GEODESIC PLANES IN THE CONVEX CORE OF AN ACYLINDRICAL 3-MANIFOLD

CURTIS T. MCMULLEN, AMIR MOHAMMADI, and HEE OH

Abstract

Let M be a convex cocompact, acylindrical hyperbolic 3-manifold of infinite volume, and let M^* denote the interior of the convex core of M. In this paper we show that any geodesic plane in M^* is either closed or dense. We also show that only countably many planes are closed. These are the first rigidity theorems for planes in convex cocompact 3-manifolds of infinite volume that depend only on the topology of M.

1. Introduction

In this paper we establish a new rigidity theorem for geodesic planes in acylindrical hyperbolic 3-manifolds.

Hyperbolic 3-manifolds. Let $M = \Gamma\backslash\mathbb{H}^3$ be a complete, oriented hyperbolic 3-manifold, presented as a quotient of hyperbolic space by the action of a discrete group

$$\Gamma \subset G = \mathrm{Isom}^+(\mathbb{H}^3).$$

Let $\Lambda \subset S^2 = \partial\mathbb{H}^3$ denote the limit set of Γ, and let $\Omega = S^2 - \Lambda$ denote the domain of discontinuity. The *convex core* of M is the smallest closed, convex subset of M containing all closed geodesics; equivalently,

$$\mathrm{core}(M) = \Gamma\backslash\,\mathrm{hull}(\Lambda) \subset M$$

is the quotient of the convex hull of the limit set Λ of Γ. Let M^* denote the interior of the convex core of M.

Geodesic planes in M^*. Let

$$f : \mathbb{H}^2 \to M$$

be a *geodesic plane*, i.e., a totally geodesic immersion of the hyperbolic plane into M. We often identify a geodesic plane with its image, $P = f(\mathbb{H}^2)$.

DUKE MATHEMATICAL JOURNAL
Vol. 171, No. 5, © 2022 DOI 10.1215/00127094-2021-0030
Received 9 June 2019. Revision received 28 January 2021.
First published online 17 March 2022.
2020 *Mathematics Subject Classification*. Primary 37A17; Secondary 57M50, 22E40.

By a geodesic plane $P^* \subset M^*$, we mean the nontrivial intersection

$$P^* = P \cap M^* \neq \emptyset$$

of a geodesic plane in M with the interior of the convex core. A plane P^* in M^* is always connected, and P^* is closed in M^* if and only if P^* is properly immersed in M^* (Section 2).

Acylindrical manifolds and rigidity. In this work, we study geodesic planes in M^* under the assumption that M is a convex cocompact, *acylindrical* hyperbolic 3-manifold. The acylindrical condition is a topological one; it means that the compact Kleinian manifold

$$\overline{M} = \Gamma \backslash (\mathbb{H}^3 \cup \Omega)$$

has incompressible boundary, and every essential cylinder in $\overline{M}$ is boundary parallel (Section 2). We will be primarily interested in the case where M is a convex cocompact manifold of infinite volume. Under this assumption, M is acylindrical if and only if Λ is a Sierpiński curve.[1]

Our main goal is to establish:

THEOREM 1.1
Let M be a convex cocompact, acylindrical, hyperbolic 3-manifold. Then any geodesic plane P^* in M^* is either closed or dense.

As a complement, we will show:

THEOREM 1.2
There are only countably many closed geodesic planes $P^* \subset M^*$.

We also establish the following topological equidistribution result:

THEOREM 1.3
If $P_i^* \subset M^*$ is an infinite sequence of distinct closed geodesic planes, then

$$\lim_{i \to \infty} P_i^* = M^*$$

in the Hausdorff topology on closed subsets of M^*.

[1] A compact set $\Lambda \subset S^2$ is a *Sierpiński curve* if $S^2 - \Lambda = \bigcup D_i$ is a dense union of Jordan disks with disjoint closures, and $\mathrm{diam}(D_i) \to 0$. Any two Sierpiński curves are homeomorphic (see [12]).

Remarks

1. We do not know of any instance of Theorem 1.1 where P^* is closed in M^* but P is not closed in M.
   Added in proof. An example of such an *exotic plane* in an acylindrical manifold has recently been constructed by Zhang. In his example, the closure of P is not even locally connected near ∂M^* (see [13]).
   Thus the rigidity of planes described in Theorem 1.1 does *not* extend beyond the convex core of M.
2. In the special case where M is compact (so $M = M^*$), Theorem 1.1 is due independently to Shah and Ratner (see [8], [9]).
3. For a general convex cocompact manifold M, there can be uncountably many distinct closed planes in M^*; see the end of Section 2.
4. Examples of acylindrical manifolds such that M^* contains infinitely many closed geodesic planes are given in [6, Corollary 11.5].
5. The study of planes P that do not meet M^* can be reduced to the case where M is a quasifuchsian manifold. This case can be analyzed via the bending lamination (cf. Section 6).

Comparison to the case of geodesic boundary. A convex cocompact hyperbolic 3-manifold M such that $\partial\,\mathrm{core}(M)$ is totally geodesic is automatically acylindrical. For these *rigid* acylindrical manifolds, the results above were obtained in our previous work [6]. While one would ultimately like to analyze planes in a large class of geometrically finite groups, our previous results covered only countably many examples (by Mostow rigidity).

The present paper makes a major step forward in this program, by developing a new argument for unipotent recurrence which works *without* geodesic boundary, which is robust enough to be invariant under quasi-isometry, and which is powerful enough to apply to the class of all convex cocompact acylindrical manifolds. The key insight is that one should work with a proper subset of the renormalized frame bundle, defined in terms of thickness of Cantor sets, where we show sufficient recurrence takes place in the acylindrical case.

The cylindrical case. The acylindrical setting is also close to optimal, since Theorem 1.1 is generally false for cylindrical manifolds.

For example, consider a quasi-Fuchsian group Γ containing a Fuchsian subgroup Γ' of the second kind with limit set $\Lambda' \subset S^1$. Given $(a,b) \in \Lambda' \times \Lambda'$, let C_{ab} denote the unique circle orthogonal to S^1 such that $C_{ab} \cap S^1 = \{a,b\}$. It is possible to choose Γ such that $C_{ab} \cap \Lambda = \{a,b\}$ for uncountably many (a,b); and further, to arrange that the corresponding hyperbolic planes $P \subset M$ and $P^* \subset M^*$ have wild closures, violating Theorem 1.1 (cf. [6, Appendix A]).

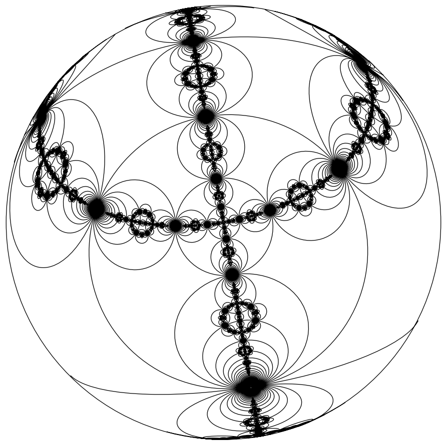

Figure 1. Limit set of a cylindrical 3-manifold.

$G = \mathrm{PSL}_2(\mathbb{C}) \cong \mathrm{Isom}^+(\mathbb{H}^3)$
$H = \mathrm{PSL}_2(\mathbb{R}) \cong \mathrm{Isom}^+(\mathbb{H}^2)$
$K = \mathrm{SU}(2)/(\pm I)$
$A = \left\{\left(\begin{smallmatrix} a & 0 \\ 0 & a^{-1} \end{smallmatrix}\right) : a > 0\right\}$
$N = \left\{n_s = \left(\begin{smallmatrix} 1 & s \\ 0 & 1 \end{smallmatrix}\right) : s \in \mathbb{C}\right\}$
$U = \{n_s : s \in \mathbb{R}\}$
$V = \{n_s : s \in i\mathbb{R}\}$
$\mathrm{F}\mathbb{H}^3 = G = \{\text{the frame bundle of } \mathbb{H}^3\}$
$\mathbb{H}^3 = G/K$
$S^2 = G/AN = \partial\mathbb{H}^3$
$\mathcal{C} = G/H = \{\text{the space of oriented circles } C \subset S^2\}$

Table 2. Notation for G and some of its subgroups and homogeneous spaces.

The same type of example can be embedded in more complicated 3-manifolds with nontrivial characteristic submanifold; an example is shown in Figure 1.

Homogeneous dynamics. Next we formulate a result in the language of Lie groups and homogeneous spaces, Theorem 1.4, that strengthens both Theorems 1.1 and 1.3.

To set the stage, we have summarized our notation for G and its subgroups in Table 2. We have similarly summarized the spaces attached to an arbitrary hyperbolic 3-manifold $M = \Gamma\backslash\mathbb{H}^3$ in Table 3. (In the definition of $\mathcal{C}^*$, a circle $C \subset S^2$ *separates* Λ if the limit set meets both components of $S^2 - C$.)

$M = \Gamma\backslash\mathbb{H}^3 =$ (the quotient hyperbolic 3-manifold)
$\overline{M} = \Gamma\backslash(\mathbb{H}^3 \cup \Omega)$
$\mathrm{core}(M) = \Gamma\backslash\,\mathrm{hull}(\Lambda)$
$M^* = \mathrm{int}\big(\mathrm{core}(M)\big)$
$\mathrm{F}\,M = \Gamma\backslash G =$ (the frame bundle of M)
$F^* = \{x \in \mathrm{F}\,M : x$ is tangent to a plane P that meets $M^*\}$
$\mathcal{C}^* = \{C \in \mathcal{C} : C$ separates $\Lambda\}$

Table 3. Spaces associated to $M = \Gamma\backslash\mathbb{H}^3$.

