



Heriot-Watt University
Research Gateway

Complete topological descriptions of certain Morse boundaries

Citation for published version:

Charney, R, Cordes, M & Sisto, A 2023, 'Complete topological descriptions of certain Morse boundaries', *Groups, Geometry, and Dynamics*, vol. 17, no. 1, pp. 157-184. <https://doi.org/10.4171/GGD/669>

Digital Object Identifier (DOI):

[10.4171/GGD/669](https://doi.org/10.4171/GGD/669)

Link:

[Link to publication record in Heriot-Watt Research Portal](#)

Document Version:

Publisher's PDF, also known as Version of record

Published In:

Groups, Geometry, and Dynamics

Publisher Rights Statement:

© 2023 European Mathematical Society.

General rights

Copyright for the publications made accessible via Heriot-Watt Research Portal is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

Heriot-Watt University has made every reasonable effort to ensure that the content in Heriot-Watt Research Portal complies with UK legislation. If you believe that the public display of this file breaches copyright please contact open.access@hw.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Complete topological descriptions of certain Morse boundaries

Ruth Charney, Matthew Cordes, and Alessandro Sisto

Abstract. We study direct limits of embedded Cantor sets and embedded Sierpiński curves. We show that under appropriate conditions on the embeddings, all limits of Cantor spaces give rise to homeomorphic spaces, called ω -Cantor spaces, and, similarly, all limits of Sierpiński curves give homeomorphic spaces, called ω -Sierpiński curves. We then show that the former occur naturally as Morse boundaries of right-angled Artin groups and fundamental groups of non-geometric graph manifolds, while the latter occur as Morse boundaries of fundamental groups of finite-volume, cusped hyperbolic 3-manifolds.

1. Introduction

Many geodesic metric spaces which are not Gromov hyperbolic nonetheless display some hyperbolic-like behavior. The Morse boundary was introduced in [5, 6] with the goal of identifying and encoding this behavior in a useful way. It is defined for any proper geodesic metric space (though in some cases it may be empty), and a key property states that it is invariant under quasi-isometries. In particular, the Morse boundary of a finitely generated group G is well defined and provides a quasi-isometry invariant for the group.

Set theoretically, the Morse boundary of X consists of equivalence classes of geodesic rays satisfying the Morse property, namely, the property that any (λ, ε) -quasi-geodesics with endpoints on the ray stay bounded distance from the ray, where the bound, $N = N(\lambda, \varepsilon)$, depends only on λ and ε . If one restricts the Morse gauge N , then the corresponding boundary points can be given a natural topology. This space is denoted by $\partial_M^N X$. The Morse boundary $\partial_M X$ is defined as the direct limit of these spaces. (See Section 2 below for formal definitions.)

It is shown in [8] that the intermediate strata $\partial_M^N X$ are, in essence, boundaries of hyperbolic spaces and thus have many nice properties. Using the direct limit topology on $\partial_M X$ has the advantage that many questions can be reduced to understanding these intermediate spaces. Indeed, the Morse boundary has been shown to have many properties analogous to boundaries of hyperbolic spaces (see, e.g., [4, 8, 19, 24, 30]). On the other hand, the limit space itself can be quite complicated. In particular, if X itself is not a

2020 *Mathematics Subject Classification.* Primary 20F67; Secondary 20F65, 57K32.

Keywords. Morse boundary, Cantor set, Sierpiński curve, right-angled Artin groups, hyperbolic 3-manifolds.

hyperbolic space, then the topology on $\partial_M X$ is neither compact nor metrizable. However, in this paper, we provide the first complete topological descriptions of non-compact Morse boundaries, showing that, at least sometimes, Morse boundaries are more “accessible” as topological spaces than previously thought.

More specifically, we study the Morse boundaries of right-angled Artin groups, fundamental groups of non-geometric graph manifolds, and fundamental groups of finite-volume cusped hyperbolic 3-manifolds. Understanding these boundaries involves not only characterizing the strata $\partial_M^N X$, but also understanding how they sit inside each other. As we will discuss below, the process is more subtle than it at first appears.

Direct limits of Cantor spaces and right-angled Artin groups

For the case of a right-angled Artin group, we prove the following theorem:

Theorem 1.1. *Let A_Γ be a right-angled Artin group, for Γ a finite graph. Then $\partial_M A_\Gamma$ satisfies exactly one of the following:*

- (1) $\partial_M A_\Gamma$ is empty,
- (2) $\partial_M A_\Gamma$ consists of two points,
- (3) $\partial_M A_\Gamma$ is a Cantor space or
- (4) $\partial_M A_\Gamma$ is an ω -Cantor space.

In addition, we give precise conditions on the defining graph Γ for when each of these cases occurs. Here, an ω -Cantor space is defined as a direct limit of a sequence of Cantor spaces $X_1 \subset X_2 \subset X_3 \cdots$ such that X_i has empty interior in X_{i+1} for all i . We call these limit spaces ω -spaces because all σ -compact Morse boundaries are examples of k_ω -spaces (for more information see [12]). Crucially, in Theorem 3.3, we show that any two ω -Cantor spaces are homeomorphic, so that (4) describes a well-defined homeomorphism type.

We also show that nice graph of group decompositions have ω -Cantor space boundaries:

Theorem 1.2. *Let G be a group with a graph of groups decomposition \mathcal{G} such that the vertex groups are undistorted in G and have empty Morse boundary, and G acts acylindrically on the Bass–Serre tree associated to the decomposition. Then $\partial_M G$ is an ω -Cantor space.*

An immediate corollary of this theorem is that non-geometric graph manifolds have ω -Cantor space boundaries:

Corollary 1.3. *Let M be a non-geometric graph manifold and let $G = \pi_1(M)$. Then $\partial_M G$ is an ω -Cantor space.*

Forthcoming work of Jacob Russell on hierarchically hyperbolic groups contains more examples of groups with ω -Cantor space Morse boundaries. In addition, as pointed out in

Remark 4.7, our arguments show that under the hypotheses of the theorem, the group G has the Morse-local-to-global property as defined in [25]. We thank Russell for pointing this out.

Theorem 1.2 in fact follows from a more general theorem:

Theorem 1.4. *Let G be a finitely generated group. Suppose that $\partial_M G$ is totally disconnected, σ -compact, and contains a Cantor subspace. Then $\partial_M G$ is either a Cantor space or an ω -Cantor space. It is a Cantor space if and only if G is hyperbolic, in which case G is virtually free.*

We note that all hypotheses on $\partial_M G$ are necessary, but for the last two there is no known counterexample; for more discussion see below.

At first sight, one might think that Theorem 1.4 should be straightforward to prove using the stratification of $\partial_M G$ given by Morse gauges. By hypothesis, each stratum is a totally disconnected compact space, which should make it a Cantor space, and hence would make $\partial_M G$ a direct limit of Cantor spaces, as required. However, the fact that strata are compact and totally disconnected does not mean that they are Cantor spaces; they could be any closed subspace of a Cantor space, of which there are uncountably many homeomorphism types. The solution is to perturb the stratification by Morse gauges to ensure that each stratum is a Cantor space. To do this we need to, roughly speaking, “add” a sequence of Cantor spaces converging to each isolated point of the stratum. This sequence of Cantor spaces is obtained as translates of the Cantor space given in the hypothesis of the theorem.

The hypothesis that $\partial_M G$ is σ -compact is necessary since the ω -Cantor space is σ -compact by construction. On the other hand, to the best of our knowledge there is no known example of a group G so that $\partial_M G$ is not σ -compact. However, we believe that such examples can be found among small-cancellation groups, and it is even possible that every infinitely presented $C'(1/6)$ -small-cancellation group has non- σ -compact Morse boundary.

The hypothesis that $\partial_M G$ contains a Cantor subspace is also a necessary condition. However, there is no known example of a group with non-compact Morse boundary that does not contain a Cantor subspace. In fact, in most of the motivating examples, one has a stable free subgroup (e.g., in acylindrically hyperbolic groups). On the other hand, there are examples by E. Fink of torsion groups with Morse rays [11], but even those probably contain a Cantor subspace. In fact, Fink finds a tree in the Cayley graph where all rays are Morse. We believe that this tree is most likely stable, thereby exhibiting a Cantor subspace in the Morse boundary. This naturally raises a question: Does there exist a group whose Morse boundary is non-compact and does not contain a Cantor subspace?

Direct limits of Sierpiński curves and cusped hyperbolic 3-manifolds

In the case of fundamental groups of cusped hyperbolic 3-manifolds, we show that their Morse boundaries are direct limits of a sequence $S_1 \subset S_2 \subset S_3 \cdots$ of Sierpiński curves.

Moreover, these curves are nicely embedded, namely the peripheral curves of S_i are disjoint from those of S_{i+1} . We call such a limit an ω -Sierpiński curve and prove that any two such are homeomorphic.

Theorem 1.5. *Let M be a finite-volume hyperbolic 3-manifold with at least one cusp, and let $G = \pi_1(M)$. Then $\partial_M G$ is an ω -Sierpiński curve.*

The difficulty here is similar to the problems that arise in the proof of Theorem 1.4. Roughly, each stratum is obtained from $\mathbb{S}^2 = \partial\mathbb{H}^3$ by removing “shadows” of horoballs, so one might at first expect that strata are Sierpiński curves. However, the shadows are not disjoint, so the situation is more complicated, and once again we have to perturb the stratification. In this case, more technology is needed to do so, namely technology from [20, 21], which allows us to carefully detour all the boundary circles of the shadows in order to make them disjoint. This is a rather delicate procedure, and in fact the technology to make something like this work in higher dimensions seems to not be available. It would be interesting to know whether Morse boundaries of fundamental groups of non-compact, finite-volume hyperbolic n -manifolds are all homeomorphic to a certain limit of the $(n-1)$ -dimensional analogues of Sierpiński curves.

One corollary of the fact the Morse boundary of a finite-volume hyperbolic 3-manifold with at least one cusp is an ω -Sierpiński curve is that the boundary is not totally disconnected. Tran conjectured that the Morse boundary of a right-angled Coxeter group is totally disconnected if and only if every induced cycle of length greater than four in the defining graph contains a pair of non-adjacent vertices that are contained in an induced 4-cycle [28]. As Tran pointed out to us, the one-skeleton of a 3-cube satisfies this condition and its associated right-angled Coxeter group is virtually a finite volume hyperbolic 3-manifold with cusps [14]. It follows from our theorem that this is a counterexample to Tran’s conjecture. We note that this is not the first counterexample to this conjecture; see [13, 16].

Outline

Section 2 contains definitions and background on Morse boundaries. In Sections 3 and 4, we investigate limits of Cantor sets and apply this to characterize Morse boundaries of right-angled Artin groups. In Sections 5 and 6, we investigate limits of Sierpiński curves and prove Theorem 1.5.

2. Background

2.1. Morse boundary

We will assume throughout that X is a proper geodesic metric space. We begin with a definition of the Morse boundary of X and some properties that will be needed for the arguments below. We refer the reader to [6] for more details.

Definition 2.1. A geodesic α in X is *Morse* if there exists a function $N : \mathbb{R}^+ \times \mathbb{R}^+ \rightarrow \mathbb{R}^+$ such that any (λ, ε) -quasi-geodesic with endpoints on α lies in the $N(\lambda, \varepsilon)$ -neighborhood of α . The function N is called a *Morse gauge* for α .

