geom. Topol.* **10** (2006), 1239–1284.

[24] A. Hatcher, On triangulations of surfaces, *Topology Appl.* **40** (1991), no. 2, 189–194.

[25] A. Hatcher and W. Thurston, Incompressible surfaces in 2-bridge knot complements, *Invent. Math.* **79** (1985), no. 2, 225–246.

[26] E. Irmak and J. D. McCarthy, Injective simplicial maps of the arc complex, *Turkish J. Math.* **34** (2010), no. 3, 339–354.

[27] M. Lackenby, Word hyperbolic Dehn surgery, *Invent. Math.* **140** (2000), no. 2, 243–282.

[28] A. Marden, *Outer Circles. An Introduction to Hyperbolic* 3*-Manifolds*, Cambridge University Press, Cambridge, 2007.

[29] Y. Minsky, The classification of Kleinian surface groups. I. Models and bounds, *Ann. of Math. (2)* **171** (2010), no. 1, 1–107.

[30] G. D. Mostow, *Strong Rigidity of Locally Symmetric Spaces*, *Ann. of Math. Stud.*, no. 78, Princeton University Press, Princeton, NJ, 1973.

[31] K. Ohshika, Kleinian groups which are limits of geometrically finite groups, *Mem. Amer. Math. Soc.* **177** (2005), no. 834, xii+116.

[32] J.-P. Otal, Le théorème d'hyperbolisation pour les variétés fibrées de dimension 3, *Astérisque* (1996), no. 235, x+159.

[33] G. Prasad, Strong rigidity of $\mathbf{Q}$-rank 1 lattices, *Invent. Math.* **21** (1973), 255–286.

[34] P. Przytycki and D. T. Wise, Mixed 3-manifolds are virtually special, preprint, http://arxiv.org/abs/1205.6742.

[35] K. Rafi and S. Schleimer, Covers and the curve complex, *Geom. Topol.* **13** (2009), no. 4, 2141–2162.

[36] R. Tang, The curve complex and covers via hyperbolic 3-manifolds, *Geom. Dedicata* **161** (2012), 233–237.

[37] W. P. Thurston, The geometry and topology of three-manifolds, Princeton Univ. Math. Dept. Notes, http://library.msri.org/books/gt3m/, 1980.

[38] ———, Three-dimensional manifolds, Kleinian groups and hyperbolic geometry, *Bull. Amer. Math. Soc. (N.S.)* **6** (1982), no. 3, 357–381.

[39] ———, On the geometry and dynamics of diffeomorphisms of surfaces, *Bull. Amer. Math. Soc. (N.S.)* **19** (1988), no. 2, 417–431.