Circles, frames, and planes. Circles, frames, and planes are closely related. In fact, if $\mathcal{P}$ denotes the set of all (oriented) planes in M, then we have the natural identifications:

$$\mathcal{P} = \Gamma\backslash\mathcal{C} = \mathrm{F}\,M/H. \tag{1.1}$$

Indeed, all three spaces can be identified with $\Gamma\backslash G/H$. We will frequently use these identifications to go back and forth between circles, frames, and planes.

When M^* is nonempty (equivalently, when Γ is Zariski-dense in G), the spaces $\mathcal{C}^*$ and F^* correspond to the set of planes $\mathcal{P}^*$ that meet M^*. In other words, we have

$$\mathcal{P}^* = \Gamma\backslash\mathcal{C}^* = F^*/H. \tag{1.2}$$

To go from a circle to a plane, let P be the image of $\mathrm{hull}(C) \subset \mathbb{H}^3$ under the covering map from $\mathbb{H}^3$ to M. To go from a frame $x \in \mathrm{F}\,M$ to a plane, take the image of xH under the natural projection $\mathrm{F}\,M \to M$.

When Λ is connected and consists of more than one point (e.g., when M is acylindrical), it is easy to see that

$$\overline{\mathcal{C}^*} = \{C \in \mathcal{C} : C \text{ meets } \Lambda\}.$$

Thus the closures of the dense sets arising in Theorem 1.4 below are quite explicit.

The closed or dense dichotomy. We can now state our main result from the perspective of homogeneous dynamics.

THEOREM 1.4
Let $M = \Gamma\backslash\mathbb{H}^3$ be a convex cocompact, acylindrical 3-manifold. Then any Γ-invariant subset of $\mathcal{C}^*$ is either closed or dense in $\mathcal{C}^*$. Equivalently, any H-invariant subset of F^* is either closed or dense in F^*.

(The equivalence is immediate from equation (1.2).)

This result sharpens Theorem 1.1 to give the following dichotomy on the level of the tangent bundles:

COROLLARY 1.5
The normal bundle to a geodesic plane $P^* \subset M^*$ is either closed or dense in the tangent bundle $\mathrm{T}\,M^*$.

Beyond the acylindrical case. This paper also establishes several results that apply outside the acylindrical setting. For example, Theorems 2.1, 4.1, 5.1, and 6.1 only require the assumption that M has incompressible boundary. In fact, the main argument pivots on a result relating Cantor sets and Sierpiński curves, Theorem 3.4, that involves no groups at all.

Discussion of the proofs. We conclude with a sketch of the proofs of Theorems 1.1 through Theorem 1.4.

Let $M = \Gamma\backslash\mathbb{H}^3$ be a convex cocompact acylindrical 3-manifold of infinite volume, with limit set Λ and domain of discontinuity Ω. The horocycle and geodesic flows on the frame bundle $\mathrm{F}\,M = \Gamma\backslash G$ are given by the right actions of U and A, respectively. The *renormalized frame bundle* of M is the compact set defined by

$$\mathrm{RF}\,M = \{x \in \mathrm{F}\,M : xA \text{ is bounded}\}. \tag{1.3}$$

In Section 2 we prove Theorem 1.2 by showing that the fundamental group of any closed plane $P^* \subset M^*$ contains a free group on two generators. We also show that Theorems 1.1 and 1.3 follow from Theorem 1.4. The remaining sections develop the proof of Theorem 1.4.

In Section 3 we show that Λ is a Sierpiński curve of positive modulus. This means there exists a $\delta > 0$ such that the modulus of the annulus between any two components D_1, D_2 of $S^2 - \Lambda$ satisfies

$$\operatorname{mod}\big(S^2 - (\overline{D}_1 \cup \overline{D}_2)\big) \geq \delta > 0.$$

We also show that if Λ is a Sierpiński curve of positive modulus, then there exists a $\delta > 0$ such that $C \cap \Lambda$ contains a Cantor set K of modulus δ, whenever C separates Λ. This means that for any disjoint components I_1 and I_2 of $C - K$, we have

$$\operatorname{mod}\big(S^2 - (\overline{I}_1 \cup \overline{I}_2)\big) \geq \delta > 0.$$

This result does not involve Kleinian groups and may be of interest in its own right.

In Section 4 we use this uniform bound on the modulus of a Cantor set to construct a compact, A-invariant set

$$\mathrm{RF}_k\, M \subset \mathrm{RF}\, M$$

with good recurrence properties for the horocycle flow on $\mathrm{F}M$. We also show that when k is sufficiently large, $\mathrm{RF}_k\, M$ meets every H-orbit in F^*.

The introduction of $\mathrm{RF}_k\, M$ is one of the central innovations of this paper that allows us to handle acylindrical manifolds with quasi-Fuchsian boundary. When M is a *rigid* acylindrical manifold, $\mathrm{RF}_k\, M = \mathrm{RF}\, M$ for all k sufficiently large, so in some sense $\mathrm{RF}_k\, M$ is a substitute for the renormalized frame bundle. For a more detailed discussion, see the end of Section 4.

In Section 5 we shift our focus to the boundary of the convex core. Using the theory of the bending lamination, we give a precise description of $C \cap \Lambda$ in the case where C comes from a supporting hyperplane for the limit set.

In Sections 6 and 7, we formulate two density theorems for hyperbolic 3-manifolds M with incompressible boundary. These results do not require that M is acylindrical. Each section gives a criterion for a sequence of circles $C_n \in \mathcal{C}^*$ to have the property that $\bigcup \Gamma C_n$ is dense in $\mathcal{C}^*$.

In Section 6 we show that density holds if $C_n \to C \notin \mathcal{C}^*$ and $\lim(C_n \cap \Lambda)$ is uncountable. The proof relies on the analysis of the convex hull given in Section 5.

In Section 7 we show that density holds if $C_n \to C \in \mathcal{C}^*$ and $C \notin \bigcup \Gamma C_n$, provided $C \cap \Lambda$ contains a Cantor set of positive modulus. The proof uses recurrence, minimal sets, and homogeneous dynamics on the frame bundle, and follows a similar argument in [6]. It also relies on the density result of Section 6.

When M is acylindrical, the Cantor set condition is automatic by Section 3. Thus Theorem 1.4 follows immediately from the density theorem of Section 7.

Question

We conclude by mentioning an open problem that goes beyond the acylindrical case. Let $P^* \subset M^*$ be a plane in a quasi-Fuchsian manifold, and suppose the corresponding circle satisfies $|C \cap \Lambda| > 2$. Does it follow that P^* is closed or dense in M^*?

2. Planes in acylindrical manifolds

In this section we will prove Theorem 1.2, and show that our other main results, Theorems 1.1 and 1.3, follow from Theorem 1.4 on the homogeneous dynamics of H acting on F^*.

Let $M = \Gamma\backslash\mathbb{H}^3$ be a convex cocompact hyperbolic 3-manifold. We first describe how the topology of $\overline{M}$ influences the shape of planes in M^*. Here are the two main results:

THEOREM 2.1
If $\overline{M}$ has incompressible boundary, then the fundamental group of any closed plane $P^* \subset M^*$ is nontrivial.

THEOREM 2.2
If $\overline{M}$ is acylindrical, then the fundamental group of any closed plane $P^* \subset M^*$ contains a free group on two generators.

The second result immediately implies Theorem 1.2, which we restate as follows:

COROLLARY 2.3
If $\overline{M}$ is acylindrical, then there are at most countably many closed planes $P^* \subset M^*$.

Proof
In this case P^* corresponds to a circle C whose stabilizer Γ^C (as discussed below) is isomorphic to the fundamental group of P^*, and contains a free group on two generators $\langle a, b\rangle$. Since C is the unique circle containing the limit set of $\langle a, b\rangle \subset \Gamma$, and there are only countably many possibilities for (a, b), there are only countable possibilities for P^*. □

In the remainder of this section, we first develop general results about planes in 3-manifolds, and prove Theorems 2.1 and 2.2. Then we derive Theorems 1.1 and 1.3 from Theorem 1.4. Finally we show by example that a cylindrical manifold can have uncountably many closed planes $P^* \subset M^*$.

Topology of 3-manifolds. We begin with some topological definitions.

Let D^2 denote a closed 2-disk, and let $C^2 \cong S^1 \times [0, 1]$ denote a closed cylinder. Let N be a compact 3-manifold with boundary. We say N has *incompressible boundary* if every continuous map

$$f : (D^2, \partial D^2) \to (N, \partial N)$$

can be deformed, as a map of pairs, so its image lies in ∂N. (This property is automatic if $\partial N = \emptyset$.)

Similarly, N is *acylindrical* if it has incompressible boundary and every continuous map

$$f : (C^2, \partial C^2) \to (N, \partial N),$$

injective on π_1, can be deformed into ∂N. That is, every incompressible disk or cylinder in N is boundary parallel.

When $N = \overline{M} = \Gamma\backslash(\mathbb{H}^3 \cup \Omega)$ is a compact Kleinian manifold, these properties are visible on the sphere at infinity: the limit set Λ of Γ is connected if and only if $\overline{M}$ has incompressible boundary, and $\overline{M}$ is acylindrical if and only if Λ is a Sierpiński curve or $\Lambda = S^2$.

For more on the topology of hyperbolic 3-manifolds, see, for example, [7], [11], and [5].

Topology of planes. Next we discuss the fundamental group of a plane $P \subset M$, and the corresponding plane $P^* \subset M^*$. These definitions apply to an arbitrary hyperbolic 3-manifold.