Lemma 2.2 ([19, Lemma 3.1]). *For each Morse gauge N , there exists an N' depending only on N , such that if α is an N -Morse geodesic in X and β is a subgeodesic of α , then β is N' -Morse.*

Remark 2.3. Taking Lemma 2.2 into consideration, we call a geodesic N -Morse if the geodesic and all subgeodesics have Morse gauge N .

It is shown in [6] that if two sides of a triangle in X are N -Morse, then the third side is N_1 -Morse where N_1 depends only on N . This also holds for ideal triangles by [4, Lemma 2.3]. We record this property here as it will be used repeatedly below.

Lemma 2.4 ([4]). *Let X be a proper geodesic metric space. Let $\Delta(x, y, z)$ be a geodesic triangle with vertices $x, y, z \in X \cup \partial_M X$ and suppose that two sides of Δ are N -Morse. Then the third side is N' -Morse where N' depends only on N .*

For two Morse rays α, β in X , say $\alpha \sim \beta$ if they have bounded Hausdorff distance. It is shown in [6] that this bound depends only on the Morse gauge N , that is, there exists C_N such that two N -Morse rays α, β are equivalent if and only if $d(\alpha(t), \beta(t)) < C_N$ for all t . The Morse boundary of X consists of the set of equivalence classes of Morse rays. To topologize this set, first choose a basepoint $e \in X$ and let N be a Morse gauge. Set

$$\partial_M^N X_e = \{[\alpha] \mid \exists \beta \in [\alpha] \text{ that is an } N\text{-Morse geodesic ray with } \beta(0) = e\}$$

with the compact-open topology. These spaces are compact. This topology is equivalent to one defined by a system of neighborhoods, $\{V_n(\alpha) \mid n \in \mathbb{N}\}$, defined as follows: $V_n(\alpha)$ is the set of $[\gamma] \in \partial_M^N X_e$ such that $d(\alpha(t), \gamma(t)) < C_N$ for all $t < n$.

Let \mathcal{M} be the set of all Morse gauges. Put a partial ordering on \mathcal{M} so that for two Morse gauges $N, N' \in \mathcal{M}$, we say $N \leq N'$ if and only if $N(\lambda, \varepsilon) \leq N'(\lambda, \varepsilon)$ for all $\lambda, \varepsilon \in \mathbb{N}$. Define the *Morse boundary* of X to be

$$\partial_M X = \varinjlim \partial_M^N X_e$$

with the induced direct limit topology; i.e., a set U is open in $\partial_M X$ if and only if $U \cap \partial_M^N X_e$ is open for all N . A change in basepoint results in a homeomorphic boundary, justifying the omission of the basepoint from the notation. More generally, we will usually assume that the basepoint is fixed and omit it from the notation $\partial_M^N X_e$ as well. The reader is warned, however, that unlike $\partial_M X$, these subspaces do depend on a choice of basepoint.

An alternate construction of the Morse boundary is given by the second author and Hume in [8]. Define $X_e^{(N)}$ to be the set of all $y \in X$ such that there exists an N -Morse geodesic $[e, y]$ in X . Then $X_e^{(N)}$ satisfies the Gromov 4-point definition of hyperbolicity, hence, we may consider its Gromov boundary, $\partial X_e^{(N)}$, and the associated visual metric $d_{(N)}$. The collection of boundaries $(\partial X_e^{(N)}, d_{(N)})$ is called the metric Morse boundary

of X . It is shown in [8] that there are natural homeomorphisms between $\partial X_e^{(N)}$ and $\partial_M^N X_e$ that induce a homeomorphism on the direct limits. Thus the Morse boundary, $\partial_M X$, can also be thought of as the direct limit of the metric spaces $\partial X_e^{(N)}$.

We now establish some useful properties of these spaces.

Lemma 2.5. *For any N, N' with $N < N'$, the inclusion $i: \partial_M^N X \hookrightarrow \partial_M^{N'} X$ is a topological embedding.*

Proof. This map is continuous by [6, Corollary 3.2]. What is left to show is that the map is closed. Let K be a closed set in $\partial_M^N X$. Let γ_i be a sequence in $i(K)$ converging to a point γ . Since the γ_i are represented by N -Morse geodesics and N -Morse geodesics converge to N -Morse geodesics in [6, Lemma 2.10], it converges to an N -Morse geodesic α in $\partial_M^{N'} X$. But since i is simply the inclusion and K is closed in $\partial_M^N X$, then $\alpha = \gamma$ and thus $i(K)$ is closed. ■

The following is a general fact about σ -compact Morse boundaries. It says that one can choose a countable exhaustion of the Morse boundary by strata.

Lemma 2.6. *If $\partial_M X$ is σ -compact, then there exists a sequence of Morse gauges N_1, N_2, \dots with $N_i \leq N_{i+1}$ so that $\partial_M X = \bigcup \partial X_e^{(N_i)} = \bigcup \partial_M^{N_i} X$.*

Proof. By definition, since $\partial_M X$ is σ -compact, then $\partial_M X = \bigcup_{i \in \mathbb{N}} K_i$ where K_i are compact and $K_i \subseteq K_{i+1}$. By [7, Lemma 4.1], we know that for any compact subset $K \subset \partial_M X$ there exists an N so that $K \subset \partial_M X_e^{(N)}$. Thus for each K_i we have a Morse gauge N_i and the ascending condition on the K_i guarantees that $N_i \leq N_{i+1}$ and since the K_i exhaust $\partial_M X$, then it follows that $\partial_M X = \bigcup \partial X_e^{(N_i)}$. ■

3. Limits of Cantor spaces

In this section, we describe particular direct limits of Cantor sets, and more precisely of sequences of Cantor spaces each having empty interior in the next one.

Definition 3.1. Let $C \subsetneq D$ be Cantor spaces. We say that C is *entwined* in D if C has empty interior in D . An ω -Cantor space is a topological space $\lim_{\rightarrow i \in \mathbb{N}} X_i$, where each X_i is a Cantor space and each X_i is entwined in X_{i+1} .

Example 3.2. Let C be the usual middle-third Cantor space. For each middle third interval I , glue in a copy C' of C by identifying the boundary points of C' with the boundary points of I . Call the resulting space D . Then D is again a Cantor space and C is entwined in D .

A similar example can be described in terms of boundaries. Let T be a (rooted) trivalent tree so that ∂T can be identified with C . Now glue a separate copy of T to each vertex v of T by identifying the root to v . The result is a larger tree T' with $T \subset T'$ such that ∂T is entwined in $\partial T'$.

The main result of this section is that there is only one ω -Cantor space up to homeomorphism:

Theorem 3.3. *Any two ω -Cantor spaces are homeomorphic.*

The rest of this section is devoted to the proof of this theorem. After the preliminary observation in the remark below, we show that, roughly speaking, when $C \subseteq D$ are Cantor spaces, with C entwined in D , and C is subdivided in two clopen sets C_0, C_1 , we can take small clopen neighborhoods of C_0, C_1 in D .

Remark 3.4. A clopen set in a Cantor space is either empty or a Cantor space.

Lemma 3.5. *Let $C \subseteq D$ be Cantor spaces, with C entwined in D . Then if $C = C_0 \sqcup C_1$ with C_i clopen and non-empty, then there exist clopen sets $D_0, D_1 \subseteq D$ so that*

- (1) $C_i = D_i \cap C$,
- (2) $D_0 \cap D_1 = \emptyset$,
- (3) $D_0 \cup D_1 \subsetneq D$,
- (4) *if d is any metric on D compatible with its topology, then for any $\varepsilon > 0$, we can choose $D_i \subseteq N_\varepsilon(C_i)$,*
- (5) C_i and D_i are Cantor spaces, and C_i is entwined in D_i .

Proof. Fix a metric d on D compatible with its topology and $\varepsilon > 0$. We can choose $\varepsilon' > 0$ so that $D'_i = N_{\varepsilon'}(C_i)$ satisfy the first four properties listed above (in fact, up to now we only need C to be properly contained in D). However, the D'_i are open, but they might be not closed; we now shrink them to make them clopen. Since every point in a Cantor space has a neighborhood basis of clopen sets, for every point x in $D' = D - (D'_0 \cup D'_1)$ there exists a clopen set U_x disjoint from $C_0 \cup C_1$ (we also use that $C_0 \cup C_1$ is closed). Since D' is compact, it is covered by finitely many U_{x_i} , and the union U of such U_{x_i} is clopen. Then, $D_i = D'_i - U$ is open (because it is obtained by removing a closed set from an open set) and closed (since the complement is $D'_{1-i} \cup U$, which is open), and it is readily seen that the first four properties are satisfied.

The fact that C_i and D_i are Cantor spaces follows now from Remark 3.4. We now argue that C_i is entwined in D_i . Notice that if U is open in D_i , then it is also open in D , since D_i is open. In particular, the interior of C_i as a subspace of D_i coincides with the interior of C_i as a subspace of D . Since $C \supseteq C_i$ has empty interior in D_i , we have that C_i has empty interior in D , whence in D_i , as required. ■

Next, we show that we can extend partial homeomorphisms of Cantor spaces defined over entwined Cantor subspaces.

Lemma 3.6. *Let $C \subseteq D$ and $C' \subseteq D'$ be Cantor spaces, with C entwined in D and C' entwined in D' . Let $\phi : C \rightarrow C'$ be a homeomorphism. Then there exists a homeomorphism $\bar{\phi} : D \rightarrow D'$ that extends ϕ .*

Proof. Let Ω be the set of finite words in the alphabet $\{0, 1\}$, and denote by e the empty word. We denote the word obtained by appending 0 (resp. 1) at the end of the word $\omega \in \Omega$ simply by $\omega 0$ (resp. $\omega 1$). Fix metrics on D and D' compatible with their respective topologies. We can choose clopen sets $C_\omega \subseteq C$ with the following properties:

- (1) $C_e = C$,
- (2) $C_{\omega 0}, C_{\omega 1} \subseteq C_\omega$,
- (3) $C_\omega = C_{\omega 0} \sqcup C_{\omega 1}$,
- (4) for any $\varepsilon > 0$, there are only finitely many C_ω whose diameter is larger than ε .

Set $C'_\omega = \phi(C_\omega)$, and notice that analogous properties hold for the C'_ω as well. Also, set $D_e = D$, and, inductively on the length of ω , use Lemma 3.5 to construct clopen sets $D_\omega \subseteq D$ satisfying

- (1) $D_\omega \cap C = C_\omega$,
- (2) $D_{\omega 0}, D_{\omega 1} \subseteq D_\omega$,
- (3) $D_{\omega 0} \cap D_{\omega 1} = \emptyset$,
- (4) $K_\omega = D_\omega - (D_{\omega 0} \cup D_{\omega 1})$ is non-empty,
- (5) for any $\varepsilon > 0$, there are only finitely many D_ω whose diameter is larger than ε .

Also, construct $D'_\omega \subseteq D'$ (and K'_ω) with the analogous properties. Notice that $D - C = \bigsqcup_{\omega \in \Omega} K_\omega$ and that each K_ω (which is clopen) is a Cantor space by Remark 3.4. Similar observations hold for $D' - C'$ and K'_ω .