For precision, it is useful to think of a plane P as being specified by an *oriented* circle $C \subset S^2$, whose convex hull covers P. More precisely, the plane attached to C is given by the map

$$\widetilde{f} : \mathrm{hull}(C) \cong \mathbb{H}^2 \subset \mathbb{H}^3 \to M = \Gamma\backslash\mathbb{H}^3$$

with image $\widetilde{f}(\mathbb{H}^2) = P$. The stabilizer of the circle C in G is a conjugate xHx^{-1} of $H = \mathrm{PSL}_2(\mathbb{R})$; hence its stabilizer in Γ is given by

$$\Gamma^C = \Gamma \cap xHx^{-1}.$$

Let

$$S = \Gamma^C \backslash \mathrm{hull}(C).$$

Then the map $\widetilde{f}$ descends to give an immersion

$$f : S \to M$$

with image P. The immersion f is generically injective if P is orientable; otherwise, it is generically two-to-one (and there is an element in Γ that reverses the orientation of C).

We refer to

$$\pi_1(S) \cong \Gamma^C$$

as the *fundamental group of* P (keeping in mind caveats about orientability).

Planes in the convex core. Now suppose $P^* = P \cap M^*$ is nonempty. In this case

$$S^* = f^{-1}(M^*)$$

is a nonempty convex subsurface of S, with $\pi_1(S^*) = \pi_1(S)$. The map

$$f : S^* \to P^* \subset M^*$$

presents S^* as the (orientable) *normalization* of P^*, i.e., as the smooth surface obtained by resolving the self-intersections of P^*. Similarly, the frame bundle of P with its branches separated is given by

$$\mathrm{F}P = xH \subset \mathrm{F}M$$

for some $x \in F^*$. (One should consistently orient C and P to define $\mathrm{F}P$.)

To elucidate the connections between these objects, we formulate:

PROPOSITION 2.4
Let M be an arbitrary hyperbolic 3-manifold. Suppose $C \in \mathcal{C}^*$ and $x \in F^*$ correspond to the same plane $P^* \subset M^*$. Then the following are equivalent:
1. *ΓC is closed in $\mathcal{C}^*$.*
2. *The inclusion $\Gamma C \subset \mathcal{C}^*$ is proper.*
3. *xH is closed in F^*.*
4. *P^* is closed in M^*.*
5. *The normalization map $f : S^* \to P^*$ is proper.*

In (2) above, ΓC is given the discrete topology.

Proof
If ΓC is not discrete in $\mathcal{C}^*$, then by homogeneity it is perfect (it has no isolated points). But a closed perfect set is uncountable, so ΓC is not closed. Thus (1) implies that $\Gamma C \subset \mathcal{C}^*$ is closed and discrete, which implies (2); and clearly (2) implies (1). The remaining equivalences are similar, using equation (1.2) to relate $\mathcal{P}^*$, $\mathcal{C}^*$ and F^*. □

Compact deformations. In the context of proper mappings, the notion of a compact deformation is also useful.

Let $f_0 : X \to Y$ be a continuous map. We say $f_1 : X \to Y$ is a *compact deformation* of f_0 if there is a continuous family of maps $f_t : X \to Y$ interpolating between them, defined for all $t \in [0, 1]$, and a compact set $X_0 \subset X$ such that $f_t(x) = f_0(x)$ for all $x \notin X_0$.

Let $P^* \subset M^*$ be a hyperbolic plane with normalization $f_0 : S^* \to M^*$. We say $Q^* \subset M^*$ is a *compact deformation* of P^* if it is the image of S^* under a compact deformation f_1 of f_0.

THEOREM 2.5
Let $M = \Gamma\backslash\mathbb{H}^3$ be an arbitrary 3-manifold, and let $K \subset M^*$ be a submanifold such that the induced map

$$\pi_1(K) \to \pi_1(M)$$

is surjective. Then K meets every geodesic plane $P^* \subset M^*$ and every compact deformation Q^* of P^*.

COROLLARY 2.6
If $\pi_1(M)$ is finitely generated, then there is a compact submanifold $K \subset M^*$ that meets every plane $P^* \subset M^*$.

Proof
Provided M^* is nonempty, $\pi_1(M^*)$ is isomorphic to $\pi_1(M)$; and since the latter group is finitely generated, there is a compact submanifold $K \subset M^*$ (say, a neighborhood of a bouquet of circles) whose fundamental group surjects onto $\pi_1(M^*)$. □

Proof of Theorem 2.5
We will use the fact that S^0 and S^1 can link in S^2.

Let P^* be a plane in M^*, arising from a circle $C \subset S^2$ with an associated map $f : S \to P$ as above. Since P meets M^*, there are points in the limit set of Γ on both sides of C. Since the endpoints of closed geodesics are dense in $\Lambda \times \Lambda$ (cf. [1]), we can find a hyperbolic element $g \in \Gamma$ such that its two fixed points

$$\mathrm{Fix}(g) = \{a_1, a_2\} \subset S^2$$

are separated by C, and the convex hull of $\{a_1, a_2\}$ in $\mathbb{H}^3$ projects to a closed geodesic $\delta \subset M$. Note that $\mathrm{Fix}(g) \cong S^0$ and $C \cong S^1$ are linked in S^2.

Since $\pi_1(K)$ maps onto $\pi_1(M)$, the loop δ is freely homotopic to a loop $\gamma \subset K$.

Let $f_0 = f|S^*$. Suppose $f_0 : S^* \to M^*$ has a compact deformation f_1 with image Q^* disjoint from K, and hence disjoint from γ. Extend this deformation trivially to the rest of S, to obtain a compact deformation f_1 of the geodesic immersion $f : S \to P$. Then $f_1(S)$ is disjoint from γ. Lifting f_1 to the universal cover of S, we obtain a continuous map

$$\widetilde{f_1} : \mathrm{hull}(C) \to \mathbb{H}^3$$

that is a bounded distance from the identity map. In particular, its image is a disk D spanning C.

Similarly, a suitable lift of γ gives a path $\widetilde{\gamma} \subset \mathbb{H}^3$, disjoint from D, that joins a_1 to a_2. This contradicts the fact that C separates a_1 from a_2 in S^2. □

We can now proceed to the:

Proof of Theorem 2.1 (The incompressible case)
For the beginning of the argument, we only use the fact that $\overline{M}$ is compact and M^* is nonempty. Using the nearest point projection, it is straightforward to show that core(M) is homeomorphic to $\overline{M}$. Thus its interior M^* deformation retracts onto a compact submanifold $K \subset M^*$, homeomorphic to $\overline{M}$, such that the inclusion is a homotopy equivalence; in particular, $\pi_1(K) \cong \pi_1(\mathrm{M}^*)$.

Consider a closed plane $P^* \subset M^*$, arising as the image of a proper map $f : S^* \to P^*$ as above. We can also arrange that K is transverse to f, so its preimage

$$S_0 = f^{-1}(K) \subset S^*$$

is a compact, smoothly bounded region in S^*. (However, S_0 need not be connected.)

We claim that, after changing f by a compact deformation, we can arrange that the inclusion of each component of S_0 into S^* is injective on π_1. This is a standard argument in 3-dimensional topology. If the inclusion is not injective on π_1, then there is a compact disk $D \subset S^*$ with $D \cap S_0 = \partial D$. The map f sends $(D, \partial D)$ into (M^*, K). Since K is a deformation retract of M^*, $f|D$ can be deformed until it maps D into K, while keeping $f|\partial D$ fixed. Then D becomes part of S_0. This deformation is compact because D is compact. Since ∂S_0 has only finitely many components, only finitely many disks of this type arise, so after finitely many compact deformations of f, the inclusion $S_0 \subset S^*$ becomes injective on π_1.

Now we use the assumption that $K \cong \overline{M}$ has incompressible boundary. Suppose that $\pi_1(S^*)$ is trivial. Then π_1 is trivial for each component of S_0, and hence each component of S_0 is a disk. By construction the deformed map f restricts to give a map of pairs

$$f : (S_0, \partial S_0) \to (K, \partial K).$$

Since K has incompressible boundary, we can further deform $f|S_0$ so it sends the whole surface S_0 into ∂K. Then the image Q^* of f gives a compact deformation of P^* that is disjoint from $K^* = K - \partial K$. But $\pi_1(K^*)$ maps onto $\pi_1(M)$, contradicting Theorem 2.5. Thus $\pi_1(S^*)$ is nontrivial. □

Proof of Theorem 2.2 (The acylindrical case)
The proof follows the same lines as the incompressible case. If $\pi_1(S^*)$ does not contain a free group on two generators, then S^* is a disk or an annulus. After a compact deformation, we can assume that the inclusion $S_0 = f^{-1}(K) \subset S^*$ is injective on π_1. Thus each component of S_0 is also a disk or an annulus. Since K is acylindrical, after a further compact deformation of f we can arrange that $f(S_0) \subset \partial K$, leading to a contradiction. □

Rigidity of planes from homogeneous dynamics. Now suppose $M = \Gamma\backslash\mathbb{H}^3$ is a convex cocompact, acylindrical 3-manifold. Assume we know Theorem 1.4, which states that under this hypothesis:

Any Γ–invariant set $E \subset \mathcal{C}^*$ is closed or dense in $\mathcal{C}^*$.

We can then prove the other two main results stated in the introduction.