We are now ready to define $\bar{\phi}$. Choose for each ω any homeomorphism $\phi_\omega : K_\omega \rightarrow K'_\omega$, and set $\bar{\phi}(x) = x$ if $x \in C$, and $\bar{\phi}(x) = \phi_\omega(x)$ if $x \in K_\omega$. Then $\bar{\phi}$ is bijective. Hence, since D, D' are compact and Hausdorff, to show that $\bar{\phi}$ is a homeomorphism we are left to show that it is continuous. Notice that by construction, we have $\bar{\phi}(D_\omega) = D'_\omega$ for each ω . It is readily checked that $\{D'_\omega\}$ is a basis for the topology of D' . In fact, for any $x \in D'$ and $\varepsilon > 0$, there exists some D'_ω of diameter at most ε that contains x , since there are infinitely many D'_ω containing x and only finitely many of them have diameter larger than ε . Hence, the preimage under $\bar{\phi}$ of any member of a neighborhood basis for D' is open, and hence $\bar{\phi}$ is continuous. ■

Finally, we are ready to prove Theorem 3.3, by iterative extensions.

Proof of Theorem 3.3. Let $X = \varinjlim_{i \in \mathbb{N}} X_i$ and $X' = \varinjlim_{i \in \mathbb{N}} X'_i$ be ω -Cantor spaces. Start with any homeomorphism $\phi_0 : X_0 \rightarrow X'_0$ and inductively define homeomorphisms $\phi_{n+1} : X_{n+1} \rightarrow X'_{n+1}$ extending ϕ_n ; those exist by Lemma 3.6. Then the ϕ_n give a well-defined bijection $X \rightarrow X'$, which is a homeomorphism by definition of the direct limit topology. ■

4. ω -Cantor boundaries

We now study totally disconnected Morse boundaries of finitely generated groups. In this case, we write $\partial_M G$ for the Morse boundary of some (hence any) Cayley graph of G .




Figure 1. If the β_i^j do not converge to γ , then the $g_i\alpha_j$ both backtrack a lot along γ . But then, this means that the α_j fellow-travel for arbitrarily long times.

Following [8], we say that a quasi-convex subspace Y of a geodesic metric space X is N -stable if every pair of points in Y can be connected by a geodesic which is N -Morse in X . We say that a subgroup is stable if it is stable as a subspace.

The main theorem in this section is the following.

Theorem 1.4. *Let G be a finitely generated group. Suppose that $\partial_M G$ is totally disconnected, σ -compact, and contains a Cantor subspace. Then it is either a Cantor space or an ω -Cantor space. It is a Cantor space if and only if G is hyperbolic, in which case G is virtually free.*

We will need the following lemma to construct sequences limiting to a specific boundary point, to show that such boundary point is not isolated.

Lemma 4.1. *Let G be a group, let $\gamma, \alpha_1, \alpha_2$ be Morse rays, with $[\alpha_1] \neq [\alpha_2]$, where all rays are based at e . Also, let $g_i = \gamma(i) \in G$ be the sequence of group elements traversed by γ . Then for some $j \in \{1, 2\}$ and some subsequence $i(n)$ we have $[g_{i(n)}\alpha_j] \rightarrow [\gamma]$.*

Proof. By Lemma 2.4, if two legs of a geodesic triangle are N -Morse, then the third leg is N' -Morse, where N' depends only on N . In our case, this says that there exists N , depending on the Morse gauges of $\gamma, \alpha_1, \alpha_2$, so that there exist N -Morse geodesic rays β_i^j based at e with $[\beta_i^j] = [g_i\alpha_j]$ (Figure 1). Moreover, there exists C , again depending only on the various Morse gauges, so that all triangles $\gamma([0, i]) \cup \beta_i^j \cup g_i\alpha_j$ are C -slim [6, Lemma 2.2].

Suppose by contradiction that for $j = 1, 2$ no subsequence of β_i^j converges to γ . This means that for all i, j the Gromov products $(\gamma, \beta_i^j)_e$ are uniformly bounded by, say, R . By slimness, this says that there exists $R' = R'(R, C)$ such that each $g_i\alpha_j$ intersects $\mathcal{N}_{R'}(e)$. As a consequence, the Gromov products $(g_i\alpha_1, g_i\alpha_2)_{g_i}$ diverge. However, all said Gromov products are equal to $(\alpha_1, \alpha_2)_e$. This is impossible unless this product is infinite, and hence $[\alpha_1] = [\alpha_2]$, contradicting one of the hypotheses. ■

We now augment stable subsets of G to make their boundary a Cantor space.

Lemma 4.2. *Let G be a group with totally disconnected Morse boundary and a Cantor space $C_0 \subset \partial_M G$. Then for every Morse gauge N , there exists an N' and a Cantor space $C \subseteq \partial_M G$ with $\partial_M^N G \subseteq C \subseteq \partial_M^{N'} G$.*

Proof. Since C_0 is compact by [7, Lemma 4.1], we know that $C_0 \subset \partial_M^{N_0} G$ for some Morse gauge N_0 . Since $\partial_M G$ is totally disconnected, it follows that for any Morse gauge N , so is $\partial_M^N G$. We will consider some N so that $\partial_M^N G$ is non-empty. Define

$$A = \partial_M^N G \cup \left(\bigcup_{g \in G_e^{(N)}} gC_0 \right) \subseteq \partial_M G.$$

In fact, by Lemma 2.4, there exists N' depending on N, N_0 only so that $A \subseteq \partial_M^{N'} G$.

Let now $C = \overline{A}$, and we claim that C is a Cantor space. To check this, we need show that C is non-empty, perfect, compact, totally disconnected, and metrizable. It is metrizable, totally disconnected, and compact because it is a closed subspace of $\partial_M^{N'} G$. It is clearly non-empty so what is left to show is that it is perfect; i.e., C has no isolated points. Any point in $C - A$ is clearly not isolated. Also, any point in some gC_0 is not isolated, since gC_0 is a Cantor space. Finally, any point in $\partial_M^N G$ is a limit of points in C (in fact, even in A) by Lemma 4.1. ■

We now show that the subsets $\partial_M^N G$ have empty interior in $\partial_M^{N'} G$ whenever $N' \gg N$.

Lemma 4.3. *If $\partial_M^N G \subsetneq \partial_M G$ and $\partial_M G$ is not compact, then there exists N' so that for all $N'' \geq N'$, $\partial_M^N G$ has empty interior in $\partial_M^{N''} G$.*

Proof. Note that it suffices to find N' so that $\partial_M^N G \subseteq \overline{\partial_M^{N'} G - \partial_M^N G}$ because then it automatically follows that for all $N'' \geq N'$, $\partial_M^N G \subseteq \overline{\partial_M^{N''} G - \partial_M^N G}$, which is the same as $\partial_M^N G$ having empty interior in $\partial_M^{N''} G$.

Let γ be a geodesic in G , based at e , representing a point $p \in \partial_M^N G$. Let $\gamma(i) = g_i \in G$. Then the sequence of points (g_i) converges to p . Passing to a subsequence if necessary, the sequence (g_i^{-1}) also converges to some point $q \in \partial_M G$. Since $\partial_M G$ is not compact, we can choose a point $z \in \partial_M G$ with $z \neq q$ and $z \notin \partial_M^{N_1} G$, where N_1 is the Morse gauge from Lemma 2.4. It follows from [19, Lemmas 5.3 and 6.9] that for some N' , $g_i z$ converges to p in $\partial_M^{N'} G$.

It remains to show that $g_i z \notin \partial_M^N G$ for all i . To see this, let β_i be a geodesic ray based at e representing $g_i z$ and let γ_i be the segment of γ from e to g_i . Then $\gamma_i, \beta_i, g_i \beta_0$ form a geodesic triangle with vertices $e, g_i, g_i z$. By Lemma 2.4, if γ_i and β_i are both N -Morse, then $g_i \beta_0$ is N_1 -Morse, and hence its translate, β_0 , is also N_1 -Morse. But $z = [\beta_0]$, so this contradicts our choice of z and we conclude that β_i is not N -Morse. Thus we have $g_i z \in \partial_M^{N'} G - \partial_M^N G$ for all i . ■

We are ready to prove Theorem 1.4.

Proof of Theorem 1.4. If $\partial_M G$ is compact, then G is hyperbolic by [7, Theorem 4.3]. Since $\partial_M G$ is a non-empty, compact, perfect, totally disconnected, metrizable metric space, it follows that its boundary is a Cantor space by [3]. Since G is a hyperbolic group with Cantor space boundary, G has to be virtually free; see e.g. [15, Theorem 8.1] and references therein.

Assume now that $\partial_M G$ is not compact, in which case G is not hyperbolic again by [7, Theorem 4.3]. Since $\partial_M G$ is σ -compact, we know by Lemma 2.6 that $\partial_M G = \varinjlim \partial_M^N G$ can be chosen to be a countable limit over gauges N_1, N_2, \dots . By Lemma 4.2 there is a Cantor space C_1 such that $\partial_M^{N_1} G \subset C_1 \subset \partial_M^{N_{j(2)}} G$ for a sufficiently large $j(2)$. In view of Lemma 4.3, we may increase $j(2)$ to ensure that $\partial_M^{N_1} G$ has empty interior in $\partial_M^{N_{j(2)}} G$. Proceeding inductively, we find a sequence $1 = j(1) < j(2) < \dots$ and Cantor spaces C_1, C_2, C_3, \dots so that $\partial_M^{N_{j(i)}} G \subset C_i \subset \partial_M^{N_{j(i+1)}} G$. Since $\partial_M^{N_{j(i+1)}} G$ has empty interior in $\partial_M^{N_{j(i+2)}} G$, it follows that C_i is entwined in C_{i+1} . Moreover, since $\partial_M G = \varinjlim C_i$, it follows that $\partial_M G$ is an ω -Cantor space. ■

4.1. Totally disconnected boundary versus totally disconnected levels

In this subsection, we show that, when $\partial_M X$ is σ -compact, then it is totally disconnected if and only if all its strata are. This will be needed in our analysis of Morse boundaries of right-angled Artin groups.

Proposition 4.4. *Let X be a proper metric space, let $e \in X$, and suppose that $\partial_M X$ is σ -compact. Then $\partial_M X$ is totally disconnected if and only if $\partial_M^N X$ is totally disconnected for every Morse gauge N .*

Proof. If $\partial_M X$ is totally disconnected, then any given stratum, $\partial_M^N X$, is totally disconnected as well since the stratum is topologically embedded in $\partial_M X$ by Lemma 2.5.

Suppose that $\partial_M^N X$ is totally disconnected for every Morse gauge N . By Lemma 2.6, there exists a sequence of Morse gauges N_1, N_2, \dots with $N_i \leq N_{i+1}$ so that $\partial_M G = \bigcup \partial_M^{N_i} X$.

We will prove that for each pair of distinct points $x, y \in \partial_M X$, there exists a clopen subset $A \subseteq \partial_M X$ so that $x \in A$ and $y \notin A$, which suffices to show that $\partial_M X$ is totally disconnected. We can assume that $x, y \in \partial_M^{N_1} X$. In fact, we will construct a sequence of clopen subsets $A_i \subseteq \partial_M^{N_i} X$ so that $x \in A_i$, $y \notin A_i$, and $A_i \cap \partial_M^{N_j} X = A_j$ for each $j < i$; we can then just set $A = \bigcup A_i$. Also, we can replace the last condition simply by $A_i \cap \partial_M^{N_{i-1}} X = A_{i-1}$ (for $i \geq 2$), and the stronger condition with $j < i$ will follow by a simple inductive argument.

Let us start with the remark that, for each i , since $\partial_M^{N_i} X$ is a totally disconnected compact metrizable space, we have that each point has a neighborhood basis of clopen subsets. In particular, there exists a clopen set $A_1 \subseteq \partial_M^{N_1} X$ with $x \in A_1$ and $y \notin A_1$ (since $\partial_M^{N_1} X$ is Hausdorff).

Suppose that we constructed A_{i-1} with the required properties, for some $i \geq 2$. Set $B_{i-1} = \partial_M^{N_{i-1}} X - A_{i-1}$, and notice that it is also clopen in $\partial_M^{N_{i-1}} X$. Both A_{i-1} and B_{i-1}

are closed in $\partial_M^{N_i} X$, since they are compact (because they are closed in the compact space $\partial_M^{N_{i-1}} X$), and $\partial_M^{N_i} X$ is Hausdorff. Since $\partial_M^{N_i} X$ is normal, there exists an open set $U \subseteq \partial_M^{N_i} X$ containing A_{i-1} , but not intersecting B_{i-1} . Since every point of $\partial_M^{N_i} X$ has a basis of clopen neighborhoods, a straightforward compactness argument gives that A_{i-1} is contained in a finite union of clopen sets of $\partial_M^{N_i} X$, each contained in U . The union A_i of said clopen sets has the required properties, and we are done. ■

4.2. Artin groups

We now show that ω -Cantor spaces arise naturally as Morse boundaries of right-angled Artin groups. Let A_Γ denote the right-angled Artin group with finite defining graph Γ .

Theorem 1.1. *Let A_Γ be a right-angled Artin group, for Γ a finite graph. Then $\partial_M A_\Gamma$ satisfies exactly one of the following:*

- (1) $\partial_M A_\Gamma$ is empty,
- (2) $\partial_M A_\Gamma$ consists of two points,
- (3) $\partial_M A_\Gamma$ is a Cantor space or
- (4) $\partial_M A_\Gamma$ is an ω -Cantor space.

Proof. If A_Γ is a direct product ($\Leftrightarrow \Gamma$ is a join), then every geodesic is contained in a flat, so the Morse boundary of A_Γ is empty. If A_Γ is a free group ($\Leftrightarrow \Gamma$ is discrete), then the Morse boundary is equal to the Gromov boundary which is either a Cantor space (if A_Γ is non-abelian) or just two points (if $A_\Gamma = \mathbb{Z}$).

In all other cases, it is shown in [17] that A_Γ contains a stable non-abelian free group and thus has a Cantor space as a subspace of the Morse boundary. Moreover, $\partial_M A_\Gamma$ is σ -compact since A_Γ is CAT(0) and thus its Morse boundary is equivalent to the contracting boundary defined in [5]. The boundary of $\partial_M^N A_\Gamma$ is totally disconnected for all Morse gauges N by [8, Theorem 5.1], and thus by Proposition 4.4 the whole Morse boundary is disconnected.

Applying Theorem 1.4 gives the desired result. ■

4.3. Graphs of groups

We now show that a certain class of graphs of groups has ω -Cantor space boundaries. We apply this to show that the Morse boundary of non-geometric graph manifolds are ω -Cantor spaces.

Theorem 1.2. *Let G be a finitely generated group with a graph of groups decomposition \mathcal{G} such that the vertex groups are finitely generated and undistorted in G and have empty Morse boundary. Assume further that G acts acylindrically on the Bass–Serre tree associated to the decomposition. Then $\partial_M G$ is a Cantor space or an ω -Cantor space. It is a Cantor space if and only if G is hyperbolic, in which case G is virtually free.*

Proof. We will apply Theorem 1.4, so we need only show that $\partial_M G$ is totally disconnected, σ -compact, and contains a Cantor subspace.

Let $\mathcal{G} = (\Gamma, \{G_v\}, \{G_e\})$ be the graph of groups decomposition. Choose a maximal tree in Γ . Choose a generating set for G consisting of a finite generating set for each vertex group G_v together with a generator t_e for each edge e of Γ not contained in the maximal tree. For edges e in the maximal tree, set $t_e = 1$. Then for any edge e and $h \in G_e$, we have $h_0 t_e = t_e h_1$ where h_0, h_1 are the images of h in the initial and terminal vertex groups of e , respectively.

It will be convenient to work with a certain quasi-isometric model X for G . To construct X , consider the Bass–Serre tree T whose vertices correspond to cosets of the vertex groups. For each vertex v in the Bass–Serre tree, take a copy X_v of the Cayley graph of the corresponding vertex group, but with vertices labeled by elements of the coset. For an edge labeled e in the Bass–Serre tree connecting v to w , attach edges between vertices of the form $gh_0 \in X_v$ and $gh_0 t_e = g t_e h_1 \in X_w$ for each h in G_e . This defines the graph X .

Note that X_v is quasi-isometrically embedded in X since we are assuming that vertex groups are undistorted, and the constants are uniform in v since there are only finitely many G -orbits of spaces X_v . Moreover, there is a 1-Lipschitz map $\pi : X \rightarrow T$ that maps X_v to v , and the preimage of the open star of a vertex v is contained in the 1-neighborhood of X_v .

We next show that if $Y \subset X$ is a stable subset, then Y quasi-isometrically embeds in the Bass–Serre tree T via the projection map $\pi : X \rightarrow T$. From this it will follow that the map on the boundaries induced by π , $\partial_M \pi : \partial_M X \rightarrow \partial T$, is continuous and injective.

Let N be such that every pair of points of Y can be connected by an N -Morse geodesic segment γ in X . For any such γ , since $\partial_M X_v$ is empty and X_v is uniformly quasi-isometrically embedded in X , we have that $\text{diam}(X_v \cap \gamma)$ is bounded by some constant $D \geq 1$ depending on N . We increase D so that a similar bound also holds when intersecting with the 1-neighborhood of X_v . We claim that $\pi \circ \gamma$ is a (D, D) -quasi-geodesic, which is readily seen to be enough to show that Y quasi-isometrically embeds in T . Since this might be of independent interest, we note that a small variation of the arguments below also proves the following more general lemma:

Lemma 4.5. *Let G be a finitely generated group with a graph of groups decomposition \mathcal{G} , and fix a word metric on G and an orbit map $\pi : G \rightarrow T$, where T is the Bass–Serre tree. Then for each $D_1, D_2 \geq 1$ there exists $K \geq 1$ such that the following holds. Let γ be a (D_1, D_1) -quasi-geodesic in G that intersects each coset of a vertex group in a subset of diameter bounded by D_2 . Then $\pi \circ \gamma$ is a (K, K) -quasi-isometric embedding.*

Let $\gamma : [a, b] \rightarrow X$ be an N -Morse geodesic segment, let p_v be the point where γ first enters the 1-neighborhood of X_v , and let q_v be the point where it last leaves said neighborhood. Call this subpath γ_v . Let $\{v_i\}_{i=1}^n \subset T$ be the collection of vertices on the geodesic from $\pi(\gamma(a))$ to $\pi(\gamma(b))$, so that

$$d_T(\pi(\gamma(a)), \pi(\gamma(b))) \geq n - 1.$$

We now argue that the subpaths γ_{v_i} must cover γ . Indeed, fix any $t \in [a, b]$. If $\gamma(t)$ projects to a vertex of the geodesic from $\pi(\gamma(a))$ to $\pi(\gamma(b))$, then this is clear. If $\gamma(t)$ projects to

the interior of an edge of said geodesic, then it is still true that $\gamma(t)$ belongs to γ_{v_i} for v_i any endpoint of the edge since $\gamma(t)$ belongs to the 1-neighborhood of X_{v_i} (recall that the π -preimage of the open star around v is contained in the 1-neighborhood of X_v).

The last remaining case is the one where $\gamma(t)$ does not project to the geodesic from $\pi(\gamma(a))$ to $\pi(\gamma(b))$. Then there is a vertex v_i separating $\pi(\gamma(t))$ from both $\pi(\gamma(a))$ and $\pi(\gamma(b))$. Hence, $\pi \circ \gamma$ has to pass through v_i before and after time t , showing that $\gamma(t)$ lies on γ_{v_i} . In view of this covering, we see

$$|b - a| \leq \sum_{i=1}^n d(p_{v_i}, q_{v_i}) \leq Dn \leq D(d_T(\pi(\gamma(a)), \pi(\gamma(b))) + 1).$$

Since π is 1-Lipschitz, we also have $d_T(\pi(\gamma(a)), \pi(\gamma(b))) \leq d(\gamma(a), \gamma(b)) = |b - a|$, so $\pi \circ \gamma$ is (D, D) -quasi-geodesic as claimed.

It follows from this that the map $\partial_M \pi: \partial_M X \rightarrow \partial T$ is continuous and injective. Since the Gromov boundary of a tree is totally disconnected, we conclude that $\partial_M X$ is also totally disconnected. To complete the proof of the theorem, it remains to show that $\partial_M X$ is σ -compact and has a Cantor subspace.

We first show that $\partial_M X$ has a Cantor subspace. Since G acts acylindrically on a tree, we know by [9, Theorem 6.14] that we have a hyperbolically embedded free group. By [26] this free subgroup is quasi-convex and thus is stable. We have thus found our stable Cantor subspace.

To see that $\partial_M X$ is σ -compact, we stratify the boundary by looking at collections of geodesics that intersect with the vertex spaces in a bounded amount: let $D \geq 0$ and consider the $x \in \partial_M X$ so that there exists a representative γ with the property that $\text{diam}(X_v \cap \gamma) \leq D$ for all v . Call this subset $\partial_D X$. As we noted earlier, any $\partial_M^N X$ is contained in some $\partial_D X$, where D depends on N .

We claim that the reverse is also true, that is, $\partial_D X$ is contained in $\partial_M^{N'} X$ for some N' . This will prove that $\partial_M X$ is σ -compact since the closures of the $\partial_D X$ provide an exhaustion of $\partial_M X$ by compact sets. We believe that this can be deduced from the arguments in the proof of [10, Theorem 4.1], but we will give an alternate proof here. Again, since this might be of independent interest, we note that similar arguments show more generally:

Lemma 4.6. *Let G be a finitely generated group acting acylindrically on a hyperbolic space Z , and fix a word metric on G and an orbit map $\pi : G \rightarrow Z$. Then for every $D_1, D_2 \geq 1$ there exists a Morse gauge N so that the following holds. If γ is a (D_1, D_1) -quasi-geodesic in G such that $\pi \circ \gamma$ is a (D_2, D_2) -quasi-geodesic in Z , then γ is N -Morse.*

To see that $\partial_D X$ is contained in $\partial_M^{N'} X$ for some N' , let γ be a representative of a point in $\partial_D X$, and let α be an (L, L) -quasi-geodesic in X with endpoints on γ . We claim that for a sufficiently large choice of t , there exists R , depending only on D, L, t such that any subsegment of α lying outside the t -neighborhood of γ has length at most R . It will then follow that α lies in the $(t + R)$ -neighborhood of γ .

To show this, we will apply Lemma 10.4 of [22]. Up to replacing α with piecewise geodesic path lying within Hausdorff distance $2L$ from it, and replacing L by $4L$, we can assume in addition that α is L -Lipschitz, a requirement for applying this lemma. (This follows from the “taming” process described in [2, Lemma III.H.11]).