Proof of Theorem 1.1
Let P^* be a geodesic plane in M^*, and let $E = \Gamma C$ be the corresponding set of circles. Then by Theorem 1.4, E is either closed or dense in $\mathcal{C}^*$, and hence P^* is either closed or dense in M^*. □

Proof of Theorem 1.3
Let P_i^* be a sequence of distinct closed planes in M^*. We wish to show that $\lim P_i^* = M^*$ in the Hausdorff topology on closed subsets of M^*. To see this, first pass to a subsequence so that P_i^* converges to $Q^* \subset M^*$. It suffices to show that $Q^* = M^*$ for every such subsequence. Since each P_i^* is nowhere dense, to show that $Q^* = M^*$ and complete the proof, it suffices to show that $\bigcup P_i^*$ is dense in M^*.

Let $E_i \subset \mathcal{C}^*$ be the Γ-orbit corresponding to P_i, and let $E = \bigcup E_i$. Since the planes P_i are distinct, the sets E_i are disjoint. By Corollary 2.6, there exists a compact set $K \subset M^*$ that meets every P_i^*, so there exists a compact set $K' \subset \mathcal{C}^*$ meeting every E_i. Thus we can choose $C_i \in E_i \cap K'$ and pass to a subsequence such that

$$C_i \to C_\infty \in K' \subset \mathcal{C}^*$$

and $C_\infty \notin E$. (If $C_\infty \in E_i = \Gamma C_i$, just drop that term from the sequence.) Since E is not closed in $\mathcal{C}^*$, it is dense in $\mathcal{C}^*$ by Theorem 1.4. Consequently $\bigcup P_i^*$ is dense in M^*, as desired. □

Example: Uncountably many geodesic cylinders. To conclude, we show that Theorem 2.2 and Corollary 2.3 do not hold for general convex cocompact manifolds with incompressible boundary.

In fact, in such a manifold one can have uncountably many distinct closed planes $P^* \subset M^*$, each with cyclic fundamental group. For a concrete example of this phenomenon, consider a closed geodesic γ and the corresponding plane P in the quasifuchsian manifold $M = M_\theta$ discussed in [6, Corollary A.2]. In this construction, γ is a simple curve in the boundary of the convex core of M, and $P \cong \gamma \times \mathbb{R}$ is a hyperbolic cylinder properly embedded in M. Consequently, $P^* \subset M^*$ is a properly immersed cylinder in M^*. By varying the angle that P meets the boundary of $\mathrm{core}(M_\theta)$ along γ, we obtain a continuous family of properly immersed planes in M^*.

3. Moduli of Cantor sets and Sierpiński curves

The rest of the paper is devoted to the proof of Theorem 1.4.

In this section we define the modulus of a Cantor set $K \subset S^1$ (or in any circle $C \subset S^2$), as well as the modulus of a Sierpiński curve $K \subset S^2$. We then prove:

THEOREM 3.1
Let Λ be the limit set of Γ, where $M = \Gamma\backslash\mathbb{H}^3$ is a convex cocompact acylindrical 3-manifold of infinite volume. Then there exists a $\delta > 0$ such that:
1. *Λ is a Sierpiński curve of modulus δ, and*
2. *$C \cap \Lambda$ contains a Cantor set of modulus δ, whenever the circle $C \subset S^2$ separates Λ.*

The modulus of a Sierpiński curve. For background on conformal invariants and quasiconformal maps, see [4].

We begin with some definitions. An *annulus* $A \subset S^2$ is an open region whose complement consists of two components. Provided neither component is a single point, A is conformally equivalent to a unique round annulus of the form

$$A_R = \{z \in \mathbb{C} : 1 < |z| < R\},$$

and its *modulus* is defined by

$$\operatorname{mod}(A) = \frac{\log R}{2\pi}.$$

(More geometrically, A is conformally equivalent to a Euclidean cylinder of radius 1 and height $\operatorname{mod}(A)$.) Since the modulus is a conformal invariant, we have

$$\operatorname{mod}(A) = \operatorname{mod}\big(g(A)\big) \quad \forall g \in G. \tag{3.1}$$

Recall that a compact set $\Lambda \subset S^2$ is a *Sierpiński curve* if its complement

$$S^2 - \Lambda = \bigcup D_i$$

is a dense union of Jordan disks D_i with disjoint closures, whose diameters tend to zero. We say Λ has *modulus* δ if

$$\inf_{i \neq j} \operatorname{mod}\big(S^2 - (\overline{D}_i \cup \overline{D}_j)\big) \geq \delta > 0.$$

The modulus of an annulus $A \subset S^1$. Let $C \subset S^2$ be a circle and let $A \subset C$ be an *annulus on C*, meaning an open set such that $C - A = I_1 \cup I_2$ is the union of two disjoint intervals (circular arcs). We extend the notion of modulus to this 1-dimensional situation by defining

$$\operatorname{mod}(A, C) = \operatorname{mod}\big(S^2 - (I_1 \cup I_2)\big).$$

Clearly $\operatorname{mod}(gA, gC) = \operatorname{mod}(A, C)$ for all $g \in G$, and consequently $\operatorname{mod}(A, C)$ depends only on the cross-ratio of the four endpoints of A. The cross-ratio is controlled by the lengths of the components A_1, A_2 of A and the components I_1, I_2 of $C - A$. From this observation and monotonicity of the modulus [4, I.6.6] it is easy to show:

PROPOSITION 3.2
There are increasing continuous functions $\delta(t), \Delta(t) > 0$ such that

$$\delta(t) < \operatorname{mod}(A, C) < \Delta(t),$$

where t is the ratio of lengths

$$t = \frac{\min(|A_1|, |A_2|)}{\min(|I_1|, |I_2|)}.$$

The same result holds with t replaced by $d(\operatorname{hull}(I_1), \operatorname{hull}(I_2))$.

For later reference we recall the following result due to Teichmüller [4, Chapter II, Theorem 1.1].

PROPOSITION 3.3
Let I_1 and I_2 be the two components of $C - A$. Then

$$\operatorname{mod}(B) \le \operatorname{mod}(A, C)$$

for any annulus $B \subset S^2$ separating the endpoints of I_1 from those of I_2.

The modulus of a Cantor set. Let $K \subset C \subset S^2$ be a compact subset of a circle, such that its complement

$$C - K = \bigcup I_i$$

is a union of open intervals with disjoint closures. Note that C is uniquely determined by K (and we allow $K = C$). We say K has *modulus* δ if we have

$$\inf_{i \neq j} \operatorname{mod}(A_{ij}, C) \ge \delta > 0, \tag{3.2}$$

where $A_{ij} = C - \overline{I_i \cup I_j}$. We will be primarily interested in the case where K is a *Cantor set*, meaning $\bigcup I_i$ is dense in C.

Slices. Next we show that circular slices of a Sierpiński curve inherit positivity of the modulus. This argument makes no reference to 3-manifolds.

THEOREM 3.4
Let $\Lambda \subset S^2$ be a Sierpiński curve of modulus $\delta > 0$. Then there exists a $\delta' > 0$ such that $C \cap \Lambda$ contains a Cantor set K of modulus δ' whenever C is a circle separating Λ.

Proof
Let $S^2 - \Lambda = \bigcup D_i$ express the complement of Λ as a union of disjoint disks. Each disk D_i meets the circle C in a collection of disjoint open intervals (see Figure 4). The proof will be based on a study of the interaction of intervals from different components.

Let $U = C - \Lambda = \bigcup U_i$, where

$$U_i = C \cap D_i.$$

Note that distinct U_i have disjoint closures, and $\operatorname{diam} U_i \to 0$, since these two properties hold for the disks D_i. The open set U_i may be empty.

We may assume U is dense in C, since otherwise we can just choose a suitable Cantor set $K \subset C - \overline{U}$. On the other hand, no U_i is dense in C; if it were, we would have $C \subset \overline{D}_i$, contrary to our assumption that C separates Λ. It follows that U_i is nonempty for infinitely many values of i.

Let us say an open interval $I = (a, b) \subset C$, with distinct endpoints, is a *bridge of type i* if $a, b \in \partial U_i$. Note that an ascending union of bridges of type i is again a bridge of type i, provided its endpoints are distinct.

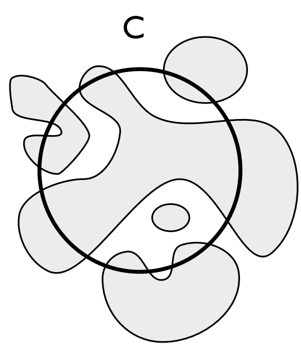

Figure 4. A circle C and some components D_i of $S^2 - \Lambda$.

Our goal is to construct a sequence of disjoint bridges $I_1, I_2, I_3, \dots$ such that $|I_1| \geq |I_2| \geq \cdots$ and $K = C - \bigcup I_i$ is a Cantor set of modulus δ'.

To start the construction, choose any bridge $I_1 \subset C$. After changing coordinates by a Möbius transformation $g \in G^C$, we can assume that I_1 fills at least half the circle; i.e., $|I_1| > |C|/2$. This will ensure that $|I_1| \geq |I_k|$ for all $k > 1$.

Next, let I_2 be a bridge of maximal length among all those which are disjoint from I_1 and of a different type from I_1. Such a bridge exists because $\operatorname{diam}(U_i) \to 0$, so only finitely many types of bridges are competing to be I_2. To complete the initial step, enlarge I_1 to a maximal interval of the same type, disjoint from I_2.

Proceeding inductively, let $I_{k+1} \subset C$ be a bridge of maximum length among all bridges disjoint from $I_1, \dots, I_k$. Since I_1 is a maximal bridge of its type among those disjoint from I_2, and vice-versa, the intervals (I_1, I_2, I_k) are of three distinct types, for all $k \geq 3$. Consequently, $|I_2| \geq |I_k|$ for all $k > 2$.