Say α' is a segment of α whose endpoints a_1, b_1 are distance t from γ and whose interior lies outside the t -neighborhood of γ . Let a_2, b_2 be points on γ at distance t from a_1, b_1 and let γ' be the segment of γ from a_2 to b_2 . Let $\bar{a} = \pi(a_2)$ and $\bar{b} = \pi(b_2)$ denote the endpoints of $\pi \circ \gamma'$. The lemma states that there exist a divergent function ρ and a constant C such that

$$\max \{l(\alpha'), l(\gamma')\} \geq (d_T(\bar{a}, \bar{b}) - 2t - C)\rho(t). \tag{*}$$

Choose t sufficiently large so that $t > C$ and $\rho(t) > 10DL$. Set $R = 10DLt$ and suppose that $l(\alpha') > R$. (Without loss of generality, we may assume that all constants appearing in this argument are ≥ 1 .) Since γ lies in $\partial_D X$, γ' intersects any X_v in a segment of diameter at most D , so the argument above shows that

$$l(\gamma') = d_X(a_2, b_2) \leq D(d_T(\bar{a}, \bar{b}) + 1).$$

And since α' is an (L, L) -quasi-geodesic,

$$\begin{aligned} l(\alpha') &\leq Ld_X(a_1, b_1) + L \\ &\leq Ld_X(a_2, b_2) + 2Lt + L \\ &\leq DLd_T(\bar{a}, \bar{b}) + DL + 2Lt + L \\ &\leq DL(d_T(\bar{a}, \bar{b}) + 2t + 2). \end{aligned} \tag{**}$$

We thus have $10DLt < l(\alpha') \leq DL(d_T(\bar{a}, \bar{b}) + 2t + 2)$, from which we deduce that $d_T(\bar{a}, \bar{b}) > 8t - 2 > 6t$, and hence

- (1) $d_T(\bar{a}, \bar{b}) + 2t + 2 < d_T(\bar{a}, \bar{b}) + 3t < 2d_T(\bar{a}, \bar{b})$, and
- (2) $d_T(\bar{a}, \bar{b}) - 2t - C > d_T(\bar{a}, \bar{b}) - 3t > d_T(\bar{a}, \bar{b})/2$.

Applying (1) to the inequality (**) gives

$$l(\alpha') < 2DLd_T(\bar{a}, \bar{b}).$$

Applying (2) to the inequality (*) gives

$$\max \{l(\alpha'), l(\gamma')\} > 5DLd_T(\bar{a}, \bar{b}).$$

This is a contradiction, so we conclude that no such segment α' exists. This completes the proof of the theorem. ■

Corollary 1.3. *Let M be a non-geometric graph manifold. Then $\partial_M(\pi_1(M))$ is an ω -Cantor space.*

Proof. Let \mathcal{G} be the graph of groups decomposition of $\pi_1(M)$ associated to the geometric decomposition of M along embedded tori and Klein bottles into finitely many Seifert manifolds. The vertices of the Bass–Serre tree associated to \mathcal{G} correspond to the Seifert manifolds, which have empty Morse boundary and are undistorted in $\pi_1(M)$. Since $\pi_1(M)$ acts acylindrically on the Bass–Serre tree [1, Theorem 7.27], we can apply the previous theorem. ■

Remark 4.7. We note that Lemmas 4.5 and 4.6 also imply that a group G satisfying these hypotheses is Morse detectable as defined in [25], and in particular it has the Morse-local-to-global property by a result in said paper.

5. Limits of Sierpiński curves

We now begin our study of limits of Sierpiński curves. First, we state the characterization of the Sierpiński curve that will be most useful for our purposes. We then define entwined Sierpiński curves, and study direct limits of entwined Sierpiński curves.

A Sierpiński curve S is a topological space homeomorphic to $\mathbb{S}^2 - \bigcup_{i \in \mathbb{N}} D_i$, where D_i is the interior of a closed disk in \mathbb{S}^2 satisfying

- $\bar{D}_i \cap \bar{D}_j = \emptyset$ all i, j ,
- $\text{diam}(D_i) \rightarrow 0$ as $i \rightarrow \infty$,
- $\bigcup \bar{D}_i$ is dense in \mathbb{S}^2 .

All Sierpiński curves are homeomorphic to each other [29]. The peripheral circles of a Sierpiński curve are the boundaries of the D_i as above, and they can be characterized as the subspaces γ of S homeomorphic to S^1 so that $S - \gamma$ is connected (this follows from the Jordan curve theorem; see the argument in [18, p. 256, case (ii)]).

Definition 5.1. Let $S \subseteq T$ be Sierpiński curves. We say that S is *entwined* in T if no peripheral circle of S intersects any peripheral circle of T .

The following lemma says the following. Suppose that $S \subseteq T$ are Sierpiński curves, with S entwined in T . Then T is obtained from S by attaching Sierpiński curves onto each peripheral circle.

Lemma 5.2. *Let $S \subseteq T$ be Sierpiński curves, with S entwined in T . Then there exist Sierpiński curves $S_i \subseteq T$ with the following properties:*

- (1) $T = S \cup \bigcup S_i$,
- (2) $S_i \cap S$ is a peripheral circle of S and S_i , and any peripheral circle of S arises in this way,
- (3) $S_i \cap S_j = \emptyset$ if $i \neq j$,
- (4) in any metric on T compatible with its topology, we have $\text{diam}(S_i) \rightarrow 0$.

Proof. We can assume that $T = \mathbb{S}^2 - \bigcup D_i$, where the D_i are open disks as described above.

Let γ be a peripheral circle of S . Then $\mathbb{S}^2 - \gamma$ is a union of two open disks D_γ, D'_γ whose boundary is γ , by the Jordan curve theorem. One of these two disks, say D_γ , is disjoint from S , for otherwise $S - \gamma$ would be disconnected. Since S is entwined in T , each D_i is contained in either D_γ or D'_γ . Thus, $\overline{D_\gamma} \cap T$ is \mathbb{S}^2 minus a union of interiors of closed disks, namely, D'_γ and the D_i contained in D_γ . Notice that the diameters of the disks D_i go to 0, and that their union is dense in \mathbb{S}^2 , and hence $\overline{D_\gamma} \cap T$ is a Sierpiński curve, which we will denote by S_γ . We claim that the Sierpiński curves $\{S_\gamma\}$ have the required properties.

(1) We first observe that $\mathbb{S}^2 - S = \bigcup D_\gamma$. In fact, the closure of any complementary components of S is a closed disk by [23, Theorem 9] and the Jordan curve theorem, and the boundary of such disk is then a peripheral circle. Thus, $\mathbb{S}^2 = S \cup \bigcup \overline{D_\gamma}$ and hence $T = S \cup \bigcup (\overline{D_\gamma} \cap T) = S \cup \bigcup S_\gamma$.

(2) We have $S_\gamma \cap S = \gamma$ by construction.

(3) It suffices to show that for two distinct peripheral circles γ, α , we have $D_\gamma \cap D_\alpha = \emptyset$. If not, since the boundary circles of γ and α are disjoint, up to switching γ and α , we would have $\overline{D_\gamma} \subseteq D_\alpha$ (given complementary components of γ, α , they are either disjoint or one of them, as well as its closure, is contained in the other). But then we would have $\gamma \subseteq D_\alpha$, in contradiction with the fact that $D_\alpha \cap S = \emptyset$.

(4) This statement does not depend on the choice metric, so for convenience we endow \mathbb{S}^2 with the metric d as a subspace of Euclidean space, and in turn endow S and T with their metrics as subspaces of (\mathbb{S}^2, d) . Since each S_γ is contained in the closure of the corresponding D_γ , it suffices to show that, given any $\varepsilon > 0$, there are only finitely many γ with $\text{diam}(D_\gamma) \geq \varepsilon$.

Fix $\varepsilon > 0$, and we can further assume that $\varepsilon < 1$. Notice that there exist only finitely many γ of diameter at least $\varepsilon/2$ (for example because S is homeomorphic to the standard Sierpiński carpet, and this is true of the peripheral circles in that case), and from now on we only consider those with diameter at most $\varepsilon/2$. The issue we have to deal with now is that D_γ can have diameter much larger than that of its boundary γ (in fact, it is always the case that one of the two components of the complement of γ has “large” diameter). However, what we know, thanks to the fact that we are working with the restriction of the Euclidean metric, is that the component D_γ of the complement of γ either has diameter at most ε , or it contains a ball in (\mathbb{S}^2, d) of radius 1. In fact, γ is contained in a closed Euclidean ball $B \subseteq \mathbb{R}^3$ of radius $\varepsilon/2$, and $B \cap \mathbb{S}^2$ is connected. Hence, either D_γ is contained in B , yielding the first case, or D_γ contains $\mathbb{S}^2 - B$, which is easily seen to contain a half-sphere (since ε is sufficiently small), yielding the second case.

The second case can occur at most finitely many times since, as we argued above, the D_γ are pairwise disjoint, and we can only fit finitely many disjoint balls of radius 1 on \mathbb{S}^2 . Hence, there are only finitely many γ with $\text{diam}(D_\gamma) \geq \varepsilon$, as required. ■

The following point-set topology lemma will allow us to check continuity of certain maps, namely the ones that we will encounter when constructing homeomorphisms between limits of entwined Sierpiński curves.

Lemma 5.3. *Let $f : X \rightarrow Y$ be a map between two metric spaces and let C_0, C_1, \dots be closed subsets of X . Suppose that*

- (1) $X = \bigcup C_i$,
- (2) $C_0 \cap C_i \neq \emptyset$ for all i ,
- (3) $\text{diam}(C_i) \rightarrow 0$,
- (4) $\text{diam}(f(C_i)) \rightarrow 0$,
- (5) $f|_{C_i}$ is continuous for every i .

Then f is continuous.

Proof. Let (x_n) be a sequence in X converging to x . It suffices to show that a subsequence of $(f(x_n))$ converges to $f(x)$.

We consider two cases, at least one of which applies, by property (1).

(1) There is a subsequence (x_{n_k}) contained in some C_j . Since C_j is closed, we also have $x \in C_j$. Then $(f(x_{n_k}))$ converges to $f(x)$ because $f|_{C_j}$ is continuous by property (5).

(2) There is a subsequence (x_{n_k}) so that no two elements of the subsequence are contained in the same C_j . Suppose that $x_{n_k} \in C_k$. Let $x'_{n_k} \in C_k \cap C_0$ (which exists by property (2)). Then (x'_{n_k}) converges to x in view of property (3), which implies that $d_X(x'_{n_k}, x_{n_k})$ tends to 0. By case (1) (with $j = 0$), we have that $(f(x'_{n_k}))$ converges to $f(x)$. Finally, by property (4), $d_Y(f(x_{n_k}), f(x'_{n_k}))$ also tends to 0, so $(f(x_{n_k}))$ also converges to $f(x)$. ■

The following proposition is, essentially, the inductive step in the definition of homeomorphisms between limits of entwined Sierpiński curves.

Proposition 5.4. *Let $S \subseteq T, S' \subseteq T'$ be Sierpiński curves, with S entwined in T and S' entwined in T' . Then any homeomorphism $\phi : S \rightarrow S'$ can be extended to a homeomorphism $\psi : T \rightarrow T'$.*

Proof. Let $S_i \subseteq T, S'_i \subseteq T'$ be the Sierpiński curves that we obtain from Lemma 5.2. Since $\{S_i \cap S\}$ is the set of peripheral circles of S and $\{S'_i \cap S'\}$ the set of peripheral circles of S' , we can choose the indices in such a way that $\phi(S_i \cap S) = S'_i \cap S'$.