Note that the bridges so constructed have disjoint closures. Indeed, if I_i and I_j were to have an endpoint a in common, with $i < j$, then $I_i \cup \{a\} \cup I_j$ would be a longer interval of the same type as I_i, contradicting to stage i of the construction.

Since U is dense in C, it follows that at any finite stage there is a bridge disjoint from all those chosen so far, and thus the inductive construction continues indefinitely. By construction, we have

$$|I_1| \geq |I_2| \geq |I_3| \cdots$$

and by disjointness, $|I_k| \to 0$. Moreover, $\bigcup I_k$ is dense in C. Otherwise, by density of U, we would be able to find a bridge J disjoint from all I_k, and longer than I_k for all k sufficiently large, contradicting the construction of I_k.

Let $K = C - \bigcup_1^\infty I_k$. Since the intervals I_k have disjoint closures, and their union is dense in C, K is a Cantor set. We have $K \subset \Lambda$ since $\partial I_k \subset \Lambda$ for all k.

Now consider any two indices $i < j$. Let

$$A = C - (\overline{I}_i \cup \overline{I}_j) = A_1 \cup A_2,$$

where the open intervals A_1 and A_2 are disjoint. If the bridges I_i and I_j have types $s \neq t$ respectively, then the annulus

$$B = S^2 - (\overline{D}_s \cup \overline{D}_t)$$

separates ∂I_i from ∂I_j, and hence

$$\operatorname{mod}(A, C) \geq \operatorname{mod}(B) \geq \delta > 0$$

by Proposition 3.3.

On the other hand, if I_i and I_j have the same type s, then $i, j > 2$, and there must be a bridge I_k, $k < i$, such that $I_1 \cup I_k$ separates I_i from I_j. Otherwise, we could have combined I_i and I_j to obtain a longer bridge at step i.

It follows that

$$t = \frac{\min(|A_1|, |A_2|)}{\min(|I_i|, |I_j|)} \geq \frac{\min(|I_1|, |I_k|)}{\min(|I_i|, |I_j|)} = \frac{|I_k|}{|I_j|} \geq 1,$$

since $k < i < j$. By Proposition 3.2, this implies that

$$\operatorname{mod}(A, C) > \delta_0 > 0$$

where δ_0 is a universal constant. Thus the Theorem holds with $\delta' = \min(\delta_0, \delta)$. □

Limit sets. We can now complete the proof of Theorem 3.1.

THEOREM 3.5
Let $M = \Gamma\backslash\mathbb{H}^3$ be a convex cocompact acylindrical 3-manifold of infinite volume. Then its limit set Λ is a Sierpiński curve of modulus δ for some $\delta > 0$.

Proof
First suppose that every component of $\Omega = S^2 - \Lambda = \bigcup D_i$ is a round disk; i.e., suppose that M is a rigid acylindrical manifold. By compactness, there exists an $L > 0$ such that the hyperbolic length of any geodesic arc $\gamma \subset \operatorname{core}(M)$ orthogonal to the boundary at its endpoints is greater than L. Consequently, $d_{ij} = d(\operatorname{hull}(D_i), \operatorname{hull}(D_j)) \geq L$ for any $i \neq j$. Since the modulus of $S^2 - (\overline{D}_i \cup \overline{D}_j)$ is given by $d_{ij}/(2\pi)$, Λ is a Sierpiński curve of modulus $\delta = L/(2\pi) > 0$.

To treat the general case, recall that for any convex cocompact acylindrical manifold M, there exists a rigid acylindrical manifold $M' = \Gamma'\backslash\mathbb{H}^3$ such that Γ' is K-quasiconformally conjugate to Γ. Since a K-quasiconformal map distorts the modulus of an annulus by at most a factor of K, and the limit set Λ' of Γ' is a Sierpiński curve with modulus $\delta' > 0$, Λ itself is a Sierpiński curve of modulus $\delta = \delta'/K > 0$. □

Proof of Theorem 3.1
Combine Theorems 3.4 and 3.5. □

4. Recurrence of horocycles

Let $M = \Gamma\backslash\mathbb{H}^3$ be an arbitrary 3-manifold. In this section we will define, for each $k > 1$, a closed, A-invariant set

$$\operatorname{RF}_k M \subset \operatorname{RF} M$$

consisting of points with good recurrence properties under the horocycle flow generated by U (for terminology, see Tables 2 and 3). We will then show the following.

THEOREM 4.1
Let $M = \Gamma\backslash\mathbb{H}^3$ be a convex cocompact acylindrical 3-manifold. We then have

$$F^* \subset (\mathrm{RF}_k\, M)H$$

for all k sufficiently large. More precisely, every plane $P^* \subset M^*$ is tangent to a frame in $\mathrm{RF}_k\, M$.

We conclude by comparing the general result above to results that hold only when ∂M^* is totally geodesic.

We remark that $(\mathrm{RF}_k\, M)H$ is usually not closed, even when M is acylindrical, because there can be circles $C \in \overline{\mathcal{C}^*}$ such that $|C \cap \Lambda| = 1$.

Thick sets. We begin by defining $\mathrm{RF}_k\, M$. Let us say a closed set $T \subset \mathbb{R}$ is *k-thick* if

$$[1,k] \cdot |T| = [0,\infty).$$

In other words, given $x \geq 0$ there exists a $t \in T$ with $|t| \in [x, kx]$. Note that if T is k-thick, then so is λT for all $\lambda \in \mathbb{R}^*$.

If the translate $T - x$ is k-thick for every $x \in T$, then we say T is *globally k-thick*. A set $K \subset U$ is (globally) k-thick if its image under an isomorphism $U \cong \mathbb{R}$ is (globally) k-thick.

Unipotent recurrence. For $x \in \mathrm{RF}\, M$, the unipotent orbit xU almost never remains in $\mathrm{RF}\, M$. Provided, however, there is a *thick set* $K \subset U$ such that $xK \subset \mathrm{RF}\, M$, we have sufficient recurrence to carry through many arguments that would be automatic if xU were bounded. The key point is to combine thickness with the polynomial behavior of unipotent flows. This theme is developed in detail in [6, Section 8], and it motivates the definition of $\mathrm{RF}_k\, M$ below.

Let

$$U(z) = \{u \in U : zu \in \mathrm{RF}\, M\} \tag{4.1}$$

denote the return times of $z \in \mathrm{F}\, M$ to the renormalized frame bundle under the horocycle flow. We define $\mathrm{RF}_k\, M$ for each $k > 1$ by

$$\mathrm{RF}_k\, M = \left\{ z \in \mathrm{RF}\, M : \begin{array}{c}\text{there exists a globally } k\text{-thick} \\ \text{set } K \text{ with } 0 \in K \subset U(z)\end{array} \right\}.$$

Let

$$U(z,k) = \{u \in U : zu \in \mathrm{RF}_k\, M\}.$$

PROPOSITION 4.2

Suppose the convex core of M is compact. Then for any $k > 1$, the set $\mathrm{RF}_k\, M$ is a compact, A-invariant subset of $\mathrm{RF}\, M$. Moreover, $U(z,k)$ is k-thick for each $z \in \mathrm{RF}_k\, M$.

Proof

Using compactness of $\mathrm{RF}\, M$, it is easily verified that if $z_n \to z$ in $\mathrm{F}\, M$, then $\limsup U(z_n) \subset U(z)$. One can also check that if $K_n \subset U$ is a sequence of globally k-thick sets with $0 \in K_n$, then $\limsup K_n$ is also globally k-thick. Consequently $\mathrm{RF}_k\, M \subset \mathrm{RF}\, M$ is closed, and hence compact.

Since $U(za)$ is a rescaling of $U(z)$ for any $a \in A$, and the notion of thickness is scale-invariant, $\mathrm{RF}_k\, M$ is A-invariant. For the final assertion, observe that $U(z,k)$ contains the thick set $K \subset U(z)$ posited in the definition of $\mathrm{RF}_k\, M$. ☐

Thickness and moduli. To complete the proof Theorem 4.1, we just need to relate thickness to the results of Section 3. For the next statement, we regard $\widehat{\mathbb{R}} = \mathbb{R} \cup \{\infty\}$ as a circle on $S^2 \cong \widehat{\mathbb{C}}$.

PROPOSITION 4.3

Let $K \subset \widehat{\mathbb{R}}$ be a Cantor set of modulus $\delta > 0$ containing ∞. Then $T = K \cap \mathbb{R}$ is a globally k-thick subset of $\mathbb{R}$, where $k > 1$ depends only on δ.

Proof

Use Proposition 3.2 to relate the modulus of K to the relative sizes of gaps in $\mathbb{R} - K$. ☐

Proof of Theorem 4.1

Since M is acylindrical, by Theorem 3.1 there exists a $\delta > 0$ such that for any $C \in \mathcal{C}^*$, there exists a Cantor set K of modulus δ with

$$K \subset C \cap \Lambda \subset S^2.$$

By Proposition 4.3, there exists a k_0 such that $T \subset \mathbb{R}$ is globally k_0-thick whenever $T \cup \infty$ is a Cantor set of modulus δ.

Let P^* be a plane in M^*. Choose $C \in \mathcal{C}^*$ such that the image of $\mathrm{hull}(C)$ in M^* contains P^*. Let $K \subset C \cap \Lambda$ be the Cantor set of modulus δ provided by Theorem 3.1.