The arguments from [29, pp. 322–323] show that any homeomorphism between peripheral circles of Sierpiński curves can be extended to a homeomorphism between the Sierpiński curves. In particular, there are homeomorphisms $\phi_i : S_i \rightarrow S'_i$ that extend $\phi|_{S_i \cap S}$. Hence, we can define $\psi : T \rightarrow T'$ by requiring $\psi|_S = \phi, \psi|_{S_i} = \phi_i$. This is a well-defined map in view of Lemma 5.2 (3). Then ψ is clearly bijective. By Lemma 5.2 (1), (2), and (4), the hypotheses of Lemma 5.3 hold, so ψ is continuous. Hence, it is a homeomorphism because T, T' are compact and Hausdorff. Thus, ψ is the desired extension of ϕ . ■

Finally, we are ready for the main theorem of this section.

Definition 5.5. An ω -Sierpiński curve is a topological space $\lim_{\rightarrow i \in \mathbb{N}} X_i$, where each X_i is a Sierpiński curve and each X_i is entwined in X_{i+1} .

Theorem 5.6. Any two ω -Sierpiński curves are homeomorphic.

Proof. The proof is identical to the proof of Theorem 3.3, using Proposition 5.4 instead of Lemma 3.6. ■

6. ω -Sierpiński boundaries

We now study Morse boundaries of fundamental groups of finite-volume hyperbolic 3-manifolds. The main theorem of this section is the following:

Theorem 1.5. Let M be a finite-volume hyperbolic 3-manifold with at least one cusp, and let $G = \pi_1(M)$. Then $\partial_M G$ is an ω -Sierpiński curve.

6.1. Setup and notation

Let M be a finite-volume hyperbolic 3-manifold with at least one cusp, and let $G = \pi_1(M)$, regarded as a subgroup of $\text{Isom}(\mathbb{H}^3)$. Lifting disjoint cuspidal neighborhoods to \mathbb{H}^3 , we obtain an equivariant family of disjoint horoballs $\{\mathcal{H}_p\}_{p \in P}$. Regarding \mathbb{S}^2 with the standard metric as the boundary of \mathbb{H}^3 , we have the collection $P \subseteq \mathbb{S}^2$ of parabolic points for G . For each p , we denote $r_p = e^{-d_{\mathbb{H}^3}(x_0, H_p)}$, where x_0 is the origin of the Poincaré disk model of \mathbb{H}^3 .

For a given $0 < \lambda \leq 1$, we denote $V_\lambda = \mathbb{S}^2 \setminus \bigcup_{p \in P} B(p, \lambda r_p)$. Set $X = \mathbb{H}^3 - \bigcup H_p$, a neutered space for M . Denote by $O_p \subseteq X$ the horosphere that bounds H_p .

We will use the fact that X is a CAT(0) space and that each O_p is a flat in X , meaning that it is convex and isometric to \mathbb{R}^2 ; see e.g. [2, Chapter II.11]. In addition, the inclusion map $\iota : X \rightarrow \mathbb{H}^3$ is 1-Lipschitz and proper.

We will use this notation throughout this section.

6.2. From the Morse boundary to $\partial \mathbb{H}^3$

In this section, we relate the Morse boundary of X with $\partial \mathbb{H}^3$; the main result is Proposition 6.4. Roughly speaking, we have to show that Morse geodesic rays in X correspond to geodesic rays in \mathbb{H}^3 that do not spend too much time in horoballs, and also relate the latter rays with points in the subspace $V_\lambda \subset \partial \mathbb{H}^3$ defined above. This can be done in greater generality than our case (groups hyperbolic relative to subgroups with empty Morse boundaries, replacing \mathbb{H}^3 with the cusped space), but we decided to use the extra structure available to us to make the proofs simpler.

First of all, we show that Morse geodesic rays in X give quasi-geodesics in \mathbb{H}^3 .

Lemma 6.1. Let γ be an N -Morse geodesic ray in X . Then $\iota(\gamma)$ is a (K, L) -quasi-geodesic ray where K, L depend only on N .

Proof. We will show that there exists a (λ, ε) -coarse Lipschitz retraction of \mathbb{H}^3 onto $\iota(\gamma)$, where λ, ε depend on N . It follows that $\iota(\gamma)$ is a (K, L) -quasi-geodesic, where K, L depend on λ, ε .

Since γ is N -Morse and X is CAT(0), by [5], we know that γ is D -strongly contracting, where D depends on N .

Let $\pi : X \rightarrow \gamma$ be the closest-point projection in X . We claim that for each horosphere O_p , the diameter of $\pi(O_p)$ is bounded in terms of N . In fact, suppose that the said diameter is much larger than D (recall that D depends on N), so that there exist points $x, y \in O_p$ with far away projections to γ . The geodesic α from x to y is contained in O_p , and by strong contraction it has to pass within controlled distance of $\pi(x), \pi(y)$. Hence, γ has a long subgeodesic, the one with endpoints $\pi(x), \pi(y)$, that lies in a controlled neighborhood of O_p . The length of this subgeodesic is, up to additive constants, the same as the diameter of $\pi(O_p)$. Since O_p is a flat, and γ is N -Morse, this length can then be controlled in terms of N , for otherwise we would find a quasi-geodesic in O_p straying too far away from γ .

Now, using the claim above, we can extend π to $\pi' : \mathbb{H}^3 \rightarrow \gamma$ by mapping $H_p - O_p$ to any point in $\pi(O_p)$. It is readily seen that π' is a coarsely Lipschitz retraction with controlled constants, as required. ■

We now show that rays that do not spend much time in any horoball are exactly those with limit points in V_λ .

Lemma 6.2. *Fix a basepoint $x_0 \in X$, let γ be a geodesic ray in \mathbb{H}^3 starting at x_0 with limit point $x \in \partial\mathbb{H}^3$, and let $p \in P$.*

- (1) *For every C , there exists λ with the following property. Suppose that γ intersects the horoball H_p in a set of diameter at most C . Then $x \notin B(p, \lambda r_p)$.*
- (2) *For every λ , there exists C with the following property. Suppose that $x \notin B(p, \lambda r_p)$. Then γ intersects the horoball H_p in a set of diameter at most C .*

Proof. Recall that the standard metric on $\mathbb{S}^2 = \partial\mathbb{H}^3$ can be written in terms of the Gromov product as $d_{\partial\mathbb{H}^3}([\alpha], [\beta]) = e^{-(\alpha, \beta)_{x_0}}$.

(1) We have to show that, for a suitable $\lambda = \lambda(C)$, if $x \in B(p, \lambda r_p)$, then $\text{diam}(\gamma \cap H_p) \geq C$. If $x \in B(p, \lambda r_p)$, then the Gromov product of γ and the geodesic ray γ_p are at least $d(x_0, H_p) - \ln(\lambda)$. Thus when λ is close to 0, γ and γ_p stay within δ of each other (where δ is the hyperbolicity constant of \mathbb{H}^3) for much longer than $d(x_0, H_p)$, so that γ intersects H_p in a set of large diameter. (Here we use that, up to bounded error, two rays δ -fellow-travel for time equal to their Gromov product; for later purposes we note that the same holds replacing “ δ ” with “ 2δ ”.)

(2) We have to show that, given λ , there exists $C = C(\lambda)$ such that if $\text{diam}(\gamma \cap H_p) \geq C$, then $x \in B(p, \lambda r_p)$. Let γ_p be as above, and let β_p be the geodesic line with endpoints at infinity x and p . Since horoballs are convex, $\gamma \cap H_p$ is a subgeodesic of γ , say with endpoints a, b . Since the Busemann function associated to p is monotonic on γ_p and β_p ,

we see that for C large compared to δ , up to switching a and b , a is δ -close to γ_p and b is δ -close to β_p (since a, b cannot be both close to, say, γ_p).

Consider the triangle with vertices a, b, p . Since a and b lie on the horosphere bounding H_p , the midpoint m of $[a, b]$ is equidistant from the other two sides (this is easily seen in the upper half space model of \mathbb{H}^3) and by the thin triangle condition, this distance is at most δ . It follows that the geodesic segment $[a, m]$ lies in the δ -neighborhood of $[a, p]$ which, in turn, lies in the δ -neighborhood of γ_p . Thus γ and γ_p stay 2δ -close for a distance of at least $d(x_0, H_p) + \frac{1}{2}C$. From this, we can see that for C sufficiently large, the Gromov product $(\gamma, \gamma_p)_{x_0}$ will be at least $d(x_0, H_p) - \ln(\lambda)$, and hence the distance between x and p will be at most λr_p . ■

The proof of the following lemma is a variation on the usual arguments to show that contracting properties of various kinds imply the Morse property. We will need it to show that certain rays in X , constructed from rays in \mathbb{H}^3 , are Morse.

Lemma 6.3. *For every diverging function $f : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, there exists a Morse gauge N with the following property. Let Z be a geodesic metric space and let $\gamma \subseteq Z$ be a geodesic. Suppose that for each $R > 0$, any path β that intersects $N_R(\gamma)$ only at its endpoints x, y has the property that $l(\beta) \geq f(R)d(x, y) - f(R)$. Then γ is N -Morse.*

Proof. Let $\alpha : [a, b] \rightarrow Z$ be a (K, C) -quasi-geodesic, so $l(\alpha|_{[t,u]}) \leq K|t - u| + C$ for each $t, u \in [a, b]$. Choose R so that $f(R) \geq K^2 + K$. We will show that any subpath β of α that intersects $N_R(\gamma)$ only at its endpoints $x = \alpha(t)$ and $y = \alpha(u)$ has length bounded by a constant D depending only on $K, C, f(R)$, showing that α is contained in the $(R + D)$ -neighborhood of γ .

We have

$$\begin{aligned} K|t - u| + C &\geq l(\beta) \geq f(R)d(x, y) - f(R) \\ &\geq \frac{f(R)}{K}|t - u| - Cf(R) - f(R) \\ &\geq (K + 1)|t - u| - Cf(R) - f(R), \end{aligned}$$

from which it follows that $|t - u| \leq Cf(R) + f(R) + C$. Setting

$$D = K Cf(R) + K f(R) + K C + C,$$

we conclude that

$$l(\beta) \leq K Cf(R) + K f(R) + K C + C = D$$

as desired. ■

Finally, we compare the spaces $\partial_M^N X$ with the spaces V_λ and then describe how; in order to study $\partial_M X$, one can study the V_λ instead. Hung Cong Tran has related results for general relatively hyperbolic groups [27]. However, for us the control on the size of the removed balls will be crucial, and this is not addressed in Tran’s paper.

Proposition 6.4. *There exists a continuous injective map $\Psi : \partial_M X \rightarrow \mathbb{S}^2$ with the following properties.*

- (1) *For each Morse gauge N , there exists $\lambda = \lambda(N)$ so that $\Psi(\partial_M^N X) \subseteq V_\lambda$.*
- (2) *For every $0 < \lambda \leq 1$, there exists a Morse gauge N_λ so that for each $N \geq N_\lambda$ we have $V_\lambda \subseteq \Psi(\partial_M^N X)$.*

Proof. Throughout the proof, we fix basepoints $x_0 = \iota(x_0)$ of X and \mathbb{H}^3 , and when discussing (quasi-)geodesic rays, we will assume that they are based at x_0 .