By a change of coordinates, we can arrange that $0, \infty \in K \subset \widehat{\mathbb{R}}$. Let $\widetilde{z} \in \mathrm{F}\mathbb{H}^3$ be any frame tangent to $\mathrm{hull}(\widehat{\mathbb{R}})$ along the geodesic γ joining zero to infinity, and let z denote its projection to $\mathrm{F}M$. Then z is tangent to P^*. It is readily verified that there exists an isomorphism $U \cong \mathbb{R}$ sending $U(z)$ to $\mathbb{R} \cap \Lambda$. Since $0 \in K \subset \mathbb{R} \cap \Lambda$ and K is globally k_0-thick, we have $z \in \mathrm{RF}_{k_0} M$ as well. Thus the Theorem holds for all $k \geq k_0$. □

Comparison with the rigid case. We conclude by comparing the case of a *general* convex cocompact acylindrical 3-manifold M, treated by Theorems 3.1 and 4.1, with the *rigid case*, studied in [6].

In the rigid case, every component D_i of $S^2 - \Lambda$ is a *round* disk; hence $C \cap D_i$ is *connected* for all $C \in \mathcal{C}^*$, and one can show:

$$K = C \cap \Lambda \textit{ is a compact set of definite modulus } \forall C \in \mathcal{C}^*.$$

See [6, Lemma 9.2]. Similarly, all horocycles passing through $\mathrm{RF}\,M$ are *recurrent*, and $\mathrm{RF}_k\, M = \mathrm{RF}\,M$ for all k sufficiently large.

On the other hand, when M is not rigid, there are cases where both these properties fail. For example, suppose the bending measure of $\mathrm{hull}(\Lambda)$ has an atom of mass θ along the geodesic γ joining $p, q \in \Lambda$. Then we can change coordinates on $S^2 \cong \widehat{\mathbb{C}}$ so that $p = 0$, $q = \infty$, and Λ is contained in the wedge defined by $|\arg(z)| < \pi - \theta/2$. Then the circle $C \in \mathcal{C}^*$ defined by $\mathrm{Re}(z) = 1$ cannot meet the limit set in a set of positive modulus, since ∞ is an isolated point of $C \cap \Lambda$.

Similarly, the horocycle in $\mathbb{H}^3 = \mathbb{C} \times \mathbb{R}_+$ defined by $\eta(t) = (it, 1)$ crosses γ when $t = 0$, and satisfies $d(\eta(t), \mathrm{hull}(\Lambda)) \to \infty$ as $|t| \to \infty$. Projecting to M, we obtain a *divergent* horocycle orbit xU with $x \in \mathrm{RF}\,M$. In particular, $x \in \mathrm{RF}\,M - \mathrm{RF}_k\, M$ for all k.

Nevertheless, $C \cap \Lambda$ can *contain* a Cantor set of positive modulus, consistent with Theorem 3.1.

5. The boundary of the convex core

In this short section we analyze the behavior of $C \cap \Lambda$ for circles that meet the limit set but do not separate it. The result we need does not require that M is acylindrical, only that its convex core is compact.

THEOREM 5.1
Let $M = \Gamma \backslash \mathbb{H}^3$ be a convex cocompact 3-manifold with limit set Λ. Let C be the boundary of a supporting hyperplane for $\mathrm{hull}(\Lambda)$. *Then:*

1. Γ^C *is a convex cocompact Fuchsian group.*

2. *There is a finite set* Λ_0 *such that*

$$C \cap \Lambda = \Lambda(\Gamma^C) \cup \Gamma^C \Lambda_0.$$

Here $\Lambda(\Gamma^C)$ denotes the limit set of $\Gamma^C = \{g \in \Gamma : g(C) = C\}$.

COROLLARY 5.2

If the projection of hull(C) *to* M *gives a plane* P *disjoint from* M^* *but tangent to a frame in* $\mathrm{RF}_k\, M$, *then* Γ^C *is nonelementary.*

Proof

The hypotheses guarantee that C does not separate Λ, and $C \cap \Lambda$ contains an (uncountable) Cantor set of positive modulus. Then by the preceding result, $\Lambda(\Gamma^C)$ is uncountable, so Γ^C is nonelementary. □

Proof of Theorem 5.1

We will use the theory of the bending lamination, developed in [2], [3], [10], and elsewhere.

If M^* is empty, then Λ is contained in a circle and the result is immediate. The desired result is also immediate if $C \cap \Lambda$ is finite, because $\Lambda(\Gamma^C) \subset C \cap \Lambda$.

Now assume $C \cap \Lambda$ is infinite and M^* is nonempty. Then $K = \partial\,\mathrm{core}(M)$ is a finite union of disjoint compact pleated surfaces with bending lamination β. Let

$$f : S = \Gamma^C \backslash \mathrm{hull}(C \cap \Lambda) \to M$$

be the natural projection. Since $|C \cap \Lambda| > 2$, S is a metrically complete hyperbolic surface with geodesic boundary, with nonempty interior S_0. The map f sends S_0 isometrically to a component of $K - \beta$; in particular, S_0 has finite area. It follows that the ends of S_0 consist of the regions between finitely many pairs of geodesics which are tangent at infinity; for an example, see Figure 5. Consequently, we can find a finite set $\Lambda_0 \subset \Lambda$ (corresponding to the finitely many ends of S_0) such that

$$C \cap \Lambda = \Lambda(\Gamma^C) \cup \Gamma^C \Lambda_0.$$

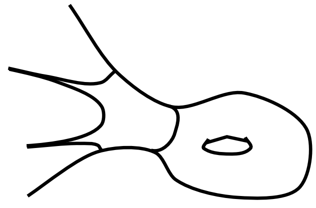

Figure 5. A surface with a crown.

The group Γ^C is convex cocompact because S has finite area and Γ contains no parabolic elements. □

6. Planes near the boundary of the convex core

In this section we take a step toward the proof of Theorem 1.4 by establishing two density results.

THEOREM 6.1
Let $M = \Gamma\backslash\mathbb{H}^3$ be a convex cocompact 3-manifold with incompressible boundary. Consider a sequence of circles $C_n \to C$ with $C_n \in \mathcal{C}^*$ but $C \notin \mathcal{C}^*$. Suppose that $L = \liminf(C_n \cap \Lambda)$ is uncountable. Then $\bigcup \Gamma C_n$ is dense in $\mathcal{C}^*$.

Under the same assumptions on M we obtain:

COROLLARY 6.2
Consider an H-invariant set $E \subset F^*$ and fix $k > 1$. If the closure of $E \cap \mathrm{RF}_k M$ contains a point outside F^*, then E is dense in F^*.

Proof
Consider a sequence $x_n \in E \cap \mathrm{RF}_k M$ such that $x_n \to x \in \mathrm{RF}_k M - F^*$. We then have a corresponding sequence of circles $C_n \in \mathcal{C}^*$ such that $C_n \to C \notin \mathcal{C}^*$. (The circles are chosen so that x_n is tangent to the image of $\mathrm{hull}(C_n)$ in M.)

Pass to a subsequence such that $U(x_n)$ (defined using equation (4.1)) converges, in the Hausdorff topology, to a closed set $K \subset U(x)$. Then $C_n \cap \Lambda$ also converges, to a compact set $L \subset C$ homeomorphic to the 1-point compactification of K. The fact that $x_n \in \mathrm{RF}_k M$ implies that K contains a globally k-thick set; hence K is uncountable, so L is as well. Then by the result above, $\bigcup \Gamma C_n$ is dense in $\mathcal{C}^*$, so E is dense in F^*. □

Roughly speaking, these results show that planes P^* that are nearly tangent to ∂M^* are also nearly dense in M^*, subject to a condition on $\mathrm{RF}_k M$ that is automatic in the acylindrical case by Theorem 4.1.

Fuchsian dynamics. The proof of Theorem 6.1 exploits the dynamics of the Fuchsian group Γ^C. Given an open round disk $D \subset S^2$ and a closed subset $E \subset \partial D$, we let $\mathrm{hull}(E, D) \subset D$ denote the convex hull of E in the hyperbolic metric on D.

The principle we will use is [6, Corollary 3.2], which we restate as follows:

THEOREM 6.3
Let $M = \Gamma\backslash\mathbb{H}^3$ be a convex cocompact hyperbolic 3-manifold. Let $D \subset S^2$ be a

round open disk that meets Λ, and let $C = \partial D$. Suppose Γ^C is a nonelementary, finitely generated group, and let $C_n \to C$ be a sequence of circles such that

$$C_n \cap \operatorname{hull}\big(\Lambda(\Gamma^C), D\big) \neq \emptyset.$$

Then the closure of $\bigcup \Gamma C_n$ in $\mathcal{C}$ contains every circle that meets Λ.

Proof of Theorem 6.1
Let D and D' denote the two components of $S^2 - C$. Since $C \notin \mathcal{C}^*$, at least one of the components, say D', is contained in Ω. Since $L \subset C \cap \Lambda$ is uncountable, Γ^C is nonelementary and finitely generated by Theorem 5.1. Consider an ideal pentagon

$$X = \operatorname{hull}(V, D) \subset \operatorname{hull}\big(\Lambda(\Gamma^C), D\big) \tag{6.1}$$

whose five vertices V lie in L. Since $L = \liminf C_n \cap \Lambda$, we can find "vertices"

$$V_n \subset C_n \cap \Lambda, \qquad |V_n| = 5,$$

such that $V_n \to V$. In particular, $|C_n \cap \Lambda| \geq 3$ for all n.

Note that C_n is the unique circle passing through any three points of V_n. If three of these points were to lie in $\overline{D}'$, then we would have $C_n \subset \overline{D}'$, and hence $|C_n \cap \Lambda| \leq 1$, since $C_n \neq C = \partial \overline{D}'$ and $D' \subset \Omega$. Hence $|V_n \cap D| \geq 3$. Since $|C_n \cap C| \leq 2$, at least two adjacent components of $C_n - V_n$ are contained in D. It follows easily that C_n meets $\operatorname{hull}(V, D)$ for all n sufficiently large. Using equation (6.1), we can then apply Theorem 6.3 to conclude that $\bigcup \Gamma C_n$ is dense in $\mathcal{C}^*$, since every $C \in \mathcal{C}^*$ meets Λ. ☐

7. Planes far from the boundary

In this section we finally prove Theorem 1.4, which we restate as Corollary 7.2. The proof rests on the following more general density theorem.