We first define $\Psi : \partial_M X \rightarrow \mathbb{S}^2$. By Lemma 6.1, given $\ell \in \partial_M^N X$ and an N -Morse geodesic ray γ representing ℓ , we can define $\Psi(\ell)$ as the limit point of $\iota(\gamma)$. The fact that Ψ is well defined and continuous follows from the fact that if two N -Morse geodesic rays γ, γ' have initial subgeodesics of length L that stay within distance C of each other, then $\iota(\gamma), \iota(\gamma')$ have the same property (since ι is 1-Lipschitz). Injectivity follows from the fact that if the distance between $\gamma(t), \gamma'(t)$ diverges as t goes to infinity, then the same is true for $\iota(\gamma), \iota(\gamma')$ since ι is a proper map.

In view of Lemma 6.2, in order to prove item (1), we have to show that given an N -Morse geodesic ray γ in X , the geodesic ray α in \mathbb{H}^3 asymptotic to $\iota(\gamma)$ intersects each horoball H_p in a set of diameter bounded by a constant depending only on N . Since the Hausdorff distance between α and $\iota(\gamma)$ is bounded in terms of N , it suffices to bound the diameter of the intersection of $\iota(\gamma)$ with a suitable neighborhood of H_p . In turn, since ι is a proper map, it suffices to do the same in X . That is, we need to show that the diameter of the intersection of γ with the R -neighborhood of O_p is bounded in terms of R and N . To see this, note that O_p is a flat, and as the diameter of this intersection increases, we can find quasi-geodesic in the flat straying farther and farther away from γ . Since γ is N -Morse, the intersection must have bounded diameter.

In view of Lemma 6.2, in order to prove item (2), given a geodesic ray α in \mathbb{H}^3 that intersects any horoball H_p in a set of diameter at most C , we have to find an N -Morse geodesic ray γ in X so that $\iota(\gamma)$ lies within finite Hausdorff distance of α , where $N = N(C)$.

We first do the case $C = 0$, and then show how to reduce to this case. If $C = 0$, then α is contained in X , and moreover it is a geodesic ray in X since ι is 1-Lipschitz. We only have to argue that α is Morse in X , with controlled Morse gauge. By Lemma 6.3, it suffices to prove that whenever β is a path (in X) intersecting the R -neighborhood of α only at its endpoints x, y , then $l(\beta) \geq f(R)d_X(x, y) - f(R)$, where f is some fixed diverging function. Consider such a path β , and regard it now as a path in \mathbb{H}^3 . Since ι is proper, β lies outside the $\rho(R)$ -neighborhood of α (in \mathbb{H}^3), where ρ is some diverging function. Moreover, ι is 1-Lipschitz and hence the endpoints of β are R -close to α . Using the hyperbolicity of \mathbb{H}^3 (or even just that α is strongly contracting), one can then show that the length of β is at least $g(\rho(R))d_{\mathbb{H}^3}(x, y) - 2R$, for some diverging function g . Notice that $|d_{\mathbb{H}^3}(x, y) - d_X(x, y)| \leq 10R$, since x, y lie R -close to α both in X and in \mathbb{H}^3 , and α is a geodesic in both X and \mathbb{H}^3 . Hence, we get the required inequality.

We are only left to reduce the case of a general C to the case $C = 0$, and we will do so by changing the neutered space. Given a geodesic ray α in \mathbb{H}^3 that intersects any horoball H_p in a set of diameter at most C , we can regard it as a geodesic ray in a neutered space X' containing X , in which we shrunk all the horoballs a uniform amount. The previous argument yields that α is a Morse geodesic ray in X' . There is a quasi-isometry $\phi : X' \rightarrow X$, whose constants depend only on C , that moves each point a bounded amount. Hence, $\phi(\alpha)$ is an N -Morse (K, C) -quasi-geodesic, with N, K, C depending only on C . It is readily seen that a geodesic ray γ within bounded Hausdorff distance from $\phi(\alpha)$ has the property that $\iota(\gamma)$ lies within finite Hausdorff distance of α , and we are done. ■

6.3. Approximating strata with Sierpiński curves

As in the case of totally disconnected boundaries, we have to wiggle the strata in order for them to be actual Sierpiński curves. The following proposition, which we prove later, is what will allow us to do this: recall that $P \subset \partial\mathbb{H}^3$ denotes the set of parabolic points for the group $G = \pi_1(M)$.

Proposition 6.5. *For each sufficiently small $0 < \lambda \leq 1$, and each $p \in P$, there exists a circle γ so that*

- (1) $\gamma \subseteq B(p, 3\lambda r_p/4) - B(p, \lambda r_p/4) = A$ and γ is homotopically non-trivial in A ,
- (2) for each $p' \in P - \{p\}$ with $r_{p'} \leq r_p$, we have $\gamma \cap B(p', 3\lambda r_{p'}/4) = \emptyset$.

Corollary 6.6. *For each sufficiently small $0 < \lambda \leq 1$, there exists a Sierpiński curve S with $V_\lambda \subseteq S \subseteq V_{\lambda/4}$, with each peripheral circle of S contained in some $B(p, 3\lambda r_p/4)$.*

Proof. Consider for each p the open disk containing p bounded by the circle γ as in the previous lemma. Such disks are either disjoint or nested, so that we see that the complement of all the disks is the required Sierpiński curve. ■

We can now prove Theorem 1.5.

Theorem 1.5. *Let M be a finite-volume hyperbolic 3-manifold with at least one cusp, and let $G = \pi_1(M)$. Then $\partial_M G$ is an ω -Sierpiński curve.*

Proof. Since G is quasi-isometric to the neutered space X , we can work with $\partial_M X$ instead. First of all, $\partial_M X$ is σ -compact since X is CAT(0) by [5, Main Theorem]. Hence, by Lemma 2.6 we have that $\partial_M X = \lim_{\rightarrow} \partial_M^N X$ can be chosen to be a countable limit over gauges N_1, N_2, \dots . By Proposition 6.4 and Corollary 6.6, and provided that N_1 is large enough, there is a Sierpiński curve S_1 so that $\partial_M^{N_1} X \subset S_1 \subset \partial_M^{N_{j(2)}} X$ for a sufficiently large $j(2)$. In fact, we can further require that $\lambda(N_{j(2)}) < \lambda(N_1)/4$, for $\lambda(N)$ as in Proposition 6.4. Applying Proposition 6.4 and Corollary 6.6 again, we find $j(3)$ and another Sierpiński curve S_2 so that $\partial_M^{N_{j(2)}} X \subset S_2 \subset \partial_M^{N_{j(3)}} X$, again with $\lambda(N_{j(3)}) < \lambda(N_{j(2)})/4$. The condition $\lambda(N_{j(2)}) < \lambda(N_1)/4$ ensures that S_1 is entwined in S_2 . Proceeding inductively, we see that $\partial_M X$ is a limit of entwined Sierpiński curves. Thus, $\partial_M X$ is an ω -Sierpiński curve, as required. ■

6.4. Proof of Proposition 6.5

We are only left to prove Proposition 6.5. We will follow arguments from [21], and the construction is roughly as follows. We start with a circle that possibly does not avoid all required balls around parabolic points, and we make detours to avoid balls of a certain size. We iterate the procedure for balls of smaller and smaller size, and then we take a limit. For technical reasons, we will work mostly with arcs rather than circles.

As in [21], in order to make this work we need machinery from [20]; we now state all relevant facts and definitions. First of all, after detouring, we would like the new circle to be close to the previous one, and this is captured by the following definition.

Definition 6.7. For any x and y in an embedded arc A , let $A[x, y]$ be the closed, possibly trivial, subarc of A that lies between them.

An arc B ι -follows an arc A if there exists a (not necessarily continuous) map $p : B \rightarrow A$, sending endpoints to endpoints, such that for all $x, y \in B$, $B[x, y]$ is in the ι -neighborhood of $A[p(x), p(y)]$; in particular, p displaces points at most ι .

In [20], Mackay defines a space X to be N -doubling if every ball can be covered by at most N balls of half the radius and L -linearly connected if for all $x, y \in X$ there exists a compact, connected subset Y containing x, y of diameter less than or equal to $Ld(x, y)$.

The following proposition will be used to remove unwanted detours from our circles (which in the limit might create unwanted topology).

Proposition 6.8 ([20, Proposition 2.1]). *Given a complete metric space X that is L -linearly connected and N -doubling, there exist constants $s = s(L, N) > 0$ and $S = S(L, N) > 0$ with the following property: for each $\iota > 0$ and each arc $A \subset X$, there exists an arc J that ι -follows A , has the same endpoints as A , and satisfies*

$$\forall x, y \in J, d(x, y) < s\iota \Rightarrow \text{diam}(J[x, y]) < S\iota. \tag{1}$$

The following lemma will allow us to take limits.

Lemma 6.9 ([20, Lemma 2.2]). *Suppose that (X, d) is an L -linearly connected, N -doubling, complete metric space, and let s, S, ε , and δ be fixed positive constants satisfying $\delta \leq \min\{\frac{s}{4+2S}, \frac{1}{10}\}$. If we have a sequence of arcs $J_1, J_2, \dots, J_n, \dots$ in X , such that for every $n \geq 1$*

- J_{n+1} $\varepsilon\delta^n$ -follows J_n , and
- J_{n+1} satisfies (1) with $\iota = \varepsilon\delta^n$ and s, S as fixed above,

then the Hausdorff limit $J = \lim_{\mathcal{H}} J_n$ exists, and is an arc. Moreover, the endpoints of J_n converge to the endpoints of J .

Finally, we need a lemma about the geometry of horoballs.

Lemma 6.10. *There exists λ_0 so that for each $\lambda_1, \lambda_2 \leq \lambda_0$, if $B(p, \lambda_1 r_p) \cap B(p', \lambda_2 r_{p'}) \neq \emptyset$, then either $r_{p'} \leq r_p/100$ or $r_p \leq r_{p'}/100$.*

Proof. In view of Lemma 6.2, a point in $B(p, \lambda_1 r_p) \cap B(p', \lambda_2 r_{p'})$ is represented by a ray from the basepoint that spends a long time in both horoballs H_p and $H_{p'}$. If the geodesic ray goes through H_p first (the other case is symmetric), then $H_{p'}$ is much further away from the basepoint than H_p . Choosing an appropriate value for λ_0 , we can make the segment inside H_p as long as we want, and thus control the distance to $H_{p'}$. In particular, we can guarantee that $r_{p'} \leq r_p/100$. ■

We are now ready to prove Proposition 6.5.

Proof of Proposition 6.5. We can assume that $\lambda \leq \lambda_0$, for λ_0 as in Lemma 6.10.

Start with the circle γ_1 of radius $\lambda/2$ around p . We think of γ_1 as the union of two arcs J_1, J'_1 joining diametrically opposite points.

We will construct a sequence of arcs J_n, J'_n with the same endpoints that have the following properties.

- J_{n+1} satisfies (1) of Proposition 6.8, with $\iota = 50^{-n} \lambda r_p$.
- J_{n+1} ($5 \times 50^{-n} \lambda r_p$)-follows J_n .
- For each $p' \in P$ with $50^{-n-1} r_p < r_{p'} \leq r_p 50^{-n}$, we have $J_{n+1} \cap B(p', \lambda 50^{-n} r_p) = \emptyset$.

And similarly for the sequence J'_n .