THEOREM 7.1
Let $M = \Gamma \backslash \mathbb{H}^3$ be a convex cocompact 3-manifold with incompressible boundary. Let $C_i \to C$ be a convergent sequence in $\mathcal{C}^*$, with $C \notin \bigcup \Gamma C_i$.

Suppose that $C \cap \Lambda$ contains a Cantor set of positive modulus. Then $\bigcup \Gamma C_i$ is dense in $\mathcal{C}^*$.

COROLLARY 7.2
If $M = \Gamma \backslash \mathbb{H}^3$ is a convex cocompact acylindrical 3-manifold, then any Γ-invariant set $E \subset \mathcal{C}^*$ is either closed or dense in $\mathcal{C}^*$.

Proof
Suppose E is not closed in $\mathcal{C}^*$. Then we can find a sequence $C_i \in E$ converging to $C \in \mathcal{C}^* - E$. Since M is acylindrical, C meets Λ in a Cantor set of positive modulus, by Theorem 3.1. Since E is Γ-invariant, the preceding result shows that $\bigcup \Gamma C_i$ is dense in $\mathcal{C}^*$, so the same is true for E. □

The proof of Theorem 7.1 follows the same lines as the proof of Theorem 7.3 in [6, Section 9]. We will freely quote results from [6] in the course of the proof. The notation from Table 2 for the subgroups U, V, A, N of G and other objects will also be in play. A generalization of Theorem 7.1 to manifolds with compressible boundary is stated at the end of this section.

Setup in the frame bundle. To prepare for the proof, we first reformulate it in terms of the frame bundle.

Let $C_i \to C$ as in the statement of Theorem 7.1. Since $C \cap \Lambda$ contains a Cantor set of positive modulus, by Proposition 4.3 we can choose $k > 1$ and $x_\infty \in \mathrm{RF}_k\, M \cap F^*$ such that $x_\infty H$ corresponds to ΓC. Let us also choose $x_i \to x_\infty$ in F^* such that $x_i H$ corresponds to ΓC_i. Since $C \notin \bigcup \Gamma C_i$, we also have

$$x_\infty \notin E = \bigcup x_i H.$$

To prove Theorem 7.1 we need to show that:

$$E \textit{ is dense in } F^*.$$

We may also assume that:

$$\textit{The set } \overline{E} \cap \mathrm{RF}_k\, M \cap F^* \textit{ is compact.} \tag{7.1}$$

Otherwise, $\overline{E} \cap F^* = F^*$ by Corollary 6.2, and hence E is dense in F^*.

Dynamics of semigroups. We say that $L \subset G$ is a 1-*parameter semigroup* if there exists a nonzero $\xi \in \mathrm{Lie}(G)$ such that

$$L = \{\exp(t\xi) : t \geq 0\}.$$

To show a set is dense in F^*, we will use the following fact.

PROPOSITION 7.3
Let $L \subset V$ be a 1-parameter semigroup. Then $\overline{xLH}$ contains F^* for all $x \in F^*$.

Proof
Let $C \in \mathcal{C}^*$ be a circle corresponding to xH. Then xLH corresponds to a family of circles C_α such that $\bigcup C_\alpha$ contains one of the components of $S^2 - C$. Since

$C \in \mathcal{C}^*$, both components meet the limit set. Hence $\overline{\Gamma C_\alpha} \supset \mathcal{C}^*$ for some α by [6, Corollary 4.2]. □

The staccato horocycle flow. Recall that the compact set $\mathrm{RF}_k\, M$ is invariant under the geodesic flow A. Moreover, Proposition 4.2 states that

$$U(z,k) = \{u \in U : zu \in \mathrm{RF}_k\, M\}$$

is a thick subset of U, for all $z \in \mathrm{RF}_k\, M$. In other words, $\mathrm{RF}_k\, M$ is also invariant under the *staccato horocycle flow*, which is interrupted outside of $U(z,k)$.

Recurrence. Next we define a compact set W with

$$x_\infty \in W \subset \overline{E} \cap F^*$$

with good recurrence properties for the horocycle flow. Namely, we let

$$W = \begin{cases} (\overline{E} - E) \cap \mathrm{RF}_k\, M \cap F^* & \text{if this set is compact, and} \\ \overline{E} \cap \mathrm{RF}_k\, M \cap F^* & \text{otherwise.} \end{cases} \tag{7.2}$$

(This definition is motivated by the proof of Lemma 7.6.)

In either case, W is compact by assumption (7.1). Since $\overline{E} \cap F^*$ is H-invariant, we have

$$WA = W \qquad \text{and} \qquad WU \cap \mathrm{RF}_k\, M \subset W.$$

The second inclusion gives good recurrence; namely, we have

$$xU(x,k) \subset W \tag{7.3}$$

for all $x \in W$; and $U(x,k)$ is thick, because $W \subset \mathrm{RF}_k\, M$.

The horocycle flow. We now exploit the fact that $\overline{E}$ is invariant under the horocycle flow. The 1-parameter horocycle subgroup $U \subset H$ is distinguished by the fact that its normalizer contains (with finite index) the large subgroup $AN \subset G$. If X is U-invariant, then so is Xg for any $g \in AN$.

Minimal sets. A closed set Y is a *U-minimal set* for $\overline{E}$ with respect to W if $Y \subset \overline{E}$, Y meets W, $YU = Y$, and

$$\overline{yU} = Y \quad \text{for all } y \in Y \cap W.$$

Note that $\overline{E}$ itself has all these properties except for the last. The existence of a minimal set Y follows from the Axiom of Choice and compactness of W. From now on we will assume that:

Y is a U-minimal set for $\overline{E}$ with respect to W.

To show that $\overline{E}$ is large, our strategy is to show it contains Yg for many $g \in AN$. To this end, we remark that for $g \in AN$:

If $(Y \cap W)g$ meets $\overline{E}$, then $Yg \subset \overline{E}$.

Indeed, in this case by minimality we have:

$$\overline{E} \supset \overline{ygU} = \overline{yU}g = Yg, \tag{7.4}$$

where $yg \in (Y \cap W)g \cap \overline{E}$.

Translation of Y inside of Y. The fact that horocycles in Y return frequently to W allows one to deduce additional invariance properties for Y itself. Note that the orbits of AV are orthogonal to the orbits of U in the Riemannian metric on $\mathrm{F}M$.

LEMMA 7.4
There exists a 1-parameter semigroup $L \subset AV$ such that

$$YL \subset Y.$$

Proof
In the rigid acylindrical case, this is Theorem 9.4 in [6] for $W = \mathrm{RF}\,M$. The only property of $\mathrm{RF}\,M$ used in the proof is the k-thickness of $\{u \in U : xu \in \mathrm{RF}\,M\}$ for any $x \in \mathrm{RF}\,M$. Hence the proof works verbatim with W replacing $\mathrm{RF}\,M$, in view of equation (7.3). In fact, $YL = Y$ since $\mathrm{id} \in L$. ☐

Translation of Y inside of $\overline{E}$. Our next goal is to find more elements $g \in G$ that satisfy $Yg \subset \overline{E}$. Consider the closed set $S(Y) \subset G$ defined by

$$S(Y) = \{g \in G : (Y \cap W)g \cap \overline{E} \neq \emptyset\}.$$

Since $\overline{E}$ is H-invariant, we have $S(Y)H = S(Y)$.

LEMMA 7.5
If $S(Y)$ contains a sequence $g_n \to \mathrm{id}$ in $G - H$, then there exists $v_n \in V - \{\mathrm{id}\}$ tending to id such that

$$Yv_n \subset \overline{E}.$$

Proof
Let $g_n \in S(Y)$ be a sequence tending to id in $G - H$. First suppose that there is a subsequence, which we continue to denote by $\{g_n\}$, of the form $g_n = v_n h_n \in VH$. Since $g_n \notin H$, we have $v_n \neq \mathrm{id}$ for all n. The claim then follows from the H-invariance of $S(Y)$ and the U-minimality of Y (see (7.4)).

Therefore, assume that $g_n \notin VH$ for all large n. Since $g_n \in S(Y)$, there exist $y_n \in Y \cap W$ such that $y_n g_n \in \overline{E}$.

Since Y is U-invariant and $WU \cap \mathrm{RF}_k M \subset \mathrm{RF}_k M$, we have $yU(y,k) \subset Y$ for all $y \in Y$, and $U(y,k)$ is a k-thick subset of U.

Therefore, by [6, Theorem 8.1], for any neighborhood G_0 of the identity in G we can choose $u_n \in U(y_n, k)$ and $h_n \in H$ such that

$$u_n^{-1} g_n h_n \to v \in V \cap G_0 - \{\mathrm{id}\}.$$

After passing to a subsequence, we have $y_n u_n \to y_0 \in Y \cap W$. Hence

$$y_n g_n h_n = (y_n u_n)(u_n^{-1} g_n h_n) \in \overline{E}$$

converges to $y_0 v \in \overline{E}$.

Since Y is U-minimal with respect to W and $y_0 \in Y \cap W$, we have

$$\overline{y_0 v U} = \overline{y_0 U} v = Y v \subset \overline{E}.$$

Since G_0 was an arbitrary neighborhood of the identity, the claim follows. ☐

Choosing Y. In general there are many possibilities for the minimal set Y, and it may be hard to describe a particular one, since the existence of a minimal set is proved using the Axiom of Choice. The next result shows that, nevertheless, we can choose Y so it remains inside $\overline{E}$ under suitable translations transverse to H but still in AN.