Assume that we have constructed J_i, J'_i for $i \leq n$. We now construct J_{n+1} and J'_{n+1} . If J_n and J'_n do not intersect any $B_{p'} = B(p', 2 \times 50^{-n} \lambda r_p)$ with $50^{-n-1} r_p \leq r_{p'} \leq r_p 50^{-n}$, we can take $J_{n+1} = J_n$ and $J'_{n+1} = J'_n$ and we are done (notice the “2” for later purposes). Otherwise, the $B_{p'}$ they intersect are disjoint by Lemma 6.10. We can then replace each maximal segment of J_n or J'_n contained in some $B_{p'}$ with an arc along the boundary of $B_{p'}$. If the segment is initial or terminal, we ensure that we move the corresponding endpoints of J_n and J'_n to the same point. In this way, we obtain immersed paths, which are not necessarily arcs, but we can restrict to arcs contained in the image of such paths. By applying this procedure, we obtain arcs K_n, K'_n that ($4 \times 50^{-n} \lambda r_p$)-follow J_n and J'_n (the constant is the upper bound on the diameters of the $B_{p'}$).

We now apply Proposition 6.8 with $\iota = 50^{-n} \lambda r_p$, obtaining arcs J_{n+1} and J'_{n+1} , satisfying (1), that ($50^{-n} \lambda r_p$)-follow K_n and K'_n , and hence ($5 \times 50^{-n} \lambda r_p$)-follow J_n and J'_n . It is straightforward to check that J_{n+1}, J'_{n+1} have the required properties.

According to Lemma 6.9, there are limit arcs J and J' , that clearly share the same endpoints. While these may intersect in multiple points, there exist a pair of subarcs sharing only endpoints, that form a circle γ around p . We now check that γ satisfies the required properties.

To check that $\gamma \subseteq B(p, 3\lambda r_p/4) - B(p, \lambda r_p/4)$, it suffices to prove the analogous containment for $J \cup J'$. Because of the second property of J_{n+1} , we have that J_{n+1} is contained in the closed neighborhood around J_1 of radius

$$5\lambda r_p (50^{-1} + 50^{-2} + \dots) \leq \lambda r_p/9.$$

A similar statement holds for J' , and since $J_1 \cup J'_1$ is the circle of radius $\lambda r_p/2$ around p , we are done.

Moreover, the same computation as above also yields that J ($\lambda r_p/9$)-follows J_1 , and similarly for J' . One can use this to show that the concatenation γ' of J and J' is homotopic to γ_1 in $A = B(p, 3\lambda r_p/4) - B(p, \lambda r_p/4)$: subdivide each into small arcs and retract arcs in γ' to corresponding arcs in γ_1 along geodesics in $\partial\mathbb{H}^3$. Also, γ' and γ are homotopic since they only differ in the $2\lambda r_p/9$ -neighborhood of the endpoints of J, J' .

Hence, γ is homotopically non-trivial.

Now let $p' \in P - \{p\}$ with $r_{p'} \leq r_p$. If $r_{p'} \geq r_p/50$. Then applying Lemma 6.10 with $\lambda_1 = 3\lambda/4$ and $\lambda_2 = \lambda$ shows that γ is disjoint from $B(p', \lambda r_{p'})$. Hence, suppose that $50^{-n-1}r_p < r_{p'} \leq 50^{-n}r_p$, for some $n \geq 1$. Then we have that $J_{n+1} \cup J'_{n+1}$ does not intersect $B(p', \lambda 50^{-n}r_p)$. Arguing as above, $J \cup J'$ is contained in the neighborhood around $J_{n+1} \cup J'_{n+1}$ of radius

$$5\lambda r_p(50^{-n-1} + 50^{-n-2} + \dots) \leq 50^{-n}\lambda r_p/10.$$

Hence, $J \cup J'$ avoids the ball of radius $(50^{-n}\lambda r_p)(1 - 1/10) \geq 3\lambda r_{p'}/4$ around p' , as required. ■

Funding. Charney was partially supported by NSF grant DMS-1607616. Cordes was partially supported by the ETH Zurich Postdoctoral Fellowship Program, cofunded by a Marie Curie Actions for People COFUND Program. Sisto was partially supported by the Swiss National Science Foundation (grant #182186), as well as partially supported by the grant 346300 for IMPAN from the Simons Foundation and the matching 2015-2019 Polish MNiSW fund.

References

- [1] C. Abbott, S. H. Balasubramanya, and D. Osin, Hyperbolic structures on groups. *Algebr. Geom. Topol.* **19** (2019), no. 4, 1747–1835 Zbl [1481.20165](#) MR [3995018](#)
- [2] M. R. Bridson and A. Haefliger, *Metric spaces of non-positive curvature*. Grundlehren Math. Wiss. 319, Springer, Berlin, 1999 Zbl [0988.53001](#) MR [1744486](#)
- [3] L. E. J. Brouwer, Over de structuur der perfecte puntverzamelingen. *Amst. Ak. Versl.* **18** (1910), 833–842 Zbl [41.0103.03](#)
- [4] R. Charney, M. Cordes, and D. Murray, Quasi-Mobius homeomorphisms of Morse boundaries. *Bull. Lond. Math. Soc.* **51** (2019), no. 3, 501–515 Zbl [1450.20008](#) MR [3964503](#)
- [5] R. Charney and H. Sultan, Contracting boundaries of CAT(0) spaces. *J. Topol.* **8** (2015), no. 1, 93–117 Zbl [1367.20043](#) MR [3339446](#)
- [6] M. Cordes, Morse boundaries of proper geodesic metric spaces. *Groups Geom. Dyn.* **11** (2017), no. 4, 1281–1306 Zbl [1423.20043](#) MR [3737283](#)
- [7] M. Cordes and M. G. Durham, Boundary convex cocompactness and stability of subgroups of finitely generated groups. *Int. Math. Res. Not. IMRN* **2019** (2019), no. 6, 1699–1724 Zbl [1446.22002](#) MR [3932592](#)
- [8] M. Cordes and D. Hume, Stability and the Morse boundary. *J. Lond. Math. Soc. (2)* **95** (2017), no. 3, 963–988 Zbl [1475.20069](#) MR [3664526](#)

- [9] F. Dahmani, V. Guirardel, and D. Osin, Hyperbolically embedded subgroups and rotating families in groups acting on hyperbolic spaces. *Mem. Amer. Math. Soc.* **245** (2017), no. 1156, v+152 Zbl [1396.20041](#) MR [3589159](#)
- [10] C. Druţu, S. Mozes, and M. Sapir, Divergence in lattices in semisimple Lie groups and graphs of groups. *Trans. Amer. Math. Soc.* **362** (2010), 2451–2505 Zbl [1260.20065](#) MR [2584607](#)
- [11] E. Fink, Morse geodesics in torsion groups. 2017, arXiv:[1710.11191](#)
- [12] S. P. Franklin and B. V. S. Thomas, A survey of k_ω -spaces. *Topology Proc.* **2** (1977), no. 1, 111–124 (1978) Zbl [0416.54027](#) MR [540599](#)
- [13] M. Graeber, A. Karrer, N. Lazarovich, and E. Stark, Surprising circles in Morse boundaries of right-angled Coxeter groups. *Topology Appl.* **294** (2021), Paper No. 107645 MR [4234246](#)
- [14] M. Haulmark, H. T. Nguyen, and H. C. Tran, The relative hyperbolicity and manifold structure of certain right-angled Coxeter groups. *Internat. J. Algebra Comput.* **30** (2020), no. 3, 501–537 Zbl [1484.20073](#) MR [4082419](#)
- [15] I. Kapovich and N. Benakli, Boundaries of hyperbolic groups. In *Combinatorial and geometric group theory (New York, 2000/Hoboken, NJ, 2001)*, pp. 39–93, Contemp. Math. 296, Amer. Math. Soc., Providence, RI, 2002 Zbl [1044.20028](#) MR [1921706](#)
- [16] A. Karrer, *Contracting boundaries of amalgamated free products of CAT(0) groups with applications for right-angled Coxeter groups*. Ph.D. thesis, Karlsruhe Institute of Technology (KIT), 2020
- [17] T. Koberda, J. Mangahas, and S. J. Taylor, The geometry of purely loxodromic subgroups of right-angled Artin groups. *Trans. Amer. Math. Soc.* **369** (2017), no. 11, 8179–8208 Zbl [1476.20042](#) MR [3695858](#)
- [18] J. Krasinkiewicz, On homeomorphisms of the Sierpiński curve. *Prace Mat.* **12** (1969), 255–257 Zbl [0235.54039](#) MR [0247618](#)
- [19] Q. Liu, Dynamics on the Morse boundary. *Geom. Dedicata* **214** (2021), 1–20 Zbl [07402671](#) MR [4308266](#)
- [20] J. M. Mackay, Existence of quasi-arcs. *Proc. Amer. Math. Soc.* **136** (2008), no. 11, 3975–3981 Zbl [1153.30019](#) MR [2425738](#)
- [21] J. M. Mackay and A. Sisto, Quasi-hyperbolic planes in relatively hyperbolic groups. *Ann. Acad. Sci. Fenn. Math.* **45** (2020), no. 1, 139–174 Zbl [1453.20058](#) MR [4056532](#)
- [22] P. Mathieu and A. Sisto, Deviation inequalities for random walks. *Duke Math. J.* **169** (2020), no. 5, 961–1036 Zbl [1465.60038](#) MR [4079419](#)
- [23] R. L. Moore, Concerning the common boundary of two domains. *Fund. Math.* **6** (1925), 203–213 Zbl [50.0138.04](#)
- [24] D. Murray, Topology and dynamics of the contracting boundary of cocompact CAT(0) spaces. *Pacific J. Math.* **299** (2019), no. 1, 89–116 Zbl [07054788](#) MR [3947271](#)
- [25] J. Russell, D. Spriano, and H. C. Tran, The local-to-global property for Morse quasi-geodesics. *Math. Z.* **300** (2022), no. 2, 1557–1602 Zbl [07461057](#) MR [4363788](#)
- [26] A. Sisto, Quasi-convexity of hyperbolically embedded subgroups. *Math. Z.* **283** (2016), no. 3–4, 649–658 Zbl [1380.20044](#) MR [3519976](#)
- [27] H. C. Tran, Divergence spectra and Morse boundaries of relatively hyperbolic groups. *Geom. Dedicata* **194** (2018), 99–129 Zbl [1469.20043](#) MR [3798063](#)
- [28] H. C. Tran, On strongly quasiconvex subgroups. *Geom. Topol.* **23** (2019), no. 3, 1173–1235 Zbl [07079057](#) MR [3956891](#)
- [29] G. T. Whyburn, Topological characterization of the Sierpiński curve. *Fund. Math.* **45** (1958), 320–324 Zbl [0081.16904](#) MR [99638](#)
- [30] A. Zalloum, A symbolic coding of the Morse boundary. 2018, arXiv:[1811.10383](#)

Received 11 September 2020.

Ruth Charney

Department of Mathematics, Brandeis University, 415 South Street, Waltham, MA 02421, USA;
charney@brandeis.edu

Matthew Cordes

Department of Mathematics, ETH Zürich, Rämistrasse 101, 8092 Zürich, Switzerland;
matthew.cordes@math.ethz.ch

Alessandro Sisto

School of Mathematical and Computer Sciences, Heriot-Watt University, Edinburgh EH14 4AS,
Scotland, UK; a.sisto@hw.ac.uk