LEMMA 7.6

There exists a U-minimal set Y for $\overline{E}$ with respect to W, and a sequence $v_n \to \mathrm{id}$ in $V - \{\mathrm{id}\}$, such that

$$Y v_n \subset \overline{E}$$

for all n.

Proof

By Lemma 7.5, it suffices to show that Y can be chosen so that $S(Y)$ contains a sequence $g_n \to \mathrm{id}$ in $G - H$. We break the analysis into two cases, depending on whether or not E meets the compact set W.

First consider the case where E is disjoint from W. Let Y be a U-minimal set for $\overline{E}$ with respect to W. Choose $y \in Y \cap W$. Since $Y \subset \overline{E}$, there exist $g_n \to \mathrm{id}$ such that $y g_n \in E$. Then $y \notin E$, and hence $g_n \in G - H$, so we are done.

Now suppose E meets W. Then $W - E$ is not closed, by equation (7.2). So in this case there exists a sequence $x_n \in W - E$ with $x_n \to x \in E \cap W$. In particular, $\overline{xH} \cap W \neq \emptyset$. Thus there exists a U-minimal set Y for $\overline{xH}$ with respect to W.

We now consider two cases. Assume first that $Y \cap W \subset xH$. Pick $y \in Y \cap W$; then $y = xh$ for some $h \in H$. Since $x_n \to x$ we have $x_n h \to y$. Now writing $yg_n = x_n h$, we have $g_n \to \mathrm{id}$. As $y \in xH \subset E$ and $x_n \notin E$, we have $g_n \in G - H$, and we are done.

Now suppose that $W \cap Y \not\subset xH$. Choose $y \in (W \cap Y) - xH$. Since we have $Y \subset \overline{xH}$, there exist $g_n \to \mathrm{id}$ with $yg_n \in xH$. Moreover, $g_n \in G - H$ since $y \notin xH$, and the proof is complete in this case as well. □

Semigroups. We are now ready to complete the proof of Theorem 7.1. We will exploit the 1-parameter semigroup $L \subset AV$ guaranteed by Lemma 7.4. To discuss the possibilities for L, let us write the elements of V and A as

$$v(s) = \begin{pmatrix} 1 & is \\ 0 & 1 \end{pmatrix} \quad \text{and} \quad a(t) = \begin{pmatrix} e^t & 0 \\ 0 & e^{-t} \end{pmatrix}.$$

We then have two semigroups in V, defined by $V_\pm = \{v(s) : \pm s \geq 0\}$, and two similar semigroups in $A_\pm$ in A. It will also be useful to introduce the interval

$$V_{[a,b]} = \big\{v(s) : s \in [a,b]\big\}.$$

In the notation above, if $L \subset AV$ is a 1-parameter semigroup, then either

(i) $L = V_\pm$;

(ii) $L = A_\pm$; or

(iii) $L = v^{-1} A_\pm v$, for some $v \in V$, $v \neq \mathrm{id}$.

Proof of Theorem 7.1

To complete the proof, it only remains to show we have $F^* \subset \overline{E}$.

Choose Y and $v_n \in V$ so that $Yv_n \subset \overline{E}$ as in Lemma 7.6. Write $v_n = v(s_n)$; then $s_n \to 0$ and $s_n \neq 0$. Passing to a subsequence, we can assume s_n has a definite sign, say, $s_n > 0$.

By Lemma 7.4, there is a 1-parameter semigroup $L \subset AV$ such that

$$YL \subset Y.$$

The rest of the argument breaks into three cases, depending on whether L is of type (i), (ii), or (iii) in the list above.

(i). If $L = V_\pm$, then we have $F^* \subset \overline{YLH} \subset \overline{E}H = \overline{E}$ by Proposition 7.3, and we are done.

(ii). Now suppose $L = A_\pm$. Let

$$B = \{\mathrm{id}\} \cup \bigcup A_\pm v_n A.$$

Since $YL \subset Y$ and $Yv_nA \subset \overline{E}A = \overline{E}$ for all n, we have

$$YB \subset \overline{E}.$$

Note that $a(t)v(s)a(-t) = v(e^{2t}s)$. Consequently, we have

$$v(e^{2t}s_n) \in B$$

for all n and all t with $a(t) \in L = A_\pm$.

Suppose $L = A_+$. Since $s_n \to 0$ and $s_n > 0$, in this case we have $V_+ \subset B$; hence $YV_+H \subset \overline{E}$ and we are done as in case (i).

Now suppose $L = A_-$. In this case at least we obtain an interval

$$V_{[0,s_1]} \subset B.$$

Choose a sequence $a_n \in A$ such that $V_+ = \bigcup a_n V_{[0,s_1]} a_n^{-1}$. Consider $y \in Y \cap W$. Since $ya_n^{-1} \in W$, and W is compact, after passing to a subsequence we can assume that

$$ya_n^{-1} \to y_0 \in W \subset F^*.$$

We then have

$$y_0 V_+ = \bigcup ya_n^{-1}(a_n V_{[0,s_1]} a_n^{-1}) \subset \overline{E},$$

which again implies that $F^* \subset \overline{E}$, by Proposition 7.3.

(iii). Finally, consider the case $L = v^{-1}A_\pm v$ for some $v \in V$, $v \neq \mathrm{id}$. We then have $YB \subset \overline{E}$, where

$$B = v^{-1}A_\pm vA.$$

By an easy computation, B contains $V_{[0,\pm s]}$ for some $s > 0$, and the argument is completed as in case (ii). □

The compressible case. In conclusion, we remark that Theorems 6.1 and 7.1 remain true without the hypothesis that M has incompressible boundary, provided we replace $\mathcal{C}^*$ with

$$\mathcal{C}^\# = \{C \in \mathcal{C}^* : C \text{ meets } \Lambda\}$$

and require that M^* is nonempty. The proofs are simple variants of those just presented.

Acknowledgments. We would like to thank Elon Lindenstrauss and Yair Minsky for useful discussions. This research was supported in part by the Alfred P. Sloan Foundation (A.M.) and the National Science Foundation.

References

[1] P. EBERLEIN, *Geodesic flows on negatively curved manifolds, I*, Ann. of Math. (2) **95** (1972), 492–510. MR 0310926. DOI 10.2307/1970869. *(1039)*

[2] D. B. A. EPSTEIN and A. MARDEN, "Convex hulls in hyperbolic space, a theorem of Sullivan, and measured pleated surfaces" in *Analytical and Geometric Aspects of Hyperbolic Space*, Cambridge University Press, Cambridge, 1987, 113–254. MR 0903852. *(1050)*

[3] Y. KAMISHIMA and SER P. TAN, "Deformation spaces on geometric structures" in *Aspects of Low Dimensional Manifolds*, Kinokuniya, Tokyo, 1992, 263–300. MR 1208313. DOI 10.2969/aspm/02010263. *(1050)*

[4] O. LEHTO and K. J. VIRTANEN, *Quasiconformal Mappings in the Plane*. Springer, New York, 1973. MR 0344463. *(1042, 1043)*

[5] A. MARDEN, *Hyperbolic Manifolds*, Cambridge Univ. Press, Cambridge, 2016, 263–300. MR 3586015. DOI 10.1017/CBO9781316337776. *(1037)*

[6] C. MCMULLEN, A. MOHAMMADI, and H. OH, *Geodesic planes in hyperbolic 3-manifolds*, Invent. Math. **209** (2017), no. 2, 425–461. MR 3674219. DOI 10.1007/s00222-016-0711-3. *(1031, 1035, 1041, 1047, 1049, 1051, 1053, 1054, 1055, 1056)*

[7] J. MORGAN, "On Thurston's uniformization theorem for three dimensional manifolds" in *The Smith Conjecture*, Academic Press, Orlando, 1984, 37–125. MR 0758464. DOI 10.1016/S0079-8169(08)61637-2. *(1037)*

[8] M. RATNER, *Raghunathan's topological conjecture and distributions of unipotent flows*, Duke Math. J. **63** (1991), no. 1, 235–280. MR 1106945. DOI 10.1215/S0012-7094-91-06311-8. *(1031)*

[9] N. A. SHAH, "Closures of totally geodesic immersions in manifolds of constant negative curvature" in *Group Theory from a Geometrical Viewpoint (Trieste, 1990)*, World Scientific, River Edge, 1991, 718–732. MR 1170382. *(1031)*

[10] W. P. THURSTON, *Geometry and topology of three-manifolds*, lecture notes, Princeton University, 1979. *(1050)*

[11] ———, *Hyperbolic structures on 3-manifolds, I: Deformations of acylindrical manifolds*, Annals of Math. (2) **124** (1986), no. 2, 203–246. MR 0855294. DOI 10.2307/1971277. *(1037)*

[12] G. T. WHYBURN, *Topological characterization of the Sierpiński curve*. Fund. Math. **45** (1958), 320–324. MR 0099638. DOI 10.4064/fm-45-1-320-324. *(1030)*

[13] Y. ZHANG, *Existence of an exotic plane in an acylindrical* 3*-manifold*, to appear in Math. Res. Lett., preprint, 2020. *(1031)*

McMullen
Mathematics Department, Harvard University, Cambridge, Massachusetts, USA

Mohammadi
Mathematics Department, UC San Diego, La Jolla, California, USA

Oh
Mathematics Department, Yale University, New Haven, Connecticut, USA; and Korea Institute for Advanced Study, Seoul, Korea